

Package: depictr (via r-universe)

October 16, 2026

Type Package

Title A Unified Toolkit for Visualising Statistical Models and Data

Version 0.3.0

Description A cohesive, publication-ready toolkit of plots that span the whole analysis workflow with one consistent look. It covers exploratory data analysis (distributions, categorical summaries, bivariate plots, scatter-plot matrices, correlation heatmaps, missing-data maps, outliers, estimation statistics and descriptive tables); multivariate analysis, clustering with diagnostics and Kaplan-Meier survival curves; time series (trends, autocorrelation, decomposition, seasonality and forecasting); model estimates and inference (forest plots, model comparison, frequentist and Bayesian estimates, predicted values, interactions, random effects and optimiser checks); diagnostics and classification (residual panels, binned residuals, influence, quantile-quantile, receiver operating characteristic (ROC) curves, calibration, threshold tuning and confusion matrices); uncertainty and power; and reporting helpers (a shared theme, colourblind-aware palettes, plot composition and saving). Every plotting function returns a 'ggplot2' object (or a 'patchwork' object for composite panels), heavier modelling back-ends are optional, and the package ships with reproducibly simulated datasets so that every example and vignette runs without further setup.

License MIT + file LICENSE

Encoding UTF-8

Language en-GB

LazyData true

Depends R (>= 4.1.0)

Imports ggplot2 (>= 3.5.2), stringr, scales, patchwork (>= 1.3.0), rlang, stats, grDevices, utils, Rdpack

Suggests lme4 (>= 1.1-19), lmerTest, MASS, broom, simr, knitr, rmarkdown, survival, ggdist, posterior, boot, cluster, colorspace, ragg, callr, spelling, testthat (>= 3.0.0), withr

VignetteBuilder knitr
Roxygen list(markdown = TRUE)
RdMacros Rdpack
Config/testthat/edition 3
Config/Needs/website pkgdown
URL <https://pablobernabeu.github.io/depictr/>,
<https://github.com/pablobernabeu/depictr>
BugReports <https://github.com/pablobernabeu/depictr/issues>
Config/roxygen2/version 8.0.0
Config/pak/sysreqs libicu-dev
Repository <https://pablobernabeu.r-universe.dev>
Date/Publication 2026-10-16 21:58:01 UTC
RemoteUrl <https://github.com/pablobernabeu/depictr>
RemoteRef HEAD
RemoteSha 3bd402d829169450aec5d44b14725b19c5670d6b

Contents

acf_plot	4
arrange_plots	5
binned_residual_plot	6
calibration_plot	7
check_figure	8
clinical_trial	13
cluster_plot	14
coefficient_plot	16
compare_models	19
confusion_matrix_plot	21
correlation_heatmap	22
crop_yield	23
decompose_plot	24
dendrogram_plot	26
depictr_options	27
depictr_palette	28
dumbbell_plot	30
ecdf_plot	32
effects_plot	33
estimation_plot	34
explore_bivariate	37
explore_categorical	38
explore_distribution	39
explore_pairs	40
format_terms	41

frequentist_bayesian_plot	43
gain_plot	45
group_comparison_plot	46
influence_plot	48
interaction_plot	48
k_diagnostic	50
lexical_decision	51
lift_plot	52
missingness_map	53
model_fit_table	54
model_report	55
monthly_sales	57
optimizer_fixef_plot	58
outlier_plot	60
palette_preview	61
palette_safety	62
pca_plot	63
posterior_plot	65
power_curve_plot	67
pr_curve_plot	69
qq_plot	70
raincloud_plot	72
random_effects_plot	73
residual_diagnostics_plot	74
ridgeline_plot	76
roc_curve_plot	77
save_plot	78
scale_colour_depictr	81
scatter_trend	82
scree_plot	84
seasonal_plot	84
silhouette_plot	86
simulate_cvd	87
simulate_figure	88
summary_table	90
survival_plot	91
theme_depictr	94
threshold_plot	95
tidy_estimates	96
timeseries_plot	97
ts_forecast	99
vif_plot	101
wellbeing_survey	103

acf_plot

*Autocorrelation plot***Description**

Plots the autocorrelation (or partial autocorrelation) function of a series, with the approximate significance bounds, as a clean lollipop chart.

Usage

```
acf_plot(
  x,
  lag_max = NULL,
  type = c("correlation", "partial"),
  conf_level = 0.95,
  title = NULL,
  x_lab = "Lag",
  y_lab = NULL
)
```

Arguments

<code>x</code>	A numeric vector or ts object holding one series. A multivariate ts or a matrix with several columns is refused.
<code>lag_max</code>	Maximum lag (passed to <code>stats::acf()</code> / <code>stats::pacf()</code>).
<code>type</code>	"correlation" for the ACF or "partial" for the PACF.
<code>conf_level</code>	Confidence level for the significance bounds; a single number strictly between 0 and 1.
<code>title, x_lab, y_lab</code>	Title and axis labels.

Details

A missing value inside the series is kept in its place, and the correlations are computed from the pairs of observations that are both present (`na.action = na.pass`), with a warning. Dropping it would join the observations on either side of the gap as if they were neighbours. The Python package refuses such a series and asks for the gap to be dropped or imputed first.

Value

A `ggplot2::ggplot` object.

Examples

```
acf_plot(AirPassengers)
acf_plot(AirPassengers, type = "partial")
```

arrange_plots*Compose several plots into one figure*

Description

A thin, friendly wrapper around `patchwork::wrap_plots()` that adds the common finishing touches: collecting duplicate legends into one, an overall title and subtitle, and automatic panel tags (A, B, C, ...).

Usage

```
arrange_plots(  
  ...,  
  ncol = NULL,  
  nrow = NULL,  
  guides = c("collect", "keep"),  
  title = NULL,  
  subtitle = NULL,  
  tag_levels = "A"  
)
```

Arguments

<code>...</code>	Plots to combine, or a single list of plots.
<code>ncol, nrow</code>	Layout dimensions (passed to patchwork).
<code>guides</code>	How to treat legends: "collect" (merge duplicates) or "keep".
<code>title, subtitle</code>	Overall title and subtitle.
<code>tag_levels</code>	Panel tag style, e.g. "A", "1" or "i"; NULL for no tags.

Value

A 'patchwork' object.

Examples

```
p1 <- explore_distribution(crop_yield, yield)  
p2 <- scatter_trend(crop_yield, fertiliser, yield)  
arrange_plots(p1, p2, ncol = 2, title = "Crop yield", tag_levels = "A")
```

binned_residual_plot *Binned-residual plot for a generalised linear model*

Description

For a fitted glm (especially binary or count models), a plot of raw residuals against fitted values is hard to read because the residuals take only a few values. Following Gelman and Hill (2007), this plot instead splits the data into equal-count bins of fitted values and plots, for each bin, the mean residual against the mean fitted value, inside a grey band of ± 2 standard errors of that mean. Under a well-specified model, the binned residuals scatter around zero, and about 95% of them should fall within the band. Systematic departures (a trend, or many points outside the band) indicate a misspecified mean structure.

Usage

```
binned_residual_plot(
  model,
  bins = NULL,
  type = c("response", "pearson"),
  point_colour = depictr_brand(),
  band_colour = depictr_accent(),
  title = NULL,
  band = c("model", "empirical")
)
```

Arguments

model	A fitted glm object (an lm is also accepted and treated as a Gaussian GLM).
bins	Number of bins. The default follows Gelman and Hill's rule of thumb of roughly $\sqrt{n}$ bins (bounded to a sensible range).
type	Residual type used for the bin means: "response" (observed minus fitted on the response scale, the Gelman-Hill default) or "pearson".
point_colour, band_colour	Colours for the bin points and the ± 2 SE band. Default to the depictr brand and accent colours.
title	Plot title.
band	Where the standard errors come from. The default, "model", takes them from the variance the fitted model implies for each residual. "empirical" takes them from the spread of the residuals in each bin, as in depictr 0.3.0 and arm: :binnedplot().

Details

By default (band = "model"), the standard error is the one the fitted model implies. The model gives the response y_i the variance $\phi V(\hat{\mu}_i)/w_i$, where V is the family's variance function, w_i the prior weight and ϕ the dispersion. The dispersion is 1 for the binomial and Poisson families, `summary(model)$dispersion` for the others and $\hat{\sigma}^2$ for an lm. Taking this as the variance of

the response residual $y_i - \hat{\mu}_i$, the band gives the mean of the n_b residuals in bin b the standard error $\sqrt{\sum_{i \in b} \phi V(\hat{\mu}_i) / w_i / n_b}$. A Pearson residual, $(y_i - \hat{\mu}_i) \sqrt{w_i / V(\hat{\mu}_i)}$, then has variance ϕ , so its bin mean has standard error $\sqrt{\phi / n_b}$. The fit absorbs a little of each observation's variance, so a residual varies slightly less than this and the band errs on the wide side. Observations with zero prior weight take no part in the fit and are left out of the bins.

`band = "empirical"` draws the band of `depictr` 0.3.0 and earlier, the one `arm::binnedplot()` draws: twice the standard deviation of the residuals in the bin, divided by $\sqrt{n_b}$. It collapses wherever the outcomes in a bin are all alike. The residuals in a logistic model's bin without events all lie just below zero, so its band nearly vanishes and the bin is flagged. Where events are rare, such bins are what the model expects. The example below, a correctly specified model, has five of its seventeen bins flagged by that band and none by the model's.

Value

A `ggplot2::ggplot` object. The per-bin summary (mean fitted value, mean residual, number of observations n , the band's half-width `se2` and an outside flag) is attached as `attr(plot, "bins")`.

References

Gelman A, Hill J (2007). *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press, Cambridge, UK. ISBN 978-0-521-68689-1. doi:[10.1017/CBO9780511790942](https://doi.org/10.1017/CBO9780511790942).

Examples

```
gfit <- glm(adverse_event ~ biomarker + age + arm,
            data = clinical_trial, family = binomial)
binned_residual_plot(gfit)

# The band of depictr 0.3.0, which flags the bins that hold no events:
binned_residual_plot(gfit, band = "empirical")
```

calibration_plot

Calibration plot

Description

Assesses how well predicted probabilities match observed frequencies. The scores are split into bins; for each bin the mean predicted probability is plotted against the observed event rate, with the diagonal marking perfect calibration. A Wilson binomial confidence interval is drawn on each bin's observed proportion so that bins backed by few observations are not over-interpreted. Pass a *named list* of models / (actual, score) pairs to overlay several colour-coded calibration curves with a legend.

Usage

```
calibration_plot(
  x,
  score = NULL,
  bins = 10,
  colour = depictr_brand(),
  conf_level = 0.95,
  title = NULL
)
```

Arguments

x	A binomial glm, the vector of observed outcomes, or a <i>named</i> list of models / (actual, score) pairs to overlay.
score	When x is an outcome vector, the matching predicted probabilities (or a list of them for the multi-model case).
bins	Number of (equal-count) bins.
colour	Point/line colour for the single-model case. Defaults to the depictr brand blue. Ignored when several models are overlaid.
conf_level	Confidence level for the per-bin Wilson interval on the observed proportion, a single number strictly between 0 and 1. Use NA to omit the intervals. Intervals are only drawn in the single-model case to keep the overlay legible.
title	Plot title.

Value

A `ggplot2::ggplot` object.

Examples

```
# Calibration is judged against the base rate, so the example uses the rare
# clinical-trial adverse event (about 10% positive).
gfit <- glm(adverse_event ~ biomarker + age + arm,
            data = clinical_trial, family = binomial)
calibration_plot(gfit, bins = 6)
```

check_figure

Audit a finished figure for accessibility and honesty

Description

Checks a figure as it will be submitted, rather than the palette it was built from. `palette_safety()` can promise that the eight colours depictr ships stay apart under colour-vision deficiency, but it knows nothing about the figure in front of you: how many of those colours it uses, what you replaced them with, how small the text will be once the figure is squeezed into a journal column, or whether the only thing separating two groups is their colour. `check_figure()` reads the built plot and answers those questions with numbers.

Usage

```
check_figure(
  plot,
  width_cm = 17.78,
  render_width_cm = 17.78,
  min_delta_e = 5,
  min_text_pt = 6,
  metric = c("ciede2000", "cie76"),
  max_text_pt = NULL
)
```

Arguments

plot	A plot, as returned by any <code>depictr</code> plotting function, including one extended afterwards with <code>+</code> , or a multi-panel composite (see "Composites").
width_cm	The width, in centimetres, that the figure will occupy in the finished document. Defaults to 17.78 cm, the seven inches <code>save_plot()</code> draws at, which means no scaling.
render_width_cm	The width, in centimetres, that the figure is drawn at. Defaults to the same 17.78 cm. The ratio of the two widths is the factor every point size is multiplied by.
min_delta_e	The smallest acceptable colour difference, on the scale of <code>metric</code> , for the colour separability checks, and the smallest acceptable difference in lightness for the greyscale check. Defaults to 5, matching <code>palette_safety()</code> .
min_text_pt	The smallest acceptable printed text size, in points, for the text around the panel and the text inside it alike. Defaults to 6, below the floor most journals set (see "Auditing at the printed size").
metric	The colour difference the colour separability checks measure: "ciede2000", the default, or "cie76", the Euclidean distance in CIELAB that earlier versions measured. "cie76" is kept so that earlier results can be reproduced, and a later release will remove it.
max_text_pt	The largest acceptable printed text size, in points, or NULL, the default, for no ceiling. When given, a <code>text_size_max</code> row is appended (see "What each check measures").

Details

Every row carries the value it measured next to the threshold it was measured against, so a verdict can be argued with rather than merely accepted. A check passes when the measured value is at least the threshold. The one exception is `text_size_max`, a ceiling added when `max_text_pt` is given, which passes when the measured value is at most its threshold. A check that has nothing to measure, such as colour separability on a figure that encodes nothing by colour, reports NA and a verdict of "not applicable" instead of a free pass.

Value

A data frame with one row per check and columns `check`, `measured`, `threshold`, `verdict` ("pass", "fail" or "not applicable") and `detail`, a short note naming what produced the measurement.

There are ten rows, and an eleventh, `text_size_max`, when `max_text_pt` is given. For a composite, the rows of each panel in turn, after a leading integer column panel.

What each check measures

Colours and text are measured as drawn. A translucent colour is composited over the background it is drawn on with the alpha it is drawn with, so it is measured as the paler colour the reader sees. As `ggplot2` draws them, bars, tiles, areas, densities, ribbons, polygons, violins, boxplots, crossbars and smooths apply a layer's alpha to their fill and leave the outline opaque. Lines, paths, segments, error bars and reference lines apply it to their colour, and points apply it to both. Any other geom is taken to apply it to both, the cautious assumption. A colour that carries its own alpha keeps it when the layer sets none, as `ggplot2` draws it. A colour drawn at several alphas within one layer is measured at the smallest. Where marks overlap one another, the colour that shows depends on what lies beneath it, and that is not measured.

`colour_separability` is the smallest colour difference between two of the figure's encoding colours, by the CIEDE2000 formula (Sharma, Wu & Dalal, 2005) unless `metric` asks for the older CIE76. The three `colour_separability_*` rows repeat that measurement after simulating each dichromacy at full severity with `simulate_cvd()`. Encoding colours are the distinct colour and fill values that tell groups apart. A layer contributes its colours when it maps colour or fill to a discrete scale, even if it draws a single colour, or when it draws two or more. Groups drawn one per layer are therefore audited as they would be in one layer. A colour set as a constant, as in `geom_point(colour = "red")`, is left out, even when another layer sets a different one, and so is a lone colour passed through `ggplot2::scale_colour_identity()`. The audit cannot tell such a colour from decoration, so a colour meant to tell groups apart should be mapped, with `ggplot2::scale_colour_manual()` to fix its values. The outline colours, across every layer, are compared among themselves, and so are the fill colours, and the row reports the closest pair from either set. An outline is therefore never compared with its own fill, nor a group with itself as another layer draws it.

A continuous colour or fill scale, such as that of a heatmap, is a smooth ramp rather than a set of codes, and none of the colour rows audits it. When nothing else in the figure encodes by colour, each of them abstains with a detail that says so. Whether a ramp reads in order, to a reader with a colour-vision deficiency and in greyscale as well, depends on how evenly its lightness changes, which the audit does not measure. Prefer a perceptually ordered map, such as `viridis` or `depictr`'s sequential and diverging ramps, over a rainbow (Crameri et al., 2020).

`greyscale_separability` is the smallest difference in CIE lightness between those same colours, which is what survives printing in black and white. It is the absolute difference in lightness, whichever `metric` the colour rows use.

`text_size` is the smallest point size of the text the theme draws around the panel: axis labels and titles, legends, facet strips, titles and captions. It is scaled by `width_cm / render_width_cm`, so a figure drawn seven inches wide and printed in an 8.9 cm column has every point size halved. Text drawn inside the panel is measured by `annotation_text_size`.

`text_contrast` is the WCAG 2.2 contrast ratio of the text that falls furthest short of its own threshold. Each text is measured against the background it is drawn on: a facet strip's fill for strip text, the legend box's fill for legend text, and the plot background for the rest. A translucent fill is measured as the colour it shows over what lies beneath it, and so is text drawn with an alpha. Text that is large at its printed size, 18 pt or 14 pt in bold, is held to 3:1, and other text to 4.5:1, as success criterion

1.4.3 asks. The row's threshold is the one that applies to the text it reports, and its detail names that text's colour, background, printed size and weight.

`geometry_contrast` is the smallest WCAG 2.2 contrast ratio of any encoded group's drawn colour against the panel background, or against the plot background where the panel paints none. Where a layer maps its outline colour to the same groups as its fill, each group counts the better of the two, since WCAG lets a border carry a graphical object's contrast. For a translucent colour, the detail names the colour, its alpha and the colour it is drawn as.

`redundant_encoding` counts how many of shape and line type also vary across the layers that draw the encoding colours. They are gathered as the colours are. A layer contributes the values it maps to a discrete scale, or the two or more it draws, and a shape or line type set as a constant is left out. Zero means the distinction between groups is carried by colour alone, which is the single most common way an otherwise careful figure becomes unreadable.

`annotation_text_size` is the smallest point size of the text drawn inside the panel, scaled as `text_size` is and held to the same `min_text_pt`. That text comes from layers such as `ggplot2::geom_text()` and `ggplot2::geom_label()`, from annotations, and from `depictr`'s own cell values and labels. Its contrast is not measured, since it sits on the marks rather than on one background. A label that a layer builds only when the figure is drawn, as `ggrepel` builds its repelled labels, is not yet in the built plot, so it is not measured either.

`text_size_max` is reported only when `max_text_pt` is given, as the last row. It is the largest point size of all the text measured, around the panel and inside it, scaled as `text_size` is. It is the one row that passes at or below its threshold.

Limitations of the default palette

The eight-colour qualitative palette clears the colour-separability checks comfortably and fails `greyscale_separability`: its orange (`#e69f00`) and sky blue (`#56b4e9`) differ by only 0.79 in lightness, so they print as the same grey. The colourblind-safety guarantee `depictr` makes is about hue confusion, and it was never a claim about greyscale. A figure that may be printed in black and white should use fewer groups, or a sequential palette, or add a redundant shape or line type; the four leading colours of the palette are also not safe in greyscale, since the bluish green and the vermillion differ by 3.55. The threshold has been left where it is rather than moved to let the package's own defaults through.

The palette falls short on contrast as well. Four of its eight colours are below the 3:1 against white that WCAG asks of a graphical object: the orange at 2.25:1, the sky blue at 2.31:1, the grey at 2.85:1 and the yellow at 1.32:1. The orange is the second colour, so a default figure on white whose colour tells two or more groups apart fails `geometry_contrast`, and drawing its marks translucent lowers the ratio further. Opaque marks, the vermillion (`#d55e00`, at 3.87:1) in the orange's place and a redundant shape clear that row and `redundant_encoding`, as the first example below shows. The contrast-first order of `depictr_palette()` puts the vermillion second and the reddish purple third, and `scale_colour_depict()` takes it as `order = "contrast"`. Up to three groups then pass every colour row, and four keep 3:1 against white.

Auditing at the printed size

Text is drawn in points, so the size it prints at depends on the two widths. A figure drawn at the size it will be printed, as `save_plot()` draws one given `width = 8.9`, `units = "cm"`, needs no scaling. Pass the saved width as both `width_cm` and `render_width_cm`, or leave both at their

defaults. Given alone, `width_cm` audits the figure as though it had been drawn seven inches wide and then shrunk to fit. `save_plot()` with `check = TRUE` audits a figure at the width it saves it at, given as both.

The default `min_text_pt` of 6 is lenient. As their guidelines stood on 30 September 2026, PLOS asked for text of 8 to 12 pt and APA Style for 8 to 14 pt. Elsevier asked for 7 pt as a rule of thumb, and Nature for 5 to 7 pt, with panel labels at 8 pt. Pass the floor of the journal the figure is for as `min_text_pt`, and its ceiling, where it sets one, as `max_text_pt`. Nature's 7 pt applies to all text but the panel labels. The tags of a composite, where patchwork puts them, are not measured, but a panel label drawn as a panel's title or annotation is measured with the rest. The four pages are listed under References.

Composites

A composite, such as `arrange_plots()`, `residual_diagnostics_plot()` and `model_report()` return, is audited panel by panel. Each panel has its own scales, theme and text, so each is audited as it would be on its own. The reports are stacked under a leading integer column, `panel`, which only the report of a composite has. The widths are those of the whole composite. A panel's text keeps its point size whatever share of the width the panel takes, so the text of every panel scales by the same ratio.

Panels are numbered depth-first through any composite nested inside, in the order the composite holds them. That is the order patchwork tags them in, and the reading order of a grid filled row by row. A spacer or a guide area draws no figure of its own and takes no number. A wrapped element, such as a grob or a table, cannot be read, so it keeps its number and is left out of the report with a warning.

A panel's legend is measured with the panel, as the panel would draw it on its own, even when `guides = "collect"` gathers it into one legend for the composite. What the composite draws around its panels is not measured: the title, subtitle and caption of `patchwork::plot_annotation()` and the panel tags.

References

- APA Style. Figures. <https://apastyle.apa.org/style-grammar-guidelines/tables-figures/figures>
- Crameri F, Shephard GE, Heron PJ (2020). "The misuse of colour in science communication." *Nature Communications*, **11**(1), 5444. doi:10.1038/s41467020191607.
- Elsevier. Artwork sizing. <https://www.elsevier.com/about/policies-and-standards/author/artwork-and-media-instructions/artwork-sizing>
- Machado GM, Oliveira MM, Fernandes LAF (2009). "A physiologically-based model for simulation of color vision deficiency." *IEEE Transactions on Visualization and Computer Graphics*, **15**(6), 1291–1298. doi:10.1109/TVCG.2009.113.
- Nature. Guide to preparing final artwork. <https://www.nature.com/documents/nature-final-artwork.pdf>
- PLOS ONE. Figures. <https://journals.plos.org/plosone/s/figures>
- Sharma G, Wu W, Dalal EN (2005). "The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations." *Color Research & Application*, **30**(1), 21–30. doi:10.1002/col.20070.

World Wide Web Consortium (2023). “Web Content Accessibility Guidelines (WCAG) 2.2.” W3C Recommendation, <https://www.w3.org/TR/WCAG22/>. Accessed 2026-08-18.

See Also

`palette_safety()` for the palette in the abstract, and `palette_preview()` to look at it.

Examples

```
library(ggplot2)

# A figure that clears every check it can be held to: two well-separated
# colours, a redundant shape, and text left at the size it was drawn. It
# draws no text inside the panel, so that row is not applicable.
good <- ggplot(crop_yield, aes(rainfall, yield, colour = treatment,
                              shape = treatment)) +
  geom_point() +
  scale_colour_manual(values = c("#005b96", "#d55e00")) +
  theme_depict()
check_figure(good)

# The same figure destined for an 8.9 cm journal column, where the text is
# half the size it looks on screen.
check_figure(good, width_cm = 8.9)

# Saved at 8.9 cm in the first place, with save_plot(width = 8.9,
# units = "cm"), it is drawn at the size it prints, so both widths are 8.9.
check_figure(good, width_cm = 8.9, render_width_cm = 8.9)

# Nature allowed no text above 7 pt but the panel labels. The axis titles
# here are 11 pt, so the ceiling row fails.
subset(check_figure(good, max_text_pt = 7), check == "text_size_max")

# Each panel of a composite is audited, and a column says which panel each
# row describes. At base size 5, the second panel's text is too small.
small <- good + theme_depict(base_size = 5)
both <- check_figure(enumerate_plots(good, small))
both[both$verdict == "fail", c("panel", "check", "measured")]
```

clinical_trial

Simulated two-arm clinical trial

Description

A reproducibly simulated randomised clinical trial, for time-to-event (survival) and imbalanced-classification examples. The two arms have a real survival difference: the treatment arm has a lower hazard, so the Kaplan-Meier curves separate clearly and the log-rank test is highly significant, and - because its event times are longer - the treatment arm is more often right-censored by the 36-month study end. age and biomarker also affect the hazard. The data feed `survival_plot()` (pass time, event and arm, or a `survival::survfit` object).

Usage

```
clinical_trial
```

Format

A data frame with 300 rows and 7 variables:

patient Patient identifier.

arm Treatment arm (factor): placebo or treatment.

age Age in years at enrolment.

biomarker Standardised baseline laboratory value.

time Follow-up time, in months.

event Event indicator: 1 event observed, 0 right-censored.

adverse_event Rare adverse-event indicator: 1 yes, 0 no (about 10% positive).

Details

adverse_event is a deliberately rare binary safety outcome (about a 10% base rate) that is predictable from arm, age and biomarker. It gives the precision-recall ([pr_curve_plot\(\)](#)), cumulative-gains ([gain_plot\(\)](#)), lift ([lift_plot\(\)](#)) and calibration ([calibration_plot\(\)](#)) demos an imbalanced target, where those charts are most informative.

The data are synthetic, generated by data-raw/generate_datasets.R with a fixed seed; they do not describe any real individuals.

Source

Simulated; see data-raw/generate_datasets.R.

cluster_plot

Cluster scatter plot

Description

Runs k-means clustering on the numeric columns of a data frame (or uses a supplied cluster assignment) and plots the observations on the first two principal components, coloured by cluster, with optional convex hulls and labelled centroids. Projecting onto principal components means the plot works for any number of variables.

Usage

```
cluster_plot(
  data,
  cols = NULL,
  k = 3,
  clusters = NULL,
  scale = TRUE,
  suggest_k = NULL,
  k_range = NULL,
  hulls = TRUE,
  label_centers = TRUE,
  point_alpha = 0.8,
  palette = NULL,
  title = NULL,
  seed = NULL,
  nstart = 10,
  iter.max = 10
)
```

Arguments

<code>data</code>	A data frame.
<code>cols</code>	Numeric columns to cluster on (default: all numeric).
<code>k</code>	Number of clusters for k-means (ignored if <code>clusters</code> is supplied, or overridden when <code>suggest_k</code> chooses a value).
<code>clusters</code>	Optional vector of cluster assignments (e.g. from <code>stats::kmeans()</code> or <code>stats::cutree()</code>); use this to plot a clustering you computed yourself. Hierarchical or PAM assignments (such as <code>stats::cutree()</code> output matching <code>dendrogram_plot()</code>) are accepted. Must have exactly one entry per row of data (rows with missing values in <code>cols</code> are dropped from both before plotting).
<code>scale</code>	Whether to scale variables to unit variance before clustering and the PCA. Columns with (near-)zero variance cannot be scaled and are dropped with a message, as in <code>correlation_heatmap()</code> .
<code>suggest_k</code>	Optionally choose <code>k</code> automatically with a cluster-quality diagnostic instead of using the supplied <code>k</code> . Either <code>TRUE</code> (uses the average-silhouette criterion), a string naming the criterion ("silhouette", "wss" or "gap"; see <code>k_diagnostic()</code>), or <code>NULL</code> (the default) to leave <code>k</code> unchanged. Ignored when <code>clusters</code> is supplied.
<code>k_range</code>	Candidate values of <code>k</code> searched when <code>suggest_k</code> is set. The default, <code>NULL</code> , searches 1:8 for "gap" and "wss" and 2:8 for "silhouette", as <code>k_diagnostic()</code> does, so the gap statistic can suggest a single cluster.
<code>hulls</code>	Whether to draw a shaded convex hull around each cluster.
<code>label_centers</code>	Whether to label the cluster centroids.
<code>point_alpha</code>	Point transparency.
<code>palette</code>	Colours for the clusters; defaults to <code>depictr_palette()</code> .
<code>title</code>	Plot title.

seed	Optional integer seed for reproducible k-means (ignored when clusters is supplied). Default NULL leaves the RNG state untouched.
nstart	Number of random starts for <code>stats::kmeans()</code> .
iter.max	Maximum iterations for <code>stats::kmeans()</code> .

Value

A `ggplot2::ggplot` object. When `suggest_k` is used, the chosen `k` and the full diagnostic are attached as the attribute "k_diagnostic".

See Also

`k_diagnostic()` and `silhouette_plot()` for cluster-quality diagnostics.

Examples

```
cluster_plot(crop_yield, cols = c("rainfall", "fertiliser", "soil_ph",
                                "yield"), k = 3, seed = 1)

# Let a silhouette diagnostic choose k:
cluster_plot(crop_yield, cols = c("rainfall", "fertiliser", "soil_ph",
                                "yield"), suggest_k = TRUE, k_range = 2:6,
            seed = 1)
```

coefficient_plot

Forest (coefficient) plot

Description

Draws a horizontal point-and-interval ("forest") plot of model estimates. The input can be a fitted model (anything `tidy_estimates()` understands) or a data frame of pre-computed estimates.

Usage

```
coefficient_plot(
  x,
  conf_level = 0.95,
  intercept = FALSE,
  order = c("none", "ascending", "descending"),
  labels = NULL,
  interaction = c("times", "asterisk", "colon", "space"),
  point_colour = depictr_brand(),
  reference_colour = depictr_reference(),
  reference_line = if (exponentiate) 1 else 0,
  point_size = 2.2,
  line_size = 0.7,
  facet = FALSE,
  scales = c("fixed", "free"),
```

```

    standardise = FALSE,
    sd_multiple = 1,
    exponentiate = FALSE,
    title = NULL,
    subtitle = NULL,
    x_lab = NULL
  )

```

Arguments

<code>x</code>	A fitted model or a tidy data frame of estimates.
<code>conf_level</code>	Confidence/credible level, passed to <code>tidy_estimates()</code> when <code>x</code> is a model: a single number strictly between 0 and 1.
<code>intercept</code>	Whether to keep the intercept term. Defaults to <code>FALSE</code> , since the intercept is seldom of interest on a forest plot and its scale often overwhelms the other terms.
<code>order</code>	Order of the terms. "ascending" puts the smallest estimate at the bottom and the largest at the top, so estimates increase up the axis. "descending" does the reverse, and "none" keeps model order from top to bottom. The faceted layout (<code>facet = TRUE</code>) stacks its panels the other way up.
<code>labels</code>	Optional display labels for the terms. An unnamed vector gives one label per term in model order (the order <code>tidy_estimates()</code> returns, after the intercept is dropped unless <code>intercept = TRUE</code>), whatever order is. A named vector mapping raw term names to labels, such as <code>c(soil_ph = "Soil pH")</code> , is safer. If <code>NULL</code> , the coefficient of a factor level is labelled with its variable and level, inside interaction terms too, so the contrast of enhanced with the reference level reads treatment: enhanced. Other names are tidied with <code>format_terms()</code> .
<code>interaction</code>	Passed to <code>format_terms()</code> to control how interaction terms are rendered (ignored when labels is supplied).
<code>point_colour, reference_colour</code>	Colours for the estimates and the reference line.
<code>reference_line</code>	Position of a vertical reference line. Defaults to 0, or to 1 when <code>exponentiate = TRUE</code> , where a ratio of 1 means no effect. Use <code>NA</code> to omit it.
<code>point_size, line_size</code>	Size of the points and interval lines.
<code>facet</code>	Whether to give each term its own panel with a free x-axis, laid out one per row. This removes the squish that occurs when terms live on very different scales (for example a large intercept alongside small slopes). Defaults to <code>FALSE</code> , preserving the shared-axis layout. A convenience alias for <code>scales = "free"</code> .
<code>scales</code>	Either "fixed" (the default, a single shared x-axis) or "free" (one free-scaled panel per term). When <code>facet = TRUE</code> this is forced to "free".
<code>standardise</code>	Whether to scale each coefficient by the spread of its inputs, following Gelman (2008). A continuous input (numeric, with more than two distinct values) is scaled by <code>sd_multiple</code> times its standard deviation in the rows the model used, so its slope becomes the change in the outcome per SD. Factor, character, logical and two-valued numeric inputs (0/1, say) are left unscaled, since a contrast already compares two groups and its per-SD value would depend on how

balanced the groups happen to be. An interaction is scaled through its inputs: `rainfall:treatment` by the SD of `rainfall`, and a product of two continuous inputs by both their SDs. Matrix-valued terms such as `poly()` and `ns()` are left unscaled, with a message. An interaction of continuous inputs draws a warning unless they are all centred. Otherwise, each input's main effect is its slope where the others are zero, and Gelman's remedy is to centre the inputs and re-fit. A continuous slope and a contrast are comparable only roughly, most nearly with `sd_multiple = 2`. Requires a fitted model (ignored, with a warning, for a tidy data frame). Defaults to `FALSE`.

<code>sd_multiple</code>	How many standard deviations of a continuous input one unit of the standardised scale spans when <code>standardise = TRUE</code> . Defaults to 1. Use 2 for Gelman's (2008) convention, under which the coefficient of a continuous input is comparable to that of a balanced binary one.
<code>exponentiate</code>	Whether to draw the exponentiated estimates and intervals on a log10 axis, which shows the coefficients of a logistic regression as odds ratios and those of a Poisson or negative binomial model as rate ratios. It applies after <code>standardise</code> , so a standardised slope becomes a ratio per SD of its predictor. A data frame <code>x</code> must hold estimates on the link scale. Defaults to <code>FALSE</code> .
<code>title, subtitle, x_lab</code>	Plot title, subtitle and x-axis label. <code>x_lab</code> defaults to "Estimate", or, when <code>standardise</code> , "Estimate per SD of predictor" ("per 2 SD" with <code>sd_multiple = 2</code>), followed by "(contrasts unscaled)" when a factor or binary term is drawn. The label names the convention on the figure itself, because only the predictors are rescaled here: a fully standardised beta would rescale the outcome too. With <code>exponentiate = TRUE</code> , "Estimate" gives way to the name of the ratio: "Odds ratio" for a logit link, "Rate ratio" for a log link in a Poisson, quasi-Poisson or negative binomial model, "Risk ratio" for a binomial model with a log link and "exp(Estimate)" otherwise. The last covers an <code>lm(log(y) ~ x)</code> slope, whose exponential is a ratio of geometric means.

Value

A `ggplot2::ggplot` object.

References

Gelman A (2008). "Scaling regression inputs by dividing by two standard deviations." *Statistics in Medicine*, 27(15), 2865–2873. doi:10.1002/sim.3107.

Examples

```
fit <- lm(yield ~ rainfall + fertiliser + soil_ph + treatment,
         data = crop_yield)
coefficient_plot(fit)

# Order terms and add a title
coefficient_plot(fit, order = "descending", title = "Drivers of crop yield")

# When an intercept or large term squishes the rest, give each term its own
```

```
# free-scaled panel:
coefficient_plot(fit, intercept = TRUE, facet = TRUE)

# Or scale the continuous slopes per SD of their predictors, leaving the
# treatment contrast in the outcome's units:
coefficient_plot(fit, standardise = TRUE)

# Odds ratios from a logistic regression, on a log axis with the line at 1:
gfit <- glm(accuracy ~ word_frequency + condition,
            data = lexical_decision, family = binomial)
coefficient_plot(gfit, exponentiate = TRUE)
```

compare_models

Compare estimates from several models or sources

Description

Overlays the estimates from two or more models (or tidy estimate tables) on a single forest plot, with one colour per source. It is the general engine behind [frequentist_bayesian_plot\(\)](#) (frequentist against Bayesian), and serves equally well for comparing nested models, several optimisers, or estimates before and after a transformation.

Usage

```
compare_models(
  ...,
  names = NULL,
  conf_level = 0.95,
  intercept = FALSE,
  order = c("none", "ascending", "descending"),
  labels = NULL,
  interaction = c("times", "asterisk", "colon", "space"),
  dodge_width = 0.6,
  reference_line = if (exponentiate) 1 else 0,
  palette = NULL,
  point_size = 2.2,
  line_size = 0.7,
  facet = FALSE,
  scales = c("fixed", "free"),
  exponentiate = FALSE,
  legend_title = "Source",
  legend_ncol = 1,
  title = NULL,
  subtitle = NULL,
  x_lab = NULL
)
```

Arguments

...	Two or more fitted models and/or tidy data frames of estimates. Name the arguments to label the sources (e.g. <code>compare_models(Frequentist = m1, Bayesian = m2)</code>).
names	Optional character vector of source labels, overriding the names of ...
conf_level	Confidence/credible level for models, a single number strictly between 0 and 1.
intercept	Whether to keep the intercept term. Defaults to FALSE.
order	Order of the terms, by their average estimate across sources. "ascending" puts the smallest estimate at the bottom and the largest at the top, so estimates increase up the axis. "descending" does the reverse, and "none" keeps the terms in the order they first appear, from top to bottom. The faceted layout (<code>facet = TRUE</code>) stacks its panels the other way up.
labels	Optional display labels. An unnamed vector gives one label per term, whatever order is. The terms are taken in the order they first appear across the sources, each source in the order <code>tidy_estimates()</code> returns, after the intercept is dropped unless <code>intercept = TRUE</code> . Terms drawn on one row, such as <code>stress</code> and <code>brms's b_stress</code> , take one label. A named vector mapping raw term names to labels, such as <code>c(soil_ph = "Soil pH")</code> , is safer. If NULL, the terms are labelled as in <code>coefficient_plot()</code> , each factor level with its variable and level (<code>treatment: enhanced</code>), from every fitted model among the sources.
interaction	Passed to <code>format_terms()</code> .
dodge_width	Vertical spacing between sources sharing a term.
reference_line	Position of a vertical reference line. Defaults to 0, or to 1 when <code>exponentiate = TRUE</code> . Use NA to omit it.
palette	Colours for the sources; defaults to <code>depictr_palette()</code> .
point_size, line_size	Point and interval-line sizes.
facet	Whether to give each term its own panel with a free x-axis, laid out one per row, so that terms on very different scales (for example a large intercept alongside small slopes) stay legible. The source overlay and its single legend are preserved within the faceted layout. Defaults to FALSE. A convenience alias for <code>scales = "free"</code> .
scales	Either "fixed" (the default, a single shared x-axis) or "free" (one free-scaled panel per term). When <code>facet = TRUE</code> this is forced to "free".
exponentiate	Whether to draw every source's exponentiated estimates and intervals on a log10 axis, as odds or rate ratios. Each model is exponentiated by <code>tidy_estimates()</code> , and a data frame among the sources must hold estimates on the link scale, since it is exponentiated too. Defaults to FALSE.
legend_title, legend_ncol	Legend title and number of columns.
title, subtitle, x_lab	Title, subtitle and x-axis label. <code>x_lab</code> defaults to "Estimate". With <code>exponentiate = TRUE</code> it defaults to the name of the ratio, as in <code>coefficient_plot()</code> , when the fitted models agree on it, and to "exp(Estimate)" when they do not.

Value

A `ggplot2::ggplot` object.

Examples

```
m1 <- lm(yield ~ rainfall + fertiliser + soil_ph, data = crop_yield)
m2 <- lm(yield ~ rainfall + fertiliser + soil_ph,
        data = crop_yield[crop_yield$treatment == "standard", ])
compare_models(`All fields` = m1, `Standard only` = m2,
              title = "Estimates by subset")

# Keep the intercept legible alongside the slopes:
compare_models(`All fields` = m1, `Standard only` = m2,
              intercept = TRUE, facet = TRUE)

# Odds ratios from two logistic regressions:
g1 <- glm(accuracy ~ word_frequency + condition,
        data = lexical_decision, family = binomial)
g2 <- glm(accuracy ~ word_frequency + condition + modality,
        data = lexical_decision, family = binomial)
compare_models(Reduced = g1, Full = g2, exponentiate = TRUE)
```

`confusion_matrix_plot` *Confusion matrix heatmap*

Description

Cross-tabulates predicted against actual classes and displays the counts as a heatmap. The input can be a fitted binomial glm (with a probability threshold) or a pair of vectors of actual and predicted classes. Each count is drawn in black or white, whichever has the higher contrast against its cell.

Usage

```
confusion_matrix_plot(
  x,
  predicted = NULL,
  threshold = 0.5,
  normalise = c("none", "row", "col"),
  title = NULL
)
```

Arguments

<code>x</code>	A binomial glm, or the vector of actual classes.
<code>predicted</code>	When <code>x</code> is an actual-class vector, the matching predicted classes.
<code>threshold</code>	When <code>x</code> is a glm, the probability threshold for the positive class. As well as a number in $[0, 1]$, you may pass the string "youden" to reuse the Youden's J optimal threshold (the same operating point <code>roc_curve_plot()</code> marks), so the confusion matrix and the ROC curve agree on the cut-off.

normalise	One of "none", "row" (by actual class) or "col" (by predicted class); controls the fill shading and the cell annotation.
title	Plot title.

Value

A `ggplot2::ggplot` object. The threshold actually used is stored in `attr(plot, "threshold")`.

Examples

```
gfit <- glm(accuracy ~ word_frequency + RT + condition,
           data = lexical_decision, family = binomial)
confusion_matrix_plot(gfit, threshold = 0.5)

# Reuse the Youden-optimal operating point.
confusion_matrix_plot(gfit, threshold = "youden")
```

correlation_heatmap	<i>Plot a correlation matrix as a heatmap</i>
---------------------	---

Description

Computes the pairwise correlations between the numeric columns of a data frame and displays them as a colour-coded heatmap, optionally annotated with the correlation values.

Usage

```
correlation_heatmap(
  data,
  cols = NULL,
  method = "pearson",
  use = "pairwise.complete.obs",
  show_values = TRUE,
  digits = 2,
  palette = depictr_palette(5, "diverging")[c(1, 3, 5)],
  reorder = FALSE,
  title = NULL
)
```

Arguments

data	A data frame.
cols	Optional character vector of columns to include, each numeric or logical. Any other column, such as a factor or a date, is an error that names it. If NULL, all numeric columns are used.
method	Correlation method: "pearson", "spearman" or "kendall".
use	Missing-value handling passed to <code>stats::cor()</code> .

show_values	Whether to annotate each cell with its correlation, drawn in black or white, whichever has the higher contrast against the cell.
digits	Number of decimal places for the annotations.
palette	Length-3 vector of colours for the lowest, mid (zero) and highest correlations. Defaults to the endpoints and midpoint of the colourblind-aware <code>depictr_palette()</code> diverging ramp (negative correlations red, zero neutral, positive correlations brand blue).
reorder	Whether to reorder the variables by hierarchical clustering of the correlation matrix (using $1 - r$ as the distance), so that blocks of mutually correlated variables sit together and structure is easier to see. Defaults to FALSE (the order of cols). Skipped with a message if any correlation is undefined (NA).
title	Plot title.

Details

Columns with (near-)zero variance cannot be correlated and are dropped automatically with an informative message, so the raw "the standard deviation is zero" warning from `stats::cor()` is not surfaced. If, after dropping them, any cells still come out NA (e.g. two variables that never co-occur under "pairwise.complete.obs"), those cells are rendered in grey and labelled n/a rather than left blank.

Value

A `ggplot2::ggplot` object.

Examples

```
correlation_heatmap(wellbeing_survey)
correlation_heatmap(crop_yield, method = "spearman", show_values = FALSE)
# Cluster correlated variables together:
correlation_heatmap(wellbeing_survey, reorder = TRUE)
```

crop_yield	<i>Simulated crop-yield field trial</i>
------------	---

Description

A reproducibly simulated agronomy field trial relating crop yield to rainfall, fertiliser, soil pH and a management treatment. Used for the regression, coefficient-plot, scatter-trend and interaction examples. The data-generating process contains a genuine fertiliser-by-treatment interaction: fertiliser raises yield far more under the enhanced treatment than under standard, so `interaction_plot(lm(yield ~ fertiliser * treatment + ..., crop_yield), "fertiliser", "treatment")` shows real, diverging slopes rather than noise. The main effects (rainfall, fertiliser, soil pH and treatment) are retained.

Usage

```
crop_yield
```

Format

A data frame with 200 rows and 6 variables:

field Field identifier.

treatment Management treatment (factor): standard or enhanced.

rainfall Seasonal rainfall, in millimetres.

fertiliser Fertiliser applied, in kilograms per hectare.

soil_ph Soil pH.

yield Crop yield, in tonnes per hectare.

Details

The data are synthetic, generated by `data-row/generate_datasets.R` with a fixed seed.

Source

Simulated; see `data-row/generate_datasets.R`.

decompose_plot	<i>Time-series decomposition plot</i>
----------------	---------------------------------------

Description

Decomposes a seasonal time series into trend, seasonal and remainder components and shows them, with the original series, as stacked panels.

Usage

```
decompose_plot(
  x,
  frequency = NULL,
  method = c("stl", "classical"),
  robust = FALSE,
  s_window = "periodic",
  confidence = FALSE,
  level = 0.95,
  title = NULL
)
```

Arguments

<code>x</code>	A <code>ts</code> object, or a numeric vector (then frequency is required), holding one series. A multivariate <code>ts</code> or a matrix with several columns is refused. Missing values at either end are trimmed, and a missing value inside the series is an error.
<code>frequency</code>	Number of observations per period (e.g. 12 for monthly data); taken from <code>x</code> when it is a <code>ts</code> .

method	"stl" (loess-based, Cleveland et al., 1990) or "classical" (<code>stats::decompose()</code>). STL needs more than two full cycles of data, and classical decomposition at least two.
robust	For method = "stl", whether to fit a robust STL (<code>stats::stl()</code> with <code>robust = TRUE</code>), which down-weights outliers in the loess fits so that an unusual point bleeds less into the trend and seasonal components. Ignored for the classical method.
s_window	For method = "stl", the seasonal window, passed to <code>stats::stl()</code> as <code>s.window</code> . The default, "periodic", estimates one seasonal value per position in the cycle and repeats it in every cycle, as classical decomposition does. An odd whole number of at least 7 lets the seasonal component change from cycle to cycle. It is the number of consecutive cycles that go into each seasonal value, so a smaller window lets the pattern change faster. Ignored for the classical method.
confidence	Whether to draw a confidence ribbon around the trend component. The band is a normal-approximation interval based on the remainder's standard deviation (see Details).
level	Coverage of the trend confidence ribbon (a single number strictly between 0 and 1).
title	Plot title.

Details

The trend confidence ribbon is a pragmatic normal-approximation band: $\text{trend} \pm z * \text{sd}(\text{remainder})$, where $z = \text{qnorm}((1 + \text{level}) / 2)$ and the remainder standard deviation is computed on the non-missing remainder values. It conveys the scale of the unexplained variation around the smoothed trend rather than a formal sampling-distribution interval.

Value

A 'patchwork' object (printable like a `ggplot2::ggplot`).

References

Cleveland RB, Cleveland WS, McRae JE, Terpenning I (1990). "STL: A seasonal-trend decomposition procedure based on loess." *Journal of Official Statistics*, **6**(1), 3–73.

See Also

`ts_forecast()`, `seasonal_plot()`

Examples

```
decompose_plot(AirPassengers)

decompose_plot(AirPassengers, method = "classical")
# Robust STL with a confidence ribbon on the trend
decompose_plot(AirPassengers, robust = TRUE, confidence = TRUE)
# A seasonal component free to grow with the level of the series
decompose_plot(AirPassengers, s_window = 13)
```

dendrogram_plot *Dendrogram*

Description

Hierarchical-clustering dendrogram, drawn with ggplot2 so it shares the package theme. Accepts a data frame (the distance matrix and clustering are computed for you), a `stats::dist` object, or an `stats::hclust` object. Optionally cut the tree into k clusters, colouring the leaf labels and drawing the cut height.

Usage

```
dendrogram_plot(
  x,
  cols = NULL,
  distance = "euclidean",
  method = "complete",
  scale = TRUE,
  k = NULL,
  horizontal = FALSE,
  labels = NULL,
  palette = NULL,
  title = NULL
)
```

Arguments

<code>x</code>	A data frame, a <code>dist</code> object, or an <code>hclust</code> object.
<code>cols</code>	When <code>x</code> is a data frame, the numeric columns to use.
<code>distance</code>	Distance measure passed to <code>stats::dist()</code> .
<code>method</code>	Linkage method passed to <code>stats::hclust()</code> .
<code>scale</code>	Whether to scale variables before computing distances (data frame input).
<code>k</code>	Optional number of clusters to highlight (between 1 and the number of leaves). The cut-height line is drawn only when $2 \leq k < n$.
<code>horizontal</code>	Whether to draw the tree horizontally.
<code>labels</code>	Whether to print the leaf labels. <code>NULL</code> (the default) chooses automatically: labels are shown for small trees (up to 40 leaves) and suppressed for larger ones, where they would otherwise collapse into an unreadable smear. <code>TRUE/FALSE</code> force them on or off. When labels are hidden and <code>k</code> is set, the cluster membership is instead conveyed by a coloured strip of leaf ticks along the bottom of the tree.
<code>palette</code>	Colours for the <code>k</code> clusters; defaults to <code>depictr_palette()</code> .
<code>title</code>	Plot title.

Value

A `ggplot2::ggplot` object.

Examples

```
# Cluster the survey regions by their wellbeing averages
d <- aggregate(cbind(stress, sleep_hours, life_satisfaction) ~ region,
               data = wellbeing_survey, FUN = mean)
rownames(d) <- d$region
dendrogram_plot(d[-1], k = 2)
```

depictr_options	<i>Get or set the depictr look-and-feel options</i>
-----------------	---

Description

depictr reads a small set of global options so that you can configure the shared look of every plot once, at the top of a script or in your `.Rprofile`, instead of passing the same arguments to each function. The options are honoured by `theme_depictr()` (base size, base family and the brand colour used for titles), by `depictr_palette()` and the `scale_colour_depictr()` family (an optional custom qualitative palette), and by every plotting function that draws a single series, a highlighted element or a reference line. Those functions list `depictr_brand()`, `depictr_accent()` or `depictr_reference()` as the default of a colour argument. The three are internal and are not called directly: each reads its option when the plotting function runs, so the colours are set here.

Usage

```
depictr_options(
  base_size,
  base_family,
  brand,
  accent,
  reference,
  palette,
  na_value
)
```

Arguments

<code>base_size</code>	Base font size in points for <code>theme_depictr()</code> .
<code>base_family</code>	Base font family for <code>theme_depictr()</code> .
<code>brand</code>	The depictr brand colour, used for plot titles and single-series geoms. It coincides with the default palette's first colour but does not alter a palette; use <code>palette</code> for that.
<code>accent</code>	A secondary highlight colour, for the one element a figure draws attention to.
<code>reference</code>	The colour used for reference / annotation lines.

palette	An optional custom qualitative palette: a character vector of hex colours used by <code>depictr_palette()</code> (type "qualitative") and the discrete scales in place of the built-in Okabe-Ito set. NULL restores the built-in palette. <code>depictr_palette(order = "contrast")</code> here draws every later plot in the contrast-first order (see <code>depictr_palette()</code>).
na_value	Colour used for NA levels by <code>scale_colour_depict()</code> and <code>scale_fill_depict()</code> .

Details

Called with no arguments, `depictr_options()` returns the currently resolved values (option if set, otherwise package default). Called with named arguments it sets the matching options(`depictr.<name> =`) entries and returns the *previous* resolved values invisibly, so the pattern `old <- depictr_options(...); on.exit(do.call` restores them. Pass NULL for an argument to clear that option and fall back to the package default.

Value

A named list of the resolved option values. When setting, the *previous* values are returned invisibly.

Examples

```
# Inspect the current settings
depictr_options()

# Set a larger base size and a different accent, then restore
old <- depictr_options(base_size = 14, accent = "#e69f00")
theme_depict()           # now uses base_size 14
do.call(depictr_options, old)

# Use a custom qualitative palette everywhere
old <- depictr_options(palette = c("#1b9e77", "#d95f02", "#7570b3"))
depictr_palette(3)
do.call(depictr_options, old)

# Or the built-in palette in its contrast-first order
old <- depictr_options(palette = depictr_palette(order = "contrast"))
depictr_palette(3)
do.call(depictr_options, old)
```

depictr_palette	<i>The depictr colour palettes</i>
-----------------	------------------------------------

Description

Colourblind-aware palettes shared by every `depictr` plot. The qualitative palette is based on the Okabe-Ito set (Okabe & Ito, 2008), a widely recommended categorical palette that stays distinguishable under the common forms of colour-vision deficiency, with the `depictr` brand blue leading. The sequential and diverging palettes are perceptually ordered single-hue and red-blue ramps.

Usage

```
depictr_palette(
  n = NULL,
  type = c("qualitative", "sequential", "diverging"),
  order = c("default", "contrast")
)
```

Arguments

<code>n</code>	Number of colours to return. If <code>NULL</code> (the default), the qualitative palette returns all of its base colours, the eight Okabe-Ito colours unless <code>options(depict.palette =)</code> supplies another set, and the sequential and diverging ramps return seven. For the qualitative palette an <code>n</code> larger than the available base colours is interpolated; the sequential and diverging palettes are ramps and accept any <code>n</code> . Interpolating the built-in qualitative palette beyond its eight base colours loses the colour-vision-deficiency guarantee and warns, so facet the groups or use the sequential ramp when many groups are needed.
<code>type</code>	Palette type: "qualitative" (categorical groups), "sequential" (ordered low-to-high) or "diverging" (a midpoint with two directions).
<code>order</code>	The order of the built-in qualitative palette: "default", or "contrast" for the contrast-first order described below. The ramps have an order of their own and refuse "contrast".

Details

The guarantee belongs to the eight Okabe-Ito colours themselves, not to any number of colours: past eight the qualitative palette has to interpolate between them, and the interpolated colours sit close enough together to fail the package's own colour-distance check. That case warns rather than handing back colours under a guarantee it cannot keep.

The qualitative palette can be overridden globally with `options(depict.palette = )` (see [depictr_options\(\)](#)); when set, that custom palette replaces the built-in Okabe-Ito set and is interpolated when more colours are requested than it provides. The sequential and diverging ramps are unaffected.

Value

A character vector of hex colour codes.

The contrast-first order

In the default order the second colour is the orange (`#e69f00`), at 2.25:1 against white. Any figure on white whose colour tells two or more groups apart therefore fails the `geometry_contrast` row of [check_figure\(\)](#). No more than three of the eight colours pass every colour row of that audit together.

`order = "contrast"` puts those three first: the brand blue, the vermillion (`#d55e00`) and the reddish purple (`#cc79a7`), then the bluish green (`#009e73`) and the other four in their default order. Up to three groups then pass every colour row, contrast against white included. They stay at least 14.01 apart in CIEDE2000 under every vision type, 6.86 apart in lightness and 3.06:1 against white. Four groups keep 3:1 against white, but the bluish green and the vermillion are 3.31 apart in lightness,

which fails greyscale_separability. Asking for more than four colours warns, since the fifth, the orange, falls below 3:1.

Under this order, the second group is drawn in the vermillion. That is also the accent (`depictr_options()`) that `depictr` uses for single highlights, such as the flagged points of `outlier_plot()` and the missing cells of `missingness_map()`. The default order is unchanged.

The order applies to the built-in palette alone. A palette set with `depictr_options(palette = )` is the user's own, so it is returned in the order it was set, with a warning. That is also how to make the order the default for every later plot: `depictr_options(palette = depictr_palette(order = "contrast"))`. The palette is then the user's own, so asking for more than four colours no longer warns. `scale_colour_depict()` and `scale_fill_depict()` take order for a single plot.

References

Okabe M, Ito K (2008). "Color Universal Design (CUD): How to make figures and presentations that are friendly to colorblind people." <https://jfly.uni-koeln.de/color/>. Accessed 2026-06-14.

World Wide Web Consortium (2023). "Web Content Accessibility Guidelines (WCAG) 2.2." W3C Recommendation, <https://www.w3.org/TR/WCAG22/>. Accessed 2026-08-18.

Examples

```
depictr_palette(3)
depictr_palette(3, order = "contrast")
depictr_palette(7, type = "sequential")
scales::show_col(depictr_palette())
```

dumbbell_plot

Dumbbell plot

Description

Compares a numeric value between exactly two groups across a set of categories: for each category the two group values are drawn as points joined by a connecting segment, so the size and direction of the gap is read at a glance. It is a clearer alternative to paired or grouped bars for before/after, two-condition or two-period comparisons.

Usage

```
dumbbell_plot(
  data,
  category,
  value,
  group,
  sort = c("gap", "value", "none"),
  point_size = 3,
  palette = NULL,
  legend_inside = FALSE,
```

```

    title = NULL,
    x_lab = NULL,
    y_lab = NULL
  )

```

Arguments

<code>data</code>	A data frame.
<code>category</code>	The categorical axis (string or unquoted name); one row of the plot per level.
<code>value</code>	The numeric value to compare (string or unquoted name).
<code>group</code>	The two-level grouping variable whose levels are the two ends of each dumbbell (string or unquoted name).
<code>sort</code>	How to order the categories up the axis: "gap" (by the signed difference between the two groups, the default), "value" (by the second group's value), or "none" (the data's own order: the order of the levels when <code>category</code> is a factor, and otherwise the order in which the categories first appear).
<code>point_size</code>	Size of the end points.
<code>palette</code>	Length-2 colours for the two groups; defaults to <code>depictr_palette()</code> .
<code>legend_inside</code>	When TRUE, draw the two-group legend inside the panel (in the top-right corner, which the default gap sort with the shortest dumbbell on top usually leaves clear) over a translucent background, instead of in a right-hand margin. Defaults to FALSE.
<code>title, x_lab, y_lab</code>	Plot title and axis labels.

Details

When a category has several rows per group the values are averaged. Rows with a missing category, value or group are dropped.

Value

A `ggplot2::ggplot` object.

Examples

```

wb <- wellbeing_survey
wb$age_group <- ifelse(wb$age < median(wb$age), "younger", "older")
dumbbell_plot(wb, region, life_satisfaction, age_group)

```

ecdf_plot

*Empirical cumulative distribution function (ECDF) plot***Description**

Draws the empirical CDF of a numeric variable (the proportion of observations at or below each value), optionally split by a grouping variable. Unlike a histogram it needs no bin-width choice and makes quantiles, medians and group shifts easy to read directly off the curve.

Usage

```
ecdf_plot(
  data,
  x,
  group = NULL,
  reference_quantiles = NULL,
  palette = NULL,
  legend_inside = FALSE,
  title = NULL,
  x_lab = NULL,
  y_lab = "Cumulative proportion"
)
```

Arguments

<code>data</code>	A data frame.
<code>x</code>	The numeric variable (string or unquoted name).
<code>group</code>	Optional grouping variable (string or unquoted name) mapped to colour, giving one ECDF per group. A numeric code counts as a group, but a numeric column with more than 20 distinct values is refused as continuous. Rows with a missing group are dropped with a message, and levels without rows are left out of the plot and its palette.
<code>reference_quantiles</code>	Optional numeric vector of probabilities in [0, 1] to mark with light horizontal guides (e.g. <code>c(0.25, 0.5, 0.75)</code>); NULL (the default) draws none.
<code>palette</code>	Colours for the groups; defaults to <code>depictr_palette()</code> .
<code>legend_inside</code>	When TRUE (and a group is given), draw the legend inside the panel (in the bottom-right corner the ECDF leaves empty once it saturates) over a translucent background, instead of in a right-hand margin. Defaults to FALSE.
<code>title, x_lab, y_lab</code>	Plot title and axis labels.

Value

A `ggplot2::ggplot` object.

Examples

```
ecdf_plot(lexical_decision, RT)
ecdf_plot(lexical_decision, RT, group = condition,
          reference_quantiles = c(0.25, 0.5, 0.75))
```

effects_plot

*Plot predicted values for one predictor***Description**

Shows the values a model predicts as one focal predictor varies, holding the other predictors at typical values (the mean for numeric predictors, the most frequent level for factors). A confidence band (numeric predictor) or confidence intervals (factor predictor) convey uncertainty.

Usage

```
effects_plot(
  model,
  predictor,
  conf_level = 0.95,
  n = 100,
  rug = TRUE,
  colour = depictr_brand(),
  title = NULL,
  x_lab = NULL,
  y_lab = NULL
)
```

Arguments

model	A fitted model (lm, glm, merMod, ...).
predictor	Name of the focal predictor (string): a variable in the model's formula, or a model-frame name such as "log(rainfall)", which stands for the one variable it transforms. A term built from several variables, such as "I(rainfall * fertiliser)", is refused.
conf_level	Confidence level for the interval, a single number strictly between 0 and 1.
n	Number of points across the range of a numeric predictor.
rug	Whether to add a rug of the observed predictor values (numeric predictors).
colour	Colour for the line/points and band. Defaults to the depictr brand blue.
title, x_lab, y_lab	Title and axis labels.

Details

The prediction grid holds the variables the formula names, on their own scale, and `stats::predict()` re-applies each transformation to it. A model with terms such as `log(rainfall)`, `poly(fertiliser, 2)`, `splines::ns()`, `scale()`, `factor()` or `I()` is therefore predicted as fitted, and the x-axis shows the raw variable. Typical values and the range of the focal predictor are taken over the rows the model was fitted to. A numeric variable that enters only through `factor()` is held at its most frequent value. A variable that enters only through a transformation is recovered from the data named in the model's call, so such a model fitted to workspace vectors without data = is refused. A term computed from statistics of its rows, such as `I(x - mean(x))`, is refused too, since `predict()` would compute it afresh over the grid. Centre such a variable in the data, or use `scale()`, which `predict()` re-applies as fitted.

An offset, whether an `offset()` term or the `offset` argument, is held at its mean over the rows the model was fitted to, on the link scale. For a Poisson rate model with `offset(log(exposure))`, the line is the expected count at the geometric-mean exposure. The offset stays at its mean when the focal predictor enters it too, as in `log(exposure) + offset(log(exposure))`, so the line then shows the predictor's effect beyond the offset. The y-axis label ends in "at the mean offset".

Predictions and standard errors come from `stats::predict()`; glm predictions are formed on the link scale and back-transformed, so a binomial model shows predicted probabilities. The band uses the multiplier of the model's own tests, as `tidy_estimates()` does. An `lm`, and a `glm` whose dispersion is estimated, get `t` on the residual degrees of freedom, so a Gaussian `glm` draws the band of the same `lm`. Binomial and Poisson fits and mixed models get `z`. Mixed models fitted with `lme4::lmer()/lme4::glmer()` are supported too: predictions use only the fixed effects (as `re.form = NA` does) and standard errors come from the fixed-effect design matrix and `vcov()`. Works with `lm`, `glm` and `merMod`; other model classes are attempted on a best-effort basis.

Value

A `ggplot2::ggplot` object.

Examples

```
fit <- lm(yield ~ rainfall + fertiliser + treatment, data = crop_yield)
effects_plot(fit, "fertiliser")
effects_plot(fit, "treatment")

gfit <- glm(accuracy ~ word_frequency + condition,
           data = lexical_decision, family = binomial)
effects_plot(gfit, "word_frequency")      # predicted probability
```

estimation_plot

Cumming estimation plot

Description

An estimation plot puts the *effect size* and its uncertainty at the centre of the comparison, rather than a p-value. The upper panel shows each group's raw data (jittered) with its mean and confidence

interval. The lower panel shows the difference between each other group's mean and a reference group's mean, with a bootstrap confidence interval. Each difference is drawn in its group's column, beneath that group's data. The two panels are stacked with 'patchwork', and each has its own outcome axis.

Usage

```
estimation_plot(
  data,
  y,
  group,
  reference = NULL,
  conf_level = 0.95,
  n_boot = 5000,
  effsize = c("hedges_g", "cohens_d", "none"),
  show_points = TRUE,
  point_alpha = 0.25,
  palette = NULL,
  title = NULL,
  y_lab = NULL,
  heights = c(2, 1.4)
)
```

Arguments

data	A data frame.
y	The numeric outcome (string or unquoted name).
group	The grouping variable (string or unquoted name).
reference	The reference (control) group that the others are compared with; defaults to the first level of group. The reference is drawn first in the upper panel and has an empty column in the lower panel, where the dashed line marks a difference of zero.
conf_level	Confidence level for both the group intervals (t-based) and the bootstrap difference intervals, a single number strictly between 0 and 1.
n_boot	Number of bootstrap resamples for the difference intervals.
effsize	Standardised effect size annotated beside each difference: "hedges_g" (the default, small-sample corrected), "cohens_d", or "none" to omit it. Every contrast against the reference is labelled, so the standardised effect is shown for both the two-group and multi-group cases.
show_points	Whether to draw the raw data behind the group means.
point_alpha	Transparency of the raw points.
palette	Colours for the groups; defaults to <code>depictr_palette()</code> .
title, y_lab	Title and outcome-axis label.
heights	Relative heights of the upper (raw data) and lower (difference) panels, passed to <code>patchwork::plot_layout()</code> .

Details

This stacked layout is the *Cumming* estimation plot (Cumming, 2012; Ho et al., 2019), with two groups as with more. Each difference is annotated with a standardised effect size: Cohen's d or, by default, the small-sample corrected Hedges' g (Hedges, 1981). The Gardner-Altman plot (Gardner & Altman, 1986) is a different layout, which sets a single difference beside the two groups on a floating axis whose zero is level with the reference group's mean.

The lower-panel interval is a non-parametric *bootstrap* of the mean difference: the two groups are resampled with replacement `n_boot` times and the requested percentile interval is read off the resampled differences. This makes no normality assumption about the sampling distribution of the difference. The bootstrap uses base R only; set a seed beforehand for reproducibility.

A group with fewer than two observations has no estimable spread, so its confidence interval is omitted (the mean is still drawn) and a difference involving it is shown as a point without a bootstrap interval; a warning is issued in both cases.

Value

A 'patchwork' object (printable like a `ggplot2::ggplot`). The computed differences and their bootstrap intervals are attached as the "differences" attribute (a data frame).

References

Cumming G (2012). *Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis*. Routledge, New York, NY. ISBN 978-0-415-87968-2.

Gardner MJ, Altman DG (1986). "Confidence intervals rather than P values: Estimation rather than hypothesis testing." *BMJ*, **292**(6522), 746–750. doi:10.1136/bmj.292.6522.746.

Hedges LV (1981). "Distribution theory for Glass's estimator of effect size and related estimators." *Journal of Educational Statistics*, **6**(2), 107–128. doi:10.3102/10769986006002107.

Ho J, Tumkaya T, Aryal S, Choi H, Claridge-Chang A (2019). "Moving beyond P values: Data analysis with estimation graphics." *Nature Methods*, **16**(7), 565–566. doi:10.1038/s415920190470-3.

Examples

```
set.seed(1)
# n_boot is kept small here for speed; use the default for real work.
estimation_plot(lexical_decision, RT, condition, n_boot = 1000)

estimation_plot(crop_yield, yield, treatment)
# More than two groups: differences vs a chosen reference
estimation_plot(wellbeing_survey, life_satisfaction, region,
  reference = "North")
```

explore_bivariate	<i>Plot any pair of variables</i>
-------------------	-----------------------------------

Description

Chooses an appropriate plot for the relationship between two variables according to their types. Two numeric variables are shown as a scatter plot with a fitted trend; a numeric and a categorical variable as box plots of the numeric variable by level; and two categorical variables as a filled bar chart of proportions.

Usage

```
explore_bivariate(  
  data,  
  x,  
  y,  
  method = "lm",  
  palette = NULL,  
  title = NULL,  
  x_lab = NULL,  
  y_lab = NULL  
)
```

Arguments

data	A data frame.
x, y	The two variables (string or unquoted name).
method	Smoothing method for the numeric-numeric case (passed to <code>ggplot2::geom_smooth()</code>); NULL for no trend.
palette	Colours used when a fill/colour is needed; defaults to <code>depictr_palette()</code> .
title, x_lab, y_lab	Title and axis labels (default to variable names).

Details

Dates, date-times and durations (Date, POSIXct and difftime columns) count as numeric. Against a numeric variable, they give a scatter plot with a trend, drawn by `scatter_trend()`, and against a categorical one box plots of their values by level. A duration is drawn as a number in its own units, named in the default axis title, as in `stay (days)`.

Value

A `ggplot2::ggplot` object.

Examples

```

explore_bivariate(crop_yield, fertiliser, yield)      # numeric ~ numeric
explore_bivariate(lexical_decision, condition, RT)  # categorical ~ numeric
explore_bivariate(wellbeing_survey, region, education) # categorical ~ categorical

```

explore_categorical *Bar chart of a categorical variable*

Description

Counts (or proportions) of the levels of a categorical variable, optionally split by a grouping variable.

Usage

```

explore_categorical(
  data,
  x,
  group = NULL,
  proportion = FALSE,
  position = c("dodge", "stack", "fill"),
  sort = TRUE,
  horizontal = FALSE,
  palette = NULL,
  title = NULL,
  x_lab = NULL
)

```

Arguments

data	A data frame.
x	The categorical variable (string or unquoted name). Numeric columns are accepted only when they have at most 20 distinct values.
group	Optional grouping variable mapped to fill, with the groups in the order of its levels. A numeric code counts as a group, but a numeric column with more than 20 distinct values is refused as continuous. Rows with a missing group are dropped with a message, and levels without rows are left out of the plot and its palette.
proportion	Whether to show proportions instead of counts. When group is set, proportions are computed within each group.
position	Bar position when group is set: "dodge", "stack" or "fill".
sort	Whether to order the bars from most to least frequent.
horizontal	Whether to draw horizontal bars, which helps when there are many levels or long labels.
palette	Colours for the groups; defaults to <code>depictr_palette()</code> .
title, x_lab	Plot title and category-axis label (defaults to the variable name).

Details

A continuous numeric column (more than 20 distinct values) is rejected with an error, since coercing it to a factor would draw one bar per value; use `explore_distribution()` for such variables instead.

Value

A `ggplot2::ggplot` object.

Examples

```
explore_categorical(wellbeing_survey, region)
explore_categorical(wellbeing_survey, education, group = region,
                    proportion = TRUE, position = "dodge")
```

`explore_distribution` *Plot the distribution of a variable*

Description

Draws a histogram, a density curve, or both, optionally split by a grouping variable. A quick first look at a continuous variable.

Usage

```
explore_distribution(
  data,
  x,
  group = NULL,
  type = c("histogram", "density", "both"),
  bins = 30,
  alpha = 0.6,
  position = NULL,
  palette = NULL,
  facet = FALSE,
  legend_inside = FALSE,
  title = NULL,
  x_lab = NULL
)
```

Arguments

<code>data</code>	A data frame.
<code>x</code>	The continuous variable to display. Either a string or an unquoted column name.
<code>group</code>	Optional grouping variable (string or unquoted name) mapped to colour/fill. A numeric code counts as a group, but a numeric column with more than 20 distinct values is refused as continuous. Rows with a missing group are dropped with a message, and levels without rows are left out of the plot and its palette.

type	One of "histogram", "density" or "both".
bins	Number of histogram bins.
alpha	Fill transparency (useful when groups overlap).
position	Histogram position adjustment, e.g. "identity", "stack" or "dodge". The default (NULL) chooses "dodge" when group is set (so overlapping bars stay readable) and "identity" otherwise.
palette	Colours for the groups; defaults to <code>depictr_palette()</code> .
facet	When a group is given, draw one panel per group instead of overlaying them. This is much clearer than an overlay once there are more than two or three groups (overlaid histograms in particular become hard to read). Defaults to FALSE. Ignored when there is no group.
legend_inside	When TRUE (and a group is given without facet), draw the colour legend inside the panel (in the top-right corner a unimodal histogram/density leaves empty) over a translucent background, instead of in a right-hand margin. Defaults to FALSE.
title, x_lab	Plot title and x-axis label (defaults to the variable name).

Value

A `ggplot2::ggplot` object.

Examples

```
explore_distribution(lexical_decision, RT)
explore_distribution(lexical_decision, RT, group = condition, type = "density")
# One panel per group keeps many groups legible:
explore_distribution(wellbeing_survey, life_satisfaction, group = region,
                     type = "both", facet = TRUE)
```

explore_pairs

Scatter-plot matrix

Description

A matrix of pairwise plots for a set of numeric variables: scatter plots below the diagonal, densities on the diagonal and correlation coefficients above the diagonal. Optionally colour everything by a grouping variable. Built with 'patchwork', so it shares the package theme and palette.

Usage

```
explore_pairs(
  data,
  cols = NULL,
  group = NULL,
  max_cols = 8,
  point_alpha = 0.5,
```

```

    method = c("pearson", "spearman", "kendall"),
    palette = NULL,
    title = NULL
  )

```

Arguments

data	A data frame.
cols	Numeric columns to include, which cannot include group. If NULL, every numeric column other than group is used, so a frame with more than <code>max_cols</code> of them needs <code>cols</code> to choose among them.
group	Optional grouping variable mapped to colour, with each level in the same colour in every panel. A numeric code counts as a group, but a numeric column with more than 20 distinct values is refused as continuous. Rows with a missing group are dropped with a message, and levels without rows are left out of the plot and its palette.
max_cols	The largest number of variables accepted, since a k-by-k matrix grows quickly. More columns than this is an error, so the matrix never leaves a variable out without saying so.
point_alpha	Point transparency in the scatter panels.
method	Correlation method for the upper-triangle coefficients, passed to <code>stats::cor()</code> : "pearson", "spearman" or "kendall". A pair whose correlation is undefined (a constant column, or too few complete cases) is labelled n/a, as in <code>correlation_heatmap()</code> .
palette	Colours for the groups; defaults to <code>depictr_palette()</code> .
title	Overall title for the matrix.

Value

A 'patchwork' object (printable like a `ggplot2::ggplot`).

Examples

```

explore_pairs(crop_yield, cols = c("rainfall", "fertiliser", "yield"))

explore_pairs(crop_yield, cols = c("rainfall", "fertiliser", "yield"),
              group = treatment, method = "spearman")

```

Description

Cleans up the term names produced by modelling functions so that they read well on a plot: the intercept is renamed, an optional b_/bs_ Bayesian prefix is stripped, interaction colons are converted to a chosen symbol, and underscores are shown as spaces (e.g. word_frequency becomes "word frequency"). The b_/bs_ prefix is stripped from *each* component of an interaction (e.g. b_x:b_y), not just the leading term. Missing values (NA) are kept as NA rather than being rendered as the literal text "NA".

Usage

```
format_terms(
  x,
  interaction = c("times", "asterisk", "colon", "space"),
  strip_prefix = TRUE,
  tidy_intercept = TRUE,
  wrap = NULL
)
```

Arguments

x	Character vector of term names.
interaction	How to render interaction colons: "times" (the default, a Unicode multiplication sign), "asterisk", "colon" (unchanged) or "space".
strip_prefix	Whether to remove a leading b_ or bs_ (as added by 'brms') from each interaction component.
tidy_intercept	Whether to replace (Intercept) with "Intercept".
wrap	Optional integer width at which to wrap long labels onto new lines (see base::strwrap()). NULL (default) leaves labels unwrapped.

Value

A character vector the same length as x, with NA preserved.

Examples

```
format_terms(c("(Intercept)", "b_conditionB", "freq:condition"))
format_terms("region:education:age", interaction = "asterisk")
format_terms(c("word_frequency", "b_sleep_hours"))
format_terms(c("b_x:b_y", NA))
```

frequentist_bayesian_plot

Plot frequentist and Bayesian estimates together

Description

Presents the estimates from a frequentist model and a Bayesian model on one plot, in the first two colours of the colourblind-safe `depictr_palette()`: brand blue for the frequentist estimates and orange for the Bayesian ones, whether the Bayesian side is drawn from draws or from summaries.

Usage

```
frequentist_bayesian_plot(
  frequentist,
  bayesian,
  conf_level = 0.95,
  labels = NULL,
  interaction = c("times", "asterisk", "colon", "space"),
  intercept = TRUE,
  facet = TRUE,
  scales = c("free", "fixed"),
  note_frequentist_no_prior = FALSE,
  vertical_line_at_x = 0,
  title = NULL,
  subtitle = NULL,
  x_lab = "Estimate",
  caption = NULL,
  ...
)
```

Arguments

<code>frequentist</code>	A frequentist model (e.g. from <code>lm</code> , <code>glm</code> or <code>lmerTest::lmer</code>) or a tidy data frame of estimates.
<code>bayesian</code>	A Bayesian model, a posterior draws object, a draws matrix/data frame, a tidy data frame of posterior summaries or a summary matrix such as <code>brms::fixef()</code> returns.
<code>conf_level</code>	Level of the intervals the function computes: the confidence intervals of a frequentist model and, when <code>bayesian</code> holds draws, the outer credible interval of each posterior (see Details). A single number strictly between 0 and 1.
<code>labels</code> , <code>interaction</code> , <code>intercept</code>	See <code>compare_models()</code> . On both paths, an unnamed <code>labels</code> vector takes the frequentist terms first, then any that only the Bayesian input holds. <code>intercept</code> defaults to <code>TRUE</code> here.

facet, scales	Layout controls. Because a Bayesian model almost always carries a large intercept alongside small slopes, the comparison defaults to a faceted, free-scaled layout (<code>facet = TRUE</code>): each term gets its own panel and free x-axis, so every posterior and its frequentist overlay stay legible. Pass <code>facet = FALSE</code> (or <code>scales = "fixed"</code>) for the classic single shared-axis plot.
note_frequentist_no_prior	If TRUE, append "(no prior)" to the frequentist legend label, which is helpful when the title names the Bayesian prior.
vertical_line_at_x	Position of the vertical reference line (NA to omit).
title, subtitle, x_lab	Title, subtitle and x-axis label.
caption	Plot caption. The default (NULL) names the level of each source's intervals. Pass NA to omit the caption.
...	Further arguments passed to <code>compare_models()</code> on the summary path (ignored on the distribution path).

Details

When `bayesian` carries posterior *draws* (a `brms/rstanarm` fit, a posterior draws object, a draws matrix, or a long/wide draws data frame), the full posterior *distribution* is drawn per term (a 'ggdist' half-eye) and the frequentist point and confidence interval is overlaid at the same position. When `bayesian` is only a table of posterior *summaries*, the function shows the familiar two-source forest plot via `compare_models()`. It reads tidy columns (`term`, `estimate`, `conf.low`, `conf.high`), the `Estimate`, `Q2.5` and `Q97.5` of the matrix `brms::fixef()` returns, and the `l-95%` CI and `u-95%` CI of `summary(fit)$fixed`.

Terms are aligned by their canonical display label, so the `brms`-style `b_` prefix is reconciled automatically against the frequentist term names. Draws are read as in `posterior_plot()`, so a draws object holding `b_` parameters keeps only those, and two terms that would share a label are an error. With draws, only the terms both sides share are drawn, and a message names any dropped, such as a `sigma` in the posterior. The forest plot of summaries draws every term of either input.

`conf_level` is the level of every interval the function computes. For a frequentist model, that is its confidence interval. For posterior draws, it is the outer of the equal-tailed credible intervals drawn beneath each half-eye, so the two sources are compared at one level. The inner interval is 66%, drawn only when `conf_level` exceeds 0.66 so that it stays the narrower of the two. A table of estimates keeps the interval limits it holds, at whatever level they were computed. By default, a caption names the level of each source's intervals or says that they are as supplied.

Value

A `ggplot2::ggplot` object.

Examples

```
# Summary path: a tidy "Bayesian" summary as a data frame. The summary here
# stands in for a regularised posterior, so a weakly informative prior pulls
# each estimate toward zero and narrows its interval; a real posterior would
# come from the model fit rather than from arithmetic on the frequentist one.
```

```

freq <- lm(life_satisfaction ~ stress + sleep_hours + exercise_days,
           data = wellbeing_survey)
bayes <- tidy_estimates(freq)
bayes$estimate <- bayes$estimate * 0.9
half <- (bayes$conf.high - bayes$conf.low) / 2 * 0.8
bayes$conf.low <- bayes$estimate - half
bayes$conf.high <- bayes$estimate + half
frequentist_bayesian_plot(freq, bayes,
                          title = "Frequentist vs. Bayesian estimates")

# Distribution path: simulated posterior draws (one column per term) drawn as
# full posteriors with the frequentist point + CI overlaid.
set.seed(1)
co <- coef(freq)
draws <- as.data.frame(lapply(co, function(m) rnorm(400, m, abs(m) * 0.1 + 0.05)))
names(draws) <- names(co)
frequentist_bayesian_plot(freq, draws,
                          title = "Posterior with frequentist overlay")

```

gain_plot

Cumulative gains chart

Description

Shows how many of the positive cases are captured as a growing share of the population is targeted in order of predicted score. It is the customary chart for judging a classifier's value for ranking and targeting, as in marketing, triage and fraud detection. The diagonal marks the no-model baseline, and a second reference line marks a perfect model. Pass a *named list* of models / (actual, score) pairs to overlay several colour-coded gains curves with a legend. The perfect-model line is then omitted, since it depends on the prevalence of the outcome and the overlaid models need not share one.

Usage

```

gain_plot(
  x,
  score = NULL,
  colour = depictr_brand(),
  legend_inside = FALSE,
  title = NULL
)

```

Arguments

x	A binomial glm; the vector of observed outcomes (0/1, logical or a two-level factor with the positive class second); or a <i>named list</i> of models / (actual, score) pairs to overlay.
score	When x is an outcome vector, the matching scores or predicted probabilities (or a list of them for the multi-model case).

colour	Curve colour for the single-model case. Defaults to the depictr brand blue. Ignored when several models are overlaid.
legend_inside	When TRUE (and several models are overlaid), draw the legend inside the panel (in the bottom-right triangle the concave curve leaves empty) over a translucent background, instead of in a right-hand margin. Defaults to FALSE.
title	Plot title.

Value

A `ggplot2::ggplot` object.

Examples

```
# Targeting only pays off when the positive class is scarce, so the example
# uses the rare clinical-trial adverse event (about 10% positive).
gfit <- glm(adverse_event ~ biomarker + age + arm,
            data = clinical_trial, family = binomial)
gain_plot(gfit)

# Compare two models.
reduced <- glm(adverse_event ~ biomarker, data = clinical_trial,
               family = binomial)
gain_plot(list(Full = gfit, Reduced = reduced))
```

`group_comparison_plot` *Compare group means with confidence intervals*

Description

An estimation-style plot of a numeric outcome across the levels of a grouping variable: each group's mean with a confidence interval, over a backdrop of the raw (jittered) data. By showing both the estimate and its uncertainty, it conveys whether the groups differ more faithfully than a bar chart does.

Usage

```
group_comparison_plot(
  data,
  y,
  group,
  conf_level = 0.95,
  show_points = TRUE,
  point_alpha = 0.25,
  differences = FALSE,
  reference = NULL,
  n_boot = 5000,
  palette = NULL,
  title = NULL,
```

```

    x_lab = NULL,
    y_lab = NULL
  )

```

Arguments

<code>data</code>	A data frame.
<code>y</code>	The numeric outcome (string or unquoted name).
<code>group</code>	The grouping variable (string or unquoted name). A numeric code counts as a group, but a numeric column with more than 20 distinct values is refused as continuous. Rows with a missing group are dropped with a message, and levels without rows are left out of the plot and its palette.
<code>conf_level</code>	Confidence level for the intervals (t-based), a single number strictly between 0 and 1.
<code>show_points</code>	Whether to draw the raw data behind the means.
<code>point_alpha</code>	Transparency of the raw points.
<code>differences</code>	If TRUE, append a lower panel showing the pairwise mean difference(s) against a reference group, each with a bootstrap confidence interval, turning the plot into a full estimation plot via estimation_plot() . The return value is then a 'patchwork' object. Defaults to FALSE, the plain group-means plot.
<code>reference</code>	Reference group for the difference panel when <code>differences = TRUE</code> ; defaults to the first level of group. Ignored otherwise.
<code>n_boot</code>	Number of bootstrap resamples for the difference intervals when <code>differences = TRUE</code> . Ignored otherwise.
<code>palette</code>	Colours for the groups; defaults to depictr_palette() .
<code>title, x_lab, y_lab</code>	Title and axis labels.

Details

A group with a single observation has no degrees of freedom for a t-based interval, so only its mean is drawn (no interval) and a warning is issued.

Value

A [ggplot2::ggplot](#) object, or a 'patchwork' object when `differences = TRUE`.

Examples

```

group_comparison_plot(lexical_decision, RT, condition)
group_comparison_plot(crop_yield, yield, treatment)
# Append the pairwise mean-difference panel (an estimation plot):
set.seed(1)
group_comparison_plot(crop_yield, yield, treatment, differences = TRUE)

```

influence_plot	<i>Influence plot</i>
----------------	-----------------------

Description

A bubble plot of leverage against the studentised residuals, with bubble area proportional to Cook’s distance. It summarises in a single picture which observations most influence a fitted model. Reference lines mark large residuals and high-leverage points, and the most influential observations are labelled.

Usage

```
influence_plot(model, n_label = 3, colour = depictr_brand(), title = NULL)
```

Arguments

model	A fitted lm or glm model.
n_label	Number of most-influential points (by Cook’s distance) to label.
colour	Bubble colour. Defaults to the depictr brand blue.
title	Plot title.

Value

A `ggplot2::ggplot` object.

References

Cook RD (1977). “Detection of influential observation in linear regression.” *Technometrics*, **19**(1), 15–18. doi:10.1080/00401706.1977.10489493.

Examples

```
fit <- lm(yield ~ rainfall + fertiliser + soil_ph, data = crop_yield)
influence_plot(fit)
```

interaction_plot	<i>Plot a two-way interaction of predicted values</i>
------------------	---

Description

Shows how the predicted relationship between a focal predictor and the response changes across the levels (or representative values) of a second, moderating predictor. Other predictors are held at typical values.

Usage

```
interaction_plot(
  model,
  predictor,
  moderator,
  moderator_values = NULL,
  conf_level = 0.95,
  n = 80,
  band = TRUE,
  palette = NULL,
  title = NULL,
  x_lab = NULL,
  y_lab = NULL
)
```

Arguments

model	A fitted model (lm, glm, merMod, ...).
predictor	Name of the focal predictor on the x-axis (string), given as in <code>effects_plot()</code> : "log(rainfall)" means rainfall.
moderator	Name of the moderating predictor, mapped to colour (string), given in the same way as predictor. It must be a different variable.
moderator_values	For a numeric moderator, the values to show, on the variable's own scale. Defaults to the 10th, 50th and 90th percentiles in the estimation sample.
conf_level	Confidence level for the bands/intervals, a single number strictly between 0 and 1.
n	Number of points across the range of a numeric focal predictor.
band	Whether to draw confidence bands (numeric focal predictor).
palette	Colours for the moderator; defaults to <code>depictr_palette()</code> .
title, x_lab, y_lab	Title and axis labels.

Details

The grid is built as in `effects_plot()`. It holds the raw variables, so transformed terms are predicted as fitted. Typical values come from the estimation sample, and an offset is held at its mean on the link scale.

Value

A `ggplot2::ggplot` object.

Examples

```
fit <- lm(yield ~ fertiliser * treatment + rainfall, data = crop_yield)
interaction_plot(fit, "fertiliser", "treatment")
```

k_diagnostic

*Suggest a number of clusters***Description**

Computes a cluster-quality diagnostic across a range of k and draws the diagnostic curve, with the suggested k highlighted. Three criteria are available: the average silhouette width (maximised), the total within-cluster sum of squares "elbow" (the point of maximum curvature), and the gap statistic (the smallest k whose gap is within one standard error of the next, using the Tibshirani et al. heuristic).

Usage

```
k_diagnostic(
  data,
  k_range = NULL,
  method = c("silhouette", "wss", "gap"),
  cols = NULL,
  scale = TRUE,
  nstart = 10,
  B = 50,
  title = NULL
)
```

Arguments

data	A data frame or numeric matrix.
k_range	Integer vector of candidate cluster counts to evaluate. The default, NULL, searches 1:8 for "gap" and "wss" and 2:8 for "silhouette". Candidates the chosen criterion cannot evaluate are skipped with a message naming them: $k = 1$ under "silhouette" (a lone cluster has no neighbouring cluster to compare against, so the width is undefined), and any k at or above the number of observations under every criterion.
method	Diagnostic: "silhouette", "wss" (within sum of squares elbow) or "gap".
cols	When data is a data frame, the numeric columns to use.
scale	Whether to scale variables to unit variance first. Columns with (near-)zero variance cannot be scaled and are dropped with a message, as in correlation_heatmap() .
nstart	Number of random starts for stats::kmeans() .
B	Number of reference bootstrap samples for the gap statistic.
title	Plot title.

Details

Only the gap statistic can conclude that the data hold no clusters, by suggesting $k = 1$, and only when the search starts at 1, as it does by default. The elbow is the candidate farthest from the straight line joining the first and last points of the curve, so it is never the first or last candidate. A range therefore has to extend one step beyond any k the elbow should be able to choose.

Value

A `ggplot2::ggplot` object showing the diagnostic value against `k`, with the suggested `k` marked and named in the subtitle. The underlying computation is attached as attributes: `attr(p, "k_table")` (a data frame with one row per `k`), `attr(p, "suggested")` (the chosen `k`) and `attr(p, "method")`.

References

Rousseeuw PJ (1987). “Silhouettes: A graphical aid to the interpretation and validation of cluster analysis.” *Journal of Computational and Applied Mathematics*, **20**, 53–65. doi:10.1016/0377-0427(87)901257.

Tibshirani R, Walther G, Hastie T (2001). “Estimating the number of clusters in a data set via the gap statistic.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **63**(2), 411–423. doi:10.1111/14679868.00293.

Examples

```
p <- k_diagnostic(crop_yield, k_range = 2:6,
                 cols = c("rainfall", "fertiliser", "soil_ph", "yield"))
attr(p, "suggested")
```

lexical_decision

*Simulated lexical-decision experiment***Description**

A reproducibly simulated psycholinguistic lexical-decision experiment, in which participants judge whether letter strings are real words. The design is fully crossed and counterbalanced: 24 participants each respond to all 40 items once (960 trials). `condition` (related/unrelated priming) and `modality` (visual/auditory) are both within-participant and within-item, assigned by a Latin-square list rotation so that, across participants, every item appears equally often in each condition and in each modality. Neither factor is therefore collinear with the item random intercept, which makes this a clean fit for crossed-random-effects models such as `lmer(RT ~ condition + modality + word_frequency + (1 | participant) + (1 | item))`, and hence for `optimizer_fixef_plot()`, `coefficient_plot()` and the distribution plots.

Usage

lexical_decision

Format

A data frame with 960 rows and 7 variables:

participant Participant identifier (factor, 24 levels).

item Item identifier (factor, 40 levels).

condition Priming condition (factor): related or unrelated.

modality Presentation modality (factor): visual or auditory.

word_frequency Item-level word frequency on the Zipf scale.

RT Reaction time, in milliseconds.

accuracy Response accuracy: 1 correct, 0 incorrect.

Details

RT is generated log-normally from participant and item random intercepts plus the fixed effects (unrelated priming about 35 ms slower; auditory a touch slower; more frequent words faster), with a 250 ms floor; RT analyses should use correct trials (`accuracy == 1`). `accuracy` is generated from the exogenous predictors (`word_frequency`, `condition`, `modality`) and random effects only - it is not derived from RT - so classification demos built on it are not circular.

The data are synthetic, generated by `data-raw/generate_datasets.R` with a fixed seed; they do not describe any real individuals.

Source

Simulated; see `data-raw/generate_datasets.R`.

lift_plot

Cumulative lift chart

Description

Shows how many times more positive cases a classifier captures, at each depth of the score-ordered population, than random targeting would. A lift of 3 at the top 10% means that decile contains three times the baseline rate of positives. The horizontal line at 1 is the no-model baseline. Pass a *named list* of models / (actual, score) pairs to overlay several colour-coded lift curves with a legend.

Usage

```
lift_plot(
  x,
  score = NULL,
  colour = depictr_brand(),
  legend_inside = FALSE,
  title = NULL
)
```

Arguments

<code>x</code>	A binomial glm; the vector of observed outcomes (0/1, logical or a two-level factor with the positive class second); or a <i>named</i> list of models / (actual, score) pairs to overlay.
<code>score</code>	When <code>x</code> is an outcome vector, the matching scores or predicted probabilities (or a list of them for the multi-model case).

colour	Curve colour for the single-model case. Defaults to the depictr brand blue. Ignored when several models are overlaid.
legend_inside	When TRUE (and several models are overlaid), draw the legend inside the panel (in the top-right corner, which the decaying lift curve leaves empty) over a translucent background, instead of in a right-hand margin. Defaults to FALSE.
title	Plot title.

Value

A `ggplot2::ggplot` object.

Examples

```
# Lift is measured against the base rate, so the example uses the rare
# clinical-trial adverse event (about 10% positive).
gfit <- glm(adverse_event ~ biomarker + age + arm,
            data = clinical_trial, family = binomial)
lift_plot(gfit)

# Compare two models.
reduced <- glm(adverse_event ~ biomarker, data = clinical_trial,
               family = binomial)
lift_plot(list(Full = gfit, Reduced = reduced))
```

missingness_map	<i>Map the missing values in a data frame</i>
-----------------	---

Description

Draws a tile map of the data frame with one column per variable and one row per observation, shading the cells that are missing. The variables are ordered by their proportion of missing values, and that proportion is shown in the axis labels, making it easy to spot variables and patterns that need attention before modelling.

Usage

```
missingness_map(
  data,
  cols = NULL,
  sort = TRUE,
  show_pct = TRUE,
  colours = c("grey85", depictr_accent()),
  legend_inside = FALSE,
  title = NULL
)
```

Arguments

<code>data</code>	A data frame.
<code>cols</code>	Optional character vector of columns to include (default: all).
<code>sort</code>	Whether to order variables by their proportion of missing values.
<code>show_pct</code>	Whether to append the percentage missing to each variable label.
<code>colours</code>	Length-2 vector: colours for present and missing cells. Defaults to a muted grey for present cells and the colourblind-safe <code>depictr_palette()</code> accent for missing cells.
<code>legend_inside</code>	When TRUE (and <code>sort = TRUE</code>), draw the legend inside the panel, in the top-right (where the most-complete columns put a solid "Present" block, so it hides no "Missing" mark) instead of in a right-hand margin. Defaults to FALSE.
<code>title</code>	Plot title.

Value

A `ggplot2::ggplot` object.

Examples

```
missingness_map(wellbeing_survey)
```

<code>model_fit_table</code>	<i>Goodness-of-fit statistics across models</i>
------------------------------	---

Description

Collects the usual model-comparison statistics into one tidy data frame, one row per model. These comprise the number of observations, the degrees of freedom, AIC, BIC, the log-likelihood, an R-squared and the root-mean-square error.

Usage

```
model_fit_table(..., digits = 3)
```

Arguments

<code>...</code>	One or more fitted models. Name the arguments to label the rows.
<code>digits</code>	Number of decimal places to round to.

Details

The R-squared depends on the kind of model. An `lm` reports the ordinary coefficient of determination. Every `glm`, including quasi-likelihood families and negative binomial fits from `MASS::glm.nb()`, reports the deviance R-squared, $1 - D/D_0$. Here D is the residual deviance of the model and D_0 the deviance of the null model, which `stats::glm()` fits to the same rows with the same prior weights and offset. The null model has an intercept only if the model has one. Both deviances are stored with the fit, so nothing is refitted. The value is the share of the null deviance that the predictors explain, the Kullback-Leibler R-squared of Cameron and Windmeijer (1997). It equals the ordinary R-squared for a Gaussian model with the identity link and McFadden's (1974) pseudo-R-squared for binary logistic regression. For Poisson, grouped binomial, Gamma and other models, the two measures differ. Mixed models (`merMod`) report an R-squared of NA, since no single definition is standard for them.

A linear mixed model fitted with `lme4::lmer()` by restricted maximum likelihood (REML, its default) is refitted by maximum likelihood with `lme4::refitML()`. Its AIC, BIC and log-likelihood come from the refit, and a message says so. A REML likelihood is the likelihood of the residuals left once the fixed effects are removed. Two models with different fixed effects therefore have REML likelihoods of different data, and their criteria cannot be compared. The `anova()` method for `lme4` fits refits them by maximum likelihood for the same reason. If a refit fails, a warning says so, and that model's df, AIC, BIC and `logLik` are NA. The other columns describe the model as supplied.

Value

A data frame with one row per model and columns `model`, `n`, `df`, `AIC`, `BIC`, `logLik`, `R2` and `RMSE`.

References

Cameron AC, Windmeijer FAG (1997). "An R-squared measure of goodness of fit for some common nonlinear regression models." *Journal of Econometrics*, **77**(2), 329–342. doi:10.1016/S0304-4076(96)018180.

McFadden D (1974). "Conditional logit analysis of qualitative choice behavior." In Zarembka P (ed.), *Frontiers in econometrics*, 105–142. Academic Press, New York, NY.

Examples

```
m1 <- lm(yield ~ rainfall, data = crop_yield)
m2 <- lm(yield ~ rainfall + fertiliser, data = crop_yield)
m3 <- lm(yield ~ rainfall + fertiliser + soil_ph + treatment,
        data = crop_yield)
model_fit_table(simple = m1, medium = m2, full = m3)
```

Description

Assembles a compact, consistent overview of a fitted model in a single figure: the coefficient estimates, the predicted effect of a focal predictor, a residual panel and a normal Q-Q plot, with a subtitle of key fit statistics. It is a convenience wrapper that composes several `depictr` plots with `arrange_plots()`, and serves well for a rapid model review or a report appendix.

Usage

```
model_report(
  model,
  predictor = NULL,
  standardise = TRUE,
  title = NULL,
  subtitle = NULL,
  glm_panels = TRUE,
  seed = NULL
)
```

Arguments

<code>model</code>	A fitted <code>lm</code> or <code>glm</code> model.
<code>predictor</code>	Focal predictor for the effect panel, given as in <code>effects_plot()</code> . If <code>NULL</code> , the first numeric variable in the formula (or, failing that, the first variable) is used, so a model with <code>log(rainfall)</code> shows the effect of <code>rainfall</code> .
<code>standardise</code>	Whether the coefficient panel scales each continuous predictor's coefficient by that predictor's standard deviation, leaving factor and binary contrasts in the outcome's units, as described for <code>standardise</code> in <code>coefficient_plot()</code> . Defaults to <code>TRUE</code> , which keeps the panel readable in this compact overview when the slopes are on very different scales. Set <code>FALSE</code> for raw estimates.
<code>title</code>	Overall title.
<code>subtitle</code>	Overall subtitle. If <code>NULL</code> , a line of fit statistics (number of observations, R-squared and AIC) is used. For a <code>glm</code> , the R-squared is the deviance R-squared described in <code>model_fit_table()</code> .
<code>glm_panels</code>	For a <code>glm</code> , whether to draw the binned-residual and quantile-residual panels (the default) instead of the classic <code>lm</code> panels. Ignored for an <code>lm</code> .
<code>seed</code>	Optional integer seed for the randomisation of discrete quantile residuals, for a reproducible Q-Q panel of a binomial or Poisson <code>glm</code> . The caller's random number stream is left as it was.

Details

For an `lm`, the residual panel plots the residuals against the fitted values, and the Q-Q panel is `qq_plot()` of the standardised residuals. A `glm` gets the panels `residual_diagnostics_plot()` draws for it (`glm_panels = TRUE`). The residual panel is a `binned_residual_plot()` (Gelman & Hill, 2007). The Q-Q panel plots randomised quantile residuals (Dunn & Smyth, 1996), which are standard normal under a correctly specified model whatever the response family, so the band of `qq_plot()` applies to them. The standardised deviance residuals of a binary or count model are

not expected to be normal, so even a correctly specified model can put many of them outside that band (about a third, in a simulated logistic model with 400 observations). The quantile residuals of a binomial or Poisson model are randomised, and even under a correct model a run of them occasionally leaves the pointwise band. A departure worth acting on persists when the panel is drawn again with another seed. For a family with no quantile residual here, such as a quasi-likelihood family, the Q-Q panel plots the deviance residuals and its axis says so.

Value

A 'patchwork' object (printable like a [ggplot2::ggplot](#)).

References

Gelman A, Hill J (2007). *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press, Cambridge, UK. ISBN 978-0-521-68689-1. doi:[10.1017/CBO9780511790942](#).

Dunn PK, Smyth GK (1996). "Randomized quantile residuals." *Journal of Computational and Graphical Statistics*, **5**(3), 236–244. doi:[10.1080/10618600.1996.10474708](#).

Examples

```
fit <- lm(yield ~ rainfall + fertiliser + soil_ph + treatment,
         data = crop_yield)
model_report(fit, title = "Crop-yield model")

# A logistic model gets binned and quantile-residual panels, and a seed
# fixes the randomisation of its quantile residuals:
gfit <- glm(accuracy ~ word_frequency + RT + condition,
           data = lexical_decision, family = binomial)
model_report(gfit, seed = 1)
```

monthly_sales

Simulated monthly sales time series

Description

A reproducibly simulated seasonal monthly time series in tidy long form, for the time-series examples. It holds two product lines (indoor and outdoor), each with six years of monthly observations carrying a linear trend, a clear twelve-month seasonal cycle (with different seasonal peaks per line) and noise. The two-series shape exercises the multi-group [timeseries_plot\(\)](#) path on first-party data; extracting one series as a monthly ts (frequency 12) lets [decompose_plot\(\)](#) and [acf_plot\(\)](#) reveal the trend and seasonality.

Usage

monthly_sales

Format

A data frame with 144 rows and 3 variables:

date First day of the month (Date), from January 2018.

series Product line (factor): indoor or outdoor.

sales Monthly sales (units).

Details

The data are synthetic, generated by `data-raw/generate_datasets.R` with a fixed seed.

Source

Simulated; see `data-raw/generate_datasets.R`.

optimizer_fixef_plot *Plot fixed effects across optimisers*

Description

Visualises how the fixed-effect estimates of a mixed model vary across the optimisers tried by `lme4::allFit()`. It offers a quick check that a model has settled on a stable solution: tight clusters of points indicate agreement between optimisers, whereas scatter signals a fragile fit.

Usage

```
optimizer_fixef_plot(  
  x,  
  intercept = TRUE,  
  select_terms = NULL,  
  interaction = c("times", "asterisk", "colon", "space"),  
  labels = NULL,  
  number_optimizers = TRUE,  
  free_y = TRUE,  
  ncol = NULL,  
  point_size = 2,  
  palette = NULL,  
  y_lab = "Fixed effect",  
  title = NULL  
)
```

Arguments

x	Either the object returned by <code>lme4::allFit()</code> , or a data frame with one row per optimiser-by-term combination (columns such as optimizer, term and value/estimate).
intercept	Whether to keep the intercept panel. Defaults to TRUE.
select_terms	Optional character vector of terms to display (the intercept is always kept when intercept = TRUE).
interaction	Passed to <code>format_terms()</code> for the panel titles.
labels	Optional named character vector renaming the term panels, e.g. <code>c(conditionunrelated = "Unrelated priming")</code> . When x is a raw <code>lme4::allFit()</code> object, a factor's coefficients are labelled with the variable and the level automatically (e.g. <code>conditionunrelated</code> becomes <code>condition: unrelated</code>) and these labels override that default.
number_optimizers	Whether to prefix each optimiser name with a number, so that the legend doubles as an index.
free_y	Whether to give each panel its own y-axis range. This is advisable, because the intercept and the slopes usually occupy very different scales.
ncol	Number of facet columns. If NULL, chosen automatically.
point_size	Point size.
palette	Colours for the optimisers; defaults to <code>depictr_palette()</code> .
y_lab, title	Axis label and plot title.

Details

Each fixed effect occupies its own panel, with its own y-axis. The function also accepts a plain data frame, so it can be used without 'lme4'.

Value

A `ggplot2::ggplot` object.

Examples

```
# Without lme4, build the input data frame directly:
set.seed(1)
df <- expand.grid(
  optimizer = c("bobyqa", "Nelder_Mead", "nlminbwrap"),
  term = c("(Intercept)", "rainfall", "fertiliser")
)
df$value <- c(5, 5.01, 4.99, 0.3, 0.31, 0.29, -0.2, -0.18, -0.21)
optimizer_fixef_plot(df)

if (requireNamespace("lme4", quietly = TRUE)) {
  m <- lme4::lmer(life_satisfaction ~ stress + (1 | region),
    data = wellbeing_survey)
  af <- lme4::allFit(m)
```

```
optimizer_fixef_plot(af)
}
```

outlier_plot

Box / violin plot highlighting outliers

Description

Shows the distribution of a numeric variable (optionally by group) as a box and/or violin plot, with points beyond the $1.5 * \text{IQR}$ fences highlighted so that outliers are easy to spot before modelling.

Usage

```
outlier_plot(
  data,
  y,
  group = NULL,
  type = c("box", "violin", "both"),
  flag = TRUE,
  outlier_colour = depictr_accent(),
  palette = NULL,
  title = NULL,
  y_lab = NULL
)
```

Arguments

data	A data frame.
y	The numeric variable (string or unquoted name).
group	Optional grouping variable on the x-axis. A numeric code counts as a group, but a numeric column with more than 20 distinct values is refused as continuous. Rows with a missing group are dropped with a message, and levels without rows are left out of the plot and its palette.
type	One of "box", "violin" or "both".
flag	Whether to highlight outliers (points beyond $1.5 * \text{IQR}$).
outlier_colour	Colour for highlighted outliers.
palette	Colours for the groups; defaults to <code>depictr_palette()</code> .
title, y_lab	Plot title and value-axis label.

Value

A `ggplot2::ggplot` object.

Examples

```
outlier_plot(crop_yield, yield)
outlier_plot(lexical_decision, RT, group = condition, type = "both")
```

palette_preview	<i>Preview the depictr palettes</i>
-----------------	-------------------------------------

Description

Displays the colours returned by `depictr_palette()` as labelled swatches. It is useful when choosing how many groups to show, selecting a palette type, or documenting a figure. Each swatch's hex label is drawn in black or white, whichever has the higher WCAG contrast against that tile, so every label reaches the 4.5:1 that WCAG 2.2 asks of text.

Usage

```
palette_preview(
  n = 8,
  type = c("qualitative", "sequential", "diverging", "all"),
  cvd = c("none", "deutan", "protan", "tritan")
)
```

Arguments

<code>n</code>	Number of colours to preview.
<code>type</code>	Palette type to preview: "qualitative", "sequential", "diverging", or "all" to show all three.
<code>cvd</code>	Colour-vision-deficiency simulation for the tiles: "none" (the default, true colours), "deutan" (red-green, deuteranopia), "protan" (red-green, protanopia) or "tritan" (blue-yellow, tritanopia).

Details

Set `cvd` to render the swatches as they would appear under a colour-vision deficiency, using the Machado et al. (2009) simulation. The tiles are recoloured to the simulated appearance while the labels keep the *original* hex code, so you can see whether two groups would still be told apart by a colourblind reader.

Value

A `ggplot2::ggplot` object.

References

Machado GM, Oliveira MM, Fernandes LAF (2009). "A physiologically-based model for simulation of color vision deficiency." *IEEE Transactions on Visualization and Computer Graphics*, **15**(6), 1291–1298. doi:10.1109/TVCG.2009.113.

Examples

```
palette_preview()
palette_preview(7, type = "sequential")
palette_preview(type = "all")
palette_preview(cvd = "deutan")
```

palette_safety

Check that a palette stays distinguishable under each deficiency

Description

For normal vision and each deficiency at full severity, the palette's colours are converted to CIELAB and the smallest colour difference between any two of them is found. The difference is CIEDE2000 (Sharma, Wu & Dalal, 2005) unless `metric` asks for the older measure. The lower this minimum, the more likely two categories are to be confused. A palette counts as safe when the minimum across all four conditions is at least `threshold`.

Usage

```
palette_safety(colours = NULL, threshold = 5, metric = c("cie2000", "cie76"))
```

Arguments

<code>colours</code>	The palette to test. Defaults to the <code>depictr</code> qualitative palette. Needs at least two colours: a pairwise distance over fewer than two has no value, rather than an infinitely safe one.
<code>threshold</code>	The smallest acceptable colour difference, on the scale of <code>metric</code> .
<code>metric</code>	The colour difference to measure: "cie2000", the default, or "cie76", the Euclidean distance in CIELAB that earlier versions measured. "cie76" is kept so that earlier results can be reproduced, and a later release will remove it.

Details

Earlier versions measured the Euclidean distance in CIELAB, known as CIE76, which overstates the difference between saturated colours: the blues #005b96 and #0060a8 are 8.20 apart by CIE76 and 2.69 by CIEDE2000. `ggplot2`'s default palette of five colours passed the check under CIE76, although under deuteranopia its sky blue and its pink come within 2.74 of each other in CIEDE2000.

The default `threshold` of 5 is calibrated against the reference colourblind-safe palette: the Okabe-Ito set's tightest pair (reddish purple against grey) sits at 6.43 under full deuteranopia, so the cut must lie below that to pass the recommended palette, while still flagging colours that become near-identical under a deficiency. The difference includes lightness, which survives colour-vision deficiency, so two colours that share a hue but differ in lightness are correctly treated as distinguishable. Full severity is the worst case; most colour-vision deficiency is milder.

This looks at a palette in the abstract. To audit a finished figure, which uses only as many colours as it has groups and has text and a background besides, see [check_figure\(\)](#).

Value

A list with `min_delta_e` (the worst case across conditions), `by_condition` (a named numeric vector, one minimum colour difference for normal vision and for each deficiency), `worst_condition` and `worst_pair` (the closest colours and where they were closest), `safe` and `threshold`.

References

Okabe M, Ito K (2008). “Color Universal Design (CUD): How to make figures and presentations that are friendly to colorblind people.” <https://jfly.uni-koeln.de/color/>. Accessed 2026-06-14.

Machado GM, Oliveira MM, Fernandes LAF (2009). “A physiologically-based model for simulation of color vision deficiency.” *IEEE Transactions on Visualization and Computer Graphics*, **15**(6), 1291–1298. doi:10.1109/TVCG.2009.113.

Sharma G, Wu W, Dalal EN (2005). “The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations.” *Color Research & Application*, **30**(1), 21–30. doi:10.1002/col.20070.

Examples

```
palette_safety()

# Two colours a hair apart are flagged.
palette_safety(c("#005b96", "#015c97"))

# ggplot2's default palette of five colours, which passed under CIE76.
hue5 <- c("#F8766D", "#A3A500", "#00BF7D", "#00B0F6", "#E76BF3")
palette_safety(hue5)$safe
palette_safety(hue5, metric = "cie76")$safe
```

pca_plot

*PCA biplot***Description**

Runs a principal component analysis on the numeric columns of a data frame (or takes an existing `stats::prcomp()` object) and draws a biplot: the observations projected onto two components, with the variable loadings shown as arrows. Optionally colour the observations by a grouping variable.

Usage

```
pca_plot(
  x,
  cols = NULL,
  group = NULL,
  components = c(1, 2),
  scale = TRUE,
```

```

loadings = TRUE,
point_alpha = 0.7,
palette = NULL,
title = NULL
)

```

Arguments

<code>x</code>	A data frame, or a <code>stats::prcomp()</code> object.
<code>cols</code>	When <code>x</code> is a data frame, the numeric columns to analyse (default: all numeric).
<code>group</code>	Optional grouping variable mapped to colour. When <code>x</code> is a data frame, give a column of <code>x</code> , unquoted or as a string, or a vector with one value per row of <code>x</code> . Rows with missing values count, although the analysis drops them. When <code>x</code> is a <code>stats::prcomp()</code> object, give a vector with one value per observation. A name that matches no column, or a vector of another length, is an error.
<code>components</code>	Length-2 integer vector: which components to plot.
<code>scale</code>	Whether to scale the variables to unit variance before the PCA. This is advisable when the variables are on different scales.
<code>loadings</code>	Whether to draw variable-loading arrows.
<code>point_alpha</code>	Point transparency.
<code>palette</code>	Colours for the groups; defaults to <code>depictr_palette()</code> .
<code>title</code>	Plot title.

Details

Each arrow is a variable's loadings on the two components, taken from the rotation matrix and multiplied by a factor that all the arrows share. The factor makes the largest coordinate of any arrow three quarters of the largest absolute score. Both axes use that one factor, as `stats::biplot()` uses one ratio for both, and the aspect ratio is fixed at 1. An arrow's angle on the page is therefore its angle in loading space, whatever the size of the figure. The panel's height and width follow the ranges of its two axes, so a figure of another shape leaves blank space around the panel. Because the scores are drawn as they are and the arrows as loadings, the angle between two arrows does not estimate the correlation between their variables. Arrows and labels are drawn in a neutral dark grey (`grey20`), which is not in the default palette, so no arrow reads as part of a group.

Value

A `ggplot2::ggplot` object.

Examples

```

pca_plot(crop_yield, cols = c("rainfall", "fertiliser", "soil_ph", "yield"),
         group = "treatment")

```

posterior_plot	<i>Plot posterior distributions</i>
----------------	-------------------------------------

Description

Displays posterior (or, more generally, bootstrap or simulation) draws as a *distribution* per parameter, in the style of a half-eye (a density slab with a point-and-interval beneath it), an interval-only forest plot, a gradient interval or a dotplot. The full shape of the posterior is shown, so skew, multimodality and the relative mass on either side of a reference value are all visible, rather than a point and two limits alone. The intervals are central (equal-tailed) quantile intervals, which leave the same share of the draws in either tail.

Usage

```
posterior_plot(
  draws,
  style = c("halfeye", "interval", "gradient", "dots"),
  point = c("median", "mean"),
  widths = c(0.66, 0.95),
  interaction = c("times", "asterisk", "colon", "space"),
  labels = NULL,
  reference_line = 0,
  rope = NULL,
  pd = FALSE,
  facet = FALSE,
  scales = c("fixed", "free"),
  colour = depictr_brand(),
  fill = depictr_brand(),
  rope_fill = depictr_reference(),
  rope_alpha = 0.15,
  title = NULL,
  x_lab = "Value",
  caption = NULL
)
```

Arguments

draws	Posterior draws: a fitted model (brms/rstanarm), a posterior draws object, a matrix, or a long/wide data frame (see Details).
style	One of "halfeye" (density slab + interval, the default), "interval" (point and two nested intervals, no slab), "gradient" (a colour-graded interval) or "dots" (a quantile dotplot). Any other value is an error. When 'ggdist' is missing, every style falls back to "interval".
point	Central summary: "median" or "mean".
widths	Two interval widths (inner and outer), as probabilities, each drawn as a central (equal-tailed) quantile interval. The outer width is used for the caption and the displayed interval mass.

interaction	Passed to <code>format_terms()</code> for the parameter labels.
labels	Optional named character vector renaming parameters, e.g. <code>c(conditionunrelated = "condition")</code> . Unmatched parameters fall back to <code>format_terms()</code> . For a brms fit, or a draws object holding <code>b_</code> parameters, the names are given without the <code>b_</code> prefix. An unnamed vector gives one label per parameter, in the order the parameters appear in draws.
reference_line	Position of a vertical reference line, or <code>NULL/NA</code> to omit it. There is no universally meaningful reference for every parameter, so this defaults to 0 (the usual "no effect" line for differences and slopes) but should be set or cleared deliberately.
rope	Optional length-2 numeric <code>c(lo, hi)</code> giving a region of practical equivalence to shade behind the distributions.
pd	If <code>TRUE</code> , annotate each parameter with its probability of direction relative to <code>reference_line</code> , the posterior probability that the parameter lies on its majority side of the reference (a value in <code>[0.5, 1]</code>). Requires a finite <code>reference_line</code> .
facet	Whether to give each parameter its own panel with a free x-axis, laid out one per row. This keeps the small parameters legible when a large one (typically the intercept) would otherwise stretch the shared axis and squish the rest. Defaults to <code>FALSE</code> . A convenience alias for <code>scales = "free"</code> .
scales	Either <code>"fixed"</code> (the default, a single shared x-axis) or <code>"free"</code> (one free-scaled panel per parameter). When <code>facet = TRUE</code> this is forced to <code>"free"</code> .
colour, fill	Colours for the point/interval and for the density slab. Default to the depictr brand blue.
rope_fill, rope_alpha	Fill colour and opacity of the ROPE band.
title, x_lab	Plot title and value-axis label.
caption	Plot caption. The default (<code>NULL</code>) auto-captions the interval mass; pass <code>NA</code> to omit a caption.

Details

Draws may be supplied in many shapes: a fitted Bayesian model (brms or rstanarm), a posterior *draws* object, a wide data frame or matrix with one column of draws per parameter, or a long data frame (a parameter column and a value column). For fitted models only the fixed-effect (population) parameters are kept, with the brms `b_` prefix stripped. A draws object that holds `b_` parameters, as `posterior::as_draws_df()` returns for a brms fit, is read the same way, with a message naming the variables it dropped. Pass it through `as.data.frame()` to keep the other variables, as a model of your own whose names happen to start with `b_` may need. The log densities `lp__` and `lprior` are never drawn as parameters. Two parameters that would share a display label, such as the `b_Intercept` and centred `Intercept` of a plain table of brms draws, are an error until one is renamed with `labels`. A matrix of posterior summaries, such as `brms::fixef()` returns, holds no draws and is refused.

The slab styles use the 'ggdist' package. When 'ggdist' is not installed the function falls back to the point-and-nested-interval display (equivalent to `style = "interval"`) and emits an informative message.

Value

A `ggplot2::ggplot` object.

Examples

```
# Wide draws: one column per parameter
set.seed(1)
draws <- data.frame(
  intercept = rnorm(400, 5, 0.3),
  slope = rnorm(400, 0.8, 0.15),
  `slope:group` = rnorm(400, -0.2, 0.2),
  check.names = FALSE
)
posterior_plot(draws)

# A region of practical equivalence and a probability-of-direction label
posterior_plot(draws, rope = c(-0.1, 0.1), pd = TRUE)

# When one parameter dwarfs the others, give each its own free-scaled panel:
posterior_plot(draws, facet = TRUE)
```

power_curve_plot

Plot a power analysis curve

Description

Draws statistical power against sample size, with a dashed line at a target power (80% by default). The input is usually a power curve produced by `simr::powerCurve()`, though a tidy data frame works equally well, allowing the plot to be redrawn without repeating a power simulation that is often slow to run.

Usage

```
power_curve_plot(
  x,
  target = 0.8,
  x_lab = "Sample size",
  x_breaks = NULL,
  x_expand = NULL,
  ribbon = TRUE,
  title = NULL,
  interaction = c("times", "asterisk", "colon", "space"),
  n = NULL,
  power = NULL,
  group = NULL,
  lower = NULL,
  upper = NULL,
```

```

    y_lab = NULL,
    solved = NULL
  )

```

Arguments

x	A powerCurve object from 'simr', or a data frame with a sample size column (n, nlevels, x, sample_size or n_subject), a power column (power, mean or p), and optional confidence limits (<power>_lo/<power>_hi, lower/upper, ci_low/ci_high or conf.low/conf.high).
target	Target power, drawn as a horizontal reference line. Use NA to omit it.
x_lab	X-axis label.
x_breaks	Approximate number of x-axis breaks.
x_expand	Optional value(s) to extend the x-axis to.
ribbon	Whether to draw the confidence band as a shaded ribbon (TRUE) or as error bars (FALSE).
title	Plot title. If NULL and x is a 'simr' power curve, the predictor name stored in the object is used.
interaction	Passed to <code>format_terms()</code> when deriving the title from a 'simr' object.
n, power	Columns holding the sample size and the power, as strings or unquoted names. By default, they are found among the names listed under x. power may name any proportion on the 0-1 scale, such as pilotr's p_meaningful.
group	Column that separates several curves, as a string or an unquoted name. Each level is drawn as its own line, in its own colour and line type. By default, an effect, param or term column with more than one level is used, with a message. Rows that share a sample size with no such column are an error, because joining them would draw a zigzag.
lower, upper	Columns holding the lower and upper confidence limits. Name both or neither. By default, they are found among the names listed under x.
y_lab	Y-axis label. By default, "Power" when the power column is one of the names listed under x, and the column's name otherwise.
solved	A sample size solved from the curve, with its confidence interval: a list with numeric value, lo and hi, as pilotr's target_n() and solve_curve() return, or a one-row data frame with the same columns. Other fields are ignored. The solved value is drawn as a dashed vertical line and its interval as a grey band behind the curve. An interval that reaches past the swept sizes widens the x-axis to show it.

Details

A table may hold several curves, one per effect or design, stacked by row. Each is drawn as its own line, told apart by colour and line type. The tables pilotr returns are read as they are: `power_curve_mixed()` gives n_subject, effect, power, power_lo and power_hi, and `precision_curve()` gives n_subject, param and the decision probabilities p_meaningful and p_equivalent with their own _lo and _hi bounds, which power can name. The sample size that pilotr's target_n() solves from such a curve, and its interval, can be marked on the plot with solved.

Value

A `ggplot2::ggplot` object.

Examples

```
pc <- data.frame(
  nlevels = c(10, 20, 30, 40, 50, 60),
  mean = c(0.18, 0.34, 0.52, 0.66, 0.79, 0.88),
  lower = c(0.10, 0.25, 0.42, 0.56, 0.70, 0.81),
  upper = c(0.28, 0.44, 0.62, 0.75, 0.86, 0.93)
)
power_curve_plot(pc, title = "Power for the condition effect")

# Two curves in one table, laid out as pilotr's power_curve_mixed() returns
# them: one row per sample size and effect.
curves <- data.frame(
  n_subject = rep(c(20, 40, 60, 80), each = 2),
  effect = rep(c("condition", "frequency"), times = 4),
  power = c(0.31, 0.12, 0.58, 0.21, 0.78, 0.33, 0.90, 0.45),
  power_lo = c(0.23, 0.07, 0.48, 0.14, 0.69, 0.25, 0.83, 0.36),
  power_hi = c(0.41, 0.19, 0.67, 0.30, 0.85, 0.43, 0.94, 0.55)
)
power_curve_plot(curves, group = "effect", x_lab = "Participants")

# The sample size at which pc reaches 80% power, with its interval, in the
# form pilotr's target_n() returns. The values here are illustrative.
power_curve_plot(pc, solved = list(value = 51, lo = 44, hi = 60))
```

pr_curve_plot

*Precision-recall curve***Description**

Plots precision against recall for a binary classifier and reports the average precision (the area under the curve). The precision-recall curve is more informative than the ROC curve when the positive class is rare, because it ignores the many true negatives. A horizontal line marks the no-skill baseline (the positive-class prevalence). Pass a *named list* of models / (actual, score) pairs to overlay several colour-coded curves with a legend and per-curve average precision.

Usage

```
pr_curve_plot(
  x,
  score = NULL,
  colour = depictr_brand(),
  f1 = FALSE,
  title = NULL
)
```

Arguments

x	A binomial glm; the vector of observed outcomes (0/1, logical or a two-level factor with the positive class second); or a <i>named</i> list of models / (actual, score) pairs to overlay.
score	When x is an outcome vector, the matching scores or predicted probabilities (or a list of them for the multi-model case).
colour	Curve colour for the single-model case. Defaults to the depictr brand blue. Ignored when several models are overlaid.
f1	Logical; if TRUE, mark the maximum-F1 operating point on each curve (the threshold maximising the harmonic mean of precision and recall).
title	Plot title.

Value

A `ggplot2::ggplot` object. The average precision is stored in `attr(plot, "average_precision")` (a named vector when several models are supplied).

Examples

```
# The clinical-trial adverse event is rare (about 10% positive), which is
# where a precision-recall curve is worth drawing.
gfit <- glm(adverse_event ~ biomarker + age + arm,
            data = clinical_trial, family = binomial)
pr_curve_plot(gfit)

# Mark the maximum-F1 operating point.
pr_curve_plot(gfit, f1 = TRUE)

# Compare two models with per-curve average precision in the legend.
reduced <- glm(adverse_event ~ biomarker, data = clinical_trial,
               family = binomial)
pr_curve_plot(list(Full = gfit, Reduced = reduced))
```

qq_plot	<i>Normal quantile-quantile plot</i>
---------	--------------------------------------

Description

A normal Q-Q plot for a numeric vector or for the standardised residuals of a fitted model, with a reference line and an optional confidence band that makes departures from normality easier to judge. Points that stray outside the band are unusual under normality; under a normal sample roughly level of the points fall inside a pointwise band.

Usage

```
qq_plot(
  x,
  colour = depictr_brand(),
  title = NULL,
  x_lab = "Theoretical quantiles",
  y_lab = NULL,
  band = TRUE,
  band_type = c("pointwise", "simulate"),
  level = 0.95,
  n_sim = 1000,
  band_fill = depictr_reference(),
  seed = NULL
)
```

Arguments

<code>x</code>	A numeric vector, or a fitted <code>lm/glm</code> model (its standardised residuals are used).
<code>colour</code>	Point colour. Defaults to the depictr brand blue.
<code>title, x_lab, y_lab</code>	Title and axis labels.
<code>band</code>	Whether to draw a confidence band/envelope.
<code>band_type</code>	Band construction: "pointwise" (analytic order-statistic standard errors, the default) or "simulate" (a Monte-Carlo envelope).
<code>level</code>	Confidence level for the band.
<code>n_sim</code>	Number of simulations for <code>band_type = "simulate"</code> .
<code>band_fill</code>	Fill colour of the band.
<code>seed</code>	Optional integer seed for the simulated envelope, for reproducibility.

Details

Two band constructions are offered. "pointwise" is analytic: it uses the large-sample standard error of the i -th order statistic, $se = \frac{\hat{\sigma}}{\phi(z_i)} \sqrt{p_i(1 - p_i)/n}$, around the fitted reference line. "simulate" builds a Monte-Carlo envelope by repeatedly drawing normal samples of the same size and taking the empirical quantiles of the simulated order statistics, which needs no large-sample approximation.

The band measures departures from a normal distribution, which the residuals of a linear model should follow. The standardised deviance residuals of a non-Gaussian `glm` are not expected to be normal, even when the model is correctly specified: those of a binary model take only two values at each fitted value. Points outside the band then say little about the fit. For a `glm`, [residual_diagnostics_plot\(\)](#) and [model_report\(\)](#) plot randomised quantile residuals (Dunn & Smyth, 1996), which are standard normal under a correct model.

Value

A `ggplot2::ggplot` object.

References

Dunn PK, Smyth GK (1996). “Randomized quantile residuals.” *Journal of Computational and Graphical Statistics*, **5**(3), 236–244. doi:[10.1080/10618600.1996.10474708](https://doi.org/10.1080/10618600.1996.10474708).

Examples

```
qq_plot(rnorm(100))
qq_plot(rt(100, df = 3), band_type = "simulate")
fit <- lm(yield ~ rainfall + fertiliser, data = crop_yield)
qq_plot(fit)
```

raincloud_plot	<i>Raincloud plot</i>
----------------	-----------------------

Description

A "raincloud" combines three views of a distribution: a half-violin density (the cloud), a narrow boxplot, and the raw jittered points (the rain). It conveys the shape, the summary and the individual observations together, giving a fuller and more transparent picture than a boxplot alone. The plot is built with base R and ggplot2 alone, with no extra package dependency.

Usage

```
raincloud_plot(
  data,
  y,
  group = NULL,
  width = 0.4,
  point_alpha = 0.4,
  palette = NULL,
  title = NULL,
  x_lab = NULL,
  y_lab = NULL
)
```

Arguments

data	A data frame.
y	The numeric variable (string or unquoted name).
group	Optional grouping variable on the x-axis. A numeric code counts as a group, but a numeric column with more than 20 distinct values is refused as continuous. Rows with a missing group are dropped with a message, and levels without rows are left out of the plot and its palette.
width	Maximum width of the half-violin.
point_alpha	Transparency of the rain points.
palette	Colours for the groups; defaults to <code>depictr_palette()</code> .
title, x_lab, y_lab	Title and axis labels.

Details

Groups with fewer than two observations cannot have a density estimated, so their half-violin is omitted (the points and box are still drawn) and a warning is issued.

Value

A `ggplot2::ggplot` object.

Examples

```
raincloud_plot(lexical_decision, RT, group = condition)
raincloud_plot(crop_yield, yield, group = treatment)
```

random_effects_plot	<i>Caterpillar plot of random effects</i>
---------------------	---

Description

Displays the conditional modes ("BLUPs") of a mixed model's random effects as a sorted point-and-interval ("caterpillar") plot. It is the usual way to inspect by-group departures from the average, and to identify unusual groups.

Usage

```
random_effects_plot(
  x,
  conf_level = 0.95,
  sort = TRUE,
  point_colour = depictr_brand(),
  title = NULL,
  x_lab = "Random effect"
)
```

Arguments

<code>x</code>	Either a mixed model fitted with 'lme4' (<code>merMod</code>), or a data frame with one row per group level and columns such as <code>level/group</code> , <code>estimate</code> , and either <code>conf.low/conf.high</code> or <code>std.error</code> . An optional <code>term</code> column facets the plot. A model may have several grouping factors and several terms per factor, correlated or not. The forms $(1 + x g)$, $(1 + x g)$ and $(1 g) + (0 + x g)$ are all drawn. The intervals are the conditional modes plus or minus the normal quantile times the conditional standard deviations that <code>lme4::ranef(condVar = TRUE)</code> reports. A model with one term in all is drawn in a single panel labelled "Random effect", and otherwise each term has a panel labelled "factor: term".
<code>conf_level</code>	Confidence level when intervals are derived from standard errors, a single number strictly between 0 and 1.
<code>sort</code>	Whether to order the levels by their estimate.

point_colour Colour for the points and intervals.
 title, x_lab Plot title and value-axis label.

Value

A `ggplot2::ggplot` object.

Examples

```
# From a data frame of by-group estimates:
re <- data.frame(
  level = paste0("G", 1:10),
  estimate = sort(rnorm(10)),
  std.error = runif(10, 0.2, 0.5)
)
random_effects_plot(re)

if (requireNamespace("lme4", quietly = TRUE)) {
  m <- lme4::lmer(RT ~ condition + (1 | participant), data = lexical_decision)
  random_effects_plot(m)
}
```

`residual_diagnostics_plot`

Residual-diagnostics panel for a fitted model

Description

Combines the classic regression diagnostic plots (residuals against fitted values, a normal Q-Q plot of the standardised residuals, a scale-location plot, and residuals against leverage) into a single panel using 'patchwork'. The function works with `lm` and `glm` objects.

Usage

```
residual_diagnostics_plot(
  model,
  which = c("resid_fitted", "qq", "scale_location", "resid_leverage"),
  ncol = 2,
  point_alpha = 0.6,
  smooth = TRUE,
  title = NULL,
  glm_panels = TRUE,
  seed = NULL
)
```

Arguments

model	A fitted lm or glm model.
which	Character vector choosing which panels to show, any of "resid_fitted", "qq", "scale_location" and "resid_leverage".
ncol	Number of columns in the panel layout.
point_alpha	Point transparency.
smooth	Whether to add a loess guide line to the residual panels.
title	Overall title for the panel.
glm_panels	For a glm, whether to use the GLM-appropriate binned and quantile-residual panels (the default) instead of the classic lm panels. Ignored for an lm.
seed	Optional integer seed for the randomisation of discrete quantile residuals, for reproducible glm Q-Q panels.

Details

For a glm, the default panels are made GLM-aware (`glm_panels = TRUE`): the "residuals vs. fitted" panel is replaced by a *binned*-residual plot (Gelman & Hill, 2007; see [binned_residual_plot\(\)](#)), which is readable even for binary or count models, and the Q-Q panel uses *randomised quantile* residuals (Dunn & Smyth, 1996), which are standard normal under a correctly specified model regardless of the response family. In a Gaussian or Gamma model, prior weights enter these residuals as they enter the fit, through each observation's variance. The scale-location and leverage panels are unchanged. Set `glm_panels = FALSE` to force the classic lm-style panels for a glm as well.

Value

A 'patchwork' object (printable like a [ggplot2::ggplot](#)).

References

- Gelman A, Hill J (2007). *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press, Cambridge, UK. ISBN 978-0-521-68689-1. doi:[10.1017/CBO9780511790942](#).
- Dunn PK, Smyth GK (1996). "Randomized quantile residuals." *Journal of Computational and Graphical Statistics*, **5**(3), 236–244. doi:[10.1080/10618600.1996.10474708](#).

Examples

```
fit <- lm(yield ~ rainfall + fertiliser + soil_ph, data = crop_yield)
residual_diagnostics_plot(fit)

residual_diagnostics_plot(fit, which = c("resid_fitted", "qq"))

# A logistic GLM gets binned and quantile-residual panels automatically:
gfit <- glm(adverse_event ~ biomarker + age + arm,
           data = clinical_trial, family = binomial)
residual_diagnostics_plot(gfit)
```

ridgeline_plot

Ridgeline plot

Description

Shows the distribution of a numeric variable across the levels of a grouping variable as a column of partially overlapping density curves, one ridge per group, stacked up the y-axis. It compares many distributions in a compact space, making shifts in location and shape easy to follow, and is a clearer choice than many overlaid densities once there are several groups.

Usage

```
ridgeline_plot(
  data,
  x,
  group,
  overlap = 1.4,
  alpha = 0.85,
  scale_height = TRUE,
  palette = NULL,
  title = NULL,
  x_lab = NULL,
  y_lab = NULL
)
```

Arguments

data	A data frame.
x	The numeric variable whose distribution is shown (string or unquoted name).
group	The grouping variable (string or unquoted name); one ridge per level. A numeric code counts as a group, but a numeric column with more than 20 distinct values is refused as continuous. Rows with a missing group are dropped with a message, and levels without rows are left out of the plot and its palette.
overlap	How far each ridge extends into the next, as a multiple of the row spacing. 1 makes the tallest ridge just touch the next baseline; larger values overlap more. Defaults to 1.4.
alpha	Fill transparency of the ridges.
scale_height	Whether to scale every ridge to the same peak height (TRUE, the default) or keep the true relative densities (FALSE).
palette	Colours for the groups; defaults to depictr_palette() .
title, x_lab, y_lab	Plot title and axis labels.

Details

Densities are computed with `stats::density()` and scaled to a common height; the `overlap` argument controls how far each ridge reaches into the one above. Groups with fewer than two non-missing values are dropped with a warning. The implementation is base R + `ggplot2`, with no extra package dependency.

Value

A `ggplot2::ggplot` object.

Examples

```
ridgeline_plot(wellbeing_survey, life_satisfaction, region)
ridgeline_plot(lexical_decision, RT, condition)
```

<code>roc_curve_plot</code>	<i>ROC curve</i>
-----------------------------	------------------

Description

Plots the receiver operating characteristic (ROC) curve for a binary classifier and reports the area under the curve (AUC). The input can be a fitted binomial glm, a pair of vectors (observed binary outcome and a continuous score), or, to compare several models, a *named list* of models or of (actual, score) pairs, which are overlaid as colour-coded curves with a legend and a per-curve AUC. An optional bootstrap confidence band can be drawn for the (single-model) curve and its AUC, and the Youden's J operating point can be marked.

Usage

```
roc_curve_plot(
  x,
  score = NULL,
  colour = depictr_brand(),
  ci = FALSE,
  conf_level = 0.95,
  youden = FALSE,
  legend_inside = FALSE,
  title = NULL
)
```

Arguments

<code>x</code>	A binomial glm; the vector of observed outcomes (0/1, logical or a two-level factor with the positive class second); or a <i>named list</i> of models / (actual, score) pairs to overlay (see Details).
<code>score</code>	When <code>x</code> is an outcome vector, the matching vector of scores or predicted probabilities. When <code>x</code> is a named list of outcome vectors, a matching named/positional list of score vectors.

colour	Curve colour for the single-model case. Defaults to the <code>depictr</code> brand blue. Ignored when several models are overlaid (the colourblind-aware <code>scale_colour_depict()</code> palette is used instead).
ci	Bootstrap confidence band for a single ROC curve and its AUC. FALSE (default) draws none; TRUE uses 2000 resamples; a positive integer sets the number of resamples. Ignored when several models are overlaid.
conf_level	Confidence level for the bootstrap band, a single number strictly between 0 and 1.
youden	Logical; if TRUE, mark the Youden's J operating point (the threshold maximising sensitivity + specificity - 1) on each curve.
legend_inside	When TRUE (and several models are overlaid), draw the legend inside the panel (in the bottom-right corner the curve leaves empty) over a translucent background, instead of in a right-hand margin. Defaults to FALSE.
title	Plot title.

Details

A named list overlays one curve per element, e.g. `roc_curve_plot(list("Full" = fit_full, "Reduced" = fit_reduced))`. Each element may be a `glm`, a length-2 list/data frame of (actual, score), or an outcome vector paired with the matching element of a score list.

Value

A `ggplot2::ggplot` object. The AUC(s) are stored in `attr(plot, "auc")` (a named vector when several models are supplied).

Examples

```
gfit <- glm(accuracy ~ word_frequency + condition + RT,
           data = lexical_decision, family = binomial)
roc_curve_plot(gfit)

# Mark the Youden operating point and add a bootstrap band.
roc_curve_plot(gfit, youden = TRUE, ci = 200)

# Compare two models with a colour-coded legend and per-curve AUC.
reduced <- glm(accuracy ~ word_frequency, data = lexical_decision,
              family = binomial)
roc_curve_plot(list(Full = gfit, Reduced = reduced))
```

Description

A convenience wrapper around `ggplot2::ggsave()` with sensible defaults for figures in papers and reports: a moderate size, 300 dpi, and the output directory created if needed. The device is inferred from the file extension, and PDF, EPS and TIFF files are written to pass the technical checks journals run on submitted figures. With `check = TRUE`, the figure is also audited with `check_figure()` at the width it is saved at.

Usage

```
save_plot(
  filename,
  plot = ggplot2::last_plot(),
  width = 7,
  height = 4.5,
  units = "in",
  dpi = 300,
  check = FALSE,
  ...
)
```

Arguments

filename	Output file path. The extension sets the device (e.g. .png, .pdf, .eps, .tiff).
plot	The plot to save; defaults to the last plot drawn.
width, height	Dimensions.
units	Units for width and height.
dpi	Resolution for raster devices. An EPS rasterises its translucent layers, and whatever they overlap, at this resolution or at 600 dpi, whichever is higher.
check	Whether to audit the figure with <code>check_figure()</code> at the width it is saved at, given as both <code>width_cm</code> and <code>render_width_cm</code> , so that the text sizes it reports are the printed ones. A warning names each check the figure fails. Defaults to FALSE. The audit runs before the file is written, so a plot it cannot read is refused without saving. For other thresholds, such as a journal's own text floor, call <code>check_figure()</code> with the saved width as both widths.
...	Passed to <code>ggplot2::ggsave()</code> , which passes what it does not use on to the device. A device or device argument given here overrides the defaults described under "Formats journals accept".

Value

The filename, invisibly. With `check = TRUE`, the audit report is attached to it as the attribute "check", which `attr(, "check")` reads.

Formats journals accept

When no device is given, three formats are written differently from `ggplot2`'s defaults.

A PDF is written with `grDevices::cairo_pdf()`, which embeds its fonts. `ggplot2`'s default, `grDevices::pdf()`, names the fonts without embedding them.

An EPS or PostScript file is written with `grDevices::cairo_ps()`, with its fonts embedded too. PostScript has no transparency, and `grDevices::postscript()` leaves every translucent layer out of the file: the points and the confidence band of the plot in the examples below vanish. `cairo_ps()` rasterises those layers, together with whatever they overlap, at `fallback_resolution`, by default `max(dpi, 600)` here. A message says so when a layer of the plot, or of any panel of a patchwork, is translucent. A PDF keeps such layers as vector graphics.

A TIFF is compressed with LZW, which is lossless. Uncompressed, a figure of 7 by 4.5 inches at 300 dpi takes 8.1 MiB, and one at the largest size PLOS accepts takes more than the 10 MB PLOS allows. The TIFF is drawn on white unless `bg` is given, so a transparent background, as in `ggplot2::theme_void()`, leaves no alpha channel in the file.

Without `cairo`, a PDF is written with `pdf()` and a message says how to embed its fonts with `grDevices::embedFonts()`, and an EPS is written with `postscript()` and a warning. R for macOS from CRAN loads `cairo` only where XQuartz (<https://www.xquartz.org>) is installed, and a Mac with neither XQuartz nor `ragg` writes the TIFF uncompressed, with a warning from `grDevices::tiff()`. A device, or a device argument such as `fallback_resolution`, `compression` or `bg`, passed in ... is used as given, so `device = "pdf"`, `device = "eps"` or `device = "tiff"` restores `ggplot2`'s own behaviour.

Which formats three publishers accepted on 2 October 2026:

Format	PLOS	Nature (final artwork)	Elsevier
PDF	No	Yes, as vector art	Yes
EPS	Yes, with fonts embedded	Yes, as vector art	Yes, with fon
TIFF	Yes, LZW-compressed RGB or greyscale, no alpha channel, at most 10 MB	No	Yes

Nature also asks for fonts embedded as TrueType 2 or 42, never Type 3, and for line art and text that are not rasterised. A figure with translucent layers therefore goes to Nature as a PDF. Elsevier asks for a graph sent as a TIFF at 1000 dpi, so pass `dpi = 1000` for one.

The Python twin takes the plot first, as `save_plot(plot, filename)`.

References

PLOS ONE. Figures. <https://journals.plos.org/plosone/s/figures>

Nature. Guide to preparing final artwork. <https://www.nature.com/documents/nature-final-artwork.pdf>

Nature research figure guide. Preparing figures: our specifications. <https://research-figure-guide.nature.com/figures/preparing-figures-our-specifications/>

Elsevier. Artwork overview. <https://www.elsevier.com/about/policies-and-standards/author/artwork-and-media-instructions/artwork-overview>

Examples

```
p <- scatter_trend(crop_yield, fertiliser, yield)
tmp <- file.path(tempdir(), "yield.png")
```

```

save_plot(tmp, p)

# A TIFF well within the 10 MB that PLOS allows.
save_plot(file.path(tempdir(), "yield.tiff"), p)

# Audited at the 8.9 cm column it is saved for, where its text prints at
# the size it is drawn.
saved <- save_plot(file.path(tempdir(), "yield-column.png"), p,
                    width = 8.9, height = 6, units = "cm", check = TRUE)
attr(saved, "check")

```

scale_colour_depictr *depictr colour and fill scales*

Description

Discrete ggplot2 scales using [depictr_palette\(\)](#). These are the canonical colour and fill scales used throughout the package. They honour the global options(`depictr.palette =`) (via [depictr_palette\(\)](#)) and options(`depictr.na_value =`) settings; see [depictr_options\(\)](#).

Usage

```

scale_colour_depictr(
  n = NULL,
  palette = NULL,
  na.value = depictr_opt("na_value"),
  order = c("default", "contrast"),
  ...
)

scale_color_depictr(
  n = NULL,
  palette = NULL,
  na.value = depictr_opt("na_value"),
  order = c("default", "contrast"),
  ...
)

scale_fill_depictr(
  n = NULL,
  palette = NULL,
  na.value = depictr_opt("na_value"),
  order = c("default", "contrast"),
  ...
)

```

Arguments

n	Optional number of colours to draw from <code>depictr_palette()</code> . By default <code>ggplot2</code> requests exactly as many colours as there are groups; pass <code>n</code> only to force a fixed slice of the palette. More than eight groups means the built-in qualitative palette is interpolated, which loses the colour-vision-deficiency guarantee and warns at draw time; facet the groups or use a sequential scale when there are that many.
palette	Optional palette override: a function of one argument (the number of colours) returning a character vector of colours. Defaults to <code>depictr_palette()</code> .
na.value	Colour for NA levels. Defaults to the resolved <code>depictr.na_value</code> option (the muted grey "grey80" unless changed).
order	"default", or "contrast" to draw the groups in the contrast-first order of <code>depictr_palette()</code> , with which two or three groups pass every colour row of <code>check_figure()</code> . More than four groups warn at draw time. It cannot be combined with <code>palette</code> .
...	Passed to <code>ggplot2::discrete_scale()</code> .

Value

A `ggplot2` scale that can be added to a plot.

Examples

```
library(ggplot2)
ggplot(crop_yield, aes(rainfall, yield, colour = treatment)) +
  geom_point() +
  scale_colour_depict() +
  theme_depict()

# The vermillion in the orange's place, at 3.87:1 against white
ggplot(crop_yield, aes(rainfall, yield, colour = treatment,
                      shape = treatment)) +
  geom_point() +
  scale_colour_depict(order = "contrast") +
  theme_depict()
```

scatter_trend	<i>Scatter plot with a fitted trend</i>
---------------	---

Description

Plots `y` against `x` with an optional fitted trend line and confidence band, optionally split by a grouping variable.

Usage

```
scatter_trend(  
  data,  
  x,  
  y,  
  group = NULL,  
  method = "lm",  
  se = TRUE,  
  point_alpha = 0.6,  
  palette = NULL,  
  title = NULL,  
  x_lab = NULL,  
  y_lab = NULL  
)
```

Arguments

data	A data frame.
x, y	The variables for the horizontal and vertical axes (string or unquoted column name).
group	Optional grouping variable mapped to colour. A numeric code counts as a group, but a numeric column with more than 20 distinct values is refused as continuous. Rows with a missing group are dropped with a message, and levels without rows are left out of the plot and its palette.
method	Smoothing method passed to <code>ggplot2::geom_smooth()</code> , e.g. "lm", "loess", or NULL for no trend line.
se	Whether to draw the confidence band around the trend.
point_alpha	Point transparency.
palette	Colours for the groups; defaults to <code>depictr_palette()</code> .
title, x_lab, y_lab	Title and axis labels (default to variable names).

Value

A `ggplot2::ggplot` object.

Examples

```
scatter_trend(crop_yield, fertiliser, yield)  
scatter_trend(crop_yield, fertiliser, yield, group = treatment,  
              method = "lm")
```

scree_plot

Scree plot

Description

Shows the proportion of variance explained by each principal component, as bars with a cumulative line. It is the customary aid for deciding how many components to retain.

Usage

```
scree_plot(x, cols = NULL, scale = TRUE, n = NULL, title = NULL)
```

Arguments

x	A data frame, or a <code>stats::prcomp()</code> object.
cols	When x is a data frame, the numeric columns to analyse.
scale	Whether to scale variables to unit variance before the PCA.
n	Maximum number of components to display.
title	Plot title.

Value

A `ggplot2::ggplot` object.

Examples

```
scree_plot(wellbeing_survey,
            cols = c("age", "income", "stress", "sleep_hours",
                    "exercise_days", "life_satisfaction"))
```

seasonal_plot

Seasonal-subseries (cycle) plot

Description

Draws a seasonal-subseries plot (also called a cycle plot): the series is split into one small panel per season (e.g. one panel per calendar month for monthly data), and within each panel the value is drawn across successive cycles (e.g. across years). A horizontal reference line in every panel marks that season's mean. This makes the seasonal pattern (differences *between* panels) and the within-season trend over time (the slope *inside* each panel) visible at the same time, which a single overlaid line cannot show.

Usage

```
seasonal_plot(
  x,
  frequency = NULL,
  style = c("subseries", "season"),
  season_labels = NULL,
  means = TRUE,
  point = TRUE,
  palette = NULL,
  title = NULL,
  x_lab = NULL,
  y_lab = NULL
)
```

Arguments

<code>x</code>	A <code>ts</code> object, or a numeric vector (then frequency is required), holding one series. A multivariate <code>ts</code> or a matrix with several columns is refused.
<code>frequency</code>	Number of observations per period (e.g. 12 for monthly data); taken from <code>x</code> when it is a <code>ts</code> . A frequency that is not a whole number, such as the $365.25 / 7$ of weekly data, is rounded down to whole seasons, as <code>stats::stl()</code> rounds it, and a message says so. The seasons are then counted from the first observation.
<code>style</code>	"subseries" for the faceted cycle plot (one panel per season) or "season" for one line per cycle across the seasons.
<code>season_labels</code>	Optional character vector of length frequency giving the season (panel/axis) labels, e.g. <code>month.abb</code> . When <code>NULL</code> (default) sensible labels are inferred (month abbreviations for frequency 12, quarter labels for frequency 4, otherwise the season index).
<code>means</code>	Whether to draw the per-season mean reference line (<code>style = "subseries"</code> only).
<code>point</code>	Whether to add points as well as the connecting line.
<code>palette</code>	Colours used for the cycle lines when <code>style = "season"</code> . Defaults to a blue ramp from <code>#4a8fc0</code> through the brand blue to <code>#08315a</code> , whose palest colour has a contrast of 3.51:1 against white. The sequential ramp of <code>depictr_palette()</code> starts paler, at <code>#e6eff5</code> , which suits a fill but leaves a line close to invisible.
<code>title, x_lab, y_lab</code>	Title and axis labels.

Details

Two layouts are offered. With `style = "subseries"` (the default) the panels are faceted side by side, each tracing one season across cycles: the classic Cleveland cycle plot, matching `feasts::gg_subseries()`. With `style = "season"` every cycle is drawn as its own line over the seasons on a shared axis (the seasonal-plot layout of `forecast::ggseasonplot()`), which is handy for spotting an unusual year. The lines run from a mid blue for the first cycle to a dark blue for the last, so they read in time order. Past about six cycles, neighbouring cycles take colours too close to tell apart, and the colours then

show the order of the cycles more than which is which. When individual cycles must be told apart, use `style = "subseries"`, where each cycle keeps its own place on the axis of every panel, or plot fewer cycles.

Value

A `ggplot2::ggplot` object.

See Also

`decompose_plot()`, `timeseries_plot()`

Examples

```
# Monthly air passengers: rising within-month trend, summer peak
seasonal_plot(AirPassengers)
seasonal_plot(AirPassengers, style = "season")
```

<code>silhouette_plot</code>	<i>Silhouette plot</i>
------------------------------	------------------------

Description

Draws a silhouette plot for a clustering: one horizontal bar per observation, grouped and sorted within cluster, with the average silhouette width shown per cluster and overall. The silhouette width $s(i)$ compares how close an observation is to its own cluster against the nearest other cluster, and lies in $[-1, 1]$ (higher is better; negative suggests the point may belong to another cluster). It is a standard internal measure of cluster quality.

Usage

```
silhouette_plot(
  data,
  clusters,
  cols = NULL,
  scale = TRUE,
  distance = "euclidean",
  palette = NULL,
  title = NULL
)
```

Arguments

<code>data</code>	A data frame, a numeric matrix, or a <code>stats::dist</code> object.
<code>clusters</code>	A vector of cluster assignments, one per observation (per row of data, or matching the objects in a <code>dist</code>). Required.
<code>cols</code>	When data is a data frame, the numeric columns to use (default: all numeric columns).

scale	Whether to scale variables to unit variance before computing the distances (ignored when data is a dist).
distance	Distance measure passed to <code>stats::dist()</code> (ignored when data is already a dist).
palette	Colours for the clusters; defaults to <code>depictr_palette()</code> .
title	Plot title.

Details

The widths are computed with `cluster::silhouette()` when the 'cluster' package is installed, and otherwise with an equivalent base-R implementation, so the function works with no extra dependency.

Value

A `ggplot2::ggplot` object. The per-observation silhouette table is attached as the attribute "silhouette" and the average width as "avg_width". The table's cluster and neighbor columns carry the assignments passed in clusters, as does each band label.

References

Rousseeuw PJ (1987). "Silhouettes: A graphical aid to the interpretation and validation of cluster analysis." *Journal of Computational and Applied Mathematics*, **20**, 53–65. doi:10.1016/0377-0427(87)901257.

Examples

```
cl <- kmeans(scale(crop_yield[c("rainfall", "fertiliser", "soil_ph",
                                "yield")])), 3)$cluster
silhouette_plot(crop_yield, cl,
  cols = c("rainfall", "fertiliser", "soil_ph", "yield"))
```

simulate_cvd

Simulate how colours appear under a colour-vision deficiency

Description

Maps each colour to the colour a reader with a given form and severity of colour-vision deficiency would perceive, using the physiologically-based model of Machado, Oliveira and Fernandes (2009). The transformation is defined on linear-light RGB, so a colour is decoded from sRGB to linear RGB (the IEC 61966-2-1 transfer functions), transformed, then re-encoded.

Usage

```
simulate_cvd(colours, deficiency, severity = 1)
```

Arguments

colours	Character vector of colours, in any form <code>grDevices::col2rgb()</code> understands.
deficiency	The deficiency to simulate: "protan" or "deutan" (the two red-green forms) or "tritan" (blue-yellow).
severity	Severity in $[0, 1]$. Zero leaves the colours unchanged and one is the full deficiency. Machado et al. publish a matrix for each tenth of severity, which is used as it stands at 0.1, 0.2 and so on. A severity between two tenths takes the linear interpolation of their matrices, as the authors suggest and as the <code>colspace</code> package does.

Details

This is the simulator behind `palette_preview()`'s `cvd` argument, behind `palette_safety()`, and behind the colour-separability rows of `check_figure()`. `simulate_figure()` applies it to every colour of a finished figure.

Value

A character vector of lower-case hex colours, the same length as `colours`.

References

Machado GM, Oliveira MM, Fernandes LAF (2009). "A physiologically-based model for simulation of color vision deficiency." *IEEE Transactions on Visualization and Computer Graphics*, **15**(6), 1291–1298. doi:10.1109/TVCG.2009.113.

Examples

```
simulate_cvd(c("#005b96", "#e69f00"), "deutan")

# Half severity moves the colours only part of the way.
simulate_cvd(c("#005b96", "#e69f00"), "deutan", severity = 0.5)
```

simulate_figure	<i>Show a finished figure as a reader with a colour-vision deficiency sees it</i>
-----------------	---

Description

`simulate_cvd()` simulates a set of colours, and this function the figure they end up in, legend, colourbar, text and background included.

Usage

```
simulate_figure(plot, deficiency, severity = 1)
```

Arguments

plot	A ggplot or a patchwork, as it would be saved.
deficiency	The deficiency to simulate: "protan" or "deutan" (the two red-green forms) or "tritan" (blue-yellow).
severity	Severity in $[0, 1]$. Zero leaves the colours unchanged and one is the full deficiency. Machado et al. publish a matrix for each tenth of severity, which is used as it stands at 0.1, 0.2 and so on. A severity between two tenths takes the linear interpolation of their matrices, as the authors suggest and as the colorspace package does.

Details

The plot is turned into the table of grobs that would be drawn, with `ggplot2::ggplotGrob()` or, for a patchwork, `patchwork::patchworkGrob()`. Every colour in that table then goes through `simulate_cvd()`. That covers the line and fill colours of each grob, at any depth, and each cell of a raster, such as a colourbar, a `ggplot2::geom_raster()` layer or an image added with `ggplot2::annotation_raster()`. Each colour keeps its alpha channel. Missing and transparent colours are left as they are, and so are gradient and pattern fills such as `grid::linearGradient()`, which keep their original colours.

The colours are changed after the plot is built. Changing the colours of the data before it is built would leave the legend keys and any colourbar in their original colours, so the legend would no longer match the figure.

The figure stays vector graphics, so it can be saved at any size. The Python twin's `simulate_figure()` uses the same matrices but rasterises the figure and simulates its pixels, so it returns an image.

Value

A gtable, which `grid::grid.draw()` draws and `save_plot()` saves.

References

Machado GM, Oliveira MM, Fernandes LAF (2009). "A physiologically-based model for simulation of color vision deficiency." *IEEE Transactions on Visualization and Computer Graphics*, **15**(6), 1291–1298. doi:10.1109/TVCG.2009.113.

See Also

`check_figure()`, which measures what this shows.

Examples

```
library(ggplot2)
p <- ggplot(crop_yield, aes(rainfall, yield, colour = treatment)) +
  geom_point() +
  scale_colour_depict() +
  theme_depict()

# The figure as a reader with deuteranopia sees it.
```

```

seen <- simulate_figure(p, "deutan")
grid::grid.newpage()
grid::grid.draw(seen)

# Saved for a side-by-side comparison with the original.
save_plot(file.path(tempdir(), "yield-deutan.png"), seen)

```

summary_table

A "Table 1" style descriptive summary

Description

Builds the kind of descriptive table that opens many empirical papers: numeric variables are summarised as mean (SD), categorical variables as counts and percentages, optionally split into one column per level of a grouping variable. The result is a plain data frame, ready to pass to `knitr::kable()` or a table package.

Usage

```

summary_table(
  data,
  vars = NULL,
  group = NULL,
  digits = 1,
  missing = TRUE,
  max_levels = 20
)

```

Arguments

<code>data</code>	A data frame.
<code>vars</code>	Columns to summarise. If <code>NULL</code> , all columns except <code>group</code> are used, with high-cardinality identifier-like columns skipped (see Details).
<code>group</code>	Optional grouping variable; one summary column is produced per level, alongside an overall column. Records whose <code>group</code> value is missing get a trailing Missing column of their own, so the group sizes still sum to the overall <code>N</code> .
<code>digits</code>	Number of decimal places for numeric summaries.
<code>missing</code>	Whether to add a Missing, <code>n (%)</code> row for variables that contain missing values. Defaults to <code>TRUE</code> .
<code>max_levels</code>	When <code>vars</code> is <code>NULL</code> , character/factor columns whose distinct-level count is at least <code>max_levels</code> <i>and</i> exceeds half the number of rows are treated as identifiers and skipped. Defaults to 20.

Details

The first row of the table always reports the sample size (N) overall and per group. By default each variable is also followed by a Missing, n (%) row whenever it contains missing values; set `missing = FALSE` to suppress these.

A date or date-time column (Date, POSIXct) gets a single Range row, written as <earliest> to <latest> in each column. The cells give the date alone when every time in the column falls at midnight, and the time to the minute, or to the second when any time has seconds. A duration (difftime) is summarised as mean (SD) in its own units, which the statistic names, as in Mean (SD), days.

When `vars` is NULL, high-cardinality character/factor columns (those whose number of distinct levels approaches the number of rows, e.g. identifier columns) are skipped automatically, with a message, rather than being expanded into hundreds of one-per-row entries. Pass such a column explicitly via `vars` to override this.

Value

A data frame with columns `variable`, `statistic`, `Overall`, one column per group level and, when the grouping variable has missing values, a trailing `Missing` column. The first row reports N.

Examples

```
summary_table(crop_yield, vars = c("yield", "rainfall", "treatment"))
summary_table(wellbeing_survey,
              vars = c("life_satisfaction", "education"),
              group = "region")
```

survival_plot

Kaplan-Meier survival plot

Description

Draws Kaplan-Meier survival curves, optionally by group, with stepwise confidence limits and censoring marks. The Kaplan-Meier estimate and its Greenwood standard error are computed with base R, so no modelling package is required; a `survfit` object from the 'survival' package is also accepted.

Usage

```
survival_plot(
  time,
  status = NULL,
  group = NULL,
  conf_level = 0.95,
  censor_marks = TRUE,
  risk_table = FALSE,
  median_line = FALSE,
  logrank = FALSE,
```

```

risk_breaks = NULL,
palette = NULL,
legend_inside = FALSE,
title = NULL,
x_lab = "Time",
y_lab = "Survival probability",
conf_type = c("log-log", "log", "plain")
)

```

Arguments

time	A numeric vector of follow-up times, a data frame with time, status and optional group columns, or a <code>survfit</code> object. Every time must be finite: a missing or infinite one is an error rather than a silent exclusion, since dropping it would move the number at risk unannounced.
status	Event indicator: a vector, or a column name when time is a data frame (by default its status or event column). Either the 0/1 convention (0 = censored, 1 = event) or the <code>survival::Surv()</code> 1/2 convention (1 = censored, 2 = event) is accepted, read over all groups together; logical values are also allowed. Other codings raise an error, among them a code that is not a whole number and a factor or character indicator, whose values do not say which one is the event. Observations whose status is missing are dropped with a message, since they record neither an event nor a censoring.
group	Optional grouping variable: a vector, or a column name when time is a data frame (by default its group, strata or arm column, if any). Observations whose group is missing are dropped with a message, since they belong to no arm.
conf_level	Confidence level for the limits, a single number strictly between 0 and 1 (NA to omit them). For a <code>survfit</code> object, the limits are recomputed at this level, even for one made with <code>conf.type = "none"</code> , since it keeps its standard errors. An object that stores no standard errors, or whose lower limits follow <code>conf.lower = "peto"</code> or <code>"modified"</code> , keeps the limits it holds, with a warning when they are not the ones asked for.
censor_marks	Whether to mark censoring times with a +.
risk_table	Whether to add a number-at-risk table beneath the curves (composed with 'patchwork'). Defaults to FALSE.
median_line	Whether to draw dashed guides to each group's median survival time and label it. Defaults to FALSE.
logrank	Whether to add a log-rank test of the group difference as a subtitle (two or more groups only). Defaults to FALSE.
risk_breaks	Optional numeric vector of times at which to report the number at risk. Defaults to the curve's x-axis breaks.
palette	Colours for the groups; defaults to <code>depictr_palette()</code> .
legend_inside	When TRUE (and there are several groups), draw the group legend inside the panel (in the bottom-left corner a monotone-decreasing survival curve always leaves empty) over a translucent background, instead of in a right-hand margin. Defaults to FALSE.

title, x_lab, y_lab	Title and axis labels.
conf_type	The scale on which the confidence limits are built: "log-log" (the default), "log" or "plain", as set out under Details.

Details

The confidence limits are pointwise. By default, they are built on the log-log scale, as $S^{\exp(\pm z\sigma / \log S)}$, where σ is the Greenwood standard error of $\log S$. These limits stay between 0 and 1, and in small samples they cover better than the untransformed limits $S \pm zS\sigma$ (Borgan & Liestøl, 1990). `conf_type = "plain"` draws the untransformed limits, clipped to that range, and `conf_type = "log"` draws $S \exp(\pm z\sigma)$, capped at 1, as `survival::survfit()` does by default. The limits are 1 while the estimate is 1, and where a curve falls to 0, they keep the values of the step before. The limits of a `survfit` object are recomputed from its standard errors at `conf_level` and `conf_type`, so they equal those that `survfit()` gives with the same `conf.int` and `conf.type` (`conf_level` names the objects that keep their own).

For publication-ready figures the plot offers the three annotations that define a "survminer-style" Kaplan-Meier display, each behind its own argument and off by default:

- **Number-at-risk table** (`risk_table = TRUE`): a small panel composed beneath the curves with 'patchwork', giving the number of subjects still at risk in each group at the x-axis breaks.
- **Median-survival guides** (`median_line = TRUE`): dashed reference lines from the 0.5 survival level down to each group's median survival time, with the median value labelled. The median is the time at which the curve falls below 0.5. Where the curve stays at 0.5 until a later drop, or until its last time, the median is the midpoint of that stretch, as in `survival::quantile.survfit()`. Groups whose curve never reaches 0.5 (median not estimable) are skipped.
- **Log-rank test** (`logrank = TRUE`): for two or more groups, the chi-squared log-rank statistic and its p-value are added as a subtitle. `survival::survdiff()` is used when the 'survival' package is installed, otherwise an equivalent base-R log-rank test is computed from the risk/event tables.

Value

A `ggplot2::ggplot` object or, when `risk_table = TRUE`, a 'patchwork' object stacking the curves above the risk table.

References

- Kaplan EL, Meier P (1958). "Nonparametric estimation from incomplete observations." *Journal of the American Statistical Association*, **53**(282), 457–481. doi:10.1080/01621459.1958.10501452.
- Greenwood M (1926). *The natural duration of cancer*, volume 33 of *Reports on Public Health and Medical Subjects*. His Majesty's Stationery Office, London.
- Borgan Ø, Liestøl K (1990). "A note on confidence intervals and bands for the survival function based on transformations." *Scandinavian Journal of Statistics*, **17**(1), 35–41.
- Mantel N (1966). "Evaluation of survival data and two new rank order statistics arising in its consideration." *Cancer Chemotherapy Reports*, **50**(3), 163–170.

Examples

```
set.seed(1)
n <- 200
grp <- sample(c("control", "treated"), n, replace = TRUE)
time <- rexp(n, rate = ifelse(grp == "treated", 0.05, 0.1))
cens <- runif(n, 0, 30)
obs <- pmin(time, cens)
event <- as.integer(time <= cens)
survival_plot(obs, event, group = grp, title = "Survival by treatment")

# survminer-style figure: risk table, median guides and a log-rank test
data(clinical_trial)
survival_plot(clinical_trial$time, clinical_trial$event,
              group = clinical_trial$arm, risk_table = TRUE,
              median_line = TRUE, logrank = TRUE,
              x_lab = "Months", title = "Overall survival")
```

theme_depict	<i>The depictr ggplot2 theme</i>
--------------	----------------------------------

Description

The minimal theme used by every plotting function in the package. It is a light modification of `ggplot2::theme_minimal()` with subtle gridlines, centred titles and comfortable margins.

Usage

```
theme_depictr(
  base_size = depictr_opt("base_size"),
  base_family = depictr_opt("base_family"),
  grid = "xy"
)
```

Arguments

<code>base_size</code>	Base font size, in points. Defaults to the <code>depictr.base_size</code> option.
<code>base_family</code>	Base font family. Defaults to the <code>depictr.base_family</code> option.
<code>grid</code>	Which major gridlines to keep: "xy", "x", "y" or "none".

Details

The default `base_size` and `base_family` come from the global options `depictr.base_size` and `depictr.base_family` (see [depictr_options\(\)](#)), so the package-wide font size can be set once; passing the arguments explicitly overrides them. The title is drawn in the brand colour, the one `options(depictr.brand = )` sets and the other functions' `depictr_brand()` defaults stand for.

Value

A ggplot2 theme object.

Examples

```
library(ggplot2)
ggplot(crop_yield, aes(fertiliser, yield)) +
  geom_point() +
  theme_depicttr()
```

threshold_plot

Classification metrics versus decision threshold

Description

Sweeps the decision threshold across the full range of scores and plots the chosen classification metrics - any of sensitivity (= recall), specificity, precision and F1 - as colour-coded curves. This is the natural companion to the ROC and precision-recall curves for *choosing* an operating point: it shows directly how each metric trades off as the cut-off moves, and (by default) marks the Youden's J and maximum-F1 optimal thresholds.

Usage

```
threshold_plot(
  x,
  score = NULL,
  metrics = c("sensitivity", "specificity", "precision", "f1"),
  mark = c("youden", "f1"),
  title = NULL
)
```

Arguments

x	A binomial glm, or the vector of observed outcomes (0/1, logical or a two-level factor with the positive class second).
score	When x is an outcome vector, the matching scores or predicted probabilities.
metrics	Which metrics to draw: any subset of "sensitivity", "specificity", "precision" and "f1". ("recall" is accepted as an alias for "sensitivity".)
mark	Which optimal operating points to mark with a dashed vertical line: any subset of "youden" (max sensitivity + specificity - 1) and "f1" (max F1). Use character(0) or NULL to mark none.
title	Plot title.

Details

At threshold t a case is predicted positive when its score is $\geq t$. The metrics are evaluated at every distinct score (the points at which the confusion matrix can change), so the curves are exact step functions rather than a coarse grid. Between two adjacent scores, each metric keeps its value at the upper one, so each step changes at a score and runs flat below it.

Value

A `ggplot2::ggplot` object. The Youden and max-F1 thresholds are stored in `attr(plot, "thresholds")`.

Examples

```
gfit <- glm(accuracy ~ word_frequency + RT + condition,
           data = lexical_decision, family = binomial)
threshold_plot(gfit)

# Sensitivity / specificity trade-off only, no markers.
threshold_plot(gfit, metrics = c("sensitivity", "specificity"),
              mark = NULL)
```

<code>tidy_estimates</code>	<i>Extract a tidy table of estimates</i>
-----------------------------	--

Description

`tidy_estimates()` turns the output of a model (or an existing data frame of results) into a single standardised table with the columns `term`, `estimate`, `std.error`, `conf.low` and `conf.high`. It is the common currency used by `coefficient_plot()` and `compare_models()`, but is useful on its own.

Usage

```
tidy_estimates(x, conf_level = 0.95, exponentiate = FALSE, ...)
```

Arguments

<code>x</code>	A fitted model or a data frame of results.
<code>conf_level</code>	Confidence (or credible) level for the interval, a single number strictly between 0 and 1.
<code>exponentiate</code>	Whether to exponentiate the estimates and interval limits, which turns log-odds into odds ratios and log rates into rate ratios. The <code>std.error</code> column stays on the link scale, as in <code>broom::tidy()</code> . A data frame is read as estimates on the link scale. Defaults to <code>FALSE</code> .
<code>...</code>	Passed to methods (and to <code>broom::tidy()</code> in the fallback). The <code>merMod</code> method additionally accepts <code>effects</code> , which currently only supports "fixed".

Details

Methods are provided for `lm`, `glm` and `merMod` (mixed models fitted with 'lme4') objects, and for data frames. For any other model class the function falls back to `broom::tidy()` when the 'broom' package is installed.

Confidence intervals are Wald intervals, estimate $\pm$ multiplier * standard error, with the multiplier the model's own tests use. An `lm` object gets the t interval of `stats::confint()`. A `glm` whose dispersion is estimated (Gaussian, Gamma, inverse Gaussian and quasi families) gets t on the residual degrees of freedom, since `stats::summary.glm()` then tests with t. One whose dispersion is fixed (binomial, Poisson and `MASS::glm.nb()`) gets z. The choice is read from the column the summary reports, "t value" or "z value". At 12 residual degrees of freedom, the t interval is about 11% wider than the z interval, and at 100 about 1% wider. A `merMod` fit gets z, since 'lme4' gives no degrees of freedom. Supply a data frame with your own intervals (for example profiled or posterior intervals) to override this.

Value

A data frame with one row per term and the columns `term`, `estimate`, `std.error`, `conf.low` and `conf.high`.

Examples

```
fit <- lm(yield ~ rainfall + fertiliser + treatment, data = crop_yield)
tidy_estimates(fit)

# Odds ratios from a logistic regression:
gfit <- glm(accuracy ~ word_frequency + condition,
           data = lexical_decision, family = binomial)
tidy_estimates(gfit, exponentiate = TRUE)

# A data frame of pre-computed estimates is standardised, not re-fitted:
df <- data.frame(
  parameter = c("a", "b"),
  Estimate = c(0.2, -0.4),
  "2.5 %" = c(0.1, -0.6),
  "97.5 %" = c(0.3, -0.2),
  check.names = FALSE
)
tidy_estimates(df)
```

timeseries_plot

Time-series plot

Description

Plots one or more series over time, with an optional moving-average overlay. The input can be a `ts` object, a numeric vector or matrix, or a data frame with a time column, a value column and an optional grouping column. A multivariate `ts`, such as `cbind(mdeaths, fdeaths)`, or a matrix with several columns gives one series per column, named after its column.

Usage

```
timeseries_plot(
  x,
  time = NULL,
  value = NULL,
  group = NULL,
  rolling = NULL,
  forecast = NULL,
  frequency = NULL,
  level = 0.95,
  lambda = NULL,
  palette = NULL,
  point = FALSE,
  title = NULL,
  x_lab = NULL,
  y_lab = NULL
)
```

Arguments

<code>x</code>	A <code>ts</code> object, a numeric vector or matrix, or a data frame.
<code>time</code>	When <code>x</code> is a data frame, the time column (string or unquoted name). When <code>x</code> is a numeric vector or matrix, an optional vector of times, one per value or per row.
<code>value</code>	When <code>x</code> is a data frame, the value column.
<code>group</code>	When <code>x</code> is a data frame, an optional grouping column mapped to colour. A factor keeps the order of its levels in the legend and the palette, and levels no row uses are left out. Other columns are ordered by first appearance.
<code>rolling</code>	Optional integer window for a centred moving-average overlay. The data are ordered by time (within each group) before the moving average is computed, so unsorted input is handled correctly. An odd window averages each observation with the same number of neighbours on either side. An even window, such as 12 for monthly data, gives the 2×12 moving average, a 12-point average followed by a 2-point one (Hyndman & Athanasopoulos, 2021, Section 3.3). The two observations at the ends of the 13 it spans count half, which centres the average on its observation, as equal weights over an even window cannot. It needs a series one observation longer than the window.
<code>forecast</code>	Optional forecast overlay for a single (non-grouped) series. Either an integer horizon (number of future steps) to forecast with the built-in STL plus random walk with drift (see <code>ts_forecast()</code>), or a pre-computed forecast supplied as a data frame with columns <code>time</code> , <code>fit</code> and (optionally) <code>lwr/upr</code> , for instance from a fitted <code>forecast::forecast()</code> object. The point forecast continues the line and a shaded prediction-interval ribbon is drawn behind it. The built-in forecast is captioned with its coverage and method, and its band widens with the horizon on the scale it is made on (see <code>lambda</code>). For a numeric <code>x</code> , it is made from the values in time order and continues the history's own time axis, one typical step at a time. Missing values at either end of the series are trimmed before forecasting,

	so the forecast starts from the last observation, and a missing value inside the series is an error.
frequency	Number of observations per period, used only to coerce a numeric <code>x</code> to a <code>ts</code> when an integer forecast horizon is requested.
level	Prediction-interval coverage for the built-in forecast (a single number strictly between 0 and 1, e.g. <code>0.95</code>).
lambda	Optional Box-Cox transformation for the built-in forecast, passed to <code>ts_forecast()</code> . <code>0</code> forecasts on the log scale, which suits a series whose seasonal swing grows with its level.
palette	Colours for the groups; defaults to <code>depictr_palette()</code> .
point	Whether to add points as well as the line.
title, x_lab, y_lab	Title and axis labels.

Value

A `ggplot2::ggplot` object.

References

Hyndman RJ, Athanasopoulos G (2021). *Forecasting: Principles and practice*, 3rd edition. OTexts, Melbourne, Australia. <https://otexts.com/fpp3/>.

See Also

`ts_forecast()`, `decompose_plot()`, `seasonal_plot()`

Examples

```
timeseries_plot(AirPassengers, rolling = 12,
                title = "Air passengers", y_lab = "Passengers")
# One line per column of a multivariate ts
timeseries_plot(cbind(mdeaths, fdeaths), y_lab = "Deaths")
# 24-month forecast with a 90% prediction interval, made on the log scale
# because the seasonal swing grows with the number of passengers
timeseries_plot(AirPassengers, forecast = 24, level = 0.9, lambda = 0)
```

ts_forecast

Forecast a seasonal series with STL and a random walk with drift

Description

A dependency-free forecaster for a single seasonal series, which forecasts by decomposition (Hyndman & Athanasopoulos, 2021, Section 5.7). `stats::stl()` with `s.window = "periodic"` estimates a seasonal component that repeats unchanged in every cycle (Cleveland et al., 1990). The seasonally adjusted series, the data minus that component, is forecast with a random walk with drift, and the last cycle of the seasonal component is added back. The result reproduces `forecast::stlf(x, method = "rwdrift", s.window = "periodic")`.

Usage

```
ts_forecast(x, h = 12, frequency = NULL, level = 0.95, lambda = NULL)
```

Arguments

<code>x</code>	A <code>ts</code> object, or a numeric vector (then <code>frequency</code> is required), holding one series. It needs more than two full cycles of observations. Missing values at either end are trimmed, so the forecast starts after the last observation, and a missing value inside the series is an error. A multivariate <code>ts</code> or a matrix with several columns is refused.
<code>h</code>	Forecast horizon: the number of future steps (a positive integer).
<code>frequency</code>	Number of observations per period; taken from <code>x</code> when it is a <code>ts</code> .
<code>level</code>	Prediction-interval coverage (a single number strictly between 0 and 1).
<code>lambda</code>	Optional Box-Cox transformation parameter. <code>NULL</code> , the default, forecasts the series as it is. <code>0</code> forecasts its logarithm, and any other number forecasts $(x^\lambda - 1) / \lambda$. Every observation must then be positive. The forecasts and limits are transformed back to the original scale, where the point forecast is the median of the forecast distribution, without the bias adjustment that would make it the mean (Hyndman & Athanasopoulos, 2021, Section 5.6). A forecast or limit beyond the range the transformation can reach comes back as <code>0</code> (for a positive <code>lambda</code>) or <code>Inf</code> (for a negative one).

Details

The prediction interval at horizon `h` is `fit +/- z * sigma_h`, with `z = qnorm((1 + level) / 2)`. The forecast standard deviation is that of a random walk with drift, `sigma_h = sigma * sqrt(h * (1 + h / (T - 1)))` (Hyndman & Athanasopoulos, 2021, Section 5.5), where `sigma` is the standard deviation of the steps of the seasonally adjusted series and `T` the length of the series. The band widens with the horizon, and the second term carries the uncertainty in the estimated drift.

The model assumes a seasonal pattern that adds to the level of the series and stays the same from cycle to cycle. When the seasonal swing grows with the level, as in `AirPassengers`, forecast on the log scale with `lambda = 0`. The band treats the seasonal component as known, so it leaves out the uncertainty in estimating it. In a rolling-origin backtest on five classic series whose seasonal swing does not grow with their level, the nominal 95% band covered 92% of one-step-ahead outcomes (`dev/ts_forecast_backtest.R` in the source repository). From three to 12 steps ahead, it covered 99%. The band therefore runs somewhat narrow at the shortest horizons and wide further out, where the uncertainty of a random walk keeps growing while these series stay close to their trend. For a different model, fit it with, for example, `forecast::forecast()` and pass its columns to `timeseries_plot()` as a data frame.

Value

A data frame with one row per forecast step and columns `time` (the future times, on the same scale as `stats::time()`), `fit`, `lwr` and `upr`.

References

Cleveland RB, Cleveland WS, McRae JE, Terpenning I (1990). “STL: A seasonal-trend decomposition procedure based on loess.” *Journal of Official Statistics*, **6**(1), 3–73.

Hyndman RJ, Athanasopoulos G (2021). *Forecasting: Principles and practice*, 3rd edition. OTexts, Melbourne, Australia. <https://otexts.com/fpp3/>.

See Also

[timeseries_plot\(\)](#), [decompose_plot\(\)](#)

Examples

```
# The seasonal swing of AirPassengers grows with the number of passengers,
# so it is forecast on the log scale.
fc <- ts_forecast(AirPassengers, h = 12, lambda = 0)
head(fc)
# On that scale, the band widens with the horizon
diff(log(fc$upr) - log(fc$lwr)) > 0
```

vif_plot

Variance inflation factor plot

Description

Computes a variance inflation factor for each *term* in a model and shows them as a bar chart, with a reference line at the usual rule of thumb (threshold). High bars flag predictors whose coefficients are unstable because they are collinear with the others. Values are computed from base R (no ‘car’ dependency).

Usage

```
vif_plot(model, threshold = 5, palette = depictr_palette(2), title = NULL)
```

Arguments

model	A fitted lm or glm model with at least two predictors.
threshold	Reference value for the ordinary VIF, drawn as a line. For models with multi-column terms it is shown on the $\text{GVIF}^{1/(2\text{df})}$ scale as $\sqrt{\text{threshold}}$.
palette	Length-2 colours encoding terms below and above the threshold. Defaults to the colourblind-safe depictr_palette() pair (blue / orange).
title	Plot title.

Details

For single-degree-of-freedom terms this is the ordinary VIF, $1/(1 - R^2)$. For terms that span several design-matrix columns (multi-level factors, or spline/polynomial bases) the function reports the generalised VIF of Fox and Monette (1992): with R the correlation matrix of the coefficient estimates other than the intercept, R_{11} the block for the term and R_{22} the block for the remaining coefficients,

$$\text{GVIF} = \frac{\det(R_{11}) \det(R_{22})}{\det(R)}.$$

When every term has a single degree of freedom the bars are the ordinary VIFs, on a plain "Variance inflation factor" axis with the reference line at threshold. If any term spans several columns the bars switch to the comparable $\text{GVIF}^{1/(2 \text{ df})}$ (which for a single-df term equals $\sqrt{\text{VIF}}$) and the reference line moves to $\sqrt{\text{threshold}}$ accordingly. The x-axis is kept tight to the data: when every bar is comfortably below the threshold the line is reported in the caption rather than drawn into a wide empty band.

The correlations come from the fit's unscaled coefficient covariance, so a `glm` or a weighted `lm` is measured with the weights of its fit, and the values agree with those of `car::vif()`. For an unweighted `lm` with an intercept, they are the values the formula gives with R the correlation matrix of the predictor columns, as Fox and Monette state it. A model with aliased coefficients has no VIF and is refused, with the terms they belong to named. A model without an intercept warns, as `car::vif()` does, because its VIFs then come from columns that are not centred.

Value

A `ggplot2::ggplot` object.

References

Fox J, Monette G (1992). "Generalized collinearity diagnostics." *Journal of the American Statistical Association*, **87**(417), 178–183. doi:10.1080/01621459.1992.10475190.

Examples

```
# Two deliberately collinear predictors: soil moisture is largely driven by
# rainfall, so both carry an inflated VIF (around 6, above the line at 5)
# while fertiliser stays near 1.
set.seed(1)
d <- crop_yield
d$soil_moisture <- 0.05 * d$rainfall + rnorm(nrow(d), sd = 2)
vif_plot(lm(yield ~ rainfall + soil_moisture + fertiliser, data = d))

# Multi-level factors get a single generalised VIF per term
fit2 <- lm(yield ~ rainfall + fertiliser + treatment, data = crop_yield)
vif_plot(fit2)
```

wellbeing_survey	<i>Simulated wellbeing survey</i>
------------------	-----------------------------------

Description

A reproducibly simulated cross-sectional wellbeing survey for descriptive, correlation, regression and missing-data examples. `life_satisfaction` responds meaningfully to its predictors: higher stress lowers it, more sleep and exercise raise it, log-income has a positive effect of about +0.4 per natural-log unit, and age has a mild inverted-U (peaking in mid-life). `education` is an ordered factor. Income is missing more often at higher stress (missing at random), and because income predicts the outcome, this missingness is informative for the missing-data examples.

Usage

```
wellbeing_survey
```

Format

A data frame with 300 rows and 9 variables:

id Respondent identifier.

region Region of residence: North, South, East or West.

age Age in years.

education Highest education level (ordered factor: secondary < undergraduate < postgraduate).

income Annual income; contains missing values (more at high stress).

stress Self-reported stress (1-7).

sleep_hours Typical hours of sleep per night; contains missing values.

exercise_days Days of exercise per week (0-7); contains missing values.

life_satisfaction Life satisfaction (1-7); contains missing values.

Details

The data are synthetic, generated by `data-raw/generate_datasets.R` with a fixed seed; they do not describe any real individuals.

Source

Simulated; see `data-raw/generate_datasets.R`.

Index

* datasets

- clinical_trial, 13
- crop_yield, 23
- lexical_decision, 51
- monthly_sales, 57
- wellbeing_survey, 103

acf_plot, 4

acf_plot(), 57

arrange_plots, 5

arrange_plots(), 12, 56

base::strwrap(), 42

binned_residual_plot, 6

binned_residual_plot(), 56, 75

broom::tidy(), 96, 97

calibration_plot, 7

calibration_plot(), 14

check_figure, 8

check_figure(), 29, 62, 79, 82, 88, 89

clinical_trial, 13

cluster::silhouette(), 87

cluster_plot, 14

coefficient_plot, 16

coefficient_plot(), 20, 51, 56, 96

compare_models, 19

compare_models(), 43, 44, 96

confusion_matrix_plot, 21

correlation_heatmap, 22

correlation_heatmap(), 15, 41, 50

crop_yield, 23

decompose_plot, 24

decompose_plot(), 57, 86, 99, 101

dendrogram_plot, 26

dendrogram_plot(), 15

depictr_options, 27

depictr_options(), 29, 30, 81, 94

depictr_palette, 28

depictr_palette(), 11, 15, 20, 23, 26–28, 31, 32, 35, 37, 38, 40, 41, 43, 47, 49, 54, 59–61, 64, 72, 76, 81–83, 85, 87, 92, 99, 101

dumbbell_plot, 30

ecdf_plot, 32

effects_plot, 33

effects_plot(), 49, 56

estimation_plot, 34

estimation_plot(), 47

explore_bivariate, 37

explore_categorical, 38

explore_distribution, 39

explore_distribution(), 39

explore_pairs, 40

format_terms, 41

format_terms(), 17, 20, 59, 66, 68

frequentist_bayesian_plot, 43

frequentist_bayesian_plot(), 19

gain_plot, 45

gain_plot(), 14

ggplot2::annotation_raster(), 89

ggplot2::discrete_scale(), 82

ggplot2::geom_label(), 11

ggplot2::geom_raster(), 89

ggplot2::geom_smooth(), 37, 83

ggplot2::geom_text(), 11

ggplot2::ggplot, 4, 7, 8, 16, 18, 21–23, 25, 27, 31, 32, 34, 36, 37, 39–41, 44, 46–49, 51, 53, 54, 57, 59–61, 64, 67, 69–71, 73–75, 77, 78, 83, 84, 86, 87, 93, 96, 99, 102

ggplot2::ggplotGrob(), 89

ggplot2::ggsave(), 79

ggplot2::scale_colour_identity(), 10

ggplot2::scale_colour_manual(), 10

ggplot2::theme_minimal(), 94

ggplot2::theme_void(), 80
 grDevices::cairo_pdf(), 80
 grDevices::cairo_ps(), 80
 grDevices::col2rgb(), 88
 grDevices::embedFonts(), 80
 grDevices::pdf(), 80
 grDevices::postscript(), 80
 grDevices::tiff(), 80
 grid::grid.draw(), 89
 grid::linearGradient(), 89
 group_comparison_plot, 46

 influence_plot, 48
 interaction_plot, 48

 k_diagnostic, 50
 k_diagnostic(), 15, 16
 knitr::kable(), 90

 lexical_decision, 51
 lift_plot, 52
 lift_plot(), 14
 lme4::allFit(), 58, 59
 lme4::glmer(), 34
 lme4::lmer(), 34, 55
 lme4::refitML(), 55

 missingness_map, 53
 missingness_map(), 30
 model_fit_table, 54
 model_fit_table(), 56
 model_report, 55
 model_report(), 12, 71
 monthly_sales, 57

 optimizer_fixef_plot, 58
 optimizer_fixef_plot(), 51
 outlier_plot, 60
 outlier_plot(), 30

 palette_preview, 61
 palette_preview(), 13, 88
 palette_safety, 62
 palette_safety(), 8, 9, 13, 88
 patchwork::patchworkGrob(), 89
 patchwork::plot_annotation(), 12
 patchwork::plot_layout(), 35
 patchwork::wrap_plots(), 5
 pca_plot, 63
 posterior_plot, 65
 posterior_plot(), 44
 power_curve_plot, 67
 pr_curve_plot, 69
 pr_curve_plot(), 14

 qq_plot, 70
 qq_plot(), 56

 raincloud_plot, 72
 random_effects_plot, 73
 residual_diagnostics_plot, 74
 residual_diagnostics_plot(), 12, 56, 71
 ridgeline_plot, 76
 roc_curve_plot, 77
 roc_curve_plot(), 21

 save_plot, 78
 save_plot(), 9, 12, 89
 scale_color_depict
 (scale_colour_depict), 81
 scale_colour_depict, 81
 scale_colour_depict(), 11, 27, 28, 30, 78
 scale_fill_depict
 (scale_colour_depict), 81
 scale_fill_depict(), 28, 30
 scatter_trend, 82
 scatter_trend(), 37
 scree_plot, 84
 seasonal_plot, 84
 seasonal_plot(), 25, 99
 silhouette_plot, 86
 silhouette_plot(), 16
 simulate_cvd, 87
 simulate_cvd(), 10, 88, 89
 simulate_figure, 88
 simulate_figure(), 88
 stats::acf(), 4
 stats::biplot(), 64
 stats::confint(), 97
 stats::cor(), 22, 23, 41
 stats::cutree(), 15
 stats::decompose(), 25
 stats::density(), 77
 stats::dist, 26, 86
 stats::dist(), 26, 87
 stats::glm(), 55
 stats::hclust, 26
 stats::hclust(), 26
 stats::kmeans(), 15, 16, 50

stats::pacf(), [4](#)
stats::prcomp(), [63](#), [64](#), [84](#)
stats::predict(), [34](#)
stats::stl(), [25](#), [85](#), [99](#)
stats::summary.glm(), [97](#)
stats::time(), [100](#)
summary_table, [90](#)
survival::quantile.survfit(), [93](#)
survival::Surv(), [92](#)
survival::survdiff(), [93](#)
survival::survfit(), [93](#)
survival_plot, [91](#)
survival_plot(), [13](#)

theme_depict, [94](#)
theme_depict(), [27](#)
threshold_plot, [95](#)
tidy_estimates, [96](#)
tidy_estimates(), [16](#), [17](#), [20](#), [34](#)
timeseries_plot, [97](#)
timeseries_plot(), [57](#), [86](#), [100](#), [101](#)
ts_forecast, [99](#)
ts_forecast(), [25](#), [98](#), [99](#)

vif_plot, [101](#)

wellbeing_survey, [103](#)