

Package ‘priorsense’

September 30, 2026

Title Prior Diagnostics and Sensitivity Analysis

Version 1.4.0

Description Provides functions for prior and likelihood sensitivity analysis in Bayesian models. Currently it implements methods to determine the sensitivity of the posterior to power-scaling perturbations of the prior and likelihood.

License GPL (>= 3)

Encoding UTF-8

Imports checkmate (>= 2.3.4), cli (>= 3.6.6), ggdist (>= 3.3.3), ggh4x (>= 0.3.1), ggplot2 (>= 4.0.3), grDevices, grid, matrixStats (>= 1.5.0), posterior (>= 1.7.0), rlang (>= 1.3.0), stats, tibble (>= 3.3.1), utils

Suggests bayesplot (>= 1.16.0), brms (>= 2.23.0), cmdstanr (>= 0.9.0), iwmms (>= 0.0.1), nimble (>= 1.4.3), philentropy (>= 0.10.0), quarto (>= 1.5.1), R2jags (>= 0.8), rstan (>= 2.32.7), testthat (>= 3.3.2), transport (>= 0.15), vdiff (>= 1.0.9)

Config/testthat/edition 3

Depends R (>= 4.1.0)

VignetteBuilder quarto

Additional_repositories <https://topipa.r-universe.dev>,
<https://stan-dev.r-universe.dev>

URL <https://mc-stan.org/priorsense/>, <https://discourse.mc-stan.org/>

BugReports <https://github.com/stan-dev/priorsense/issues>

Config/Needs/website quarto

Config/roxygen2/version 8.1.0

NeedsCompilation no

Author Noa Kallioinen [aut, cre] (ORCID:

<<https://orcid.org/0000-0003-1586-8382>>),

Topi Paananen [aut] (ORCID: <<https://orcid.org/0000-0002-6542-407X>>),

Paul-Christian Bürkner [aut] (ORCID:

<<https://orcid.org/0000-0001-5765-8995>>),

Aki Vehtari [aut] (ORCID: <<https://orcid.org/0000-0003-2164-9469>>),
 Frank Weber [ctb] (ORCID: <<https://orcid.org/0000-0002-4842-7922>>),
 Simon Taylor [rev],
 Dylan Dijk [rev]

Maintainer Noa Kallioinen <noa.kallioinen@helsinki.fi>

Repository CRAN

Date/Publication 2026-09-30 18:50:02 UTC

Contents

priorsense-package	2
cjs_dist	4
create-priorsense-data	5
example_powerscale_model	7
log_lik_draws	7
log_prior_draws	8
powerscale-gradients	9
powerscale-overview	11
powerscale-plots	15
powerscale-sensitivity	18
powerscale_derivative	21
predictions_as_draws	22
Index	24

priorsense-package	<i>priorsense: Prior (and likelihood) diagnostics and sensitivity analysis</i>
--------------------	--

Description

The **priorsense** package provides functions for prior and likelihood sensitivity analysis of Bayesian models. Currently it implements methods to determine the sensitivity of the posterior to power-scaling perturbations of the prior and likelihood and is the first implementation of the method described in Kallioinen et al. (2023).

Details

The main diagnostic function provided by **priorsense** is `powerscale_sensitivity`. Given a fitted model or draws object, it computes the powerscaling sensitivity diagnostic described in Kallioinen et al. (2023). It does so by perturbing the prior and likelihood and computing the effect on the posterior, without needing to refit the model (using Pareto smoothed importance sampling and importance weighted moment matching; Vehtari et al. 2022, Paananen et al. 2021).

In addition, visual diagnostics are available by first using `powerscale_sequence` to create a sequence of perturbed posteriors, and then a plot function such as `powerscale_plot_ecdf` to visualise the change.

For an introduction, see `vignette("priorsense")`.

The following global options are available:

- `priorsense.plot_help_text`: If TRUE (the default), priorsense plots will include a title and explanatory text. If FALSE they will not.
- `priorsense.plot_variables_per_page`: Number specifying the maximum number of variables to be plotted on one page of a plot.
- `priorsense.plot_ask`: If TRUE (the default), when multiple pages are plotted input is required before each subsequent page is rendered. If FALSE no input is required.

Author(s)

Maintainer: Noa Kallioinen <noa.kallioinen@helsinki.fi> ([ORCID](#))

Authors:

- Noa Kallioinen <noa.kallioinen@helsinki.fi> ([ORCID](#))
- Topi Paananen ([ORCID](#))
- Paul-Christian Bürkner ([ORCID](#))
- Aki Vehtari ([ORCID](#))

Other contributors:

- Frank Weber ([ORCID](#)) [contributor]
- Simon Taylor [reviewer]
- Dylan Dijk [reviewer]

References

- Kallioinen, N., Paananen, T., Bürkner, P-C., Vehtari, A. (2023). Detecting and diagnosing prior and likelihood sensitivity with power-scaling perturbations. *Statistics and Computing*. 34(57). doi:10.1007/s11222-023-10366-5
- Vehtari, A., Simpson, D., Gelman, A., Yao, Y., and Gabry, J. (2024). Pareto smoothed importance sampling. *Journal of Machine Learning Research*. 25(72). <https://jmlr.org/papers/v25/19-556.html>
- Paananen, T., Piironen, J., Bürkner, P-C., Vehtari, A. (2021). Implicitly adaptive importance sampling. *Statistics and Computing*. 31(16). doi:10.1007/s11222-020-09982-2

See Also

[powerscale-overview](#), [powerscale-plots](#)

cjs_dist	<i>Cumulative Jensen-Shannon divergence</i>
----------	---

Description

Computes the cumulative Jensen-Shannon distance between two samples.

Usage

```
cjs_dist(
  x,
  y,
  x_weights = NULL,
  y_weights = NULL,
  metric = TRUE,
  unsigned = TRUE,
  ...
)
```

Arguments

x	numeric vector of draws from first distribution
y	numeric vector of draws from second distribution
x_weights	numeric vector (same length as x) of weights for the draws of the first distribution
y_weights	numeric vector (same length as y) of weights for the draws of the second distribution
metric	Logical; if TRUE, return square-root of CJS. Default is TRUE
unsigned	Logical; if TRUE then return max of CJS(P(x) Q(x)) and CJS(P(-x) Q(-x)). This ensures invariance to transformations such as PCA. Default is TRUE
...	unused

Details

The Cumulative Jensen-Shannon distance is a symmetric metric based on the cumulative Jensen-Shannon divergence. The divergence $CJS(P \parallel Q)$ between two cumulative distribution functions P and Q is defined as:

$$CJS(P \parallel Q) = \sum P(x) \log_2 \frac{P(x)}{0.5(P(x) + Q(x))} + \frac{1}{2 \ln 2} \sum (Q(x) - P(x))$$

The symmetric metric is defined as:

$$CJS_{dist}(P \parallel Q) = \sqrt{CJS(P \parallel Q) + CJS(Q \parallel P)}$$

This has an upper bound of $\sqrt{\sum (P(x) + Q(x))}$

Value

distance value based on CJS computation.

References

Nguyen H-V., Vreeken J. (2015). Non-parametric Jensen-Shannon Divergence. In: Appice A., Rodrigues P., Santos Costa V., Gama J., Jorge A., Soares C. (eds) Machine Learning and Knowledge Discovery in Databases. ECML PKDD 2015. Lecture Notes in Computer Science, vol 9285. Springer, Cham. doi:10.1007/978-3-319-23525-7_11

Examples

```
x <- rnorm(100)
y <- rnorm(100, 2, 2)
cjs_dist(x, y, x_weights = NULL, y_weights = NULL)
```

create-priorsense-data

Create data structure for priorsense

Description

Create a data structure that contains all required data and functions for priorsense

Usage

```
create_priorsense_data(x, ...)

## Default S3 method:
create_priorsense_data(
  x,
  fit = NULL,
  log_prior_fn = log_prior_draws,
  log_lik_fn = log_lik_draws,
  log_prior = NULL,
  log_lik = NULL,
  log_ratio_fn = NULL,
  log_prior_name = "lprior",
  log_lik_name = "log_lik",
  ...
)

## S3 method for class 'stanfit'
create_priorsense_data(x, ...)

## S3 method for class 'CmdStanFit'
create_priorsense_data(x, ...)
```

```
## S3 method for class 'draws'
create_priorsense_data(x, ...)

## S3 method for class 'rjags'
create_priorsense_data(x, ...)

## S3 method for class 'jagsUI'
create_priorsense_data(x, ...)

## S3 method for class 'mcmc.list'
create_priorsense_data(x, ...)
```

Arguments

<code>x</code>	an object for which the method is defined or an object coercible to a posterior: : draws object
<code>...</code>	arguments passed to methods
<code>fit</code>	a model fit object (only used if <code>x</code> is not a fit object)
<code>log_prior_fn</code>	function to derive log prior from <code>x</code> or <code>fit</code> (if not NULL)
<code>log_lik_fn</code>	function to derive log likelihood from <code>x</code> or <code>fit</code> (if not NULL)
<code>log_prior</code>	draws object from log prior, must be numeric and not include NA, NaN, Inf, -Inf or be constant
<code>log_lik</code>	draws from log likelihood, must be numeric and not include NA, NaN, Inf, -Inf or be constant
<code>log_ratio_fn</code>	function for moment matching
<code>log_prior_name</code>	Character (case sensitive) specifying name of the variable storing the log prior evaluations
<code>log_lik_name</code>	Character (case sensitive) specifying name of the variable storing the log likelihood evaluations

Value

A `priorsense_data` object, which contains the data and functions to run sensitivity analyses.

Examples

```
x <- example_powerscale_model()
drw <- x$draws

psd <- create_priorsense_data(drw)
```

example_powerscale_model

Example models for power-scaling sensitivity

Description

Provides example models (with data) that are ready for use with power-scaling.

Usage

```
example_powerscale_model(model = "univariate_normal", language = "stan")
```

Arguments

model	Character specifying which model code to return. Currently "univariate_normal" and "eight_schools" are implemented.
language	Character specifying which modelling language to return. One of "stan", "jags", "nimble". Default is "stan".

Value

List containing model code and corresponding data.

Examples

```
ex_normal <- example_powerscale_model(model = "univariate_normal", language = "stan")
ex_eightschools <- example_powerscale_model(model = "eight_schools", language = "jags")
```

log_lik_draws

Extract log likelihood draws

Description

Extract log likelihood from fitted model and return as a draws object.

Usage

```
log_lik_draws(x, ...)

## S3 method for class 'stanfit'
log_lik_draws(x, joint = FALSE, log_lik_name = "log_lik", ...)

## S3 method for class 'CmdStanFit'
log_lik_draws(x, joint = FALSE, log_lik_name = "log_lik", ...)

## S3 method for class 'draws'
log_lik_draws(x, joint = FALSE, log_lik_name = "log_lik", ...)
```

Arguments

x	Model fit or draws object.
...	Arguments passed to individual methods.
joint	Logical indicating whether to return the joint log likelihood or array. Default is FALSE.
log_lik_name	Name of parameter in Stan model corresponding to log likelihood, default is "log_lik".

Value

A draws_array object containing log_lik values.

Examples

```
ex <- example_powerscale_model()
drw <- ex$draws

log_lik_draws(drw)
```

log_prior_draws	<i>Extract log prior draws</i>
-----------------	--------------------------------

Description

Extract log prior from fitted model and return as a draws object.

Usage

```
log_prior_draws(x, ...)

## S3 method for class 'stanfit'
log_prior_draws(x, joint = FALSE, log_prior_name = "lprior", ...)

## S3 method for class 'CmdStanFit'
log_prior_draws(x, joint = FALSE, log_prior_name = "lprior", ...)

## S3 method for class 'draws'
log_prior_draws(x, joint = FALSE, log_prior_name = "lprior", ...)
```

Arguments

x	Model fit or draws object.
...	Arguments passed to individual methods.
joint	Logical indicating whether to return the joint log prior or array. Default is FALSE.
log_prior_name	Name of parameter in Stan model corresponding to log prior, default is "lprior".

Value

A `draws_array` object containing `log_prior` values.

Examples

```
ex <- example_powerscale_model()
drw <- ex$draws

log_prior_draws(drw)
```

powerscale-gradients *Power-scale gradients*

Description

Calculate the numerical derivative of posterior quantities/divergence with respect to power-scaling the specified component (prior or likelihood). This is done using importance sampling (and optionally moment matching).

Usage

```
powerscale_gradients(x, ...)
```

Default S3 method:

```
powerscale_gradients(
  x,
  log_prior_name = "lprior",
  log_lik_name = "log_lik",
  ...
)
```

S3 method for class 'priorsense_data'

```
powerscale_gradients(
  x,
  variable = NULL,
  component = c("prior", "likelihood"),
  type = c("quantities", "divergence"),
  lower_alpha = 0.99,
  upper_alpha = 1.01,
  div_measure = "cjs_dist",
  measure_args = list(),
  moment_match = FALSE,
  k_threshold = NULL,
  resample = FALSE,
  transform = NULL,
  prediction = NULL,
```

```

    scale = FALSE,
    prior_selection = NULL,
    likelihood_selection = NULL,
    ...
)

```

Arguments

<code>x</code>	Model fit or draws object.
<code>...</code>	Further arguments passed to functions.
<code>log_prior_name</code>	Character (case sensitive) specifying name of the variable storing the log prior evaluations
<code>log_lik_name</code>	Character (case sensitive) specifying name of the variable storing the log likelihood evaluations
<code>variable</code>	Variables to compute sensitivity of. If NULL (default) sensitivity is computed for all variables.
<code>component</code>	Component to power-scale (prior or likelihood).
<code>type</code>	type of sensitivity to measure ("distance", "quantity"). Multiple options can be specified at the same time.
<code>lower_alpha</code>	lower power to scale component by, should be < 1 (default is 0.9).
<code>upper_alpha</code>	upper power to scale component by, should be > 1 (default is 1.1).
<code>div_measure</code>	Character (case sensitive) specifying the divergence measure to use. The following methods are implemented: • <code>"cjs_dist"</code>: Cumulative Jensen-Shannon distance. Default method. See function <code>cjs_dist()</code> for more details. • <code>"js_dist"</code>: Jensen-Shannon distance. First estimates density using <code>stats::density()</code>. Requires the <code>philentropy</code> package. See <code>philentropy::jensen_shannon()</code>. • <code>"js_div"</code>: Jensen-Shannon divergence. First estimates density using <code>stats::density()</code>. Requires the <code>philentropy</code> package. See <code>philentropy::jensen_shannon()</code>. • <code>"hellinger_dist"</code>: Hellinger distance. First estimates density using <code>stats::density()</code>. Requires the <code>philentropy</code> package. See <code>philentropy::hellinger()</code>. • <code>"kl_dist"</code>: Kullback-Leibler distance. First estimates density using <code>stats::density()</code>. Requires the <code>philentropy</code> package. See <code>philentropy::kullback_leibler_distance()</code>. • <code>"kl_div"</code>: Kullback-Leibler divergence. First estimates density using <code>stats::density()</code>. Requires the <code>philentropy</code> package. See <code>philentropy::kullback_leibler_distance()</code>. • <code>"ks_dist"</code>: Kolmogorov-Smirnov distance. See <code>stats::ks.test()</code>. • <code>"ws_dist"</code>: Wasserstein distance (pass <code>measure_args = list(p = N)</code>) for a different order, where N is the order. Requires the <code>transport</code> package. See <code>transport::wasserstein1d()</code>.
<code>measure_args</code>	Named list of further arguments passed to divergence measure functions.
<code>moment_match</code>	Logical; Indicate whether or not moment matching should be performed. Can only be TRUE if <code>is_method</code> is "psis".
<code>k_threshold</code>	Threshold value for Pareto k values above which the moment matching algorithm is used. Default is 0.5.

resample	Logical; Indicate whether or not draws should be resampled based on calculated importance weights.
transform	Indicate a transformation of posterior draws to perform before sensitivity analysis. Either "scale" or "whiten".
prediction	Function taking the model fit and returning a draws_df of predictions to be appended to the posterior draws
scale	logical scale quantity gradients by base posterior standard deviation.
prior_selection	Numeric vector specifying which priors to consider.
likelihood_selection	Numeric vector specifying which likelihoods to consider.

Value

Maximum of the absolute derivatives above and below $\alpha = 1$.

Examples

```
ex <- example_powerscale_model()
drw <- ex$draws

powerscale_gradients(drw)
```

powerscale-overview *Prior/likelihood power-scaling perturbation*

Description

Estimate posterior draws based on power-scaling perturbations of prior or likelihood using importance sampling (and optionally moment matching).

Usage

```
powerscale(x, ...)

## Default S3 method:
powerscale(
  x,
  component,
  alpha,
  moment_match = FALSE,
  k_threshold = NULL,
  resample = FALSE,
  transform = NULL,
  prediction = NULL,
  variable = NULL,
```

```

        variables = NULL,
        selection = NULL,
        log_prior_name = "lprior",
        log_lik_name = "log_lik",
        separator = "_",
        ...
    )

## S3 method for class 'priorsense_data'
powerscale(
    x,
    component,
    alpha,
    moment_match = FALSE,
    k_threshold = NULL,
    resample = FALSE,
    transform = NULL,
    prediction = NULL,
    variable = NULL,
    variables = NULL,
    selection = NULL,
    separator = "_",
    ...
)

powerscale_sequence(x, ...)

## Default S3 method:
powerscale_sequence(
    x,
    lower_alpha = 0.8,
    upper_alpha = 1/lower_alpha,
    length = 3,
    variable = NULL,
    variables = NULL,
    component = c("prior", "likelihood"),
    moment_match = FALSE,
    k_threshold = 0.5,
    resample = FALSE,
    transform = NULL,
    prediction = NULL,
    auto_alpha_range = FALSE,
    symmetric = TRUE,
    prior_selection = NULL,
    likelihood_selection = NULL,
    ...
)

```

```
## S3 method for class 'priorsense_data'
powerscale_sequence(
  x,
  lower_alpha = 0.8,
  upper_alpha = 1/lower_alpha,
  length = 3,
  variable = NULL,
  variables = NULL,
  component = c("prior", "likelihood"),
  moment_match = FALSE,
  k_threshold = NULL,
  resample = FALSE,
  transform = NULL,
  prediction = NULL,
  auto_alpha_range = FALSE,
  symmetric = TRUE,
  prior_selection = NULL,
  likelihood_selection = NULL,
  ...
)
```

Arguments

x	A fitted model object.
...	Further arguments passed to internal functions.
component	Component to be power-scaled (either "prior" or "likelihood"). For power-scale_sequence, this can be both "prior" and "likelihood".
alpha	Value by which to power-scale specified component. (likelihood/prior).
moment_match	Logical; Indicate whether or not moment matching should be performed. Can only be TRUE if is_method is "psis".
k_threshold	Threshold value for Pareto k values above which the moment matching algorithm is used. Default is 0.5.
resample	Logical; Indicate whether or not draws should be resampled based on calculated importance weights.
transform	Indicate a transformation of posterior draws to perform before sensitivity analysis. Either "scale" or "whiten".
prediction	Function taking the model fit and returning a draws_df of predictions to be appended to the posterior draws
variable	Vector of variable names to return estimated posterior draws for. If NULL all variables will be included.
variables	Alias of variable.
selection	Vector specifying partitions of component to be included in power-scaling. Default is NULL, which takes all partitions. If this is a character, then it is appended to the variable name (log_prior_name or log_lik_name) with an _ between them.

<code>log_prior_name</code>	Character (case sensitive) specifying name of the variable storing the log prior evaluations
<code>log_lik_name</code>	Character (case sensitive) specifying name of the variable storing the log likelihood evaluations
<code>separator</code>	Character specifying the separator between the component name and the partition. The default is "_".
<code>lower_alpha</code>	Lower power-scaling alpha value in sequence.
<code>upper_alpha</code>	Upper power-scaling alpha value in sequence.
<code>length</code>	Total length of alpha sequence, including 1.0.
<code>auto_alpha_range</code>	Boolean. Restrict range to ensure Pareto-k values below threshold?
<code>symmetric</code>	Boolean. Should the alpha range be symmetrical around alpha = 1 on log-scale? If TRUE, the length of the range will be always be odd, and will be equal to length - 1 if specified length is even.
<code>prior_selection</code>	Vector specifying partitions of component to be included in power-scaling. Default is NULL, which takes all partitions. If this is a character, then it is appended to the variable name (specified by <code>log_prior_name</code>) with a separator between them. If numeric, then it is appended inside [].
<code>likelihood_selection</code>	Vector specifying partitions of component to be included in power-scaling. Default is NULL, which takes all partitions. If this is a character, then it is appended to the variable name (specified by <code>log_lik_name</code>) with a separator between them. If numeric, then it is appended inside [].

Details

Power-scaling of the prior or likelihood is defined as exponentiating the prior or likelihood to some value alpha. The effect on the posterior is then estimated through Pareto-smoothed importance sampling and optionally importance weighted moment matching. In order to perform the calculations, the log-prior and log-likelihood evaluations must be present in the object. For a general introduction to the package see `vignette("priorsense")`, and for details on the method see Kallioinen et al. (2023). For details on Pareto-smoothed importance sampling and importance weighted moment matching, see Vehtari et al. (2024) and Paananen et al. (2021), respectively.

Value

A `powerscaled_draws` or `powerscaled_sequence` object, which contains the estimated posterior draws resulting from the power-scaling perturbations and details of the perturbation and estimation methods.

References

Kallioinen, N., Paananen, T., Bürkner, P-C., Vehtari, A. (2023). Detecting and diagnosing prior and likelihood sensitivity with power-scaling perturbations. *Statistics and Computing*. 34(57). doi:10.1007/s11222-023-10366-5

Vehtari, A., Simpson, D., Gelman, A., Yao, Y., and Gabry, J. (2024). Pareto smoothed importance sampling. *Journal of Machine Learning Research*. 25(72). <https://jmlr.org/papers/v25/19-556.html>

Paananen, T., Piironen, J., Bürkner, P-C., Vehtari, A. (2021). Implicitly adaptive importance sampling. *Statistics and Computing*. 31(16). doi:10.1007/s11222-020-09982-2

See Also

[powerscale-plots](#)

Examples

```
ex <- example_powerscale_model()

powerscale(ex$draws, component = "prior", alpha = 0.5)

powerscale_sequence(ex$draws)
```

powerscale-plots	<i>Diagnostic plots for power-scaling sensitivity</i>
------------------	---

Description

Various diagnostic plots for power-scaling sensitivity. See **Plot Descriptions** below for details.

Usage

```
powerscale_plot_dens(x, ...)

powerscale_plot_ecdf(x, ...)

## S3 method for class 'powerscaled_sequence'
powerscale_plot_ecdf(
  x,
  variable = NULL,
  variables = NULL,
  resample = FALSE,
  length = 3,
  facet_rows = "component",
  help_text = getOption("priorsense.plot_help_text", TRUE),
  colors = NULL,
  colours = NULL,
  variables_per_page = getOption("priorsense.plot_variables_per_page", 6),
  ...
)

powerscale_plot_quantities(x, ...)
```

```
## S3 method for class 'powerscaled_sequence'
powerscale_plot_quantities(
  x,
  variable = NULL,
  variables = NULL,
  quantity = c("mean", "sd"),
  div_measure = "cjs_dist",
  resample = FALSE,
  measure_args = NULL,
  mcse = TRUE,
  quantity_args = NULL,
  help_text = getOption("priorsense.plot_help_text", TRUE),
  colors = NULL,
  colours = NULL,
  variables_per_page = getOption("priorsense.plot_variables_per_page", 6),
  ...
)

## S3 method for class 'powerscaled_sequence'
plot(x, type = c("dens", "ecdf", "quantities"), ...)
```

Arguments

<code>x</code>	An object of class <code>powerscaled_sequence</code> or an object for which <code>powerscale_sequence</code> will first be run on.
<code>...</code>	Arguments passed to <code>powerscale_sequence</code> if <code>x</code> is not of class <code>powerscaled_sequence</code> .
<code>variable</code>	A character vector of variable names. If <code>NULL</code> (the default) all variables will be plotted.
<code>variables</code>	Alias of <code>variable</code> .
<code>resample</code>	Logical; Indicate whether or not draws should be resampled based on calculated importance weights.
<code>length</code>	Numeric specifying how many alpha values should be used. Ignored if the object is of class <code>powerscaled_sequence</code> .
<code>facet_rows</code>	Character defining the rows of the plot facets, either "variable" or "component". Default is "variable".
<code>help_text</code>	Logical indicating whether title and subtitle with explanatory description should be included in the plot. Default is <code>TRUE</code> . Can be set via option "priorsense.show_help_text".
<code>colors</code>	Character vector of colors to be used for plots. Either length 3 for <code>powerscale_plot_ecdf</code> and <code>powerscale_plot_dens</code> with order lowest, base, highest; or length 2 for <code>powerscale_plot_quantities</code> with order low Pareto k, high Pareto k. If <code>NULL</code> the defaults will be used.
<code>colours</code>	Alias of <code>colors</code> .
<code>variables_per_page</code>	Number specifying the maximum number of variables to show on each page of the plot. Default is 6. If <code>NULL</code> or <code>Inf</code> , all variables will be plotted on the same page.

quantity	A character vector specifying one or several quantities to plot. Options are "mean", "median", "sd", "mad", "quantile".
div_measure	Character (case sensitive) specifying the divergence measure to use. The following methods are implemented: • "cjs_dist": Cumulative Jensen-Shannon distance. Default method. See function <code>cjs_dist()</code> for more details. • "js_dist": Jensen-Shannon distance. First estimates density using <code>stats::density()</code>. Requires the <code>philentropy</code> package. See <code>philentropy::jensen_shannon()</code>. • "js_div": Jensen-Shannon divergence. First estimates density using <code>stats::density()</code>. Requires the <code>philentropy</code> package. See <code>philentropy::jensen_shannon()</code>. • "hellinger_dist": Hellinger distance. First estimates density using <code>stats::density()</code>. Requires the <code>philentropy</code> package. See <code>philentropy::hellinger()</code>. • "kl_dist": Kullback-Leibler distance. First estimates density using <code>stats::density()</code>. Requires the <code>philentropy</code> package. See <code>philentropy::kullback_leibler_distance()</code>. • "kl_div": Kullback-Leibler divergence. First estimates density using <code>stats::density()</code>. Requires the <code>philentropy</code> package. See <code>philentropy::kullback_leibler_distance()</code>. • "ks_dist": Kolmogorov-Smirnov distance. See <code>stats::ks.test()</code>. • "ws_dist": Wasserstein distance (pass <code>measure_args = list(p = N)</code>) for a different order, where N is the order. Requires the <code>transport</code> package. See <code>transport::wasserstein1d()</code>.
measure_args	Named list of further arguments passed to divergence measure functions.
mcse	Boolean; If TRUE will plot $\pm 2 \times$ Monte Carlo standard error of the base quantity on the quantities plot.
quantity_args	Named list of further arguments passed to quantity functions. Passed as <code>.args</code> to <code>[posterior::summarise_draws]</code> .
type	Character specifying type of plot, one of "dens", "ecdf", or "quantities".

Value

A ggplot object (or a `priorsense_plot` object which is a list of ggplot objects if there is more than one page) that can be further customized using the **ggplot2** package.

Plot Descriptions

`powerscale_plot_dens()` Kernel density plot of power-scaled posterior draws with respect to power-scaling.

`powerscale_plot_ecdf()` Empirical cumulative distribution function plot of power-scaled posterior draws with respect to power-scaling.

`powerscale_plot_quantities()` Plot of posterior quantities with respect to power-scaling.

Examples

```
ex <- example_powerscale_model()

powerscale_plot_dens(ex$draws)
```

powerscale-sensitivity

Power-scaling sensitivity analysis

Description

Calculates the prior/likelihood sensitivity based on power-scaling perturbations. This is done using importance sampling (and optionally moment matching).

Usage

```
powerscale_sensitivity(x, ...)

## Default S3 method:
powerscale_sensitivity(
  x,
  variable = NULL,
  variables = NULL,
  lower_alpha = 0.99,
  upper_alpha = 1.01,
  div_measure = "cjs_dist",
  measure_args = list(),
  component = c("prior", "likelihood"),
  sensitivity_threshold = 0.05,
  moment_match = FALSE,
  k_threshold = 0.5,
  resample = FALSE,
  transform = NULL,
  prediction = NULL,
  prior_selection = NULL,
  likelihood_selection = NULL,
  log_prior_name = "lprior",
  log_lik_name = "log_lik",
  separator = "_",
  num_args = NULL,
  ...
)

## S3 method for class 'priorsense_data'
powerscale_sensitivity(
  x,
  variable = NULL,
  variables = NULL,
  lower_alpha = 0.99,
  upper_alpha = 1.01,
  div_measure = "cjs_dist",
  measure_args = list(),
```

```

    component = c("prior", "likelihood"),
    sensitivity_threshold = 0.05,
    moment_match = FALSE,
    k_threshold = 0.5,
    resample = FALSE,
    transform = NULL,
    prediction = NULL,
    prior_selection = NULL,
    likelihood_selection = NULL,
    separator = "_",
    num_args = NULL,
    ...
)

```

```

## S3 method for class 'CmdStanFit'
powerscale_sensitivity(x, ...)

```

```

## S3 method for class 'stanfit'
powerscale_sensitivity(x, ...)

```

Arguments

<code>x</code>	Model fit object or <code>priorsense_data</code> object.
<code>...</code>	Further arguments passed to functions.
<code>variable</code>	Character vector of variables to check.
<code>variables</code>	alias of <code>variable</code> .
<code>lower_alpha</code>	Lower alpha value for gradient calculation.
<code>upper_alpha</code>	Upper alpha value for gradient calculation.
<code>div_measure</code>	Character (case sensitive) specifying the divergence measure to use. The following methods are implemented: • <code>"cjs_dist"</code>: Cumulative Jensen-Shannon distance. Default method. See function cjs_dist() for more details. • <code>"js_dist"</code>: Jensen-Shannon distance. First estimates density using stats::density(). Requires the philentropy package. See philentropy::jensen_shannon(). • <code>"js_div"</code>: Jensen-Shannon divergence. First estimates density using stats::density(). Requires the philentropy package. See philentropy::jensen_shannon(). • <code>"hellinger_dist"</code>: Hellinger distance. First estimates density using stats::density(). Requires the philentropy package. See philentropy::hellinger(). • <code>"kl_dist"</code>: Kullback-Leibler distance. First estimates density using stats::density(). Requires the philentropy package. See philentropy::kullback_leibler_distance(). • <code>"kl_div"</code>: Kullback-Leibler divergence. First estimates density using stats::density(). Requires the philentropy package. See philentropy::kullback_leibler_distance(). • <code>"ks_dist"</code>: Kolmogorov-Smirnov distance. See stats::ks.test(). • <code>"ws_dist"</code>: Wasserstein distance (pass <code>measure_args = list(p = N)</code>) for a different order, where <code>N</code> is the order. Requires the transport package. See transport::wasserstein1d().

measure_args	Named list of further arguments passed to divergence measure functions.
component	Character vector specifying component(s) to scale (default is both "prior" and "likelihood").
sensitivity_threshold	Threshold for flagging variable as sensitive to power-scaling.
moment_match	Logical; Indicate whether or not moment matching should be performed. Can only be TRUE if is_method is "psis".
k_threshold	Threshold value for Pareto k values above which the moment matching algorithm is used. Default is 0.5.
resample	Logical; Indicate whether or not draws should be resampled based on calculated importance weights.
transform	Indicate a transformation of posterior draws to perform before sensitivity analysis. Either "scale" or "whiten".
prediction	Function taking the model fit and returning a draws_df of predictions to be appended to the posterior draws
prior_selection	Vector specifying partitions of component to be included in power-scaling. Default is NULL, which takes all partitions. If this is a character, then it is appended to the variable name (specified by log_prior_name) with a separator between them. If numeric, then it is appended inside [].
likelihood_selection	Vector specifying partitions of component to be included in power-scaling. Default is NULL, which takes all partitions. If this is a character, then it is appended to the variable name (specified by log_lik_name) with a separator between them. If numeric, then it is appended inside [].
log_prior_name	Character (case sensitive) specifying name of the variable storing the log prior evaluations
log_lik_name	Character (case sensitive) specifying name of the variable storing the log likelihood evaluations
separator	Character specifying the separator between the component name and the partition. The default is "_".
num_args	(named list) Optional arguments passed to <code>num()</code> for pretty printing of summaries. Can be controlled globally via the <code>posterior.num_args</code> option.

Value

Table of sensitivity values for each specified variable.

References

- Kallioinen, N., Paananen, T., Bürkner, P-C., Vehtari, A. (2023). Detecting and diagnosing prior and likelihood sensitivity with power-scaling perturbations. *Statistics and Computing*. 34(57). doi:10.1007/s11222-023-10366-5
- Vehtari, A., Simpson, D., Gelman, A., Yao, Y., and Gabry, J. (2024). Pareto smoothed importance sampling. *Journal of Machine Learning Research*. 25(72). <https://jmlr.org/papers/v25/19-556.html>
- Paananen, T., Piironen, J., Bürkner, P-C., Vehtari, A. (2021). Implicitly adaptive importance sampling. *Statistics and Computing*. 31(16). doi:10.1007/s11222-020-09982-2

Examples

```
ex <- example_powerscale_model()
powerscale_sensitivity(ex$draws)
```

`powerscale_derivative` *Derivative with respect to power-scaling*

Description

Calculate the analytical derivative of a quantity with respect to power-scaling prior or likelihood.

Usage

```
powerscale_derivative(x, log_component, quantity = "mean", ...)
```

Arguments

<code>x</code>	draws object of posterior draws
<code>log_component</code>	numeric vector of log likelihood or log prior values
<code>quantity</code>	Character specifying quantity of interest (default is "mean"). Options are "mean", "sd", "var".
<code>...</code>	unused

Value

Derivative of the quantity with respect to log2 of the power-scaling factor (alpha).

Examples

```
example_model <- example_powerscale_model()
draws <- example_model$draws
log_prior <- log_prior_draws(draws, joint = TRUE)
posterior::summarise_draws(
  posterior::subset_draws(draws, variable = c("mu", "sigma")),
  mean,
  mean_sens = ~powerscale_derivative(.x, log_prior, quantity = "mean")
)
```

predictions_as_draws *brms predictions as draws*

Description

Create predictions using brms functions and convert them into draws format

Usage

```
predictions_as_draws(
  x,
  predict_fn,
  prediction_names = NULL,
  warn_dims = getOption("priorsense.warn", TRUE),
  ...
)
```

Arguments

x	brmsfit object
predict_fn	function for predictions
prediction_names	optional names of the predictions
warn_dims	throw a warning when coercing predict_fn's output from 3 margins to 2 margins?
...	further arguments passed to predict_fn

Value

draws array of predictions

Examples

```
## Not run:
library(brms)

if ("log_prior_draws.brmsfit" %in% methods(log_prior_draws) &&
    "log_lik_draws.brmsfit" %in% methods(log_lik_draws)) {
  fit <- brm(
    yield ~ N * P * K,
    data = npk,
    prior = prior(normal(0, 1), class = "b"),
    refresh = 0
  )

  powerscale_sensitivity(
    fit,
```

```
    variable = "_pred",  
    prediction = function(x) predictions_as_draws(  
      x, brms::posterior_epred  
    )  
  }  
  
## End(Not run)
```

Index

cjs_dist, 4
cjs_dist(), 10, 17, 19
create-priorsense-data, 5
create_priorsense_data
 (create-priorsense-data), 5

example_powerscale_model, 7

log_lik_draws, 7
log_prior_draws, 8

num(), 20

option, 20

philentropy::hellinger(), 10, 17, 19
philentropy::jensen_shannon(), 10, 17, 19
philentropy::kullback_leibler_distance(), 10, 17, 19
plot.powerscaled_sequence
 (powerscale-plots), 15
powerscale (powerscale-overview), 11
powerscale-gradients, 9
powerscale-overview, 3, 11
powerscale-plots, 3, 15, 15
powerscale-sensitivity, 18
powerscale.default
 (powerscale-overview), 11
powerscale.priorsense_data
 (powerscale-overview), 11
powerscale_derivative, 21
powerscale_gradients
 (powerscale-gradients), 9
powerscale_plot_dens
 (powerscale-plots), 15
powerscale_plot_ecdf, 2
powerscale_plot_ecdf
 (powerscale-plots), 15
powerscale_plot_quantities
 (powerscale-plots), 15

powerscale_sensitivity, 2
powerscale_sensitivity
 (powerscale-sensitivity), 18
powerscale_sequence, 2
powerscale_sequence
 (powerscale-overview), 11
predictions_as_draws, 22
priorsense (priorsense-package), 2
priorsense-package, 2

stats::density(), 10, 17, 19
stats::ks.test(), 10, 17, 19

transport::wasserstein1d(), 10, 17, 19