

Package ‘tidyposterior’

November 15, 2018

Version 0.0.2

Title Bayesian Analysis to Compare Models using Resampling Statistics

Description Bayesian analysis used here to answer the question: “when looking at resampling results, are the differences between models ‘real?’” To answer this, a model can be created where the performance statistic is the resampling statistics (e.g. accuracy or RMSE). These values are explained by the model types. In doing this, we can get parameter estimates for each model's affect on performance and make statistical (and practical) comparisons between models. The methods included here are similar to Benavoli et al (2017) <<http://jmlr.org/papers/v18/16-305.html>>.

License GPL-2

URL <https://tidymodels.github.io/tidyposterior>

BugReports <https://github.com/tidymodels/tidyposterior/issues>

VignetteBuilder knitr

Encoding UTF-8

LazyData true

ByteCompile true

RoxygenNote 6.1.0.9000

Depends R (>= 2.10)

Imports rsample (>= 0.0.2), tidyr (>= 0.7.1), dplyr, rstanarm (>= 2.15.3), rlang, utils, purrr, tibble, generics, ggplot2,

Suggests knitr, testthat, covr

NeedsCompilation no

Author Max Kuhn [aut, cre],
RStudio [cph]

Maintainer Max Kuhn <max@rstudio.com>

Repository CRAN

Date/Publication 2018-11-15 20:20:03 UTC

R topics documented:

contrast_models	2
ggplot.posterior	3
ggplot.posterior_diff	4
no_trans	5
perf_mod	6
precise_example	8
summary.posterior	9
summary.posterior_diff	10
tidy.perf_mod	11
Index	12

contrast_models	<i>Estimate the Difference Between Models</i>
-----------------	---

Description

The posterior distributions created by `perf_mod()` can be used to obtain the posterior distribution of the difference(s) between models. One or more comparisons can be computed at the same time.

Usage

```
contrast_models(x, list_1 = NULL, list_2 = NULL,
  seed = sample.int(10000, 1))
```

Arguments

<code>x</code>	An object produced by <code>perf_mod()</code> .
<code>list_1</code> , <code>list_2</code>	Character vectors of equal length that specify the specific pairwise contrasts. The contrast is parameterized as <code>list_1[i] - list_2[i]</code> . If the defaults are left to <code>NULL</code> , all combinations are evaluated.
<code>seed</code>	A single integer for sampling from the posterior.

Details

If a transformation was used when `x` was created, the inverse is applied *before* the difference is computed.

Value

A data frame of the posterior distribution(s) of the difference(s). The object has an extra class of "posterior_diff".

Examples

```
# Example objects from the "Getting Started" vignette at
# https://topepo.github.io/tidy posterior/articles/Getting_Started.html

# File for pre-run model is at
ex_dat <- "https://bit.ly/2S1v6H9"

# load(load(url(ex_dat))

# head(glm_v_nnet)
```

ggplot.posterior

Visualize the Posterior Distributions of Model Statistics

Description

A simple violin plot is created by the function.

Usage

```
## S3 method for class 'posterior'
ggplot(data, mapping = NULL, ...,
        environment = NULL, reorder = TRUE)
```

Arguments

data An object produced by `tidy.perf_mod()`.
mapping, ..., environment Not currently used.
reorder A logical; should the model column be reordered by the average of the posterior distribution?

Value

A `ggplot2::ggplot()` object using `ggplot2::geom_violin()` for the posteriors.

Examples

```
# Example objects from the "Getting Started" vignette at
# https://topepo.github.io/tidy posterior/articles/Getting_Started.html

# File for pre-run model is at
ex_dat <- "https://bit.ly/20Jdv11"

# load(load(url(ex_dat))

# posterior_values <- tidy(roc_model)
```

```
# library(ggplot2)
# ggplot(posterior_values) + theme_bw()
```

ggplot.posterior_diff *Visualize the Posterior Distributions of Model Differences*

Description

A density is created for each contrast in a faceted grid.

Usage

```
## S3 method for class 'posterior_diff'
ggplot(data, mapping = NULL, ...,
        environment = NULL, size = 0)
```

Arguments

data	An object produced by <code>contrast_models()</code> .
mapping, ..., environment	Not currently used.
size	The size of an effective difference. For example, a 5% increase in accuracy between two models might be considered a "real" difference.

Value

A `ggplot2::ggplot()` object using `geom_density` faceted by the models being contrasted (when there are 2 or more contrasts).

Examples

```
# Example objects from the "Getting Started" vignette at
# https://topepo.github.io/tidy posterior/articles/Getting_Started.html

# File for pre-run model is at
ex_dat <- "https://bit.ly/2S1v6H9"

# load(load(url(ex_dat))

# library(ggplot2)
# ggplot(glm_v_nnet, size = 0.02) + theme_bw()
```

Description

A set of objects are contained here to easily facilitate the use of outcome transformations for modeling. For example, if there is a large amount of variability in the resampling results for the Kappa statistics, which lies between -1 and 1, assuming normality may produce posterior estimates outside of the natural bound. One way to solve this is to use a link function or assume a prior that is appropriately bounded. Another approach is to transform the outcome values prior to modeling using a Gaussian prior and reverse-transforming the posterior estimates prior to visualization and summarization. These object can help facilitate this last approach.

Usage

no_trans

logit_trans

Fisher_trans

ln_trans

inv_trans

Format

An object of class `list` of length 2.

Details

The `logit_trans` object is useful for model performance statistics bounds in zero and one, such as accuracy or the area under the ROC curve.

`ln_trans` and `inv_trans` can be useful when the statistics are right-skewed and strictly positive.

`Fisher_trans` was originally used for correlation statistics but can be used here for an metrics falling between -1 and 1, such as Kappa.

Examples

```
logit_trans$func(.5)
logit_trans$inv(0)
```

Description

Bayesian analysis used here to answer the question: "when looking at resampling results, are the differences between models 'real?'" To answer this, a model can be created were the *outcome* is the resampling statistics (e.g. accuracy or RMSE). These values are explained by the model types. In doing this, we can get parameter estimates for each model's affect on performance and make statistical (and practical) comparisons between models.

Usage

```
perf_mod(object, ...)

## S3 method for class 'rset'
perf_mod(object, transform = no_trans,
         hetero_var = FALSE, ...)

## S3 method for class 'vfold_cv'
perf_mod(object, transform = no_trans,
         hetero_var = FALSE, ...)

## S3 method for class 'resamples'
perf_mod(object, transform = no_trans,
         hetero_var = FALSE, metric = object$metrics[1], ...)

## S3 method for class 'data.frame'
perf_mod(object, transform = no_trans,
         hetero_var = FALSE, ...)
```

Arguments

object	A data frame or an rset object (such as <code>rsample::vfold_cv()</code>) containing the id column(s) and at least two numeric columns of model performance statistics (e.g. accuracy). Additionally, an object from <code>caret::resamples</code> can be used.
...	Additional arguments to pass to <code>rstanarm::stan_glmer()</code> such as <code>verbose</code> , <code>prior</code> , <code>seed</code> , <code>family</code> , etc.
transform	An named list of transformation and inverse transformation fuctions. See <code>logit_trans()</code> as an example.
hetero_var	A logical; if TRUE, then different variances are estimated for each model group. Otherwise, the same variance is used for each group. Estimating heterogeneous variances may slow or prevent convergence.
metric	A single character value for the statstic from the <code>resamples</code> object that should be analyzed.

Details

These functions can be used to process and analyze matched resampling statistics from different models using a Bayesian generalized linear model with effects for the model and the resamples.

By default, a generalized linear model with Gaussian error and an identity link is fit to the data and has terms for the predictive model grouping variable. In this way, the performance metrics can be compared between models.

Additionally, random effect terms are also used. For most resampling methods (except repeated V -fold cross-validation), a simple random intercept model is used with an exchangeable (i.e. compound-symmetric) variance structure. In the case of repeated cross-validation, two random intercept terms are used; one for the repeat and another for the fold within repeat. These also have exchangeable correlation structures.

The above model specification assumes that the variance in the performance metrics is the same across models. However, this is unlikely to be true in some cases. For example, for simple binomial accuracy, it is well known that the variance is highest when the accuracy is near 50 percent. When the argument `hetero_var = TRUE`, the variance structure uses random intercepts for each model term. This may produce more realistic posterior distributions but may take more time to converge.

Also, as shown in the package vignettes, the Gaussian assumption may be unrealistic. In this case, there are at least two approaches that can be used. First, the outcome statistics can be transformed prior to fitting the model. For example, for accuracy, the logit transformation can be used to convert the outcome values to be on the real line and a model is fit to these data. Once the posterior distributions are computed, the inverse transformation can be used to put them back into the original units. The `transform` argument can be used to do this.

The second approach would be to use a different error distribution from the exponential family. For RMSE values, the Gamma distribution may produce better results at the expense of model computational complexity. This can be achieved by passing the `family` argument to `perf_mod` as one might with the `glm` function.

Value

An object of class `perf_mod`.

Examples

```
# Example objects from the "Getting Started" vignette at
# https://topepo.github.io/tidy posterior/articles/Getting_Started.html

# File for pre-run model is at
ex_dat <- "https://bit.ly/20Jdv11"

# load(load(url(ex_dat))

# roc_model

# Summary method shows the underlying `stan` model
# summary(roc_model)
```

precise_example	<i>Example Data Sets</i>
-----------------	--------------------------

Description

Example Data Sets

Details

Several data sets are contained in the package as examples. Each *simulates* an `rset` object but the `splits` columns are not included to save space.

- `precise_example` contains the results of the classification analysis of a real data set using 10-fold CV. The holdout data sets contained thousands of examples and have precise performance estimates. Three models were fit to the original data and several performance metrics are included.
- `noisy_example` was also generated from a regression data simulation. The original data set was small (50 samples) and 10-repeated of 10-fold CV were used with four models. There is an excessive of variability in the results (probably more than the resample-to-resample variability). The RMSE distributions show fairly right-skewed distributions.
- `concrete_example` contains the results of the regression case study from the book *Applied Predictive Modeling*. The original data set contained 745 samples in the training set. 10-repeats of 10-fold CV was also used and 13 models were fit to the data.
- `ts_example` is from a data set where rolling-origin forecast resampling was used. Each assessment set is the summary of 14 observations (i.e. 2 weeks). The analysis set consisted of a base of about 5,500 samples plus the previous assessment sets. Four regression models were applied to these data.

Value

Tibbles with the additional class `rset`

Examples

```
data(precise_example)
precise_example
```

summary.posterior	<i>Summarize the Posterior Distributions of Model Statistics</i>
-------------------	--

Description

Numerical summaries are created for each model including the posterior mean and upper and lower credible intervals (aka uncertainty intervals).

Usage

```
## S3 method for class 'posterior'  
summary(object, prob = 0.9,  
        seed = sample.int(10000, 1), ...)
```

Arguments

object	An object produced by <code>tidy.perf_mod()</code> .
prob	A number p ($0 < p < 1$) indicating the desired probability mass to include in the intervals.
seed	A single integer for sampling from the posterior.
...	Not currently used

Value

A data frame with summary statistics and a row for each model.

Examples

```
# Example objects from the "Getting Started" vignette at  
# https://topepo.github.io/tidyposterior/articles/Getting\_Started.html  
  
# File for pre-run model is at  
ex_dat <- "https://bit.ly/20Jdv11"  
  
# load(load(url(ex_dat))  
  
# posterior_values <- tidy(roc_model)  
# summary(posterior_values)
```

`summary.posterior_diff`*Summarize Posterior Distributions of Model Differences*

Description

Credible intervals are created for the differences. Also, region of practical equivalence (ROPE) statistics are computed when the effective size of a difference is given.

Usage

```
## S3 method for class 'posterior_diff'  
summary(object, prob = 0.9, size = 0, ...)
```

Arguments

<code>object</code>	An object produced by <code>contrast_models()</code> .
<code>prob</code>	A number p ($0 < p < 1$) indicating the desired probability mass to include in the intervals.
<code>size</code>	The size of an effective difference. For example, a 5% increase in accuracy between two models might be considered a "real" difference.
<code>...</code>	Not currently used

Details

The ROPE estimates included in the results are the columns `pract_neg`, `pract_equiv`, and `pract_pos`. `pract_neg` integrates the portion of the posterior below `-size` (and `pract_pos` is the upper integral starting at `size`). The interpretation depends on whether the metric being analyzed is better when larger or smaller. `pract_equiv` integrates between `[-size, size]`. If this is close to one, the two models are unlikely to be practically different relative to `size`.

Value

A data frame with interval and ROPE statistics for each comparison.

Examples

```
# Example objects from the "Getting Started" vignette at  
# https://topepo.github.io/tidy posterior/articles/Getting\_Started.html  
  
# File for pre-run model is at  
ex_dat <- "https://bit.ly/2S1v6H9"  
  
# load(load(url(ex_dat))  
  
# summary(glm_v_nnet, size = 0.02)
```

`tidy.perf_mod`*Extract Posterior Distributions for Models*

Description

`tidy` can be used on an object produced by `perf_mod()` to create a data frame with a column for the model name and the posterior predictive distribution values.

Usage

```
## S3 method for class 'perf_mod'
tidy(x, seed = sample.int(10000, 1), ...)
```

Arguments

<code>x</code>	An object from <code>perf_mod()</code>
<code>seed</code>	A single integer for sampling from the posterior.
<code>...</code>	Not currently used

Details

Note that this posterior only reflects the variability of the groups (i.e. the fixed effects). This helps answer the question of which model is best *for this data set*. It does not answer the question of which model would be best on a new resample of the data (which would have greater variability).

Value

A data frame with the additional class "posterior"

Examples

```
# Example objects from the "Getting Started" vignette at
# https://topepo.github.io/tidyposterior/articles/Getting_Started.html

# File for pre-run model is at
ex_dat <- "https://bit.ly/20Jdv11"

# load(load(url(ex_dat))

# roc_model
# posterior_values <- tidy(roc_model)
# head(posterior_values)
# class(posterior_values)
```

Index

*Topic **datasets**

- no_trans, [5](#)
- precise_example, [8](#)

concrete_example (precise_example), [8](#)
contrast_models, [2](#)
contrast_models(), [4](#), [10](#)

Fisher_trans (no_trans), [5](#)

ggplot.posterior, [3](#)
ggplot.posterior_diff, [4](#)
ggplot2::geom_violin(), [3](#)
ggplot2::ggplot(), [3](#), [4](#)

inv_trans (no_trans), [5](#)

ln_trans (no_trans), [5](#)
logit_trans (no_trans), [5](#)
logit_trans(), [6](#)

no_trans, [5](#)
noisy_example (precise_example), [8](#)

perf_mod, [6](#)
perf_mod(), [2](#), [11](#)
precise_example, [8](#)

rsample::vfold_cv(), [6](#)
rstanarm::stan_glmer(), [6](#)

summary.posterior, [9](#)
summary.posterior_diff, [10](#)

tidy.perf_mod, [11](#)
tidy.perf_mod(), [3](#), [9](#)
ts_example (precise_example), [8](#)