

Package: parsnip (via r-universe)

September 27, 2026

Title A Common API to Modeling and Analysis Functions

Version 1.6.1.9000

Maintainer Max Kuhn <max@posit.co>

Description A common interface is provided to allow users to specify a model without having to remember the different argument names across different functions or computational engines (e.g. 'R', 'Spark', 'Stan', 'H2O', etc).

License MIT + file LICENSE

URL <https://github.com/tidymodels/parsnip>,
<https://parsnip.tidymodels.org/>

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

Depends R (>= 4.1)

Imports cli, dplyr (>= 1.1.0), generics (>= 0.1.2), ggplot2, globals, glue, hardhat (>= 1.4.1), lifecycle, magrittr, pillar, prettyunits, purrr (>= 1.0.0), rlang (>= 1.1.0), sparsevctrs (>= 0.2.0), stats, tibble (>= 2.1.1), tidyr (>= 1.3.0), utils, vctrs (>= 0.6.0)

Suggests bench, C50, covr, dials (>= 1.4.3), earth, ggrepel, keras, keras3, kernlab, kknn, knitr, LiblineaR, MASS, Matrix, methods, mgcv, modeldata, nlme, prodlm, ranger (>= 0.12.0), remotes, rmarkdown, rpart, sparklyr (>= 1.0.0), survival, tensorflow, testthat (>= 3.0.0), withr, xgboost (>= 1.5.0.1)

VignetteBuilder knitr

ByteCompile true

Config/Needs/website parsnip, tidymodels/tidymodels, tidyverse/tidytemplate, xgboost, rmarkdown, torch

Config/rcmdcheck/ignore-inconsequential-notes true

Config/testthat/edition 3

Encoding UTF-8

LazyData true

Roxygen list(markdown = TRUE)
Config/roxygen2/version 8.1.0
Config/pak/sysreqs libicu-dev
Repository <https://tidymodels.r-universe.dev>
Date/Publication 2026-09-27 15:09:39 UTC
RemoteUrl <https://github.com/tidymodels/parsnip>
RemoteRef HEAD
RemoteSha c010004ac4acd59f890324b9848b25cdb68a5039

Contents

.extract_surv_status	3
.extract_surv_time	4
.get_prediction_column_names	4
add_rowindex	5
augment.model_fit	6
auto_ml	8
autoplot.model_fit	9
bag_mars	10
bag_mlp	11
bag_tree	12
bart	13
boost_tree	16
C5_rules	18
case_weights	19
case_weights_allowed	20
control_parsnip	20
ctree_train	21
cubist_rules	22
decision_tree	24
descriptors	26
discrim_flexible	28
discrim_linear	29
discrim_quad	30
discrim_regularized	31
extract-parsnip	33
fit.model_spec	34
gen_additive_mod	37
glance.model_fit	38
glm_grouped	39
linear_reg	40
logistic_reg	42
mars	44
max_mtry_formula	46
maybe_matrix	46

<code>min_cols</code>	47
<code>mlp</code>	48
<code>model_fit</code>	49
<code>model_formula</code>	50
<code>model_spec</code>	52
<code>multi_predict</code>	54
<code>multinom_reg</code>	56
<code>naive_Bayes</code>	58
<code>nearest_neighbor</code>	59
<code>null_model</code>	60
<code>parsnip_addin</code>	62
<code>pls</code>	63
<code>poisson_reg</code>	64
<code>rand_forest</code>	65
<code>repair_call</code>	67
<code>req_pkgs</code>	68
<code>required_pkgs.model_spec</code>	69
<code>rule_fit</code>	70
<code>set_args</code>	72
<code>set_engine</code>	73
<code>show_engines</code>	74
<code>sparse_data</code>	75
<code>svm_linear</code>	75
<code>svm_poly</code>	76
<code>svm_rbf</code>	78
<code>tabular_auto_int</code>	79
<code>tabular_chronos</code>	81
<code>tabular_icl</code>	83
<code>tabular_pfn</code>	84
<code>tabular_resnet</code>	85
<code>tabular_rln</code>	87
<code>tabular_saint</code>	89
<code>tidy.model_fit</code>	91
<code>translate</code>	92
<code>update.bag_mars</code>	93

Index**108**

<code>.extract_surv_status</code>	<i>Extract survival status</i>
-----------------------------------	--------------------------------

Description

Extract the status from a `survival::Surv()` object.

Usage

```
.extract_surv_status(surv)
```

Arguments

`surv` A single `survival::Surv()` object.

Value

A numeric vector.

`.extract_surv_time` *Extract survival time*

Description

Extract the time component(s) from a `survival::Surv()` object.

Usage

```
.extract_surv_time(surv)
```

Arguments

`surv` A single `survival::Surv()` object.

Value

A vector when the type is "right" or "left" and a tibble otherwise.

`.get_prediction_column_names`
Obtain names of prediction columns for a fitted model or workflow

Description

`.get_prediction_column_names()` returns a list that has the names of the columns for the primary prediction types for a model.

Usage

```
.get_prediction_column_names(x, syms = FALSE)
```

Arguments

`x` A fitted parsnip model (class "model_fit") or a fitted workflow.

`syms` Should the column names be converted to symbols? Defaults to FALSE.

Value

A list with elements "estimate" and "probabilities".

Examples

```
library(dplyr)
library(modeldata)
data("two_class_dat")

levels(two_class_dat$Class)
lr_fit <- logistic_reg() |> fit(Class ~ ., data = two_class_dat)

.get_prediction_column_names(lr_fit)
.get_prediction_column_names(lr_fit, syms = TRUE)
```

add_rowindex

Add a column of row numbers to a data frame

Description

Add a column of row numbers to a data frame

Usage

```
add_rowindex(x)
```

Arguments

x A data frame

Value

The same data frame with a column of 1-based integers named `.row`.

Examples

```
mtcars |> add_rowindex()
```

augment.model_fit	<i>Augment data with predictions</i>
-------------------	--------------------------------------

Description

augment() will add column(s) for predictions to the given data.

Usage

```
## S3 method for class 'model_fit'
augment(x, new_data, eval_time = NULL, ...)
```

Arguments

x	A model fit produced by fit.model_spec() or fit_xy.model_spec() .
new_data	A data frame or matrix.
eval_time	For censored regression models, a vector of time points at which the survival probability is estimated.
...	Not currently used.

Details

Regression:

For regression models, a .pred column is added. If x was created using [fit.model_spec\(\)](#) and new_data contains a regression outcome column, a .resid column is also added.

Classification:

For classification models, the results can include a column called .pred_class as well as class probability columns named .pred_{level}. This depends on what type of prediction types are available for the model.

Censored Regression:

For these models, predictions for the expected time and survival probability are created (if the model engine supports them). If the model supports survival prediction, the eval_time argument is required.

If survival predictions are created and the survival outcome can be resolved from new_data (either as a literal [survival::Surv\(\)](#) column or from the variables referenced in the model's formula LHS), additional columns are added for inverse probability of censoring weights (IPCW) (see [tidymodels.org](#) page in the references below). This enables the user to compute performance metrics in the [yardstick](#) package. The same columns can be obtained directly from [predict.model_fit\(\)](#) by setting add_censoring_weights = TRUE.

Quantile Regression:

For quantile regression models, a .pred_quantile column is added that contains the quantile predictions for each row. This column has a special class "quantile_pred" and can be unnested using [tidyr::unnest\(\)](#)

References

<https://www.tidymodels.org/learn/statistics/survival-metrics/>

Examples

```
car_trn <- mtcars[11:32,]
car_tst <- mtcars[ 1:10,]

reg_form <-
  linear_reg() |>
  set_engine("lm") |>
  fit(mpg ~ ., data = car_trn)
reg_xy <-
  linear_reg() |>
  set_engine("lm") |>
  fit_xy(car_trn[, -1], car_trn$mpg)

augment(reg_form, car_tst)
augment(reg_form, car_tst[, -1])

augment(reg_xy, car_tst)
augment(reg_xy, car_tst[, -1])

# -----

data(two_class_dat, package = "modeldata")
cls_trn <- two_class_dat[-(1:10), ]
cls_tst <- two_class_dat[ 1:10 , ]

cls_form <-
  logistic_reg() |>
  set_engine("glm") |>
  fit(Class ~ ., data = cls_trn)
cls_xy <-
  logistic_reg() |>
  set_engine("glm") |>
  fit_xy(cls_trn[, -3],
    cls_trn$Class)

augment(cls_form, cls_tst)
augment(cls_form, cls_tst[, -3])

augment(cls_xy, cls_tst)
augment(cls_xy, cls_tst[, -3])

# -----

if (rlang::is_installed("quantreg")) {
  # Quantile regression example
  qr_form <-
    linear_reg() |>
    set_engine("quantreg") |>
```

```

set_mode("quantile regression", quantile_levels = c(0.25, 0.5, 0.75)) |>
fit(mpg ~ ., data = car_trn)

augment(qr_form, car_tst)
augment(qr_form, car_tst[, -1])
}

```

auto_ml

Automatic Machine Learning

Description

auto_ml() defines an automated searching and tuning process where many models of different families are trained and ranked given their performance on the training data.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [h2o](#)^{1 2}

¹ The default engine. ² Requires a parsnip extension package for classification and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
auto_ml(mode = "unknown", engine = "h2o")
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
engine	A single character string specifying what computational engine to use for fitting.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```

value <- 1
auto_ml(argument = !!value)

```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [h2o engine details](#)

autoplot.model_fit	Create a ggplot for a model object
--------------------	------------------------------------

Description

This method provides a good visualization method for model results. Currently, only methods for glmnet models are implemented.

Usage

```
## S3 method for class 'model_fit'
autoplot(object, ...)

## S3 method for class 'glmnet'
autoplot(object, ..., min_penalty = 0, best_penalty = NULL, top_n = 3L)
```

Arguments

object	A model fit object.
...	For autoplot.glmnet() , options to pass to ggrepel::geom_label_repel() . Otherwise, this argument is ignored.
min_penalty	A single, non-negative number for the smallest penalty value that should be shown in the plot. If left NULL, the whole data range is used.
best_penalty	A single, non-negative number that will show a vertical line marker. If left NULL, no line is shown. When this argument is used, the ggrepel package is required.
top_n	A non-negative integer for how many model predictors to label. The top predictors are ranked by their absolute coefficient value. For multinomial or multivariate models, the top_n terms are selected within class or response, respectively.

Details

The **glmnet** package will need to be attached or loaded for its `autoplot()` method to work correctly.

Value

A ggplot object with penalty on the x-axis and coefficients on the y-axis. For multinomial or multivariate models, the plot is faceted.

bag_mars

*Ensembles of MARS models***Description**

`bag_mars()` defines an ensemble of generalized linear models that use artificial features for some predictors. These features resemble hinge functions and the result is a model that is a segmented regression in small dimensions. This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [earth](#)^{1 2}

¹ The default engine. ² Requires a `parSNIP` extension package for classification and regression.

More information on how **parSNIP** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
bag_mars(
  mode = "unknown",
  num_terms = NULL,
  prod_degree = NULL,
  prune_method = NULL,
  engine = "earth"
)
```

Arguments

<code>mode</code>	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
<code>num_terms</code>	The number of features that will be retained in the final model, including the intercept.
<code>prod_degree</code>	The highest possible interaction degree.
<code>prune_method</code>	The pruning method.
<code>engine</code>	A single character string specifying what computational engine to use for fitting.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See `set_engine()` for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the `fit()` function is used with the data.

Each of the arguments in this function other than `mode` and `engine` are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
bag_mars(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [earth engine details](#)

bag_mlp

Ensembles of neural networks

Description

`bag_mlp()` defines an ensemble of single layer, feed-forward neural networks. This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [nnet](#)^{1 2}

¹ The default engine. ² Requires a parsnip extension package for classification and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
bag_mlp(
  mode = "unknown",
  hidden_units = NULL,
  penalty = NULL,
  epochs = NULL,
  engine = "nnet"
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
hidden_units	An integer for the number of units in the hidden model.
penalty	A non-negative numeric value for the amount of weight decay.
epochs	An integer for the number of training iterations.
engine	A single character string specifying what computational engine to use for fitting.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
bag_mlp(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [nnet engine details](#)

bag_tree

Ensembles of decision trees

Description

`bag_tree()` defines an ensemble of decision trees. This function can fit classification, regression, and censored regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [rpart](#)¹²
- [C5.0](#)²

¹ The default engine. ² Requires a parsnip extension package for censored regression, classification, and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
bag_tree(
  mode = "unknown",
  cost_complexity = 0,
  tree_depth = NULL,
  min_n = 2,
  class_cost = NULL,
  engine = "rpart"
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", "classification", or "censored regression".
cost_complexity	A positive number for the the cost/complexity parameter (a.k.a. Cp) used by CART models (specific engines only).
tree_depth	An integer for the maximum depth of the tree (i.e. number of splits) (specific engines only).
min_n	An integer for the minimum number of data points in a node that is required for the node to be split further.
class_cost	A non-negative scalar for a class cost (where a cost of 1 means no extra cost). This is useful for when the first level of the outcome factor is the minority class. If this is not the case, values between zero and one can be used to bias to the second level of the factor.
engine	A single character string specifying what computational engine to use for fitting.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
bag_tree(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [rpart engine details](#), [C5.0 engine details](#)

Description

`bart()` defines a tree ensemble model that uses Bayesian analysis to assemble the ensemble. This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [dbarts](#)¹

¹ The default engine.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
bart(
  mode = "unknown",
  engine = "dbarts",
  trees = NULL,
  prior_terminal_node_coef = NULL,
  prior_terminal_node_expo = NULL,
  prior_outcome_range = NULL
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
engine	A single character string specifying what computational engine to use for fitting.
trees	An integer for the number of trees contained in the ensemble.
prior_terminal_node_coef	A coefficient for the prior probability that a node is a terminal node. Values are usually between 0 and one with a default of 0.95. This affects the baseline probability; smaller numbers make the probabilities larger overall. See Details below.
prior_terminal_node_expo	An exponent in the prior probability that a node is a terminal node. Values are usually non-negative with a default of 2. This affects the rate that the prior probability decreases as the depth of the tree increases. Larger values make deeper trees less likely.
prior_outcome_range	A positive value that defines the width of a prior that the predicted outcome is within a certain range. For regression it is related to the observed range of the data; the prior is the number of standard deviations of a Gaussian distribution defined by the observed range of the data. For classification, it is defined as the range of +/-3 (assumed to be on the logit scale). The default value is 2.

Details

The prior for the terminal node probability is expressed as $\text{prior} = a * (1 + d)^{-b}$ where d is the depth of the node, a is `prior_terminal_node_coef` and b is `prior_terminal_node_expo`. See the Examples section below for an example graph of the prior probability of a terminal node for different values of these parameters.

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See `set_engine()` for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the `fit()` function is used with the data.

Each of the arguments in this function other than `mode` and `engine` are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
bart(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [dbarts engine details](#)

Examples

```
show_engines("bart")

bart(mode = "regression", trees = 5)

# -----
# Examples for terminal node prior

library(ggplot2)
library(dplyr)

prior_test <- function(coef = 0.95, expo = 2, depths = 1:10) {
  tidyr::crossing(coef = coef, expo = expo, depth = depths) |>
    mutate(
      `terminal node prior` = coef * (1 + depth)^(-expo),
      coef = format(coef),
      expo = format(expo)
    )
}

prior_test(coef = c(0.05, 0.5, .95), expo = c(1/2, 1, 2)) |>
  ggplot(aes(depth, `terminal node prior`, col = coef)) +
  geom_line() +
  geom_point() +
  facet_wrap(~ expo)
```

boost_tree

*Boosted trees***Description**

`boost_tree()` defines a model that creates a series of decision trees forming an ensemble. Each tree depends on the results of previous trees. All trees in the ensemble are combined to produce a final prediction. This function can fit classification, regression, and censored regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [xgboost](#)¹
- [C5.0](#)
- [catboost](#)²
- [h2o](#)²
- [lightgbm](#)²
- [mboost](#)²
- [spark](#)

¹ The default engine. ² Requires a **parsnip** extension package for censored regression, classification, and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
boost_tree(
  mode = "unknown",
  engine = "xgboost",
  mtry = NULL,
  trees = NULL,
  min_n = NULL,
  tree_depth = NULL,
  learn_rate = NULL,
  loss_reduction = NULL,
  sample_size = NULL,
  stop_iter = NULL
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", "classification", or "censored regression".
engine	A single character string specifying what computational engine to use for fitting.

mtry	A number for the number (or proportion) of predictors that will be randomly sampled at each split when creating the tree models (specific engines only).
trees	An integer for the number of trees contained in the ensemble.
min_n	An integer for the minimum number of data points in a node that is required for the node to be split further.
tree_depth	An integer for the maximum depth of the tree (i.e. number of splits) (specific engines only).
learn_rate	A number for the rate at which the boosting algorithm adapts from iteration-to-iteration (specific engines only). This is sometimes referred to as the shrinkage parameter.
loss_reduction	A number for the reduction in the loss function required to split further (specific engines only).
sample_size	A number for the number (or proportion) of data that is exposed to the fitting routine. For xgboost, the sampling is done at each iteration while C5.0 samples once during training.
stop_iter	The number of iterations without improvement before stopping (specific engines only).

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
boost_tree(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [xgboost engine details](#), [C5.0 engine details](#), [catboost engine details](#), [h2o engine details](#), [lightgbm engine details](#), [mboost engine details](#), [spark engine details](#), [xgb_train\(\)](#), [C5.0_train\(\)](#)

Examples

```
show_engines("boost_tree")

boost_tree(mode = "classification", trees = 20)
```

C5_rules

*C5.0 rule-based classification models***Description**

C5_rules() defines a model that derives feature rules from a tree for prediction. A single tree or boosted ensemble can be used. This function can fit classification models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [C5.0](#)^{1,2}

¹ The default engine. ² Requires a parsnip extension package.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
C5_rules(mode = "classification", trees = NULL, min_n = NULL, engine = "C5.0")
```

Arguments

mode	A single character string for the type of model. The only possible value for this model is "classification".
trees	A non-negative integer (no greater than 100) for the number of members of the ensemble.
min_n	An integer greater between zero and nine for the minimum number of data points in a node that are required for the node to be split further.
engine	A single character string specifying what computational engine to use for fitting.

Details

C5.0 is a classification model that is an extension of the C4.5 model of Quinlan (1993). It has tree- and rule-based versions that also include boosting capabilities. C5_rules() enables the version of the model that uses a series of rules (see the examples below). To make a set of rules, an initial C5.0 tree is created and flattened into rules. The rules are pruned, simplified, and ordered. Rule sets are created within each iteration of boosting.

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
C5_rules(argument = !!value)
```

References

Quinlan R (1993). *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers.

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[C50::C5.0\(\)](#), [C50::C5.0Control\(\)](#), [fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [C5.0 engine details](#)

Examples

```
show_engines("C5_rules")
```

```
C5_rules()
```

case_weights	<i>Using case weights with parsnip</i>
--------------	--

Description

Case weights are positive numeric values that influence how much each data point has during the model fitting process. There are a variety of situations where case weights can be used.

Details

tidymodels packages differentiate *how* different types of case weights should be used during the entire data analysis process, including preprocessing data, model fitting, performance calculations, etc.

The tidymodels packages require users to convert their numeric vectors to a vector class that reflects how these should be used. For example, there are some situations where the weights should not affect operations such as centering and scaling or other preprocessing operations.

The types of weights allowed in tidymodels are:

- Frequency weights via [hardhat::frequency_weights\(\)](#)
- Importance weights via [hardhat::importance_weights\(\)](#)

More types can be added by request.

For parsnip, the [fit\(\)](#) and [fit_xy\(\)](#) functions contain a `case_weight` argument that takes these data. For Spark models, the argument value should be a character value.

See Also

[frequency_weights\(\)](#), [importance_weights\(\)](#), [fit\(\)](#), [fit_xy\(\)](#)

case_weights_allowed	<i>Determine if case weights are used</i>
----------------------	---

Description

Not all modeling engines can incorporate case weights into their calculations. This function can determine whether they can be used.

Usage

```
case_weights_allowed(spec)
```

Arguments

spec	A parsnip model specification .
------	---

Value

A single logical.

Examples

```
case_weights_allowed(linear_reg())
case_weights_allowed(linear_reg(engine = "keras"))
```

control_parsnip	<i>Control the fit function</i>
-----------------	---------------------------------

Description

Pass options to the [fit.model_spec\(\)](#) function to control its output and computations

Usage

```
control_parsnip(verbosity = 1L, catch = FALSE)
```

Arguments

verbosity	An integer to control how verbose the output is. For a value of zero, no messages or output are shown when packages are loaded or when the model is fit. For a value of 1, package loading is quiet but model fits can produce output to the screen (depending on if they contain their own verbose-type argument). For a value of 2 or more, any output at all is displayed and the execution time of the fit is recorded and printed.
catch	A logical where a value of TRUE will evaluate the model inside of <code>try(, silent = TRUE)</code> . If the model fails, an object is still returned (without an error) that inherits the class "try-error".

Value

An S3 object with class "control_parsnip" that is a named list with the results of the function call

Examples

```
control_parsnip(verbosity = 2L)
```

ctree_train	<i>A wrapper function for conditional inference tree models</i>
-------------	---

Description

These functions are slightly different APIs for `partykit::ctree()` and `partykit::cforest()` that have several important arguments as top-level arguments (as opposed to being specified in `partykit::ctree_control()`).

Usage

```
ctree_train(  
  formula,  
  data,  
  weights = NULL,  
  minsplit = 20L,  
  maxdepth = Inf,  
  teststat = "quadratic",  
  testtype = "Bonferroni",  
  mincriterion = 0.95,  
  ...  
)  
  
cforest_train(  
  formula,  
  data,  
  weights = NULL,  
  minsplit = 20L,  
  maxdepth = Inf,  
  teststat = "quadratic",  
  testtype = "Univariate",  
  mincriterion = 0,  
  mtry = ceiling(sqrt(ncol(data) - 1)),  
  ntree = 500L,  
  ...  
)
```

Arguments

formula	A symbolic description of the model to be fit.
data	A data frame containing the variables in the model.
weights	A vector of weights whose length is the same as <code>nrow(data)</code> . For <code>partykit::ctree()</code> models, these are required to be non-negative integers while for <code>partykit::cforest()</code> they can be non-negative integers or doubles.
minsplit	The minimum sum of weights in a node in order to be considered for splitting.
maxdepth	maximum depth of the tree. The default <code>maxdepth = Inf</code> means that no restrictions are applied to tree sizes.
teststat	A character specifying the type of the test statistic to be applied.
testtype	A character specifying how to compute the distribution of the test statistic.
mincriterion	The value of the test statistic (for <code>testtype == "Teststatistic"</code>), or 1 - p-value (for other values of <code>testtype</code>) that must be exceeded in order to implement a split.
...	Other options to pass to <code>partykit::ctree()</code> or <code>partykit::cforest()</code> .
mtry	Number of input variables randomly sampled as candidates at each node for random forest like algorithms. The default <code>mtry = Inf</code> means that no random selection takes place.
ntree	Number of trees to grow in a forest.

Value

An object of class `party` (for `ctree`) or `cforest`.

Examples

```
if (rlang::is_installed(c("modeldata", "partykit"))) {
  data(bivariate, package = "modeldata")
  ctree_train(Class ~ ., data = bivariate_train)
  ctree_train(Class ~ ., data = bivariate_train, maxdepth = 1)
}
```

cubist_rules

Cubist rule-based regression models

Description

`cubist_rules()` defines a model that derives simple feature rules from a tree ensemble and creates regression models within each rule. This function can fit regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [Cubist](#)¹²

¹ The default engine. ² Requires a `parsnip` extension package.

More information on how **`parsnip`** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```

cubist_rules(
  mode = "regression",
  committees = NULL,
  neighbors = NULL,
  max_rules = NULL,
  engine = "Cubist"
)

```

Arguments

mode	A single character string for the type of model. The only possible value for this model is "regression".
committees	A non-negative integer (no greater than 100) for the number of members of the ensemble.
neighbors	An integer between zero and nine for the number of training set instances that are used to adjust the model-based prediction.
max_rules	The largest number of rules.
engine	A single character string specifying what computational engine to use for fitting.

Details

Cubist is a rule-based ensemble regression model. A basic model tree (Quinlan, 1992) is created that has a separate linear regression model corresponding for each terminal node. The paths along the model tree are flattened into rules and these rules are simplified and pruned. The parameter `min_n` is the primary method for controlling the size of each tree while `max_rules` controls the number of rules.

Cubist ensembles are created using *committees*, which are similar to boosting. After the first model in the committee is created, the second model uses a modified version of the outcome data based on whether the previous model under- or over-predicted the outcome. For iteration m , the new outcome y^* is computed using

$$y_{(m)}^* = y - (\hat{y}_{(m-1)} - y)$$

If a sample is under-predicted on the previous iteration, the outcome is adjusted so that the next time it is more likely to be over-predicted to compensate. This adjustment continues for each ensemble iteration. See Kuhn and Johnson (2013) for details.

After the model is created, there is also an option for a post-hoc adjustment that uses the training set (Quinlan, 1993). When a new sample is predicted by the model, it can be modified by its nearest neighbors in the original training set. For K neighbors, the model-based predicted value is adjusted by the neighbor using:

$$\frac{1}{K} \sum_{\ell=1}^K w_{\ell} [t_{\ell} + (\hat{y} - \hat{t}_{\ell})]$$

where $\hat{t}$ is the training set prediction and w is a weight that is inverse to the distance to the neighbor.

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See `set_engine()` for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the `fit()` function is used with the data.

Each of the arguments in this function other than `mode` and `engine` are captured as `quosures`. To pass values programmatically, use the `injection operator` like so:

```
value <- 1
cubist_rules(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

Quinlan R (1992). "Learning with Continuous Classes." Proceedings of the 5th Australian Joint Conference On Artificial Intelligence, pp. 343-348.

Quinlan R (1993). "Combining Instance-Based and Model-Based Learning." Proceedings of the Tenth International Conference on Machine Learning, pp. 236-243.

Kuhn M and Johnson K (2013). *Applied Predictive Modeling*. Springer.

See Also

`Cubist::cubist()`, `Cubist::cubistControl()`, `fit()`, `set_engine()`, `update()`, [Cubist engine details](#)

decision_tree

Decision trees

Description

`decision_tree()` defines a model as a set of if/then statements that creates a tree-based structure. This function can fit classification, regression, and censored regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [rpart](#)^{1,2}
- [C5.0](#)

- [partykit](#)²
- [rpartScore](#)²
- [spark](#)

¹ The default engine. ² Requires a parsnip extension package for censored regression, classification, and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
decision_tree(
  mode = "unknown",
  engine = "rpart",
  cost_complexity = NULL,
  tree_depth = NULL,
  min_n = NULL
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", "classification", or "censored regression".
engine	A single character string specifying what computational engine to use for fitting.
cost_complexity	A positive number for the the cost/complexity parameter (a.k.a. Cp) used by CART models (specific engines only).
tree_depth	An integer for maximum depth of the tree.
min_n	An integer for the minimum number of data points in a node that are required for the node to be split further.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
decision_tree(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [rpart engine details](#), [C5.0 engine details](#), [partykit engine details](#), [rpartScore engine details](#), [spark engine details](#)

Examples

```
show_engines("decision_tree")

decision_tree(mode = "classification", tree_depth = 5)
```

descriptors

Data Set Characteristics Available when Fitting Models

Description

When using the `fit()` functions there are some variables that will be available for use in arguments. For example, if the user would like to choose an argument value based on the current number of rows in a data set, the `.obs()` function can be used. See Details below.

Usage

```
.cols()

.preds()

.obs()

.lvls()

.facts()

.x()

.y()

.dat()
```

Details

Existing functions:

- `.obs()`: The current number of rows in the data set.
- `.preds()`: The number of columns in the data set that is associated with the predictors prior to dummy variable creation.
- `.cols()`: The number of predictor columns available after dummy variables are created (if any).

- `.facts()`: The number of factor predictors in the data set.
- `.lvls()`: If the outcome is a factor, this is a table with the counts for each level (and NA otherwise).
- `.x()`: The predictors returned in the format given. Either a data frame or a matrix.
- `.y()`: The known outcomes returned in the format given. Either a vector, matrix, or data frame.
- `.dat()`: A data frame containing all of the predictors and the outcomes. If `fit_xy()` was used, the outcomes are attached as the column, `..y`.

For example, if you use the model formula `circumference ~ .` with the built-in Orange data, the values would be

```
.preds() = 2          (the 2 remaining columns in `Orange`)
.cols()  = 5          (1 numeric column + 4 from Tree dummy variables)
.obs()   = 35
.lvls()  = NA         (no factor outcome)
.facts() = 1          (the Tree predictor)
.y()     = <vector>    (circumference as a vector)
.x()     = <data.frame> (The other 2 columns as a data frame)
.dat()   = <data.frame> (The full data set)
```

If the formula `Tree ~ .` were used:

```
.preds() = 2          (the 2 numeric columns in `Orange`)
.cols()  = 2          (same)
.obs()   = 35
.lvls()  = c("1" = 7, "2" = 7, "3" = 7, "4" = 7, "5" = 7)
.facts() = 0
.y()     = <vector>    (Tree as a vector)
.x()     = <data.frame> (The other 2 columns as a data frame)
.dat()   = <data.frame> (The full data set)
```

To use these in a model fit, pass them to a model specification. The evaluation is delayed until the time when the model is run via `fit()` (and the variables listed above are available). For example:

```
library(modeldata)
data("lending_club")

rand_forest(mode = "classification", mtry = .cols() - 2)
```

When no descriptors are found, the computation of the descriptor values is not executed.

discrim_flexible *Flexible discriminant analysis*

Description

discrim_flexible() defines a model that fits a discriminant analysis model that can use nonlinear features created using multivariate adaptive regression splines (MARS). This function can fit classification models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [earth](#)¹²

¹ The default engine. ² Requires a parsnip extension package.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
discrim_flexible(
  mode = "classification",
  num_terms = NULL,
  prod_degree = NULL,
  prune_method = NULL,
  engine = "earth"
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
num_terms	The number of features that will be retained in the final model, including the intercept.
prod_degree	The highest possible interaction degree.
prune_method	The pruning method.
engine	A single character string specifying what computational engine to use for fitting.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
discrim_flexible(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [earth engine details](#)

discrim_linear	<i>Linear discriminant analysis</i>
----------------	-------------------------------------

Description

`discrim_linear()` defines a model that estimates a multivariate distribution for the predictors separately for the data in each class (usually Gaussian with a common covariance matrix). Bayes' theorem is used to compute the probability of each class, given the predictor values. This function can fit classification models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [MASS](#)^{1,2}
- [mda](#)²
- [sda](#)²
- [sparsediscrim](#)²

¹ The default engine. ² Requires a parsnip extension package.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
discrim_linear(
  mode = "classification",
  penalty = NULL,
  regularization_method = NULL,
  engine = "MASS"
)
```

Arguments

mode	A single character string for the type of model. The only possible value for this model is "classification".
penalty	An non-negative number representing the amount of regularization used by some of the engines.
regularization_method	A character string for the type of regularized estimation. Possible values are: "diagonal", "min_distance", "shrink_cov", and "shrink_mean" (sparsediscrim engine only).
engine	A single character string specifying what computational engine to use for fitting.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
discrim_linear(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [MASS engine details](#), [mda engine details](#), [sda engine details](#), [sparsediscrim engine details](#)

discrim_quad	<i>Quadratic discriminant analysis</i>
--------------	--

Description

`discrim_quad()` defines a model that estimates a multivariate distribution for the predictors separately for the data in each class (usually Gaussian with separate covariance matrices). Bayes' theorem is used to compute the probability of each class, given the predictor values. This function can fit classification models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [MASS](#)^{1 2}
- [sparsediscrim](#)²

¹ The default engine. ² Requires a parsnip extension package.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
discrim_quad(
  mode = "classification",
  regularization_method = NULL,
  engine = "MASS"
)
```

Arguments

mode	A single character string for the type of model. The only possible value for this model is "classification".
regularization_method	A character string for the type of regularized estimation. Possible values are: "diagonal", "shrink_cov", and "shrink_mean" (sparsediscrim engine only).
engine	A single character string specifying what computational engine to use for fitting.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See `set_engine()` for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the `fit()` function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
discrim_quad(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [MASS engine details](#), [sparsediscrim engine details](#)

discrim_regularized *Regularized discriminant analysis*

Description

`discrim_regularized()` defines a model that estimates a multivariate distribution for the predictors separately for the data in each class. The structure of the model can be LDA, QDA, or some amalgam of the two. Bayes' theorem is used to compute the probability of each class, given the predictor values. This function can fit classification models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [klaR](#)^{1,2}

¹ The default engine. ² Requires a parsnip extension package.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
discrim_regularized(
  mode = "classification",
  frac_common_cov = NULL,
  frac_identity = NULL,
  engine = "klaR"
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
frac_common_cov, frac_identity	Numeric values between zero and one.
engine	A single character string specifying what computational engine to use for fitting.

Details

There are many ways of regularizing models. For example, one form of regularization is to penalize model parameters. Similarly, the classic James–Stein regularization approach shrinks the model structure to a less complex form.

The model fits a very specific type of regularized model by Friedman (1989) that uses two types of regularization. One modulates how class-specific the covariance matrix should be. This allows the model to balance between LDA and QDA. The second regularization component shrinks the covariance matrix towards the identity matrix.

For the penalization approach, `discrim_linear()` with a mda engine can be used. Other regularization methods can be used with `discrim_linear()` and `discrim_quad()` can be used via the `sparseDiscrim` engine for those functions.

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See `set_engine()` for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the `fit()` function is used with the data.

Each of the arguments in this function other than mode and engine are captured as `quosures`. To pass values programmatically, use the `injection operator` like so:

```
value <- 1
discrim_regularized(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

Friedman, J (1989). Regularized Discriminant Analysis. *Journal of the American Statistical Association*, 84, 165-175.

See Also

`fit()`, `set_engine()`, `update()`, `klaR` engine details

extract-parsnip	<i>Extract elements of a parsnip model object</i>
-----------------	---

Description

These functions extract various elements from a parsnip object. If they do not exist yet, an error is thrown.

- `extract_spec_parsnip()` returns the parsnip [model specification](#).
- `extract_fit_engine()` returns the engine specific fit embedded within a parsnip model fit. For example, when using [linear_reg\(\)](#) with the "lm" engine, this returns the underlying lm object.
- `extract_parameter_dials()` returns a single dials parameter object.
- `extract_parameter_set_dials()` returns a set of dials parameter objects.
- `extract_fit_time()` returns a tibble with fit times. The fit times correspond to the time for the parsnip engine to fit and do not include other portions of the elapsed time in [fit.model_spec\(\)](#).

Usage

```
## S3 method for class 'model_fit'
extract_spec_parsnip(x, ...)

## S3 method for class 'model_fit'
extract_fit_engine(x, ...)

## S3 method for class 'model_spec'
extract_parameter_set_dials(x, ...)

## S3 method for class 'model_spec'
extract_parameter_dials(x, parameter, ...)

## S3 method for class 'model_fit'
extract_fit_time(x, summarize = TRUE, ...)
```

Arguments

<code>x</code>	A parsnip <code>model_fit</code> object or a parsnip <code>model_spec</code> object.
<code>...</code>	Not currently used.
<code>parameter</code>	A single string for the parameter ID.
<code>summarize</code>	A logical for whether the elapsed fit time should be returned as a single row or multiple rows. Doesn't support FALSE for parsnip models.

Details

Extracting the underlying engine fit can be helpful for describing the model (via `print()`, `summary()`, `plot()`, etc.) or for variable importance/explainers.

However, users should not invoke the `predict()` method on an extracted model. There may be preprocessing operations that `parsnip` has executed on the data prior to giving it to the model. Bypassing these can lead to errors or silently generating incorrect predictions.

Good:

```
parsnip_fit |> predict(new_data)
```

Bad:

```
parsnip_fit |> extract_fit_engine() |> predict(new_data)
```

Value

The extracted value from the `parsnip` object, `x`, as described in the description section.

Examples

```
lm_spec <- linear_reg() |> set_engine("lm")
lm_fit <- fit(lm_spec, mpg ~ ., data = mtcars)

lm_spec
extract_spec_parsnip(lm_fit)

extract_fit_engine(lm_fit)
lm(mpg ~ ., data = mtcars)
```

fit.model_spec

Fit a Model Specification to a Dataset

Description

`fit()` and `fit_xy()` take a model specification, translate the required code by substituting arguments, and execute the model fit routine.

Usage

```
## S3 method for class 'model_spec'
fit(
  object,
  formula,
  data,
  case_weights = NULL,
  control = control_parsnip(),
```

```

    ...
  )

  ## S3 method for class 'model_spec'
  fit_xy(object, x, y, case_weights = NULL, control = control_parsnip(), ...)

```

Arguments

object	An object of class <code>model_spec</code> that has a chosen engine (via set_engine()).
formula	An object of class <code>formula</code> (or one that can be coerced to that class): a symbolic description of the model to be fitted.
data	Optional, depending on the interface (see Details below). A data frame containing all relevant variables (e.g. outcome(s), predictors, case weights, etc). Note: when needed, a <i>named argument</i> should be used.
case_weights	An optional classed vector of numeric case weights. This must return TRUE when hardhat::is_case_weights() is run on it. See hardhat::frequency_weights() and hardhat::importance_weights() for examples.
control	A named list with elements <code>verbosity</code> and <code>catch</code> . See control_parsnip() .
...	Must be empty; an error is raised if any arguments are passed here. Options required to fit the model should be passed to set_engine() , and case weights to the <code>case_weights</code> argument.
x	A matrix, sparse matrix, or data frame of predictors. Only some models have support for sparse matrix input. See parsnip::get_encoding() for details. <code>x</code> should have column names.
y	A vector, matrix or data frame of outcome data.

Details

`fit()` and `fit_xy()` substitute the current arguments in the model specification into the computational engine's code, check them for validity, then fit the model using the data and the engine-specific code. Different model functions have different interfaces (e.g. `formula` or `x/y`) and these functions translate between the interface used when `fit()` or `fit_xy()` was invoked and the one required by the underlying model.

When possible, these functions attempt to avoid making copies of the data. For example, if the underlying model uses a `formula` and `fit()` is invoked, the original data are references when the model is fit. However, if the underlying model uses something else, such as `x/y`, the `formula` is evaluated and the data are converted to the required format. In this case, any calls in the resulting model objects reference the temporary objects used to fit the model.

If the model engine has not been set, the model's default engine will be used (as discussed on each model page). If the `verbosity` option of [control_parsnip\(\)](#) is greater than zero, a warning will be produced.

If you would like to use an alternative method for generating contrasts when supplying a `formula` to `fit()`, set the global option `contrasts` to your preferred method. For example, you might set it to: `options(contrasts = c(ordered = "contr.helmert", ordered = "contr.poly"))`. See the help page for [stats::contr.treatment\(\)](#) for more possible contrast types.

For models with "censored regression" modes, an additional computation is executed and saved in the `parsnip` object. The `ensor_probs` element contains a "reverse Kaplan-Meier" curve that models the probability of censoring. This may be used later to compute inverse probability censoring weights for performance measures.

Sparse data is supported, with the use of the `x` argument in `fit_xy()`. See `allow_sparse_x` column of `get_encoding()` for sparse input compatibility.

Value

A `model_fit` object that contains several elements:

- `lvl`: If the outcome is a factor, this contains the factor levels at the time of model fitting.
- `ordered`: If the outcome is a factor, was it an ordered factor?
- `spec`: The model specification object (object in the call to `fit`)
- `fit`: when the model is executed without error, this is the model object. Otherwise, it is a try-error object with the error message.
- `preproc`: any objects needed to convert between a formula and non-formula interface (such as the terms object)

The return value will also have a class related to the fitted model (e.g. `"_glm"`) before the base class of `"model_fit"`.

See Also

`set_engine()`, `control_parsnip()`, `model_spec`, `model_fit`

Examples

```
# Although `glm()` only has a formula interface, different
# methods for specifying the model can be used

library(dplyr)
library(modeldata)
data("lending_club")

lr_mod <- logistic_reg()

using_formula <-
  lr_mod |>
  set_engine("glm") |>
  fit(Class ~ funded_amnt + int_rate, data = lending_club)

using_xy <-
  lr_mod |>
  set_engine("glm") |>
  fit_xy(x = lending_club[, c("funded_amnt", "int_rate")],
        y = lending_club$Class)

using_formula
using_xy
```

gen_additive_mod	<i>Generalized additive models (GAMs)</i>
------------------	---

Description

gen_additive_mod() defines a model that can use smoothed functions of numeric predictors in a generalized linear model. This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [mgcv](#)¹
- [vgam](#)²

¹ The default engine. ² Requires a parsnip extension package for classification.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
gen_additive_mod(
  mode = "unknown",
  select_features = NULL,
  adjust_deg_free = NULL,
  engine = "mgcv"
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
select_features	TRUE or FALSE. If TRUE, the model has the ability to eliminate a predictor (via penalization). Increasing adjust_deg_free will increase the likelihood of removing predictors.
adjust_deg_free	If select_features = TRUE, then acts as a multiplier for smoothness. Increase this beyond 1 to produce smoother models.
engine	A single character string specifying what computational engine to use for fitting.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
gen_additive_mod(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [mgcv engine details](#), [vgam engine details](#)

Examples

```
show_engines("gen_additive_mod")

gen_additive_mod()
```

glance.model_fit	<i>Construct a single row summary "glance" of a model, fit, or other object</i>
------------------	---

Description

This method glances the model in a parsnip model object, if it exists.

Usage

```
## S3 method for class 'model_fit'
glance(x, ...)
```

Arguments

x	model or other R object to convert to single-row data frame
...	other arguments passed to methods

Value

a tibble

glm_grouped

*Fit a grouped binomial outcome from a data set with case weights***Description**

`stats::glm()` assumes that a tabular data set with case weights corresponds to "different observations have different dispersions" (see `?glm`).

In some cases, the case weights reflect that the same covariate pattern was observed multiple times (i.e., *frequency weights*). In this case, `stats::glm()` expects the data to be formatted as the number of events for each factor level so that the outcome can be given to the formula as `cbind(events_1, events_2)`.

`glm_grouped()` converts data with integer case weights to the expected "number of events" format for binomial data.

Usage

```
glm_grouped(formula, data, weights, ...)
```

Arguments

formula	A formula object with one outcome that is a two-level factors.
data	A data frame with the outcomes and predictors (but not case weights).
weights	An integer vector of weights whose length is the same as the number of rows in data. If it is a non-integer numeric, it will be converted to integer (with a warning).
...	Options to pass to <code>stats::glm()</code> . If family is not set, it will automatically be assigned the basic binomial family.

Value

A object produced by `stats::glm()`.

Examples

```
#-----
# The same data set formatted three ways

# First with basic case weights that, from ?glm, are used inappropriately.
ucb_weighted <- as.data.frame(UCBAdmissions)
ucb_weighted$Freq <- as.integer(ucb_weighted$Freq)
head(ucb_weighted)
nrow(ucb_weighted)

# Format when yes/no data are in individual rows (probably still inappropriate)
library(tidyr)
ucb_long <- uncount(ucb_weighted, Freq)
head(ucb_long)
```

```

nrow(ucb_long)

# Format where the outcome is formatted as number of events
ucb_events <-
  ucb_weighted |>
  tidyr::pivot_wider(
    id_cols = c(Gender, Dept),
    names_from = Admit,
    values_from = Freq,
    values_fill = 0L
  )
head(ucb_events)
nrow(ucb_events)

#-----
# Different model fits

# Treat data as separate Bernoulli data:
glm(Admit ~ Gender + Dept, data = ucb_long, family = binomial)

# Weights produce the same statistics
glm(
  Admit ~ Gender + Dept,
  data = ucb_weighted,
  family = binomial,
  weights = ucb_weighted$Freq
)

# Data as binomial "x events out of n trials" format. Note that, to get the same
# coefficients, the order of the levels must be reversed.
glm(
  cbind(Rejected, Admitted) ~ Gender + Dept,
  data = ucb_events,
  family = binomial
)

# The new function that starts with frequency weights and gets the correct place:
glm_grouped(Admit ~ Gender + Dept, data = ucb_weighted, weights = ucb_weighted$Freq)

```

linear_reg

Linear regression

Description

`linear_reg()` defines a model that can predict numeric values from predictors using a linear function. This function can fit regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- `lm`¹
- `brulee`
- `gee`²
- `glm`
- `glmer`²
- `glmnet`
- `gls`²
- `h2o`²
- `keras`
- `keras3`
- `lme`²
- `lmer`²
- `quantreg`
- `spark`
- `stan`
- `stan_glmer`²

¹ The default engine. ² Requires a `parsnip` extension package for regression.

More information on how **`parsnip`** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
linear_reg(mode = "regression", engine = "lm", penalty = NULL, mixture = NULL)
```

Arguments

<code>mode</code>	A single character string for the type of model. The only possible value for this model is "regression".
<code>engine</code>	A single character string specifying what computational engine to use for fitting. Possible engines are listed below. The default for this model is "lm".
<code>penalty</code>	A non-negative number representing the total amount of regularization (specific engines only).
<code>mixture</code>	A number between zero and one (inclusive) denoting the proportion of L1 regularization (i.e. lasso) in the model. • <code>mixture = 1</code> specifies a pure lasso model, • <code>mixture = 0</code> specifies a ridge regression model, and • <code>0 < mixture < 1</code> specifies an elastic net model, interpolating lasso and ridge.

Available for specific engines only.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See `set_engine()` for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the `fit()` function is used with the data.

Each of the arguments in this function other than mode and engine are captured as `quosures`. To pass values programmatically, use the `injection operator` like so:

```
value <- 1
linear_reg(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

`fit()`, `set_engine()`, `update()`, `lm` engine details, `brulee` engine details, `gee` engine details, `glm` engine details, `glmer` engine details, `glmnet` engine details, `gl`s engine details, `h2o` engine details, `keras` engine details, `keras3` engine details, `lme` engine details, `lmer` engine details, `quantreg` engine details, `spark` engine details, `stan` engine details, `stan_glmer` engine details

Examples

```
show_engines("linear_reg")

linear_reg()
```

logistic_reg

Logistic regression

Description

`logistic_reg()` defines a generalized linear model for binary outcomes. A linear combination of the predictors is used to model the log odds of an event. This function can fit classification models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- `glm`¹
- `brulee`
- `gee`²
- `glmer`²
- `glmnet`

- `h2o`²
- `keras`
- `keras3`
- `LiblineaR`
- `spark`
- `stan`
- `stan_glmer`²

¹ The default engine. ² Requires a parsnip extension package.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
logistic_reg(
  mode = "classification",
  engine = "glm",
  penalty = NULL,
  mixture = NULL
)
```

Arguments

mode	A single character string for the type of model. The only possible value for this model is "classification".
engine	A single character string specifying what computational engine to use for fitting. Possible engines are listed below. The default for this model is "glm".
penalty	A non-negative number representing the total amount of regularization (specific engines only). For keras models, this corresponds to purely L2 regularization (aka weight decay) while the other models can be either or a combination of L1 and L2 (depending on the value of mixture).
mixture	A number between zero and one (inclusive) giving the proportion of L1 regularization (i.e. lasso) in the model. • <code>mixture = 1</code> specifies a pure lasso model, • <code>mixture = 0</code> specifies a ridge regression model, and • <code>0 < mixture < 1</code> specifies an elastic net model, interpolating lasso and ridge. Available for specific engines only. For LiblineaR models, mixture must be exactly 1 or 0 only.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See `set_engine()` for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the `fit()` function is used with the data.

Each of the arguments in this function other than `mode` and `engine` are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
logistic_reg(argument = !!value)
```

This model fits a classification model for binary outcomes; for multiclass outcomes, see [multinom_reg\(\)](#).

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [glm engine details](#), [brulee engine details](#), [gee engine details](#), [glmer engine details](#), [glmnet engine details](#), [h2o engine details](#), [keras engine details](#), [keras3 engine details](#), [LiblineaR engine details](#), [spark engine details](#), [stan engine details](#), [stan_glmer engine details](#)

Examples

```
show_engines("logistic_reg")

logistic_reg()
```

`mars`

Multivariate adaptive regression splines (MARS)

Description

`mars()` defines a generalized linear model that uses artificial features for some predictors. These features resemble hinge functions and the result is a model that is a segmented regression in small dimensions. This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [earth](#)¹

¹ The default engine.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
mars(  
  mode = "unknown",  
  engine = "earth",  
  num_terms = NULL,  
  prod_degree = NULL,  
  prune_method = NULL  
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
engine	A single character string specifying what computational engine to use for fitting.
num_terms	The number of features that will be retained in the final model, including the intercept.
prod_degree	The highest possible interaction degree.
prune_method	The pruning method.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1  
mars(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [earth engine details](#)

Examples

```
show_engines("mars")  
  
mars(mode = "regression", num_terms = 5)
```

max_mtry_formula	<i>Determine largest value of mtry from formula. This function potentially caps the value of mtry based on a formula and data set. This is a safe approach for survival and/or multivariate models.</i>
------------------	---

Description

Determine largest value of mtry from formula. This function potentially caps the value of mtry based on a formula and data set. This is a safe approach for survival and/or multivariate models.

Usage

```
max_mtry_formula(mtry, formula, data)
```

Arguments

mtry	An initial value of mtry (which may be too large).
formula	A model formula.
data	The training set (data frame).

Value

A value for mtry.

Examples

```
# should be 9
max_mtry_formula(200, cbind(wt, mpg) ~ ., data = mtcars)
```

maybe_matrix	<i>Fuzzy conversions</i>
--------------	--------------------------

Description

These are substitutes for `as.matrix()` and `as.data.frame()` that leave a sparse matrix as-is.

Usage

```
maybe_matrix(x)

maybe_data_frame(x)
```

Arguments

x	A data frame, matrix, or sparse matrix.
---	---

Value

A data frame, matrix, or sparse matrix.

min_cols

Execution-time data dimension checks

Description

For some tuning parameters, the range of values depend on the data dimensions (e.g. mtry). Some packages will fail if the parameter values are outside of these ranges. Since the model might receive resampled versions of the data, these ranges can't be set prior to the point where the model is fit. These functions check the possible range of the data and adjust them if needed (with a warning).

Usage

```
min_cols(num_cols, source)

min_rows(num_rows, source, offset = 0)
```

Arguments

num_cols, num_rows	The parameter value requested by the user.
source	A data frame for the data to be used in the fit. If the source is named "data", it is assumed that one column of the data corresponds to an outcome (and is subtracted off).
offset	A number subtracted off of the number of rows available in the data.

Value

An integer (and perhaps a warning).

Examples

```
nearest_neighbor(neighbors= 100) |>
  set_engine("kkn") |>
  set_mode("regression") |>
  translate()

library(ranger)
rand_forest(mtry = 2, min_n = 100, trees = 3) |>
  set_engine("ranger") |>
  set_mode("regression") |>
  fit(mpg ~ ., data = mtcars)
```

mlp

*Single layer neural network***Description**

`mlp()` defines a multilayer perceptron model (a.k.a. a single layer, feed-forward neural network). This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- `nnet`¹
- `brulee`
- `brulee_two_layer`
- `h2o`²
- `keras`
- `keras3`
- `qrnn`

¹ The default engine. ² Requires a `parsnip` extension package for classification and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
mlp(
  mode = "unknown",
  engine = "nnet",
  hidden_units = NULL,
  penalty = NULL,
  dropout = NULL,
  epochs = NULL,
  activation = NULL,
  learn_rate = NULL
)
```

Arguments

<code>mode</code>	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
<code>engine</code>	A single character string specifying what computational engine to use for fitting.
<code>hidden_units</code>	An integer for the number of units in the hidden model.
<code>penalty</code>	A non-negative numeric value for the amount of weight decay.
<code>dropout</code>	A number between 0 (inclusive) and 1 denoting the proportion of model parameters randomly set to zero during model training.

epochs	An integer for the number of training iterations.
activation	A character vector denoting the type of relationship between the different layers. The activation function between the hidden and output layers is automatically set to either "linear" or "softmax" depending on the type of outcome. Possible values depend on the engine being used.
learn_rate	A number for the rate at which the boosting algorithm adapts from iteration-to-iteration (specific engines only). This is sometimes referred to as the shrinkage parameter.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
mlp(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [nnet engine details](#), [brulee engine details](#), [brulee_two_layer engine details](#), [h2o engine details](#), [keras engine details](#), [keras3 engine details](#), [qrnn engine details](#)

Examples

```
show_engines("mlp")

mlp(mode = "classification", penalty = 0.01)
```

model_fit

Model Fit Objects

Description

Model fits are trained [model specifications](#) that are ready to [predict](#) on new data. Model fits have class `model_fit` and, usually, a subclass referring to the engine used to fit the model.

Details

An object with class "model_fit" is a container for information about a model that has been fit to the data.

The main elements of the object are:

- `lvl`: A vector of factor levels when the outcome is a factor. This is `NULL` when the outcome is not a factor vector.
- `spec`: A `model_spec` object.
- `fit`: The object produced by the fitting function.
- `preproc`: This contains any data-specific information required to process new a sample point for prediction. For example, if the underlying model function requires arguments `x` and `y` and the user passed a formula to `fit`, the `preproc` object would contain items such as the terms object and so on. When no information is required, this is `NA`.

As discussed in the documentation for [model_spec](#), the original arguments to the specification are saved as quosures. These are evaluated for the `model_fit` object prior to fitting. If the resulting model object prints its call, any user-defined options are shown in the call preceded by a tilde (see the example below). This is a result of the use of quosures in the specification.

This class and structure is the basis for how **parsnip** stores model objects after seeing the data and applying a model.

Examples

```
# Keep the `x` matrix if the data are not too big.
spec_obj <-
  linear_reg() |>
  set_engine("lm", x = .obs() < 500)
spec_obj

fit_obj <- fit(spec_obj, mpg ~ ., data = mtcars)
fit_obj

nrow(fit_obj$fit$x)
```

model_formula

Formulas with special terms in tidymodels

Description

In R, formulas provide a compact, symbolic notation to specify model terms. Many modeling functions in R make use of "specials", or nonstandard notations used in formulas. Specials are defined and handled as a special case by a given modeling package. For example, the `mgcv` package, which provides support for [generalized additive models](#) in R, defines a function `s()` to be in-lined into formulas. It can be used like so:

```
mgcv::gam(mpg ~ wt + s(displ, k = 5), data = mtcars)
```

In this example, the `s()` special defines a smoothing term that the `mgcv` package knows to look for when preprocessing model input.

The `parsnip` package can handle most specials without issue. The analogous code for specifying this generalized additive model [with the `parsnip` "mgcv" engine](#) looks like:

```
gen_additive_mod() |>
  set_mode("regression") |>
  set_engine("mgcv") |>
  fit(mpg ~ wt + s(displ, k = 5), data = mtcars)
```

However, `parsnip` is often used in conjunction with the greater `tidymodels` package ecosystem, which defines its own pre-processing infrastructure and functionality via packages like `hardhat` and `recipes`. The specials defined in many modeling packages introduce conflicts with that infrastructure.

To support specials while also maintaining consistent syntax elsewhere in the ecosystem, **`tidymodels` delineates between two types of formulas: preprocessing formulas and model formulas**. Preprocessing formulas specify the input variables, while model formulas determine the model structure.

Example

To create the preprocessing formula from the model formula, just remove the specials, retaining references to input variables themselves. For example:

```
model_formula <- mpg ~ wt + s(displ, k = 5)
preproc_formula <- mpg ~ wt + displ
```

- **With `parsnip`**, use the model formula:

```
model_spec <-
  gen_additive_mod() |>
  set_mode("regression") |>
  set_engine("mgcv")

model_spec |>
  fit(model_formula, data = mtcars)
```

- **With `recipes`**, use the preprocessing formula only:

```
library(recipes)

recipe(preproc_formula, mtcars)
```

The `recipes` package supplies a large variety of preprocessing techniques that may replace the need for specials altogether, in some cases.

- **With `workflows`**, use the preprocessing formula everywhere, but pass the model formula to the `formula` argument in `add_model()`:

```
library(workflows)

wflow <-
  workflow() |>
  add_formula(preproc_formula) |>
  add_model(model_spec, formula = model_formula)

fit(wflow, data = mtcars)
```

The workflow will then pass the model formula to `parsnip`, using the preprocessor formula elsewhere. We would still use the preprocessing formula if we had added a recipe preprocessor using `add_recipe()` instead a formula via `add_formula()`.

model_spec

Model Specifications

Description

The `parsnip` package splits the process of fitting models into two steps:

1. Specify how a model will be fit using a *model specification*
2. Fit a model using the model specification

This is a different approach to many other model interfaces in R, like `lm()`, where both the specification of the model and the fitting happens in one function call. Splitting the process into two steps allows users to iteratively define model specifications throughout the model development process.

This intermediate object that defines how the model will be fit is called a *model specification* and has class `model_spec`. Model type functions, like `linear_reg()` or `boost_tree()`, return `model_spec` objects.

Fitted model objects, resulting from passing a `model_spec` to `fit()` or `fit_xy()`, have class `model_fit`, and contain the original `model_spec` objects inside them. See `?model_fit` for more on that object type, and `?extract_spec_parsnip` to extract `model_specs` from `model_fits`.

Details

An object with class `"model_spec"` is a container for information about a model that will be fit.

The main elements of the object are:

- `args`: A vector of the main arguments for the model. The names of these arguments may be different from their counterparts in the underlying model function. For example, for a `glmnet` model, the argument name for the amount of the penalty is called "penalty" instead of "lambda" to make it more general and usable across different types of models (and to not be specific to a particular model function). The elements of `args` can be tuned with the use of the [tune package](https://www.tidymodels.org/start/tuning/). For more information see <https://www.tidymodels.org/start/tuning/>. If left to their defaults (NULL), the arguments will use the underlying model functions default value. As discussed below, the arguments in `args` are captured as quosures and are not immediately executed.

- ...: Optional model-function-specific parameters. As with args, these will be quosures and can be tune().
- mode: The type of model, such as "regression" or "classification". Other modes will be added once the package adds more functionality.
- method: This is a slot that is filled in later by the model's constructor function. It generally contains lists of information that are used to create the fit and prediction code as well as required packages and similar data.
- engine: This character string declares exactly what software will be used. It can be a package name or a technology type.

This class and structure is the basis for how parsnip stores model objects prior to seeing the data.

Argument Details

An important detail to understand when creating model specifications is that they are intended to be functionally independent of the data. While it is true that some tuning parameters are *data dependent*, the model specification does not interact with the data at all.

For example, most R functions immediately evaluate their arguments. For example, when calling `mean(dat_vec)`, the object `dat_vec` is immediately evaluated inside of the function.

parsnip model functions do not do this. For example, using

```
rand_forest(mtry = ncol(mtcars) - 1)
```

does not execute `ncol(mtcars) - 1` when creating the specification. This can be seen in the output:

```
> rand_forest(mtry = ncol(mtcars) - 1)
Random Forest Model Specification (unknown)
```

Main Arguments:

```
mtry = ncol(mtcars) - 1
```

The model functions save the argument *expressions* and their associated environments (a.k.a. a quosure) to be evaluated later when either `fit.model_spec()` or `fit_xy.model_spec()` are called with the actual data.

The consequence of this strategy is that any data required to get the parameter values must be available when the model is fit. The two main ways that this can fail is if:

1. The data have been modified between the creation of the model specification and when the model fit function is invoked.
2. If the model specification is saved and loaded into a new session where those same data objects do not exist.

The best way to avoid these issues is to not reference any data objects in the global environment but to use data descriptors such as `.cols()`. Another way of writing the previous specification is

```
rand_forest(mtry = .cols() - 1)
```

This is not dependent on any specific data object and is evaluated immediately before the model fitting process begins.

One less advantageous approach to solving this issue is to use quasiquotation. This would insert the actual R object into the model specification and might be the best idea when the data object is small. For example, using

```
rand_forest(mtry = ncol(!mtcars) - 1)
```

would work (and be reproducible between sessions) but embeds the entire mtcars data set into the mtry expression:

```
> rand_forest(mtry = ncol(!mtcars) - 1)
Random Forest Model Specification (unknown)
```

Main Arguments:

```
mtry = ncol(structure(list(Sepal.Length = c(5.1, 4.9, 4.7, 4.6, 5, <snip>
```

However, if there were an object with the number of columns in it, this wouldn't be too bad:

```
> mtry_val <- ncol(mtcars) - 1
> mtry_val
[1] 10
> rand_forest(mtry = !mtry_val)
Random Forest Model Specification (unknown)
```

Main Arguments:

```
mtry = 10
```

More information on quosures and quasiquotation can be found at <https://adv-r.hadley.nz/quasiquotation.html>.

multi_predict

Model predictions across many sub-models

Description

For some models, predictions can be made on sub-models in the model object.

Usage

```
multi_predict(object, ...)

## Default S3 method:
multi_predict(object, ...)

## S3 method for class `"_xgb.Booster"`
```

```

multi_predict(object, new_data, type = NULL, trees = NULL, ...)

## S3 method for class '`_C5.0`'
multi_predict(object, new_data, type = NULL, trees = NULL, ...)

## S3 method for class '`_elnet`'
multi_predict(object, new_data, type = NULL, penalty = NULL, ...)

## S3 method for class '`_lognet`'
multi_predict(object, new_data, type = NULL, penalty = NULL, ...)

## S3 method for class '`_multnet`'
multi_predict(object, new_data, type = NULL, penalty = NULL, ...)

## S3 method for class '`_glmnetfit`'
multi_predict(object, new_data, type = NULL, penalty = NULL, ...)

## S3 method for class '`_earth`'
multi_predict(object, new_data, type = NULL, num_terms = NULL, ...)

## S3 method for class '`_torch_mlp`'
multi_predict(object, new_data, type = NULL, epochs = NULL, ...)

## S3 method for class '`_train.kknn`'
multi_predict(object, new_data, type = NULL, neighbors = NULL, ...)

```

Arguments

object	A model fit .
...	Optional arguments to pass to <code>predict.model_fit(type = "raw")</code> such as <code>type</code> .
new_data	A rectangular data object, such as a data frame.
type	A single character value or NULL. Possible values are: • regression: "numeric" • classification: "class", "prob" • censored regression: "survival", "time", "hazard", "linear_pred" • quantile regression: "quantile" • interval estimates: "conf_int", "pred_int" • other: "raw" When NULL, <code>predict()</code> will choose an appropriate value based on the model's mode.
trees	An integer vector for the number of trees in the ensemble.
penalty	A numeric vector of penalty values.
num_terms	An integer vector for the number of MARS terms to retain.
epochs	An integer vector for the number of training epochs.
neighbors	An integer vector for the number of nearest neighbors.

Value

A tibble with the same number of rows as the data being predicted. There is a list-column named `.pred` that contains tibbles with multiple rows per sub-model. Note that, within the tibbles, the column names follow the usual standard based on prediction type (i.e. `.pred_class` for `type = "class"` and so on).

`type = "raw"` is the exception: it returns the engine's own prediction object with no parsnip post-processing, so it is not a tibble with a `.pred` list-column. For `glmnet` engines this is a matrix with one column per penalty value, or a three-dimensional array for `multinom_reg()`. Because `...` is passed to the engine but `type` is a named argument of `multi_predict()` itself, an engine-level prediction type cannot be supplied this way; the engine's default is used.

<code>multinom_reg</code>	<i>Multinomial regression</i>
---------------------------	-------------------------------

Description

`multinom_reg()` defines a model that uses linear predictors to predict multiclass data using the multinomial distribution. This function can fit classification models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [nnet](#)¹
- [brulee](#)
- [glmnet](#)
- [h2o](#)²
- [keras](#)
- [keras3](#)
- [spark](#)

¹ The default engine. ² Requires a parsnip extension package.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
multinom_reg(
  mode = "classification",
  engine = "nnet",
  penalty = NULL,
  mixture = NULL
)
```

Arguments

mode	A single character string for the type of model. The only possible value for this model is "classification".
engine	A single character string specifying what computational engine to use for fitting. Possible engines are listed below. The default for this model is "nnet".
penalty	A non-negative number representing the total amount of regularization (specific engines only). For keras models, this corresponds to purely L2 regularization (aka weight decay) while the other models can be a combination of L1 and L2 (depending on the value of mixture).
mixture	A number between zero and one (inclusive) giving the proportion of L1 regularization (i.e. lasso) in the model. • mixture = 1 specifies a pure lasso model, • mixture = 0 specifies a ridge regression model, and • $0 < \text{mixture} < 1$ specifies an elastic net model, interpolating lasso and ridge. Available for specific engines only.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
multinom_reg(argument = !!value)
```

This model fits a classification model for multiclass outcomes; for binary outcomes, see [logistic_reg\(\)](#).

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [nnet engine details](#), [brulee engine details](#), [glmnet engine details](#), [h2o engine details](#), [keras engine details](#), [keras3 engine details](#), [spark engine details](#)

Examples

```
show_engines("multinom_reg")

multinom_reg()
```

naive_Bayes

Naive Bayes models

Description

`naive_Bayes()` defines a model that uses Bayes' theorem to compute the probability of each class, given the predictor values. This function can fit classification models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [klaR](#)^{1,2}
- [h2o](#)²
- [naivebayes](#)²

¹ The default engine. ² Requires a `parsnip` extension package.

More information on how **`parsnip`** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
naive_Bayes(
  mode = "classification",
  smoothness = NULL,
  Laplace = NULL,
  engine = "klaR"
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
smoothness	An non-negative number representing the the relative smoothness of the class boundary. Smaller examples result in model flexible boundaries and larger values generate class boundaries that are less adaptable
Laplace	A non-negative value for the Laplace correction to smoothing low-frequency counts.
engine	A single character string specifying what computational engine to use for fitting.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than `mode` and `engine` are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
naive_Bayes(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [klaR engine details](#), [h2o engine details](#), [naivebayes engine details](#)

nearest_neighbor	<i>K-nearest neighbors</i>
------------------	----------------------------

Description

`nearest_neighbor()` defines a model that uses the K most similar data points from the training set to predict new samples. This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [kknn](#)¹

¹ The default engine.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
nearest_neighbor(
  mode = "unknown",
  engine = "kknn",
  neighbors = NULL,
  weight_func = NULL,
  dist_power = NULL
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
engine	A single character string specifying what computational engine to use for fitting.
neighbors	A single integer for the number of neighbors to consider (often called k). For kknn , a value of 5 is used if neighbors is not specified.
weight_func	A <i>single</i> character for the type of kernel function used to weight distances between samples. Valid choices are: "rectangular", "triangular", "epanechnikov", "biweight", "triweight", "cos", "inv", "gaussian", "rank", or "optimal".
dist_power	A single number for the parameter used in calculating Minkowski distance.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See `set_engine()` for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the `fit()` function is used with the data.

Each of the arguments in this function other than mode and engine are captured as `quosures`. To pass values programmatically, use the `injection operator` like so:

```
value <- 1
nearest_neighbor(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

`fit()`, `set_engine()`, `update()`, `kkn` engine details

Examples

```
show_engines("nearest_neighbor")

nearest_neighbor(neighbors = 11)
```

<code>null_model</code>	<i>Null model</i>
-------------------------	-------------------

Description

Fit a single mean, set of quantiles, or largest class model. `null_model()` is the user-facing function that relies on the underlying computational function, `nullmodel()`.

Usage

```
null_model(mode = "classification", engine = "parsnip")
```

Arguments

- | | |
|---------------------|--|
| <code>mode</code> | A single character string for the type of model. The only possible values for this model are "regression", "quantile regression", and "classification". |
| <code>engine</code> | A single character string specifying what computational engine to use for fitting. Possible engines are listed below. The default for this model is "parsnip". |

Details

`null_model()` defines a simple, non-informative model. It doesn't have any main arguments. This function can fit classification, regression, and quantile regression models.

`null_model()` emulates other model building functions, but returns the simplest model possible given a training set: a single mean or set of quantiles for numeric outcomes and the most prevalent class for factor outcomes. When class probabilities are requested, the percentage of the training set samples with the most prevalent class is returned.

Engine Details

Engines may have pre-set default arguments when executing the model fit call. For this type of model, the template of the fit calls are below:

parsnip:

```
null_model() |>
  set_engine("parsnip") |>
  set_mode("regression") |>
  translate()

## Null Model Specification (regression)
##
## Computational engine: parsnip
##
## Model fit template:
## parsnip::nullmodel(x = missing_arg(), y = missing_arg())

null_model() |>
  set_engine("parsnip") |>
  set_mode("quantile regression", quantile_levels = c(0.25, 0.5, 0.75)) |>
  translate()

## Null Model Specification (quantile regression)
##
## Computational engine: parsnip
##
## Model fit template:
## parsnip::nullmodel(x = missing_arg(), y = missing_arg(), quantile_levels = quantile_levels)

## Quantile levels: 0.25, 0.5, and 0.75.

null_model() |>
  set_engine("parsnip") |>
  set_mode("classification") |>
  translate()

## Null Model Specification (classification)
##
## Computational engine: parsnip
##
```

```
## Model fit template:
## parsnip::nullmodel(x = missing_arg(), y = missing_arg())
```

Prediction types:

```
parsnip::get_from_env("null_model_predict") |>
  dplyr::filter(engine == "parsnip") |>
  dplyr::select(mode, type)
```

```
## # A tibble: 6 x 2
##   mode                type
##   <chr>              <chr>
## 1 quantile regression quantile
## 2 regression         numeric
## 3 regression         raw
## 4 classification     class
## 5 classification     prob
## 6 classification     raw
```

See Also

[fit.model_spec\(\)](#)

Examples

```
null_model(mode = "regression")
```

parsnip_addin

Start an RStudio Addin that can write model specifications

Description

`parsnip_addin()` starts a process in the RStudio IDE Viewer window that allows users to write code for parsnip model specifications from various R packages. The new code is written to the current document at the location of the cursor.

Usage

```
parsnip_addin()
```

pls *Partial least squares (PLS)*

Description

`pls()` defines a partial least squares model that uses latent variables to model the data. It is similar to a supervised version of principal component. This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [mixOmics](#)¹²

¹ The default engine. ² Requires a `parsnip` extension package for classification and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
pls(
  mode = "unknown",
  predictor_prop = NULL,
  num_comp = NULL,
  engine = "mixOmics"
)
```

Arguments

<code>mode</code>	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
<code>predictor_prop</code>	The maximum proportion of original predictors that can have <i>non-zero</i> coefficients for each PLS component (via regularization). This value is used for all PLS components for X.
<code>num_comp</code>	The number of PLS components to retain.
<code>engine</code>	A single character string specifying what computational engine to use for fitting.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the `fit()` function is used with the data.

Each of the arguments in this function other than `mode` and `engine` are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
pls(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [mixOmics](#) engine details

poisson_reg

Poisson regression models

Description

`poisson_reg()` defines a generalized linear model for count data that follow a Poisson distribution. This function can fit regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [glm](#)^{1 2}
- [gee](#)²
- [glmer](#)²
- [glmnet](#)²
- [h2o](#)²
- [hurdle](#)²
- [stan](#)²
- [stan_glmer](#)²
- [zeroinfl](#)²

¹ The default engine. ² Requires a parsnip extension package.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
poisson_reg(
  mode = "regression",
  penalty = NULL,
  mixture = NULL,
  engine = "glm"
)
```

Arguments

mode	A single character string for the type of model. The only possible value for this model is "regression".
penalty	A non-negative number representing the total amount of regularization (glmnet only).
mixture	A number between zero and one (inclusive) giving the proportion of L1 regularization (i.e. lasso) in the model. • mixture = 1 specifies a pure lasso model, • mixture = 0 specifies a ridge regression model, and • $0 < \text{mixture} < 1$ specifies an elastic net model, interpolating lasso and ridge. Available for glmnet and spark only.
engine	A single character string specifying what computational engine to use for fitting.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
poisson_reg(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [glm engine details](#), [gee engine details](#), [glmer engine details](#), [glmnet engine details](#), [h2o engine details](#), [hurdle engine details](#), [stan engine details](#), [stan_glmer engine details](#), [zeroinfl engine details](#)

Description

`rand_forest()` defines a model that creates a large number of decision trees, each independent of the others. The final prediction uses all predictions from the individual trees and combines them. This function can fit classification, regression, and censored regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [ranger](#)^{1 2}
- [aorsf](#)²
- [grf](#)
- [h2o](#)²
- [ordinalForest](#)²
- [partykit](#)²
- [randomForest](#)
- [randomForestSRC](#)²
- [spark](#)

¹ The default engine. ² Requires a `parsnip` extension package for censored regression, classification, and regression.

More information on how **`parsnip`** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
rand_forest(
  mode = "unknown",
  engine = "ranger",
  mtry = NULL,
  trees = NULL,
  min_n = NULL
)
```

Arguments

<code>mode</code>	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", "classification", or "censored regression".
<code>engine</code>	A single character string specifying what computational engine to use for fitting.
<code>mtry</code>	An integer for the number of predictors that will be randomly sampled at each split when creating the tree models.
<code>trees</code>	An integer for the number of trees contained in the ensemble.
<code>min_n</code>	An integer for the minimum number of data points in a node that are required for the node to be split further.

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See `set_engine()` for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the `fit()` function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
rand_forest(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, [searchable table of parsnip models](#)

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [ranger engine details](#), [aorsf engine details](#), [grf engine details](#), [h2o engine details](#), [ordinalForest engine details](#), [partykit engine details](#), [randomForest engine details](#), [randomForestSRC engine details](#), [spark engine details](#)

Examples

```
show_engines("rand_forest")

rand_forest(mode = "classification", trees = 2000)
```

repair_call	<i>Repair a model call object</i>
-------------	-----------------------------------

Description

When the user passes a formula to `fit()` *and* the underlying model function uses a formula, the call object produced by `fit()` may not be usable by other functions. For example, some arguments may still be quosures and the data portion of the call will not correspond to the original data.

Usage

```
repair_call(x, data)
```

Arguments

x	A fitted parsnip model. An error will occur if the underlying model does not have a call element.
data	A data object that is relevant to the call. In most cases, this is the data frame that was given to parsnip for the model fit (i.e., the training set data). The name of this data object is inserted into the call.

Details

`repair_call()` call can adjust the model objects call to be usable by other functions and methods.

Value

A modified `parsnip` fitted model.

Examples

```
fitted_model <-
  linear_reg() |>
  set_engine("lm", model = TRUE) |>
  fit(mpg ~ ., data = mtcars)

# In this call, note that `data` is not `mtcars` and the `model = ~TRUE`
# indicates that the `model` argument is an rlang quosure.
fitted_model$fit$call

# All better:
repair_call(fitted_model, mtcars)$fit$call
```

req_pkgs

Determine required packages for a model

Description

[Deprecated]

Usage

```
req_pkgs(x, ...)
```

Arguments

<code>x</code>	A model specification or fit .
<code>...</code>	Not used.

Details

This function has been deprecated in favor of `required_pkgs()`.

Value

A character string of package names (if any).

`required_pkgs.model_spec`*Determine required packages for a model*

Description

Determine required packages for a model

Usage

```
## S3 method for class 'model_spec'  
required_pkgs(x, infra = TRUE, ...)
```

```
## S3 method for class 'model_fit'  
required_pkgs(x, infra = TRUE, ...)
```

Arguments

<code>x</code>	A model specification or fit .
<code>infra</code>	Should parsnip itself be included in the result?
<code>...</code>	Not used.

Value

A character vector

Examples

```
should_fail <- try(required_pkgs(linear_reg(engine = NULL)), silent = TRUE)  
should_fail
```

```
linear_reg() |>  
  set_engine("glmnet") |>  
  required_pkgs()
```

```
linear_reg() |>  
  set_engine("glmnet") |>  
  required_pkgs(infra = FALSE)
```

```
linear_reg() |>  
  set_engine("lm") |>  
  fit(mpg ~ ., data = mtcars) |>  
  required_pkgs()
```

rule_fit

*RuleFit models***Description**

rule_fit() defines a model that derives simple feature rules from a tree ensemble and uses them as features in a regularized model. This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [xrf](#)^{1 2}
- [h2o](#)²

¹ The default engine. ² Requires a parsnip extension package for classification and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
rule_fit(
  mode = "unknown",
  mtry = NULL,
  trees = NULL,
  min_n = NULL,
  tree_depth = NULL,
  learn_rate = NULL,
  loss_reduction = NULL,
  sample_size = NULL,
  stop_iter = NULL,
  penalty = NULL,
  engine = "xrf"
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
mtry	A number for the number (or proportion) of predictors that will be randomly sampled at each split when creating the tree models (specific engines only).
trees	An integer for the number of trees contained in the ensemble.
min_n	An integer for the minimum number of data points in a node that is required for the node to be split further.
tree_depth	An integer for the maximum depth of the tree (i.e. number of splits) (specific engines only).
learn_rate	A number for the rate at which the boosting algorithm adapts from iteration-to-iteration (specific engines only). This is sometimes referred to as the shrinkage parameter.

loss_reduction	A number for the reduction in the loss function required to split further (specific engines only).
sample_size	A number for the number (or proportion) of data that is exposed to the fitting routine. For xgboost, the sampling is done at each iteration while C5.0 samples once during training.
stop_iter	The number of iterations without improvement before stopping (specific engines only).
penalty	L1 regularization parameter.
engine	A single character string specifying what computational engine to use for fitting.

Details

The RuleFit model creates a regression model of rules in two stages. The first stage uses a tree-based model that is used to generate a set of rules that can be filtered, modified, and simplified. These rules are then added as predictors to a regularized generalized linear model that can also conduct feature selection during model training.

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
rule_fit(argument = !!value)
```

References

Friedman, J. H., and Popescu, B. E. (2008). "Predictive learning via rule ensembles." *The Annals of Applied Statistics*, 2(3), 916-954.

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[xrf::xrf.formula\(\)](#), [fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [xrf engine details](#), [h2o engine details](#)

Examples

```
show_engines("rule_fit")

rule_fit()
```

set_args	<i>Change elements of a model specification</i>
----------	---

Description

set_args() can be used to modify the arguments of a model specification while set_mode() is used to change the model's mode.

Usage

```
set_args(object, ...)

set_mode(object, mode, ...)

## S3 method for class 'model_spec'
set_mode(object, mode, quantile_levels = NULL, ...)
```

Arguments

object	A model specification .
...	One or more named model arguments.
mode	A character string for the model type (e.g. "classification" or "regression")
quantile_levels	A vector of values between zero and one (only for the "quantile regression" mode); otherwise, it is NULL. The model uses these values to appropriately train quantile regression models to make predictions for these values (e.g., quantile_levels = 0.5 is the median).

Details

set_args() will replace existing values of the arguments.

Value

An updated model object.

Examples

```
rand_forest()

rand_forest() |>
  set_args(mtry = 3, importance = TRUE) |>
  set_mode("regression")

linear_reg() |>
  set_mode("quantile regression", quantile_levels = c(0.2, 0.5, 0.8))
```

set_engine

Declare a computational engine and specific arguments

Description

set_engine() is used to specify which package or system will be used to fit the model, along with any arguments specific to that software.

Usage

```
set_engine(object, engine, ...)
```

Arguments

object	A model specification .
engine	A character string for the software that should be used to fit the model. This is highly dependent on the type of model (e.g. linear regression, random forest, etc.).
...	Any optional arguments associated with the chosen computational engine. These are captured as quosures and can be tuned with tune().

Details

In parsnip,

- the model **type** differentiates basic modeling approaches, such as random forests, logistic regression, linear support vector machines, etc.,
- the **mode** denotes in what kind of modeling context it will be used (most commonly, classification or regression), and
- the computational **engine** indicates how the model is fit, such as with a specific R package implementation or even methods outside of R like Keras or Stan.

Use [show_engines\(\)](#) to get a list of possible engines for the model of interest.

Modeling functions in parsnip separate model arguments into two categories:

- *Main arguments* are more commonly used and tend to be available across engines. These names are standardized to work with different engines in a consistent way, so you can use the parsnip main argument `trees`, instead of the heterogeneous arguments for this parameter from **ranger** and **randomForest** packages (`num.trees` and `ntree`, respectively). Set these in your model type function, like `rand_forest(trees = 2000)`.
- *Engine arguments* are either specific to a particular engine or used more rarely; there is no change for these argument names from the underlying engine. The `...` argument of `set_engine()` allows any engine-specific argument to be passed directly to the engine fitting function, like `set_engine("ranger", importance = "permutation")`.

Value

An updated model specification.

Examples

```
# First, set main arguments using the standardized names
logistic_reg(penalty = 0.01, mixture = 1/3) |>
  # Now specify how you want to fit the model with another argument
  set_engine("glmnet", nlambdas = 10) |>
  translate()

# Many models have possible engine-specific arguments
decision_tree(tree_depth = 5) |>
  set_engine("rpart", parms = list(prior = c(.65,.35))) |>
  set_mode("classification") |>
  translate()
```

show_engines

Display currently available engines for a model

Description

The possible engines for a model can depend on what packages are loaded. Some `parsnip` extension add engines to existing models. For example, the **poissonreg** package adds additional engines for the `poisson_reg()` model and these are not available unless **poissonreg** is loaded.

Usage

```
show_engines(x)
```

Arguments

`x` The name of a `parsnip` model (e.g., "linear_reg", "mars", etc.)

Value

A tibble.

Examples

```
show_engines("linear_reg")
```

sparse_data

*Using sparse data with parsnip***Description**

You can figure out whether a given model engine supports sparse data by calling `get_encoding("name of model")` and looking at the `allow_sparse_x` column.

Details

Using sparse data for model fitting and prediction shouldn't require any additional configurations. Just pass in a sparse matrix such as `dgCMatrix` from the `Matrix` package or a sparse tibble from the `sparsevctrs` package to the `data` argument of `fit()`, `fit_xy()`, and `predict()`.

Models that don't support sparse data will try to convert to non-sparse data with warnings. If conversion isn't possible, an informative error will be thrown.

svm_linear

*Linear support vector machines***Description**

`svm_linear()` defines a support vector machine model. For classification, the model tries to maximize the width of the margin between classes (using a linear class boundary). For regression, the model optimizes a robust loss function that is only affected by very large model residuals and uses a linear fit. This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [Liblinear](#)¹
- [kernlab](#)

¹ The default engine.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
svm_linear(mode = "unknown", engine = "Liblinear", cost = NULL, margin = NULL)
```

Arguments

<code>mode</code>	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
<code>engine</code>	A single character string specifying what computational engine to use for fitting.
<code>cost</code>	A positive number for the cost of predicting a sample within or on the wrong side of the margin
<code>margin</code>	A positive number for the epsilon in the SVM insensitive loss function (regression only)

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
svm_linear(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [Liblinear engine details](#), [kernlab engine details](#)

Examples

```
show_engines("svm_linear")

svm_linear(mode = "classification")
```

svm_poly

Polynomial support vector machines

Description

`svm_poly()` defines a support vector machine model. For classification, the model tries to maximize the width of the margin between classes using a polynomial class boundary. For regression, the model optimizes a robust loss function that is only affected by very large model residuals and uses polynomial functions of the predictors. This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [kernlab](#)¹

¹ The default engine.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
svm_poly(  
  mode = "unknown",  
  engine = "kernlab",  
  cost = NULL,  
  degree = NULL,  
  scale_factor = NULL,  
  margin = NULL  
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
engine	A single character string specifying what computational engine to use for fitting.
cost	A positive number for the cost of predicting a sample within or on the wrong side of the margin
degree	A positive number for polynomial degree.
scale_factor	A positive number for the polynomial scaling factor.
margin	A positive number for the epsilon in the SVM insensitive loss function (regression only)

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See [set_engine\(\)](#) for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the [fit\(\)](#) function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1  
svm_poly(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [kernlab engine details](#)

Examples

```
show_engines("svm_poly")  
  
svm_poly(mode = "classification", degree = 1.2)
```

Description

`svm_rbf()` defines a support vector machine model. For classification, the model tries to maximize the width of the margin between classes using a nonlinear class boundary. For regression, the model optimizes a robust loss function that is only affected by very large model residuals and uses nonlinear functions of the predictors. The function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [kernlab](#)¹

¹ The default engine.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
svm_rbf(
  mode = "unknown",
  engine = "kernlab",
  cost = NULL,
  rbf_sigma = NULL,
  margin = NULL
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
engine	A single character string specifying what computational engine to use for fitting. Possible engines are listed below. The default for this model is "kernlab".
cost	A positive number for the cost of predicting a sample within or on the wrong side of the margin
rbf_sigma	A positive number for radial basis function.
margin	A positive number for the epsilon in the SVM insensitive loss function (regression only)

Details

This function only defines what *type* of model is being fit. Once an engine is specified, the *method* to fit the model is also defined. See `set_engine()` for more on setting the engine, including how to set engine arguments.

The model is not trained or fit until the `fit()` function is used with the data.

Each of the arguments in this function other than mode and engine are captured as [quosures](#). To pass values programmatically, use the [injection operator](#) like so:

```
value <- 1
svm_rbf(argument = !!value)
```

References

<https://www.tidymodels.org>, *Tidy Modeling with R*, searchable table of parsnip models

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [kernlab engine details](#)

Examples

```
show_engines("svm_rbf")

svm_rbf(mode = "classification", rbf_sigma = 0.2)
```

tabular_auto_int

AutoInt: Automatic Feature Interaction Learning

Description

`tabular_auto_int()` uses an attention mechanism to automatically learn embedding co-representations for tabular data. This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [brulee](#)¹²

¹ The default engine. ² Requires a parsnip extension package for classification and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
tabular_auto_int(
  mode = "unknown",
  engine = "brulee",
  epochs = NULL,
  num_embedding = NULL,
  hidden_units = NULL,
  hidden_activations = NULL,
  num_attn_feat = NULL,
  num_attn_heads = NULL,
  num_attn_blocks = NULL,
```

```

    activation = NULL,
    dropout = NULL,
    dropout_attn = NULL,
    dropout_embedding = NULL,
    penalty = NULL,
    mixture = NULL,
    learn_rate = NULL,
    rate_schedule = NULL,
    momentum = NULL,
    batch_size = NULL,
    class_weights = NULL,
    stop_iter = NULL
)

```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
engine	A single character string specifying what computational engine to use for fitting.
epochs	An integer for the number of training iterations.
num_embedding	An integer for the dimensionality of the embedding space for features.
hidden_units	An integer vector for the number of units in the hidden layers after the attention mechanism.
hidden_activations	A character vector denoting the activation functions for the hidden layers.
num_attn_feat	An integer for the number of attention features.
num_attn_heads	An integer for the number of attention heads in the multi-head attention mechanism.
num_attn_blocks	An integer for the number of sequential attention blocks.
activation	A character vector denoting the type of relationship between the different layers. The activation function between the hidden and output layers is automatically set to either "linear" or "softmax" depending on the type of outcome. Possible values depend on the engine being used.
dropout	A number between 0 (inclusive) and 1 denoting the proportion of model parameters randomly set to zero during model training.
dropout_attn	A number between 0 (inclusive) and 1 denoting the proportion of attention weights set to zero during model training.
dropout_embedding	A number between 0 (inclusive) and 1 denoting the proportion of embedding values set to zero during model training.
penalty	A non-negative numeric value for the amount of weight decay.
mixture	A number between zero and one (inclusive) denoting the proportion of L1 regularization (i.e. lasso) in the model. • mixture = 1 specifies a pure lasso model,

	• <code>mixture = 0</code> specifies a ridge regression model, and • $0 < \text{mixture} < 1$ specifies an elastic net model, interpolating lasso and ridge.
	Available for specific engines only.
<code>learn_rate</code>	A number for the rate at which the boosting algorithm adapts from iteration-to-iteration (specific engines only). This is sometimes referred to as the shrinkage parameter.
<code>rate_schedule</code>	A character string for the learning rate schedule.
<code>momentum</code>	A number for the momentum parameter in optimizers that use it.
<code>batch_size</code>	An integer for the number of training instances in each batch.
<code>class_weights</code>	Numeric class weights for imbalanced data (classification only).
<code>stop_iter</code>	The number of iterations without improvement before stopping (specific engines only).

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [brulee engine details](#), [brulee::brulee_auto_int\(\)](#)

Examples

```
show_engines("tabular_auto_int")

tabular_auto_int(mode = "classification", num_attn_blocks = 4)
```

tabular_chronos	<i>Chronos-2 pretrained forecasting model</i>
-----------------	---

Description

`tabular_chronos()` defines a pretrained time-series forecasting model that produces quantile (distributional) forecasts. The network has fixed pretrained weights, so no training is performed; the historical ("context") data is ingested at fit time and the model forecasts a fixed horizon. This function can fit quantile regression and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [brulee](#)^{1,2}

¹ The default engine. ² Requires a `parsnip` extension package for quantile regression and regression. More information on how **`parsnip`** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
tabular_chronos(mode = "unknown", engine = "brulee")
```

Arguments

mode	A single character string for the type of model. The possible values for this model are "quantile regression" (the natural mode, which returns a <code>hardhat::quantile_pred()</code>) and "regression" (which returns the median point forecast). The mode must be set before fitting; for "quantile regression" it is set with <code>set_mode("quantile regression", quantile_levels = ...)</code> .
engine	A single character string specifying what computational engine to use for fitting. The only valid value is "brulee".

Details

Unlike the other models in this package, Chronos-2 is pretrained and has no tuning parameters. Forecast configuration is supplied through the engine with `set_engine()`, e.g. `set_engine("brulee", prediction_length = 14)`. The available engine arguments mirror `brulee::brulee_chronos()`: `prediction_length`, `id_column`, `timestamp_column`, `model_id`, `revision`, `device`, and `cache_dir`. The `quantile_levels` are taken from the mode (via `set_mode()`) and forwarded to the fit automatically.

On first use the engine downloads the pretrained weights (about 500MB) and caches them locally.

The **parsnip** interface forecasts a **single series**. `predict()` returns one row per horizon step (`.pred_quantile` or `.pred`), which cannot unambiguously represent more than one series; supplying data with multiple `id_column` values is therefore an error. For multi-series forecasting, call `brulee::brulee_chronos()` directly, where the `id` column is retained in the output.

References

Ansari, A. F., Shchur, O., Küken, J., Auer, A., Han, B., Mercado, P., et al. (2025). "Chronos-2: From univariate to universal forecasting." *arXiv preprint* arXiv:2510.15821.

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [brulee engine details](#), [brulee::brulee_chronos\(\)](#)

Examples

```
show_engines("tabular_chronos")

# Quantile (distributional) forecast
tabular_chronos() |>
  set_engine("brulee", prediction_length = 14) |>
  set_mode("quantile regression", quantile_levels = (1:9) / 10)

# Median point forecast
tabular_chronos() |>
  set_engine("brulee", prediction_length = 14) |>
  set_mode("regression")
```

tabular_icl

TabICL: prior data fitted networks

Description

This function uses a pre-trained deep learning network that emulates Bayesian inference. The model was trained on a large number of simulated data sets and an attention mechanism is used to make relevant predictions for specific (i.e., real) data sets.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [brulee](#)¹²

¹ The default engine. ² Requires a **parsnip** extension package for classification and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
tabular_icl(
  mode = "unknown",
  engine = "brulee",
  num_estimators = NULL,
  softmax_temperature = NULL
)
```

Arguments

mode	A single character value for the type of model. The possible values for this model are "classification" and "regression".
engine	A single character string specifying what computational engine to use for fitting. Possible engines are listed below. The default for this model is "brulee".
num_estimators	An integer for the ensemble size. Default is 8L.
softmax_temperature	An adjustment factor that is a divisor in the exponents of the softmax function (see brulee::brulee_tab_icl()). Defaults to 0.9.

Details

This function fits classification and regression models.

References

<https://github.com/soda-inria/tabicl>

Qu, J., Holzmüller, D., Varoquaux, G., & Morvan, M. L. (2025). Tabicl: A tabular foundation model for in-context learning on large data. arXiv preprint arXiv:2502.05564.

Qu, J., Holzmüller, D., Varoquaux, G., & Morvan, M. L. (2026). TabICLv2: A better, faster, scalable, and open tabular foundation model. arXiv preprint arXiv:2602.11139.

See Also

`fit()`, `set_engine()`, `update()`, `brulee` engine details, `brulee::brulee_tab_icl()`

Examples

```
show_engines("tabular_icl")

tabular_icl()
```

tabular_pfn

TabPFN: prior data fitted networks

Description

This function uses a pre-trained deep learning network that emulates Bayesian inference. The model was trained on a large number of simulated data sets and an attention mechanism is used to make relevant predictions for specific (i.e., real) data sets.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- `tabpfn`¹²

¹ The default engine. ² Requires a `parsnip` extension package for classification and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
tabular_pfn(
  mode = "unknown",
  engine = "tabpfn",
  num_estimators = NULL,
  softmax_temperature = NULL,
  balance_probabilities = NULL,
  average_before_softmax = NULL
)
```

Arguments

<code>mode</code>	A single character value for the type of model. The possible values for this model are "classification" and "regression".
<code>engine</code>	A single character string specifying what computational engine to use for fitting. Possible engines are listed below. The default for this model is "tabpfn".
<code>num_estimators</code>	An integer for the ensemble size. Default is 8L.
<code>softmax_temperature</code>	An adjustment factor that is a divisor in the exponents of the softmax function (see <code>tabpfn::tab_pfn()</code>). Defaults to 0.9.

balance_probabilities

A logical to adjust the prior probabilities in cases where there is a class imbalance. Default is FALSE. Classification only.

average_before_softmax

A logical. For cases where num_estimators > 1, should the average be done before using the softmax function or after? Default is FALSE.

Details

This function fits classification and regression models.

References

<https://github.com/PriorLabs/TabPFN>

Hollmann, Noah, Samuel Müller, Lennart Purucker, Arjun Krishnakumar, Max Körfer, Shi Bin Hoo, Robin Tibor Schirrmeister, and Frank Hutter. "Accurate predictions on small data with a tabular foundation model." *Nature* 637, no. 8045 (2025): 319-326.

Hollmann, Noah, Samuel Müller, Katharina Eggersperger, and Frank Hutter. "Tabpfn: A transformer that solves small tabular classification problems in a second." *arXiv preprint arXiv:2207.01848* (2022).

Müller, Samuel, Noah Hollmann, Sebastian Pineda Arango, Josif Grabocka, and Frank Hutter. "Transformers can do Bayesian inference." *arXiv preprint arXiv:2112.10510* (2021).

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [tabpfn engine details](#), [tabpfn::tab_pfn\(\)](#)

Examples

```
show_engines("tabular_pfn")
```

```
tabular_pfn()
```

tabular_resnet

Residual Neural Network for Tabular Data

Description

tabular_resnet() ... This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [brulee](#)^{1,2}

¹ The default engine. ² Requires a parsnip extension package for classification and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```

tabular_resnet(
  mode = "unknown",
  engine = "brulee",
  hidden_units = NULL,
  bottleneck_units = NULL,
  residual_at = NULL,
  penalty = NULL,
  mixture = NULL,
  dropout = NULL,
  epochs = NULL,
  activation = NULL,
  learn_rate = NULL,
  rate_schedule = NULL,
  momentum = NULL,
  batch_size = NULL,
  class_weights = NULL,
  stop_iter = NULL
)

```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
engine	A single character string specifying what computational engine to use for fitting.
hidden_units	An integer vector for the number of units in the hidden model.
bottleneck_units	The number of embeddings that are produced by batch normalization.
residual_at	An integer vector with the layer number should use a residual connection (i.e., skip layer).
penalty	A non-negative numeric value for the amount of weight decay.
mixture	A number between zero and one (inclusive) giving the proportion of L1 regularization (i.e. lasso) in the model. <code>mixture = 1</code> is a pure lasso model while <code>mixture = 0</code> indicates ridge regression (a.k.a weight decay).
dropout	A number between 0 (inclusive) and 1 denoting the proportion of model parameters randomly set to zero during model training.
epochs	An integer for the number of training iterations.
activation	A vector character strings denoting the type of relationship between the layers. The activation function between the hidden and output layers is automatically set to either "linear" or "softmax" depending on the type of outcome. Possible values depend on the engine being used.
learn_rate	A number for the rate at which the boosting algorithm adapts from iteration-to-iteration (specific engines only). This is sometimes referred to as the shrinkage parameter.
rate_schedule	A character string for the learning rate schedule.

momentum	A number for the momentum parameter in optimizers that use it.
batch_size	An integer for the number of training instances in each batch.
class_weights	Numeric class weights for imbalanced data (classification only).
stop_iter	The number of iterations without improvement before stopping (specific engines only).

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [brulee engine details](#), [brulee::brulee_resnet\(\)](#)

Examples

```
show_engines("tabular_resnet")

tabular_resnet(mode = "classification", penalty = 0.01)
```

tabular_rln	<i>Regularization Learning Network</i>
-------------	--

Description

`tabular_rln()` defines a single-hidden-layer neural network where each weight learns its own adaptive regularization coefficient. This function can fit regression models only.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [brulee](#)¹²

¹ The default engine. ² Requires a `parsnip` extension package.

More information on how **`parsnip`** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
tabular_rln(
  mode = "regression",
  engine = "brulee",
  hidden_units = NULL,
  penalty_type = NULL,
  penalty_average = NULL,
  step_rate = NULL,
  activation = NULL,
  epochs = NULL,
  learn_rate = NULL,
  rate_schedule = NULL,
  momentum = NULL,
  batch_size = NULL,
  stop_iter = NULL
)
```

Arguments

mode	A single character string for the type of model. The only valid value is "regression".
engine	A single character string specifying what computational engine to use for fitting. The only valid value is "brulee".
hidden_units	An integer for the number of units in the single hidden layer (≥ 1).
penalty_type	A string for the regularization norm: "L1" (default in brulee) or "L2". L1 is recommended by the original paper.
penalty_average	A positive numeric value for the target geometric mean of the per-weight regularization coefficients, on the natural scale. Best tuned on the log10 scale via dials::penalty_average() .
step_rate	A positive numeric value for the step size used to update the per-weight regularization coefficients, on the natural scale. Best tuned on the log10 scale via dials::step_rate() .
activation	A character string for the activation function between the hidden and output layers (e.g., "relu", "elu", "tanh").
epochs	An integer for the number of training iterations.
learn_rate	A positive number for the learning rate.
rate_schedule	A character string for the learning rate schedule (e.g., "none", "decay_time", "cyclic").
momentum	A number between 0 (inclusive) and 1 for the momentum parameter used by the optimizer.
batch_size	An integer for the number of training samples used per gradient update step.
stop_iter	A non-negative integer for the number of iterations without improvement before stopping training early.

References

Shavitt, I., & Segal, E. (2018). Regularization learning networks: Deep learning for tabular datasets. *Advances in Neural Information Processing Systems*, 31, 1379-1389.

See Also

[fit\(\)](#), [set_engine\(\)](#), [update\(\)](#), [brulee engine details](#), [brulee::brulee_rln\(\)](#)

Examples

```
show_engines("tabular_rln")

tabular_rln(hidden_units = 10L, penalty_average = 1e-8)
```

tabular_saint

SAINT: Self-Attention and Inter-sample Attention Transformer

Description

tabular_saint() uses self-attention and inter-sample attention mechanisms to learn feature interactions for tabular data. This function can fit classification and regression models.

There are different ways to fit this model, and the method of estimation is chosen by setting the model *engine*. The engine-specific pages for this model are listed below.

- [brulee](#)¹²

¹ The default engine. ² Requires a parsnip extension package for classification and regression.

More information on how **parsnip** is used for modeling is at <https://www.tidymodels.org/>.

Usage

```
tabular_saint(
  mode = "unknown",
  engine = "brulee",
  epochs = NULL,
  num_embedding = NULL,
  attention_type = NULL,
  num_attn_heads = NULL,
  num_attn_blocks = NULL,
  dropout_attn = NULL,
  dropout_hidden = NULL,
  dropout_last = NULL,
  hidden_units = NULL,
  hidden_activations = NULL,
  target_token = NULL,
  penalty = NULL,
  mixture = NULL,
  learn_rate = NULL,
  rate_schedule = NULL,
  momentum = NULL,
  batch_size = NULL,
  class_weights = NULL,
  stop_iter = NULL
)
```

Arguments

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
engine	A single character string specifying what computational engine to use for fitting.

epochs	An integer for the number of training iterations.
num_embedding	An integer for the dimensionality of the embedding space for features.
attention_type	A character string for the type of attention to use. Options are "column" (SAINT-s), "row" (SAINT-i), or "both" (full SAINT).
num_attn_heads	An integer for the number of attention heads in the multi-head attention mechanism.
num_attn_blocks	An integer for the number of sequential attention blocks.
dropout_attn	A number between 0 (inclusive) and 1 denoting the proportion of attention weights set to zero during model training.
dropout_hidden	A number between 0 (inclusive) and 1 denoting the proportion of values in the feed-forward layers set to zero during training.
dropout_last	A number between 0 (inclusive) and 1 denoting the proportion of values set to zero between the last hidden layer and the output head.
hidden_units	An integer vector for the number of units in the hidden layers after the attention mechanism.
hidden_activations	A character vector denoting the activation functions for the hidden layers.
target_token	A logical for whether to use a learnable target token (CLS-like embedding) to aggregate information for prediction. When TRUE, the model appends a special target token to the input that attends to all features via the attention mechanism.
penalty	A non-negative numeric value for the amount of weight decay.
mixture	A number between zero and one (inclusive) denoting the proportion of L1 regularization (i.e. lasso) in the model. • mixture = 1 specifies a pure lasso model, • mixture = 0 specifies a ridge regression model, and • $0 < \text{mixture} < 1$ specifies an elastic net model, interpolating lasso and ridge. Available for specific engines only.
learn_rate	A number for the rate at which the boosting algorithm adapts from iteration-to-iteration (specific engines only). This is sometimes referred to as the shrinkage parameter.
rate_schedule	A character string for the learning rate schedule.
momentum	A number for the momentum parameter in optimizers that use it.
batch_size	An integer for the number of training instances in each batch.
class_weights	Numeric class weights for imbalanced data (classification only).
stop_iter	The number of iterations without improvement before stopping (specific engines only).

Details

Row attention at prediction time:

When `attention_type` is "row" or "both", SAINT applies inter-sample (row) attention across the samples in a batch. The **brulee** engine keeps this on at prediction time by default, so the prediction for a given row depends on which other rows are passed to `predict()` in the same call. To obtain batch-independent predictions (where a row's prediction does not change with its neighbors), bypass row attention at predict time with `set_engine("brulee", row_attention_on_predict = FALSE)`.

References

Somepalli, G., Goldblum, M., Schwarzschild, A., Bruss, C. B., & Goldstein, T. (2021). SAINT: Improved Neural Networks for Tabular Data via Row Attention and Contrastive Pre-Training. arXiv:2106.01342.

See Also

`fit()`, `set_engine()`, `update()`, [brulee engine details](#), `brulee::brulee_saint()`

Examples

```
show_engines("tabular_saint")

tabular_saint(mode = "classification", num_attn_blocks = 4)
```

tidy.model_fit

Turn a parsnip model object into a tidy tibble

Description

This method tidies the model in a parsnip model object, if it exists.

Usage

```
## S3 method for class 'model_fit'
tidy(x, ...)
```

Arguments

`x` An object to be converted into a tidy `tibble::tibble()`.

`...` Additional arguments to tidying method.

Value

a tibble

translate

Resolve a Model Specification for a Computational Engine

Description

translate() will translate a [model specification](#) into a code object that is specific to a particular engine (e.g. R package). It translates generic parameters to their counterparts.

Usage

```
translate(x, ...)

## Default S3 method:
translate(x, engine = x$engine, ...)
```

Arguments

x	A model specification .
...	Not currently used.
engine	The computational engine for the model (see ?set_engine).

Details

translate() produces a *template* call that lacks the specific argument values (such as data, etc). These are filled in once fit() is called with the specifics of the data for the model. The call may also include tune() arguments if these are in the specification. To handle the tune() arguments, you need to use the [tune package](#). For more information see <https://www.tidymodels.org/start/tuning/>

It does contain the resolved argument names that are specific to the model fitting function/engine.

This function can be useful when you need to understand how parsnip goes from a generic model specific to a model fitting function.

Note: this function is used internally and users should only use it to understand what the underlying syntax would be. It should not be used to modify the model specification.

Examples

```
lm_spec <- linear_reg(penalty = 0.01)

# `penalty` is translated to `lambda`
translate(lm_spec, engine = "glmnet")

# `penalty` not applicable for this model.
translate(lm_spec, engine = "lm")

# `penalty` is translated to `reg_param`
translate(lm_spec, engine = "spark")
```

```
# with a placeholder for an unknown argument value:
translate(linear_reg(penalty = tune(), mixture = tune()), engine = "glmnet")
```

update.bag_mars

Updating a model specification

Description

If parameters of a model specification need to be modified, `update()` can be used in lieu of recreating the object from scratch.

Usage

```
## S3 method for class 'bag_mars'
update(
  object,
  parameters = NULL,
  num_terms = NULL,
  prod_degree = NULL,
  prune_method = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'bag_mlp'
update(
  object,
  parameters = NULL,
  hidden_units = NULL,
  penalty = NULL,
  epochs = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'bag_tree'
update(
  object,
  parameters = NULL,
  cost_complexity = NULL,
  tree_depth = NULL,
  min_n = NULL,
  class_cost = NULL,
  fresh = FALSE,
  ...
)
```

```
## S3 method for class 'bart'
update(
  object,
  parameters = NULL,
  trees = NULL,
  prior_terminal_node_coef = NULL,
  prior_terminal_node_expo = NULL,
  prior_outcome_range = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'boost_tree'
update(
  object,
  parameters = NULL,
  mtry = NULL,
  trees = NULL,
  min_n = NULL,
  tree_depth = NULL,
  learn_rate = NULL,
  loss_reduction = NULL,
  sample_size = NULL,
  stop_iter = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'C5_rules'
update(
  object,
  parameters = NULL,
  trees = NULL,
  min_n = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'cubist_rules'
update(
  object,
  parameters = NULL,
  committees = NULL,
  neighbors = NULL,
  max_rules = NULL,
  fresh = FALSE,
  ...
)
```

```
)

## S3 method for class 'decision_tree'
update(
  object,
  parameters = NULL,
  cost_complexity = NULL,
  tree_depth = NULL,
  min_n = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'discrim_flexible'
update(
  object,
  num_terms = NULL,
  prod_degree = NULL,
  prune_method = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'discrim_linear'
update(
  object,
  penalty = NULL,
  regularization_method = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'discrim_quad'
update(object, regularization_method = NULL, fresh = FALSE, ...)

## S3 method for class 'discrim_regularized'
update(
  object,
  frac_common_cov = NULL,
  frac_identity = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'gen_additive_mod'
update(
  object,
  select_features = NULL,
```

```
    adjust_deg_free = NULL,
    parameters = NULL,
    fresh = FALSE,
    ...
)

## S3 method for class 'linear_reg'
update(
  object,
  parameters = NULL,
  penalty = NULL,
  mixture = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'logistic_reg'
update(
  object,
  parameters = NULL,
  penalty = NULL,
  mixture = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'mars'
update(
  object,
  parameters = NULL,
  num_terms = NULL,
  prod_degree = NULL,
  prune_method = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'mlp'
update(
  object,
  parameters = NULL,
  hidden_units = NULL,
  penalty = NULL,
  dropout = NULL,
  epochs = NULL,
  activation = NULL,
  learn_rate = NULL,
  fresh = FALSE,
```

```
    ...
  )

## S3 method for class 'multinom_reg'
update(
  object,
  parameters = NULL,
  penalty = NULL,
  mixture = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'naive_Bayes'
update(object, smoothness = NULL, Laplace = NULL, fresh = FALSE, ...)

## S3 method for class 'nearest_neighbor'
update(
  object,
  parameters = NULL,
  neighbors = NULL,
  weight_func = NULL,
  dist_power = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'ordinal_reg'
update(
  object,
  parameters = NULL,
  ordinal_link = NULL,
  odds_link = NULL,
  threshold_structure = NULL,
  parallel_reg = NULL,
  penalty = NULL,
  mixture = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'pls'
update(
  object,
  parameters = NULL,
  predictor_prop = NULL,
  num_comp = NULL,
  fresh = FALSE,
```

```
    ...
)

## S3 method for class 'poisson_reg'
update(
  object,
  parameters = NULL,
  penalty = NULL,
  mixture = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'proportional_hazards'
update(
  object,
  parameters = NULL,
  penalty = NULL,
  mixture = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'rand_forest'
update(
  object,
  parameters = NULL,
  mtry = NULL,
  trees = NULL,
  min_n = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'rule_fit'
update(
  object,
  parameters = NULL,
  mtry = NULL,
  trees = NULL,
  min_n = NULL,
  tree_depth = NULL,
  learn_rate = NULL,
  loss_reduction = NULL,
  sample_size = NULL,
  penalty = NULL,
  fresh = FALSE,
  ...
)
```

```
)

## S3 method for class 'surv_reg'
update(object, parameters = NULL, dist = NULL, fresh = FALSE, ...)

## S3 method for class 'survival_reg'
update(object, parameters = NULL, dist = NULL, fresh = FALSE, ...)

## S3 method for class 'svm_linear'
update(
  object,
  parameters = NULL,
  cost = NULL,
  margin = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'svm_poly'
update(
  object,
  parameters = NULL,
  cost = NULL,
  degree = NULL,
  scale_factor = NULL,
  margin = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'svm_rbf'
update(
  object,
  parameters = NULL,
  cost = NULL,
  rbf_sigma = NULL,
  margin = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'tabular_auto_int'
update(
  object,
  parameters = NULL,
  epochs = NULL,
  num_embedding = NULL,
  hidden_units = NULL,
```

```

        hidden_activations = NULL,
        num_attn_feat = NULL,
        num_attn_heads = NULL,
        num_attn_blocks = NULL,
        activation = NULL,
        dropout = NULL,
        dropout_attn = NULL,
        dropout_embedding = NULL,
        penalty = NULL,
        mixture = NULL,
        learn_rate = NULL,
        rate_schedule = NULL,
        momentum = NULL,
        batch_size = NULL,
        class_weights = NULL,
        stop_iter = NULL,
        fresh = FALSE,
        ...
    )

## S3 method for class 'tabular_chronos'
update(object, parameters = NULL, fresh = FALSE, ...)

## S3 method for class 'tabular_icl'
update(
  object,
  parameters = NULL,
  num_estimators = NULL,
  softmax_temperature = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'tabular_pfn'
update(
  object,
  parameters = NULL,
  num_estimators = NULL,
  softmax_temperature = NULL,
  balance_probabilities = NULL,
  average_before_softmax = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'tabular_resnet'
update(
  object,

```

```
parameters = NULL,
hidden_units = NULL,
bottleneck_units = NULL,
residual_at = NULL,
penalty = NULL,
mixture = NULL,
dropout = NULL,
epochs = NULL,
activation = NULL,
learn_rate = NULL,
rate_schedule = NULL,
momentum = NULL,
batch_size = NULL,
class_weights = NULL,
stop_iter = NULL,
fresh = FALSE,
...
)

## S3 method for class 'tabular_rln'
update(
  object,
  parameters = NULL,
  hidden_units = NULL,
  penalty_type = NULL,
  penalty_average = NULL,
  step_rate = NULL,
  activation = NULL,
  epochs = NULL,
  learn_rate = NULL,
  rate_schedule = NULL,
  momentum = NULL,
  batch_size = NULL,
  stop_iter = NULL,
  fresh = FALSE,
  ...
)

## S3 method for class 'tabular_saint'
update(
  object,
  parameters = NULL,
  epochs = NULL,
  num_embedding = NULL,
  attention_type = NULL,
  num_attn_heads = NULL,
  num_attn_blocks = NULL,
  dropout_attn = NULL,
```

```

dropout_hidden = NULL,
dropout_last = NULL,
hidden_units = NULL,
hidden_activations = NULL,
target_token = NULL,
penalty = NULL,
mixture = NULL,
learn_rate = NULL,
rate_schedule = NULL,
momentum = NULL,
batch_size = NULL,
class_weights = NULL,
stop_iter = NULL,
fresh = FALSE,
...
)

```

Arguments

object	A model specification .
parameters	A 1-row tibble or named list with <i>main</i> parameters to update. Use either parameters or the main arguments directly when updating. If the main arguments are used, these will supersede the values in parameters. Also, using engine arguments in this object will result in an error.
num_terms	The number of features that will be retained in the final model, including the intercept.
prod_degree	The highest possible interaction degree.
prune_method	The pruning method.
fresh	A logical for whether the arguments should be modified in-place or replaced wholesale.
...	Not used for update().
hidden_units	An integer vector for the number of units in the hidden layers after the attention mechanism.
penalty	An non-negative number representing the amount of regularization used by some of the engines.
epochs	An integer for the number of training iterations.
cost_complexity	A positive number for the the cost/complexity parameter (a.k.a. Cp) used by CART models (specific engines only).
tree_depth	An integer for maximum depth of the tree.
min_n	An integer for the minimum number of data points in a node that are required for the node to be split further.
class_cost	A non-negative scalar for a class cost (where a cost of 1 means no extra cost). This is useful for when the first level of the outcome factor is the minority class. If this is not the case, values between zero and one can be used to bias to the second level of the factor.

trees	An integer for the number of trees contained in the ensemble.
prior_terminal_node_coef	A coefficient for the prior probability that a node is a terminal node.
prior_terminal_node_expo	An exponent in the prior probability that a node is a terminal node.
prior_outcome_range	A positive value that defines the width of a prior that the predicted outcome is within a certain range. For regression it is related to the observed range of the data; the prior is the number of standard deviations of a Gaussian distribution defined by the observed range of the data. For classification, it is defined as the range of ± 3 (assumed to be on the logit scale). The default value is 2.
mtry	A number for the number (or proportion) of predictors that will be randomly sampled at each split when creating the tree models (specific engines only).
learn_rate	A number for the rate at which the boosting algorithm adapts from iteration-to-iteration (specific engines only). This is sometimes referred to as the shrinkage parameter.
loss_reduction	A number for the reduction in the loss function required to split further (specific engines only).
sample_size	A number for the number (or proportion) of data that is exposed to the fitting routine. For xgboost, the sampling is done at each iteration while C5.0 samples once during training.
stop_iter	The number of iterations without improvement before stopping (specific engines only).
committees	A non-negative integer (no greater than 100) for the number of members of the ensemble.
neighbors	An integer between zero and nine for the number of training set instances that are used to adjust the model-based prediction.
max_rules	The largest number of rules.
regularization_method	A character string for the type of regularized estimation. Possible values are: "diagonal", "min_distance", "shrink_cov", and "shrink_mean" (sparsediscrim engine only).
frac_common_cov, frac_identity	Numeric values between zero and one.
select_features	TRUE or FALSE. If TRUE, the model has the ability to eliminate a predictor (via penalization). Increasing adjust_deg_free will increase the likelihood of removing predictors.
adjust_deg_free	If select_features = TRUE, then acts as a multiplier for smoothness. Increase this beyond 1 to produce smoother models.
mixture	A number between zero and one (inclusive) denoting the proportion of L1 regularization (i.e. lasso) in the model. • mixture = 1 specifies a pure lasso model,

	• <code>mixture = 0</code> specifies a ridge regression model, and • $0 < \text{mixture} < 1$ specifies an elastic net model, interpolating lasso and ridge.
	Available for specific engines only.
<code>dropout</code>	A number between 0 (inclusive) and 1 denoting the proportion of model parameters randomly set to zero during model training.
<code>activation</code>	A character vector denoting the type of relationship between the different layers. The activation function between the hidden and output layers is automatically set to either "linear" or "softmax" depending on the type of outcome. Possible values depend on the engine being used.
<code>smoothness</code>	An non-negative number representing the the relative smoothness of the class boundary. Smaller examples result in model flexible boundaries and larger values generate class boundaries that are less adaptable
<code>Laplace</code>	A non-negative value for the Laplace correction to smoothing low-frequency counts.
<code>weight_func</code>	A <i>single</i> character for the type of kernel function used to weight distances between samples. Valid choices are: "rectangular", "triangular", "epanechnikov", "biweight", "triweight", "cos", "inv", "gaussian", "rank", or "optimal".
<code>dist_power</code>	A single number for the parameter used in calculating Minkowski distance.
<code>ordinal_link</code>	The ordinal link function.
<code>odds_link</code>	The odds or probability link function.
<code>threshold_structure</code>	The threshold structure for the cutpoints (specific engines only).
<code>parallel_reg</code>	Logical; whether predictor effects are shared across thresholds (TRUE) or category-specific effects (FALSE). The default, NULL, uses the engine default. Available for specific engines only.
<code>predictor_prop</code>	The maximum proportion of original predictors that can have <i>non-zero</i> coefficients for each PLS component (via regularization). This value is used for all PLS components for X.
<code>num_comp</code>	The number of PLS components to retain.
<code>dist</code>	A character string for the probability distribution of the outcome. The default is "weibull".
<code>cost</code>	A positive number for the cost of predicting a sample within or on the wrong side of the margin
<code>margin</code>	A positive number for the epsilon in the SVM insensitive loss function (regression only)
<code>degree</code>	A positive number for polynomial degree.
<code>scale_factor</code>	A positive number for the polynomial scaling factor.
<code>rbf_sigma</code>	A positive number for radial basis function.
<code>num_embedding</code>	An integer for the dimensionality of the embedding space for features.
<code>hidden_activations</code>	A character vector denoting the activation functions for the hidden layers.

num_attn_feat	An integer for the number of attention features.
num_attn_heads	An integer for the number of attention heads in the multi-head attention mechanism.
num_attn_blocks	An integer for the number of sequential attention blocks.
dropout_attn	A number between 0 (inclusive) and 1 denoting the proportion of attention weights set to zero during model training.
dropout_embedding	A number between 0 (inclusive) and 1 denoting the proportion of embedding values set to zero during model training.
rate_schedule	A character string for the learning rate schedule.
momentum	A number for the momentum parameter in optimizers that use it.
batch_size	An integer for the number of training instances in each batch.
class_weights	Numeric class weights for imbalanced data (classification only).
num_estimators	An integer for the ensemble size. Default is 8L.
softmax_temperature	An adjustment factor that is a divisor in the exponents of the softmax function (see brulee::brulee_tab_icl()). Defaults to 0.9.
balance_probabilities	A logical to adjust the prior probabilities in cases where there is a class imbalance. Default is FALSE. Classification only.
average_before_softmax	A logical. For cases where num_estimators > 1, should the average be done before using the softmax function or after? Default is FALSE.
bottleneck_units	The number of embeddings that are produced by batch normalization.
residual_at	An integer vector with the layer number should use a residual connection (i.e., skip layer).
penalty_type	A string for the regularization norm: "L1" (default in brulee) or "L2". L1 is recommended by the original paper.
penalty_average	A positive numeric value for the target geometric mean of the per-weight regularization coefficients, on the natural scale. Best tuned on the log10 scale via dials::penalty_average() .
step_rate	A positive numeric value for the step size used to update the per-weight regularization coefficients, on the natural scale. Best tuned on the log10 scale via dials::step_rate() .
attention_type	A character string for the type of attention to use. Options are "column" (SAINT-s), "row" (SAINT-i), or "both" (full SAINT).
dropout_hidden	A number between 0 (inclusive) and 1 denoting the proportion of values in the feed-forward layers set to zero during training.
dropout_last	A number between 0 (inclusive) and 1 denoting the proportion of values set to zero between the last hidden layer and the output head.
target_token	A logical for whether to use a learnable target token (CLS-like embedding) to aggregate information for prediction. When TRUE, the model appends a special target token to the input that attends to all features via the attention mechanism.

Value

An updated model specification.

Examples

```
# -----

model <- C5_rules(trees = 10, min_n = 2)
model
update(model, trees = 1)
update(model, trees = 1, fresh = TRUE)

# -----

model <- cubist_rules(committees = 10, neighbors = 2)
model
update(model, committees = 1)
update(model, committees = 1, fresh = TRUE)

model <- pls(predictor_prop = 0.1)
model
update(model, predictor_prop = 1)
update(model, predictor_prop = 1, fresh = TRUE)

# -----

model <- rule_fit(trees = 10, min_n = 2)
model
update(model, trees = 1)
update(model, trees = 1, fresh = TRUE)

model <- boost_tree(mtry = 10, min_n = 3)
model
update(model, mtry = 1)
update(model, mtry = 1, fresh = TRUE)

param_values <- tibble::tibble(mtry = 10, tree_depth = 5)

model |> update(param_values)
model |> update(param_values, mtry = 3)

param_values$verbose <- 0
# Fails due to engine argument
# model |> update(param_values)

model <- linear_reg(penalty = 10, mixture = 0.1)
```

```
model  
update(model, penalty = 1)  
update(model, penalty = 1, fresh = TRUE)
```

Index

.cols (descriptors), 26
.dat (descriptors), 26
.extract_surv_status, 3
.extract_surv_time, 4
.facts (descriptors), 26
.get_prediction_column_names, 4
.get_prediction_column_names(), 4
.lvls (descriptors), 26
.obs (descriptors), 26
.preds (descriptors), 26
.x (descriptors), 26
.y (descriptors), 26
?extract_spec_parsnip, 52
?model_fit, 52

add_rowindex, 5
aorsf, 66
augment.model_fit, 6
auto_ml, 8
autoplot.glmnet (autoplot.model_fit), 9
autoplot.glmnet(), 9
autoplot.model_fit, 9

bag_mars, 10
bag_mlp, 11
bag_tree, 12
bart, 13
boost_tree, 16
boost_tree(), 52
brulee, 41, 42, 48, 56, 79, 81, 83, 85, 87, 89
brulee::brulee_auto_int(), 81
brulee::brulee_chronos(), 82
brulee::brulee_resnet(), 87
brulee::brulee_rln(), 88
brulee::brulee_saint(), 91
brulee::brulee_tab_icl(), 83, 84, 105
brulee_two_layer, 48

C5.0, 12, 16, 18, 24
C5.0_train(), 17

C50::C5.0(), 19
C50::C5.0Control(), 19
C5_rules, 18
case_weights, 19
case_weights_allowed, 20
catboost, 16
cforest_train (ctree_train), 21
control_parsnip, 20
control_parsnip(), 35, 36
ctree_train, 21
Cubist, 22
Cubist::cubist(), 24
Cubist::cubistControl(), 24
cubist_rules, 22

dbarts, 14
decision_tree, 24
descriptors, 26
dials::penalty_average(), 88, 105
dials::step_rate(), 88, 105
discrim_flexible, 28
discrim_linear, 29
discrim_linear(), 32
discrim_quad, 30
discrim_quad(), 32
discrim_regularized, 31

earth, 10, 28, 44
extract_parsnip, 33
extract_fit_engine.model_fit
 (extract_parsnip), 33
extract_fit_time.model_fit
 (extract_parsnip), 33
extract_parameter_dials.model_spec
 (extract_parsnip), 33
extract_parameter_set_dials.model_spec
 (extract_parsnip), 33
extract_spec_parsnip.model_fit
 (extract_parsnip), 33

- fit, 68, 69
- fit(), 8–13, 15, 17–19, 24–26, 28–32, 37, 38, 42–45, 49, 52, 57–60, 63–65, 67, 71, 75–79, 81, 82, 84, 85, 87, 88, 91
- fit.model_spec, 34
- fit.model_spec(), 6, 20, 33, 53, 62
- fit_xy, 52
- fit_xy(), 19, 75
- fit_xy.model_spec(fit.model_spec), 34
- fit_xy.model_spec(), 6, 53
- frequency_weights(), 19

- gee, 41, 42, 64
- gen_additive_mod, 37
- generalized additive models, 50
- get_encoding(), 36
- ggrepel::geom_label_repel(), 9
- glance.model_fit, 38
- glm, 41, 42, 64
- glm_grouped, 39
- glm_grouped(), 39
- glmer, 41, 42, 64
- glmnet, 41, 42, 56, 64
- glis, 41
- grf, 66

- h2o, 8, 16, 41, 43, 48, 56, 58, 64, 66, 70
- hardhat::frequency_weights(), 19, 35
- hardhat::importance_weights(), 19, 35
- hardhat::is_case_weights(), 35
- hurdle, 64

- importance_weights(), 19
- injection operator, 8, 10, 12, 13, 15, 17, 18, 24, 25, 28, 30–32, 37, 42, 44, 45, 49, 57, 58, 60, 63, 65, 67, 71, 76, 77, 79

- keras, 41, 43, 48, 56
- keras3, 41, 43, 48, 56
- kernlab, 75, 76, 78
- kknn, 59
- klaR, 31, 58

- LiblineaR, 43, 75
- lightgbm, 16
- linear_reg, 40
- linear_reg(), 33, 52
- lm, 41
- lme, 41
- lmer, 41
- logistic_reg, 42
- logistic_reg(), 42, 57

- mars, 44
- MASS, 29, 30
- max_mtry_formula, 46
- maybe_data_frame(maybe_matrix), 46
- maybe_matrix, 46
- mboost, 16
- mda, 29
- mgcv, 37
- min_cols, 47
- min_rows(min_cols), 47
- mixOmics, 63
- mlp, 48
- model_fit, 6, 55
- model_specification, 20, 33, 68, 69, 72, 73, 92, 102
- model_specifications, 49
- model_fit, 49
- model_formula, 50
- model_spec, 50, 52
- multi_predict, 54
- multinom_reg, 56
- multinom_reg(), 44

- naive_Bayes, 58
- naivebayes, 58
- nearest_neighbor, 59
- nnet, 11, 48, 56
- null_model, 60

- ordinalForest, 66

- parsnip_addin, 62
- parsnip_update(update.bag_mars), 93
- partykit, 25, 66
- partykit::cforest(), 21, 22
- partykit::ctree(), 21, 22
- partykit::ctree_control(), 21
- pls, 63
- poisson_reg, 64
- poisson_reg(), 74
- predict, 49
- predict(), 75, 91
- predict.model_fit(), 6
- qrnn, 48

- quantreg, [41](#)
- quosures, [8](#), [10](#), [12](#), [13](#), [15](#), [17](#), [18](#), [24](#), [25](#), [28](#),
[30–32](#), [37](#), [42](#), [44](#), [45](#), [49](#), [57](#), [58](#), [60](#),
[63](#), [65](#), [67](#), [71](#), [76](#), [77](#), [79](#)
- rand_forest, [65](#)
- randomForest, [66](#)
- randomForestSRC, [66](#)
- ranger, [66](#)
- repair_call, [67](#)
- req_pkgs, [68](#)
- required_pkgs.model_fit
 (required_pkgs.model_spec), [69](#)
- required_pkgs.model_spec, [69](#)
- rpart, [12](#), [24](#)
- rpartScore, [25](#)
- rule_fit, [70](#)
- sda, [29](#)
- set_args, [72](#)
- set_engine, [73](#)
- set_engine(), [8–13](#), [15](#), [17–19](#), [24–26](#),
[28–32](#), [35–38](#), [42–45](#), [49](#), [57–60](#),
[63–65](#), [67](#), [71](#), [76–79](#), [81](#), [82](#), [84](#), [85](#),
[87](#), [88](#), [91](#)
- set_mode(set_args), [72](#)
- set_mode(), [82](#)
- show_engines, [74](#)
- show_engines(), [73](#)
- spark, [16](#), [25](#), [41](#), [43](#), [56](#), [66](#)
- sparse_data, [75](#)
- sparsediscrim, [29](#), [30](#)
- stan, [41](#), [43](#), [64](#)
- stan_glmer, [41](#), [43](#), [64](#)
- stats::contr.treatment(), [35](#)
- stats::glm(), [39](#)
- survival::Surv(), [3](#), [4](#), [6](#)
- svm_linear, [75](#)
- svm_poly, [76](#)
- svm_rbf, [78](#)
- tabpfn, [84](#)
- tabpfn::tab_pfn(), [84](#), [85](#)
- tabular_auto_int, [79](#)
- tabular_chronos, [81](#)
- tabular_icl, [83](#)
- tabular_pfn, [84](#)
- tabular_resnet, [85](#)
- tabular_rln, [87](#)
- tabular_saint, [89](#)
- tibble::tibble(), [91](#)
- tidy.model_fit, [91](#)
- tidyr::unnest(), [6](#)
- translate, [92](#)
- update(), [9](#), [11–13](#), [15](#), [17](#), [19](#), [24](#), [26](#), [29–32](#),
[38](#), [42](#), [44](#), [45](#), [49](#), [57](#), [59](#), [60](#), [64](#), [65](#),
[67](#), [71](#), [76](#), [77](#), [79](#), [81](#), [82](#), [84](#), [85](#), [87](#),
[88](#), [91](#)
- update.bag_mars, [93](#)
- update.bag_mlp(update.bag_mars), [93](#)
- update.bag_tree(update.bag_mars), [93](#)
- update.bart(update.bag_mars), [93](#)
- update.boost_tree(update.bag_mars), [93](#)
- update.C5_rules(update.bag_mars), [93](#)
- update.cubist_rules(update.bag_mars),
[93](#)
- update.decision_tree(update.bag_mars),
[93](#)
- update.discrim_flexible
 (update.bag_mars), [93](#)
- update.discrim_linear
 (update.bag_mars), [93](#)
- update.discrim_quad(update.bag_mars),
[93](#)
- update.discrim_regularized
 (update.bag_mars), [93](#)
- update.gen_additive_mod
 (update.bag_mars), [93](#)
- update.linear_reg(update.bag_mars), [93](#)
- update.logistic_reg(update.bag_mars),
[93](#)
- update.mars(update.bag_mars), [93](#)
- update.mlp(update.bag_mars), [93](#)
- update.multinom_reg(update.bag_mars),
[93](#)
- update.naive_Bayes(update.bag_mars), [93](#)
- update.nearest_neighbor
 (update.bag_mars), [93](#)
- update.ordinal_reg(update.bag_mars), [93](#)
- update.pls(update.bag_mars), [93](#)
- update.poisson_reg(update.bag_mars), [93](#)
- update.proportional_hazards
 (update.bag_mars), [93](#)
- update.rand_forest(update.bag_mars), [93](#)
- update.rule_fit(update.bag_mars), [93](#)
- update.surv_reg(update.bag_mars), [93](#)

`update.survival_reg(update.bag_mars),`
 93
`update.svm_linear(update.bag_mars),` 93
`update.svm_poly(update.bag_mars),` 93
`update.svm_rbf(update.bag_mars),` 93
`update.tabular_auto_int`
 `(update.bag_mars),` 93
`update.tabular_chronos`
 `(update.bag_mars),` 93
`update.tabular_icl(update.bag_mars),` 93
`update.tabular_pfn(update.bag_mars),` 93
`update.tabular_resnet`
 `(update.bag_mars),` 93
`update.tabular_rln(update.bag_mars),` 93
`update.tabular_saint(update.bag_mars),`
 93

`vgam,` 37

with the `parsnip` `mgcv` engine, 51

`xgb_train()`, 17
`xgboost,` 16
`xrf,` 70
`xrf::xrf.formula()`, 71

`zeroinfl,` 64