

# Package ‘parsnip’

March 16, 2022

**Title** A Common API to Modeling and Analysis Functions

**Version** 0.2.1

**Maintainer** Max Kuhn <max@rstudio.com>

**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', etc).

**License** MIT + file LICENSE

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

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

**Depends** R (>= 2.10)

**Imports** cli,  
dplyr (>= 0.8.0.1),  
generics (>= 0.1.0.9000),  
ggplot2,  
globals,  
glue,  
hardhat (>= 0.1.6.9001),  
lifecycle,  
magrittr,  
prettyunits,  
purrr,  
rlang (>= 0.3.1),  
stats,  
tibble (>= 2.1.1),  
tidyr (>= 1.0.0),  
utils,  
vctrs (>= 0.2.0),  
withr

**Suggests** C50,  
covr,  
dials (>= 0.0.10.9001),  
earth,  
tensorflow,  
ggrepel,

keras,  
 kernlab,  
 kknn,  
 knitr,  
 LiblineaR,  
 MASS,  
 Matrix,  
 mgcv,  
 modeldata,  
 nlme,  
 ranger ( $\geq 0.12.0$ ),  
 remotes,  
 rmarkdown,  
 rpart,  
 sparklyr ( $\geq 1.0.0$ ),  
 survival,  
 testthat,  
 xgboost ( $\geq 1.5.0.1$ )

**VignetteBuilder** knitr

**ByteCompile** true

**Config/Needs/website** C50,

dbarts,  
 earth,  
 glmnet,  
 keras,  
 kernlab,  
 kknn,  
 LiblineaR,  
 mgcv,  
 nnet,  
 parsnip,  
 randomForest,  
 ranger,  
 rpart,  
 rstanarm,  
 tidymodels/tidymodels,  
 tidyverse/tidytemplate,  
 rstudio/reticulate,  
 xgboost

**Config/rcmdcheck/ignore-inconsequential-notes** true

**Encoding** UTF-8

**LazyData** true

**Roxygen** list(markdown = TRUE)

**RoxygenNote** 7.1.2

## R topics documented:

add_rowindex . . . . .	4
augment.model_fit . . . . .	4

autoplot.model_fit . . . . .	5
bag_mars . . . . .	6
bag_tree . . . . .	7
bart . . . . .	8
boost_tree . . . . .	10
C5_rules . . . . .	12
control_parsnip . . . . .	13
contr_one_hot . . . . .	13
cubist_rules . . . . .	15
decision_tree . . . . .	17
descriptors . . . . .	18
discrim_flexible . . . . .	19
discrim_linear . . . . .	20
discrim_quad . . . . .	22
discrim_regularized . . . . .	23
extract_parsnip . . . . .	24
fit.model_spec . . . . .	25
gen_additive_mod . . . . .	27
glance.model_fit . . . . .	28
linear_reg . . . . .	29
logistic_reg . . . . .	30
mars . . . . .	32
maybe_matrix . . . . .	33
min_cols . . . . .	33
mlp . . . . .	34
model_fit . . . . .	35
model_spec . . . . .	36
multinom_reg . . . . .	38
multi_predict . . . . .	39
naive_Bayes . . . . .	41
nearest_neighbor . . . . .	42
null_model . . . . .	43
parsnip_addin . . . . .	44
pls . . . . .	44
poisson_reg . . . . .	45
rand_forest . . . . .	47
repair_call . . . . .	48
required_pkgs.model_spec . . . . .	49
req_pkgs . . . . .	50
rule_fit . . . . .	51
set_args . . . . .	52
set_engine . . . . .	53
show_engines . . . . .	54
svm_linear . . . . .	55
svm_poly . . . . .	56
svm_rbf . . . . .	57
tidy.model_fit . . . . .	58
translate . . . . .	59
update.bag_mars . . . . .	60

---

add_rowindex	<i>Add a column of row numbers to a data frame</i>
--------------	----------------------------------------------------

---

**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, ...)
```

**Arguments**

x	A <code>model_fit</code> object produced by <code>fit.model_spec()</code> or <code>fit_xy.model_spec()</code>
.	.
new_data	A data frame or matrix.
...	Not currently used.

**Details**

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

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.

**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])

```

---

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 <code>autoplot.glmnet()</code> , options to pass to <code>ggrepel::geom_label_repel()</code> . 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 <b>ggrepel</b> 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	<i>Ensembles of MARS models</i>
----------	---------------------------------

---

### 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](#)<sup>1,2</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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**

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.

**References**

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

**See Also**

[fit\(\)](#), [set\\_engine\(\)](#), [update\(\)](#), [earth 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](#)<sup>1,2</sup>
- [C5.0](#)<sup>2</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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", or "classification".
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.
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.

## 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](#)

---

bart

*Bayesian additive regression trees (BART)*

---

## 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](#)<sup>1</sup>

<sup>1</sup> 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  $prior = a * (1 + d)^{-b}$  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.

**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](#)<sup>1</sup>
- [C5.0](#)
- [mboost](#)<sup>2</sup>
- [spark](#)

<sup>1</sup> The default engine. <sup>2</sup> Requires a `parsnip` extension package for censored regression.

More information on how **parsonip** 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", or "classification".
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).
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.

### 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](#), [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](#)<sup>12</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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

<code>mode</code>	A single character string for the type of model. The only possible value for this model is "classification".
<code>trees</code>	A non-negative integer (no greater than 100) for the number of members of the ensemble.
<code>min_n</code>	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.
<code>engine</code>	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.

## 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()
```

---

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)
```

---

contr_one_hot	<i>Contrast function for one-hot encodings</i>
---------------	------------------------------------------------

---

**Description**

This contrast function produces a model matrix with indicator columns for each level of each factor.

**Usage**

```
contr_one_hot(n, contrasts = TRUE, sparse = FALSE)
```

## Arguments

n	A vector of character factor levels or the number of unique levels.
contrasts	This argument is for backwards compatibility and only the default of TRUE is supported.
sparse	This argument is for backwards compatibility and only the default of FALSE is supported.

## Details

By default, `model.matrix()` generates binary indicator variables for factor predictors. When the formula does not remove an intercept, an incomplete set of indicators are created; no indicator is made for the first level of the factor.

For example, `species` and `island` both have three levels but `model.matrix()` creates two indicator variables for each:

```
library(dplyr)
library(modeldata)
data(penguins)

levels(penguins$species)

## [1] "Adelie" "Chinstrap" "Gentoo"

levels(penguins$island)

## [1] "Biscoe" "Dream" "Torgersen"

model.matrix(~ species + island, data = penguins) %>%
  colnames()

## [1] "(Intercept)" "speciesChinstrap" "speciesGentoo" "islandDream"
## [5] "islandTorgersen"
```

For a formula with no intercept, the first factor is expanded to indicators for *all* factor levels but all other factors are expanded to all but one (as above):

```
model.matrix(~ 0 + species + island, data = penguins) %>%
  colnames()

## [1] "speciesAdelie" "speciesChinstrap" "speciesGentoo" "islandDream"
## [5] "islandTorgersen"
```

For inference, this hybrid encoding can be problematic.

To generate all indicators, use this contrast:

```
# Switch out the contrast method
old_contr <- options("contrasts")$contrasts
new_contr <- old_contr
new_contr["unordered"] <- "contr_one_hot"
options(contrasts = new_contr)

model.matrix(~ species + island, data = penguins) %>%
  colnames()
```

```
## [1] "(Intercept)"      "speciesAdelie"      "speciesChinstrap"  "speciesGentoo"
## [5] "islandBiscoe"      "islandDream"       "islandTorgersen"

options(contrasts = old_contr)
```

Removing the intercept here does not affect the factor encodings.

### Value

A diagonal matrix that is n-by-n.

---

cubist_rules	<i>Cubist rule-based regression models</i>
--------------	--------------------------------------------

---

### 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](#)<sup>1,2</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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

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

## 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	<i>Decision trees</i>
---------------	-----------------------

---

## 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](#)<sup>1,2</sup>
- [C5.0](#)
- [party](#)<sup>2</sup>
- [spark](#)

<sup>1</sup> The default engine. <sup>2</sup> Requires a `parsnip` extension package for censored 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

<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_complexity</code>	A positive number for the the cost/complexity parameter (a.k.a. Cp) used by CART models (specific engines only).
<code>tree_depth</code>	An integer for maximum depth of the tree.
<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.

## 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](#), [party 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](#)<sup>1,2</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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.

## 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

*Linear discriminant analysis*

---

## 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](#)<sup>12</sup>
- [mda](#)<sup>2</sup>
- [sda](#)<sup>2</sup>

- [sparsediscrim](#)<sup>2</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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.

## 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`*Quadratic discriminant analysis*

---

## 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](#)<sup>1,2</sup>
- [sparsediscrim](#)<sup>2</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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

<code>mode</code>	A single character string for the type of model. The only possible value for this model is "classification".
<code>regularization_method</code>	A character string for the type of regularized estimation. Possible values are: "diagonal", "shrink_cov", and "shrink_mean" (sparsediscrim engine only).
<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.

## 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](#)<sup>1,2</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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

<code>mode</code>	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
<code>frac_common_cov</code> , <code>frac_identity</code>	Numeric values between zero and one.
<code>engine</code>	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.

## 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.

## 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, ...)
```

## 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.



**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, control = control_parsnip(), ...)

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

**Arguments**

object	An object of class <code>model_spec</code> that has a chosen engine (via <code>set_engine()</code> ).
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.
control	A named list with elements <code>verbosity</code> and <code>catch</code> . See <code>control_parsnip()</code> .
...	Not currently used; values passed here will be ignored. Other options required to fit the model should be passed using <code>set_engine()</code> .
x	A matrix, sparse matrix, or data frame of predictors. Only some models have support for sparse matrix input. See <code>parsnip::get_encoding()</code> 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.

**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.
- `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

*Generalized additive models (GAMs)*

---

**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](#)<sup>1</sup>

<sup>1</sup> The default engine.

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.

**References**

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

**See Also**

`fit()`, `set_engine()`, `update()`, `mgcv 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

---

linear_reg	<i>Linear regression</i>
------------	--------------------------

---

**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](#)<sup>1</sup>
- [brulee](#)
- [gee](#)<sup>2</sup>
- [glm](#)
- [glmnet](#)
- [gls](#)<sup>2</sup>
- [keras](#)
- [lme](#)<sup>2</sup>
- [lmer](#)<sup>2</sup>
- [spark](#)
- [stan](#)
- [stan\\_glmer](#)<sup>2</sup>

<sup>1</sup> The default engine. <sup>2</sup> Requires a `parsnip` extension package.

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**

mode	A single character string for the type of model. The only possible value for this model is "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 "lm".
penalty	A non-negative number representing the total amount of regularization (specific engines only).
mixture	A number between zero and one (inclusive) that is the proportion of L1 regularization (i.e. lasso) in the model. When <code>mixture = 1</code> , it is a pure lasso model while <code>mixture = 0</code> indicates that ridge regression is being used (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.

## 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](#), [glmnet engine details](#), [gls engine details](#), [keras engine details](#), [lme engine details](#), [lmer 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](#)<sup>1</sup>
- [brulee](#)
- [gee](#)<sup>2</sup>
- [glmer](#)<sup>2</sup>
- [glmnet](#)
- [keras](#)
- [LiblineaR](#)
- [spark](#)
- [stan](#)
- [stan\\_glmer](#)<sup>2</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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) that is the proportion of L1 regularization (i.e. lasso) in the model. When mixture = 1, it is a pure lasso model while mixture = 0 indicates that ridge regression is being used. (specific engines only). For Liblinear models, mixture must be exactly 0 or 1 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.

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](#), [keras 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](#)<sup>1</sup>

<sup>1</sup> 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

<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>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.

## 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.

## 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)
```

---

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	<i>Execution-time data dimension checks</i>
----------	---------------------------------------------

---

**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("kknn") %>%
  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](#)<sup>1</sup>
- [brulee](#)
- [keras](#)

<sup>1</sup> The default engine.

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**

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 for the number of units in the hidden model.
penalty	A non-negative numeric value for the amount of 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 single character string denoting the type of relationship between the original predictors and the hidden unit layer. 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 are: "linear", "softmax", "relu", and "elu"
learn_rate	A number for the rate at which the boosting algorithm adapts from iteration-to-iteration (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.

**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](#), [keras engine details](#)

**Examples**

```
show_engines("mlp")

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

---

 model\_fit

*Model Fit Object Information*


---

**Description**

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

## Details

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 = ifelse(.obs() < 500, TRUE, FALSE))
spec_obj

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

nrow(fit_obj$fit$x)
```

---

model\_spec

*Model Specification Information*

---

## Description

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

## Details

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 tune() with the use of the [tune package](#). 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://tidyeval.tidyverse.org>.

---

multinom\_reg

*Multinomial regression*

---

## 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](#)<sup>1</sup>
- [brulee](#)
- [glmnet](#)
- [keras](#)
- [spark](#)

<sup>1</sup>The default engine.

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) that is the proportion of L1 regularization (i.e. lasso) in the model. When mixture = 1, it is a pure lasso model while mixture = 0 indicates that ridge regression is being used. (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.

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](#), [keras engine details](#), [spark engine details](#)

**Examples**

```
show_engines("multinom_reg")

multinom_reg()
```

---

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 '`_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 '`_multnet`'
multi_predict(object, new_data, type = NULL, penalty = NULL, ...)

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

```

**Arguments**

object	A <code>model_fit</code> object.
...	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 <code>NULL</code> . Possible values are "numeric", "class", "prob", "conf_int", "pred_int", "quantile", or "raw". When <code>NULL</code> , <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).



## 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](#)<sup>1,2</sup>
- [naivebayes](#)<sup>2</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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

<code>mode</code>	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
<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>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.

## References

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

## See Also

[fit\(\)](#), [set\\_engine\(\)](#), [update\(\)](#), [klaR engine details](#), [naivebayes engine details](#)

---

nearest\_neighbor      *K-nearest neighbors*

---

## 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](#)<sup>1</sup>

<sup>1</sup> 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 <b>kknn</b> , 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.

## 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)
```

---

null\_model

*Null model*


---

**Description**

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

**Usage**

```
null_model(mode = "classification")
```

**Arguments**

`mode` A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".

**Details**

The model can be created using the `fit()` function using the following *engines*:

- **R:** "parsnip"

**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()

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

```

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

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

```

**See Also**

[fit.model\\_spec\(\)](#)

**Examples**

```

null_model(mode = "regression")

```

---

parsnip_addin	<i>Start an RStudio Addin that can write model specifications</i>
---------------	-------------------------------------------------------------------

---

**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	<i>Partial least squares (PLS)</i>
-----	------------------------------------

---

**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](#)<sup>1,2</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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**

mode	A single character string for the prediction outcome mode. Possible values for this model are "unknown", "regression", or "classification".
predictor_prop	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.
num_comp	The number of PLS components to retain.
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.

**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](#)<sup>1,2</sup>
- [gee](#)<sup>2</sup>
- [glmer](#)<sup>2</sup>
- [glmnet](#)<sup>2</sup>
- [hurdle](#)<sup>2</sup>

- [stan](#)<sup>2</sup>
- [stan\\_glmer](#)<sup>2</sup>
- [zeroinfl](#)<sup>2</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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

<code>mode</code>	A single character string for the type of model. The only possible value for this model is "regression".
<code>penalty</code>	A non-negative number representing the total amount of regularization (glmnet only).
<code>mixture</code>	A number between zero and one (inclusive) that is the proportion of L1 regularization (i.e. lasso) in the model. When <code>mixture = 1</code> , it is a pure lasso model while <code>mixture = 0</code> indicates that ridge regression is being used. (glmnet and spark only).
<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.

## 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](#), [hurdle engine details](#), [stan engine details](#), [stan\\_glmer engine details](#), [zeroinfl engine details](#)

---

rand_forest	<i>Random forest</i>
-------------	----------------------

---

## 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](#)<sup>1</sup>
- [party](#)<sup>2</sup>
- [randomForest](#)
- [spark](#)

<sup>1</sup> The default engine. <sup>2</sup> Requires a `parsnip` extension package for censored 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", or "classification".
<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.

## References

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

## See Also

[fit\(\)](#), [set\\_engine\(\)](#), [update\(\)](#), [ranger engine details](#), [party engine details](#), [randomForest 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
```

---

```
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

x	A model specification or fit.
infra	Should parsnip itself be included in the result?
...	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()
```

---

`req_pkgs`*Determine required packages for a model*

---

## Description

**[Deprecated]**

## Usage

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

## Arguments

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

## Details

For a model specification, the engine must be set. The list produced by `req_pkgs()` does not include the `parsnip` package while `required_pkgs()` does.

## Value

A character string of package names (if any).

## Examples

```
should_fail <- try(req_pkgs(linear_reg()), silent = TRUE)  
should_fail  
  
linear_reg() %>%  
  set_engine("glmnet") %>%  
  req_pkgs()  
  
linear_reg() %>%  
  set_engine("lm") %>%  
  fit(mpg ~ ., data = mtcars) %>%  
  req_pkgs()
```

---

rule_fit	<i>RuleFit models</i>
----------	-----------------------

---

### 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](#)<sup>12</sup>

<sup>1</sup> The default engine. <sup>2</sup> 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,
  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).
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 <code>xgboost</code> , the sampling is done at each iteration while <code>C5.0</code> samples once during training.
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.

## 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](#)

## Examples

```
show_engines("rule_fit")

rule_fit()
```

---

set\_args

*Change elements of a model specification*

---

## 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)
```

## Arguments

object	A model specification.
...	One or more named model arguments.
mode	A character string for the model type (e.g. "classification" or "regression")

## 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")
```

---

set_engine	<i>Declare a computational engine and specific arguments</i>
------------	--------------------------------------------------------------

---

**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. Set these in `set_engine()`, 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")
```

---

`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](#)<sup>1</sup>
- [kernlab](#)

<sup>1</sup>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.

## 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<sup>1</sup>

<sup>1</sup> 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.

**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)
```

svm\_rbf

*Radial basis function support vector machines***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](#)<sup>1</sup>

<sup>1</sup>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.

## 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)
```

---

tidy.model_fit	<i>Turn a parsnip model object into a tidy tibble</i>
----------------	-------------------------------------------------------

---

## 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 <code>tibble::tibble()</code> .
...	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

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

## 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](https://www.tidymodels.org/start/tuning/). 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	<i>Updating a model specification</i>
-----------------	---------------------------------------

---

## 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_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 '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,
  ...
)

```

### Arguments

object	A model specification.
parameters	A 1-row tibble or named list with <i>main</i> parameters to update. Use <b>either</b> parameters <b>or</b> 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().
cost_complexity	A positive number for the the cost/complexity parameter (a.k.a. $C_p$ ) 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 $\pm 3$ (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).
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.
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).

frac_common_cov	Numeric values between zero and one.
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) that is the proportion of L1 regularization (i.e. lasso) in the model. When mixture = 1, it is a pure lasso model while mixture = 0 indicates that ridge regression is being used (specific engines only).
hidden_units	An integer for the number of units in the hidden model.
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 single character string denoting the type of relationship between the original predictors and the hidden unit layer. 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 are: "linear", "softmax", "relu", and "elu"
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.
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.
predictor_prop	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.
num_comp	The number of PLS components to retain.
dist	A character string for the probability distribution of the outcome. The default is "weibull".
cost	A positive number for the cost of predicting a sample within or on the wrong side of the margin
margin	A positive number for the epsilon in the SVM insensitive loss function (regression only)
degree	A positive number for polynomial degree.
scale_factor	A positive number for the polynomial scaling factor.
rbf_sigma	A positive number for radial basis function.

**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), 18  
.dat (descriptors), 18  
.facts (descriptors), 18  
.lvl (descriptors), 18  
.obs (descriptors), 18  
.preds (descriptors), 18  
.x (descriptors), 18  
.y (descriptors), 18

add\_rowindex, 4  
augment.model\_fit, 4  
autoplot.glmnet (autoplot.model\_fit), 5  
autoplot.glmnet(), 6  
autoplot.model\_fit, 5

bag\_mars, 6  
bag\_tree, 7  
bart, 8  
boost\_tree, 10  
brulee, 29, 30, 34, 38  
brulee engine details, 30, 31, 35, 39

C5.0, 7, 10, 12, 17  
C5.0 engine details, 8, 11, 12, 18  
C5.0\_train(), 11  
C50::C5.0(), 12  
C50::C5.0Control(), 12  
C5\_rules, 12  
contr\_one\_hot, 13  
control\_parsnip, 13  
control\_parsnip(), 26, 27  
Cubist, 15  
Cubist engine details, 16  
Cubist::cubist(), 16  
Cubist::cubistControl(), 16  
cubist\_rules, 15

dbarts, 8  
dbarts engine details, 9  
decision\_tree, 17  
descriptors, 18  
discrim\_flexible, 19  
discrim\_linear, 20  
discrim\_linear(), 23

discrim\_quad, 22  
discrim\_quad(), 23  
discrim\_regularized, 23

earth, 6, 19, 32  
earth engine details, 7, 20, 32  
extract\_parsnip, 24  
extract\_fit\_engine.model\_fit  
  (extract\_parsnip), 24  
extract\_parameter\_dials.model\_spec  
  (extract\_parsnip), 24  
extract\_parameter\_set\_dials.model\_spec  
  (extract\_parsnip), 24  
extract\_spec\_parsnip.model\_fit  
  (extract\_parsnip), 24

fit(), 7–9, 11, 12, 16–18, 20–24, 28, 30–32,  
  35, 39, 41–43, 45–48, 52, 55–58  
fit.model\_spec, 25  
fit.model\_spec(), 4, 13, 37, 44  
fit\_xy.model\_spec (fit.model\_spec), 25  
fit\_xy.model\_spec(), 4, 37

gee, 29, 30, 45  
gee engine details, 30, 31, 46  
gen\_additive\_mod, 27  
ggrepel::geom\_label\_repel(), 6  
glance.model\_fit, 28  
glm, 29, 30, 45  
glm engine details, 30, 31, 46  
glmer, 30, 45  
glmer engine details, 31, 46  
glmnet, 29, 30, 38, 45  
glmnet engine details, 30, 31, 39, 46  
gls, 29  
gls engine details, 30

hurdle, 45  
hurdle engine details, 46

keras, 29, 30, 34, 38  
keras engine details, 30, 31, 35, 39  
kernlab, 55–57  
kernlab engine details, 55, 57, 58  
kknn, 42

- kknn engine details, [43](#)
- klaR, [23](#), [41](#)
- klaR engine details, [24](#), [41](#)
- LiblineaR, [30](#), [55](#)
- LiblineaR engine details, [31](#), [55](#)
- linear\_reg, [29](#)
- linear\_reg(), [24](#)
- lm, [29](#)
- lm engine details, [30](#)
- lme, [29](#)
- lme engine details, [30](#)
- lmer, [29](#)
- lmer engine details, [30](#)
- logistic\_reg, [30](#)
- logistic\_reg(), [30](#), [39](#)
- mars, [32](#)
- MASS, [20](#), [22](#)
- MASS engine details, [21](#), [22](#)
- maybe\_data\_frame (maybe\_matrix), [33](#)
- maybe\_matrix, [33](#)
- mboost, [10](#)
- mboost engine details, [11](#)
- mda, [20](#)
- mda engine details, [21](#)
- mgcv, [27](#)
- mgcv engine details, [28](#)
- min\_cols, [33](#)
- min\_rows (min\_cols), [33](#)
- mixOmics, [44](#)
- mixOmics engine details, [45](#)
- mlp, [34](#)
- model\_fit, [35](#)
- model\_spec, [36](#), [36](#)
- multi\_predict, [39](#)
- multinom\_reg, [38](#)
- multinom\_reg(), [31](#)
- naive\_Bayes, [41](#)
- naivebayes, [41](#)
- naivebayes engine details, [41](#)
- nearest\_neighbor, [42](#)
- nnet, [34](#), [38](#)
- nnet engine details, [35](#), [39](#)
- null\_model, [43](#)
- parsnip\_addin, [44](#)
- parsnip\_update (update.bag\_mars), [60](#)
- party, [17](#), [47](#)
- party engine details, [18](#), [48](#)
- pls, [44](#)
- poisson\_reg, [45](#)
- poisson\_reg(), [54](#)
- rand\_forest, [47](#)
- randomForest, [47](#)
- randomForest engine details, [48](#)
- ranger, [47](#)
- ranger engine details, [48](#)
- repair\_call, [48](#)
- req\_pkgs, [50](#)
- required\_pkgs.model\_fit  
(required\_pkgs.model\_spec), [49](#)
- required\_pkgs.model\_spec, [49](#)
- rpart, [7](#), [17](#)
- rpart engine details, [8](#), [18](#)
- rule\_fit, [51](#)
- sda, [20](#)
- sda engine details, [21](#)
- set\_args, [52](#)
- set\_engine, [53](#)
- set\_engine(), [7–9](#), [11](#), [12](#), [16–18](#), [20–24](#),  
[26–28](#), [30–32](#), [35](#), [39](#), [41–43](#), [45–48](#),  
[52](#), [55–58](#)
- set\_mode (set\_args), [52](#)
- show\_engines, [54](#)
- show\_engines(), [53](#)
- spark, [10](#), [17](#), [29](#), [30](#), [38](#), [47](#)
- spark engine details, [11](#), [18](#), [30](#), [31](#), [39](#), [48](#)
- sparsediscrim, [21](#), [22](#)
- sparsediscrim engine details, [21](#), [22](#)
- stan, [29](#), [30](#), [46](#)
- stan engine details, [30](#), [31](#), [46](#)
- stan\_glmer, [29](#), [30](#), [46](#)
- stan\_glmer engine details, [30](#), [31](#), [46](#)
- stats::contr.treatment(), [26](#)
- svm\_linear, [55](#)
- svm\_poly, [56](#)
- svm\_rbf, [57](#)
- tibble::tibble(), [58](#)
- tidy.model\_fit, [58](#)
- translate, [59](#)
- update(), [7–9](#), [11](#), [12](#), [16](#), [18](#), [20–22](#), [24](#), [28](#),  
[30–32](#), [35](#), [39](#), [41](#), [43](#), [45](#), [46](#), [48](#), [52](#),  
[55](#), [57](#), [58](#)
- update.bag\_mars, [60](#)
- update.bag\_tree (update.bag\_mars), [60](#)
- update.bart (update.bag\_mars), [60](#)
- update.boost\_tree (update.bag\_mars), [60](#)
- update.C5\_rules (update.bag\_mars), [60](#)
- update.cubist\_rules (update.bag\_mars),  
[60](#)

update.decision\_tree (update.bag\_mars),  
60  
update.discrim\_flexible  
(update.bag\_mars), 60  
update.discrim\_linear  
(update.bag\_mars), 60  
update.discrim\_quad (update.bag\_mars),  
60  
update.discrim\_regularized  
(update.bag\_mars), 60  
update.gen\_additive\_mod  
(update.bag\_mars), 60  
update.linear\_reg (update.bag\_mars), 60  
update.logistic\_reg (update.bag\_mars),  
60  
update.mars (update.bag\_mars), 60  
update.mlp (update.bag\_mars), 60  
update.multinom\_reg (update.bag\_mars),  
60  
update.naive\_Bayes (update.bag\_mars), 60  
update.nearest\_neighbor  
(update.bag\_mars), 60  
update.pls (update.bag\_mars), 60  
update.poisson\_reg (update.bag\_mars), 60  
update.proportional\_hazards  
(update.bag\_mars), 60  
update.rand\_forest (update.bag\_mars), 60  
update.rule\_fit (update.bag\_mars), 60  
update.surv\_reg (update.bag\_mars), 60  
update.survival\_reg (update.bag\_mars),  
60  
update.svm\_linear (update.bag\_mars), 60  
update.svm\_poly (update.bag\_mars), 60  
update.svm\_rbf (update.bag\_mars), 60  
  
xgb\_train(), 11  
xgboost, 10  
xgboost engine details, 11  
xrf, 51  
xrf engine details, 52  
xrf::xrf.formula(), 52  
  
zeroinfl, 46  
zeroinfl engine details, 46