

Package ‘parsnip’

November 4, 2018

Version 0.0.1

Title A Common API to Modeling and Analysis Functions

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).

Maintainer Max Kuhn <max@rstudio.com>

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

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

License GPL-2

Encoding UTF-8

LazyData true

ByteCompile true

VignetteBuilder knitr

Depends R ($i=$ 2.10)

Imports dplyr,
 rlang ($i=$ 0.3.0.1),
 purrr,
 utils,
 tibble,
 generics,
 glue,
 magrittr,
 stats,
 tidyr,
 globals

Roxygen list(markdown = TRUE)

RoxygenNote 6.1.0.9000

Suggests testthat,
 knitr,
 rmarkdown,
 survival,
 keras,
 xgboost,
 covr,
 sparklyr

R topics documented:

boost_tree	2
C5.0_train	5
decision_tree	5
descriptors	7
fit.model_spec	9
fit_control	11
keras_mlp	11
lending_club	12
linear_reg	13
logistic_reg	15
mars	17
mlp	19
model_fit	21
model_spec	22
multinom_reg	24
nearest_neighbor	26
predict.model_fit	27
rand_forest	29
rpart_train	31
set_args	32
set_engine	32
surv_reg	33
svm_poly	35
svm_rbf	36
translate	38
varying	39
varying_args	39
wa_churn	40
xgb_train	41
Index	42

boost_tree	<i>General Interface for Boosted Trees</i>
------------	--

Description

boost_tree is a way to generate a *specification* of a model before fitting and allows the model to be created using different packages in R or via Spark. The main arguments for the model are:

- `mtry`: The number of predictors that will be randomly sampled at each split when creating the tree models.
- `trees`: The number of trees contained in the ensemble.
- `min_n`: The minimum number of data points in a node that are required for the node to be split further.
- `tree_depth`: The maximum depth of the tree (i.e. number of splits).
- `learn_rate`: The rate at which the boosting algorithm adapts from iteration-to-iteration.
- `loss_reduction`: The reduction in the loss function required to split further.

- `sample_size`: The amount of data exposed to the fitting routine.

These arguments are converted to their specific names at the time that the model is fit. Other options and argument can be set using the `set_engine` function. If left to their defaults here (NULL), the values are taken from the underlying model functions. If parameters need to be modified, `update` can be used in lieu of recreating the object from scratch.

Usage

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

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

Arguments

<code>mode</code>	A single character string for the type of model. Possible values for this model are "unknown", "regression", or "classification".
<code>mtry</code>	An number for the number (or proportion) of predictors that will be randomly sampled at each split when creating the tree models (xgboost only).
<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.
<code>tree_depth</code>	An integer for the maximum depth of the tree (i.e. number of splits) (xgboost only).
<code>learn_rate</code>	A number for the rate at which the boosting algorithm adapts from iteration-to-iteration (xgboost only).
<code>loss_reduction</code>	A number for the reduction in the loss function required to split further (xgboost only).
<code>sample_size</code>	An number for the number (or proportion) of data that is exposed to the fitting routine. For xgboost, the sampling is done at at each iteration while C5.0 samples once during training.
<code>object</code>	A boosted tree model specification.
<code>fresh</code>	A logical for whether the arguments should be modified in-place of or replaced wholesale.
<code>...</code>	Not used for <code>update</code> .

Details

The data given to the function are not saved and are only used to determine the *mode* of the model. For `boost_tree`, the possible modes are "regression" and "classification".

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

- **R**: "xgboost", "C5.0"
- **Spark**: "spark"

Value

An updated model specification.

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:

xgboost classification

```
parsnip::xgb_train(x = missing_arg(), y = missing_arg(), nthread = 1,
  verbose = 0)
```

xgboost regression

```
parsnip::xgb_train(x = missing_arg(), y = missing_arg(), nthread = 1,
  verbose = 0)
```

C5.0 classification

```
parsnip::C5.0_train(x = missing_arg(), y = missing_arg(), weights = missing_arg())
```

spark classification

```
sparklyr::ml_gradient_boosted_trees(x = missing_arg(), formula = missing_arg(),
  type = "classification", seed = sample.int(10^5, 1))
```

spark regression

```
sparklyr::ml_gradient_boosted_trees(x = missing_arg(), formula = missing_arg(),
  type = "regression", seed = sample.int(10^5, 1))
```

Note

For models created using the spark engine, there are several differences to consider. First, only the formula interface to via `fit` is available; using `fit_xy` will generate an error. Second, the predictions will always be in a spark table format. The names will be the same as documented but without the dots. Third, there is no equivalent to factor columns in spark tables so class predictions are returned as character columns. Fourth, to retain the model object for a new R session (via `save`), the `model$fit` element of the `parsnip` object should be serialized via `ml_save(object$fit)` and separately saved to disk. In a new session, the object can be reloaded and reattached to the `parsnip` object.

See Also

[varying\(\)](#), [fit\(\)](#), [set_engine\(\)](#)

Examples

```
boost_tree(mode = "classification", trees = 20)
# Parameters can be represented by a placeholder:
boost_tree(mode = "regression", mtry = varying())
model <- boost_tree(mtry = 10, min_n = 3)
model
update(model, mtry = 1)
update(model, mtry = 1, fresh = TRUE)
```

C5.0_train

Boosted trees via C5.0

Description

C5.0_train is a wrapper for the C5.0() function in the **C50** package that fits tree-based models where all of the model arguments are in the main function.

Usage

```
C5.0_train(x, y, weights = NULL, trials = 15, minCases = 2,
           sample = 0, ...)
```

Arguments

x	A data frame or matrix of predictors.
y	A factor vector with 2 or more levels
weights	An optional numeric vector of case weights. Note that the data used for the case weights will not be used as a splitting variable in the model (see http://www.rulequest.com/see5-win.html#CASEWEIGHT for Quinlan's notes on case weights).
trials	An integer specifying the number of boosting iterations. A value of one indicates that a single model is used.
minCases	An integer for the smallest number of samples that must be put in at least two of the splits.
sample	A value between (0, .999) that specifies the random proportion of the data should be used to train the model. By default, all the samples are used for model training. Samples not used for training are used to evaluate the accuracy of the model in the printed output.
...	Other arguments to pass.

Value

A fitted C5.0 model.

decision_tree

General Interface for Decision Tree Models

Description

decision_tree is a way to generate a *specification* of a model before fitting and allows the model to be created using different packages in R or via Spark. The main arguments for the model are:

- **cost_complexity**: The cost/complexity parameter (a.k.a. Cp) used by CART models (rpart only).
- **tree_depth**: The *maximum* depth of a tree (rpart and spark only).

- `min_n`: The minimum number of data points in a node that are required for the node to be split further.

These arguments are converted to their specific names at the time that the model is fit. Other options and argument can be set using `set_engine`. If left to their defaults here (NULL), the values are taken from the underlying model functions. If parameters need to be modified, `update` can be used in lieu of recreating the object from scratch.

Usage

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

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

Arguments

<code>mode</code>	A single character string for the type of model. Possible values for this model are "unknown", "regression", or "classification".
<code>cost_complexity</code>	A positive number for the the cost/complexity parameter (a.k.a. Cp) used by CART models (<code>rpart</code> 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.
<code>object</code>	A random forest model specification.
<code>fresh</code>	A logical for whether the arguments should be modified in-place of or replaced wholesale.
<code>...</code>	Not used for <code>update</code> .

Details

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

- **R**: "rpart" or "C5.0" (classification only)
- **Spark**: "spark"

Note that, for `rpart` models, but `cost_complexity` and `tree_depth` can be both be specified but the package will give precedence to `cost_complexity`. Also, `tree_depth` values greater than 30 `rpart` will give nonsense results on 32-bit machines.

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::

rpart classification

```
rpart::rpart(formula = missing_arg(), data = missing_arg(), weights = missing_arg())
```

rpart regression

```
rpart::rpart(formula = missing_arg(), data = missing_arg(), weights = missing_arg())
```

C5.0 classification

```
parsnip::C5.0_train(x = missing_arg(), y = missing_arg(), weights = missing_arg(),  
  trials = 1)
```

spark classification

```
sparklyr::ml_decision_tree_classifier(x = missing_arg(), formula = missing_arg(),  
  seed = sample.int(10^5, 1), type = "classification")
```

spark regression

```
sparklyr::ml_decision_tree_classifier(x = missing_arg(), formula = missing_arg(),  
  seed = sample.int(10^5, 1), type = "regression")
```

Note

For models created using the spark engine, there are several differences to consider. First, only the formula interface to via `fit` is available; using `fit_xy` will generate an error. Second, the predictions will always be in a spark table format. The names will be the same as documented but without the dots. Third, there is no equivalent to factor columns in spark tables so class predictions are returned as character columns. Fourth, to retain the model object for a new R session (via `save`), the `model$fit` element of the `parsnip` object should be serialized via `ml_save(object$fit)` and separately saved to disk. In a new session, the object can be reloaded and reattached to the `parsnip` object.

See Also

[varying\(\)](#), [fit\(\)](#)

Examples

```
decision_tree(mode = "classification", tree_depth = 5)  
# Parameters can be represented by a placeholder:  
decision_tree(mode = "regression", cost_complexity = varying())  
model <- decision_tree(cost_complexity = 10, min_n = 3)  
model  
update(model, cost_complexity = 1)  
update(model, cost_complexity = 1, fresh = TRUE)
```

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 are 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 dat 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 `Sepal.Width ~ .` with the `iris` data, the values would be

```
.preds() = 4          (the 4 columns in `iris`)
.cols()  = 5          (3 numeric columns + 2 from Species dummy variables)
.obs()   = 150
.lvls()  = NA         (no factor outcome)
.facts() = 1          (the Species predictor)
.y()     = <vector>   (Sepal.Width as a vector)
.x()     = <data.frame> (The other 4 columns as a data frame)
.dat()   = <data.frame> (The full data set)
```

If the formula `Species ~ .` where used:

```
.preds() = 4          (the 4 numeric columns in `iris`)
.cols()  = 4          (same)
```

```
.obs() = 150
.lvls() = c(setosa = 50, versicolor = 50, virginica = 50)
.facts() = 0
.y() = <vector> (Species as a vector)
.x() = <data.frame> (The other 4 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:

```
data("lending_club")

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

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

<code>fit.model_spec</code>	<i>Fit a Model Specification to a Dataset</i>
-----------------------------	---

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 = NULL, data = NULL,
     control = fit_control(), ...)

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

Arguments

<code>object</code>	An object of class <code>model_spec</code> that has a chosen engine (via <code>set_engine()</code>).
<code>formula</code>	An object of class "formula" (or one that can be coerced to that class): a symbolic description of the model to be fitted.
<code>data</code>	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.
<code>control</code>	A named list with elements <code>verbosity</code> and <code>catch</code> . See <code>fit_control()</code> .
<code>...</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> .
<code>x</code>	A matrix or data frame of predictors.
<code>y</code>	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, checks them for validity, then fits 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` were 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.

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\(\)](#), [fit_control\(\)](#), `model_spec`, `model_fit`

Examples

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

library(dplyr)
data("lending_club")

lr_mod <- logistic_reg()

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

fit_control *Control the fit function*

Description

Options can be passed to the `fit()` function that control the output and computations

Usage

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

Arguments

verbosity	An integer where a value of zero indicates that no messages or output should be shown when packages are loaded or when the model is fit. A value of 1 means that package loading is quiet but model fits can produce output to the screen (depending on if they contain their own verbose-type argument). A value of 2 or more indicates that any output should be seen.
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

A named list with the results of the function call

keras_mlp *Simple interface to MLP models via keras*

Description

Instead of building a `keras` model sequentially, `keras_mlp` can be used to create a feed-forward network with a single hidden layer. Regularization is via either weight decay or dropout.

Usage

```
keras_mlp(x, y, hidden_units = 5, decay = 0, dropout = 0,
          epochs = 20, act = "softmax", seeds = sample.int(10^5, size = 3),
          ...)
```

Arguments

x	A data frame or matrix of predictors
y	A vector (factor or numeric) or matrix (numeric) of outcome data.
hidden_units	An integer for the number of hidden units.
decay	A non-negative real number for the amount of weight decay. Either this parameter <i>or</i> dropout can specified.
dropout	The proportion of parameters to set to zero. Either this parameter <i>or</i> decay can specified.
epochs	An integer for the number of passes through the data.
act	A character string for the type of activation function between layers.
seeds	A vector of three positive integers to control randomness of the calculations.
...	Currently ignored.

Value

A keras model object.

lending_club	<i>Loan Data</i>
--------------	------------------

Description

Loan Data

Details

These data were downloaded from the Lending Club access site (see below) and are from the first quarter of 2016. A subset of the rows and variables are included here. The outcome is in the variable `Class` and is either "good" (meaning that the loan was fully paid back or currently on-time) or "bad" (charged off, defaulted, of 21-120 days late). A data dictionary can be found on the source website.

Value

lending_club a data frame

Source

Lending Club Statistics <https://www.lendingclub.com/info/download-data.action>

Examples

```
data(lending_club)
str(lending_club)
```

Description

`linear_reg` is a way to generate a *specification* of a model before fitting and allows the model to be created using different packages in R, Stan, keras, or via Spark. The main arguments for the model are:

- `penalty`: The total amount of regularization in the model. Note that this must be zero for some engines.
- `mixture`: The proportion of L1 regularization in the model. Note that this will be ignored for some engines.

These arguments are converted to their specific names at the time that the model is fit. Other options and argument can be set using `set_engine`. If left to their defaults here (NULL), the values are taken from the underlying model functions. If parameters need to be modified, `update` can be used in lieu of recreating the object from scratch.

Usage

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

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

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>	An non-negative number representing the total amount of regularization (<code>glmnet</code> , <code>keras</code> , and <code>spark</code> only). For <code>keras</code> 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 <code>mixture</code>).
<code>mixture</code>	A number between zero and one (inclusive) that represents the proportion of regularization that is used for the L2 penalty (i.e. weight decay, or ridge regression) versus L1 (the lasso) (<code>glmnet</code> and <code>spark</code> only).
<code>object</code>	A linear regression model specification.
<code>fresh</code>	A logical for whether the arguments should be modified in-place or replaced wholesale.
<code>...</code>	Not used for <code>update</code> .

Details

The data given to the function are not saved and are only used to determine the *mode* of the model. For `linear_reg`, the mode will always be "regression".

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

- **R**: "lm" or "glmnet"

- **Stan:** "stan"
- **Spark:** "spark"
- **keras:** "keras"

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:

lm

```
stats::lm(formula = missing_arg(), data = missing_arg(), weights = missing_arg())
```

glmnet

```
glmnet::glmnet(x = missing_arg(), y = missing_arg(), weights = missing_arg(),
  family = "gaussian")
```

stan

```
rstanarm::stan_glm(formula = missing_arg(), data = missing_arg(),
  weights = missing_arg(), family = stats::gaussian)
```

spark

```
sparklyr::ml_linear_regression(x = missing_arg(), formula = missing_arg(),
  weight_col = missing_arg())
```

keras

```
parsnip::keras_mlp(x = missing_arg(), y = missing_arg(), hidden_units = 1,
  act = "linear")
```

When using `glmnet` models, there is the option to pass multiple values (or no values) to the `penalty` argument. This can have an effect on the model object results. When using the `predict` method in these cases, the return object type depends on the value of `penalty`. If a single value is given, the results will be a simple numeric vector. When multiple values or no values for `penalty` are used in `linear_reg`, the `predict` method will return a data frame with columns `values` and `lambda`.

For prediction, the `stan` engine can compute posterior intervals analogous to confidence and prediction intervals. In these instances, the units are the original outcome and when `std_error = TRUE`, the standard deviation of the posterior distribution (or posterior predictive distribution as appropriate) is returned.

Note

For models created using the `spark` engine, there are several differences to consider. First, only the formula interface to `via fit` is available; using `fit_xy` will generate an error. Second, the predictions will always be in a spark table format. The names will be the same as documented but without the dots. Third, there is no equivalent to factor columns in spark tables so class predictions are returned as character columns. Fourth, to retain the model object for a new R session (via `save`), the `model$fit` element of the `parsnip` object should be serialized via `ml_save(object$fit)` and separately saved to disk. In a new session, the object can be reloaded and reattached to the `parsnip` object.

See Also

[varying\(\)](#), [fit\(\)](#), [set_engine\(\)](#)

Examples

```
linear_reg()
# Parameters can be represented by a placeholder:
linear_reg(penalty = varying())
model <- linear_reg(penalty = 10, mixture = 0.1)
model
update(model, penalty = 1)
update(model, penalty = 1, fresh = TRUE)
```

logistic_reg

General Interface for Logistic Regression Models

Description

`logistic_reg` is a way to generate a *specification* of a model before fitting and allows the model to be created using different packages in R, Stan, keras, or via Spark. The main arguments for the model are:

- `penalty`: The total amount of regularization in the model. Note that this must be zero for some engines.
- `mixture`: The proportion of L1 regularization in the model. Note that this will be ignored for some engines.

These arguments are converted to their specific names at the time that the model is fit. Other options and argument can be set using `set_engine`. If left to their defaults here (NULL), the values are taken from the underlying model functions. If parameters need to be modified, `update` can be used in lieu of recreating the object from scratch.

Usage

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

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

Arguments

- | | |
|----------------------|---|
| <code>mode</code> | A single character string for the type of model. The only possible value for this model is "classification". |
| <code>penalty</code> | An non-negative number representing the total amount of regularization (glmnet, keras, and spark 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 <code>mixture</code>). |
| <code>mixture</code> | A number between zero and one (inclusive) that represents the proportion of regularization that is used for the L2 penalty (i.e. weight decay, or ridge regression) versus L1 (the lasso) (glmnet and spark only). |

<code>object</code>	A logistic regression model specification.
<code>fresh</code>	A logical for whether the arguments should be modified in-place of or replaced wholesale.
<code>...</code>	Not used for <code>update</code> .

Details

For `logistic_reg`, the mode will always be "classification".

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

- **R:** "glm" or "glmnet"
- **Stan:** "stan"
- **Spark:** "spark"
- **keras:** "keras"

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:

glm

```
stats::glm(formula = missing_arg(), data = missing_arg(), weights = missing_arg(),
  family = stats::binomial)
```

glmnet

```
glmnet::glmnet(x = missing_arg(), y = missing_arg(), weights = missing_arg(),
  family = "binomial")
```

stan

```
rstanarm::stan_glm(formula = missing_arg(), data = missing_arg(),
  weights = missing_arg(), family = stats::binomial)
```

spark

```
sparklyr::ml_logistic_regression(x = missing_arg(), formula = missing_arg(),
  weight_col = missing_arg(), family = "binomial")
```

keras

```
parsnip::keras_mlp(x = missing_arg(), y = missing_arg(), hidden_units = 1,
  act = "linear")
```

When using `glmnet` models, there is the option to pass multiple values (or no values) to the `penalty` argument. This can have an effect on the model object results. When using the `predict` method in these cases, the return object type depends on the value of `penalty`. If a single value is given, the results will be a simple numeric vector. When multiple values or no values for `penalty` are used in `logistic_reg`, the `predict` method will return a data frame with columns `values` and `lambda`.

For prediction, the `stan` engine can compute posterior intervals analogous to confidence and prediction intervals. In these instances, the units are the original outcome and when `std_error = TRUE`, the standard deviation of the posterior distribution (or posterior predictive distribution as appropriate) is returned. For `glm`, the standard error is in logit units while the intervals are in probability units.

Note

For models created using the spark engine, there are several differences to consider. First, only the formula interface to via `fit` is available; using `fit_xy` will generate an error. Second, the predictions will always be in a spark table format. The names will be the same as documented but without the dots. Third, there is no equivalent to factor columns in spark tables so class predictions are returned as character columns. Fourth, to retain the model object for a new R session (via `save`), the `model$fit` element of the `parsnip` object should be serialized via `ml_save(object$fit)` and separately saved to disk. In a new session, the object can be reloaded and reattached to the `parsnip` object.

See Also

`varying()`, `fit()`

Examples

```
logistic_reg()
# Parameters can be represented by a placeholder:
logistic_reg(penalty = varying())
model <- logistic_reg(penalty = 10, mixture = 0.1)
model
update(model, penalty = 1)
update(model, penalty = 1, fresh = TRUE)
```

mars

General Interface for MARS

Description

`mars` is a way to generate a *specification* of a model before fitting and allows the model to be created using R. The main arguments for the model are:

- `num_terms`: The number of features that will be retained in the final model.
- `prod_degree`: The highest possible degree of interaction between features. A value of 1 indicates an additive model while a value of 2 allows, but does not guarantee, two-way interactions between features.
- `prune_method`: The type of pruning. Possible values are listed in `?earth`.

These arguments are converted to their specific names at the time that the model is fit. Other options and argument can be set using `set_engine`. If left to their defaults here (NULL), the values are taken from the underlying model functions. If parameters need to be modified, `update` can be used in lieu of recreating the object from scratch.

Usage

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

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

Arguments

<code>mode</code>	A single character string for the type of model. Possible values for this model are "unknown", "regression", or "classification".
<code>num_terms</code>	The number of features that will be retained in the final model, including the intercept.
<code>prod_degree</code>	The highest possible interaction degree.
<code>prune_method</code>	The pruning method.
<code>object</code>	A MARS model specification.
<code>fresh</code>	A logical for whether the arguments should be modified in-place of or replaced wholesale.
<code>...</code>	Not used for <code>update</code> .

Details

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

- **R:** "earth"

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:

earth classification

```
earth::earth(x = missing_arg(), y = missing_arg(), weights = missing_arg(),
             glm = list(family = stats::binomial), keepxy = TRUE)
```

earth regression

```
earth::earth(x = missing_arg(), y = missing_arg(), weights = missing_arg(),
             keepxy = TRUE)
```

Note that, when the model is fit, the **earth** package only has its namespace loaded. However, if `multi_predict` is used, the package is attached.

See Also

[varying\(\)](#), [fit\(\)](#)

Examples

```
mars(mode = "regression", num_terms = 5)
model <- mars(num_terms = 10, prune_method = "none")
model
update(model, num_terms = 1)
update(model, num_terms = 1, fresh = TRUE)
```

Description

`mlp`, for multilayer perceptron, is a way to generate a *specification* of a model before fitting and allows the model to be created using different packages in R or via keras. The main arguments for the model are:

- `hidden_units`: The number of units in the hidden layer (default: 5).
- `penalty`: The amount of L2 regularization (aka weight decay, default is zero).
- `dropout`: The proportion of parameters randomly dropped out of the model (keras only, default is zero).
- `epochs`: The number of training iterations (default: 20).
- `activation`: The type of function that connects the hidden layer and the input variables (keras only, default is softmax).

If parameters need to be modified, this function can be used in lieu of recreating the object from scratch.

Usage

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

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

Arguments

<code>mode</code>	A single character string for the type of model. Possible values for this model are "unknown", "regression", or "classification".
<code>hidden_units</code>	An integer for the number of units in the hidden model.
<code>penalty</code>	A non-negative numeric value for the amount of weight decay.
<code>dropout</code>	A number between 0 (inclusive) and 1 denoting the proportion of model parameters randomly set to zero during model training.
<code>epochs</code>	An integer for the number of training iterations.
<code>activation</code>	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".
<code>object</code>	A random forest model specification.
<code>fresh</code>	A logical for whether the arguments should be modified in-place or replaced wholesale.
<code>...</code>	Not used for <code>update</code> .

Details

These arguments are converted to their specific names at the time that the model is fit. Other options and argument can be set using `set_engine`. If left to their defaults here (see above), the values are taken from the underlying model functions. One exception is `hidden_units` when `nnet::nnet` is used; that function's `size` argument has no default so a value of 5 units will be used. Also, unless otherwise specified, the `linout` argument to `nnet::nnet` will be set to `TRUE` when a regression model is created. If parameters need to be modified, `update` can be used in lieu of recreating the object from scratch.

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

- **R**: "nnet"
- **keras**: "keras"

An error is thrown if both `penalty` and `dropout` are specified for `keras` models.

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:

keras classification

```
parsnip::keras_mlp(x = missing_arg(), y = missing_arg())
```

keras regression

```
parsnip::keras_mlp(x = missing_arg(), y = missing_arg())
```

nnet classification

```
nnet::nnet(formula = missing_arg(), data = missing_arg(), weights = missing_arg(),
  size = 5, trace = FALSE, linout = FALSE)
```

nnet regression

```
nnet::nnet(formula = missing_arg(), data = missing_arg(), weights = missing_arg(),
  size = 5, trace = FALSE, linout = TRUE)
```

See Also

[varying\(\)](#), [fit\(\)](#)

Examples

```
mlp(mode = "classification", penalty = 0.01)
# Parameters can be represented by a placeholder:
mlp(mode = "regression", hidden_units = varying())
model <- mlp(hidden_units = 10, dropout = 0.30)
model
update(model, hidden_units = 2)
update(model, hidden_units = 2, fresh = TRUE)
```

model_fit	<i>Model Fit Object Information</i>
-----------	-------------------------------------

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 to 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 be `varying()`. 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 `varying()`.
- **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(iris) - 1)
```

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

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

Main Arguments:

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

The model functions save the argument *expressions* and their associated environments (a.k.a. a quosure) to be evaluated later when either `fit()` or `fit_xy()` 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(!iris) - 1)
```

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

```
> rand_forest(mtry = ncol(!iris) - 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(iris) - 1
> mtry_val
[1] 4
> rand_forest(mtry = !!mtry_val)
Random Forest Model Specification (unknown)
```

Main Arguments:

```
mtry = 4
```

More information on quosures and quasiquotation can be found at <https://tidyeval.tidyverse.org>.

multinom_reg

General Interface for Multinomial Regression Models

Description

`multinom_reg` is a way to generate a *specification* of a model before fitting and allows the model to be created using different packages in R, keras, or Spark. The main arguments for the model are:

- **penalty**: The total amount of regularization in the model. Note that this must be zero for some engines.
- **mixture**: The proportion of L1 regularization in the model. Note that this will be ignored for some engines.

These arguments are converted to their specific names at the time that the model is fit. Other options and argument can be set using `set_engine`. If left to their defaults here (NULL), the values are taken from the underlying model functions. If parameters need to be modified, `update` can be used in lieu of recreating the object from scratch.

Usage

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

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

Arguments

<code>mode</code>	A single character string for the type of model. The only possible value for this model is "classification".
<code>penalty</code>	An non-negative number representing the total amount of regularization (<code>glmnet</code> , <code>keras</code> , and <code>spark</code> only). For <code>keras</code> 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 <code>mixture</code>).
<code>mixture</code>	A number between zero and one (inclusive) that represents the proportion of regularization that is used for the L2 penalty (i.e. weight decay, or ridge regression) versus L1 (the lasso) (<code>glmnet</code> only).
<code>object</code>	A multinomial regression model specification.
<code>fresh</code>	A logical for whether the arguments should be modified in-place of or replaced wholesale.
<code>...</code>	Not used for <code>update</code> .

Details

For `multinom_reg`, the mode will always be "classification".

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

- **R**: "glmnet"
- **Stan**: "stan"
- **keras**: "keras"

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:

glmnet

```
glmnet::glmnet(x = missing_arg(), y = missing_arg(), weights = missing_arg(),
  family = "multinomial")
```

spark

```
sparklyr::ml_logistic_regression(x = missing_arg(), formula = missing_arg(),
  weight_col = missing_arg(), family = "multinomial")
```

keras

```
parsnip::keras_mlp(x = missing_arg(), y = missing_arg(), hidden_units = 1,
  act = "linear")
```

When using `glmnet` models, there is the option to pass multiple values (or no values) to the `penalty` argument. This can have an effect on the model object results. When using the `predict` method in these cases, the return object type depends on the value of `penalty`. If a single value is given, the results will be a simple numeric vector. When multiple values or no values for `penalty` are used in `multinom_reg`, the `predict` method will return a data frame with columns `values` and `lambda`.

Note

For models created using the spark engine, there are several differences to consider. First, only the formula interface to `via fit` is available; using `fit_xy` will generate an error. Second, the predictions will always be in a spark table format. The names will be the same as documented but without the dots. Third, there is no equivalent to factor columns in spark tables so class predictions are returned as character columns. Fourth, to retain the model object for a new R session (via `save`), the `model$fit` element of the `parsnip` object should be serialized via `ml_save(object$fit)` and separately saved to disk. In a new session, the object can be reloaded and reattached to the `parsnip` object.

See Also

[varying\(\)](#), [fit\(\)](#)

Examples

```
multinom_reg()
# Parameters can be represented by a placeholder:
multinom_reg(penalty = varying())
model <- multinom_reg(penalty = 10, mixture = 0.1)
model
update(model, penalty = 1)
update(model, penalty = 1, fresh = TRUE)
```

Description

`nearest_neighbor()` is a way to generate a *specification* of a model before fitting and allows the model to be created using different packages in R. The main arguments for the model are:

- `neighbors`: The number of neighbors considered at each prediction.
- `weight_func`: The type of kernel function that weights the distances between samples.
- `dist_power`: The parameter used when calculating the Minkowski distance. This corresponds to the Manhattan distance with `dist_power = 1` and the Euclidean distance with `dist_power = 2`.

These arguments are converted to their specific names at the time that the model is fit. Other options and argument can be set using `set_engine`. If left to their defaults here (NULL), the values are taken from the underlying model functions. If parameters need to be modified, `update()` can be used in lieu of recreating the object from scratch.

Usage

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

Arguments

<code>mode</code>	A single character string for the type of model. Possible values for this model are "unknown", "regression", or "classification".
<code>neighbors</code>	A single integer for the number of neighbors to consider (often called <i>k</i>).
<code>weight_func</code>	A <i>single</i> character for the type of kernel function used to weight distances between samples. Valid choices are: "rectangular", "triangular", "epanechnikov", "biweight", "triweight", "cos", "inv", "gaussian", "rank", or "optimal".
<code>dist_power</code>	A single number for the parameter used in calculating Minkowski distance.

Details

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

- **R**: "kknn"

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:

kknn (classification or regression)

```
kknn::train.kknn(formula = missing_arg(), data = missing_arg(),
  kmax = missing_arg())
```

Note

For `kkn`, the underlying modeling function used is a restricted version of `train.kknn()` and not `kknn()`. It is set up in this way so that `parsnip` can utilize the underlying `predict.train.kknn` method to predict on new data. This also means that a single value of that function's `kernel` argument (a.k.a `weight_func` here) can be supplied

See Also

[varying\(\)](#), [fit\(\)](#)

Examples

```
nearest_neighbor(neighbors = 11)
```

`predict.model_fit` *Model predictions*

Description

Apply a model to create different types of predictions. `predict` can be used for all types of models and used the "type" argument for more specificity.

Usage

```
## S3 method for class 'model_fit'
predict(object, new_data, type = NULL,
        opts = list(), ...)

## S3 method for class 'model_fit'
predict_raw(object, new_data, opts = list(), ...)

predict_raw(object, ...)
```

Arguments

<code>object</code>	An object of class <code>model_fit</code>
<code>new_data</code>	A rectangular data object, such as a data frame.
<code>type</code>	A single character value or <code>NULL</code> . Possible values are "numeric", "class", "probs", "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.
<code>opts</code>	A list of optional arguments to the underlying predict function that will be used when <code>type = "raw"</code> . The list should not include options for the model object or the new data being predicted.
<code>...</code>	Ignored. To pass arguments to pass to the underlying function when <code>predict.model_fit(type = "raw")</code> , use the <code>opts</code> argument.

Details

If "type" is not supplied to `predict`, then a choice is made (`type = "numeric"` for regression models and `type = "class"` for classification).

`predict` is designed to provide a tidy result (see "Value" section below) in a tibble output format.

When using `type = "conf_int"` and `type = "pred_int"`, the options `level` and `std_error` can be used. The latter is a logical for an extra column of standard error values (if available).

Value

With the exception of `type = "raw"`, the results of `predict.model_fit` will be a tibble as many rows in the output as there are rows in `new_data` and the column names will be predictable.

For numeric results with a single outcome, the tibble will have a `.pred` column and `.pred_yname` for multivariate results.

For hard class predictions, the column is named `.pred_class` and, when `type = "prob"`, the columns are `.pred_classlevel`.

`type = "conf_int"` and `type = "pred_int"` return tibbles with columns `.pred_lower` and `.pred_upper` with an attribute for the confidence level. In the case where intervals can be produced for class probabilities (or other non-scalar outputs), the columns will be named `.pred_lower_classlevel` and so on.

Quantile predictions return a tibble with a column `.pred`, which is a list-column. Each list element contains a tibble with columns `.pred` and `.quantile` (and perhaps other columns).

Using `type = "raw"` with `predict.model_fit` (or using `predict_raw`) will return the unadulterated results of the prediction function.

In the case of Spark-based models, since table columns cannot contain dots, the same convention is used except 1) no dots appear in names and 2) vectors are never returned but type-specific prediction functions.

Examples

```
library(dplyr)

lm_model <-
  linear_reg() %>%
  set_engine("lm") %>%
  fit(mpg ~ ., data = mtcars %>% slice(11:32))

pred_cars <-
  mtcars %>%
  slice(1:10) %>%
  select(-mpg)

predict(lm_model, pred_cars)

predict(
  lm_model,
  pred_cars,
  type = "conf_int",
  level = 0.90
)
```

```

predict(
  lm_model,
  pred_cars,
  type = "raw",
  opts = list(type = "terms")
)

```

rand_forest

General Interface for Random Forest Models

Description

rand_forest is a way to generate a *specification* of a model before fitting and allows the model to be created using different packages in R or via Spark. The main arguments for the model are:

- `mtry`: The number of predictors that will be randomly sampled at each split when creating the tree models.
- `trees`: The number of trees contained in the ensemble.
- `min_n`: The minimum number of data points in a node that are required for the node to be split further.

These arguments are converted to their specific names at the time that the model is fit. Other options and argument can be set using `set_engine`. If left to their defaults here (NULL), the values are taken from the underlying model functions. If parameters need to be modified, `update` can be used in lieu of recreating the object from scratch.

Usage

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

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

Arguments

<code>mode</code>	A single character string for the type of model. Possible values for this model are "unknown", "regression", or "classification".
<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.
<code>object</code>	A random forest model specification.
<code>fresh</code>	A logical for whether the arguments should be modified in-place of or replaced wholesale.
<code>...</code>	Not used for <code>update</code> .

Details

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

- **R**: "ranger" or "randomForest"
- **Spark**: "spark"

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::

ranger classification

```
ranger::ranger(formula = missing_arg(), data = missing_arg(),
  case.weights = missing_arg(), num.threads = 1, verbose = FALSE,
  seed = sample.int(10^5, 1), probability = TRUE)
```

ranger regression

```
ranger::ranger(formula = missing_arg(), data = missing_arg(),
  case.weights = missing_arg(), num.threads = 1, verbose = FALSE,
  seed = sample.int(10^5, 1))
```

randomForests classification

```
randomForest::randomForest(x = missing_arg(), y = missing_arg())
```

randomForests regression

```
randomForest::randomForest(x = missing_arg(), y = missing_arg())
```

spark classification

```
sparklyr::ml_random_forest(x = missing_arg(), formula = missing_arg(),
  type = "classification", seed = sample.int(10^5, 1))
```

spark regression

```
sparklyr::ml_random_forest(x = missing_arg(), formula = missing_arg(),
  type = "regression", seed = sample.int(10^5, 1))
```

For **ranger** confidence intervals, the intervals are constructed using the form `estimate +/- z * std.error`. For classification probabilities, these values can fall outside of `[0, 1]` and will be coerced to be in this range.

Note

For models created using the spark engine, there are several differences to consider. First, only the formula interface to via `fit` is available; using `fit_xy` will generate an error. Second, the predictions will always be in a spark table format. The names will be the same as documented but without the dots. Third, there is no equivalent to factor columns in spark tables so class predictions are returned as character columns. Fourth, to retain the model object for a new R session (via `save`), the `model$fit` element of the `parsnip` object should be serialized via `ml_save(object$fit)` and separately saved to disk. In a new session, the object can be reloaded and reattached to the `parsnip` object.

See Also

[varying\(\)](#), [fit\(\)](#)

Examples

```
rand_forest(mode = "classification", trees = 2000)
# Parameters can be represented by a placeholder:
rand_forest(mode = "regression", mtry = varying())
model <- rand_forest(mtry = 10, min_n = 3)
model
update(model, mtry = 1)
update(model, mtry = 1, fresh = TRUE)
```

rpart_train

Decision trees via rpart

Description

rpart_train is a wrapper for rpart() tree-based models where all of the model arguments are in the main function.

Usage

```
rpart_train(formula, data, weights = NULL, cp = 0.01, minsplit = 20,
  maxdepth = 30, ...)
```

Arguments

formula	A model formula.
data	A data frame.
weights	Optional case weights.
cp	A non-negative number for complexity parameter. Any split that does not decrease the overall lack of fit by a factor of <code>cp</code> is not attempted. For instance, with anova splitting, this means that the overall R-squared must increase by <code>cp</code> at each step. The main role of this parameter is to save computing time by pruning off splits that are obviously not worthwhile. Essentially, the user informs the program that any split which does not improve the fit by <code>cp</code> will likely be pruned off by cross-validation, and that hence the program need not pursue it.
minsplit	An integer for the minimum number of observations that must exist in a node in order for a split to be attempted.
maxdepth	An integer for the maximum depth of any node of the final tree, with the root node counted as depth 0. Values greater than 30 rpart will give nonsense results on 32-bit machines. This function will truncate <code>maxdepth</code> to 30 in those cases.
...	Other arguments to pass to either rpart or rpart.control.

Value

A fitted rpart model.

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

Description

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

Usage

```
set_args(object, ...)
```

```
set_mode(object, mode)
```

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 varying().

Value

An updated model specification.

Examples

```
# First, set general arguments using the standardized names
mod <-
  logistic_reg(mixture = 1/3) %>%
  # now say how you want to fit the model and another other options
  set_engine("glmnet", nlambda = 10)
translate(mod, engine = "glmnet")
```

surv_reg

General Interface for Parametric Survival Models

Description

surv_reg is a way to generate a *specification* of a model before fitting and allows the model to be created using R. The main argument for the model is:

- **dist**: The probability distribution of the outcome.

This argument is converted to its specific names at the time that the model is fit. Other options and argument can be set using `set_engine`. If left to its default here (NULL), the value is taken from the underlying model functions.

If parameters need to be modified, this function can be used in lieu of recreating the object from scratch.

Usage

```
surv_reg(mode = "regression", dist = NULL)

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

Arguments

mode	A single character string for the type of model. The only possible value for this model is "regression".
dist	A character string for the outcome distribution. "weibull" is the default.
object	A survival regression model specification.
fresh	A logical for whether the arguments should be modified in-place of or replaced wholesale.
...	Not used for update.

Details

The data given to the function are not saved and are only used to determine the *mode* of the model. For `surv_reg`, the mode will always be "regression".

Since survival models typically involve censoring (and require the use of `survival::Surv()` objects), the `fit()` function will require that the survival model be specified via the formula interface.

Also, for the `flexsurv::flexsurvfit` engine, the typical `strata` function cannot be used. To achieve the same effect, the extra parameter roles can be used (as described above).

For `surv_reg`, the mode will always be "regression".

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

- R: "flexsurv", "survreg"

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:

flexsurv

```
flexsurv::flexsurvreg(formula = missing_arg(), data = missing_arg(),
  weights = missing_arg())
```

survreg

```
survival::survreg(formula = missing_arg(), data = missing_arg(),
  weights = missing_arg(), model = TRUE)
```

Note that `model = TRUE` is needed to produce quantile predictions when there is a stratification variable and can be overridden in other cases.

References

Jackson, C. (2016). `flexsurv`: A Platform for Parametric Survival Modeling in R. *Journal of Statistical Software*, 70(8), 1 - 33.

See Also

[varying\(\)](#), [fit\(\)](#), [survival::Surv\(\)](#)

Examples

```
surv_reg()
# Parameters can be represented by a placeholder:
surv_reg(dist = varying())

model <- surv_reg(dist = "weibull")
model
update(model, dist = "lnorm")
```

svm_poly

General interface for polynomial support vector machines

Description

svm_poly is a way to generate a *specification* of a model before fitting and allows the model to be created using different packages in R or via Spark. The main arguments for the model are:

- **cost**: The cost of predicting a sample within or on the wrong side of the margin.
- **degree**: The polynomial degree.
- **scale_factor**: A scaling factor for the kernel.
- **margin**: The epsilon in the SVM insensitive loss function (regression only)

These arguments are converted to their specific names at the time that the model is fit. Other options and argument can be set using `set_engine`. If left to their defaults here (NULL), the values are taken from the underlying model functions. If parameters need to be modified, `update` can be used in lieu of recreating the object from scratch.

Usage

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

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

Arguments

mode	A single character string for the type of model. Possible values for this model are "unknown", "regression", or "classification".
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)
object	A polynomial SVM model specification.
fresh	A logical for whether the arguments should be modified in-place of or replaced wholesale.
...	Not used for update.

Details

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

- **R**: "kernlab"

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::

kernlab classification

```
kernlab::ksvm(x = missing_arg(), y = missing_arg(), kernel = "polydot",
  prob.model = TRUE)
```

kernlab regression

```
kernlab::ksvm(x = missing_arg(), y = missing_arg(), kernel = "polydot")
```

See Also

[varying\(\)](#), [fit\(\)](#)

Examples

```
svm_poly(mode = "classification", degree = 1.2)
# Parameters can be represented by a placeholder:
svm_poly(mode = "regression", cost = varying())
model <- svm_poly(cost = 10, scale_factor = 0.1)
model
update(model, cost = 1)
update(model, cost = 1, fresh = TRUE)
```

svm_rbf

General interface for radial basis function support vector machines

Description

svm_rbf is a way to generate a *specification* of a model before fitting and allows the model to be created using different packages in R or via Spark. The main arguments for the model are:

- **cost**: The cost of predicting a sample within or on the wrong side of the margin.
- **rbf_sigma**: The precision parameter for the radial basis function.
- **margin**: The epsilon in the SVM insensitive loss function (regression only)

These arguments are converted to their specific names at the time that the model is fit. Other options and argument can be set using `set_engine`. If left to their defaults here (NULL), the values are taken from the underlying model functions. If parameters need to be modified, `update` can be used in lieu of recreating the object from scratch.

Usage

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

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

Arguments

mode	A single character string for the type of model. Possible values for this model are "unknown", "regression", or "classification".
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)
object	A radial basis function SVM model specification.
fresh	A logical for whether the arguments should be modified in-place or replaced wholesale.
...	Not used for update.

Details

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

- **R**: "kernlab"

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::

kernlab classification

```
kernlab::ksvm(x = missing_arg(), y = missing_arg(), kernel = "rbfdot",
  prob.model = TRUE)
```

kernlab regression

```
kernlab::ksvm(x = missing_arg(), y = missing_arg(), kernel = "rbfdot")
```

See Also

[varying\(\)](#), [fit\(\)](#)

Examples

```
svm_rbf(mode = "classification", rbf_sigma = 0.2)
# Parameters can be represented by a placeholder:
svm_rbf(mode = "regression", cost = varying())
model <- svm_rbf(cost = 10, rbf_sigma = 0.1)
model
update(model, cost = 1)
update(model, cost = 1, fresh = TRUE)
```

`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, ...)
```

Arguments

<code>x</code>	A model specification.
<code>...</code>	Not currently used.

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 `varying` arguments if these are in the specification.

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(mixture = varying()), engine = "glmnet")
```

varying	<i>A placeholder function for argument values</i>
---------	---

Description

`varying()` is used when a parameter will be specified at a later date.

Usage

```
varying()
```

varying_args	<i>Determine varying arguments</i>
--------------	------------------------------------

Description

`varying_args` takes a model specification and lists all of the arguments along with whether they are fully specified or not.

Usage

```
varying_args(x, id, ...)
```

```
## S3 method for class 'model_spec'
```

```
varying_args(x, id = NULL, ...)
```

```
## S3 method for class 'recipe'
```

```
varying_args(x, id = NULL, ...)
```

```
## S3 method for class 'step'
```

```
varying_args(x, id = NULL, ...)
```

Arguments

<code>x</code>	An object
<code>id</code>	A string describing the object <code>x</code> .
<code>...</code>	Not currently used.

Value

A tibble with columns for the parameter name (`name`), whether it contains *any* varying value (`varying`), the `id` for the object, and the class that was used to call the method (`type`).

Examples

```

library(dplyr)
library(rlang)

rand_forest() %>% varying_args(id = "plain")

rand_forest(mtry = varying()) %>% varying_args(id = "one arg")

rand_forest() %>%
  set_engine("ranger", sample.fraction = varying()) %>%
  varying_args(id = "only eng_args")

rand_forest() %>%
  set_engine(
    "ranger",
    strata = expr(Class),
    sampsize = c(varying(), varying())
  ) %>%
  varying_args(id = "add an expr")

rand_forest() %>%
  set_engine("ranger", classwt = c(class1 = 1, class2 = varying())) %>%
  varying_args(id = "list of values")

```

wa_churn

Watson Churn Data

Description

Watson Churn Data

Details

These data were downloaded from the IBM Watson site (see below) in September 2018. The data contain a factor for whether a customer churned or not. Alternatively, the `tenure` column presumably contains information on how long the customer has had an account. A survival analysis can be done on this column using the `churn` outcome as the censoring information. A data dictionary can be found on the source website.

Value

wa_churn a data frame

Source

IBM Watson Analytics <https://ibm.co/2sOvyvy>

Examples

```

data(wa_churn)
str(wa_churn)

```

`xgb_train`*Boosted trees via xgboost*

Description

`xgb_train` is a wrapper for `xgboost` tree-based models where all of the model arguments are in the main function.

Usage

```
xgb_train(x, y, max_depth = 6, nrounds = 15, eta = 0.3,  
          colsample_bytree = 1, min_child_weight = 1, gamma = 0,  
          subsample = 1, ...)
```

Arguments

<code>x</code>	A data frame or matrix of predictors
<code>y</code>	A vector (factor or numeric) or matrix (numeric) of outcome data.
<code>max_depth</code>	An integer for the maximum depth of the tree.
<code>nrounds</code>	An integer for the number of boosting iterations.
<code>eta</code>	A numeric value between zero and one to control the learning rate.
<code>colsample_bytree</code>	Subsampling proportion of columns.
<code>min_child_weight</code>	A numeric value for the minimum sum of instance weights needed in a child to continue to split.
<code>gamma</code>	An number for the minimum loss reduction required to make a further partition on a leaf node of the tree
<code>subsample</code>	Subsampling proportion of rows.
<code>...</code>	Other options to pass to <code>xgb.train</code> .

Value

A fitted `xgboost` object.

Index

- *Topic **datasets**
 - `lending_club`, 12
 - `wa_churn`, 40
 - `.cols` (descriptors), 7
 - `.dat` (descriptors), 7
 - `.facts` (descriptors), 7
 - `.lvls` (descriptors), 7
 - `.obs` (descriptors), 7
 - `.preds` (descriptors), 7
 - `.x` (descriptors), 7
 - `.y` (descriptors), 7
- `boost_tree`, 2
- `C5.0_train`, 5
- `decision_tree`, 5
- descriptors, 7
- `fit()`, 4, 7, 11, 15, 17, 18, 20, 23, 25, 27, 31, 34, 36, 37
- `fit.model_spec`, 9
- `fit_control`, 11
- `fit_control()`, 9, 10
- `fit_xy()`, 23
- `fit_xy.model_spec` (`fit.model_spec`), 9
- `keras_mlp`, 11
- `lending_club`, 12
- `linear_reg`, 13
- `logistic_reg`, 15
- `mars`, 17
- `mlp`, 19
- `model_fit`, 21
- `model_spec`, 21, 22
- `multinom_reg`, 24
- `nearest_neighbor`, 26
- `predict.model_fit`, 27
- `predict_raw` (`predict.model_fit`), 27
- `rand_forest`, 29
- `rpart_train`, 31
- `set_args`, 32
- `set_engine`, 32
- `set_engine()`, 4, 9, 10, 15
- `set_mode` (`set_args`), 32
- `surv_reg`, 33
- `survival::Surv()`, 34
- `svm_poly`, 35
- `svm_rbf`, 36
- `translate`, 38
- `update.boost_tree` (`boost_tree`), 2
- `update.decision_tree` (`decision_tree`), 5
- `update.linear_reg` (`linear_reg`), 13
- `update.logistic_reg` (`logistic_reg`), 15
- `update.mars` (`mars`), 17
- `update.mlp` (`mlp`), 19
- `update.multinom_reg` (`multinom_reg`), 24
- `update.rand_forest` (`rand_forest`), 29
- `update.surv_reg` (`surv_reg`), 33
- `update.svm_poly` (`svm_poly`), 35
- `update.svm_rbf` (`svm_rbf`), 36
- `varying`, 39
- `varying()`, 4, 7, 15, 17, 18, 20, 25, 27, 31, 34, 36, 37, 39
- `varying_args`, 39
- `wa_churn`, 40
- `xgb_train`, 41