

# Package ‘svmodt’

June 30, 2026

**Type** Package

**Title** Linear SVM-Based Recursive Decision Trees

**Version** 0.1.0

**Description** Implements Support Vector Machine Oblique Decision Trees (SVMODT).  
Recursively builds classification trees using linear Support Vector Machines (SVM) hyper-planes at each node instead of axis-parallel splits, creating oblique decision boundaries. Features include multiple feature selection methods, dynamic feature subset strategies, class weight support for imbalanced datasets, pruning and feature penalization.

**License** GPL (>= 3)

**Encoding** UTF-8

**LazyData** true

**Suggests** knitr, rmarkdown, bookdown, testthat (>= 3.0.0), rpart, rsample, gridExtra, tidyr, kableExtra, palmerpenguins, dplyr

**VignetteBuilder** knitr

**Depends** R (>= 3.5)

**Imports** rlang, e1071, FSelectorRcpp, ggplot2

**RoxygenNote** 7.3.3

**Config/testthat/edition** 3

**URL** <https://github.com/AneeshAgarwala/svmodt>

**BugReports** <https://github.com/AneeshAgarwala/svmodt/issues>

**NeedsCompilation** no

**Author** Aneesh Agarwal [aut, cre, cph],  
Jack Jewson [aut, ths],  
Erik Sverdrup [aut, ths]

**Maintainer** Aneesh Agarwal <aaga0022@student.monash.edu>

**Repository** CRAN

**Date/Publication** 2026-06-30 11:10:02 UTC

## Contents

plot.svmodt_node . . . . .	2
predict.svmodt_node . . . . .	3
print.svmodt_node . . . . .	4
svm_split . . . . .	5
trace_path . . . . .	8
wdbc . . . . .	9
wine . . . . .	10

<b>Index</b>	<b>12</b>
--------------	-----------

---

plot.svmodt_node	<i>Plot method for svmodt_node objects</i>
------------------	--

---

## Description

Thin S3 wrapper that dispatches to [plot\\_boundary](#) or [plot\\_surface](#) depending on `plot.type`.

## Usage

```
## S3 method for class 'svmodt_node'
plot(
  x,
  y = NULL,
  ...,
  data = NULL,
  response = NULL,
  plot.type = c("surface", "boundary"),
  features = NULL,
  max_depth = NULL,
  check_accuracy = TRUE,
  resolution = NULL
)
```

## Arguments

<code>x</code>	An <code>svmodt_node</code> returned by <a href="#">svm_split</a> .
<code>y</code>	Ignored; present only to satisfy the <code>graphics::plot</code> generic signature.
<code>...</code>	Currently unused.
<code>data</code>	The original training data frame (required).
<code>response</code>	Character string naming the response column (required).
<code>plot.type</code>	One of "surface" (default) or "boundary".
<code>features</code>	Length-2 character vector of axis features ("surface" only; default uses root node features).
<code>max_depth</code>	Maximum depth to visualize ("boundary" only; default NULL = full tree).
<code>check_accuracy</code>	Logical; show per-node accuracy ("boundary" only; default TRUE).
<code>resolution</code>	Grid resolution per axis. Default 100 for "boundary", 200 for "surface".

**Value**

- "boundary": invisibly returns the list from [plot\\_boundary](#).
- "surface": invisibly returns the **ggplot2** object from [plot\\_surface](#).

**Examples**

```
tree <- svm_split(wdbc, response = "diagnosis", max_depth = 3)

# All-node boundary panels - prints first, returns list
viz <- plot(tree,
  data = wdbc, response = "diagnosis",
  plot.type = "boundary"
)
viz$plots[[2]] # second node

# Global decision surface
plot(tree,
  data = wdbc, response = "diagnosis",
  plot.type = "surface"
)

# Surface with explicit feature axes
plot(tree,
  data = wdbc, response = "diagnosis",
  plot.type = "surface",
  features = c("radius_mean", "concavity_mean")
)
```

---

predict.svmodt\_node     *Predict method for svmodt\_node objects*

---

**Description**

Predict method for svmodt\_node objects

**Usage**

```
## S3 method for class 'svmodt_node'
predict(object, newdata, return_probs = FALSE, calibrate_probs = TRUE, ...)
```

**Arguments**

object	An object of class svmodt_node.
newdata	A data frame of new predictor values.
return_probs	Logical; if TRUE, returns predictions and probabilities.

calibrate\_probs Logical; if TRUE, uses logistic calibration on decision values.  
 ... Currently unused.

### Value

If return\_probs = FALSE (the default), a character vector of predicted class labels, one element per row of newdata.

If return\_probs = TRUE, a named list with two elements:

**predictions** Character vector of predicted class labels (length = nrow(newdata)).

**probabilities** Numeric matrix of class probabilities with nrow(newdata) rows and one column per class. Column names are the class labels; each row sums to 1. When calibrate\_probs = TRUE, probabilities are derived from the SVM decision value via logistic calibration; otherwise empirical class frequencies at the leaf node are used.

### Examples

```
# Train DTSVM tree
tree <- svm_split(
  data = wdbc,
  response = "diagnosis",
  max_depth = 3,
  max_features = 2,
  feature_method = "cor"
)

# Predict on WDBC data (returns a character vector of class labels)
preds <- predict(tree, newdata = wdbc)

# Predict with probabilities and logistic calibration
result <- predict(tree, newdata = wdbc,
  return_probs = TRUE, calibrate_probs = TRUE
)
head(result$predictions)
head(result$probabilities)
```

---

```
print.svmodt_node      ' Print method for svmodt_node objects'
```

---

### Description

' Print method for svmodt\_node objects

### Usage

```
## S3 method for class 'svmodt_node'
print(x, ...)
```

**Arguments**

- x                    An object of class `svmodt_node`.  
...                  Further arguments passed to `print_svm_tree`.

**Value**

Invisibly returns `x` (the `svmodt_node` object), called for its side effect of printing a human-readable summary of the tree structure to the console.

**Examples**

```
tree <- svm_split(  
  data = wdbc,  
  response = "diagnosis",  
  max_features = 2,  
  max_depth = 3,  
  min_samples = 5,  
  feature_method = "random",  
  verbose = TRUE  
)  
print(tree)
```

---

`svm_split`*Build an Oblique Decision Tree Using SVM Splits*

---

**Description**

Constructs a decision tree where each internal node uses a Support Vector Machine (SVM) to determine the split. Supports dynamic feature selection, feature penalization, scaling, and class weighting.

**Usage**

```
svm_split(  
  data,  
  response,  
  depth = 1,  
  max_depth = 10,  
  min_samples = 5,  
  max_features = NULL,  
  feature_method = c("random", "mutual", "cor"),  
  impurity_measure = c("entropy", "gini"),  
  max_features_strategy = c("constant", "random", "decrease"),  
  max_features_decrease_rate = 0.8,  
  max_features_random_range = c(0.3, 1),  
  penalize_used_features = FALSE,
```

```

feature_penalty_weight = 0.5,
n_subsets = 1,
used_features = character(0),
class_weights = c("none", "balanced", "custom"),
custom_class_weights = NULL,
min_impurity_decrease = 0.001,
verbose = FALSE,
all_classes = NULL,
...
)

```

## Arguments

<code>data</code>	A data frame containing predictors and the response variable.
<code>response</code>	Character string specifying the response column in ‘data’. All other columns are treated as predictors.
<code>depth</code>	Integer indicating the current recursion depth (used internally; default is 1).
<code>max_depth</code>	Maximum depth of the tree.
<code>min_samples</code>	Minimum number of samples required to attempt a split.
<code>max_features</code>	Maximum number of features to consider at each split.
<code>feature_method</code>	Feature selection method at each node. One of: <ul style="list-style-type: none"> <li>“random”: randomly select features,</li> <li>“mutual”: select based on mutual information with the response,</li> <li>“cor”: select based on correlation with the response.</li> </ul>
<code>impurity_measure</code>	Information Gain evaluation criteria <ul style="list-style-type: none"> <li>“gini”: use Gini ratio</li> <li>“entropy”: use Shannon entropy</li> </ul>
<code>max_features_strategy</code>	Strategy to adjust the number of features per node: <ul style="list-style-type: none"> <li>“constant”: keep ‘max_features’ constant,</li> <li>“decrease”: reduce features with depth,</li> <li>“random”: randomly vary number of features within a range.</li> </ul>
<code>max_features_decrease_rate</code>	Numeric fraction for decreasing features if ‘max_features_strategy = “decrease”’.
<code>max_features_random_range</code>	Numeric vector of length 2 specifying min and max fraction of features if ‘max_features_strategy = “random”’.
<code>penalize_used_features</code>	Logical; if TRUE, features used in ancestor nodes are penalized to encourage diversity.
<code>feature_penalty_weight</code>	Numeric (0 < $\lambda$ < 1) weight for penalizing previously used features.

<code>n_subsets</code>	Number of Evaluated Random Feature combinations at each node when <code>'feature_method = "random"'</code>
<code>used_features</code>	Character vector of features already used in ancestor nodes (used internally).
<code>class_weights</code>	Character string specifying how to handle class imbalance. One of: <ul style="list-style-type: none"> <li>• <code>"none"</code>: no weighting,</li> <li>• <code>"balanced"</code>: weight classes inversely proportional to their frequency,</li> <li>• <code>"custom"</code>: use <code>'custom_class_weights'</code>.</li> </ul>
<code>custom_class_weights</code>	Optional named numeric vector specifying custom weights per class.
<code>min_impurity_decrease</code>	Required decrease in impurity by a split to be considered valid
<code>verbose</code>	Logical; if TRUE, prints information about each node during tree construction.
<code>all_classes</code>	Optional character vector of all possible response classes (used internally).
<code>...</code>	Additional arguments passed to the underlying SVM fitting function.

## Details

This function recursively splits the dataset using an SVM at each node. Splitting stops when maximum depth is reached, the node contains fewer than `'min_samples'`, or all samples belong to the same class. Features are scaled and selected dynamically at each node, and previously used features can be penalized to promote diversity. Class weighting schemes support handling imbalanced datasets. This approach allows construction of an *oblique decision tree*, where splits are linear hyperplanes rather than axis-aligned.

## Value

A nested list representing the decision tree. Each node contains:

**is\_leaf** Logical; TRUE if the node is a leaf.

**model** Fitted SVM model at this node (for internal nodes).

**features** Vector of features selected for this node.

**scaler** Scaling information used at this node.

**left** Left child node (decision value  $> 0$ ).

**right** Right child node (decision value  $< 0$ ).

**depth** Depth of this node in the tree.

**n** Number of samples at this node.

**max\_features\_used** Number of features considered at this node.

**penalty\_applied** Logical; TRUE if feature penalization was applied.

**class\_weights\_used** Class weights applied at this node.

**Examples**

```
data(wdbc)
tree <- svm_split(
  data = wdbc,
  response = "diagnosis",
  max_depth = 3,
  min_samples = 5,
  feature_method = "random",
  verbose = TRUE
)
```

---

trace\_path

*Trace the prediction path of a sample through an svmmodt tree*


---

**Description**

Generic function that walks the tree for a single row of new data, printing the SVM decision value and chosen branch at every internal node and the final predicted class at the leaf.

**Usage**

```
trace_path(object, ...)
```

```
## S3 method for class 'svmmodt_node'
trace_path(object, sample_data, sample_idx = 1, ...)
```

**Arguments**

object	An svmmodt_node returned by <a href="#">svm_split</a> .
...	Currently unused.
sample_data	A data frame of new predictor values (one or more rows).
sample_idx	Integer; which row to trace (default 1).

**Value**

Invisibly returns the predicted class label (character string).

**Methods (by class)**

- trace\_path(svmmodt\_node): Method for svmmodt\_node objects.

**Examples**

```
tree <- svm_split(wdbc, response = "diagnosis", max_depth = 3)
trace_path(tree, wdbc, sample_idx = 5)
```

---

wdbc

*Wisconsin Diagnostic Breast Cancer Dataset*

---

## Description

The WDBC dataset contains quantitative measurements from digitized images of fine needle aspirates (FNA) of breast masses. It is commonly used for classification tasks to distinguish between benign and malignant tumors.

## Usage

wdbc

## Format

A data frame with 569 rows and 32 columns:

**radius\_mean** Mean of radius  
**radius\_se** Standard error of radius  
**radius\_worst** Worst (largest) radius  
**texture\_mean** Mean of texture  
**texture\_se** Standard error of texture  
**texture\_worst** Worst texture  
**perimeter\_mean** Mean of perimeter  
**perimeter\_se** Standard error of perimeter  
**perimeter\_worst** Worst perimeter  
**area\_mean** Mean area  
**area\_se** Standard error of area  
**area\_worst** Worst area  
**smoothness\_mean** Mean smoothness  
**smoothness\_se** Standard error of smoothness  
**smoothness\_worst** Worst smoothness  
**compactness\_mean** Mean compactness  
**compactness\_se** Standard error of compactness  
**compactness\_worst** Worst compactness  
**concavity\_mean** Mean concavity  
**concavity\_se** Standard error of concavity  
**concavity\_worst** Worst concavity  
**concave.points\_mean** Mean concave points  
**concave.points\_se** Standard error of concave points

**concave.points\_worst** Worst concave points  
**symmetry\_mean** Mean symmetry  
**symmetry\_se** Standard error of symmetry  
**symmetry\_worst** Worst symmetry  
**fractal\_dimension\_mean** Mean fractal dimension  
**fractal\_dimension\_se** Standard error of fractal dimension  
**fractal\_dimension\_worst** Worst fractal dimension  
**diagnosis** Factor with levels 'B' and 'M'

### Source

Dr. William H. Wolberg, W. Nick Street, and Olvi L. Mangasarian, University of Wisconsin-Madison.  
 Original dataset available at: <https://archive.ics.uci.edu/dataset/17/breast+cancer+wisconsin+diagnostic>

---

wine

*Wine Dataset*

---

### Description

The Wine dataset contains the results of a chemical analysis of wines derived from three different cultivars grown in the same region of Italy. The dataset is commonly used for multiclass classification tasks, where the objective is to identify the cultivar of origin based on physicochemical properties.

### Usage

wine

### Format

A data frame with 178 rows and 14 columns:

**class** Factor with levels 1, 2, and 3 indicating cultivar  
**alcohol** Alcohol content  
**malic\_acid** Malic acid concentration  
**ash** Ash content  
**alkalinity\_of\_ash** Alkalinity of ash  
**magnesium** Magnesium content  
**total\_phenols** Total phenols  
**flavonoids** Flavonoid content  
**nonflavonoid\_phenols** Nonflavonoid phenols  
**proanthocyanins** Proanthocyanin content  
**color\_intensity** Color intensity  
**hue** Hue  
**od280\_od315** OD280/OD315 of diluted wines  
**proline** Proline concentration

**Source**

Aeberhard, S. & Forina, M. (1992). Wine Dataset. UCI Machine Learning Repository. Original dataset available at: <<https://archive.ics.uci.edu/dataset/109/wine>>

# Index

## \* datasets

wdbc, 9  
wine, 10

plot.svmodt\_node, 2  
plot\_boundary, 2, 3  
plot\_surface, 2, 3  
predict.svmodt\_node, 3  
print.svmodt\_node, 4  
print\_svm\_tree, 5

svm\_split, 2, 5, 8

trace\_path, 8

wdbc, 9  
wine, 10