

Package ‘AddiVortes’

January 13, 2026

Title (Bayesian) Additive Voronoi Tessellations

Version 0.4.7

Date 2026-01-09

Description Implements the Bayesian Additive Voronoi Tessellation model for non-parametric regression and machine learning as introduced in Stone and Gosling (2025) <[doi:10.1080/10618600.2024.2414104](https://doi.org/10.1080/10618600.2024.2414104)>. This package provides a flexible alternative to BART (Bayesian Additive Regression Trees) using Voronoi tessellations instead of trees. Users can fit Bayesian regression models, estimate posterior distributions, and visualise the resulting tessellations. It is particularly useful for spatial data analysis, machine learning regression, complex function approximation and Bayesian modeling where the underlying structure is unknown. The method is well-suited to capturing spatial patterns and non-linear relationships.

License GPL (>= 3)

Encoding UTF-8

RoxygenNote 7.3.2

Depends R (>= 3.5.0)

Imports parallel (>= 4.0.0), pbapply (>= 1.6)

Suggests knitr, rmarkdown, testthat (>= 3.0.0), xml2

Config/testthat/edition 3

URL <https://johnpaulgosling.github.io/AddiVortes/>

VignetteBuilder knitr

BugReports <https://github.com/johnpaulgosling/AddiVortes/issues>

NeedsCompilation yes

Author Adam Stone [aut] (ORCID: <<https://orcid.org/0009-0004-0058-6117>>),
John Paul Gosling [aut, cre] (ORCID:
<<https://orcid.org/0000-0002-4072-3022>>)

Maintainer John Paul Gosling <john-paul.gosling@durham.ac.uk>

Repository CRAN

Date/Publication 2026-01-13 18:40:02 UTC

Contents

AddiVortes-package	2
AddiVortes	3
new_AddiVortesFit	4
plot.AddiVortesFit	5
predict.AddiVortesFit	6
print.AddiVortesFit	8
summary.AddiVortesFit	9

Index	10
--------------	-----------

AddiVortes-package	<i>AddiVortes: Bayesian Additive Voronoi Tessellations for Machine Learning</i>
--------------------	---

Description

AddiVortes implements Bayesian Additive Voronoi Tessellation models for machine learning regression and non-parametric statistical modeling. This package provides a flexible alternative to BART (Bayesian Additive Regression Trees), using Voronoi tessellations instead of trees for spatial partitioning. The method is particularly effective for spatial data analysis, complex function approximation, and Bayesian regression modeling.

Details

Key features include:

- Machine learning regression with Bayesian inference
- Alternative to BART using Voronoi tessellations
- Spatial data analysis and modeling
- Non-parametric regression capabilities
- Complex function approximation
- Uncertainty quantification through posterior inference

Author(s)

Maintainer: John Paul Gosling <john-paul.gosling@durham.ac.uk> ([ORCID](#))

Authors:

- Adam Stone <adam.stone2@durham.ac.uk> ([ORCID](#))

References

Stone, A. and Gosling, J.P. (2025). AddiVortes: (Bayesian) additive Voronoi tessellations. *Journal of Computational and Graphical Statistics*.

See Also

<https://johnpaulgosling.github.io/AddiVortes/>

AddiVortes

AddiVortes

Description

The AddiVortes function is a Bayesian nonparametric regression model that uses a tessellation to model the relationship between the covariates and the output values. The model uses a backfitting algorithm to sample from the posterior distribution of the output values for each tessellation. The function returns the RMSE value for the test samples.

Usage

```
AddiVortes(
  y,
  x,
  m = 200,
  totalMCMCIter = 1200,
  mcmcBurnIn = 200,
  nu = 6,
  q = 0.85,
  k = 3,
  sd = 0.8,
  Omega = min(3, ncol(x)),
  LambdaRate = 25,
  IntialSigma = "Linear",
  thinning = 1,
  showProgress = TRUE
)
```

Arguments

y	A vector of the output values.
x	A matrix of the covariates.
m	The number of tessellations.
totalMCMCIter	The number of iterations.
mcmcBurnIn	The number of burn in iterations.
nu	The degrees of freedom.
q	The quantile.
k	The number of centres.
sd	The standard deviation.
Omega	Omega/(number of covariates) is the prior probability of adding a dimension.

LambdaRate	The rate of the Poisson distribution for the number of centres.
InitialSigma	The method used to calculate the initial variance.
thinning	The thinning rate.
showProgress	Logical; if TRUE (default), progress bars and messages are shown during fitting.

Value

An AddiVortesFit object containing the posterior samples of the tessellations, dimensions and predictions.

Examples

```
# Simple example with simulated data
set.seed(123)
x <- matrix(rnorm(50), 10, 5)
y <- rnorm(10)
# Fit model with reduced iterations for quick example
fit <- AddiVortes(y, x, m = 5, totalMCMCIter = 50, mcmcBurnIn = 10)
```

new_AddiVortesFit *Create an AddiVortesFit Object*

Description

A constructor for the AddiVortesFit class.

Usage

```
new_AddiVortesFit(
  posteriorTess,
  posteriorDim,
  posteriorSigma,
  posteriorPred,
  xCentres,
  xRanges,
  yCentre,
  yRange,
  inSampleRmse
)
```

Arguments

posteriorTess A list of the posterior samples of the tessellations.
 posteriorDim A list of the posterior samples of the dimensions.
 posteriorSigma A list of the posterior samples of the error variance.
 posteriorPred A list of the posterior samples of the predictions.
 xCentres The centres of the covariates.
 xRanges The ranges of the covariates.
 yCentre The centre of the output values.
 yRange The range of the output values.
 inSampleRmse The in-sample RMSE.

Value

An object of class AddiVortesFit.

plot.AddiVortesFit *Plot Method for AddiVortesFit***Description**

Generates comprehensive diagnostic plots for a fitted AddiVortesFit object. This function creates multiple diagnostic plots including residuals, MCMC traces for sigma, and tessellation complexity over iterations.

Usage

```
## S3 method for class 'AddiVortesFit'
plot(
  x,
  x_train,
  y_train,
  sigma_trace = NULL,
  which = c(1, 2, 3),
  ask = FALSE,
  ...
)
```

Arguments

x An object of class AddiVortesFit, typically the result of a call to AddiVortes().
 x_train A matrix of the original training covariates.
 y_train A numeric vector of the original training true outcomes.
 sigma_trace An optional numeric vector of sigma values from MCMC samples. If not provided, the method will attempt to extract it from the model object.

which	A numeric vector specifying which plots to generate: 1 = Residuals plot, 2 = Sigma trace, 3 = Tessellation complexity trace, 4 = Predicted vs Observed. Default is c(1, 2, 3).
ask	Logical; if TRUE, the user is asked to press Enter before each plot.
...	Additional arguments passed to plotting functions.

Details

The function generates up to four diagnostic plots:

1. **Residuals Plot:** Residuals vs fitted values with smoothed trend line
2. **Sigma Trace:** MCMC trace plot for the error variance parameter
3. **Tessellation Complexity:** Trace of average tessellation size over iterations
4. **Predicted vs Observed:** Scatter plot with credible intervals

Value

This function is called for its side effect of creating plots and returns NULL invisibly.

Examples

```
## Not run:
# Assuming 'fit' is a trained AddiVortesFit object
plot(fit, x_train = x_train_data, y_train = y_train_data)

# Show only specific plots
plot(fit, x_train = x_train_data, y_train = y_train_data, which = c(1, 3))

# With custom sigma trace
plot(fit, x_train = x_train_data, y_train = y_train_data,
      sigma_trace = my_sigma_samples)

## End(Not run)
```

`predict.AddiVortesFit` *Predict Method for AddiVortesFit*

Description

Predicts outcomes for new data using a fitted AddiVortesFit model object. It can return mean predictions, quantiles and optionally calculate the Root Mean Squared Error (RMSE) if true outcomes are provided.

Usage

```
## S3 method for class 'AddiVortesFit'
predict(
  object,
  newdata,
  type = c("response", "quantile"),
  quantiles = c(0.025, 0.975),
  interval = c("credible", "prediction"),
  showProgress = TRUE,
  parallel = TRUE,
  cores = NULL,
  ...
)
```

Arguments

object	An object of class AddiVortesFit, typically the result of a call to AddiVortes().
newdata	A matrix of covariates for the new test set. The number of columns must match the original training data.
type	The type of prediction required. The default "response" gives the mean prediction. The alternative "quantile" returns the quantiles specified by the quantiles argument.
quantiles	A numeric vector of probabilities to compute for the predictions when type = "quantile".
interval	The type of interval calculation. The default "credible" accounts only for uncertainty in the mean (similar to lm's confidence interval). The alternative "prediction" also includes the model's error variance, producing wider intervals (similar to lm's prediction interval).
showProgress	Logical; if TRUE (default), a progress bar is shown during prediction.
parallel	Logical; if TRUE (default), predictions are computed in parallel.
cores	The number of CPU cores to use for parallel processing. If NULL (default), it defaults to one less than the total number of available cores.
...	Further arguments passed to or from other methods (currently unused).

Details

This function relies on the internal helper function applyScaling_internal being available in the environment, which is used by the main AddiVortes function.

When interval = "prediction" and type = "quantile", the function samples additional Gaussian noise with variance equal to the sampled sigma squared from the posterior. This accounts for the inherent variability in individual predictions, not just uncertainty in the mean function. The noise is added in the scaled space before unscaling predictions.

Value

If type = "response", a numeric vector of mean predictions. If type = "quantile", a matrix where each row corresponds to an observation in newdata and each column to a quantile.

Examples

```

# Fit a model
set.seed(123)
X <- matrix(rnorm(100), 20, 5)
Y <- rnorm(20)
fit <- AddiVortes(Y, X, m = 5, totalMCMCIter = 50, mcmcBurnIn = 10)

# New data for prediction
X_new <- matrix(rnorm(25), 5, 5)

# Mean predictions
pred_mean <- predict(fit, X_new, type = "response")

# Credible intervals (uncertainty in mean only)
pred_conf <- predict(fit, X_new, type = "quantile",
                      interval = "credible",
                      quantiles = c(0.025, 0.975))

# Prediction intervals (includes error variance)
pred_pred <- predict(fit, X_new, type = "quantile",
                      interval = "prediction",
                      quantiles = c(0.025, 0.975))

# Prediction intervals are wider than credible intervals
mean(pred_pred[, 2] - pred_pred[, 1]) > mean(pred_conf[, 2] - pred_conf[, 1])

```

print.AddiVortesFit *Print Method for AddiVortesFit*

Description

Prints a summary of a fitted AddiVortesFit object, providing information about the model structure, dimensions, and fit quality similar to the output of a linear model summary.

Usage

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

Arguments

- x An object of class AddiVortesFit, typically the result of a call to AddiVortes().
- ... Further arguments passed to or from other methods (currently unused).

Details

The print method displays:

- The model formula representation
- Number of covariates and posterior samples
- Number of tessellations used
- In-sample RMSE
- Covariate scaling information

Value

The function is called for its side effect of printing model information and returns the input object `x` invisibly.

`summary.AddiVortesFit` *Summary Method for AddiVortesFit*

Description

Provides a detailed summary of a fitted `AddiVortesFit` object, including more comprehensive information than the `print` method.

Usage

```
## S3 method for class 'AddiVortesFit'  
summary(object, ...)
```

Arguments

<code>object</code>	An object of class <code>AddiVortesFit</code> , typically the result of a call to <code>AddiVortes()</code> .
<code>...</code>	Further arguments passed to or from other methods (currently unused).

Value

The function is called for its side effect of printing detailed model information and returns the input object `object` invisibly.

Index

- * **BART**
 AddiVortes-package, 2
- * **bayesian**
 AddiVortes-package, 2
- * **machine-learning**
 AddiVortes-package, 2
- * **package**
 AddiVortes-package, 2
- * **regression**
 AddiVortes-package, 2
- * **spatial**
 AddiVortes-package, 2
- * **tessellation**
 AddiVortes-package, 2

 AddiVortes, 3
 AddiVortes-package, 2

 new_AddiVortesFit, 4

 plot.AddiVortesFit, 5
 predict.AddiVortesFit, 6
 print.AddiVortesFit, 8

 summary.AddiVortesFit, 9