

Package ‘spCP’

May 9, 2026

Type Package

Title Spatially Varying Change Points

Version 1.4.0

Date 2025-09-29

Description Implements a spatially varying change point model with unique intercepts, slopes, variance intercepts and slopes, and change points at each location. Inference is within the Bayesian setting using Markov chain Monte Carlo (MCMC). The response variable can be modeled as Gaussian (no nugget), probit or Tobit link and the five spatially varying parameter are modeled jointly using a multivariate conditional autoregressive (MCAR) prior. The MCAR is a unique process that allows for a dissimilarity metric to dictate the local spatial dependencies. Full details of the package can be found in the accompanying vignette. Furthermore, the details of the package can be found in the corresponding paper published in *Spatial Statistics* by Berchuck et al (2019): “A spatially varying change points model for monitoring glaucoma progression using visual field data”, <doi:10.1016/j.spasta.2019.02.001>.

License GPL (>= 2)

Encoding UTF-8

LazyData true

RoxygenNote 7.3.2

NeedsCompilation yes

Depends R (>= 3.0.2)

Imports graphics, grDevices, msm (>= 1.0.0), mvtnorm (>= 1.0-0), Rcpp (>= 0.12.9), stats, utils

Suggests coda, classInt, knitr, rmarkdown, womblR (>= 1.0.3)

LinkingTo Rcpp, RcppArmadillo (>= 0.7.500.0.0)

VignetteBuilder knitr

Author Samuel I. Berchuck [aut, cre]

Maintainer Samuel I. Berchuck <sib2@duke.edu>

Repository CRAN

Date/Publication 2025-09-29 20:20:02 UTC

Contents

diagnostics	2
GarwayHeath	3
HFAII_Queen	4
HFAII_QueenHF	4
HFAII_Rook	5
is.spCP	5
PlotCP	6
predict.spCP	8
spCP	8
VFSeries	12

Index **13**

diagnostics	<i>diagnostics</i>
-------------	--------------------

Description

Calculates diagnostic metrics using output from the [spCP](#) model.

Usage

```
diagnostics(
  obj,
  diags = c("dic", "dinf", "waic"),
  keepDeviance = FALSE,
  keepPPD = FALSE
)
```

Arguments

obj	A spCP model object for which diagnostics are desired from.
diags	A vector of character strings indicating the diagnostics to compute. Options include: Deviance Information Criterion ("dic"), d-infinity ("dinf") and Watanabe-Akaike information criterion ("waic"). At least one option must be included. Note: The probit model cannot compute the DIC or WAIC diagnostics due to computational issues with computing the multivariate normal CDF.
keepDeviance	A logical indicating whether the posterior deviance distribution is returned (default = FALSE).
keepPPD	A logical indicating whether the posterior predictive distribution at each observed location is returned (default = FALSE).

Details

To assess model fit, DIC, d-infinity and WAIC are used. DIC is based on the deviance statistic and penalizes for the complexity of a model with an effective number of parameters estimate p_D (Spiegelhalter et al 2002). The d-infinity posterior predictive measure is an alternative diagnostic tool to DIC, where $d\text{-infinity}=P+G$. The G term decreases as goodness of fit increases, and P, the penalty term, inflates as the model becomes over-fit, so small values of both of these terms and, thus, small values of d-infinity are desirable (Gelfand and Ghosh 1998). WAIC is invariant to parametrization and is asymptotically equal to Bayesian cross-validation (Watanabe 2010). $WAIC = -2 * (lppd - p_waic_2)$. Where $lppd$ is the log pointwise predictive density and p_waic_2 is the estimated effective number of parameters based on the variance estimator from Vehtari et al. 2016. (p_waic_1 is the mean estimator).

Value

`diagnostics` returns a list containing the diagnostics requested and possibly the deviance and/or posterior predictive distribution objects.

Author(s)

Samuel I. Berchuck

References

- Gelfand, A. E., & Ghosh, S. K. (1998). Model choice: a minimum posterior predictive loss approach. *Biometrika*, 1-11.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4), 583-639.
- Vehtari, A., Gelman, A., & Gabry, J. (2016). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 1-20.
- Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11(Dec), 3571-3594.

GarwayHeath

Garway-Heath angles for the HFA-II

Description

These Garway-Heath angles are used as the dissimilarity metric when implementing the boundary detection model for a longitudinal series of visual fields.

Usage

`data(GarwayHeath)`

Format

A vector with length 54, where each entry represents the angle (in degrees) that the underlying retinal nerve fiber enters the optic nerve head. The measure ranges from 0-360, where 0 is designated at the 9-o'clock position (right eye) and angles are counted counter clockwise. These angles are estimates for the Humphrey Field Analyzer-II (Carl Zeiss Meditec Inc., Dublin, CA). The 26th and 35th entries are missing as they correspond to a natural blind spot.

References

Garway-Heath, et al. (2000). Ophthalmology 107:10:1809–1815. ([PubMed](#))

HFAII_Queen

HFAII Queen Adjacency Matrix

Description

Binary adjacency matrix for the Humphrey Field Analyzer-II (Carl Zeiss Meditec Inc., Dublin, CA)

Usage

`data(HFAII_Queen)`

Format

This adjacency matrix is formatted using queen neighbor criteria, meaning two locations on the visual field are only considered neighbors if they share an edge or corner. The adjacency matrix is a 54 x 54 dimensional binary object with zeros on the diagonal and the column and row sums are equal to the number of neighbors.

HFAII_QueenHF

HFAII Queen Hemisphere Adjacency Matrix

Description

Binary adjacency matrix for the Humphrey Field Analyzer-II (Carl Zeiss Meditec Inc., Dublin, CA)

Usage

`data(HFAII_QueenHF)`

Format

This adjacency matrix is formatted using queen neighbor criteria, meaning two locations on the visual field are only considered neighbors if they share an edge or corner. An additional criterion is included that locations are not considered neighbors if they fall within different hemispheres on the visual field. The adjacency matrix is a 54 x 54 dimensional binary object with zeros on the diagonal and the column and row sums are equal to the number of neighbors.

`HFAII_Rook`*HFAII Rook Adjacency Matrix*

Description

Binary adjacency matrix for the Humphrey Field Analyzer-II (Carl Zeiss Meditec Inc., Dublin, CA)

Usage

```
data(HFAII_Rook)
```

Format

This adjacency matrix is formatted using rook neighbor criteria, meaning two locations on the visual field are only considered neighbors if they share an edge. The adjacency matrix is a 54 x 54 dimensional binary object with zeros on the diagonal and the column and row sums are equal to the number of neighbors.

`is.spCP`*is.spCP*

Description

`is.spCP` is a general test of an object being interpretable as a `spCP` object.

Usage

```
is.spCP(x)
```

Arguments

`x` object to be tested.

Details

The `spCP` class is defined as the regression object that results from the `spCP` regression function.

 PlotCP

PlotCP

Description

Plots estimated visual field sensitivities using change point model.

Usage

```
PlotCP(
  object,
  data,
  location = "Location",
  time = "Time",
  dls = "DLS",
  line = TRUE,
  ci = TRUE,
  lwd = 1,
  lty = 1,
  col = 2,
  ci.lwd = 1,
  ci.lty = 2,
  ci.col = 2,
  cp.line = FALSE,
  cp.ci = FALSE,
  cp.lwd = 1,
  cp.lty = 1,
  cp.col = 4,
  cp.ci.lwd = 1,
  cp.ci.lty = 2,
  cp.ci.col = 4,
  main = "Estimated visual field sensitivity\nusing change points",
  xlab = "Years from baseline visit",
  ylab = "Sensitivity (dB)"
)
```

Arguments

object	a spCP regression object.
data	a dataframe containing the raw sensitivities.
location	a character string indicating the column name in data of the variable of locations on the visual field.
time	a character string indicating the column name in data of the variable of visual field testing times.
dls	a character string indicating the column name in data of the variable of the raw visual field sensitivities.

line	logical, determines if there are regression lines printed (default = TRUE).
ci	logical, determines if there are confidence intervals printed (default = TRUE).
lwd	integer, specifies the width of the regression line (default = 1).
lty	integer, specifies the type of regression line (default = 1).
col	color for the regression line, either character string corresponding to a color or a integer (default = "red").
ci.lwd	integer, specifies the width of the confidence intervals (default = 1).
ci.lty	integer, specifies the type of confidence intervals (default = 2).
ci.col	color for the confidence intervals for the regression line, either character string corresponding to a color or a integer (default = "red").
cp.line	logical, determines if there the change point is printed (default = FALSE).
cp.ci	logical, determines if there the change point confidence intervals are printed (default = FALSE).
cp.lwd	integer, specifies the width of the change point line (default = 1).
cp.lty	integer, specifies the type of change point line (default = 1).
cp.col	color for the optional change point, either character string corresponding to a color or a integer (default = "blue").
cp.ci.lwd	integer, specifies the width of the change point confidence intervals (default = 1).
cp.ci.lty	integer, specifies the type of change point confidence intervals (default = 2).
cp.ci.col	color for the confidence intervals of the optional change point, either character string corresponding to a color or a integer (default = "blue").
main	an overall title for the plot.
xlab	a title for the x axis.
ylab	a title for the y axis.

Details

PlotCP is used in the application of glaucoma progression. The function is capable of plotting the observed DLS values across the visual field, along with the estimated mean process (with 95 percent credible intervals) and the estimated mean posterior change point location (with 95 percent credible intervals).

Author(s)

Samuel I. Berchuck

predict.spCP	<i>predict.spCP</i>
--------------	---------------------

Description

Predicts future observations from the [spCP](#) model.

Usage

```
## S3 method for class 'spCP'
predict(object, NewTimes, ...)
```

Arguments

object	a spCP model object for which predictions are desired from.
NewTimes	a numeric vector including desired time(s) points for prediction.
...	other arguments.

Details

predict.spCP uses Bayesian krigging to obtain posterior samples from future time points.

Value

predict.spCP returns a list containing the following objects.

Y A list containing a matrix of predictions for each future time point. Each matrix has one column for each location and contains posterior samples obtained by Bayesian krigging.

Author(s)

Samuel I. Berchuck

spCP	<i>MCMC sampler for spatially varying change point model.</i>
------	---

Description

spCP is a Markov chain Monte Carlo (MCMC) sampler for a spatially varying change point model with spatially varying slopes, intercepts, and unique variances at each spatial-temporal location. The model is implemented using a Bayesian hierarchical framework.

Implements a spatially varying change point model with unique intercepts, slopes, variance intercepts and slopes, and change points at each location. Inference is within the Bayesian setting using Markov chain Monte Carlo (MCMC). The response variable can be modeled as Gaussian (no nugget), probit or Tobit link and the five spatially varying parameter are modeled jointly using a

multivariate conditional autoregressive (MCAR) prior. The MCAR is a unique process that allows for a dissimilarity metric to dictate the local spatial dependencies. Full details of the package can be found in the accompanying vignette. Furthermore, the details of the package can be found in the corresponding paper on arXiv by Berchuck et al (2018): "A spatially varying change points model for monitoring glaucoma progression using visual field data", [arXiv:1811.11038](https://arxiv.org/abs/1811.11038).

Usage

```
spCP(
  Y,
  DM,
  W,
  Time,
  Starting = NULL,
  Hypers = NULL,
  Tuning = NULL,
  MCMC = NULL,
  Family = "tobit",
  Weights = "continuous",
  Distance = "circumference",
  Rho = 0.99,
  ScaleY = 10,
  ScaleDM = 100,
  Seed = 54
)
```

Arguments

Y	An N dimensional vector containing the observed outcome data. Here, $N = M * Nu$, where M represents the number of spatial locations and Nu the number of temporal visits. The observations in Y must be first ordered spatially and then temporally, meaning the first M observations in Y should come from the initial time point.
DM	An M dimensional vector containing a dissimilarity metric for each spatial location. The order of the spatial locations must match the order from Y.
W	An M x M dimensional binary adjacency matrix for dictating the spatial neighborhood structure.
Time	A Nu dimensional vector containing the observed time points for each vector of outcomes in increasing order.
Starting	Either NULL or a list containing starting values to be specified for the MCMC sampler. If NULL is not chosen then none, some or all of the starting values may be specified. When NULL is chosen then default starting values are automatically generated. Otherwise a list must be provided with names Delta, Alpha or Sigma containing appropriate objects. Delta must be a 5 dimensional vector, Sigma a 5 x 5 dimensional matrix and Alpha a scalar.

Hypers	<p>Either NULL or a list containing hyperparameter values to be specified for the MCMC sampler. If NULL is not chosen then none, some or all of the hyperparameter values may be specified.</p> <p>When NULL is chosen then default hyperparameter values are automatically generated. These default hyperparameters are described in detail in (Berchuck et al.). Otherwise a list must be provided with names Delta, Sigma or Alpha containing further hyperparameter information. These objects are themselves lists and may be constructed as follows.</p> <p>Delta is a list with one object, Kappa2. Kappa2 represents the prior variance and must be a scalar. Sigma is a list with two objects, Xi and Psi. Xi represents the degrees of freedom parameter for the inverse-Wishart hyperprior and must be a real number scalar, while Psi represents the scale matrix and must be a 5 x 5 dimensional positive definite matrix.</p> <p>Alpha is a list with two objects, AAlpha and BAlpha. AAlpha represents the lower bound for the uniform hyperprior, while BAlpha represents the upper bound. The bounds must be specified carefully.</p>
Tuning	<p>Either NULL or a list containing tuning values to be specified for the MCMC Metropolis steps. If NULL is not chosen then all of the tuning values must be specified.</p> <p>When NULL is chosen then default tuning values are automatically generated to 1. Otherwise a list must be provided with names Lambda0Vec, Lambda1Vec, EtaVec and Alpha. Lambda0Vec, Lambda1Vec, and EtaVec must be M dimensional vectors and Alpha a scalar. Each containing tuning variances for their corresponding Metropolis updates.</p>
MCMC	<p>Either NULL or a list containing input values to be used for implementing the MCMC sampler. If NULL is not chosen then all of the MCMC input values must be specified.</p> <p>NBurn: The number of sampler scans included in the burn-in phase. (default = 10,000)</p> <p>NSims: The number of post-burn-in scans for which to perform the sampler. (default = 100,000)</p> <p>NThin: Value such that during the post-burn-in phase, only every NThin-th scan is recorded for use in posterior inference (For return values we define, NKeep = NSims / NThin (default = 10).</p> <p>NPilot: The number of times during the burn-in phase that pilot adaptation is performed (default = 20)</p>
Family	<p>Character string indicating the distribution of the observed data. Options include: "normal", "probit", "tobit".</p>
Weights	<p>Character string indicating the type of weight used. Options include: "continuous" and "binary".</p>
Distance	<p>Character string indicating the distance metric for computing the dissimilarity metric. Options include: "euclidean" and "circumference".</p>
Rho	<p>A scalar in $(0, 1)$ that dictates the magnitude of local spatial sharing. By default it is fixed at 0.99 as suggested by Lee and Mitchell (2012).</p>

ScaleY	A positive scalar used for scaling the observed data, Y. This is used to aid numerically for MCMC convergence, as scaling large observations often stabilizes chains. By default it is fixed at 10.
ScaledM	A positive scalar used for scaling the dissimilarity metric distances, DM. This is used to aid numerically for MCMC convergence. as scaling spatial distances is often used for improved MCMC convergence. By default it is fixed at 100.
Seed	An integer value used to set the seed for the random number generator (default = 54).

Details

Details of the underlying statistical model proposed by proposed by Berchuck et al. 2018. are forthcoming.

Value

spCP returns a list containing the following objects

beta0 NKeep x M matrix of posterior samples for beta0. The s-th column contains posterior samples from the the s-th location.

beta1 NKeep x M matrix of posterior samples for beta1. The s-th column contains posterior samples from the the s-th location.

lambda0 NKeep x M matrix of posterior samples for lambda0. The s-th column contains posterior samples from the the s-th location.

lambda1 NKeep x M matrix of posterior samples for lambda1. The s-th column contains posterior samples from the the s-th location.

eta NKeep x M matrix of posterior samples for eta. The s-th column contains posterior samples from the the s-th location.

theta NKeep x M matrix of posterior samples for theta. The s-th column contains posterior samples from the the s-th location.

delta NKeep x 5 matrix of posterior samples for delta. The columns have names that describe the samples within them.

sigma NKeep x 15 matrix of posterior samples for Sigma. The columns have names that describe the samples within them. The row is listed first, e.g., Sigma32 refers to the entry in row 3, column 2.

alpha NKeep x 1 matrix of posterior samples for Alpha.

metropolis $(3 * M + 1) \times 3$ matrix of metropolis acceptance rates, updated tuners, and original tuners that result from the pilot adaptation. The first M correspond to the Lambda0Vec parameters, the next M correspond to the Lambda1Vec, the next M correspond to the EtaVec parameters and the last row give the Alpha values.

runtime A character string giving the runtime of the MCMC sampler.

datobj A list of data objects that are used in future spCP functions and should be ignored by the user.

dataug A list of data augmentation objects that are used in future spCP functions and should be ignored by the user.

Author(s)

Samuel I. Berchuck <sib2@duke.edu>

References

Reference for Berchuck et al. 2018 is forthcoming.

VFSeries

Visual field series for one patient.

Description

A dataset containing 9 visual field series from a patient of the Vein Pulsation Study Trial in Glaucoma and the Lions Eye Institute trial registry, Perth, Western Australia.

Usage

```
data(VFSeries)
```

Format

A data frame with 486 rows and 4 variables:

Visit The visual field test visit number, (1, 2, ... , 9).

DLS The observed outcome variable, differential light sensitivity (DLS).

Time The time of the visual field test (in days from baseline).

Location The location on the visual field of a Humphrey Field Analyzer-II (Carl Zeiss Meditec Inc., Dublin, CA) (1, 2, ... , 54).

Source

<https://anzctr.org.au/Trial/Registration/TrialReview.aspx?ACTRN=12608000274370>

Index

* datasets

- GarwayHeath, [3](#)
- HFAII_Queen, [4](#)
- HFAII_QueenHF, [4](#)
- HFAII_Rook, [5](#)
- VFSeries, [12](#)

diagnostics, [2](#)

GarwayHeath, [3](#)

HFAII_Queen, [4](#)

HFAII_QueenHF, [4](#)

HFAII_Rook, [5](#)

is.spCP, [5](#)

PlotCP, [6](#)

predict.spCP, [8](#)

spCP, [2](#), [5](#), [8](#), [8](#)

spCP-package (spCP), [8](#)

VFSeries, [12](#)