Package: hkevp (via r-universe)

October 25, 2026

Type Package

Title Spatial Extreme Value Analysis with the Hierarchical Model of Reich and Shaby (2012)

Version 1.1.6

Description Several procedures for the hierarchical kernel extreme value process of Reich and Shaby (2012) <DOI:10.1214/12-AOAS591>, including simulation, estimation and spatial extrapolation. The spatial latent variable model <DOI:10.1214/11-STS376> is also included.

BugReports https://github.com/lbelzile/hkevp/issues/

License GPL

Imports Rcpp (>= 0.12.16)

LinkingTo Rcpp, RcppArmadillo

Encoding UTF-8

RoxygenNote 7.3.3

NeedsCompilation yes

Author Quentin Sebille [aut], Leo Belzile [cre]

Maintainer Leo Belzile <belzilel@gmail.com>

Repository https://cran.r-universe.dev

Date/Publication 2025-09-26 11:20:02 UTC

RemoteUrl https://github.com/cran/hkevp

RemoteRef HEAD

RemoteSha 49db3e70d30c7d480be46c53b2d126a8b8392959

Contents

extrapol.gev 2
extrapol.return.level 3
hkevp 5
hkevp.expmeasure 8

hkevp.fit 9
hkevp.predict 13
hkevp.rand 15
latent.fit 17
mcmc.fun 20
mcmc.plot 22
return.level 23

Index **24**

`extrapol.gev` *Spatial extrapolation of GEV parameters with the HKEVP*

Description

Predictive distributions of the GEV parameters at a set of ungauged sites (targets), given the output from the MCMC procedures `hkevp.fit` or `latent.fit`. See details.

Usage

```
extrapol.gev(fit, targets, targets.covariates)
```

Arguments

`fit`	Output from the `hkevp.fit` procedure.
`targets`	A matrix of real values giving the spatial coordinates of the ungauged positions. Each row corresponds to an ungauged position.
`targets.covariates`	A matrix of real values giving the spatial covariates of the ungauged positions. Must match with the covariates used in `hkevp.fit` or `latent.fit`.

Details

Since the GEV parameters are modelled with latent Gaussian processes, spatial extrapolation of the marginal distributions at target positions $(s_1^*, ..., s_k^*)$ is performed with simple kriging. Estimation is done at each MCMC step to produce a sample of the predictive distribution.

Value

A named list of three elements: `loc`, `scale`, `shape`. Each one is a matrix with columns corresponding to targets positions.

Author(s)

Quentin Sebille

See Also

`extrapol.return.level`

Examples

```
# Simulation of HKEVP:
sites <- as.matrix(expand.grid(1:3,1:3))
loc <- sites[,1]*10
scale <- 3
shape <- 0
alpha <- .4
tau <- 1
ysim <- hkevp.rand(10, sites, sites, loc, scale, shape, alpha, tau)

# HKEVP fit:
fit <- hkevp.fit(ysim, sites, niter = 1000)

## Extrapolation:
targets <- matrix(1.5, 1, 2)
gev.targets <- extrapol.gev(fit, targets)

## True vs predicted:
predicted <- sapply(gev.targets, median)
sd.predict <- sapply(gev.targets, sd)
true <- c(targets[,1]*10, scale, shape)
# cbind(true, predicted, sd.predict)
```

---

extrapol.return.level *Spatial extrapolation of a return level.*

Description

Predictive distribution of a T-years return level at ungauged positions (targets), given the output from the MCMC procedures `hkevp.fit` or `latent.fit`.

Usage

```
extrapol.return.level(period, fit, targets, targets.covariates)
```

Arguments

`period`	An integer indicating the wished return period T.
`fit`	Output from the `hkevp.fit` procedure.
`targets`	A matrix of real values giving the spatial coordinates of the ungauged positions. Each row corresponds to an ungauged position.
`targets.covariates`	A matrix of real values giving the spatial covariates of the ungauged positions. Must match with the covariates used in `hkevp.fit` or `latent.fit`.

Details

Spatial extrapolation of the return level at target positions $(s_1^*, ..., s_k^*)$ is a two-step procedure:

- Estimation of the predictive distribution for GEV parameters at $(s_1^*, ..., s_k^*)$, by using `{extrapol.gev}`.
- Computation of the associated return level for each state of the predictive distribution.

Value

A matrix of predictive sample. Each column corresponds to a target position and each row to a predictive draw.

Author(s)

Quentin Sebille

See Also

`extrapol.gev`

Examples

```
# Simulation of HKEVP:
sites <- as.matrix(expand.grid(1:3,1:3))
knots <- sites
loc <- sites[,1]*10
scale <- 1
shape <- .2
alpha <- .4
tau <- 1
ysim <- hkevp.rand(10, sites, knots, loc, scale, shape, alpha, tau)

# HKEVP fit:
fit <- hkevp.fit(ysim, sites, niter = 1000)

## Extrapolation of the 100-years return level (may need more iterations and burn-in/nthin):
targets <- as.matrix(expand.grid(1.5:2.5,1.5:2.5))
pred.sample <- extrapol.return.level(100, fit, targets)
pred.mean <- apply(pred.sample, 2, mean)
pred.sd <- apply(pred.sample, 2, sd)
true <- return.level(100, targets[,1]*10, scale, shape)
# cbind(true, pred.mean, pred.sd)
```

hkevp *A hierarchical model for spatial extreme values*

Description

The HKEVP of *Reich and Shaby (2012)* is a hierarchical model which can be fitted on point-referenced block maxima across a region of space. Its acronym stands for Hierarchical Kernel Extreme Value Model.

This model fits both marginal GEV parameters and a conditional spatial dependence structure in a Bayesian framework. Estimation of all parameters is performed with a Metropolis-within-Gibbs algorithm that returns samples of posterior distributions, given prior distributions and data. See details.

Simulation and fitting procedure for this model are provided along with several other tools such as spatial extrapolation and conditional simulation, which is used for instance in *Shaby and Reich (2012)*.

A particular simpler case of the HKEVP, defined in Davison et al. (2012), is also available. This model, referred as the latent variable model, assumes conditional independence of the block maxima data (i.e. no dependence structure).

Spatial modelling of extreme values with general max-stable proceses is taken care of in other R libraries such as `RandomFields` and `SpatialExtremes`.

Details

The functions included in `hkevp` package are listed below:

1. `hkevp.fit` (resp. `latent.fit`): fits the HKEVP (resp. the latent variable model) to spatial block maxima data.
2. `hkevp.rand`: simulates data from the HKEVP.
3. `hkevp.expmeasure`: computes the exponent measure of the HKEVP. See details below.
4. `mcmc.fun`: applies a function to the main Markov chains obtained by `hkevp.fit`. Useful to compute the posterior means or quantiles for instance.
5. `mcmc.plot`: plots the resulting Markov chains, in order to visually assess convergence.
6. `extrapol.gev` (resp. `extrapol.return.level`) : computes the predictive distribution of the GEV parameters (resp. a return level) at ungauged positions.
7. `hkevp.predict`: predictive distribution of the spatial process at ungauged stes with the HKEVP given observations.

The HKEVP of *Reich and Shaby (2012)* is a hierarchical spatial max-stable model for extreme values. Max-stable models arise as the limiting distribution of renormalized maxima of stochastic processes and they generalize the Extreme Value Theory (EVT) to the infinite-dimensional case. For more information about EVT and max-stable processes, see for instance *Beirlant et al. (2004)*, *de Haan and Ferreira (2006)* and *Coles (2001)*. For an emphasis on statistical inference on spatial extremes, see for instance *Cooley et al. (2012)* and *Davison et al. (2012)*.

Let $Y(\cdot)$ be the process of block maxima recorded over a spatial region. Assume this process is max-stable, then all marginal distributions are necessarily GEV, i.e:

$$Y(s) \sim GEV\{\mu(s), \sigma(s), \xi(s)\}$$

with location, scale and shape parameters $\mu(s)$, $\sigma(s)$ and $\xi(s)$ respectively, indexed by position s in space. Without loss of generality, one may look at the *simple max-stable process* $Z(\cdot)$ with $GEV(1, 1, 1)$ margins, where spatial dependence is contained unconditionally of the marginals.

The HKEVP is thus defined by assuming that there exists $\alpha \in (0, 1]$ and a set of knots $\{v_1, ..., v_L\}$ with associated kernels $\{\omega_1(\cdot), ..., \omega_L(\cdot)\}$ such that:

$$Z(s) = U(s)\theta(s)$$

where $U(s)$ is a spatially-independent process with $GEV(1, \alpha, \alpha)$ margins and

$$\theta(s) = \left[\sum_{\ell=1}^{L} A_\ell \omega_\ell(s)^{1/\alpha}\right]^\alpha$$

is the *residual dependence process*, defined with a random variable $A_\ell \sim PS(\alpha)$, the positive stable distribution with characteristic exponent α. Note that the kernels must satisfy the condition $\sum_{\ell=1}^{L} \omega_\ell(s) = 1$, for all s.

Under those assumptions, *Reich and Shaby (2012)* showed that this model lead to an explicit formula for the distribution of a joint vector of maxima, which is not the case for max-stable processes since no general parametric representation exists (cf. *de Haan and Ferreira (2006)*). The joint distribution of the vector $\{Z(s_1), \ldots, Z(s_n)\}$ under the HKEVP assumptions can indeed be written:

$$P\{Z_1(s_1) < z_1, \ldots, Z_n(s_n) < z_n\} = \sum_{\ell=1}^{L} \left[\sum_{i=1}^{n} \left(\frac{\omega_\ell(s_i)}{z_i}\right)^{1/\alpha}\right]^\alpha .$$

The HKEVP can be seen as an approximation of the *Smith (1990)* model, but with an additional dependence parameter α which controls the strength of spatial dependence. Low value for α leads to a strong dependence structure while $\alpha = 1$ means independence.

In the article of *Reich and Shaby (2012)*, the kernels are chosen to be standardized Gaussian kernels, i.e:

$$\omega_\ell(s) = \frac{K(s_\ell | v_\ell, \tau)}{\sum_{j=1}^{L} K(s_j | v_j, \tau)} ,$$

where $K(\cdot|v, \tau)$ is the Gaussian kernel centered at v_ℓ and τ is a bandwidth parameter.

Conditionally on the marginal and the dependence parameters, we obtain independent responses which allow the computation of the likelihood. Since the model is structured via multiple layers, Bayesian inference is preferred. Details of the MCMC algorithm can be found in *Reich and Shaby (2012)*.

Marginal parameters $\mu(s)$, $\sigma(s)$ and $\xi(s)$ are modelled through latent Gaussian processes.

A simpler version of the HKEVP, namely the latent variable model of Davison et al. (2012), is also available in this package. This model assumes conditional independence and can therefore be seen as a special case of the HKEVP, with the condition $\alpha = 1$.

Note

I would like to thank Brian Reich and Benjamin Shaby for their help regarding the implementation of the inference procedure of the HKEVP, and Mathieu Ribatet for introducing me to the Bayesian world. I also acknowledge Electricite de France (EDF) for the financial support and the helpful discussions around the HKEVP with Anne Dutfoy, Marie Gallois, Thi Thu Huong Hoang and Sylvie Parey. Finally, I thank my two PhD supervisors Anne-Laure Fougeres and Cecile Mercadier for their reviews of this package.

Author(s)

Quentin Sebille

References

Beirlant, J., Goegebeur, Y., Segers, J. J. J., & Teugels, J. (2004). Statistics of Extremes: Theory and Applications. <DOI:10.1002/0470012382>

Coles, S. (2001). An Introduction to Statistical Modeling of Extreme Values. Springer Science & Business Media. <10.1007/978-1-4471-3675-0>

Cooley, D., Cisewski, J., Erhardt, R. J., Jeon, S., Mannshardt, E., Omolo, B. O., & Sun, Y. (2012). A survey of spatial extremes: Measuring spatial dependence and modeling spatial effects. Revstat, 10(1), 135-165.

Davison, A. C., Padoan, S. A., & Ribatet, M. (2012). Statistical modeling of spatial extremes. Statistical Science, 27(2), 161-186. <DOI:10.1214/11-STS376>

de Haan, L., & Ferreira, A. (2006). Extreme Value Theory: An Introduction. Springer Science & Business Media. <DOI:10.1007/0-387-34471-3>

Reich, B. J., & Shaby, B. A. (2012). A hierarchical max-stable spatial model for extreme precipitation. The annals of applied statistics, 6(4), 1430. <DOI:10.1214/12-AOAS591>

Shaby, B. A., & Reich, B. J. (2012). Bayesian spatial extreme value analysis to assess the changing risk of concurrent high temperatures across large portions of European cropland. Environmetrics, 23(8), 638-648. <DOI:10.1002/env.2178>

Smith, R. L. (1990). Max-stable processes and spatial extremes. Unpublished manuscript.

See Also

Useful links:

- Report bugs at https://github.com/lbelzile/hkevp/issues/

`hkevp.expmeasure` *Exponent measure of the HKEVP*

Description

Exponent measure $V(z_1, ..., z_n)$ of the HKEVP of *Reich and Shaby (2012)*, with given model parameters or output from `hkevp.fit` or `latent.fit`.

Usage

```
hkevp.expmeasure(z, sites, knots, alpha, tau, fit)
```

Arguments

`z`	The vector $(z_1, ..., z_n)$ where the exponent measure is computed. Can be of length one and thus corresponds then to $(z, ..., z)$.
`sites`	The coordinates of the sites where the data are observed. Each row correspond to a site position.
`knots`	The coordinates of the knots in the HKEVP. By default, the positions of the knots coincide with the positions of the sites.
`alpha`	The dependence parameter α of the HKEVP: a single value in (0,1].
`tau`	The bandwidth parameter τ of the kernel functions in the HKEVP: a positive value.
`fit`	Output from the `hkevp.fit` procedure.

Details

The exponent measure describes the spatial dependence structure of a max-stable process, independently from the values of the marginal parameters. If $Z(\cdot)$ is a simple max-stable process, i.e. with unit GEV(1,1,1) margins, recorded at the set of sites $(s_1, \ldots, s_n)$, its joint cumulative probability density function is given by:

$$P\{Z(s_1) \leq z_1, \ldots, Z(s_n) \leq z_n\} = \exp(-V(z_1, \ldots, z_n)) \, ,$$

where V is the so-called exponent measure. For the HKEVP, the exponent measure is explicit for any number n of sites:

$$V(z_1, \ldots, z_n) = \sum_{\ell=1}^{L} \left[\sum_{i=1}^{n} \left(\frac{\omega_\ell(s_i)}{z_i} \right)^{1/\alpha} \right]^{\alpha} \, .$$

If argument `fit` is provided, the predictive distribution of

$$V(z_1, \ldots, z_n)$$

is computed. If not, the function uses arguments `sites`, `knots`, `alpha`, and `tau`.

Value

Either a vector if argument `fit` is provided, or a single value.

Author(s)

Quentin Sebille

References

Reich, B. J., & Shaby, B. A. (2012). A hierarchical max-stable spatial model for extreme precipitation. The annals of applied statistics, 6(4), 1430. <DOI:10.1214/12-AOAS591>

Examples

```
sites <- as.matrix(expand.grid(1:3,1:3))
loc <- sites[,1]*10
scale <- 3
shape <- 0
alpha <- .4
tau <- 1
ysim <- hkevp.rand(10, sites, sites, loc, scale, shape, alpha, tau)

# HKEVP fit:
fit <- hkevp.fit(ysim, sites, niter = 1000)

predict.em <- hkevp.expmeasure(1, fit = fit)
true.em <- hkevp.expmeasure(1, sites, sites, alpha, tau)
# plot(predict.em, ylim = range(predict.em, true.em), type = "l")
# abline(h = true.em, col = 2, lwd = 2)
```

`hkevp.fit` *Fitting procedure of the HKEVP with MCMC algorithm*

Description

Metropolis-within-Gibbs algorithm that returns samples from posterior distribution of all the parameters of the HKEVP.

Most of the input parameters have default values, so that the procedure can be easily handled. However, convergence of the Markov chains should be assessed by using `mcmc.plot` for instance. The experimented user can set initial states, prior hyperparameters along with the magnitude of the MCMC jumps.

Usage

```
hkevp.fit(
  y,
  sites,
  knots,
  niter,
  nburn,
  nthin,
  quiet,
  trace,
  fit.margins,
  gev.vary,
  spatial.covariates,
  log.scale,
  correlation,
  mcmc.init,
  mcmc.prior,
  mcmc.jumps
)
```

Arguments

y	A matrix of observed block maxima. Each column corresponds to a site position.
sites	The coordinates of the sites where the data are observed. Each row corresponds to a site position.
knots	The coordinates of the knots in the HKEVP. By default, the positions of the knots coincide with the positions of the sites.
niter	The number of MCMC iterations.
nburn	The number of first MCMC iterations that are discarded. Zero by default.
nthin	The size of the MCMC thinning. One by default (i.e. no thinning).
quiet	A logical indicating if the progression of the routine should be displayed. TRUE by default.
trace	If quiet is FALSE, the log-likelihood of the model is displayed each block of trace MCMC steps to observe fitting progression.
fit.margins	A logical that indicates if the GEV parameters should be fitted along with the dependence structure. TRUE by default.
gev.vary	A logical vector of size three indicating if the GEV parameters (respectively the location, the scale and the shape) are spatially-varying. If not (by default for the shape), the parameter is the same at each position.
spatial.covariates	A numerical matrix of spatial covariates. Each row corresponds to a site position. See details.
log.scale	A logical value indicating if the GEV scale parameter σ is modelled by its log. FALSE by default. See details.

`correlation`	A character string indicating the form of the correlation function associated to the latent Gaussian processes that describes the marginal parameters. Must be one of `"expo"`, `"gauss"`, `"mat32"` (By default) and `"mat52"`, respectively corresponding to the exponential, Gaussian, Matern-3/2 and Matern-5/2 correlation functions.
`mcmc.init`	A named list indicating the initial states of the chains. See details.
`mcmc.prior`	A named list indicating the hyperparameters of the prior distributions. See details.
`mcmc.jumps`	A named list indicating the amplitude of the jumps to propose the MCMC candidates. See details.

Details

Details of the MCMC procedure are presented in *Reich and Shaby (2012)*. This function follows the indications and the choices of the authors, with the exception of several small changes:

- The scale parameter σ can be modelled like the two other marginal parameters as in *Davison et al. (2012)* or by its logarithm as in *Reich and Shaby (2012)*. For this, use the argument `log.scale`, set to FALSE by default.
- The Inverse-Gamma prior distributions defined for the bandwith parameter τ and for the ranges λ of the latent processes are replaced by a Beta distribution over the interval $[0, 2D_{max}]$, where D_{max} stands for the maximum distance between two sites.

The procedure can be used normally with `fit.margins = TRUE` (default) or by assuming that the observed process had GEV(1,1,1) margins already and thus ignoring the marginal estimation.

If the margins are estimated and the parameters are assumed spatially-varying, the user can provide spatial covariates to fit the mean of the latent Gaussian processes. Recall for instance for the GEV location parameter that:

$$\mu(s) = \beta_{0,\mu} + \beta_{1,\mu} c_1(s) + ... + \beta_{p,\mu} c_p(s) \ .$$

The given matrix `spatial.covariates` that represents the $c_i(s)$ elements should have the first column filled with ones to account for the intercept β_0.

The arguments `mcmc.init`, `mcmc.prior` and `mcmc.jumps` are named list that have default values. The user can make point changes in these arguments, by setting `mcmc.init = list(alpha = .5)` for instance, but must respect the constraints of each element:

- `mcmc.init`. All elements are of length one. The possibilities are:
  - `loc`, `scale` and `shape` (GEV parameters).
  - `range` and `sill` of the correlation functions.
  - `alpha`, `tau`, `A` and `B`, the dependence parameters and conditional variables of the HKEVP.
- mcmc.prior. The possible elements are:
  - `constant.gev`: a 2×3 matrix of normal parameters for spatially-constant μ, σ and ξ. The first row are the means, the second are the standard deviations.
  - `beta.sd`: the normal sd prior of all β parameters (a single value).
  - `range`, `alpha` and `tau`: the two Beta parameters.
  - `sill`: the two Inverse-Gamma parameters.

- mcmc.jumps. The possible elements are:
  - gev and range: a vector of length 3 (for each GEV parameter).
  - tau, alpha, A, B: single values for each.

Value

A named list with following elements:

- GEV: the Markov chains associated to the GEV parameters. The dimensions of the array correspond respectively to the sites positions, the three GEV parameters and the states of the Markov chains.
- alpha: the Markov chain associated to the dependence parameter α.
- tau: the Markov chain associated to the dependence parameter τ.
- A: the Markov chains associated to the positive stable random effect per site and per block. The dimensions correspond respectively to the indices of blocks, the knots positions and the states of the Markov chains.
- llik: the log-likelihood of the model for each step of the algorithm.
- time: time (in sec) spent for the fit.
- spatial: a named list with four elements linked to the GEV spatially-varying parameters:
  - vary: the argument gev.vary.
  - beta: the β parameters for each GEV parameter. The dimensions correspond respectively to the steps of the Markov chains, the p spatial covariates and the GEV parameters
  - sills: the Markov chains associated to the sills in the correlation functions of the latent Gaussian processes.
  - ranges: the Markov chains associated to the ranges in the correlation functions of the latent Gaussian processes.
- data: the data fitted.
- sites: the sites where the data are observed.
- knots: the set of knots.
- spatial.covariates: the spatial covariates.
- correlation: the type of correlation function for the marginal latent processes.
- nstep: the number of steps at the end of the routine after burn-in and thinning.
- log.scale: a boolean indicating if the scale parameter has been modelled via its logarithm.
- fit.type: either "hkevp" or "dep-only" character string to specify the type of fit.

If fit.margins is false, only the dependence-related elements are returned.

Author(s)

Quentin Sebille

References

Reich, B. J., & Shaby, B. A. (2012). A hierarchical max-stable spatial model for extreme precipitation. The annals of applied statistics, 6(4), 1430. <DOI:10.1214/12-AOAS591>

Stephenson, A. G. (2009) High-dimensional parametric modelling of multivariate extreme events. Aust. N. Z. J Stat, 51, 77-88. <DOI:10.1111/j.1467-842X.2008.00528.x>

Davison, A. C., Padoan, S. A., & Ribatet, M. (2012). Statistical modeling of spatial extremes. Statistical Science, 27(2), 161-186. <DOI:10.1214/11-STS376>

See Also

latent.fit

Examples

```
# Simulation of HKEVP:
set.seed(1)
sites <- as.matrix(expand.grid(1:3,1:3))
loc <- sites[,1]*10
scale <- 3
shape <- 0
alpha <- .4
tau <- 1
ysim <- hkevp.rand(10, sites, sites, loc, scale, shape, alpha, tau)

# HKEVP fit:
fit <- latent.fit(ysim, sites, niter = 1000)
```

`hkevp.predict` *Predictive distribution of the max-stable process at target positions.*

Description

Computes the predictive distribution of $Y(\cdot)$ at a set of ungauged positions $(s_1^*, ..., s_k^*)$, given data at gauged positions $(s_1, ..., s_n)$, by using the output of *latent.fit* or `hkevp.fit`.

Two types of prediction are available for the HKEVP, as described in *Shaby and Reich (2012)*. See details.

Usage

```
hkevp.predict(fit, targets, targets.covariates, predict.type = "kriging")
```

Arguments

`fit`	Output from the `hkevp.fit` procedure.
`targets`	A matrix of real values giving the spatial coordinates of the ungauged positions. Each row corresponds to an ungauged position.
`targets.covariates`	A matrix of real values giving the spatial covariates of the ungauged positions. Must match with the covariates used in `hkevp.fit` or `latent.fit`.
`predict.type`	Character string specifying the type of prediction. Must be one of "`kriging`" (default) or "`climat`". See details.

Details

The spatial prediction of $Y_t(s^*)$ for a target site s^* and a realisation t of the process is described in *Shaby and Reich (2012)*. This method involves a three-step procedure:

1. Computation of the residual dependence process $\theta(\cdot)$ at the target positions.
2. Computation of the conditional GEV parameters (μ^*, σ^*, ξ^*) at the target sites. See the definition of the HKEVP in *Reich and Shaby (2012)*.
3. Generation of $Y_t(s^*)$ from an independent GEV distribution with parameters (μ^*, σ^*, ξ^*).

As sketched in *Shaby and Reich (2012)*, two types of prediction are possible: the kriging-type and the climatological-type. These two types differ when the residual dependence process θ is computed (first step of the prediction):

- The kriging-type takes the actual value of A in the MCMC algorithm to compute the residual dependence process. The prediction will be the distribution of the maximum recorded at the specified targets.
- The climatological-type generates A by sampling from the positive stable distribution with characteristic exponent α, where α is the actual value of the MCMC step. The prediction in climatological-type will be the distribution of what could happen in the conditions of the HKEVP dependence structure.

Posterior distribution for each realisation t of the process and each target position s^* is represented with a sample where each element corresponds to a step of the MCMC procedure.

Value

A three-dimensional array where:

- Each row corresponds to a different realisation of the process (a block).
- Each column corresponds to a target position.
- Each slice corresponds to a MCMC step.

Author(s)

Quentin Sebille

References

Reich, B. J., & Shaby, B. A. (2012). A hierarchical max-stable spatial model for extreme precipitation. The annals of applied statistics, 6(4), 1430. <DOI:10.1214/12-AOAS591>

Shaby, B. A., & Reich, B. J. (2012). Bayesian spatial extreme value analysis to assess the changing risk of concurrent high temperatures across large portions of European cropland. Environmetrics, 23(8), 638-648. <DOI:10.1002/env.2178>

Examples

```
# Simulation of HKEVP:
sites <- as.matrix(expand.grid(1:3,1:3))
targets <- as.matrix(expand.grid(1.5:2.5,1.5:2.5))
all.pos <- rbind(sites, targets)
knots <- sites
loc <- all.pos[,1]*10
scale <- 3
shape <- 0
alpha <- .4
tau <- 1
ysim <- hkevp.rand(10, all.pos, knots, loc, scale, shape, alpha, tau)
yobs <- ysim[,1:9]

# HKEVP fit (omitting first site, used as target):
fit <- hkevp.fit(yobs, sites, niter = 1000)

# Extrapolation:
ypred <- hkevp.predict(fit, targets, predict.type = "kriging")

# Plot of the density and the true value for 4 first realizations:
# par(mfrow = c(2, 2))
# plot(density(ypred[1,1,]), main = "Target 1 / Year 1")
# abline(v = ysim[1,10], col = 2, lwd = 2)
# plot(density(ypred[2,1,]), main = "Target 1 / Year 2")
# abline(v = ysim[2,10], col = 2, lwd = 2)
# plot(density(ypred[1,2,]), main = "Target 2 / Year 1")
# abline(v = ysim[1,11], col = 2, lwd = 2)
# plot(density(ypred[2,2,]), main = "Target 2 / Year 2")
# abline(v = ysim[2,11], col = 2, lwd = 2)
```

---

hkevp.rand *Simulation of the HKEVP*

Description

Simulation procedure of the HKEVP with given sites and knots positions and marginal and spatial dependence parameters.

Usage

```
hkevp.rand(nrep, sites, knots, loc, scale, shape, alpha, tau)
```

Arguments

nrep	A positive integer. Number of realisations of the block mashapema process.
sites	The coordinates of the sites where the data are observed. Each row corresponds to a site position.
knots	The coordinates of the knots in the HKEVP. By default, the positions of the knots coincide with the positions of the sites.
loc	A numerical value or a vector of real values for the GEV location parameter. If a vector, its length must coincide with the number of sites. The value by default is 1.
scale	A numerical value or a vector of real values for the GEV scale parameter. If a vector, its length must coincide with the number of sites. The value by default is 1.
shape	A numerical value or a vector of real values for the GEV shape parameter. If a vector, its length must coincide with the number of sites. The value by default is 1.
alpha	The dependence parameter α of the HKEVP: a single value in (0,1].
tau	The bandwidth parameter τ of the kernel functions in the HKEVP: a positive value.

Details

Simulating one realisation of the block mashapema process $Y(\cdot)$ from the HKEVP involves three steps:

1. The *nugget process* $U(\cdot)$ is generated independently at each position, by simulating a random variable with $GEV(1, \alpha, \alpha)$ distribution.
2. The *residual dependence process* $\theta(\cdot)$ is computed by using the kernel functions centered at the set of knots, the bandwidth parameter τ and the simulations of the positive stable $PS(\alpha)$ random effect A.
3. The process $Z = U\theta$ is computed and its margins are transformed to the general GEV distribution with $\mu(s), \sigma(s)$ and $\xi(s)$ parameters.

Value

A numerical matrix of real values. Each column corresponds to a position and each row to a realisation of the process.

Author(s)

Quentin Sebille

Examples

```
# Simulation of HKEVP:
sites <- as.matrix(expand.grid(1:3,1:3))
loc <- sites[,1]*10
scale <- 3
shape <- 0
alpha <- .4
tau <- 1
ysim <- hkevp.rand(10, sites, sites, loc, scale, shape, alpha, tau)
```

`latent.fit` *Fitting procedure of the latent variable model*

Description

Metropolis-within-Gibbs algorithm that returns posterior distribution (as Markov chains) for the marginal GEV parameters of an observed spatial process. This function is close to `hkevp.fit` but with less parameters since conditional independence is assumed and only the margins are estimated. In `SpatialExtremes` library, a similar function can be found under the name `latent`.

Usage

```
latent.fit(
  y,
  sites,
  niter,
  nburn,
  nthin,
  quiet,
  trace,
  gev.vary,
  spatial.covariates,
  log.scale,
  correlation,
  mcmc.init,
  mcmc.prior,
  mcmc.jumps
)
```

Arguments

`y`	A matrix of observed block maxima. Each column corresponds to a site position.

`sites`	The coordinates of the sites where the data are observed. Each row corresponds to a site position.
`niter`	The number of MCMC iterations.
`nburn`	The number of first MCMC iterations that are discarded. Zero by default.
`nthin`	The size of the MCMC thinning. One by default (i.e. no thinning).
`quiet`	A logical indicating if the progression of the routine should be displayed. TRUE by default.
`trace`	If `quiet` is FALSE, the log-likelihood of the model is displayed each block of `trace` MCMC steps to observe fitting progression.
`gev.vary`	A logical vector of size three indicating if the GEV parameters (respectively the location, the scale and the shape) are spatially-varying. If not (by default for the shape), the parameter is the same at each position.
`spatial.covariates`	A numerical matrix of spatial covariates. Each row corresponds to a site position. See details.
`log.scale`	A logical value indicating if the GEV scale parameter σ is modelled by its log. FALSE by default. See details.
`correlation`	A character string indicating the form of the correlation function associated to the latent Gaussian processes that describes the marginal parameters. Must be one of `"expo"`, `"gauss"`, `"mat32"` (By default) and `"mat52"`, respectively corresponding to the exponential, Gaussian, Matern-3/2 and Matern-5/2 correlation functions.
`mcmc.init`	A named list indicating the initial states of the chains. See details.
`mcmc.prior`	A named list indicating the hyperparameters of the prior distributions. See details.
`mcmc.jumps`	A named list indicating the amplitude of the jumps to propose the MCMC candidates. See details.

Details

Details of the MCMC procedure are presented in *Davison et al. (2012)*. This function follows the indications and the choices of the authors, with the exception of several small changes:

- The scale parameter σ can be modelled like the two other marginal parameters as in *Davison et al. (2012)* or by its logarithm as in *Reich and Shaby (2012)*. For this, use the argument `log.scale`, set to FALSE by default.
- The Inverse-Gamma prior distributions defined for the bandwith parameter τ and for the ranges λ of the latent processes are replaced by a Beta distribution over the interval $[0, 2D_{max}]$, where D_{max} stands for the maximum distance between two sites.

If the the parameters are assumed spatially-varying, the user can provide spatial covariates to fit the mean of the latent Gaussian processes. Recall for instance for the GEV location parameter that:

$$\mu(s) = \beta_{0,\mu} + \beta_{1,\mu} c_1(s) + ... + \beta_{p,\mu} c_p(s) \ .$$

The given matrix `spatial.covariates` that represents the $c_i(s)$ elements should have the first column filled with ones to account for the intercept β_0.

The arguments mcmc.init, mcmc.prior and mcmc.jumps are named list that have default values. The user can make point changes in these arguments, by setting mcmc.init = list(loc = .5) for instance, but must respect the constraints of each element:

- mcmc.init: all elements are of length one. The possible elements are:
  - loc, scale and shape (GEV parameters).
  - range and sill of the correlation functions.
- mcmc.prior: the possible elements are:
  - constant.gev: a 2×3 matrix of normal parameters for spatially-constant μ, σ and ξ. The first row are the means, the second are the standard deviations.
  - beta.sd: the normal sd prior of all β parameters (a single value).
  - range: the two Beta parameters.
  - sill: the two Inverse-Gamma parameters.
- mcmc.jumps: the possible elements are gev and range, vectors of length 3 (for each GEV parameter).

Value

A named list with following elements (less elements if fit.margins is FALSE):

- GEV: the Markov chains associated to the GEV parameters. The dimensions of the array correspond respectively to the sites positions, the three GEV parameters and the states of the Markov chains.
- llik: the log-likelihood of the model for each step of the algorithm.
- time: time (in sec) spent for the fit.
- spatial: a named list with four elements linked to the GEV spatially-varying parameters:
  - vary: the argument gev.vary.
  - beta: the β parameters for each GEV parameter. The dimensions correspond respectively to the steps of the Markov chains, the p spatial covariates and the GEV parameters
  - sills: the Markov chains associated to the sills in the correlation functions of the latent Gaussian processes.
  - ranges: the Markov chains associated to the ranges in the correlation functions of the latent Gaussian processes.
- data: the data fitted.
- sites: the sites where the data are observed.
- spatial.covariates: the spatial covariates.
- correlation: the type of correlation function for the marginal latent processes.
- nstep: the number of steps at the end of the routine after burn-in and thinning.
- log.scale: a boolean indicating if the scale parameter has been modelled via its logarithm.
- fit.type: "latent" character string to specify the type of fit.

Author(s)

Quentin Sebille

References

Reich, B. J., & Shaby, B. A. (2012). A hierarchical max-stable spatial model for extreme precipitation. The annals of applied statistics, 6(4), 1430. <DOI:10.1214/12-AOAS591>

Davison, A. C., Padoan, S. A., & Ribatet, M. (2012). Statistical modeling of spatial extremes. Statistical Science, 27(2), 161-186. <DOI:10.1214/11-STS376>

See Also

hkevp.fit

Examples

```
# Simulation of HKEVP:
sites <- as.matrix(expand.grid(1:4,1:4))
loc <- sites[,1]*10
scale <- 3
shape <- .2
alpha <- 1
tau <- 2
ysim <- hkevp.rand(15, sites, sites, loc, scale, shape, alpha, tau)

# Latent Variable Model fit:
set.seed(1)
fit <- latent.fit(ysim, sites, niter = 10000, nburn = 5000, nthin = 5)
mcmc.plot(fit, TRUE)
par(mfrow = c(2,2))
apply(fit$GEV[,1,], 1, acf)
```

---

mcmc.fun *Point estimates of HKEVP fit*

Description

Application of a function to the main Markov chains resulting from the procedure `hkevp.fit`. May be used to obtain point estimates of the posterior distribution (e.g., the mean or the median). See details.

Usage

```
mcmc.fun(fit, FUN, ...)
```

Arguments

fit	A named list. Output from the hkevp.fit procedure.
FUN	The function applied to the Markov chains in fit. The median by default. The output from FUN must be a single value.
...	Optional arguments of the function to be applied on the Markov chains (e.g. na.rm = FALSE).

Details

A function is applied to the main Markov chains resulting from the MCMC procedures hkevp.fit or latent.fit. These chains correspond to the three GEV parameters, the dependence parameter α and the bandwidth τ.

The value returned by FUN must be a single value.

Value

If fitted model is the HKEVP, a named list with three elements:

- GEV: A numerical matrix. Result of the function FUN for each GEV parameter (columns) and each site position (rows).
- alpha: A numerical value. Result of the function FUN on the Markov chain associated to the dependence parameter α.
- tau: A numerical value. Result of the function FUN on the Markov chain associated to the bandwidth parameter τ.

If fitted model is the latent variable model, the functions returns the GEV matrix only.

Author(s)

Quentin Sebille

Examples

```
# Simulation of HKEVP:

sites <- as.matrix(expand.grid(1:3,1:3))
knots <- sites
loc <- sites[,1]*10
scale <- 3
shape <- .2
alpha <- .4
tau <- 1
ysim <- hkevp.rand(10, sites, knots, loc, scale, shape, alpha, tau)

# HKEVP fit:
fit <- hkevp.fit(ysim, sites, niter = 1000)

# Posterior median and standard deviation:
# mcmc.fun(fit, median)
```

```
# mcmc.fun(fit, sd)
```

`mcmc.plot` *Markov chains plotting*

Description

Plots of the resulting Markov chains obtained by the MCMC procedures `hkevp.fit` or `latent.fit`. May be used to assess graphically convergence of the chains.

Usage

```
mcmc.plot(fit, plot.spatial, mfrow)
```

Arguments

`fit`	Output from the `hkevp.fit` procedure.
`plot.spatial`	Logical indicating if the Markov chains of the sills and ranges hyperparameters should be plotted. FALSE by default.
`mfrow`	Optional vector of two numerical values indicating the parameter of the window plotting called by the `plot(...)` function.

Author(s)

Quentin Sebille

Examples

```
# Simulation of HKEVP:
sites <- as.matrix(expand.grid(1:3,1:3))
knots <- sites
loc <- sites[,1]*10
scale <- 3
shape <- .2
alpha <- .4
tau <- 1
ysim <- hkevp.rand(10, sites, knots, loc, scale, shape, alpha, tau)

# HKEVP fit:
fit <- hkevp.fit(ysim, sites, niter = 1000)

# Markov chains plot:
mcmc.plot(fit, TRUE)
```

`return.level` *The associated return level*

Description

Computation of the associated return level with given period and GEV parameters.

Usage

```
return.level(period, loc, scale, shape)
```

Arguments

`period`	An integer indicating the wished return period T.
`loc`	A numerical value or vector for the GEV location parameter. Must be of length one or same length as `scale` and/or `shape`.
`scale`	A numerical value or vector for the GEV scale parameter. Must be of length one or same length as `loc` and/or `shape`.
`shape`	A numerical value or vector for the GEV shape parameter. Must be of length one or same length as `loc` and/or `scale`.

Details

The T-year return level is a common value of risk in Extreme Value Theory. It represents the value that is expected to be exceeded once over T years by the annual maxima. Given the parameters μ, σ and ξ of the GEV distribution associated to the yearly maxima, we can compute the associated T-return level y_T by:

$$y_T := \mu + \frac{\sigma}{\xi}\left[\log\left(\frac{T}{T-1}\right)^{-\xi} - 1\right] .$$

Value

A numerical value or a numerical vector, depending on the input arguments `loc`, `scale`, `shape`

Author(s)

Quentin Sebille

Examples

```
return.level(period = 100, loc = 1, scale = 1, shape = 1)
return.level(period = 200, loc = 1:10, scale = 1, shape = 0)
```

Index

extrapol.gev, 2, *4, 5*
extrapol.return.level, *2*, 3, *5*

hkevp, 5
hkevp-package *(*hkevp*)*, 5
hkevp.expmeasure, *5*, 8
hkevp.fit, *2, 3*, *5*, 9, *14*
hkevp.predict, *5*, 13
hkevp.rand, *5*, 15

latent.fit, *2, 3*, *5*, *14*, 17

mcmc.fun, *5*, 20
mcmc.plot, *5*, *9*, 22

return.level, 23