SpatialExtremes: An R package for modelling spatial extremes

Mathieu Ribatet

École Polytechnique Fédérale de Lausanne,
Chaire de Statistique

22nd June 2009

Outline

Introduction

Unconditional Simulation of Max-stable Processes

Spatial Dependence of Max-Stable Random Fields

Fitting Max-stable Processes to Data

Model Selection

Assuming non unit Fréchet margins

Model Checking

Predictions

Conclusion

Outline

Geostatistics of Extremes

- Our goal is to model spatial extremes
- Conventional geostatistics are not relevant as
  - Extremes are far from being Normal
  - Variogram based approaches may not even exist - $\mathbb{E}[Y(x)] = +\infty$, $\text{Var}[Y(x)] = +\infty$
- We want to extend the EVT to the spatial case

Definition (Max-stable processes)

A max-stable process $Z(\cdot)$ is the limit process of maxima of i.i.d. random fields $Y_i(x)$, $x \in \mathbb{R}^d$. Namely, for suitable $a_n(x) > 0$ and $b_n(x) \in \mathbb{R}$,

$$Z(x) = \lim_{n \to +\infty} \frac{\max_{i=1}^n Y_i(x) - b_n(x)}{a_n(x)}, \qquad x \in \mathbb{R}^d$$

- We hope that, as the GEV/GPD, max-stable processes will be good candidates for modelling spatial extremes

Towards Parametric Models

- Schlather [2002] introduced a very useful representation of max-stable processes

Theorem

Let $\{\xi_i\}_{i \geq 1}$ *be the points of a homogeneous Poisson process on* $\mathbb{R}_+$ *with intensity* $d\Lambda(\xi) = \xi^{-2} d\xi$, *and* $\{Y_i(\cdot)\}_{i \geq 1}$ *be i.i.d. replicates of a stationary process on* $\mathbb{R}^d$ *such that* $\mathbb{E}[\max\{0, Y(x)\}] = 1$.
Then

$$Z(x) = \max_i \xi_i \max\{0, Y_i(x)\}$$

is a stationary max-stable process with unit Fréchet margins

- Different choices for $Y(\cdot)$ lead to different max-stable processes

The Smith model

- Smith [1990] proposed to take $Y_i(x) = \varphi(x - X_i)$ where φ is a zero mean multivariate normal density with covariance matrix Σ and $\{X_i\}_{i \geq 1}$ is a homogeneous Poisson process, both on $\mathbb{R}^d$.

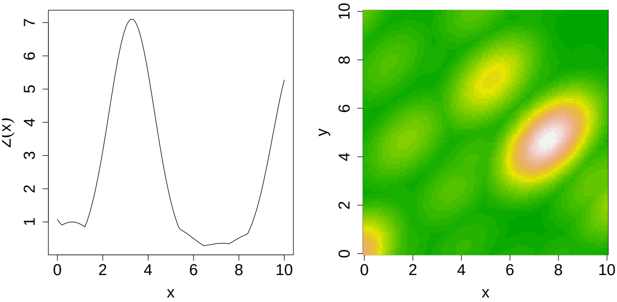

Figure: Two realisations of the Smith model in $\mathbb{R}$ and $\mathbb{R}^2$. $\mathbb{R} : \sigma^2 = 1$. $\mathbb{R}^2 : \Sigma = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 1 \end{bmatrix}$

The Schlather Model

- Schlather [2002] proposed to take $Y_i(\cdot)$ as an appropriately scaled stationary Gaussian processes with correlation function ρ.

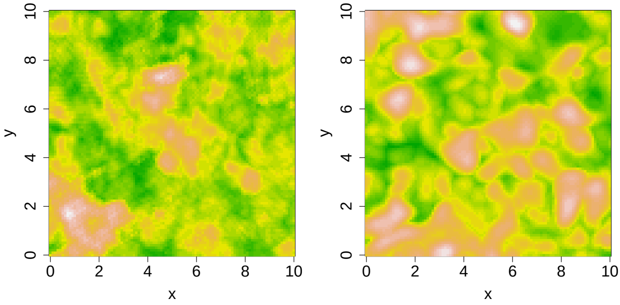

Figure: Two realisations of the Schlather model in $\mathbb{R}^2$. Left: Whittle-Matérn. Right: Cauchy.

The Schlather Model (2)

Table: Correlation function families implemented in the package.

Whittle–Matérn	$\rho(h) = c_1 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{h}{c_2}\right)^{\nu} K_\nu \left(\frac{h}{c_2}\right)$,	$\nu > 0$
Cauchy	$\rho(h) = c_1 \left[1 + \left(\frac{h}{c_2}\right)^2\right]^{-\nu}$,	$\nu > 0$
Powered Exponential	$\rho(h) = c_1 \exp\left[-\left(\frac{h}{c_2}\right)^{\nu}\right]$,	$0 < \nu \leq 2$
Bessel	$\rho(h) = c_1 \left(\frac{2c_2}{h}\right)^{\nu} \Gamma(\nu + 1) J_\nu \left(\frac{h}{c_2}\right)$,	$\nu \geq \frac{d-2}{2}$

where Γ is the gamma function, K_ν is the modified Bessel function of the third kind with degree ν and J_ν is the Bessel function of order ν.

- c_1 is the sill parameter, $0 < c_1 \leq 1$
- c_2 is the range parameter, $c_2 > 0$
- ν is the smooth parameter

Plotting covariance functions

Example (File covariance.R)

```
covariance(sill = 1, range = 1, smooth = 1, cov.mod = "whitmat",
           xlim = c(0, 10), ylim = c(-.2, 1))
covariance(sill = 1, range = 1, smooth = 1, cov.mod = "cauchy",
           add = TRUE, col = 2, xlim = c(0, 10), ylim = c(-.2, 1))
covariance(sill = 1, range = 1, smooth = 1, cov.mod = "powexp",
           add = TRUE, col = 3, xlim = c(0, 10), ylim = c(-.2, 1))
covariance(sill = 1, range = 1, smooth = 1, cov.mod = "bessel",
           add = TRUE, col = 4, xlim = c(0, 10), ylim = c(-.2, 1))
legend("topright", c("Whittle-Matern", "Cauchy", "Pow. Exp.", "Bessel"),
       lty = 1, col = 1:4, inset = 0.05)
```

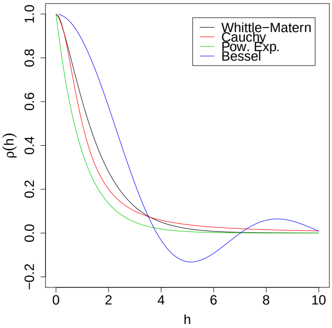

Outline

Simulation Algorithms

- Remind Schlather's representation of max-stable processes
- It involves the maximum over an infinite number of random processes replications - ouch!!!

Theorem (Schlather (2002))

Let Π be a P.P., $d\Lambda(y,\xi) = \xi^{-2} dy d\xi$. Assume that Y is uniformly bounded by $C \in (0, +\infty)$ and has support in the ball $b(o, r)$, $r < +\infty$. Let B be a compact set, Y_i be i.i.d. replications of Y, U_i be i.i.d. uniformly distributed on $B_r = \cup_{x \in B} b(x, r)$, ξ_i be i.i.d. standard exponential r.v. and Π, Y_i, ξ_i, U_i be mutually independent. Then, on B,

$$Z_*(x) = |B_r| \sup \left\{ \frac{Y_i(x - U_i)}{\sum_{k=1}^{i} \xi_k} : i = 1, \ldots, m \right\}$$

where m is such that $\frac{C}{\sum_{k=1}^{m} \xi_k} \leq \max_{1 \leq i \leq m} \frac{Y_i(x - U_i)}{\sum_{k=1}^{i} \xi_k}$, equals $Z(\cdot)$ almost surely.

- The beauty of the previous theorem is that we only have to take the maximum over *m* replications, *m* being finite!!!
- The function *rmaxstab* uses this algorithm to generate max-stable random fields

Example (File: simMaxStab.R)

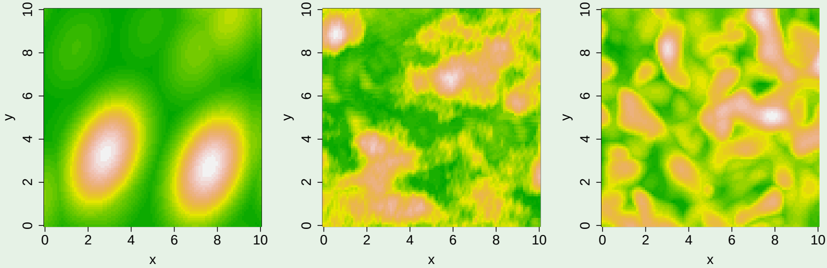

Figure: Ouput of simMaxStab.R

Comments on *rmaxstab* and grid locations

- Simulation from the Smith model is pretty fast
- Simulation from the Schlather model isn't (turning bands)
- I'll try to improve it → circulant embedding method
- For large grids, you should prefer using the *RandomFields* package

Example (*RandomFields* package)

```
x <- seq(0, 10, length = 200)
y <- x
data <- MaxStableRF(x, y, grid=TRUE, model="wh",
                    param=c(0,1,0,1, 1),
                    maxstable="extr", n = 1)
data <- t(data)
image(x, y, sqrt(data), col = terrain.colors(30))
```

Outline

- For stationary gaussian processes, the variogram is useful to assess how evolves dependence in space

$$\gamma(x_1 - x_2) = \frac{1}{2}\text{Var}\left[Y(x_1) - Y(x_2)\right] = \sigma^2\left\{1 - \rho(x_1 - x_2)\right\}$$

- For extreme observations, the variance (and even the mean) might be infinite
- There's a pressing need to know how evolves the spatial dependence of extremes
- The extremal coefficient function is what we need

Extremal coefficient function

- Let $Z(\cdot)$ be a stationary max-stable random field with unit Fréchet margins.
- The extremal coefficient function $\theta(\cdot)$ is defined by

$$\Pr\left[Z(x_1) \leq z, Z(x_2) \leq z\right] = \exp\left\{-\frac{\theta(x_1 - x_2)}{z}\right\}$$

- If the random field is isotropic, this simplifies to

$$\theta(h) = \theta(||x_1 - x_2||),$$

where h is the euclidean distance between x_1 and x_2.
- $\theta(h) = 1$ is equivalent to complete dependence as

$$\Pr[Z(x_1) \leq z, Z(x_2) \leq z] = \exp(-1/z) = \Pr[Z(x_1) \leq z]$$

- $\theta(h) = 2$ is equivalent to independence as

$$\begin{aligned} \Pr[Z(x_1) \leq z, Z(x_2) \leq z] &= \exp(-2/z) \\ &= \Pr[Z(x_1) \leq z] \Pr[Z(x_2) \leq z] \end{aligned}$$

The *fitextcoeff* function

- Smith [1990] and Schlather and Tawn [2003] proposed two estimators for $\theta(h_{ij})$, where h_{ij} is the euclidean distance between locations x_i and x_j.
- These estimators are implemented in the *fitextcoeff* function

Example (file extCoeff.R)

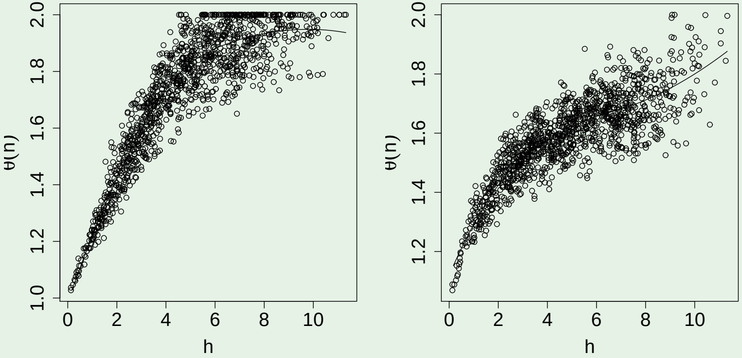

Figure: Pairwise extremal coefficient estimates and lowess curves. Left: Smith. Right: Schlather (Whittle-Matérn).

Variogram based approaches

- Cooley, Guillou, Naveau and Poncet proposed variogram based approaches especially designed for extremes
- The madogram [Matheron, 1987] is

$$\nu(x_1 - x_2) = \mathbb{E}[|Z(x_1) - Z(x_2)|]$$

- As stated earlier, the mean might be infinite
- Cooley et al. proposed modified madograms. For instance, the F-madogram

$$\nu_F(x_1 - x_2) = \frac{1}{2}\mathbb{E}[|F\{Z(x_1)\} - F\{Z(x_2)\}|]$$

  where $Z(\cdot)$ is a stationary max-stable random field with unit Fréchet margins and $F(z) = \exp(-1/z)$.
- Hence $F\{Z(x_1)\} \sim U(0, 1)$ and the F-madogram is well defined

Connections with the Extremal Coefficient

- It is not hard to show that

$$\theta(x_1 - x_2) = \frac{1 + 2\nu_F(x_1 - x_2)}{1 - 2\nu_F(x_1 - x_2)}$$

Example (file madogram.R)

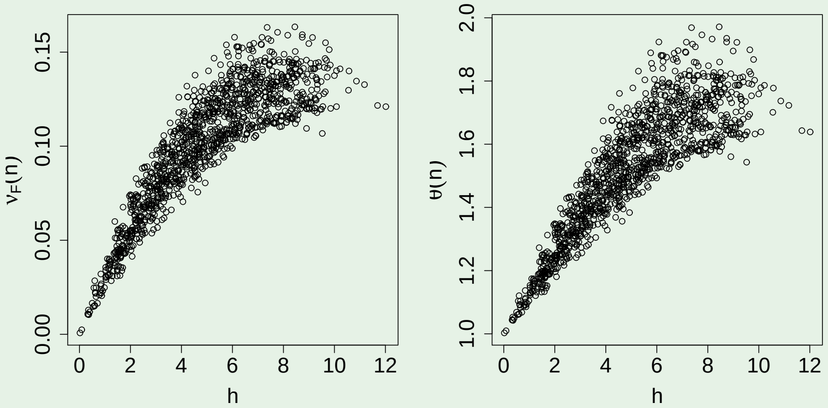

Connections with the Extremal Coefficient

- It is not hard to show that

$$\theta(x_1 - x_2) = \frac{1 + 2\nu_F(x_1 - x_2)}{1 - 2\nu_F(x_1 - x_2)}$$

Example (file madogram.R)

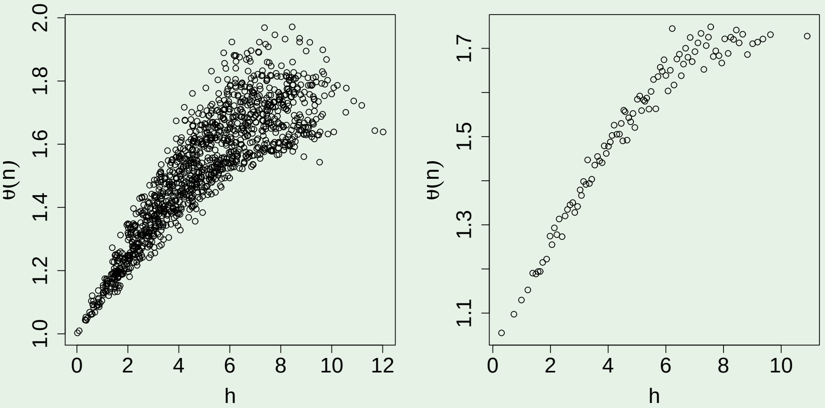

λ-madogram

- $\theta(\cdot)$ doesn't fully characterise the spatial dependence - only $\Pr[Z(x_1) \leq z, Z(x_2) \leq z]$
- Naveau et al [2009] introduce the λ-madogram as follows

$$\nu_\lambda(x_1 - x_2) = \frac{1}{2}\mathbb{E}[|F^\lambda\{Z(x_1)\} - F^{1-\lambda}\{Z(x_2)\}|], \qquad 0 \leq \lambda \leq 1$$

- The idea is to consider $\Pr[Z(x_1) \leq z_1, Z(x_2) \leq z_2]$ where $z_1 = \lambda z$ and $z_2 = (1 - \lambda)z$.
- It is not hard to show that

$$\nu_\lambda(x_1 - x_2) = \frac{V_{x_1,x_2}(\lambda, 1-\lambda)}{1 + V_{x_1,x_2}(\lambda, 1-\lambda)} - \frac{3}{(1+\lambda)(2-\lambda)}$$

  where $\Pr[Z(x_1) \leq z_1, Z(x_2) \leq z_2] = \exp\{-V_{x_1,x_2}(z_1, z_2)\}$.
- Remark: The F-madogram is somehow similar to the λ-madogram when $\lambda = 0.5$

λ-madogram (2)

- The λ-madogram is a really nice tool
- It's not easy how to interpret it though

Example (File lmadogram.R)

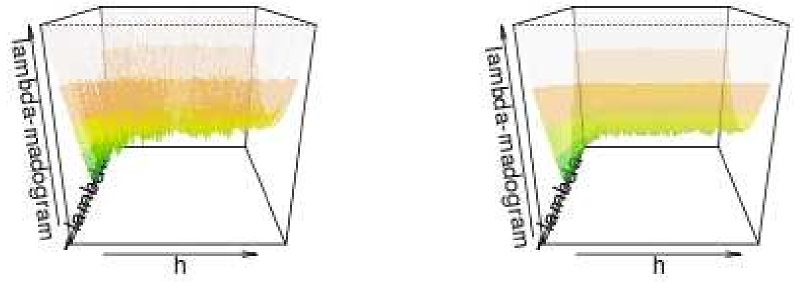

Figure: λ-madogram (left) and its binned version (right).

Outline

Why is it so difficult?

- So far, we haven't talk too much about parametric models
- It is almost impossible to get analytical CDF for these models
- Indeed as $Z(x) = \max_i \xi_i Y_i^+(x)$, $Y_i^+(x) = \max\{0, Y_i(x)\}$
- Hence to get the k-variate CDF we need to compute

$$\begin{aligned}
F(z_1, \dots, z_k) &= \Pr[\max_i \xi_i Y_i^+(x_j) \leq z_j, j = 1, \dots, k] \\
&= \Pr\left[\xi_i \leq \frac{z_j}{Y_i^+(x_j)}, \forall i, j = 1, \dots, k\right] \\
&= \exp\left\{-\int_{\mathbb{R}^d} \int_{\mathbb{R}} \mathbb{I}\left(\xi \leq \min_j \frac{z_j}{Y^+(x_j)}\right) \xi^{-2} \mathrm{d}\xi \mathrm{dP}\{Y(\cdot)\}\right\} \\
&= \exp\left\{-\int_{\mathbb{R}^d} \max_j \frac{Y^+(x_j)}{z_j} \mathrm{dP}\{Y(\cdot)\}\right\}
\end{aligned}$$

- The number of possible cases becomes quickly intractable when k gets large

Least Squares

- This led Smith [1990] to use least squares
- Find ψ minimizing

$$C(\psi) = \sum_{i<j} \left(\frac{\theta(x_i - x_j) - \tilde{\theta}(x_i - x_j)}{s\{\tilde{\theta}(x_i - x_j)\}} \right)^2$$

Example (File leastSquares.R)

```
        Estimator: Least Square
            Model: Schlather
  Objective Value: 1025.681
Covariance Family: Powered Exponential

Estimates
  Marginal Parameters:
  Assuming unit Frechet.
  Dependence Parameters:
  sill   range  smooth
 1.000   1.612   1.129

Optimization Information
  Convergence: successful
  Function Evaluations: 172
```

Maximum pairwise likelihood estimator

- As already stated the k-variate densities aren't analytically known
- One can maximise pairwise likelihood instead of full likelihood

$$\ell_p(\mathbf{z}; \psi) = \sum_{i<j} \sum_{k=1}^{n} \log f(z_k^{(i)}, z_k^{(j)}; \psi)$$

- The MPLE $\hat{\psi}_p$ shares similar properties with the MLE

$$\hat{\psi}_p \dot{\sim} N\left(\psi, H(\psi)^{-1} J(\psi) H(\psi)^{-1}\right)$$

where $H(\psi) = \mathbb{E}[\nabla^2 \ell_p(\psi; \mathbf{Z})]$ and $J(\psi) = \mathrm{Var}[\nabla \ell_p(\psi; \mathbf{Z})]$ and the expectations are w.r.t. the "full" density.

Maximum pairwise likelihood estimator (2)

Example (File pairwise.R)

```
       Estimator: MPLE
           Model: Schlather
  Pair. Deviance: 964496.2
             TIC: 964808.7
Covariance Family: Powered Exponential

Estimates
  Marginal Parameters:
  Assuming unit Frechet.
  Dependence Parameters:
  sill    range   smooth
0.9777  1.8524  0.5953

Standard Error Type: score

Standard Errors
  sill    range   smooth
0.2393  0.6415  0.2629

Asymptotic Variance Covariance
        sill      range     smooth
sill     0.05726  -0.13735  -0.06208
range   -0.13735   0.41152   0.14204
smooth  -0.06208   0.14204   0.06910

Optimization Information
  Convergence: successful
  Function Evaluations: 60
```

```
       Estimator: MPLE
           Model: Smith
  Pair. Deviance: 996556.5
             TIC: 996682.5
Covariance Family: Gaussian

Estimates
  Marginal Parameters:
  Not estimated.
  Dependence Parameters:
 cov11   cov12   cov22
0.9063  0.3624  3.2020

Standard Error Type: score

Standard Errors
  cov11    cov12    cov22
0.06922  0.09537  0.27087

Asymptotic Variance Covariance
       cov11     cov12     cov22
cov11  0.004791  0.002209  0.003831
cov12  0.002209  0.009095  0.006482
cov22  0.003831  0.006482  0.073370

Optimization Information
  Convergence: successful
  Function Evaluations: 56
```

Outline

- The use of the MPLE enables us to “easily” perform model selection such as AIC, likelihood ratio tests
- One has to pay attention that we now deal with misspecified models i.e. we use pairwise likelihood instead of full likelihood

Takeuchi Information Criterion

- Recall that $\text{AIC} = -2\{\ell(\hat{\psi}) - p\}$, where $p = \dim(\psi)$
- In presence of misspecification, the use of AIC is not justified as the second Bartlett idendity is not satisfied i.e.

$$\mathbb{E}[\nabla^2 \ell_p(\psi; \mathbf{Z})] + \text{Var}[\nabla \ell_p(\psi; \mathbf{Z})] \neq \mathbf{0}$$

- One should prefer the TIC

$$\text{TIC} = -2\ell_p(\hat{\psi}_p) - 2\text{tr}\left\{J(\hat{\psi}_p)H(\hat{\psi}_p)^{-1}\right\}$$

- AIC is just a special case of TIC where the second Bartlett idendity holds i.e.

$$J(\psi)H(\psi)^{-1} = -\text{Id}_p \Longrightarrow \text{TIC} = \text{AIC}, \qquad \text{Id}_p = p \times p \text{ id. matrix}$$

Example (file TIC.R)

```
TIC(M0, M1)
     M0      M1
1022201 1022257
```

Likelihood ratio test under misspecification

- Recall that the likelihood ratio test is based on

$$2\left\{\ell(\hat{\psi}) - \ell(\kappa_0, \hat{\phi}_{\kappa_0})\right\} \longrightarrow \chi^2_p, \qquad n \to +\infty$$

  where $\psi = (\kappa, \phi)$, $\hat{\phi}_{\kappa_0}$ is the MLE under the restriction $\kappa = \kappa_0$ and $p = \dim(\kappa)$
- In presence of misspecification, this result slightly differs

$$2\left\{\ell_p(\hat{\psi} - p) - \ell_p(\kappa_0, \hat{\phi}_{\kappa_0})\right\} \longrightarrow \sum_{i=1}^{p} \lambda_i X_i, \qquad n \to +\infty$$

  where the λ_is are the eigenvalues of $(H^{-1}JH^{-1})_\kappa\{-(H^{-1})_\kappa\}^{-1}$ and $X_i \overset{iid}{\sim} \chi^2_1$
- There exist two different ways to perform model selection
  - RJ Approximate the distribution of $\sum_{i=1}^{p} \lambda_i X_i$ [Rotnitzki and Jewell, 1990]
  - CB Adjust $\ell_p(\cdot)$ to have the appropriate curvature [Chandler and Bate, 2007]

Likelihood ratio test under misspecification (2)

- Rotnitzki and Jewell suggest

$$2p\left\{\ell_p(\hat{\psi}_p) - \ell_p(\kappa_0, \hat{\phi}_{\kappa_0})\right\} \sim \chi^2_p$$

  but this is only an approximation that matches only the first moment
- Chandler and Bate suggest replacing $\ell_p(\cdot)$ by

$$\ell_A(\psi) = \ell_p(\psi_*), \qquad \psi_* = \hat{\psi}_p + M^{-1}M_A(\psi - \hat{\psi}_p)$$

  where $MM^T = H$ and $M_A M_A^T = H^{-1}JH^{-1}$

so that $\ell_A(\hat{\psi}_p)$ has the appropriate curvature $H^{-1}JH^{-1}$

Example (file likratiotest.R)

```
anova(M0, M1)
Eigenvalue(s):
7.47

Analysis of Variance Table
   MDf Deviance Df   Chisq Pr(> sum lambda Chisq)
M0   2  1036388
M1   3  1036369  1 19.082                  0.1100
```

```
anova(M0, M1, method = "CB")

Analysis of Variance Table
   MDf Deviance Df   Chisq Pr(> sum lambda Chisq)
M0   2   111027
M1   3   111025  1 1.5569                  0.2121
```

Outline

- So far only max-stable processes with unit Fréchet margins were considered
- For real world analysis, the unit Fréchet assumption is too restrictive
- We need a more flexible inferential procedure
- The MPLE, as we will see later, is perfectly adapted for this
- A first idea consists in transforming data to the unit Fréchet scale i.e.

$$Z(x_i) = -\frac{1}{\log F_{x_i}\{Y(x_i)\}}$$

  where $F_x(\cdot)$ is either the empirical CDF or any appropriate CDF
- Then fit max-stable processes as before
- But this is not satisfactory as
  - predictions at ungauged locations won't be possible
  - standard errors will be underestimated as we suppose that our data were originally unit Fréchet

How do we allow for unknown GEV margins?

- To transform GEV data to the unit Fréchet scale we need the mapping

$$t : Y(x) \mapsto \left(1 + \xi(x)\frac{Y(x) - \mu(x)}{\sigma(x)}\right)^{1/\xi(x)}$$

- Hence

$$\Pr[Y(x_1) \leq y_1, Y(x_2) \leq y_2] = \Pr[Z(x_1) \leq t(y_1), Z(x_2) \leq t(y_2)]$$

- And the log-pairwise likelihood becomes

$$\ell_p(y;\psi) = \sum_{i<j}\sum_{k=1}^{n}\left[\log f\{t(y_k^{(i)}), t(y_k^{(j)});\psi\} + \log|J(y_k^{(i)})J(y_k^{(j)})|\right]$$

where $|J(y_k^{(i)})|$ is the jacobian of the mapping t i.e.

$$|J(y_k^{(i)})| = \frac{1}{\sigma(x_i)}\left(1 + \xi(x_i)\frac{y_k^{(i)} - \mu(x_i)}{\sigma(x_i)}\right)^{1/\xi(x_i)-1}$$

Dimensional curse

- Fitting one GEV to each location will lead to $3K + p$ parameters to be estimated
- We thus need response surfaces on the GEV parameters to get more parsimonious models

$$\mu = X_\mu \beta_\mu, \qquad \sigma = X_\sigma \beta_\sigma, \qquad \xi = X_\xi \beta_\xi$$

- The current trend surfaces implemented are linear models and p-splines with radial basis functions

Example (File fitmaxstab.R)

```
        Estimator: MPLE
            Model: Schlather
   Pair. Deviance: 2589198
              TIC: NA
Covariance Family: Powered Exponential

Estimates
  Marginal Parameters:
    Location Parameters:
locCoeff1  locCoeff2
  -10.722      2.229
       Scale Parameters:
scaleCoeff1  scaleCoeff2  scaleCoeff3
      4.139        2.060        1.042
```

```
          Shape Parameters:
shapeCoeff1
      0.275
  Dependence Parameters:
  sill   range  smooth
0.9990  1.4927  0.6946

Optimization Information
  Convergence: successful
  Function Evaluations: 327
```

Outline

- Once we have fitted our max-stable process
- One often want to check if our model is appropriate or not
- This amounts to check if
  1. The margins are appropriately modelled
  2. The spatial dependence structure is satisfactory
- This can be done through two different functions
  1. *qqgev*
  2. *madogram*, *fmadogram*

Checking the margins

- The idea is to check if the GEV parameters predicted by the trend surfaces are relevant
- For this we will compare them to the GEV MLE at each location

Example (File modelCheck.R)

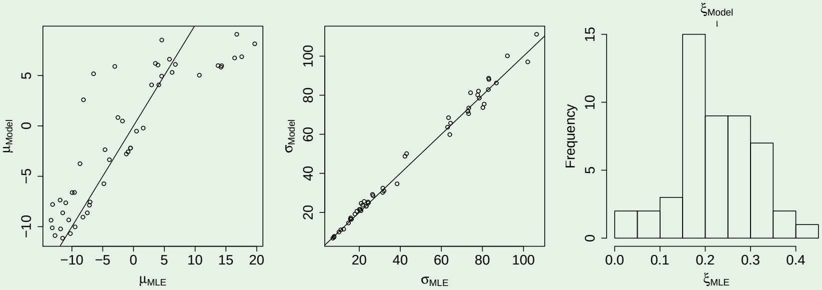

Figure: Checking the margins using the qqgev function.

Checking the spatial dependence parameters

- The idea is to check if the model is able to reproduce the spatial dependence structure
- This is done by comparing semi-parametric estimates for the extremal coefficients and the ones predicted from the model

Example (File modelCheck.R)

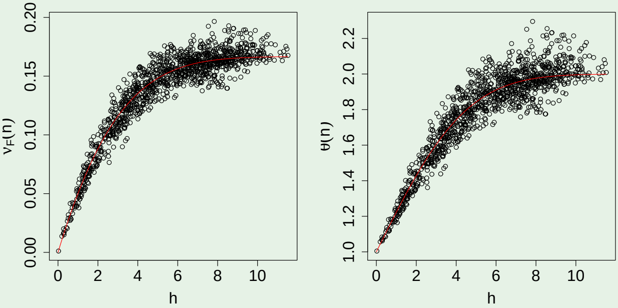

Figure: Checking the spatial dependence structure using the fmadogram function.

Outline

- Knowing the spatial dependence structure is not enough
- Sonner or later, one will be interested in prediction
- We will see how it is possible to get prediction at ungauged locations
- Currently, there are two possible types of predictions
  1. Pointwise predictions
  2. Conditional predictions
- This is achieved with the following functions
  1. *predict*, *map*
  2. *condmap*

Pointwise quantile

- As we fit trend surfaces for each GEV parameters, it is possible to know the distribution of extreme for each coordinates within the region
- The return level with return period T-year is

$$z_T(x) = \mu(x) + \sigma(x) \frac{\{-\log(1 - 1/T)\}^{-\xi(x)} - 1}{\xi(x)}$$

- This is the level which is expected to be exceeded once every T-year

Example (File predict.R)

```
predict(fitted, new.coord, ret.per = 50)
       lon      lat        loc    scale     shape      Q50
1 4.399032 8.202387 -0.9648912 43.11199 0.2473414 282.2889
2 3.392812 2.442806 -3.0530930 22.96123 0.2473414 147.8065
3 4.278938 3.011325 -1.2141221 31.52409 0.2473414 205.9050
```

Maps of pointwise estimates

- The *map* function produces maps for the GEV parameters as well as return levels.

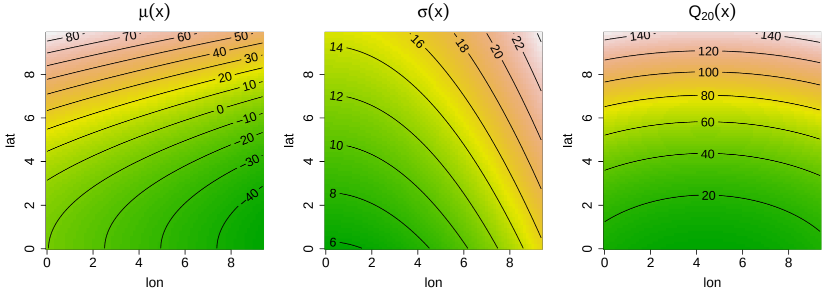

Figure: Maps of the pointwise estimates for the location, scale and 20-year return level.

Maps of conditional quantiles

- A conditional return level is defined as follows

$$\Pr[Z(x_2) \geq z_2 | Z(x_1) \geq z_1] = \frac{1}{T_2}, \quad \Pr[Z(x_1) \geq z_1] = 1 - \frac{1}{T_1}$$

Example (File predict.R)

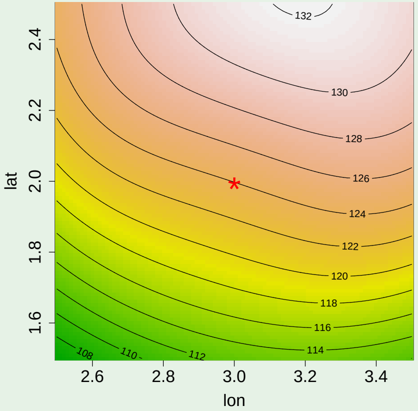

Outline

Conclusion

- Max-stable processes are asymptotically justified models for spatial extremes
- The *SpatialExtremes* provides functions to fit and analyse max-stable processes to spatial extremes
- I hope you will find it useful

Weak points

- The package is currently in intensive development - some bugs might (well almost surely) exist
- The optimisation might fail for complex data - need more robust optimisation. Double check your estimates!
- Computing the asymptotic covariance matrix is unstable too - finite difference may fail
- Other max-stable models will be available soon
- Other approaches for spatial extremes modelling are needed

- If you need more details, a package vignette has been written
- You can have a look by invoking *vignette("SpatialExtremesGuide")*
- You can also have a look at its web page

  `http://spatialextremes.r-forge.r-project.org/`

THANK YOU FOR YOUR ATTENTION!