Multivariate Extreme Value Thresholding (for Environmental Hazards)

K. Abayomi M. Lacer

[1] Environmental Studies Program & Department of Mathematical Sciences
Binghamton University
[2] ISYE, GaTech

Institute for Operations Research and Management Sciences, 2012

Outline

High Risk Hotspot

Between 1994-1998: Volcano eruption in Rabaul, Cyclone Justin in the Milne Bay (SE from map selection), and El Niño-induced drought

One Variable

- In the simplest of cases, thresholding is straightforward.

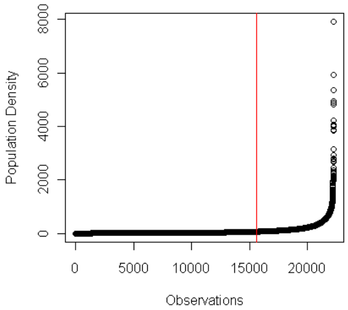

- This will separate data into regular-valued and extreme-valued portions.

One Variable

- In the simplest of cases, thresholding is straightforward.

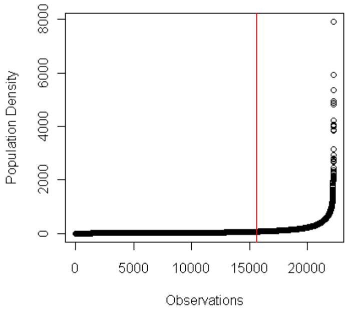

- This will separate data into regular-valued and extreme-valued portions.

Thresholding

Taking multivariate **q** we want to return the set $\mathcal{Q}$ such that

$$\mathcal{Q}_c = \{q | F(\mathbf{Q} > \mathbf{q}) > 1 - c\} \tag{1}$$

The 'independent' approach is to censor the data:

$$\mathcal{Q} \supset \mathcal{Q}_* = \{\mathbf{q} \mid q_i > c, \forall i\} \tag{2}$$

And the output is

$$\hat{F}(\mathbf{q}) = \mathbb{P}(\widehat{\mathbf{Q} \leq \mathbf{q}}) = \prod_{i=1}^{k} \mathbb{P}_n(Q_i \leq q_i)$$

Multi Variable

- But this doesn't yield a compact hull...

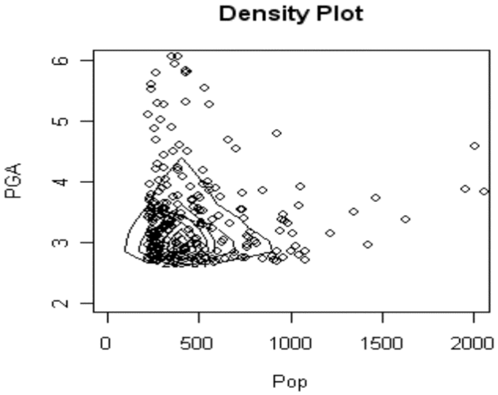

- ...if any subset of $\mathcal{Q}_c$ is non-monotone in the (real) F

Multi Variable

- But this doesn't yield a compact hull...

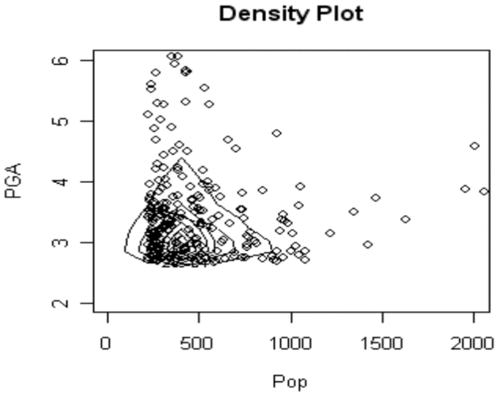

- ...if any subset of $\mathcal{Q}_c$ is non-monotone in the (real) F

Global Natural Disaster Risk Hotspots

Worldwide data has been gridded to $1\frac{1}{2}^{\circ}$ boxes for 8 predictor variables.

- GDP
- Population
- Peak Ground Acceleration (PGA)
- Floods
- Cyclones
- Drought
- Volcanoes
- Landslides

Global Natural Disaster Risk Hotspots

Worldwide data has been gridded to $1\frac{1}{2}^{\circ}$ boxes for 8 predictor variables.

- GDP
- Population
- Peak Ground Acceleration (PGA)
- Floods
- Cyclones
- Drought
- Volcanoes
- Landslides

Floods

.9 ptile of Flood counts

Volcanos

'.9' ptile of Volcano incidence

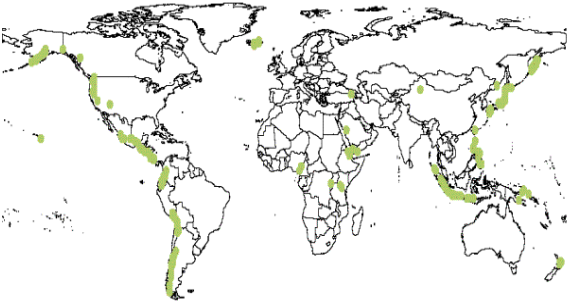

Droughts

Droughts: Classifying a drought.

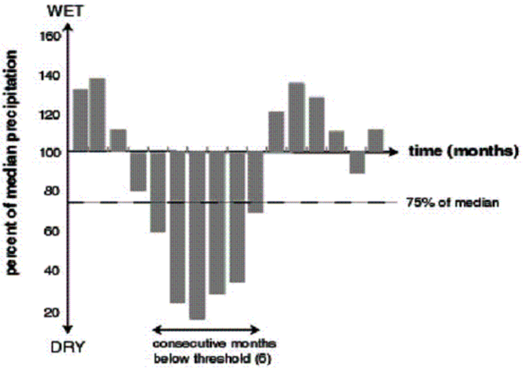

Example of a drought event defined by monthly precipitation being below a threshold of 75% of the long-term median value for at least 3 consecutive months. In this case, the duration of the event was 6 months.

Droughts

50 pct Weighted Anomaly Standardized Precipitation (WASP)

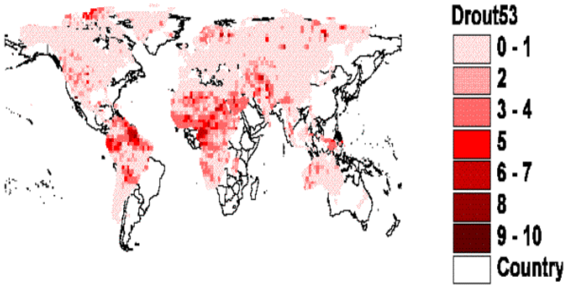

Droughts

75 pct Weighted Anomaly Standardized Precipitation (WASP)

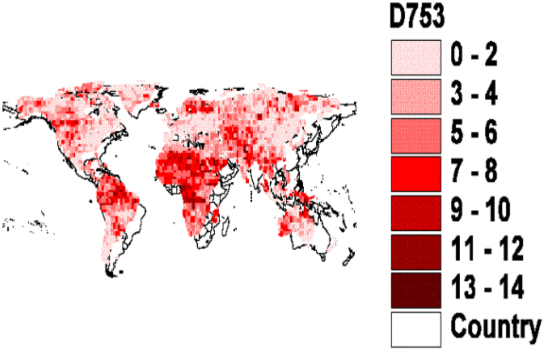

Droughts

Drought declaration vs. Drought classification

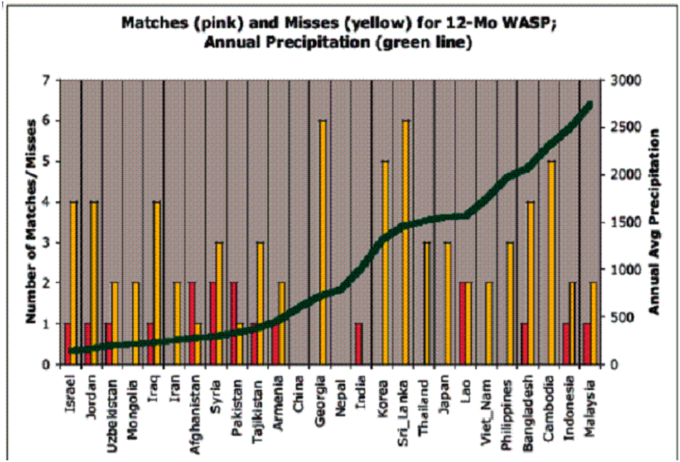

Quakes

Peak Ground Acceleration

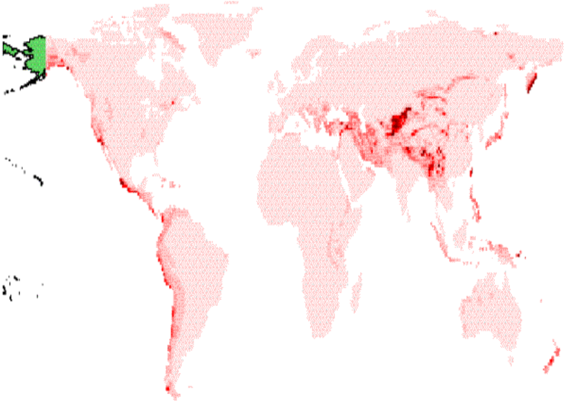

Population

Population Density

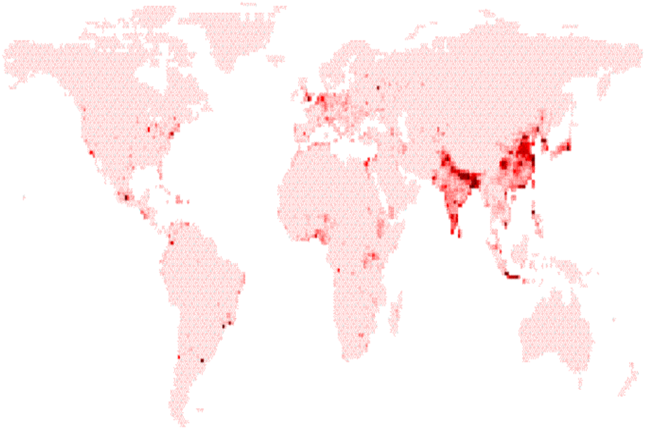

Income

GNP

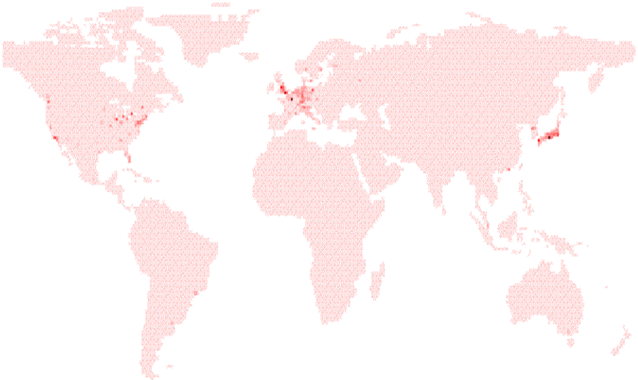

Select Bivariate Plots

- PGA vs. Floods
- PGA vs. Volcanoes

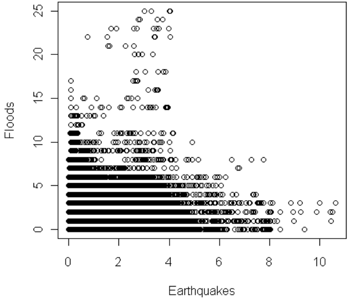

Select Bivariate Plots

- PGA vs. Volcanoes

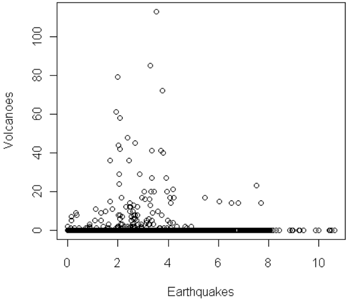

Multivariate Extreme Value Thresholding

A brute force approach is to:

- Select a thresholding level c
- Fit an extreme-valued (parametric) model to the data tail
- Measure a distance between the (parametric) model and an empirical distribution function

Multivariate Extreme Value Thresholding

A brute force approach is to:

- Select a thresholding level c
- Fit an extreme-valued (parametric) model to the data tail
- Measure a distance between the (parametric) model and an empirical distribution function

Multivariate Extreme Value Thresholding

A brute force approach is to:

- Select a thresholding level c
- Fit an extreme-valued (parametric) model to the data tail
- Measure a distance between the (parametric) model and an empirical distribution function

Multivariate Extreme Value Thresholding

A brute force approach is to:

- Select a thresholding level c
- Fit an extreme-valued (parametric) model to the data tail
- Measure a distance between the (parametric) model and an empirical distribution function

Parametric Model

Asymmetric Logistic Distribution (Tawn 1990):

$$G(x_1, \ldots, x_d) = \exp\left[-\sum_{b \in B}\left[\sum_{j \in b}\left(\frac{\theta_{j,b}}{y_j}\right)^{1/\alpha_b}\right]^{\alpha_b}\right]$$

- $j \in \{1, \ldots, d\}$, and y_j is the transformed data
- $B = \text{PowerSet}\{1, \ldots, d\} \setminus \emptyset$. Hence, $|B| = 2^d - 1$
- Say, $b = \{2, 4, 7\}$, then the inner sum is over $j = 2, 4, 7$
- $\alpha_b \in (0, 1]\ \forall\ b \in B \setminus B_1$ are dependence parameters
- $\theta_{j,b}$ are asymmetry parameters, with the constraint: $\sum_{b \in B_{(j)}} \theta_{j,b} = 1$ for $j = 1, \ldots, d$ to force univariate margins to be of the correct form. Here, $B_{(j)} = \{b \in B\ :\ j \in b\}$.
- $|B_{(j)}| = 2^{d-1}$.

Parametric Model

Asymmetric Logistic Distribution (Tawn 1990):

$$G(x_1, \ldots, x_d) = \exp\left[-\sum_{b\in B}\left[\sum_{j\in b}\left(\frac{\theta_{j,b}}{y_j}\right)^{1/\alpha_b}\right]^{\alpha_b}\right]$$

- $j \in \{1, \ldots, d\}$, and y_j is the transformed data
- $B = \text{PowerSet}\{1, \ldots, d\} \setminus \emptyset$. Hence, $|B| = 2^d - 1$
- Say, $b = \{2, 4, 7\}$, then the inner sum is over $j = 2, 4, 7$
- $\alpha_b \in (0, 1]\ \forall\ b \in B \setminus B_1$ are dependence parameters
- $\theta_{j,b}$ are asymmetry parameters, with the constraint: $\sum_{b\in B_{(j)}} \theta_{j,b} = 1$ for $j = 1, \ldots, d$ to force univariate margins to be of the correct form. Here, $B_{(j)} = \{b \in B\ :\ j \in b\}$.
- $|B_{(j)}| = 2^{d-1}$.

Parametric Model

Asymmetric Logistic Distribution (Tawn 1990):

$$G(x_1, \ldots, x_d) = \exp\left[-\sum_{b \in B}\left[\sum_{j \in b}\left(\frac{\theta_{j,b}}{y_j}\right)^{1/\alpha_b}\right]^{\alpha_b}\right]$$

- $j \in \{1, \ldots, d\}$, and y_j is the transformed data
- $B = \text{PowerSet}\{1, \ldots, d\} \setminus \emptyset$. Hence, $|B| = 2^d - 1$
- Say, $b = \{2, 4, 7\}$, then the inner sum is over $j = 2, 4, 7$
- $\alpha_b \in (0, 1]\ \forall\ b \in B \setminus B_1$ are dependence parameters
- $\theta_{j,b}$ are asymmetry parameters, with the constraint: $\sum_{b \in B_{(j)}} \theta_{j,b} = 1$ for $j = 1, \ldots, d$ to force univariate margins to be of the correct form. Here, $B_{(j)} = \{b \in B\ :\ j \in b\}$.
- $|B_{(j)}| = 2^{d-1}$.

Parametric Model

Asymmetric Logistic Distribution (Tawn 1990):

$$G(x_1,\ldots,x_d) = \exp\left[-\sum_{b\in B}\left[\sum_{j\in b}\left(\frac{\theta_{j,b}}{y_j}\right)^{1/\alpha_b}\right]^{\alpha_b}\right]$$

- $j \in \{1,\ldots,d\}$, and y_j is the transformed data
- $B = \text{PowerSet}\{1,\ldots,d\} \setminus \emptyset$. Hence, $|B| = 2^d - 1$
- Say, $b = \{2,4,7\}$, then the inner sum is over $j = 2,4,7$
- $\alpha_b \in (0,1]\ \forall\ b \in B \setminus B_1$ are dependence parameters
- $\theta_{j,b}$ are asymmetry parameters, with the constraint: $\sum_{b\in B_{(j)}} \theta_{j,b} = 1$ for $j = 1,\ldots,d$ to force univariate margins to be of the correct form. Here, $B_{(j)} = \{b \in B : j \in b\}$.
- $|B_{(j)}| = 2^{d-1}$.

Parametric Model

Asymmetric Logistic Distribution (Tawn 1990):

$$G(x_1, \ldots, x_d) = \exp\left[-\sum_{b \in B}\left[\sum_{j \in b}\left(\frac{\theta_{j,b}}{y_j}\right)^{1/\alpha_b}\right]^{\alpha_b}\right]$$

- $j \in \{1, \ldots, d\}$, and y_j is the transformed data
- $B = \text{PowerSet}\{1, \ldots, d\} \setminus \emptyset$. Hence, $|B| = 2^d - 1$
- Say, $b = \{2, 4, 7\}$, then the inner sum is over $j = 2, 4, 7$
- $\alpha_b \in (0, 1]\ \forall\ b \in B \setminus B_1$ are dependence parameters
- $\theta_{j,b}$ are asymmetry parameters, with the constraint: $\sum_{b \in B_{(j)}} \theta_{j,b} = 1$ for $j = 1, \ldots, d$ to force univariate margins to be of the correct form. Here, $B_{(j)} = \{b \in B\ :\ j \in b\}$.
- $|B_{(j)}| = 2^{d-1}$.

Parametric Model

Asymmetric Logistic Distribution (Tawn 1990):

$$G(x_1, \ldots, x_d) = \exp\left[- \sum_{b \in B} \left[\sum_{j \in b} \left(\frac{\theta_{j,b}}{y_j} \right)^{1/\alpha_b} \right]^{\alpha_b} \right]$$

- $j \in \{1, \ldots, d\}$, and y_j is the transformed data
- $B = \text{PowerSet}\{1, \ldots, d\} \setminus \emptyset$. Hence, $|B| = 2^d - 1$
- Say, $b = \{2, 4, 7\}$, then the inner sum is over $j = 2, 4, 7$
- $\alpha_b \in (0, 1] \; \forall \; b \in B \setminus B_1$ are dependence parameters
- $\theta_{j,b}$ are asymmetry parameters, with the constraint: $\sum_{b \in B_{(j)}} \theta_{j,b} = 1$ for $j = 1, \ldots, d$ to force univariate margins to be of the correct form. Here, $B_{(j)} = \{b \in B \; : \; j \in b\}$.
- $|B_{(j)}| = 2^{d-1}$.

Parametric Model

Asymmetric Logistic Distribution (Tawn 1990):

$$G(x_1, \ldots, x_d) = \exp\left[-\sum_{b \in B} \left[\sum_{j \in b} \left(\frac{\theta_{j,b}}{y_j} \right)^{1/\alpha_b} \right]^{\alpha_b} \right]$$

- $j \in \{1, \ldots, d\}$, and y_j is the transformed data
- $B = \text{PowerSet}\{1, \ldots, d\} \setminus \emptyset$. Hence, $|B| = 2^d - 1$
- Say, $b = \{2, 4, 7\}$, then the inner sum is over $j = 2, 4, 7$
- $\alpha_b \in (0, 1] \ \forall \ b \in B \setminus B_1$ are dependence parameters
- $\theta_{j,b}$ are asymmetry parameters, with the constraint: $\sum_{b \in B_{(j)}} \theta_{j,b} = 1$ for $j = 1, \ldots, d$ to force univariate margins to be of the correct form. Here, $B_{(j)} = \{b \in B \ : \ j \in b\}$.
- $|B_{(j)}| = 2^{d-1}$.

Conditional Representation

To derive the pdf, we make use of the positive stable (PS) distribution and its Laplace transform (Stephenson 2009):

- $\int_0^\infty h_1(s)\exp(-st)ds = \exp(-t^\alpha)$
- Take $S_b \sim \mathrm{PS}(\alpha_b)\ \forall\ b \in \mathrm{B} \setminus \mathrm{B}_1$, and $\mathbf{S} = \{S_b \mid b \in \mathrm{B} \setminus \mathrm{B}_1\}$.
- Then we have for $j = 1, \ldots, d$:

$$\Pr(X_j < x_j \mid \mathbf{S} = \mathbf{s}) = \exp\left[-\sum_{b \in B_{(j)}} s_b \left(\frac{\theta_{j,b}}{y_j}\right)^{1/\alpha_b}\right]$$

while $X_1, \ldots, X_d$ are conditionally independent given $\mathbf{S} = \mathbf{s}$

- Thus, each marginal asymmetric logistic pdf can be given by:

$$f_j(x_j|s) = \sigma_j^{-1} y_j^{-x i_j} \left[\sum_{b \in B_{(j)}} (z_{j,b}/\alpha_b)\right] \exp\left(-\sum_{b \in B_{(j)}} z_j, b\right)$$

where $z_{j,b} = s_b(\theta_{j,b}/y_j)^{1/\alpha_b}$

Conditional Representation

To derive the pdf, we make use of the positive stable (PS) distribution and its Laplace transform (Stephenson 2009):

- $\int_0^\infty h_1(s)\exp(-st)ds = \exp(-t^\alpha)$
- Take $S_b \sim \text{PS}(\alpha_b)\ \forall\ b \in \text{B} \setminus \text{B}_1$, and $\mathbf{S} = \{S_b \mid b \in \text{B} \setminus \text{B}_1\}$.
- Then we have for $j = 1, \ldots, d$:

$$\Pr(X_j < x_j \mid \mathbf{S} = \mathbf{s}) = \exp\left[-\sum_{b \in B_{(j)}} s_b \left(\frac{\theta_{j,b}}{y_j}\right)^{1/\alpha_b}\right]$$

  while $X_1, \ldots, X_d$ are conditionally independent given $\mathbf{S} = \mathbf{s}$

- Thus, each marginal asymmetric logistic pdf can be given by:

$$f_j(x_j|s) = \sigma_j^{-1} y_j^{-x_j} \left[\sum_{b \in B_{(j)}} (z_{j,b}/\alpha_b)\right] \exp\left(-\sum_{b \in B_{(j)}} z_j, b\right)$$

where $z_{j,b} = s_b(\theta_{j,b}/y_j)^{1/\alpha_b}$

Conditional Representation

To derive the pdf, we make use of the positive stable (PS) distribution and its Laplace transform (Stephenson 2009):

- $\int_0^\infty h_1(s)\exp(-st)ds = \exp(-t^\alpha)$
- Take $S_b \sim \text{PS}(\alpha_b)\ \forall\ b \in \text{B} \setminus \text{B}_1$, and $\mathbf{S} = \{S_b \mid b \in \text{B} \setminus \text{B}_1\}$.
- Then we have for $j = 1, \ldots, d$:

$$\Pr(X_j < x_j \mid \mathbf{S} = \mathbf{s}) = \exp\left[-\sum_{b \in B_{(j)}} s_b \left(\frac{\theta_{j,b}}{y_j}\right)^{1/\alpha_b}\right]$$

  while $X_1, \ldots, X_d$ are conditionally independent given $\mathbf{S} = \mathbf{s}$
- Thus, each marginal asymmetric logistic pdf can be given by:

$$f_j(x_j|s) = \sigma_j^{-1} y_j^{-x_j} \left[\sum_{b \in B_{(j)}} (z_{j,b}/\alpha_b)\right] \exp\left(-\sum_{b \in B_{(j)}} z_j, b\right)$$

  where $z_{j,b} = s_b(\theta_{j,b}/y_j)^{1/\alpha_b}$

Conditional Representation

To derive the pdf, we make use of the positive stable (PS) distribution and its Laplace transform (Stephenson 2009):

- $\int_0^\infty h_1(s)\exp(-st)ds = \exp(-t^\alpha)$
- Take $S_b \sim \text{PS}(\alpha_b)\ \forall\ b \in \text{B} \setminus \text{B}_1$, and $\mathbf{S} = \{S_b \mid b \in \text{B} \setminus \text{B}_1\}$.
- Then we have for $j = 1, \ldots, d$:

  $$\Pr(X_j < x_j \mid \mathbf{S} = \mathbf{s}) = \exp\left[-\sum_{b \in B_{(j)}} s_b \left(\frac{\theta_{j,b}}{y_j}\right)^{1/\alpha_b}\right]$$

  while $X_1, \ldots, X_d$ are conditionally independent given $\mathbf{S} = \mathbf{s}$
- Thus, each marginal asymmetric logistic pdf can be given by:

  $$f_j(x_j|s) = \sigma_j^{-1} y_j^{-xi_j}\left[\sum_{b \in B_{(j)}} (z_{j,b}/\alpha_b)\right] \exp\left(-\sum_{b \in B_{(j)}} z_j, b\right)$$

  where $z_{j,b} = s_b(\theta_{j,b}/y_j)^{1/\alpha_b}$

Parameter Estimation

- We begin by estimating the marginal parameters ($\mu_j, \sigma_j,$ and ξ_j) from univariate data and keep them fixed throughout.
- Simplifying assumptions: we consider high-dimensional (5 and more) asymmetry parameters to be trivial; also, we assume a non-informative prior.
- To obtain estimates for α and θ, we use Metropolis-Hastings within Gibbs to calculate conditional posterior means.

Parameter Estimation

- We begin by estimating the marginal parameters $(\mu_j, \sigma_j,$ and $\xi_j)$ from univariate data and keep them fixed throughout.
- Simplifying assumptions: we consider high-dimensional (5 and more) asymmetry parameters to be trivial; also, we assume a non-informative prior.
- To obtain estimates for α and θ, we use Metropolis-Hastings within Gibbs to calculate conditional posterior means.

Parameter Estimation

- We begin by estimating the marginal parameters ($\mu_j, \sigma_j,$ and ξ_j) from univariate data and keep them fixed throughout.
- Simplifying assumptions: we consider high-dimensional (5 and more) asymmetry parameters to be trivial; also, we assume a non-informative prior.
- To obtain estimates for α and θ, we use Metropolis-Hastings within Gibbs to calculate conditional posterior means.

Thresholding

To select the best threshold, we minimize distances between our parametric fit $G_{\hat{\theta}}$ and the empirical distribution function $\hat{F}_n$:

$$\hat{F}_n(t_1, \ldots, t_d) = \frac{1}{nk} \sum_{j=1}^{d} \sum_{i=1}^{n} \mathbf{1}\{x_{ij} < t_j\}$$

Thresholding

To select the best threshold, we minimize distances between our parametric fit $G_{\hat{\theta}}$ and the empirical distribution function $\hat{F}_n$:

$$\hat{F}_n(t_1, \ldots, t_d) = \frac{1}{nk} \sum_{j=1}^{d} \sum_{i=1}^{n} \mathbf{1}\{x_{ij} < t_j\}$$

Thresholding

To select the best threshold, we minimize distances between our parametric fit $G_{\hat{\theta}}$ and the empirical distribution function $\hat{F}_n$:

$$\hat{F}_n(t_1, \ldots, t_d) = \frac{1}{nk} \sum_{j=1}^{d} \sum_{i=1}^{n} \mathbf{1}\{x_{ij} < t_j\}$$

Pickands Type

Pickands suggesting minimizing KS distance

$$d_k = sup_{\mathbf{q}} |\hat{F}_n(\mathbf{q}) - \hat{G}_\theta(\mathbf{q})|$$

with $k = 1, 2, ...[n/4]$

Joe Type

Joe suggests computing measure of association and setting cutoff to maximize tail dependence

$$max_k \; \tau_{1-k/n} = max \; \tau(\mathbf{q}|\mathbf{q} > \mathbf{C}_k)$$

where $\mathbf{C}_k$ is the (multivariate) kth cutoff...

...which (we show) is equivalent to

$$= max_k \; 4E[C_\theta(\mathbf{q}|\mathbf{q} > \mathbf{C}_k)] - 1$$

where C_θ is the copula induced by G_θ

Joe Type

Joe suggests computing measure of association and setting cutoff to maximize tail dependence

$$max_k \; \tau_{1-k/n} = max \; \tau(\mathbf{q}|\mathbf{q} > \mathbf{C}_k)$$

where $\mathbf{C}_k$ is the (multivariate) kth cutoff...
...which (we show) is equivalent to

$$= max_k \; 4E[C_\theta(\mathbf{q}|\mathbf{q} > \mathbf{C}_k)] - 1$$

where C_θ is the copula induced by G_θ

Joe Type

Joe suggests computing measure of association and setting cutoff to maximize tail dependence

$$max_k\ \tau_{1-k/n} = max\ \tau(\mathbf{q}|\mathbf{q} > \mathbf{C}_k)$$

where $\mathbf{C}_k$ is the (multivariate) kth cutoff...
...which (we show) is equivalent to

$$= max_k\ 4E[C_\theta(\mathbf{q}|\mathbf{q} > \mathbf{C}_k)] - 1$$

where C_θ is the copula induced by G_θ

Our generalization of Joe Type

Which we can generalize, for a particular G_θ as a likelihood problem (in θ and k): $max_\theta \; min_k \; d_\theta(\mathbf{q}, \mathbf{C}_{k,\theta})$ is

$$\mathcal{L}_{(\theta,k)} = max_\theta \; min_k \; E[ln(\frac{dG_\theta(\mathbf{q})}{dG_\theta(\mathbf{C}_k)})]$$

Since the RHS is (essentially)

$$E_{G_\theta}[ln(\mathbb{P}(\mathbf{q} \perp\!\!\!\perp \mathbf{C}_k))]$$

Kendall's Taus on tails

$\tau_{1-k/n}$	$\tau_{.9}$	$\tau_{.95}$	$\tau_{.99}$
Pop-Pga	.072	.186	.472
GNP-Flood	.113	.270	.326
GNP-Drought	.208	.290	.168

70-percentile

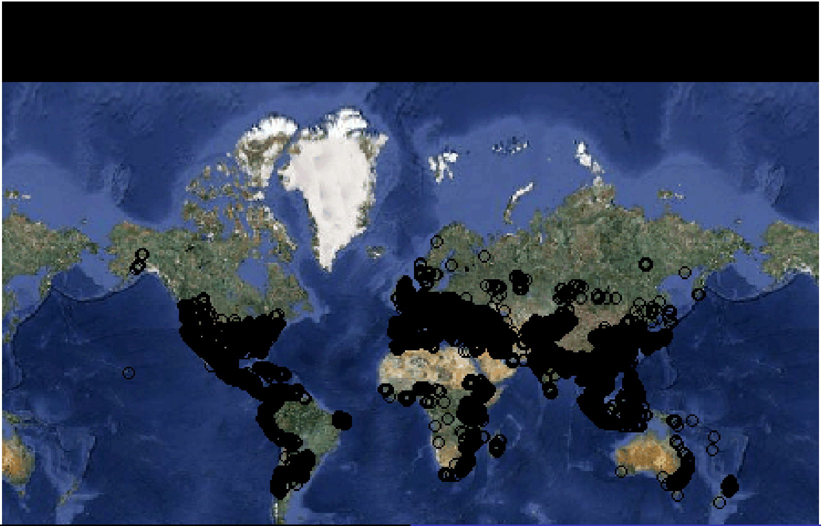

75-percentile

80-percentile

85-percentile

90-percentile

95-percentile

99-percentile

Comments

- For $dim(\mathbf{Q}) \geq 2$, d_k is distributionally unspecified
- However $\mathcal{L}_{(\theta,k)} \sim \chi^2_{dim(\mathbf{Q})-1} \cdots$
- ...to say nothing of the computational time to estimate G_θ

Comments

- For $dim(\mathbf{Q}) \geq 2$, d_k is distributionally unspecified
- However $\mathcal{L}_{(\theta,k)} \sim \chi^2_{dim(\mathbf{Q})-1}$ ···
- ...to say nothing of the computational time to estimate G_θ

Comments

- For $dim(\mathbf{Q}) \geq 2$, d_k is distributionally unspecified
- However $\mathcal{L}_{(\theta,k)} \sim \chi^2_{dim(\mathbf{Q})-1}$ ···
- ...to say nothing of the computational time to estimate G_θ

Comments

- For $dim(\mathbf{Q}) \geq 2$, d_k is distributionally unspecified
- However $\mathcal{L}_{(\theta,k)} \sim \chi^2_{dim(\mathbf{Q})-1}$ ···
- ...to say nothing of the computational time to estimate G_θ

Comments: Another approach?

Since the extreme part of $\mathbf{Q}$ is distributed G_θ
it has Fisher Inf: $E[[\partial_\theta[dG_\theta(\mathbf{q})]]^2]$

and the negentropy

$$J(\mathbf{Q}) = E[[\partial_\theta[dG_\theta(\mathbf{q})] - \partial_\theta[d\Phi_{\theta*}(\mathbf{q})]]^2]$$

Comments: Another approach?

Since the extreme part of $\mathbf{Q}$ is distributed G_θ
it has Fisher Inf: $E[[\partial_\theta[dG_\theta(\mathbf{q})]]^2]$

and the negentropy

$$J(\mathbf{Q}) = E[[\partial_\theta[dG_\theta(\mathbf{q})] - \partial_\theta[d\Phi_{\theta*}(\mathbf{q})]]^2]$$

Comments: Another approach?

Since the extreme part of $\mathbf{Q}$ is distributed G_θ
it has Fisher Inf: $E[[\partial_\theta[dG_\theta(\mathbf{q})]]^2]$

and the negentropy

$$J(\mathbf{Q}) = E[[\partial_\theta[dG_\theta(\mathbf{q})] - \partial_\theta[d\Phi_{\theta*}(\mathbf{q})]]^2]$$

Comments: Another approach?

Since the extreme part of $\mathbf{Q}$ is distributed G_θ
it has Fisher Inf: $E[[\partial_\theta[dG_\theta(\mathbf{q})]]^2]$

and the negentropy

$$J(\mathbf{Q}) = E[[\partial_\theta[dG_\theta(\mathbf{q})] - \partial_\theta[d\Phi_{\theta*}(\mathbf{q})]]^2]$$

Comments: Another approach?

which by linearity of the partial differential is

$$= E[[\partial_\theta[dG_\theta(\mathbf{q}) - d\Phi_{\theta*}(\mathbf{q})]]^2]$$

and if we can find a $\Gamma = \Gamma(\theta, \theta^*)$ we can write

$$= E[[\partial_\theta[d[G_\Gamma(\mathbf{q}) - \Phi_\Gamma(\mathbf{q})]]]^2]$$

and, replacing $\mathbf{q}$ with $\mathbf{q}_k$ this is a version of the Fisher information of the extreme valued data under the Normal distributional assumption.

Comments: Another approach?

which by linearity of the partial differential is

$$= E[[\partial_\theta[dG_\theta(\mathbf{q}) - d\Phi_{\theta^*}(\mathbf{q})]]^2]$$

and if we can find a $\Gamma = \Gamma(\theta, \theta^*)$ we can write

$$= E[[\partial_\theta[d[G_\Gamma(\mathbf{q}) - \Phi_\Gamma(\mathbf{q})]]]^2]$$

and, replacing $\mathbf{q}$ with $\mathbf{q}_k$ this is a version of the Fisher information of the extreme valued data under the Normal distributional assumption.

Comments: Another approach?

which by linearity of the partial differential is

$$= E[[\partial_\theta[dG_\theta(\mathbf{q}) - d\Phi_{\theta^*}(\mathbf{q})]]^2]$$

and if we can find a $\Gamma = \Gamma(\theta, \theta^*)$ we can write

$$= E[[\partial_\theta[d[G_\Gamma(\mathbf{q}) - \Phi_\Gamma(\mathbf{q})]]]^2]$$

and, replacing $\mathbf{q}$ with $\mathbf{q}_k$ this is a version of the Fisher information of the extreme valued data under the Normal distributional assumption.

Summary

- We fit a flexible model to high-dimensional data.
- This framework allows for the identification of multivariate extremes via thresholding...
- ...for multivariate distances...
- ...one which can be expressed as likelihood via the copula approach, and another candidate as Fisher information.

Summary

- We fit a flexible model to high-dimensional data.
- This framework allows for the identification of multivariate extremes via thresholding...
- ...for multivariate distances...
- ...one which can be expressed as likelihood via the copula approach, and another candidate as Fisher information.

Summary

- We fit a flexible model to high-dimensional data.
- This framework allows for the identification of multivariate extremes via thresholding...
- ...for multivariate distances...
- ...one which can be expressed as likelihood via the copula approach, and another candidate as Fisher information.

Summary

- We fit a flexible model to high-dimensional data.
- This framework allows for the identification of multivariate extremes via thresholding...
- ...for multivariate distances...
- ...one which can be expressed as likelihood via the copula approach, and another candidate as Fisher information.

Summary

- The likelihood-copula distance method, on data ending in 2003, identified several, post hoc, locations. Significantly, Haiti.
- Compare thresholded 'hotspots' with disaster record from 2003-2012.

Summary

- The likelihood-copula distance method, on data ending in 2003, identified several, post hoc, locations. Significantly, Haiti.
- Compare thresholded ‘hotspots’ with disaster record from 2003-2012.

Summary

- The likelihood-copula distance method, on data ending in 2003, identified several, post hoc, locations. Significantly, Haiti.
- Compare thresholded ‘hotspots’ with disaster record from 2003-2012.