Conditional sampling for max-stable random fields[1]

Department of Statistics
The University of Michigan
Ann Arbor, MI, USA.

Yizao Wang and Stilian A. Stoev
{yizwang, sstoev}@umich.edu

Summary

Max-stable random fields play a central role in modeling extreme value phenomena. We obtain an explicit formula for the conditional probability in general max-linear models, which include a large class of max-stable random fields. As a consequence, we develop an algorithm for efficient and exact sampling from the conditional distributions. Our method provides a computational solution to the prediction problem for spectrally discrete max-stable random fields. This work offers new tools and a new perspective to many statistical inference problems for spatial extremes, arising, for example, in meteorology, geology, and environmental applications.

Introduction

A max–stable random field with $\alpha-$Fréchet marginals has the representation

$$X_t = \int^{e}_{U} f_t(u) M_\alpha(du), \quad f_t \in L^\alpha_+(U,\mu) \tag{1}$$

with $t \in T$, where T can be $\mathbb{R}$, $\mathbb{R}^2$, $\mathbb{R}^d$, etc. (see e.g. [4] & [7]).
The *finite–dimensional distributions* of $X = \{X_t\}_{t\in T}$ are (for $t_i \in T,\ x_i > 0$) as follows:

$$\mathbb{P}\{X_{t_i} \le x_i,\ 1 \le i \le n\} = \exp\Big\{-\int_U \max_{1\le i\le n}\Big(\frac{f_{t_i}(u)}{x_i}\Big)^\alpha \mu(du)\Big\} = \exp\Big\{-\int_{\mathbb{S}_+^{n-1}}\Big(\max_{1\le i\le n} w_i/x_i\Big)^\alpha \Gamma(\mathrm{d}w)\Big\}.$$

Here $\mathbb{S}_+^{n-1} = \{w \in \mathbb{R}^n_+ : \max_{1\le i\le n} w_i = 1\}$ is the positive unit sphere in the sup-norm, and Γ is a unique finite measure on $\mathbb{S}_+^{n-1}$ called the *spectral measure* of the distribution (see e.g. [5]).

The Max–Linear Model

Let $a_{i,j} \ge 0$ and $Z_j,\ 1 \le j \le p$ be i.i.d. $\alpha-$Fréchet.

$$X_i := \bigvee_{j=1}^{p} a_{i,j} Z_j, \quad (1 \le i \le n) \tag{2}$$

denoted by

$$X = A \odot Z, \quad \text{with } X = (X_i)_{i=1}^n,\ Z = (Z_j)_{j=1}^p,\ A = (a_{i,j})_{n\times p} \tag{3}$$

is called a max–linear model.

Example Let $\phi_t(j),\ t \in T,\ 1 \le j \le p$ be non–negative functions. Consider the max–stable ($\alpha-$Féchet) random field:

$$X_t := \bigvee_{j=1}^{p} \phi_t(j) Z_j, \quad t \in T. \tag{4}$$

The random vectors $(X_{t_i})_{i=1}^n$ follow the max–linear model in (3) with $A = (\phi_{t_i}(j))_{n\times p}$.

Remarks

- Since the Z_i's are i.i.d. $\alpha-$Fréchet, the *max-linear model* has the representation (1) with *simple spectral functions* f_t's.
- The spectral measure Γ for the max–linear model is *discrete*. Hence the models (3) and (4) are spectrally discrete.
- The max–linear model in (3) can approximate (in $\mathbb{P}$) any general model (1).

Prediction for Max–Stable Random fields

The Problem: A max-stable random field $\{X_t\}_{t\in T}$ is observed at locations $t_1, \cdots, t_n \in T$. The goal is to predict, $X_{s_1}, \cdots, X_{s_m}$.

Remarks

- Previously addressed by [3, 1, 2] among others.
- Important but very challenging problem.
- No hope for an *exact* analytical solution.

[1] The authors were partially supported by NSF grant DMS-0806094 at the University of Michigan.
[2] We thank Dan Cooley and Météo France for sharing and acquiring the rainfall data set.

A Computational Solution

In our recent work [8], we obtain a computational solution of the general prediction problem for the *max–linear* model.

- We focus on the spectrally discrete max–linear model (3).
- Obtain an exact formula (see [8]) of the regular conditional probability distribution of

$$Z|X = x \tag{5}$$

- Thus, obtain the conditional distribution:

$$Y|X = x, \quad \text{where } Y = B \odot Z, \tag{6}$$

with $Y = (Y_k)_{k=1}^m,\ B = (b_{k,j})_{m\times p}$, where B is a given non–negative matrix.
- Sample efficiently and exactly (with a computer) from (5) and hence (6).
- Thus, obtain empirical (Monte Carlo) approximations for functionals of $Y|X = x$.
- Can obtain L^1, L^2, etc. *optimal predictors* such as conditional medians, means, etc.; *conditional quantiles* and *confidence sets* (not just intervals!).

An algorithm for efficient and exact sampling from the conditional distribution is available in C/C++ with interface as an R-package from Yizao Wang's web page:

http://www.stat.lsa.umich.edu/~yizwang/software/maxLinear/

Prediction for MARMA

The max–autoregressive moving averages (MARMA(m,q)) are the stationary time series solutions to:

$$X_t = \phi_1 X_{t-1} \vee \cdots \vee \phi_m X_{t-m} \vee Z_t \vee \theta_1 Z_{t-1} \vee \cdots \vee \theta_q Z_{t-q}, \quad t \in \mathbb{Z}, \tag{7}$$

with $\phi_i \ge 0, \theta_j \ge 0, 1 \le i \le m, 1 \le j \le q$ and i.i.d. $1-$Fréchet Z_t's.
If $\phi^* = \bigvee_{i=1}^m \phi_i < 1$, then (7) has a unique solution:

$$X_t = \bigvee_{j=0}^{\infty} \psi_j Z_{t-j}, \quad \text{with } \psi_j = \bigvee_{k=0}^{j\wedge q} \alpha_{j-k}\theta_k,$$

where $\alpha_j = 0,\ j < 0,\ \alpha_0 = 1$ and $\alpha_j = \phi_1\alpha_{j-1} \vee \phi_2\alpha_{j-2} \vee \cdots \vee \phi_m\alpha_{j-m},\ j \ge 1$.
Since the ψ_j's decay exponentially, by truncation, one can approximate well

$$X_t \overset{\mathbb{P}}{\approx} \widetilde{X}_t = \bigvee_{j=0}^{M} \psi_j Z_{t-j},$$

by a max–linear model.
The graph illustrates prediction of a MARMA(3,0) with $\phi_1 = 0.7, \phi_2 = 0.5$ and $\phi_3 = 0.3$. Observed are the first 100 values and predicted are the next 50.

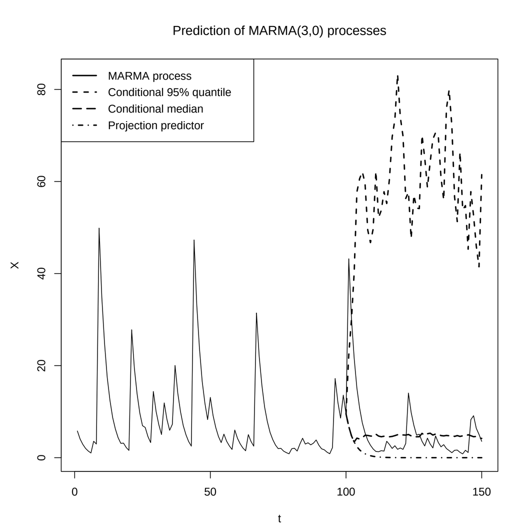

Lag	Coverage	Width	Proj Pred
1	0.956	13.06	0.706
2	0.952	26.6	0.503
3	0.954	37.8	0.356
4	0.957	45.6	0.253
5	0.966	51.2	0.178
10	0.947	62.8	0.029
20	0.943	66.0	0.001
30	0.951	66.2	0.000
40	0.955	65.4	0.000

The table shows the coverage and width of 95% upper prediction intervals for MARMA. The third column reports

$$\mathbb{P}\{X_{100+s} \le \widehat{X}^{(DR)}_{100+s} | X_t,\ 1 \le t \le 100\},$$

where $\widehat{X}^{(DR)}_{100+s}$ are the Projection Predictors of Davis & Resnick [3].

Remarks

- The accurate coverages demonstrate the validity of our algorithm.
- Naturally, the width of the prediction intervals grows with the lag.
- One can numerically quantify the prediction error of previous projection predictors of [3].

The Discrete Smith Model

Take in (1) $T = \mathbb{R}^2$, $U = \mathbb{R}^2$ and $f_t(u) := \phi(t-u)$ with

$$\phi(t_1,t_2) := \frac{\beta_1\beta_2}{2\pi\sqrt{1-\rho^2}}\exp\Big\{-\frac{1}{2(1-\rho^2)}\Big\{\beta_1^2 t_1^2 - 2\rho\beta_1\beta_2 t_1 t_2 + \beta_2^2 t_2^2\Big\}\Big\}.$$

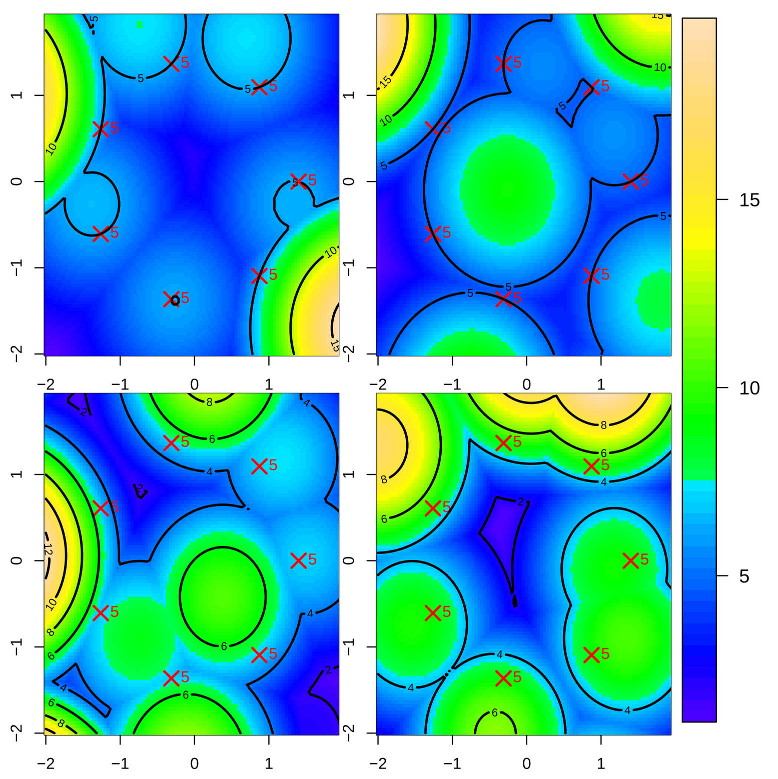

Smith moving maxima random field model

- Parameters: $\beta_1, \beta_2 > 0$ and $-1 < \rho < 1$.
- To apply our conditional sampling methodology, we discretize the extremal integral:

$$X_t := \bigvee_{-q\le j_1,j_2\le q-1} h^{2/\alpha}\phi(t-u_{j_1j_2})Z_{j_1j_2},$$

where $u_{j_1j_2} = ((j_1+1/2)h, (j_2+1/2)h)$.
- Figure on the left: 4 conditional samples with $\beta_1 = \beta_2 = 1$, $\rho = 0$, given 7 observed values (all equal to 5).

Rainfall Maxima in Bourgogne[2]

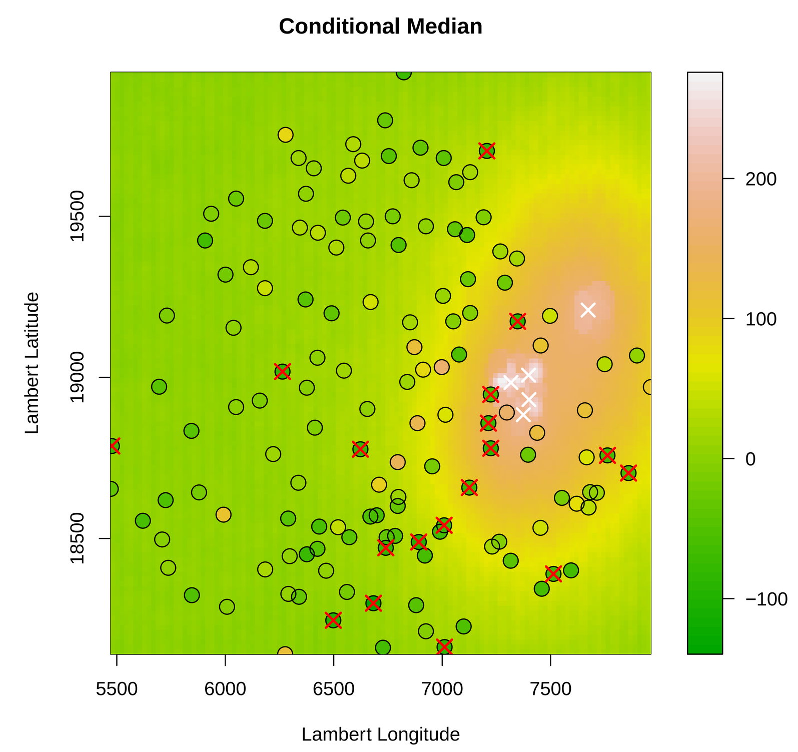

- The discrete Smith model is applied to 50–year maxima at 146 stations in Bourgogne, France.
- Data marginally transformed to a (non-standard) Gumbel scale.
- The heatmap shows the conditional median based on 5 observations with white crosses.
- The shades of the circles indicate true values.
- Crossed circles are outside a two–sided 95% prediction interval.
- The Smith model was fit using cross–validation.

References

[1] D. Cooley, Ph. Naveau, and P. Poncet. Variograms for spatial max-stable random fields. in *Lecture Notes in Statist.*, Vol. 187, pp 373–390. Springer, New York, 2006.

[2] D. Cooley, D. Nychka, and Ph. Naveau. Bayesian spatial modeling of extreme precipitation return levels. *J. Amer. Statist. Assoc.*, 102(479):824–840, 2007.

[3] R. A. Davis and S. I. Resnick. Basic properties and prediction of max-ARMA processes. *Adv. in Appl. Probab.*, 21(4):781–803, 1989.

[4] L. de Haan. A spectral representation for max–stable processes. *Annals of Probability*, 12(4):1194–1204, 1984.

[5] S. I. Resnick. *Extreme Values, Regular Variation and Point Processes.* Springer, New York, 1987.

[6] R. L. Smith. Max–stable processes and spatial extremes. unpublished manuscript, 1990.

[7] S.A. Stoev and M. S. Taqqu. Extremal stochastic integrals: a parallel between max–stable processes and $\alpha-$stable processes. *Extremes*, 8:237–266, 2005.

[8] Y. Wang and S.A. Stoev. Conditional sampling for max-stable random fields. Technical Report 509. Department of Statistics, University of Michigan, 2010. http://arxiv.org/abs/1005.0312.