Bayesian inference for functional extreme events defined via partially unobserved processes

Max Thannheimer[1] and Marco Oesting[1,2]

[1]University of Stuttgart, Institute for Stochastic and Applications, Allmandring 5b, 70569 Stuttgart, Germany
[2]University of Stuttgart, Stuttgart Center for Simulation Science (SC SimTech), Allmandring 5b, 70569 Stuttgart, Germany

In order to describe the extremal behaviour of some stochastic process X approaches from univariate extreme value theory are typically generalized to the spatial domain. Besides max-stable processes that can be used in analogy to the block maxima approach, a generalized peaks-over-threshold approach can be used allowing us to consider single extreme events. These can be flexibly defined as exceedances of a risk functional ℓ such as a spatial average applied to X. Inference for the resulting limit process, the so-called ℓ-Pareto process, requires the evaluation of $\ell(X)$ and thus the knowledge of the whole process X. In practical applications, we face the challenge that observations of X are only available at single sites.

To overcome this issue, we propose a two-step MCMC-algorithm in a Bayesian framework. In a first step, we sample from X conditionally on the observations in order to evaluate which observations lead to ℓ-exceedances. In a second step, we use these exceedances to sample from the posterior distribution of the parameters of the limiting ℓ-Pareto process. Alternating these steps results in a full Bayesian model for the extremes of X.

We show that, under appropriate assumptions, the probability of classifying an observation as ℓ-exceedance in the first step converges to the desired probability. Furthermore, given the first step, the distribution of the Markov chain constructed in the second step converges to the posterior distribution of interest. Our procedure is compared to the Bayesian version of the standard procedure in a simulation study.