Tracking evolving environmental risk with Bayesian structural time series

Bastiaan Van Velthoven[1] and Stijn Luca[1]

[1]Ghent University, Department of Data Analysis and Mathematical Modelling

Environmental risk in a changing climate is a time-series problem: return periods evolve through time and should be monitored with uncertainty. We present a Bayesian structural time series framework for extremes that extends the Gaussian dynamic linear model to block extremes by combining latent Markov state-space models with Generalized Extreme Value observations. The framework decomposes non-stationarity into interpretable components such as trend, seasonality/cycle, interventions, and covariate effects, while allowing bulk and tail behaviour to evolve differently. This is crucial in environmental applications, where changes in extremes need not and do not mirror changes in means. Since climate signals often evolve slowly, state innovations can be much smaller than observation noise, making posterior computation difficult. We address this with a non-centred reparameterization, Bayesian sparsity priors for component selection, tailored particle MCMC for accurate inference, and a fast Laplace-based pseudo-Gaussian approximation for scalable inference and online updating.

We illustrate the methodology in simulation studies and on long temperature and precipitation records and show how to deal with the contrasting inferential and computational challenges of temperature extremes, which are bounded and primarily location-driven, versus precipitation extremes, which are heavy-tailed and may exhibit time-varying scale dynamics suggestive of stochastic-volatility-type behaviour.