

We propose a novel statistical model to describe spatio-temporal extreme events. The model can be used, for instance, to estimate extreme-value temporal pattern such as seasonality and trend, and further to predict the distribution of extreme events in the future. Such model usually involves thousands or even millions of variables in the spatio-temporal domain, whereas only one single observation is available for each location and time point. To address this challenge, previous works usually employ learning and inference methods that are computationally burdensome, and therefore are prohibitive for large-scale data. Moreover, they assume that the shape and scale parameters of the extreme-value distributions are constant across the spatio-temporal domain, which is often too restrictive in practice. In this paper, we break through these limitations by exploring graphical models to capture the highly structured dependencies among the parameters of extreme-value distributions. Furthermore, we develop an efficient stochastic variational inference (SVI) algorithm to learn the parameters of the resulting non-Gaussian graphical model. The computational complexity of the SVI algorithm is sublinear in the number of variables, thus enabling the proposed model to tackle large-scale spatio-temporal data in real-life applications. Results of both synthetic and real data demonstrate the effectiveness of the proposed approach.