Computers & Geosciences

Acceleration of hidden Markov model fitting using graphical processing units, with application to low-frequency tremor classification

Stoltz, M; Stoltz, George G; Obara, K; Wang, T and Bryant, D

Abstract
Hidden Markov models (HMMs) are general purpose models for time-series data widely used across the sciences because of their flexibility and elegance. Fitting HMMs can often be computationally demanding and time consuming, particularly when the number of hidden states is large or the Markov chain itself is long. Here we introduce a new Graphical Processing Unit (GPU)-based algorithm designed to fit long-chain HMMs, applying our approach to a model for low-frequency tremor events. Even on a modest GPU, our implementation resulted in an increase in speed of several orders of magnitude compared to the standard single processor algorithm. This permitted a full Bayesian inference of uncertainty related to model parameters and forecasts based on posterior predictive distributions. Similar improvements would be expected for HMM models given large number of observations and moderate state spaces (states with current hardware). We discuss the model, general GPU architecture and algorithms and report performance of the method on a tremor dataset from the Shikoku region, Japan. The new approach led to improvements in both computational performance and forecast accuracy, compared to existing frequentist methodology.