Local Gaussian process approximation for large computer experiments

Robert B. Gramacy
The University of Chicago, Chicago IL, USA - rbgramacy@chicagobooth.edu

We provide a new approach to approximate emulation of large computer experiments. By focusing expressly on desirable properties of the predictive equations, we derive a family of local sequential design schemes that dynamically define the support of a Gaussian process predictor based on a local subset of the data. We further derive expressions for fast sequential updating of all needed quantities as the local designs are built-up iteratively. Then we show how independent application of our local design strategy across the elements of a vast predictive grid facilitates a trivially parallel implementation. The end result is a global predictor able to take advantage of modern multicore architectures, GPUs, and cluster computing, while at the same time allowing for a non stationary modeling feature as a bonus. We demonstrate our method on two examples utilizing designs sized in the tens of thousands to over a million data points. Comparisons are made to the method of compactly supported covariances, and we present an application to computer model calibration.
Keywords: sequential design, sequential updating, active learning, surrogate model, emulator, compactly supported covariance, local kriging neighborhoods