What is the dimension of a stochastic process?

Victor Panaretos (EPFL Lausanne)

How can we determine whether a mean-square continuous stochastic process is, in fact, finite-dimensional, and if so, what its actual dimension is? And is it possible to do so at a given level of confidence? This question is central to a great deal of methods for functional data analysis, which require low-dimensional representations whether by functional PCA or other methods. The difficulty is that the determination is to be made on the basis of iid replications of the process observed discretely and with measurement error contamination. This adds a ridge to the empirical covariance, obfuscating the underlying dimension. We build a matrix-completion-inspired test procedure that circumvents this issue by measuring the best possible least square fit of the empirical covariance's off-diagonal elements, optimised over covariances of given finite rank. For a fixed grid of sufficient size, we determine the statistic's asymptotic null distribution as the number of replications grows. We then use it to construct a bootstrap implementation of a stepwise testing procedure corresponding to the collection of hypothesis formalising the question at hand.

This is based on joint work with Anirvan Chakraborty (EPFL).