Maximum likelihood estimation of a multidimensional log-concave density

by

Richard Samworth University of Cambridge Centre for Mathematical Sciences Wilberforce Road Cambridge CB3 0WB, UK rjs57@cam.ac.uk

Abstract

If $X_1, ..., X_n$ are a random sample from a log-concave density f in \mathbb{R}^d , then with probability one there exists a unique maximum likelihood estimator \hat{f}_n of f. The use of this estimator is attractive because, unlike kernel density estimation, the estimator is fully automatic, with no smoothing parameters to choose. We exhibit an iterative algorithm for computing the estimator and show how the method can be combined with the EM algorithm to fit finite mixtures of log-concave densities. Applications and recent theoretical results on the performance of the estimator will be discussed, and the talk will be illustrated with pictures from the R package LogConcDEAD.