Computing highest density regions for continuous univariate distributions with known probability functions

Ben O'Neill[1]

Abstract

We examine the problem of computing the highest density region (HDR) in a computational context where the user has access to a density function and quantile function for the distribution (e.g., in the statistical language R). We examine several common classes of continuous univariate distributions based on the shape of the density function; this includes monotone densities, quasi-concave and quasi-convex densities, and general multimodal densities. In each case we show how the user can compute the HDR from the quantile and density functions by framing the problem as a nonlinear optimisation problem. We implement these methods in R to obtain general functions to compute HDRs for classes of distributions, and for commonly used families of distributions. We compare our method to existing R packages for computing HDRs and we show that our method performs favourably in terms of both accuracy and average speed.