

STEIN POINTS

François-Xavier BRIOL
University of Warwick, UK

An important task in computational statistics is to approximate a posterior distribution with an empirical measure supported on a set of representative points. This paper focuses on methods where the selection of points is essentially deterministic, with an emphasis on achieving accurate approximation when the number of points is small. To this end, we present an algorithm called Stein Points. The idea is to exploit either a greedy or a conditional gradient method to iteratively minimise a Stein discrepancy between the empirical measure and the posterior distribution. Our empirical results demonstrate that Stein Points enable accurate approximation of the posterior at modest computational cost. In addition, theoretical results are provided to establish convergence of the method.

This talk is based on joint work with Wilson Y. Chen, Lester Mackey, Jackson Gorham and Chris J. Oates.