

The Intrinsic Dimension and Complexity of Importance Sampling

A. M. STUART

The University of Warwick, UK

Email: a.m.stuart@warwick.ac.uk

The basic idea of importance sampling is to use independent samples from one measure in order to approximate expectations with respect to another measure. Understanding how many samples are needed is key to understanding the efficiency of the method, and hence to understanding when it will be effective and when it will not. It is intuitive that the size of the difference between the measure which is sampled, and the measure against which expectations are to be computed, is key to the efficiency. An implicit challenge in many of the published works in this area is to find useful quantities which measure this difference in terms of parameters which are pertinent for the practitioner. The subject has attracted substantial interest recently from within a variety of communities, including atmospheric sciences, machine learning, statistics and subsurface geophysics. The objective of this talk is to overview and unify the resulting literature in the area, and to create an overarching framework. The general setting is studied in some detail, followed by deeper development in the context of Bayesian inverse problems and filtering.

This is joint work with S. Agapiou, O. Papaspiliopoulos and D. Sanz-Alonso. it may be found as [arXiv.1511.06196](https://arxiv.org/abs/1511.06196)