Available online at www.sciencedirect.com

ScienceDirect

Stochastic Processes and their Applications 126 (2016) 138–170

stochastic processes and their applications

www.elsevier.com/locate/spa

Large deviations for weighted empirical measures arising in importance sampling

Henrik Hult, Pierre Nyquist*

Department of Mathematics, KTH, 100 44 Stockholm, Sweden

Received 29 August 2014; received in revised form 1 June 2015; accepted 14 August 2015
Available online 21 August 2015

Abstract

In this paper the efficiency of an importance sampling algorithm is studied by means of large deviations for the associated weighted empirical measure. The main result, stated as a Laplace principle for these weighted empirical measures, can be viewed as an extension of Sanov's theorem. The main theorem is used to quantify the performance of an importance sampling algorithm over a collection of subsets of a given target set as well as quantile estimates. The analysis yields an estimate of the sample size needed to reach a desired precision and of the reduction in cost compared to standard Monte Carlo.

MSC: primary 60F10; 65C05; secondary 60G57; 68U20

Keywords: Large deviations; Empirical measures; Importance sampling; Monte Carlo

1. Introduction

Computational methods based on stochastic simulation – the collective term for simulating a physical system, involving random coefficients, on a computer – are fundamental in a wide range of applications, including the natural sciences, computer science, economics and finance, operations research, engineering sciences such as power grids, reliability, solid mechanics, etc. As the complexity of the mathematical models used in these subjects has increased, so has the

* Corresponding author.
E-mail addresses: hult@kth.se (H. Hult), pierre_nyquist@brown.edu (P. Nyquist).

http://dx.doi.org/10.1016/j.spa.2015.08.002
0304-4149/

need for fast and reliable computational methods. The aim of this paper is to introduce a new approach, based on the theory of large deviations for empirical measures, for analyzing the efficiency of stochastic simulation methods. In the present paper the emphasis is on importance sampling algorithms and the results can be viewed as a first step towards a new method for analyzing efficiency that can complement the existing variance-based approach.

The basic idea in stochastic simulation is to generate a population of particles that moves randomly according to the laws of the physical system. Each particle carries an individual weight, which may be updated during the simulation, and quantities of the underlying system are computed by averaging the particles' weights depending on their position. The canonical example is (standard) Monte Carlo simulation where all particles are independent and statistically identical, their weights being constant and equal.

Although the standard Monte Carlo method is widely used it is by no means universally applicable. One reason is that particles may wander off to irrelevant parts of the state space, leaving only a small fraction of relevant particles that contribute to the computational task at hand. Therefore, standard Monte Carlo may come with a computational cost that is too high for practical purposes. To remedy this, a control mechanism is needed that forces the particles to move to the relevant part of the state space, thereby reducing the computational cost. Such a control mechanism may come in different forms depending on the type of algorithm under consideration. In importance sampling, see e.g. [1], the control is the choice of sampling dynamics. In splitting algorithms [3] and genealogical particle systems [4,6] the control mechanism come, roughly speaking, in the form of a birth/death mechanism, which controls that important particles give birth to new particles and irrelevant particles are killed. For more on these and other related methods see [1,12] and the references therein.

Much of the theoretical work on the efficiency of stochastic simulation algorithms in general, and importance sampling algorithms in particular, is based on analyzing the variance of the resulting estimator. As long as an estimator is unbiased, variance is indeed the canonical measure of efficiency. However, when one considers situations in which the estimator is no longer unbiased, for example non-linear functionals of the underlying distribution, variance is no longer the obvious choice for determining performance and additional measures of efficiency can be useful.

Monte Carlo estimators are of the plug-in type where one first constructs an empirical measure associated with the distribution of interest and then plug this into whatever functional one is trying to estimate. A large deviation analysis for the underlying empirical measure will thus be useful in understanding the properties, such as rate of convergence, of the estimator. Therefore, much like for many other statistical estimators, a natural progression in the study of importance sampling algorithms and their performance is to consider the large deviation properties of the underlying weighted empirical measure. In the context of bootstrap estimators [13] derives such large deviation principles. Therein the sum of the weights is specified a priori and they are either deterministic or independent of the (possibly random) locations of the point masses in the empirical measure, making the setting quite different from that considered in this paper.

We propose in this paper that large deviation theory for the empirical measures arising in importance sampling can be used to define new performance criteria, based on the associated rate function, to complement the variance analysis. In addition to deriving the relevant large deviation result, we introduce one such criterion, which we refer to as evaluating performance over subsets. The intuition is that if a simulation algorithm is to approximate the original system in a certain region of the state space, then it should do so equally well over all subsets of that region that are not, in a sense, too small.