

Adaptive importance sampling of stochastic processes with graph-based mean hitting times

Guillaume Chennetier^{†,1,2}, Hassane Chraïbi¹, Anne Dutfoy¹, Josselin Garnier^{§,2}

[†] PhD student (presenting author). [§] PhD supervisor

PhD expected duration: Apr. 2021 – May. 2024

¹ EDF Lab Paris-Saclay, Boulevard Gaspard Monge, 91120 Palaiseau, France
`{guillaume.chennetier,hassane.chraibi,anne.dutfoy}@edf.fr`

² CMAP, Ecole Polytechnique, Institut Polytechnique de Paris, 91128 Palaiseau Cedex, France
`{guillaume.chennetier,josselin.garnier}@polytechnique.edu`

Abstract

Importance sampling is one of the most famous and perhaps the most direct method of variance reduction in simulation-based inference when estimating quantities of the form $\mathbb{E}_{X \sim \mathbf{p}} [\varphi(X)]$.

$$\bar{\varphi} := \mathbb{E}_{X \sim \mathbf{p}} [\varphi(X)] = \mathbb{E}_{X \sim \mathbf{g}} \left[\varphi(X) \frac{\mathbf{p}(X)}{\mathbf{g}(X)} \right] \approx \frac{1}{n} \sum_{k=1}^n \varphi(X_k) \frac{\mathbf{p}(X_k)}{\mathbf{g}(X_k)} \quad \text{with } X_1, \dots, X_n \sim \mathbf{g}. \quad (1)$$

Generating a sample X_1, \dots, X_n under an importance distribution \mathbf{g} as close as possible to $\mathbf{g}_{\text{opt}} \propto |\varphi| \times \mathbf{p}$ results in a low-variance estimator of $\bar{\varphi}$. A central use case is the rare event setting $\varphi(X) = \mathbb{1}_{X \in \mathbf{F}}$ (with $\bar{\varphi}$ very small) for which Monte Carlo methods are too expensive.

Simulating X according to an efficient importance distribution is notoriously challenging in high dimension. Various approaches have been proposed to address this problem (notably within the MASCOT-NUM community) such as projection into well-chosen subspaces [1, 2] or the use of generative models [3]. The dimension of the space is also an issue when X is not a vector but the trajectory of a stochastic process. This is a common case in reliability assessment where the operation of an industrial system is modeled by a stochastic process (see our previous work [4]).

We propose a new family of importance distributions tailored for dynamic rare event simulation with any non-diffusive Markov process, i.e any piecewise deterministic Markov process (PDMP). These processes evolve according to deterministic differential equations whose parameters are subject to random jumps. Their simulation cost prohibits the learning of a good importance distribution in a non-parametric framework, and the hybrid geometry of their state space prevents the application of most known dimension reduction methods.

The optimal importance distribution for PDMPs was characterized in [5] from the so-called “committor function” of the process. We present a methodology that combines the approximation of this committor function and a cross-entropy procedure to sequentially estimate $\bar{\varphi}$:

1. The PDMP is approximated by a time-homogeneous Markov random walk on a graph.
2. Mean hitting times of a target region of the graph are explicit for the simplified process, and form the basis of a family of approximations of the committor function.

