Adaptive importance sampling for reliability assessment of an industrial system modeled by a Piecewise Deterministic Markov Process

Guillaume Chennetier [1,2] · Anne Dutfoy [1] · Hassane Chraibi [1] · Josselin Garnier [2]

[1] EDF Lab Paris-Saclay, Boulevard Gaspard Monge, 91120 Palaiseau, France.
[2] CMAP, Ecole Polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France.

Industrial goal

Estimation of the probability of failure of systems involved in the operation of nuclear power plants and dams.

- A computer code simulates the real time operation of the system.
  PyCATSHOO $\longrightarrow$ **P**iecewise **D**eterministic **M**arkov **P**rocesses.
- Typical probabilities of failure are very small (about 10^{-5}).
- Each simulation is numerically expensive.

$\hookrightarrow$ Crude Monte-Carlo methods are not feasible.

Modeling with PDMP

Definition of a PDMP

Piecewise Deterministic Markov Process

(M.H.A Davis 1984)

Hybrid process: $Z_t = (X_t, M_t) \in E$

- position X_t is continuous
- mode M_t is discrete

1. **Flow** $\Phi \rightarrow$ deterministic dynamics between two jumps
2. **Jump intensity** $\lambda \rightarrow$ law of the time of the random jumps
3. **Jump kernel** $K \rightarrow$ law of the state of the process after a jump

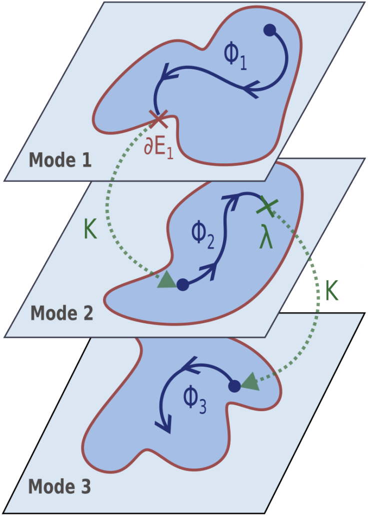

Monte-Carlo estimation

Objective: estimate $P = \mathbb{P}_{f_0}(\mathcal{Z} \in \mathcal{D})$

- f_0 nominal density[1] of a PDMP trajectory $\mathcal{Z}$ of fixed duration $t_{\max}$

$$(Z_t)_{t\in[0,t_{\max}]} =: \mathcal{Z} \sim f_0$$

- $\mathcal{D}$ subset of possible trajectories (in practice set of faulty trajectories)

Crude Monte-Carlo :

$$\widehat{P}_N^{\text{CMC}} = \frac{1}{N}\sum_{k=1}^{N} \mathbb{1}_{\mathcal{Z}_k \in \mathcal{D}} \quad \text{with } \mathcal{Z}_1, \ldots, \mathcal{Z}_N \overset{\text{i.i.d.}}{\sim} f_0 \tag{1}$$

$\hookrightarrow$ High relative variance of $\widehat{P}^{\text{CMC}}$ when P is small

[1] The density of a PDMP trajectory is mathematically sophisticated but analytically known and inexpensive to evaluate.

Importance sampling for PDMP

Importance sampling

Idea:

1. simulate PDMP trajectory $\mathcal{Z}$ according to an alternative distribution g which gives more weight on $\mathcal{D}$ than f_0
2. fix the bias with the likelihood ratio f_0/g

Importance sampling trick with alternative distribution g :

$$P = \mathbb{E}_{f_0}\left[\mathbb{1}_{\mathcal{Z}\in\mathcal{D}}\right] = \int \mathbb{1}_{\mathcal{Z}\in\mathcal{D}} \frac{f_0(\mathcal{Z})}{g(\mathcal{Z})} g(\mathcal{Z}) d\zeta(\mathcal{Z}) = \mathbb{E}_g\left[\mathbb{1}_{\mathcal{Z}\in\mathcal{D}} \frac{f_0(\mathcal{Z})}{g(\mathcal{Z})}\right] \quad (2)$$

IS estimator :

$$\widehat{P}_N^{\text{IS}} = \frac{1}{N}\sum_{k=1}^{N} \mathbb{1}_{\mathcal{Z}_k\in\mathcal{D}} \frac{f_0(\mathcal{Z}_k)}{g(\mathcal{Z}_k)} \quad \text{with } \mathcal{Z}_1, \dots, \mathcal{Z}_N \overset{\text{i.i.d.}}{\sim} g \quad (3)$$

$\hookrightarrow$ Variance of $\widehat{P}_N^{\text{IS}}$ relies on the choice of g

Optimal importance sampling

Optimal IS distribution produces IS estimator with zero variance.

General case:

$$g_{\text{opt}} : \mathcal{Z} \mapsto \frac{\mathbb{1}_{\mathcal{Z}\in\mathcal{D}}\, f_0(\mathcal{Z})}{P} \equiv f_0\left(\mathcal{Z} \,|\, \mathcal{Z} \in \mathcal{D}\right)$$

$\hookrightarrow$ untractable distribution.

PDMP case: (Thomas Galtier 2019)

g_{opt} fully determined by optimal jump intensity λ_{opt} and optimal jump kernel K_{opt} of the form:

$$\lambda_{\text{opt}} \equiv \lambda_0 \times \frac{U^-_{\text{opt}}}{U_{\text{opt}}} \quad \text{and} \quad K_{\text{opt}} \equiv K_0 \times \frac{U_{\text{opt}}}{U^-_{\text{opt}}} \qquad (4)$$

where

1. λ_0, K_0 are jump intensity and jump kernel of PDMP of distribution f_0
2. U_{opt} and U^-_{opt} are the so-called **committor functions** of the process

Committor functions

- U_{opt} probability of realizing the rare event $\{\mathcal{Z} \in \mathcal{D}\}$ knowing that at a fixed time $s > 0$ the process is in a given state z.

$$U_{\text{opt}}(z,s) = \mathbb{P}_{f_0}\left(\mathcal{Z} \in \mathcal{D} \,|\, Z_s = z\right), \tag{5}$$

- U^{-}_{opt} is the probability of realizing the rare event $\{\mathcal{Z} \in \mathcal{D}\}$ knowing that at a fixed time $s > 0$ the process jumps from a given state z^{-}.

$$U^{-}_{\text{opt}}(z^{-},s) = \text{"} \sum_{z \in E} U_{\text{opt}}(z,s)\, K\left(z^{-}, z\right) \text{"}. \tag{6}$$

Knowing U_{opt} is sufficient to build the optimal IS estimator.

Adaptive algorithm

Our contribution

Recently submitted article: *Adaptive importance sampling based on fault tree analysis for piecewise deterministic Markov process.*

1. Fault tree analysis methods are used to construct a family of approximations of the committor function U_{opt}.
2. The best representative of this family is sequentially determined using a cross-entropy procedure coupled with a recycling scheme for past samples.
3. A consistent and asymptotically normal post-processing estimator of the final probability P is returned.

Performances on an industrial case

We tested this importance sampling approach on a complex case from nuclear industry and we compared it to a massive crude Monte-Carlo method.

Method	N	$\widehat{P}$	$\widehat{\sigma}/\widehat{P}$	95% confidence interval
CMC	10^5	2×10^{-5}	223.60	$[0 ; 4.77 \times 10^{-5}]$
	10^6	1.3×10^{-5}	277.35	$[5.93 \times 10^{-6} ; 2.01 \times 10^{-5}]$
	10^7	1.77×10^{-5}	237.68	$[1.51 \times 10^{-5} ; 2.03 \times 10^{-5}]$
AIS	10^2	2.18×10^{-5}	4.69	$[1.76 \times 10^{-5} ; 4.18 \times 10^{-5}]$
	10^3	2.19×10^{-5}	3.01	$[1.78 \times 10^{-5} ; 2.60 \times 10^{-5}]$
	10^4	1.99×10^{-5}	1.01	$[1.96 \times 10^{-5} ; 2.03 \times 10^{-5}]$

Table 1: **Comparison between crude Monte-Carlo (CMC) and our adaptive importance sampling method (AIS).**

$\hookrightarrow$ Variance reduction by a factor of 10,000.

Stability of the method

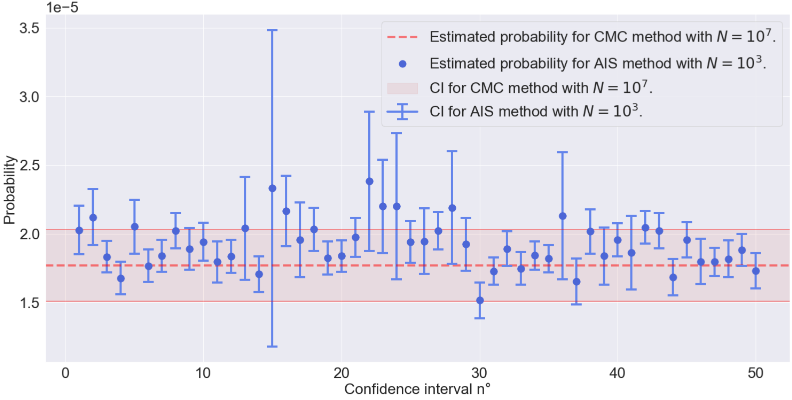

Figure 1: **50 confidence intervals with AIS method and sample size of 1000 vs 1 confidence interval with CMC method and sample size of 10^7.**

(New?) Bandit problem

Best arm identification

Context:

- Several nominal densities $f_1, \ldots, f_d$
- $P_i := \mathbb{P}_{f_i}(\mathcal{Z} \in \mathcal{D}) = \mathbb{E}_{f_i}[\mathbb{1}_{\mathcal{Z}\in\mathcal{D}}]$ for $i = 1, \ldots, d$.

Objective: find the most reliable distribution

$$\underset{i\in\{1,\ldots,d\}}{\arg\min}\ \mathbb{E}_{f_i}[\mathbb{1}_{\mathcal{Z}\in\mathcal{D}}] \tag{7}$$

Best arm identification (BAI) framework:

- *sampling rule*: at iteration k, draw $\mathcal{Z}_k \sim f_{i_k}$ with $i_k \in \{1, \ldots, d\}$
- *stopping rule*: fixed-budget setting (stop when $k = k_{\max}$), fixed-confidence setting (stop when the error probability is small enough), etc.
- *recommendation rule*: which distribution to bet on at the end.

Off-policy best arm identification

Difference between our case and standard BAI:

- the P_i are very small $\rightarrow$ we do not draw from "true arms" $\{f_1, \ldots, f_d\}$
- we generate $\mathcal{Z}$ from $\mathcal{G}$ a family of alternative distributions (IS = "off-policy" method)
- each draw gives information on every P_i thanks to reverse IS

Existing contributions:

- "Optimal" algorithms for standard BAI
- "Off-policy" methods (with IS) for multi-armed bandit with a regret minimization objective (not BAI)

What about optimal off-policy best arm identification?

References

H. Chraibi, A. Dutfoy, T. Galtier et J. Garnier (2019), On the optimal importance process for piecewise deterministic Markov process. *ESAIM: PS. 23:893–921*.

M. H. A. Davis (1984), Piecewise-Deterministic Markov Processes: A general class of non-diffusion stochastic models. *J. R. Stat. Soc. B 46:353-388*.

B. Delyon and F. Portier (2018), Asymptotic optimality of adaptive importance sampling. *NIPS 2018*.

B. de Saporta, F. Dufour and H. Zhang (2015), Numerical methods for simulation and optimization of piecewise deterministic Markov processes : application to reliability. *Mathematics and statistics series Wiley-ISTE*.

R. Rubinstein (1999), The Cross-Entropy method for combinatorial and continuous optimization. *Methodology and Computing in Applied Probability 1:127–190*.

E. Kaufmann, O. Cappé, and A. Garivier (2016), On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research 17.1 (2016): 1-42*.

Thank you for your attention

Questions?

Supplementary material

PDMP material

Characterization of a PDMP

The **flow** Φ, solution of differential equations, gives the deterministic dynamic. If there is no jump between time s and time $s+t$ then:

$$Z_{s+t} = \Phi_{Z_s}(t). \tag{8}$$

The deterministic jumps occur when the process reaches the **boundaries of the state space** E.

$$t_z^{\partial} = \inf\{t > 0 : \Phi_z(t) \in \partial E\}. \tag{9}$$

The **jump intensity** λ gives the distribution of the time T_z of the next random jump knowing current state z.

$$\mathbb{P}(T_z > t \mid Z_s = z) = \mathbb{1}_{t<t_z^{\partial}} \exp\left(-\int_0^t \lambda\left(\Phi_z(u)\right) du \right). \tag{10}$$

The **jump kernel** $\mathcal{K}$ gives the law of the post-jump location. Jumping from z^-, the arrival state z is randomly chosen by the jump kernel $\mathcal{K}_{z^-}$ of probability density function $z \mapsto K(z^-, z)$ with respect to some measure ν_{z^-}.

Likelihood of a PDMP trajectory

Probability density function of a PDMP trajectory (*Thomas Galtier 2019*)

There is a dominant measure ζ for which a PDMP trajectory $\mathcal{Z}$ with $n_{\mathcal{Z}}$ jumps, inter-jump times $t_1, \ldots, t_{n_{\mathcal{Z}}}$ and arrival states $z_1, \ldots, z_{n_{\mathcal{Z}}}$ admits a probability density function π.

$$\pi\left(\mathcal{Z}\right) = \prod_{k=0}^{n_{\mathcal{Z}}} \left[\lambda\left(\Phi_{z_k}\left(t_k\right)\right)\right]^{\mathbb{1}_{t_k < t^{\partial}_{z_k}}} \exp\left[-\int_0^{t_k} \lambda\left(\Phi_{z_k}(u)\right) \mathrm{d}u\right] \prod_{k=0}^{n_{\mathcal{Z}}-1} K\left(\Phi_{z_k}(t_k), z_{k+1}\right). \quad (11)$$

Take home message:

- explicit computation of the pdf of a PDMP trajectory,
- no need to recalculate the flow.

Test case: the spent fuel pool

If the system does not cool the pool, the nuclear fuel evaporates the water then damages the structure and contaminates the outside.

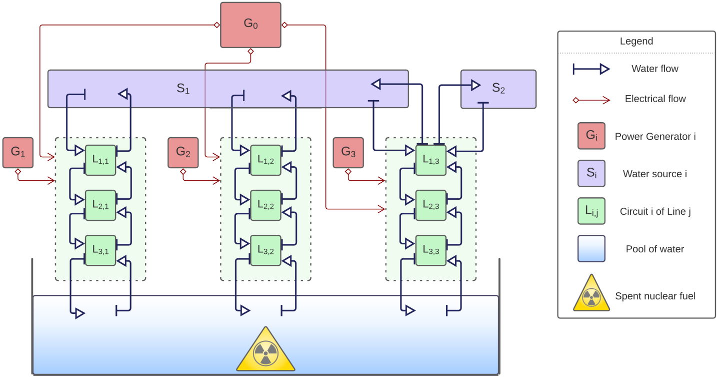

Aim: estimating the probability of the water level falling below a set threshold.

Approximation with MPS

Approximation of the committor function with minimal path sets

The path sets of a system are the sets of components such that:

1. keeping all components of any path set intact prevents system failure.
2. keeping one component broken in each path set ensures system failure.

A **M**inimal **P**ath **S**et is a path set that does not contain any other path set.

We note:

- d_{MPS} the number of MPS (they are unique if the system is coherent),
- $\beta^{(\text{MPS})}(z)$ the number of MPS with at least one broken component.

A good U_α should therefore be increasing in $\beta^{(\text{MPS})}(z)$.

Minimal path sets: the spent fuel pool case

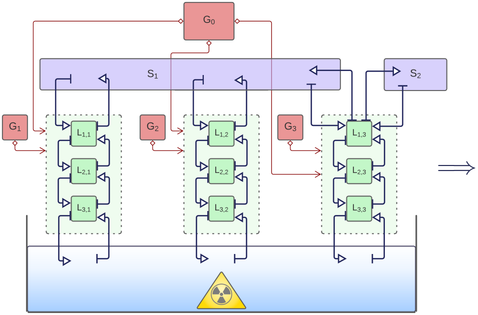

Figure 2: Physical representation of the SFP

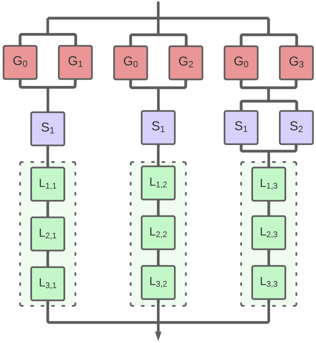

Figure 3: Functionnal diagram of the SFP

8 MPS in the spent fuel pool system: (with $\boldsymbol{L_j} = (L_{i,j})_{i=1}^{3}$ for $j = 1, 2, 3$)

$$(G_0, S_1, \boldsymbol{L_1}), (G_1, S_1, \boldsymbol{L_1}), (G_0, S_1, \boldsymbol{L_2}), (G_2, S_1, \boldsymbol{L_2}),$$

$$(G_0, S_1, \boldsymbol{L_3}), (G_3, S_1, \boldsymbol{L_3}), (G_0, S_2, \boldsymbol{L_3}), (G_3, S_2, \boldsymbol{L_3}).$$

Our MPS-based proposition

For $\boldsymbol{\alpha} \in \mathbb{R}_+^{d_{\mathbf{MPS}}}$ we propose:

$$U_{\boldsymbol{\alpha}}^{(\mathrm{MPS})}(z) = \exp\left[\left(\sum_{i=1}^{\beta^{(\mathbf{MPS})}(z)} \alpha_i\right)^2\right]. \tag{12}$$

Flexible dimension of α: imposing equality on some coordinates of α reduce its effective dimension and simplify the search for a good α when d_{MPS} is large.

→ Example for dimension 1 with $\alpha_1 = \cdots = \alpha_{d_{\mathbf{MPS}}}$:

$$U_{\boldsymbol{\alpha}}^{(\mathrm{MPS})}(z) = \exp\left[\left(\alpha_1\, \beta^{(\mathrm{MPS})}(z)\right)^2\right]. \tag{13}$$

The form $x \mapsto \exp(x^2)$ garantees that the ratios $U_{\boldsymbol{\alpha}}^-/U_{\boldsymbol{\alpha}}$ are strictly increasing in $\beta^{(\mathrm{MPS})}$. Without this condition, it is increasingly difficult to break new components and they are repaired faster and faster as they are lost.

Minimal cut sets

Minimal cut sets: smallest sets of components that if left broken ensure system failure. (permanent repair of one component in each group prevents the failure)

In this system: there is 69 minimal cut sets for 15 components.

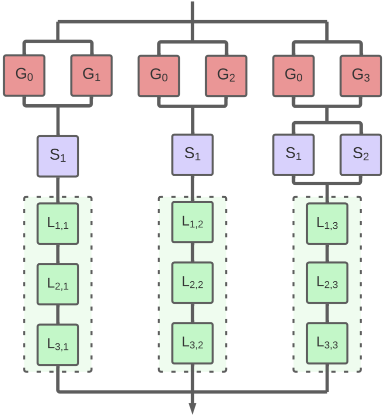

Figure 4: Functionnal diagram of the SFP

Examples: $(G_0, G_1, G_2, G_3), (S_1, S_2), (C_1L_1, C_3L_2, C_1L_3), (G_0, G_3, S_1), \ldots$

Recycling adaptive IS

The cross entropy procedure

How to find the best candidate within the family $(U_{\alpha})_{\alpha \in \mathbb{A}}$?

To each $\alpha \in \mathbb{A} \subset \mathbb{R}^{d_\alpha}$ corresponds an approximation U_α and an associated importance distribution g_α. We look for the closest distribution g_α to f_{opt} in the sense of the Kullback-Leibler divergence.

$$
\begin{aligned}
\underset{\alpha \in \mathbb{A}}{\arg\min}\, \mathcal{D}_{\text{KL}}\left(f_{\text{opt}} \| g_\alpha\right) &= \underset{\alpha \in \mathbb{A}}{\arg\min}\, \mathbb{E}_{f_{\text{opt}}}\left[\log\left(\frac{f_{\text{opt}}(\mathcal{Z})}{g_\alpha(\mathcal{Z})}\right)\right] \\
&= \underset{\alpha \in \mathbb{A}}{\arg\min} \int -\log\left(g_\alpha(\mathcal{Z})\right) \frac{\mathbb{1}_{\mathcal{Z} \in \mathcal{D}}\, f_0(\mathcal{Z})}{P} d\mathcal{Z} \\
&= \underset{\alpha \in \mathbb{A}}{\arg\min} \left\{-\mathbb{E}_{f_0}\left[\mathbb{1}_{\mathcal{Z} \in \mathcal{D}} \log\left(g_\alpha\left(\mathcal{Z}\right)\right)\right]\right\}
\end{aligned}
$$

This last quantity does not depend on f_{opt}, it can be minimized iteratively by successive Monte-Carlo approximations with importance sampling.

Adaptive algorithm with recycling of past samples

Start with an initial parameter $\boldsymbol{\alpha}^{(1)}$. At iteration $q = 1, \ldots, Q$:

1. **Simulation step**: generate a new sample of n_q trajectories

$$\mathcal{Z}_1^{(q)}, \ldots, \mathcal{Z}_{n_q}^{(q)} \overset{\text{i.i.d.}}{\sim} g_{\boldsymbol{\alpha}^{(q)}}$$

2. **Optimization step**: compute the next iterate $\boldsymbol{\alpha}^{(q+1)}$ by solving (14):

$$\boldsymbol{\alpha}^{(q+1)} = \underset{\boldsymbol{\alpha} \in \mathbb{A}}{\arg\min} \left\{ - \sum_{r=1}^{q} \sum_{k=1}^{n_r} \mathbb{1}_{\mathcal{Z}_k^{(r)} \in \mathcal{D}} \frac{f_0\left(\mathcal{Z}_k^{(r)}\right)}{g_{\boldsymbol{\alpha}^{(r)}}\left(\mathcal{Z}_k^{(r)}\right)} \log \left[g_{\boldsymbol{\alpha}}\left(\mathcal{Z}_k^{(r)}\right) \right] \right\} \quad (14)$$

Estimation step: at iteration Q, the final estimator of the probability P is:

$$\widehat{P} = \frac{1}{\sum_{q=1}^{Q} n_q} \sum_{q=1}^{Q} \sum_{k=1}^{n_q} \mathbb{1}_{\mathcal{Z}_k^{(q)} \in \mathcal{D}} \frac{f_0\left(\mathcal{Z}_k^{(q)}\right)}{g_{\boldsymbol{\alpha}^{(q)}}\left(\mathcal{Z}_k^{(q)}\right)} \quad (15)$$

Past samples are reused at each optimization step and at estimation step.

We proved the consistency and asymptotic normality of the estimator (15).

Cross-Entropy: initialization and optimization routine

Initialization: finding $\boldsymbol{\alpha}^{(0)}$ to start CE

1. Fix $\widetilde{p} \in [0,1]$ and $\widetilde{t} > 0$ (example $\widetilde{p} = 0.95$ and $\widetilde{t} = t_{\max}$).
2. Find the smallest $\alpha \in \mathbb{R}_+$ such that the probability that the time of the first failure occurs before $\widetilde{t}$ is greater than $\widetilde{p}$.

$$\widetilde{\alpha} = \inf \left\{ \alpha \in \mathbb{R}_+ \ : \ \mathbb{P}_{g_{\boldsymbol{\alpha}}} \left(T \leq \widetilde{t} \mid Z = z_0 \right) \geq \widetilde{p} \right\}. \qquad (16)$$

3. Start CE with $\boldsymbol{\alpha}^{(0)} = (\widetilde{\alpha}, \dots, \widetilde{\alpha})$.

Optimization routine

Since the gradient of $\boldsymbol{\alpha} \mapsto g_{\boldsymbol{\alpha}}$ is known, we have an explicit gradient for the objective function of the CE minimization.

We used the `BFGS` method from the Python library `scipy.optimize`.

The norm of the gradient in the stopping criterion must be very small (in our case at most 10^{-30}) because the probability density functions are themselves very small.

Asymptotic optimality

Theorem

If $\mathbb{A}$ is compact and if moreover:

1. *The functions λ, K, and $(U_\alpha)_{\alpha\in\mathbb{A}}$ are bounded on their support below and above by strictly positive constants,*
2. *$\alpha_{opt} \in \mathbb{A}$ is the unique minimizer of $\left(\alpha \mapsto -\mathbb{E}_{f_0}\left[\mathbb{1}_{\mathcal{Z}\in D} \log\left(g_\alpha\left(\mathcal{Z}\right)\right)\right]\right)$,*
3. *there is $t_\varepsilon > 0$ such that $t_z^\partial \geq t_\varepsilon$ for any $z^- \in \partial E$ and any $z \in \text{supp } K\left(z^-, \cdot\right)$,*

then, with $V(\alpha) = \mathbb{E}_{f_0}\left[\mathbb{1}_{\mathcal{Z}\in D}\frac{f_0(\mathcal{Z})}{g_\alpha(\mathcal{Z})}\right] - P^2$ we have :

$$\alpha^{(Q)} \xrightarrow[Q\to\infty]{a.s} \alpha_{opt} \quad \text{and} \quad \sqrt{N_Q}\left(\widehat{P}_{N_Q} - P\right) \xrightarrow[Q\to\infty]{\mathcal{L}} \mathcal{N}\left(0, V\left(\alpha_{opt}\right)\right).$$

Asymptotic confidence interval of level $1-a$ for P:

$$\mathbb{P}\left(P \in \left[\widehat{P}_{N_Q} - q_{1-\alpha/2}\sqrt{\frac{\widehat{\sigma}^2_{N_Q}}{N_Q}}\,;\ \widehat{P}_{N_Q} + q_{1-\alpha/2}\sqrt{\frac{\widehat{\sigma}^2_{N_Q}}{N_Q}}\right]\right) \xrightarrow[Q\to\infty]{} 1-a.$$

Sensitivity analysis

Assumptions:

- The distribution of the PDMP depends on a parameter vector $\boldsymbol{\theta}$.
- We estimated the probability of failure by importance sampling from a sample of N trajectories of distribution g.

We would like to measure the sensitivity of the probability of critical failure to these parameters without generating new trajectories.

Reverse importance sampling trick:

$$\mathbb{E}_{f_{\boldsymbol{\theta}}}\left[\mathbb{1}_{\mathcal{Z}\in \boldsymbol{D}}\right] = \int \mathbb{1}_{\mathcal{Z}\in \boldsymbol{D}} \frac{f_{\boldsymbol{\theta}}(\mathcal{Z})}{g(\mathcal{Z})} g(\mathcal{Z}) d\zeta(\mathcal{Z}) = \mathbb{E}_{g}\left[\mathbb{1}_{\mathcal{Z}\in \boldsymbol{D}} \frac{f_{\boldsymbol{\theta}}(\mathcal{Z})}{g(\mathcal{Z})}\right]. \qquad (17)$$

Input/output dataset $\left(\boldsymbol{\theta}^{(i)}, \widehat{P}_{f_{\boldsymbol{\theta}^{(i)}}}\right)_{i=1,\dots,n}$ with:

$$\widehat{P}_{f_{\boldsymbol{\theta}^{(i)}}} = \frac{1}{N} \sum_{k=1}^{N} \mathbb{1}_{\mathcal{Z}_k\in \boldsymbol{D}} \frac{f_{\boldsymbol{\theta}^{(i)}}(\mathcal{Z}_k)}{g(\mathcal{Z}_k)} \quad \text{with} \quad Z_k \sim g. \qquad (18)$$

Post-processing sensitivity indices from this dataset.