



**HAL**  
open science

## Adaptive Reduced Multilevel Splitting

Frédéric Cérou, Patrick Héas, Mathias Rousset

► **To cite this version:**

Frédéric Cérou, Patrick Héas, Mathias Rousset. Adaptive Reduced Multilevel Splitting. 2023. hal-04364553

**HAL Id: hal-04364553**

**<https://inria.hal.science/hal-04364553v1>**

Preprint submitted on 13 Jan 2025

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# ADAPTIVE REDUCED MULTILEVEL SPLITTING

FRÉDÉRIC CÉROU\*, PATRICK HÉAS\*, MATHIAS ROUSSET\*

## Abstract.

This paper considers the classical problem of sampling with Monte Carlo methods a target rare event distribution defined by a score function that is very expensive to compute. We assume we can build using evaluations of the true score, an approximate surrogate score certified with error bounds. This work proposes a fully adaptive algorithm to sequentially sample surrogate rare event distributions with increasing target levels. An essential contribution consists in sampling at each iteration the surrogate rare event at a critical level corresponding to a specific cost. This cost is related to importance sampling for a target for a given budget. The critical level is calculated solely from the reduced score and its error bound. From a practical point of view, sampling the proposal sequence is performed by extending the framework of the popular adaptive multilevel splitting algorithm to the use of score approximations. Numerical experiments evaluate the proposed importance sampling algorithm in terms of computational complexity versus squared error. In particular, we investigate the performance of the algorithm when simulating rare events related to the solution of a parametric PDE, which is approximated by a reduced basis.

**Key words.** Rare event simulation, reduced modeling, importance sampling, adaptive multilevel splitting, relative entropy.

## 1. Introduction.

**1.1. Context and objective.** Let  $\pi(dx)$  denotes a given, easily simulable, reference probability distribution on a state space  $\mathcal{X}$ . A random variable with distribution  $\pi$  may describe for instance unknown physical parameters, uncertainty on initial/boundary conditions, or additional noise modeling either a chaotic behavior or the interaction with an environment. We also consider a real-valued *score* function  $S^* : \mathcal{X} \rightarrow \mathbb{R}$  which is assumed to be of the form

$$S^*(x) = \text{score}(\Psi^*(x)),$$

where  $\text{score}$  is an easily computable function of physical interest, and  $\Psi^*(x)$  is the outcome of a presumably exact numerical computation of a complex physical system parametrized by  $x \in \mathcal{X}$ . With the previous ingredients, one can define the following family of conditional probability distributions on  $\mathcal{X}$ :

$$\eta_l^*(dx) \stackrel{\text{def}}{=} \frac{1}{p_l^*} \mathbf{1}_{S^*(x) > l} \pi(dx) \quad l \in \mathbb{R} \quad (1.1)$$

defined by conditioning the distribution  $\pi$  on the event  $\{S^* > l\}$  which may be interpreted as a certain failure region/event. The quantity  $p_l^* \stackrel{\text{def}}{=} \int_{\mathcal{X}} \mathbf{1}_{S^* > l} d\pi = \pi(\{S^* > l\})$  is the associated normalization, here the probability of the considered (possibly rare) event  $\{S^* > l\}$ , for  $l \in \mathbb{R}$ .  $\eta_l^*$  is the probability distribution of the system conditioned to the failure event.

We consider in this paper the classical problem of designing an algorithmic Monte Carlo procedure that samples according to  $\eta_{l_{\max}}^*$  – for one specific large  $l_{\max} \in \mathbb{R}$  – and estimate the associated rare event probability  $p_{l_{\max}}^*$ . Note that depending on the problem at hand, the final event of interest might be either  $\{S^* > l_{\max}\}$  or  $\{S^* \geq l_{\max}\}$ , the intermediate sets still being defined as  $\{S^* > l\}$  for  $l < l_{\max}$ .

\*Inria and Irmar, University of Rennes, France.  
(frederic.cerou@inria.fr,patrick.heas@inria.fr,mathias.rousset@inria.fr)

In this work, we will focus on scenarios with *two* types of difficulty. A first difficulty arises when  $l_{\max}$  is large, generating a very small (yet strictly positive) probability  $p_{l_{\max}}^* > 0$ . This rare event context confronts the practitioner with a first computational challenge: the majority of the probability mass of the conditional distribution  $\eta_{l_{\max}}^*$  is concentrated in specific areas of  $\mathcal{X}$ , which are roughly described by the maxima of  $S^*$  which are usually unknown. A direct sampling with a Monte Carlo simulation of  $\pi$  requires a sample size of order at least  $1/p_{l_{\max}}^*$ , which is often infeasible. In that context, a popular efficient strategy to simulate a sample according to  $\eta_{l_{\max}}^*$  and to estimate the associated normalizing constant is to resort to 'Sequential Monte Carlo' (SMC) methods. Such methods start with a Monte Carlo sample of clones (also called below 'particles') of the system distributed according to  $\pi$  and then sequentially samples  $\eta_l^*$  for increasing values of  $l$ . To do so, those methods resort to a combination of Importance Sampling (IS), re-sampling (selection) of weighted samples, and mutations of particles based on suitable Markov Chain Monte Carlo transitions. In this work we will focus on a specific instance called *Adaptive Multilevel Splitting* (AMS, see [5, 6]) which moreover chooses adaptively the sequence of levels  $l$  where re-sampling is performed.

A second difficulty appears when the numerical evaluation of  $\Psi^*$  is extremely expensive, so that the number of evaluations required in, say, a usual SMC or AMS approach becomes prohibitive. In order to circumvent this issue, we assume that the physical model  $\Psi^*$  can be approximated by a *reduced model* denoted  $\Psi$ , the cost of evaluating  $\Psi$  being small as compared to one evaluation of the true score with  $\Psi^*$ . The reduced model  $\Psi$  we consider is assumed to come with two key features. To explain these, let us denote from now on the approximate score function by

$$S(x) = \text{score}(\Psi(x)) \in \mathbb{R}.$$

The first feature is an *a posteriori* pointwise error quantification associated with  $S$ . This error quantification comes in the form of an explicit error bound denoted  $E(x) = \text{error}(\Psi(x)) \in \mathbb{R}$ , and satisfying the pointwise estimate:

$$|S(x) - S^*(x)| \leq E(x), \quad \forall x \in \mathcal{X}. \quad (1.2)$$

The second key feature assumes that we are able to update the reduced model, through a procedure of the following form

$$\{(X_1, \Psi^*(X_1)), \dots, (X_k, \Psi^*(X_k))\} \xrightarrow{\text{reduced model}} \Psi = \Psi^{(X_1, \dots, X_k)},$$

which takes as an input any sequence of states in  $\mathcal{X}$  – called here *a sample of snapshots* – together with the solution of the associated true model  $\Psi^*$  evaluated for each of those snapshots. The output is the reduced model  $\Psi$ . To improve the readability of the document, we will use a superscript  $(k)$  to specify the dependency on  $(X_1, \dots, X_k)$ , or even omit this dependency if the context is unambiguous,

$$(S^{(X_1, \dots, X_k)}, E^{(X_1, \dots, X_k)}) = (S^{(k)}, E^{(k)}) = (S, E).$$

## 1.2. Previous works.

*Reduced models for rare events.* The idea of using a simpler surrogate model in Monte-Carlo simulation for target distributions is not new and there is a large amount of literature on the subject. Most references are concerned with either kriging on the one hand, and reduced basis methods on the other hand. For kriging approaches

[12, 20, 23, 27, 2, 18, 3], the surrogate score function  $S : \mathcal{X} \rightarrow \mathbb{R}$  is a *random* field, distributed according to the posterior of a prior Gaussian field conditioned by fitting the exact values of the evaluated snapshots. The estimate (1.2) is however not satisfied almost surely, and one has to resort to averaged error estimates. For *reduced basis* approaches [9, 17, 14, 16], the surrogate score function is deterministic, and constructed by projection of the function  $\Psi^*$  on a subspace where it is efficiently approximated and characterized with precise error bounds of the form (1.2). This model reduction framework is known to be very efficient for the numerical approximation of problems involving the repeated solution of parametric *Partial Differential Equations* (PDEs) [22]. We note that besides these two main approaches, other families of reduced models [4, 13, 26, 21] or combination of them [13, 26, 21] have been proposed for rare event simulation.

Arguably, the most sensitive feature of any of these reduced model approaches is the specific collection of snapshots  $(X_k, \Psi^*(X_k))$ ,  $k \geq 1$  with which the reduced model is updated along the iterations. A natural approach consists in sampling sequentially and adaptively the collection of snapshots, the reduced model being updated after each new evaluation of the true model at iteration  $k$ . The updated reduced model (and the associated error estimates) can then be used to pick the next snapshot in presumably *interesting* areas. Sampling the snapshot  $X_k$  rely on two ingredients: i) a probability distribution used to generate candidates, and ii) a weighting function, also called *learning function*, used to pick the presumably best candidate.

The idea of ingredient ii) can be found in the seminal work of [12] where selecting snapshots is called *active learning*. In the latter approach, snapshots are selected by optimizing a criterion (the learning function) that will most likely reduce the uncertainty of the surrogate model. An exhaustive set of references and a very quantitative review of methods related to active learning and kriging is provided in [20]. Some other references, as [17], pick the snapshots  $(X_k)_{k \geq 1}$  uniformly among candidates using a decreasing sequence of subsets that contain the target rare event, constructed using the iteratively updated error estimates.

However, as we will argue in the present work, ingredient i) is also of high importance. Indeed, a key general phenomenon in a rare-event and high-dimensional setting is the following: the true target distribution  $\eta_l^*$  for  $l$  large is extremely concentrated around very specific but unknown points in  $\mathcal{X}$ , so that, unless the pool of candidates among which snapshots  $(X_1, \dots, X_k)$  are picked is partly also located in such regions, the information given by the reduced model may eventually lack the accuracy required in order to describe the rare event. As a consequence, unless one is able to sample an extremely large amount of candidates, it is critical to select such candidates in areas that carefully track the true rare event. To make the picture clearer, consider at iteration  $k$  the following surrogates of the target conditional distributions (1.1) built using the reduced score:

$$\mu_l \stackrel{\text{def}}{=} \frac{1}{Z_l} \mathbf{1}_{S>l} d\pi \quad l \in \mathbb{R}, \quad (1.3)$$

where the normalization  $Z_l \stackrel{\text{def}}{=} \int_{\mathcal{X}} \mathbf{1}_{S>l} d\pi$ , represents an approximation of the probability of the rare event. Then, for large  $l$ , the reduced rare event is also rare, hence very concentrated, and thus usually too far away from the true target rare event. If we want to guarantee sufficient mass in the region of the target rare event, we propose to specify a *maximum possible level*  $l$  below which the true conditional distribution

$\eta_l^*$  and the reduced one  $\mu_l$  do match  $\eta_l^* \simeq \mu_l$ . This delicate issue is only rarely addressed in the literature. Most references avoid it and generate pools of candidates for the choice of snapshot relying on the reference distribution  $\pi$ . To the best of our knowledge, the only related pieces of work deal with the reverse problem of selecting (in a hierarchy of pre-computed surrogates) [16] or refining [3] the reduced score  $S$ , so that it is appropriate for a certain level  $l$  according to some error quantification.

The main contribution of the present work is to use the following guidelines to design an efficient and adaptive algorithm that synchronizes: i) the sequential Monte Carlo sampling of rare events using a reduced model (that is sampling the flow  $l \mapsto \mu_l$ ), and ii) the sequential update of the reduced model with evaluations of the true model.

*Reduced Monte-Carlo simulation.* Strategies for Monte-Carlo sampling of the reduced conditional distributions  $\mu_l$  can be grouped in two broad categories: on the one hand *importance sampling* (IS), usually through a parametric approach, and on the other hand splitting methods, *i.e.*, *sequential Monte Carlo* [11] methods, which are non-parametric, and in particular AMS [5] also called *subset simulation* [1].

Many approaches to IS of the reduced distribution  $\mu_l$  have been proposed in recent years. In particular, *cross-entropy methods* [10] have been proposed relying on kriging [2] or reduced basis approximations [14, 16]. In order to increase the effectiveness of the cross-entropy method, it has also been suggested to approximate the parameter variable  $X$  by constraining its support in well-chosen low dimensional subspaces of  $\mathcal{X}$  [25]. Other recent contributions aim to build specific adaptive proposals for IS [23, 27] or for stratified sampling [4] (that can be interpreted as a form of IS). However, the effectiveness of IS approaches depends heavily on the choice of parametric family of proposals, for which we don't necessarily have any prior knowledge.

On the other hand, splitting methods have also been used in conjunction with (adaptively constructed) kriging-based surrogate models [18, 3]. We refer again to the review on methods related to active learning and kriging provided in [20]. However, in these works, the splitting method samples an approximation  $\mu_l$  of the target probability (1.1), which may be biased. We note nevertheless that in [13] subset simulation is performed using nested multiresolution sets (provided by an accuracy hierarchy of the surrogate model), yielding an estimate with an associated *a posteriori* error.

Adapting splitting methods to reduced scores can be optionally supplemented with a *bridging* procedure. The goal of this bridge is to sample a new proposal  $\mu_{l'}$  defined with some updated reduced score  $S'$  and level  $l'$ , starting from the samples of the current proposal  $\mu_l$ . In general, this sampling is not trivial. For example, the nesting underlying a AMS simulation is broken, since in general updating the reduced score and the level will imply  $\{S' > l'\} \not\subset \{S > l\}$ . There is very little work in the literature on this bridging problem. Let us mention the work [26], which proposes a bridging procedure with a constant  $l' = l$  level. Specifically, the authors use the sequential Monte Carlo method with soft levels and tempering instead of targeting distributions defined by conditioning on surrogate rare events. The authors mention that such an approach outperforms the multilevel subset simulation [24] algorithm they previously proposed, which dealt with the nesting problem by computing correction terms and performing Markov chain Monte Carlo simulations.

**1.3. Contributions.** In this work, we propose a sequential algorithm relying on reduced scores and error bounds. The algorithm provides an answer to the questions and limitations listed in Section 1.2. We consider the family of proposal distributions (1.3) (built using the reduced score) for sampling the target conditional distributions (1.1). The sketch of the proposed algorithm is the following. At each

iteration  $k$ , the algorithm

- computes adaptively a maximum level denoted  $l^{(k)}$  such that the proposal  $\mu_{l^{(k)}}^{(k)}$  matches the true conditional distribution  $\eta_{l^{(k)}}^*$ , as enabled by the accuracy of the current reduced model  $S^{(k)}$
- simulates (with a population Monte Carlo procedure) the proposal  $\mu_l^{(k)}$  from an initial level  $l = l_b^{(k)}$  to the maximum level  $l = l^{(k)}$ ;
- samples the new snapshot  $(X^{(k)}, \Psi^*(X^{(k)}))$  according to a weighted version of the distribution  $\mu_{l^{(k)}}^{(k)}$ . The weights (or *learning function*) may favor states  $x$  with large reduced score error  $E^{(k)}(x)$  (see (1.2)).
- updates the reduced model score  $S^{(k)} \rightarrow S^{(k+1)}$ .
- computes (using records of previous iterations) the next initial level  $l_b^{(k+1)}$  so that the nesting (inclusion) of the support of the considered proposals are guaranteed.
- bridges with (population) Monte Carlo simulation records of previous iterations to the next proposal  $\mu_{l_b^{(k+1)}}^{(k+1)}$ .

The Monte Carlo aspect of the algorithm is implemented using an AMS simulation methodology, in which  $\mu_l^{(k)}$  is approximated by the empirical distribution of a sample (population) of particles (clones).

Our *first main methodological contribution* is the specific criterion used to adaptively compute the maximum level of iteration  $k$  denoted  $l^{(k)}$ . At each iteration  $k$ , we will use the information given by the error quantification (1.2) so that  $\mu_{l^{(k)}}^{(k)}$  remains a reasonable importance distribution for sampling the true target  $\eta_{l^{(k)}}^*$ . The evaluation of the accuracy of an importance distribution will be done using a quantification of the logarithmic cost of importance sampling by relative entropy:

$$\ln \text{cost} \simeq \text{Ent}(\eta_{l^{(k)}} | \mu_{l^{(k)}}^{(k)}) \stackrel{\text{def}}{=} \eta_{l^{(k)}} \left( \ln \frac{\eta_{l^{(k)}}}{\mu_{l^{(k)}}^{(k)}} \right).$$

This choice is discussed and justified in Appendix A. The latter cannot be evaluated exactly, but we can leverage the existence of error estimates in order to obtain (more or less pessimistic) approximations of it. Relying on this cost, we are able to stop the increase of rareness of the targets when the reduced modeling becomes critically inaccurate, the associated maximum level being  $l^{(k)}$ .

Snapshot sampling is then performed using the distribution  $\mu_{l^{(k)}}^{(k)}$ , which we may reweight by favoring states with higher error quantification. Remark that snapshot sampling is performed at the precise level  $l^{(k)}$  of the flow of reduced proposal distributions  $(\mu_l^{(k)})_{l \in \mathbb{R}}$  where increasing the quality of these proposals becomes necessary: this prevents a poor sampling of a new snapshot inconsistent with the unknown target while keeping the associated level  $l^{(k)}$  as large as possible.

The proposed method can also yield importance sampling estimators. If for instance the snapshots are sampled from the start without re-weighting according to proposal  $X^{k+1} \sim \mu_{l^{(k)}}^{(k)}$ , it naturally leads to the following estimator:

$$\hat{\pi}^{(k)} \stackrel{\text{def}}{=} \frac{1}{k} \sum_{k'=1}^k \frac{d\pi}{d\mu_{l^{(k'-1)}}^{(k'-1)}}(X^{(k')}) \delta_{X^{(k')}}, \quad (1.4)$$

which by construction, is an *unbiased* empirical estimator of the reference distribution  $\pi$  (see Section 2.6 for a proof), up to sampling errors of the proposal  $\mu_l$ , which will

be done in practice using AMS. This leads to an IS estimator *e.g.* of the rare event probability

$$\hat{p} = \hat{\pi}^{(k)}(\mathbf{1}_{S^* > l}).$$

As mentioned above, the sequence of importance distribution  $\mu_{l^{(k)}}^{(k)}$  will be simulated by a population of particles and an AMS methodology in the spirit of current rare event simulation algorithms. The only difference is that the target distribution is constructed with a reduced score that is adaptively enriched by few evaluations of the true score throughout the simulation. As remarked in Section 1.2, updating the reduced score will in general break the nesting of the sequence of distributions. A simple way of addressing this problem is to perform a fresh AMS simulation each time that a new snapshot is picked and that the reduced score is updated. To significantly reduce the number of evaluations of the reduced model, another option is to propose a *bridging* procedure. This is the *second main methodological contribution* of this work. At iteration  $k$ , this procedure searches in the simulation history for a well-chosen previous proposal denoted  $\mu_{l^{(k_b)}}^{(k_b)}$  with  $k_b \leq k$ , so that the updated proposal denoted  $\mu_{l_0^{(k+1)}}^{(k+1)}$ , with  $l_0^{(k+1)}$  a well-designed level, is dominated and sufficiently close to the former. The latter proposal is then simulated starting from the former one using an AMS framework.

In numerical experiments, we study the proposed algorithm, called *Adaptive Reduced Multilevel Splitting* (ARMS) and show the empirical convergence of the ARMS algorithm on a toy model and a rare-event problem based on a realistic elliptic PDE and a *reduced basis* approach for model reduction. Our study confirms that the proposed relative entropy criterion plays a crucial role in ARMS and prevent the simulation to be taken into spurious regions of the state-space remote from the region of interest, resulting in a large variance of the IS estimator. Moreover, by evaluating the computational complexity in relation to the squared error of the estimate, we show empirically that the bridging procedure reduces in a certain regime the simulation cost while achieving a comparable squared error.

The paper is organized as follows. The main methodological concepts underlying the ARMS algorithm are presented in Section 2. We end this section by introducing an idealized version of the ARMS algorithm, which assumes the exact evaluation of expectations and exact sampling of distributions. In Section 3, we define ARMS, which consists in a practical implementation of this idealized algorithm using AMS simulations. In Section 4, we numerically evaluate the proposed algorithm for various rare event problems. A final section deals with conclusions and perspectives. The appendix collects some technical proofs and details.

## 2. Underlying Methodological Concepts.

**2.1. Importance sampling cost.** Let  $l_{\max}$  denote a large level of interest, and let us assume for the sake of the argument within the present section, that one is able to simulate a random variable  $X$  distributed according to the proposal distribution  $X \sim \mu_l = \frac{1}{Z_l} \mathbf{1}_{S > l} d\pi$ , the normalization  $Z_l$  being known. *Importance sampling* of the target rare event  $\{S^* > l\}$  can be performed in our context if the distribution  $\mu_l$  dominates the measure  $\mathbf{1}_{S^* > l} d\pi$ , and simply amounts to remark that one can estimate the averages

$$Z_l \mathbb{E} [\mathbf{1}_{S^*(X) > l} \varphi(X)] = \pi(\varphi \mathbf{1}_{S^* > l}),$$

where  $\varphi$  denotes a generic test function. For instance taking  $\varphi = \mathbf{1}$ , the above yields an unbiased estimator of the normalization probability  $p_l^*$ .

The domination requirement, denoted  $\mu_l \gg \eta_l^*$ , holds true by definition if and only if one has the inclusion  $\{S^* > l\} \subseteq \{S > l\}$ . This is not guaranteed here as  $S$  is an approximation of the true score  $S^*$ , and since  $S^*$  is not computed for each sample states, one cannot check it in a systematic fashion. Moreover, even if the above domination condition holds, importance sampling becomes quickly infeasible if the target distribution diverges too widely from the proposal, a typical phenomenon in high dimension. As discussed in Appendix A, it has been thoroughly argued in [8] (see also [7]) that the cost of importance sampling – in terms of the required sample size  $K$  – is quite generically and roughly given by the exponential of the relative entropy between  $\eta_l^*$  and  $\mu_l$ :

$$K \simeq e^{\text{Ent}(\eta_l^* | \mu_l)}. \quad (2.1)$$

This fact will justify the choice of various heuristics that will be found in this paper. We will call  $\text{Ent}(\eta | \mu)$  the *log cost of importance sampling* of the target  $\eta$  by the proposal  $\mu$ . In the present rare event case, note that relative entropy is simply the logarithm of small probability ratios:

$$\text{Ent}(\eta_l^* | \mu_l) = \ln \frac{Z_l}{p_l^*}$$

Note that domination between measures is implied by finite relative entropy, and will rather use throughout the paper the condition  $\text{Ent}(\eta | \mu) < +\infty$  in order to check the domination condition  $\mu \gg \eta$ .

**2.2. First non-achievable level.** In this section, we describe how to choose the target level  $l^{(k)}$ , which is suited to the current score function  $S^{(k)}$  at iteration  $k$  of the algorithm –  $k$  being the number of sampled snapshots of the true model. The idea is to tune  $l^{(k)}$  with a criteria evaluating that surpassing this level should require, in some sense, a better surrogate model, and thus new snapshots. The adjusted proposal  $\mu_{l^{(k)}}^{(k)}$  will then be used to sample a new snapshot, and obtain an updated score function  $S^{(k+1)}$  for the next iteration  $k + 1$ . In short, we are going to develop a heuristic criterion on levels to detect when a new snapshot is needed.

The general insight enabling to define  $l^{(k)}$  is the following. Consider the following problem: one is given a family of rare event target distributions  $(\eta_l^*)_{l \in \mathbb{R}}$  in the form of (1.1), a given computational budget, and a family of importance sampling proposal distributions  $(\mu_l^{(k)})_{l \in \mathbb{R}}$  in the form of (1.3). One wants to perform some sequential algorithm to sample the distributions  $\mu_l$  along an increasing sequence of levels  $l$  as high as possible, and then use importance sampling. However, we are given limited resources preventing sampling for levels that are unreasonably high. As an idealized criteria formalizing achievable levels, we propose the importance sampling cost (2.1) of the true target. The (ideal) first non-achievable (hence critical) level is thus given by the smallest level for which this cost matches a certain given numerical threshold  $c$ , that is:

$$\text{Find first } l : \quad \text{Ent}(\eta_l^* | \mu_l^{(k)}) \geq c.$$

By doing so, we ensure that the snapshot sampled with respect to  $\mu_{l^{(k)}}^{(k)}$  will not diverge too much from the target distribution  $\eta_{l^{(k)}}^*$ . As an example, if, as it is the



case in the present context,  $\mu_l^{(k)}$  has been constructed using a reduced model based on  $k$  snapshots  $(X_1, \dots, X_k)$ , the first non-achievable level  $l^{(k)}$  will typically be of order  $\max(S^*(X_1), \dots, S^*(X_k))$ , as the quality of the approximation  $S^* - S^{(k)}$  will deteriorate in areas with higher scores.

In a practical context, the log cost  $\text{Ent}(\eta_l^* | \mu_l^{(k)})$  has to be approximated using some prior information on  $\eta_l^*$ , for instance using some form of error quantification. We propose in this work the following method, which is by no means unique. We consider a *pessimistic* (respectively an *optimistic*) surrogate of  $\eta_l^*$ , denoted by  $\check{\eta}_l^{(k)}$  (respectively by  $\hat{\eta}_l^{(k)}$ ), constructed with the reduced score  $S^{(k)}$  and its error quantification  $E^{(k)}$ , and typically associated with a much lower (respectively higher) rare event probability than  $\eta_l^*$ . In the present work, we will choose (2.3)-(2.4) below which satisfy:

$$\check{\eta}_l^{(k)} \ll \eta_l^* \ll \hat{\eta}_l^{(k)}.$$

The first non-achievable level will be estimated using the *pessimistic* surrogate by searching the closest level such that the log cost  $\text{Ent}(\check{\eta}_l^{(k)} | \mu_l^{(k)})$  matches the threshold  $c$ . Given an initial level  $l_b$ , this yields the following rigorous definition of the (critical) first non-achievable level:

$$l^{(k)}(c, l_b) = \min\{l \geq l_b : \text{Ent}(\check{\eta}_l^{(k)} | \mu_l^{(k)}) \geq c \text{ or } \text{Ent}(\hat{\eta}_{l_{\max}}^{(k)} | \mu_l^{(k)}) = +\infty\}. \quad (2.2)$$

$l^{(k)}(c, l_b)$  is the smallest level above  $l_b$  such that, in the (unfortunate) case where the true target is given by the pessimistic choice  $\check{\eta}_l^{(k)}$ , a log cost of  $c$  is necessary in order to perform importance sampling. For convenience, the dependency on  $c$  and  $l_b$  of  $l^{(k)}$  will be removed if the context makes it clear. Note that the additional inequality constraint  $\text{Ent}(\hat{\eta}_{l_{\max}}^{(k)} | \mu_l^{(k)}) < +\infty$  in (2.2) ensures that the support of  $\mu_l^{(k)}$  contains the support of the final target  $\eta_{l_{\max}}^*$ .

In the present rare event case, the importance distribution is (1.3), the pointwise error bound is (1.2) and we propose the pessimistic and optimistic surrogate of the target (1.1)

$$\check{\eta}_l^{(k)} \stackrel{\text{def}}{=} \frac{1}{\check{Z}_l^{(k)}} \mathbf{1}_{S^{(k)} - E^{(k)} > l} d\pi \quad \text{with} \quad \check{Z}_l^{(k)} \stackrel{\text{def}}{=} \pi(\mathbf{1}_{S^{(k)} - E^{(k)} > l}), \quad (2.3)$$

and

$$\hat{\eta}_l^{(k)} \stackrel{\text{def}}{=} \frac{1}{\hat{Z}_l^{(k)}} \mathbf{1}_{S^{(k)} + E^{(k)} > l} d\pi \quad \text{with} \quad \hat{Z}_l^{(k)} \stackrel{\text{def}}{=} \pi(\mathbf{1}_{S^{(k)} + E^{(k)} > l}). \quad (2.4)$$

REMARK 2.1. *Using the surrogate (2.3), the particularization of the worst-case log cost in (2.2) is  $\text{Ent}(\check{\eta}_l^{(k)} | \mu_l^{(k)}) = \ln \frac{Z_l^{(k)}}{\check{Z}_l^{(k)}}$ , thus the logarithm of the ratio of the approximate probability  $\pi(\{S^{(k)} > l\})$  over the smaller pessimistic probability  $\pi(\{S^{(k)} - E^{(k)} > l\})$ .*

REMARK 2.2. *It can easily be checked that the critical level  $l^{(k)}(c, l_b)$  tends to infinity when the error  $E^{(k)}$  vanishes pointwise.*

**2.3. Snapshot sampling and learning function.** Given the importance distribution  $\mu_{l^{(k)}}^{(k)}$ , we now discuss the procedure used to sample a new snapshot. A possible choice used in the present work is the penalization of the importance distribution as follows:

$$X \sim \frac{1}{\mu_{l^{(k)}}^{(k)}(e^{\tau^{(k)} E^{(k)}})} e^{\tau^{(k)} E^{(k)}} \mu_{l^{(k)}}^{(k)},$$

with  $\tau^{(k)} \geq 0$ . The above penalization is similar to the so-called *learning function* in the literature on 'active learning' [20]. In this work, we simply choose an exponential function of the error  $c$  with a penalty parameter  $\tau^{(k)}$  adaptively chosen at iteration  $k$  of the algorithm. In particular,  $\tau^{(k)}$  will be chosen to be high in the beginning of the algorithm and then be reduced to 0 at some point. The idea here is that  $\tau^{(k)}$  enables to find a compromise between two objectives: first i) the snapshots are used to refine the reduced model  $S^{(k)}$ , and ii) the snapshots are also used in estimation by importance sampling. A large  $\tau^{(k)}$  is associated with the objective i). Indeed, favoring candidates with high errors in a  $\mu_{l^{(k)}}^{(k)}$ -distributed sample enables a faster learning of the true score  $S^*$ . A small or vanishing  $\tau^{(k)}$  is associated with objective ii). Indeed, the distribution  $\mu_{l^{(k)}}^{(k)}$  is constructed to be a robust importance sampling distribution for the target  $\eta_{l^{(k)}}^*$  so that states with smaller error are needed for estimation, ensuring a limited variance for importance sampling.

In the numerical experiments of Section 4, we will consider a very large  $\tau^{(k)}$  in a first phase of the algorithm (until a given number of snapshots are taken in the rare event  $\{S^* > l_{\max}\}$  of interest), and then, because estimation is then possible, we will take  $\tau^{(k)} = 0$ .

No attempt is made to optimize the learning function and the choice of the penalty parameter; those studies are left for future work, taking into account the thoroughful studies on active learning methods as found in the survey [20].

**2.4. Bridging two distributions with updated score approximation (optional).** We will now focus on what happens after the update of the reduced model that follows the sampling of a new snapshot. We are given a certain level  $l^{(k)}$  and an associated importance distribution  $\mu_{l^{(k)}}^{(k)}$ . The problem is to find a new initial level which we will denote  $l_b(k+1)$  and then sample the new importance distribution  $\mu_{l_b(k+1)}^{(k+1)}$  associated with the updated score approximation  $S^{(k+1)}$ .

We emphasize that this bridging methodology is only optional. It is perfectly possible (and we will provide some numerical comparisons) to set  $l_b(k+1) = -\infty$  which amounts to start at the next iteration with an importance distribution  $mu_{l_b(k+1)}^{(k+1)} = \pi$  given by the reference distribution.

This initial level  $l_b(k+1)$  will then be used in the subsequent  $k+1$  iteration to determine (with definition (2.2)) the critical level at iteration  $k+1$ :  $l^{(k+1)}(c, l_b(k+1))$  for  $S^{(k+1)}$ . This will complete the main loop of the algorithm.

However, if one tries to sample  $\mu_{l_b(k+1)}^{(k+1)}$  starting from  $\mu_{l^{(k)}}^{(k)}$ , using some form of importance or sequential sampling, there is an issue with inclusion. For instance, simply setting the initial level  $l_b(k+1)$  to  $l^{(k)}$ , the inclusion  $\{S^{(k+1)} > l^{(k)}\} \subseteq \{S^{(k)} > l^{(k)}\}$  will generally not be valid, making the task infeasible.

In order to circumvent this issue, we propose to keep in memory the importance samples  $\mu_{l^{(k')}}^{(k')}$  for  $k' \leq k$ . We then set a quantile  $\theta = \frac{M}{N} \in (0, 1)$  associated with an

effective sample size  $N - M \in \mathbb{N}$ , and find an index

$$\text{Find: } k_b \leq k$$

and a level

$$\text{Find: } l_b(k + 1)$$

verifying:

$$\text{Ent} \left( \mu_{l_b(k+1)}^{(k+1)} \mid \mu_{l^{(k_b)}}^{(k_b)} \right) \leq -\ln(1 - \theta). \quad (2.5)$$

This condition ensures that a relatively straightforward importance sampling procedure enables to obtain the initial importance sampling distribution  $\mu_{l_b(k+1)}^{(k+1)}$  of the forthcoming iteration, from a previous importance distribution  $\mu_{l^{(k_b)}}^{(k_b)}$ . In particular, the inclusion condition  $\{S^{(k+1)} \geq l_b(k + 1)\} \subseteq \{S^{(k_b)} > l^{(k_b)}\}$  will hold true.

In addition of the condition (2.5), we also want to ensure that the two conditions in (2.2) used to define the maximal achievable level  $l^{(k+1)}$  hold for this initial level  $l_b(k + 1)$ . This is necessary to ensure formally that  $l_b(k + 1) \leq l^{(k+1)}$ . Explicitly:

$$\text{Ent}(\tilde{\eta}_{l_b(k+1)}^{(k+1)} \mid \mu_{l_b(k+1)}^{(k+1)}) \leq c, \quad (2.6)$$

$$\text{Ent}(\hat{\eta}_{l_{\max}}^{(k+1)} \mid \mu_{l_b(k+1)}^{(k+1)}) < +\infty. \quad (2.7)$$

A pair  $(k_b, l_b(k + 1)) \in \llbracket 0, k \rrbracket \times \mathbb{R}$  satisfying (2.5), (2.6) and (2.7) determine two consecutive proposal distributions  $\mu_{l^{(k_b)}}^{(k_b)}$  and  $\mu_{l_b(k+1)}^{(k+1)}$  “bridging” the score approximations  $S^{(k_b)}$  and  $S^{(k+1)}$ . This pair can be determined according to various strategies.

- A trivial feasible pair is to set  $k_b = 0$  such that  $l^{(k_b)} = l_b(k + 1) = -\infty$ . As explained in Section 3, this will account in practice to restart the AMS simulation with the new score  $S^{(k+1)}$  from the beginning until it reaches the critical level  $l^{(k+1)}(c, -\infty)$ . However, this option can lead to a significant computational overhead if the complexity of the reduced score evaluation is not negligible.
- A non trivial feasible pair of interest is given by the largest (lexicographic) pair  $(k_b, l_b(k + 1)) \in \llbracket 0, k \rrbracket \times \mathbb{R}$  satisfying the three conditions (2.5), (2.6) and (2.7). This means first selecting the importance sampling distribution  $\mu_{l^{(k_b)}}^{(k_b)}$ , among the already computed distributions, from the most recent to the oldest, and then to search the largest level  $l_b(k + 1)$  such that the three conditions hold. We will see in Section 3 that in practice, this solution will lead to an important computation saving, but will require nevertheless to save the history of the previous AMS simulations.

**2.5. Stopping criteria.** In practical situations, a maximum level  $l_{\max}$  is setting the range of the computation of interest: one is interested in computing the rare event probability  $p_{l_{\max}}^*$ . The level  $l^{(k)}(c, l_b)$  defined in (2.2) is constrained to stay below  $l_{\max}$ , and the definition of  $l^{(k)}(c, l_b)$  is modified as the minimum between  $l_{\max}$  and the critical level given by (2.2). The algorithm can thus continue as long as one wishes, for instance until enough snapshot samples are obtained.

In addition, in order to avoid unnecessarily enriching an already accurate reduced score and obtain computational savings, it is possible to stop updating the reduced score if the worst-case log cost at the critical level  $l^{(k)}$  is almost zero and if this critical level is just below  $l_{\max}$ . More precisely we define the update stopping index<sup>1</sup> as the first index  $k$  such that

$$\text{Ent} \left( \tilde{\eta}_{l^{(k)}}^{(k)} \mid \mu_{l^{(k)}}^{(k)} \right) \leq \epsilon \quad \& \quad l^{(k)} \simeq l_{\max}, \quad (2.8)$$

where  $\epsilon > 0$  is a given threshold parameter close to the machine precision. This stop-update rule must however be implemented only when using a non-trivial bridging procedure. We will indeed see in Section 3 that using such a stopping of update of the reduced score, the normalization constant is no longer updated resulting in a sustained variance while new snapshots are sampled. Yet this sustained variance always occurs when using a non trivial bridging procedure. We will hence see in Section 3 that using the latter stop-update rule with a non-trivial bridge usually largely compensate additional variance with computational (and memory) savings.

**2.6. The estimator.** We are now ready to discuss the estimation of the probability  $p_{l_{\max}}^*$ . It is more convenient to work with the un-normalized version of the target distribution  $\eta_{l_{\max}}^*$  defined in (1.1), which will be denoted from now on by  $\gamma_{l_{\max}}^*$  and defined by

$$\gamma_{l_{\max}}^*(\varphi) \stackrel{\text{def}}{=} \pi(\varphi \mathbf{1}_{S^* > l_{\max}}),$$

for the level of interest  $l_{\max} \in \mathbb{R}$  and some test function  $\varphi$ . Note that  $p_{l_{\max}}^* = \gamma_{l_{\max}}^*(\mathbf{1})$ .

The importance sampling estimator  $\hat{\gamma}_{l_{\max}}^H$  of  $\gamma_{l_{\max}}^*$  is defined using a random iteration index which we will denote by  $K_{j_0}$ . Only the snapshots sampled at and after iteration  $K_{j_0} + 1$  will be used to build our importance sampling estimator. More precisely,  $K_{j_0}$  must be of *optional type*, in the sense that  $K_{j_0}$  is defined as the first iteration index  $k$  for which a certain property depending on the snapshots  $X_1, \dots, X_k$  holds true.

We propose the following option:  $K_{j_0}$  is the first iteration index  $k$  such that the score of the past snapshots have reached the level  $l_{\max}$  (the rare event)  $j_0$  times, for some prescribed  $j_0 \geq 1$ . At iteration  $K_{j_0} + H$  (for some  $H > 0$ ) of the algorithm, the estimator of  $\gamma_{l_{\max}}^*$  is then given by:

$$\hat{\gamma}_{l_{\max}}^H \stackrel{\text{def}}{=} \frac{1}{H} \sum_{k'=K_{j_0}+1}^{K_{j_0}+H} f_{l_{\max}}^{(k')}(X_{k'+1}) \delta_{X_{k'+1}}, \quad (2.9)$$

where, if we choose for  $k' \geq K_{j_0}$  to set the learning parameter as  $\tau = 0$ , the IS density ratio can be simplified as

$$f_{l_{\max}}^{(k')}(x) \stackrel{\text{def}}{=} \mathbf{1}_{S^* \geq l_{\max}}(x) \frac{d\pi}{d\mu_{l^{(k')}}^{(k')}}(x).$$

Setting a learning parameter  $\tau > 0$  in the early stages of the algorithm is crucial as the proposal law has a low probability mass over the set of rare events, which implies

<sup>1</sup>We call a stopping index any exit criteria of the algorithm that is independent of the future snapshots.

that snapshots taken so far are still insufficient to learn an accurate reduced model in this rare event regime. However, once  $j_0$  snapshots have already been sampled from the set of rare events, the choice of  $\tau = 0$  implies that the snapshot  $X_{k'+1}$  is drawn according to the proposition  $d\mu_{l^{(k')}}^{(k')}$ , thus favoring the robustness of the IS estimator rather than the accuracy of an already refined reduced model. Note that : i) the entropic criterion in (2.2) and (2.7) ensures that  $\eta_{l_{\max}}^* \ll \mu_{l^{(k)}}^{(k)}$ , and ii)  $f_{l_{\max}}^{(k)}(x)$  can be evaluated pointwise as long as we have an estimate of the normalizations constant of the proposal  $\mu_{l^{(k)}}^{(k)}$ .

As shown in Appendix B with a martingale argument, because of the optional property of  $K_{j_0}$ , this estimator is unbiased

$$\mathbb{E} \hat{\gamma}_{l_{\max}}^H = \gamma_{l_{\max}}^*,$$

and its variance decreases in  $\mathcal{O}(H^{-1})$ . Note that samples  $X_k$  in (2.9) are assumed to be drawn exactly from the sequence of importance distributions  $\mu_{l^{(k-1)}}^{(k-1)}$ . In practice, this assumption will not hold exactly because of the adaptive features used in the sampling of the proposals by the Monte Carlo routine, yielding a small bias of order the inverse of the sample size. This issue is discussed in Section 3.

**2.7. The idealized importance sampling algorithm.** Based on the concepts introduced in Sections 2.1 to 2.6, we can now describe in this section the algorithm that will lead to an estimation of  $\gamma_{l_{\max}}^*$  based on importance sampling.

**Algorithm 1 (Idealized ARMS)**

**Require:** level of interest  $l_{\max}$ , initial reduced score  $S^{(1)}$  and error  $E^{(1)}$  functions, budget of snapshots  $K$ , worst-case log cost threshold  $c$ , reduced score precision  $\epsilon \ll e$ , number of hits of rare event before estimation  $j_0$ , quantile  $\theta = M/N \in (0, 1)$  for the bridge between proposals, boolean *restart* variable for restarting or bridging after a reduced model update, initial learning parameter  $\tau_0$

- 1: Set  $l^{(0)} = l_b(1) = -\infty$ ,  $H = 0$ ,  $j = 0$ ,  $k = 0$ ,  $\tau = \tau_0$ ,  $Z_l^{(1)} = 1$  ▷ (initialization)
- 2: **while**  $K > 0$  **do**
- 3:      $k \leftarrow k + 1$  ▷ (searching the critical level, see Section 2.2)
- 4:     Simulate the flow of proposal distributions  $\mu_l^{(k)}$  for  $l \leq l^{(k)}$ . The maximal level  $l^{(k)}$  must be a level  $l \in [l_b(k), l_{\max}]$  satisfying both of the following two conditions:

$$\begin{aligned} \text{Ent}(\tilde{\eta}_l^{(k)} \mid \mu_l^{(k)}) &\leq c \\ \text{Ent}(\hat{\eta}_{l_{\max}}^{(k)} \mid \mu_l^{(k)}) &< +\infty. \end{aligned}$$

$l^{(k)}$  is defined as the smallest  $l \in [l_b(k), l_{\max}]$  (in practice  $l^{(k)} = l - \delta$  for some  $\delta > 0$  very small) violating one the above two conditions; if search fails, then  $l^{(k)} = l_{\max}$ .

- 5:     Compute<sup>2</sup> the associated normalisation  $Z_{l^{(k)}}^{(k)}$ . ▷ (snapshot sampling & score update, see Section 2.3)
- 6:     Draw
 
$$X \sim \frac{e^{\tau E^{(k)}}}{\mu_{l^{(k)}}^{(k)}(e^{\tau E^{(k)}})} \mu_{l^{(k)}}^{(k)}$$
- 7:     Compute the snapshot  $S^*(X) = \text{score}(\Psi(X))$
- 8:     **if**  $\text{Ent}(\tilde{\eta}_{l^{(k)}}^{(k)} \mid \mu_{l^{(k)}}^{(k)}) > \epsilon$  or  $l^{(k)} < l_{\max}$  **then**

---

<sup>2</sup>With a Sequential Monte Carlo approach for instance, see next sections.

```

9:      Update reduced model: build  $(S^{(k+1)}, E^{(k+1)})$  by enriching the previous surrogate
model  $(S^{(k)}, E^{(k)})$  with  $(X, S^*(X))$ .
10:    else
11:      Don't update the reduced model: set  $(S^{(k+1)}, E^{(k+1)}) \leftarrow (S^{(k)}, E^{(k)})$ .
12:    end if
13:    if  $S^*(X) > l_{\max}$  then  $j \leftarrow j + 1$  end if
14:    if  $j < j_0$  then update the learning function parameter  $\tau$  else set  $\tau = 0$  end if
▷ (importance sampling estimation, see Section 2.6)
15:    if  $j > j_0$  then
16:       $H \leftarrow H + 1$ 
17:

$$\hat{\gamma}_{l_{\max}}^H \leftarrow \frac{1}{H} \left( (H - 1) \hat{\gamma}_{l_{\max}}^{H-1} + Z_{l^{(k)}}^{(k)} \mathbf{1}_{S^*(X) > l_{\max}} \delta_X \right)$$

18:    end if
19:    Set  $K \leftarrow K - 1$ 
20:    if  $(S^{(k+1)}, E^{(k+1)}) \neq (S^{(k)}, E^{(k)})$  then
▷ (bridging, see Section 2.4)
21:      if restart then
22:        Find the largest (lexicographic) pair  $(k', l') \in \llbracket 0, k \rrbracket \times \mathbb{R}$  satisfying the following
three conditions:

```

$$\begin{aligned} \text{Ent}(\mu_{l'}^{(k+1)} \mid \mu_{l^{(k')}}^{(k')}) &\leq -\ln(1 - \theta) \\ \text{Ent}(\hat{\eta}_{l'}^{(k+1)} \mid \mu_{l'}^{(k+1)}) &\leq c, \\ \text{Ent}(\hat{\eta}_{l_{\max}}^{(k+1)} \mid \mu_{l'}^{(k+1)}) &< +\infty, \end{aligned}$$

and denote the solution pair by  $(k_b(k+1), l_b(k+1)) \leftarrow (k', l')$ .

```

23:      Simulate the proposal  $\mu_{l_b(k+1)}^{(k+1)}$  using  $\mu_{l^{k_b(k+1)}}^{(k_b(k+1))}$ ; then compute the associated
normalization  $Z_{l_b(k+1)}^{(k+1)}$ .
24:      else
25:        Set  $l_b(k+1) = -\infty$ 
26:        Set  $Z^{(k+1)} = 1$  and reinitialize  $\mu_{l_b(k+1)}^{(k+1)}$  with  $\pi$ 
27:      end if
28:      else
29:        Set  $l_b(k+1) = l^{(k)}$ 
30:        Set  $Z^{(k+1)} = Z^{(k)}$  and  $\mu_{l_b(k+1)}^{(k+1)} = \mu_{l^{(k)}}^{(k)}$ 
31:      end if
32:    end while
33:  Return Estimate  $\hat{\gamma}_{l_{\max}}^H$ 

```

This algorithm is an idealized concept and, in practice, it is not directly computable. Indeed it assumes exact evaluation of expectations and exact sampling according to the proposal and the surrogate distributions which are defined up to a normalization constant.

**3. Practical Implementation: AMS Simulations.** In this section, we propose a practical implementation of the idealized algorithm introduced in Section 2.7 using an AMS methodology. The latter will simulate the flow of proposals for growing level parameters  $l$  by providing Monte Carlo samples with approximately correct distributions for large sample sizes.

**3.1. AMS.** We wish to draw a sample of size  $N$  according to the importance sampling distribution  $\mu_{l^{(k)}}^{(k)}$  in the form (1.3), where, for ease of presentation, we temporarily assume that the target level  $l^{(k)}$  is known. We suppose we start with

a given sample of size  $N$  approximately distributed according to the law  $\mu_{l_b}^{(k)}$  with  $l_b < l^{(k)}$ , as well as with a given estimator of the associated normalization  $Z_{l_b}^{(k)}$ . The AMS algorithm [5] will be used here to generate a finite sequence of  $N$ -samples approximately distributed according to a sequence of nested distributions  $\{\mu_{l_i}^{(k)}\}_{i \geq 1}$ , as well as estimators of the associated normalizations  $\{Z_{l_i}^{(k)}\}_{i \geq 1}$ , *i.e.*, such that  $l_b < l_1 < \dots < l_m = l^{(k)}$ . At iteration  $i \geq 0$ , the  $N$ -sample of particles (clones) will be denoted

$$(\Xi_1^{(i)}, \dots, \Xi_N^{(i)})$$

and we omit in notation the dependence in  $(k)$ . The associated empirical distribution will be denoted

$$\mu_{l_i}^{N,(k)} := \frac{1}{N} \sum_{n=1}^N \delta_{\Xi_n^{(i)}} \simeq \mu_{l_i}^{(k)},$$

and the estimator of the normalization constant  $Z_{l_m}^{(k)}$  will be denoted  $Z_{l_m}^{N,(k)}$ . As a consequence, AMS will yield an approximation of the reference distribution  $\pi$  restricted to the current reduced score  $S^{(k)}$  as follows:

$$Z_{l_m}^{N,(k)} \mu_{l_m}^{N,(k)}(\varphi \mathbf{1}_{S^{(k)} > l^{(k)}}) \simeq \pi(\varphi \mathbf{1}_{S^{(k)} > l^{(k)}}),$$

for  $\varphi$  any given test function.

Let us describe briefly the AMS algorithm. At the  $i$ -th iteration, the algorithm performs successively: *i*) a weighting and selection step, and *ii*) a mutation step. The weighting step *i*) assigns weights to the empirical distribution  $\mu_{l_i}^{N,(k)}$  so that it approximates  $\mu_{l_{i+1}}^{(k)}$ , up to a normalization constant given by an empirical approximation of the normalization ratios. More precisely, the empirical distribution at level  $l_{i+1}$  and the associated normalizations are obtained by weighting the current  $N$ -sample:

$$\frac{Z_{l_{i+1}}^{N,(k)}}{Z_{l_i}^{N,(k)}} \mu_{l_{i+1}}^{N,(k)} := \frac{1}{N} \sum_{j=1}^N \mathbf{1}_{S^{(k)} > l_{i+1}} \delta_{\Xi_j^{(i)}},$$

where the equality is well-defined because  $l_i < l_{i+1}$ . This empirical distribution of course requires the determination (or a choice) of  $l_{i+1}$ . AMS performs an adaptive choice of  $l_{i+1}$  as follows. AMS starts by sorting the  $N$  particles according to their scores in ascending order. AMS then considers an adaptive target distribution defined by a fixed number  $M < N$  of particles. The latter are assigned a weight equal to zero, and the level  $l_{i+1}$  is defined by the  $M$ -th smallest particle score

$$l_{i+1} = S^{(k)}(\Xi_{(M)}^{(i)}),$$

in which  $(1), \dots, (N)$  is the usual shorthand notation for increasing order statistics (based on scores, here). In other words, this corresponds to set  $l_{i+1}$  as the highest level  $l$  in a set given by a maximal relative entropy

$$l_{i+1} := \sup \left\{ l \in [l_i, l_{\max}] : \text{Ent}(\mu_l^{(k)} \mid \mu_{l_i}^{(k)}) \simeq -\ln \mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} > l}) < -\ln \left( \frac{N-M}{N} \right) \right\},$$

with the quantile (proportion of killed particles) satisfies  $M = \lfloor \theta N \rfloor$ .

In the selection step, the first  $M$  particles with zero-weights are removed and surviving particles are duplicated. The  $M$  resampled particles are chosen uniformly among the surviving particles. The AMS approximation of the normalization constant thus follows the recursion<sup>3</sup>

$$Z_{l_{i+1}}^{N,(k)} = Z_{l_i}^{N,(k)} \mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} > l_{i+1}}) = Z_{l_i}^{N,(k)}(1 - M/N), \quad (3.1)$$

yielding

$$Z_{l^{(k)}}^{N,(k)} = Z_{l_m}^{N,(k)} = Z_{l_b}^{N,(k)}(1 - M/N)^m.$$

The mutation step  $ii$ ) then consists of a Markovian transition that leaves  $\mu_{l_{i+1}}^{(k)}$  invariant and is applied individually and independently a prescribed number of times to the  $M$  particles after resampling, the goal being to bring diversity in the  $N$ -sample to counterbalance the effect of selection. In practice, Markovian transitions can be performed using a Metropolis-Hasting algorithm or other techniques studied in the Markov Chain Monte Carlo (MCMC) literature. We mention that the Markovian kernel is often chosen to be adaptive in order to maintain a given Metropolis acceptance rate [15].

Note that AMS typically involves many evaluations of the score function underlying the sequence of targeted distributions. This paper precisely focuses on cases where the true score  $S^*$  too expensive for so many evaluation, is substituted by but a reduced score  $S^{(k)}$ , much cheaper to evaluate.

**3.2. Reaching the critical level with AMS.** Let us now describe the main Monte Carlo sampling routine of the algorithm. We recall that after  $k$  iterations, one is given with a current reduced score  $S^{(k)}$ , an associated pointwise error estimate  $E^{(k)}$ , an initial level  $l_b(k)$  and a sample of size  $N$  approximately distributed according to  $\mu_{l_b(k)}^{(k)}$ , that we denote using a sample-size superscript:  $\mu_{l_b(k)}^{N,(k)}$ . One want to perform an AMS simulation targeting the distribution  $\mu_{l^{(k)}}^{(k)}$ , where we recall that  $l^{(k)}$  is the critical, non-achievable level the closest to  $l_b(k)$ . The latter is formally defined by (2.2) (for  $N = +\infty$ ). When using sample-based empirical distributions, one will need to adapt the latter conditions.

The AMS simulation will proceed as explained in Section 3.1, creating a sequence  $(l_i)_{i \geq 0}$  of increasing levels starting from  $l_b = l_b(k)$ , creating a sequence of empirical distributions  $\mu_{l_i}^{N,(k)}$  for  $i \geq 0$ . We then check for each proposed level increase from  $l_i$  to  $l_{i+1}$  that the entropic conditions (2.2) using the empirical distribution  $\mu_{l_i}^{N,(k)}$  to approximate the latter. More precisely, we remark that

$$\text{Ent} \left( \tilde{\eta}_{l_{i+1}}^{(k)} \mid \mu_{l_{i+1}}^{(k)} \right) = \ln \frac{\mu_{l_i}^{(k)}(\mathbf{1}_{S^{(k)} > l_{i+1}})}{\mu_{l_i}^{(k)}(\mathbf{1}_{S^{(k)} - E^{(k)} > l_{i+1}}} \simeq \ln \frac{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} > l_{i+1}})}{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} - E^{(k)} > l_{i+1}})}, \quad (3.2)$$

$$\text{Ent} \left( \hat{\eta}_{l_{\max}}^{(k)} \mid \mu_{l_{i+1}}^{(k)} \right) = \ln \frac{\mu_{l_i}^{(k)}(\mathbf{1}_{S^{(k)} > l_{i+1}})}{\mu_{l_i}^{(k)}(\mathbf{1}_{S^{(k)} + E^{(k)} > l_{\max}})} \simeq \ln \frac{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} > l_{i+1}})}{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} + E^{(k)} > l_{\max}})}, \quad (3.3)$$

<sup>3</sup>To ease the presentation we have assumed here that particles have different scores. However, in the case  $\exists \mathcal{J} \subset \mathbb{N}$  such that  $S^{(k)}(\Xi_{(M)}^{(i)}) = S^{(k)}(\Xi_{(M+j)}^{(i)})$  for  $j \in \mathcal{J}$ , a good practice is to kill in addition the set of particles  $\{S^{(k)}(\Xi_{(M+j)}^{(i)})\}_{j \in \mathcal{J}}$ , and locally change the value of  $M$  in (3.1) to  $M + \text{card}(\mathcal{J})$  accordingly.



so that (2.2) is re-written as:

$$l^{(k)} = l_m :=$$

$$\min\{l_i, i \geq 0 : \ln \frac{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} > l_{i+1}})}{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} - E^{(k)} > l_{i+1}})} \geq c \text{ or } \ln \frac{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} > l_{i+1}})}{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} + E^{(k)} > l_{\max}})} = +\infty\}.$$

The above means that when the AMS simulation proposes to increment the level  $l_i$  to  $l_{i+1}$  and that at least one of these conditions is violated, then the critical level set to  $l^{(k)} = l_i$  is reached.

REMARK 3.1. *If the logarithm of the ratio in (3.3) is finite when  $N \rightarrow \infty$  for all intermediate levels  $l_j$  from  $-\infty$  to  $l_i$  then we can verify recursively that, at a formal level,*

$$\lim_{N \rightarrow \infty} \ln \frac{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} > l_{i+1}})}{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} + E^{(k)} > l_{\max}})} < +\infty \Leftrightarrow \text{Ent} \left( \hat{\eta}_{l_{\max}}^{(k)} \mid \mu_{l_{i+1}}^{(k)} \right) < +\infty,$$

which precisely means that the rare event of interest is a subset of the support of the current proposal distribution  $\mu_{l_{i+1}}^{(k)}$ .

The procedure to reach the critical level in practice is summarized by the following algorithm.

**Algorithm 2 (AMS simulation up to the critical level)**

**Require:** level of interest  $l_{\max}$ , initial distribution  $\mu_{l_b}^{N,(k)}$  and normalization  $Z_{l_b}^{N,(k)}$ , related reduced score  $S^{(k)}$  and error  $E^{(k)}$ , quantile  $\theta \in [1/N, 1[$ , worst-case log cost threshold  $c$ .

- 1:
- 2: Set  $i = 0, c_{try} = 0$
- 3: **while**  $l_i < l_{\max}$  &  $c_{try} \leq c$  &  $d_{try} < +\infty$   
**do**
- 4: Order scores  $S^{(k)}$  relative to particles of  $\mu_{l_i}^{N,(k)}$
- 5: Set  $l_{i+1} = \min(l_M, l_{\max})$  with  $l_M$  the  $M$ -th smallest score
- 6: Compute

$$c_{try} = \ln \frac{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} > l_{i+1}})}{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} - E^{(k)} > l_{i+1}})} \simeq \text{Ent} \left( \hat{\eta}_{l_{i+1}}^{(k)} \mid \mu_{l_{i+1}}^{(k)} \right)$$

as well as

$$d_{try} = \ln \frac{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} > l_{i+1}})}{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} + E^{(k)} > l_{i+1}})} \simeq \text{Ent} \left( \hat{\eta}_{l_{\max}}^{(k)} \mid \mu_{l_{i+1}}^{(k)} \right)$$

as defined in (3.2) and (3.3)

- 7: **if**  $l_{i+1} < l_{\max}$  &  $c_{try} \leq c$  &  $d_{try} < +\infty$  **then**
- 8: Kill the  $M \geq \lceil \theta N \rceil$  particles with the lowest score or with score equal to  $l_{i+1}$
- 9: Update  $Z_{l_{i+1}}^{N,(k)}$  with recursion (3.1)
- 10: Perform  $M$  duplications (uniformly among the survivors)
- 11: Mutate the duplicates using a MCMC leaving target  $\mu_{l_{i+1}}^{(k)}$  invariant.
- 12: Denote the resulting sample  $\mu_{l_{i+1}}^{N,(k)}$ .
- 13: **else**
- 14: Set  $l^{(k)} = l_i$
- 15: **end if**
- 16: Set  $i \leftarrow i + 1$

17: **end while**

18: **Return** Empirical measure  $\mu_{l_i^{(k)}}^{N,(k)}$  and normalizing constant  $Z_{l_i^{(k)}}^{N,(k)}$

The following remark is useful in practice, as it provides some prior information on the worst-case log cost threshold parameter  $c$ .

**REMARK 3.2** (Worst-case log cost: a lower bound in practice). *Consider, in Algorithm 2 above, the empirical distribution  $\mu_{l_i}^{N,(k)}$ . By construction of AMS,  $l_{i+1}$  is the  $M$ -th smallest score among the  $N$  scores. Set  $c_{\min}(M, N) \stackrel{\text{def}}{=} \ln \frac{N - M}{N - M - 1}$ . As long as  $E^{(k)}$  is non zero, at least one particle of the sample, the  $M$ -th particle, lies outside the level set obtained by subtracting the error to the score.  $c_{\min}(M, N)$  is thus the lower bound of the entropy between the empirical measures with and without the error term, that is*

$$c_{\min}(M, N) \leq \ln \frac{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} > l_{i+1}})}{\mu_{l_i}^{N,(k)}(\mathbf{1}_{S^{(k)} - E^{(k)} > l_{i+1}}}.$$

Setting the numerical threshold to the minimal possible value  $c = c_{\min}(M, N)$  can be considered a safe strategy, leading to a slower but safer convergence of the algorithm.

**3.3. Bridging with an AMS step.** As exposed in Section 2.4, the bridging procedure at the end of iteration  $k$  consist in selecting a past index  $= k_b := k_b(k) \leq k$  and a new initial level  $l_b(k + 1)$  which will be used in the AMS simulation with the new score approximation  $S^{(k+1)}$  starting from a past (appropriately recorded) empirical distribution  $\mu_{l_b(k_b)}^{N,(k_b)}$ . As defined previously, bridging consists in choosing, using a lexicographic search, a feasible pair  $(k_b, l_b(k + 1)) \in \llbracket 0, k \rrbracket \times \mathbb{R}$ , where feasibility means satisfying the quantile condition (2.5) and the two entropic conditions (2.6) and (2.7). It is now necessary to define empirical counterparts to the latter conditions. As previously, the two entropic conditions are approximated by

$$\text{Ent} \left( \tilde{\eta}_{l'}^{(k+1)} \mid \mu_{l'}^{(k+1)} \right) \simeq \ln \frac{\mu_{l^{(k')}}^{N,(k')}(\mathbf{1}_{S^{(k+1)} > l'})}{\mu_{l^{(k')}}^{N,(k')}(\mathbf{1}_{S^{(k+1)} - E^{(k+1)} > l'})} \leq c, \quad (3.4)$$

$$\text{Ent} \left( \tilde{\eta}_{l_{\max}}^{(k+1)} \mid \mu_{l'}^{(k+1)} \right) \simeq \ln \frac{\mu_{l^{(k')}}^{N,(k')}(\mathbf{1}_{S^{(k+1)} > l'})}{\mu_{l^{(k')}}^{N,(k')}(\mathbf{1}_{S^{(k+1)} + E^{(k+1)} > l_{\max}})} < +\infty, \quad (3.5)$$

where in the above  $(k', l')$  denote feasible pairs that are candidates for the solution  $(k_b, l_b(k + 1))$ . Note that the well-posedness of (3.5) can be justified by Remark 3.1. Concerning the quantile condition, we remark that, given some  $k'$  and  $l'$ , and assuming that the domination relation stands

$$\mu_{l^{(k')}}^{(k')} \gg \mu_{l'}^{(k+1)} \Leftrightarrow \left\{ S^{(k+1)} > l' \right\} \subset \left\{ S^{(k')} > l^{(k')} \right\}, \quad (3.6)$$

an empirical approximation of the quantile condition (2.5) is

$$\mu_{l'}^{N,(k')}(\mathbf{1}_{S^{(k+1)} \leq l'}) \leq \theta. \quad (3.7)$$

In order to verify (3.6), we propose to browse the history of past AMS simulations and check the inclusion of the support of  $\mu_{l'}^{(k+1)}$  in the one of  $\mu_{l^{(k')}}^{(k')}$  using all the recorded

$N$ -samples targeting the past proposal distributions  $\mu_{l(\tilde{k})}^{N,(\tilde{k})}$  with  $\tilde{k} < k_b$ , whose samples are scattered between the reference distribution and the current distribution. This check must be carried out for all  $\tilde{k} < k_b$  and is given by

$$\mu_{l(\tilde{k})}^{N,(\tilde{k})} (\mathbf{1}_{S^{(k')} > l^{(k')}} \mathbf{1}_{S^{(k+1)} > l'} + \mathbf{1}_{S^{(k+1)} \leq l'}) = 1. \quad (3.8)$$

To sum up, bridging consist in selecting some feasible pair, *i.e.*, some pair  $(k_b, l_b(k+1))$  satisfying (3.4), (3.5), (3.7) and (3.8). Note that once the pair has been determined, the normalization estimate can be updated as

$$Z_{l_b(k+1)}^{N,(k+1)} = Z_{l^{(k_b)}}^{N,(k_b)} \mu_{l^{(k_b)}}^{N,(k_b)} (\mathbf{1}_{S^{(k+1)} > l_b(k+1)}). \quad (3.9)$$

Then the  $\tilde{M}(k_b)$  particles (at max given by  $\tilde{M}(k_b) = \lfloor \theta N \rfloor$  when (3.7) is saturated) with updated scores below or equal to  $l_b(k+1)$  are replaced by duplicating surviving particles and performing mutation targeting  $\mu_{l_b(k+1)}^{(k+1)}$  (in an AMS line of thinking). The next AMS simulation with  $S^{(k+1)}$  can then be started from the level  $l_b(k+1)$ . The following iteration then starts up to the critical level  $l^{(k+1)}(c, l_b(k+1))$ , an so on.

As mentioned in Section 2.4, among the feasible set, we focus on either the trivial pair  $(k_b, l_b(k+1)) = (0, -\infty)$ , or the largest (lexicographic) pair  $(k_b, l_b(k+1)) \in \llbracket 0, k \rrbracket \times \mathbb{R}$ . A simple way to solve this maximisation problem is to begin with the highest possible index  $k' := k$  and then to decrease  $k'$  by a unit if no feasible level  $l'$  has been found for this index. To search for the largest feasible level  $l'$  at a given index  $k'$ , one can proceed as follows. Begin by setting  $l'$  to the highest level satisfying (3.7), then lower this level until (3.4) and (3.5) hold and finally perform for all  $\tilde{k} < k_b$  the check (3.8).

The non-trivial bridging procedure is summarized by the following algorithm.

### Algorithm 3 (Bridging)

**Require:** level of interest  $l_{\max}$ , current index  $k$ , for all  $k' \leq k$ , particles of distribution  $\mu_{l^{(k')}}^{N,(k')}$ , normalization  $Z_{l^{(k')}}^{N,(k')}$  and related reduced score  $S^{(k')}$  and error  $E^{(k')}$ , reduced score  $S^{(k+1)}$ , the quantile  $\theta \in [1/N, 1[$ , worst-case log cost threshold  $c$ .

- 1:
- 2: Set  $k' = k$  and  $feasible = false$  ▷ (find largest feasible pair)
- 3: **while**  $!feasible$  &  $k' > 0$  **do**
- 4:     Compute  $S^{k+1}$  for particles in  $\mu_{l^{(k')}}^{(k')}$  and sort their scores
- 5:     Set  $\tilde{M} = \lfloor N\theta \rfloor$
- 6:     **while**  $!feasible$  &  $\tilde{M} > 0$  **do**
- 7:         Set  $l' = \min(l_{\tilde{M}}, l_{\max})$  with  $l_{\tilde{M}}$  the  $\tilde{M}$ -th smallest score
- 8:         Compute (3.4):  $c_{try} = \ln \frac{\mu_{l^{(k')}}^{N,(k')} (\mathbf{1}_{S^{(k+1)} > l'})}{\mu_{l^{(k')}}^{N,(k')} (\mathbf{1}_{S^{(k+1)} - E^{(k+1)} > l'})} \simeq \text{Ent} \left( \hat{\eta}_{l'}^{(k+1)} \mid \mu_{l'}^{(k+1)} \right)$
- 9:         Compute (3.5):  $d_{try} = \ln \frac{\mu_{l^{(k')}}^{N,(k')} (\mathbf{1}_{S^{(k+1)} > l'})}{\mu_{l^{(k')}}^{N,(k')} (\mathbf{1}_{S^{(k+1)} + E^{(k+1)} > l_{\max}})}$   $\simeq \text{Ent} \left( \hat{\eta}_{l_{\max}}^{(k+1)} \mid \mu_{l'}^{(k+1)} \right)$
- 10:        **if**  $c_{try} \leq c$  &  $d_{try} < +\infty$  **then**
- 11:            Set  $feasible = true$
- 12:            **for**  $\tilde{k} \in \llbracket 0, k' - 1 \rrbracket$  **do**
- 13:                Compute  $S^{(k+1)}$  and  $S^{(k')}$  for particles in  $\mu_{l(\tilde{k})}^{(k')}$
- 14:                **if**  $\mu_{l(\tilde{k})}^{N,(\tilde{k})} (\mathbf{1}_{S^{(k')} > l^{(k')}} \mathbf{1}_{S^{(k+1)} > l'} + \mathbf{1}_{S^{(k+1)} < l'}) \neq 1$
- 15:                    Set  $feasible = false$

```

16:           end if
17:       end for
18:       Set  $\tilde{M} \leftarrow \tilde{M} - 1$ 
19:   end while
20:   Set  $k' \leftarrow k' - 1$ 
21: end while
22: if feasible then                                ▷ ( $k' > 0$ : perform an AMS step)
23:     Set the largest pair as  $(k_b, l_b(k+1)) = (k', l')$ 
24:     Compute  $Z_{l_b(k+1)}^{N, (k+1)}$  using (3.9)
25:     Kill the  $\tilde{M}$  particles in  $\mu_{l(k_b)}^{N, (k_b)}$  of smallest score  $S^{(k+1)}$  or with score equal to  $l_b(k+1)$ 
26:     Perform duplications (uniformly among the survivors) to replace the killed particles
27:     Simulate  $\mu_{l_b(k+1)}^{N, (k+1)}$  by mutating the duplicates to target  $\mu_{l_b(k+1)}^{(k+1)}$ 
28: else                                                ▷ ( $k' = 0$ : initialize a fresh AMS)
29:   Set  $l_b(k+1) = -\infty$ 
30:   Set  $Z_{l_b(k+1)}^{N, (k+1)} = 1$ 
31:   Simulate  $\mu_{l_b(k+1)}^{N, (k+1)}$  (by sampling the reference distribution  $\pi$ )
32: end if
33: Return Empirical measure  $\mu_{l_b(k+1)}^{N, (k+1)}$  and normalizing constant  $Z_{l_b(k+1)}^{N, (k+1)}$ 

```

**3.4. AMS-based importance sampling estimator.** Using the sequence of empirical distributions  $\{\mu_{l^{(k)}}^{N, (k)}\}_{k=1}^K$  generated by AMS simulations, we can write a closed-form approximation of the importance sampling estimator  $\hat{\gamma}_{l_{\max}}^H$  of  $\gamma_{l_{\max}}^*$  defined by (2.9). The AMS approximation of the estimator  $\hat{\gamma}_{l_{\max}}^H$  is defined by

$$\hat{\gamma}_{l_{\max}}^{H, N}(\varphi) \stackrel{\text{def}}{=} \frac{1}{H} \sum_{k=K_{j_0}+1}^{K_{j_0}+H} Z_{l^{(k)}}^{N, (k)} \varphi(X_{k+1}) \mathbf{1}_{S^* > l_{\max}}(X_{k+1}), \quad (3.10)$$

In particular, the AMS approximation of the estimator  $\hat{\gamma}_{l_{\max}}^H(\mathbf{1})$  of  $p^* = \gamma_{l_{\max}}^*(\mathbf{1})$  is

$$\hat{p}_{IS} \stackrel{\text{def}}{=} \hat{\gamma}_{l_{\max}}^{H, N}(\mathbf{1}), \quad (3.11)$$

AMS also provides as by-products a sequence of approximate estimators of  $\pi(\{S^{(k)} \geq l_{\max}\})$  given by

$$\hat{p}_{AMS} \stackrel{\text{def}}{=} Z_{l^{(k)}}^{N, (k)} \mu_{l^{(k)}}^{N, (k)}(\mathbf{1}_{S^{(k)} > l_{\max}}), \quad (3.12)$$

This estimator can be an increasingly accurate approximation of  $p^*$ , but as our numerical experiments show, it can also be highly biased, even in the case where the reduced score  $S^{(k)}$  converges uniformly to  $S^*$  as  $k$  increases.

**3.5. The ARMS algorithm.** The complete methodology proposed in this work is summarized in the complete algorithm presented below named adaptive reduced multilevel splitting (ARMS). ARMS is a practical implementation of the idealized importance sampling estimator (Algorithm 1), built on the AMS simulation framework described above.

**Algorithm 4 (ARMS)**

**Require:** level of interest  $l_{\max}$ , initial reduced score  $S^{(1)}$  and error  $E^{(1)}$ , budget of snapshots  $K$ , the quantile  $\theta \in (0, 1)$  proposing the bridge level and level increases in AMS, number of hits of rare event target  $j_0$ , number of particles  $N$  for AMS, worst-case log cost threshold  $c$ , reduced score precision  $\epsilon \ll e$ , boolean *restart* variable for restarting or bridging after a reduced model update, initial learning parameter  $\tau_0$ .

▷ (**initialization**)

1: Set  $l^{(0)} = l_b(1) = -\infty$ ,  $H = 0$ ,  $j = 0$ ,  $k = 0$ ,  $\tau = \tau_0$ ,  $Z_{l^{(0)}}^{N,(0)} = 1$   
2: Set  $\mu_{l_b(1)}^{N,(1)}$  by drawing a  $N$ -sample of the reference distribution  $\pi$   
3: **while**  $K > 0$  **do**  
4:      $k \leftarrow k + 1$  ▷ (AMS simulation to critical level)  
5:     Use **Algorithm 2** (see Section 3.2) from  $l_b = l_b(k)$  to:  

- estimate the critical level  $l^{(k)} \in [l_b(k), l_{\max}]$
- simulate  $\mu_{l^{(k)}}^{N,(k)}$  starting from  $\mu_{l_b(k)}^{N,(k)}$
- compute the associated normalisation  $Z_{l^{(k)}}^{N,(k)}$ . ▷ (snapshot sampling & score update)

6:     Draw  

$$X \sim \frac{\mu_{l^{(k)}}^{N,(k)} e^{\tau E^{(k)}}}{\mu_{l^{(k)}}^{N,(k)} (e^{\tau E^{(k)}})}$$
7:     Compute the snapshot  $S^*(X)$   
8:     **if**  $l^{(k)} < l_{\max}$  or  $\text{Ent}(\hat{\eta}_{l^{(k)}}^{(k)} | \mu_{l^{(k)}}^{(k)}) \simeq -\ln \mu_{l^{(k)}}^{N,(k)} (\mathbf{1}_{S^{(k)} - E^{(k)} > l^{(k)}}) > \epsilon$  **then**  
9:         Build  $(S^{(k+1)}, E^{(k+1)})$  by enriching  $(S^{(k)}, E^{(k)})$  with  $(X, S^*(X))$ .  
10:     **else**  
11:         Set  $(S^{(k+1)}, E^{(k+1)}) \leftarrow (S^{(k)}, E^{(k)})$ .  
12:     **end if**  
13:     **if**  $S^*(X) > l_{\max}$  **then**  $j \leftarrow j + 1$  **end if**  
14:     **if**  $j < j_0$  **then** update the learning function parameter  $\tau$  **else** set  $\tau = 0$  **end if** ▷ (importance sampling)  
15:     **if**  $j > j_0$  **then**  
16:          $H \leftarrow H + 1$   
17:         Update estimate (see Section 3.4)  

$$\hat{\gamma}_{l_{\max}}^{H,N} \leftarrow \frac{1}{H} \left( (H-1) \hat{\gamma}_{l_{\max}}^{H-1,N} + Z_{l^{(k)}}^{N,(k)} \mathbf{1}_{S^*(X) > l_{\max}} \delta_X \right)$$
18:     **end if**  
19:     Set  $K \leftarrow K - 1$   
20:     **if**  $(S^{(k+1)}, E^{(k+1)}) \neq (S^{(k)}, E^{(k)})$  **then**  
21:         **if**  $!restart$  **then** ▷ (AMS step for bridging)  
22:             Use **Algorithm 3** (see Section 3.3) to:  

- find the largest feasible pair  $(k_b, l_b(k+1))$
- simulate  $\mu_{l_b(k+1)}^{N,(k+1)}$  from  $\mu_{l^{(k_b)}}^{N,(k_b)}$
- compute the associated normalisation  $Z_{l_b(k+1)}^{N,(k+1)}$ . ▷ (or AMS reinitialization)

23:             **else**  
24:                 Set  $l_b(k+1) = -\infty$   
25:                 Set  $Z_{l_b(k+1)}^{N,(k+1)} = 1$   
26:                 Set  $\mu_{l_b(k+1)}^{N,(k+1)}$  by drawing a  $N$ -sample of the reference distribution  $\pi$   
27:             **end if**  
28:             **else**  
29:                 Set  $l_b(k+1) = l^{(k)}$   
30:                 Set  $Z_{l_b(k+1)}^{N,(k+1)} = Z_{l^{(k)}}^{N,(k)}$ ,  $\mu_{l_b(k+1)}^{N,(k+1)} = \mu_{l^{(k)}}^{N,(k)}$   
31:             **end if**  
32:     **end while**  
33: **Return** Estimate  $\hat{\gamma}_{l_{\max}}^{H,N}$

We comment hereafter the ARMS algorithm. The simulation starts from an i.i.d. sample with the reference distribution. ARMS then iterates through four blocks until the computing budget allocated to snapshots is exhausted.

The first block (lines 4 to 5) consists in increasing the level until reaching the first non-achievable level presented in Section 2.2 and implemented in practice by an AMS routine in Section 3.2. In short, it simulates the distribution flow  $l \mapsto \mu_l^{N,(k)}$  with  $l \in [l_b(k), l^{(k)}]$ , where  $l^{(k)}$  is the critical level where IS of the true target  $\eta_{l^{(k)}}^*$  by the current reduced model  $S^{(k)}$  is expected to be too expensive.

The second block (lines 6 to 14) samples a new sample according to the proposal  $\mu_{l^{(k)}}^{N,(k)}$ , penalized by some learning function in the early iterations. It then computes the related snapshot, updates the reduced score and finally updates the learning function. This block follows the guideline presented in Section 2.3.

The third block (lines 15 to 19) is the estimation phase beginning when the algorithm has reached the rare event in some sense. This block performs importance sampling with the proposal  $\mu_{l^{(k)}}^{N,(k)}$  and updates the IS estimate (3.10).

The fourth and final block (lines 20 to 31) is the restart or bridge routine presented in Section 2.4 and implemented in practice by an AMS step in Section 3.3. In short, if the restart option is untriggered ( $restart = false$ ), then the bridging procedure searches in past simulations to find a certain pair  $(k_b, l_b(k+1))$  which enables to simulate from  $\mu_{l^{(k_b)}}^{N,(k_b)}$  towards  $\mu_{l_b(k+1)}^{N,(k+1)}$  in a single AMS step. Otherwise (in the case  $restart = true$ ), a fresh AMS simulation is initialized at each score approximation update. We mention that the bridging routine is avoided in the case where the stopping criterion (2.8) is reached. Indeed, in this case, no bridging is needed, since the reduced score is not updated anymore and the next initial level can simply be set to  $l^{(k)}$ .

**4. Numerical Simulations.** We first consider a toy example given by an explicit one dimensional true score. Reduced modeling is performed by cubic spline interpolation. The rare event simulation is in this case one-dimensional but nonetheless challenging, as the score presents several maxima and moreover splines may produce highly oscillating approximations. We then focus our attention on a more realistic scenario in which the rare event is defined through the high-dimensional solution of an elliptic PDE with random inputs. The PDE solution is approximated by a reduced basis approach which yields the reduced models. All our numerical simulations are carried out using an implementation of the ARMS algorithm in the Julia language, available at <https://gitlab.inria.fr/pheas/arms>.

**4.1. Examples.** We will consider the problem of estimation of  $p^* = \pi(\mathbf{1}_{S^* > l_{\max}})$  with  $S^*(X) = \text{score}(\Psi^*(X))$ , for various models  $\Psi^*$  and reduced models  $\Psi$ , various dimensions  $q = \dim(X)$  of the parameter space, various model dimensions  $h = \dim(\Psi^*(X))$  and different score functions.

	$\Psi^*$	$\Psi$	$\dim(X)$	$\dim(\Psi^*(X))$	score	$p^*$	$l_{\max}$
example #1	analytical	spline	1	1	$ \cdot $	2.18e-08	90
example #2a	PDE	RB	4	5101	$L^1$	6.31e-04	0.5
example #2b	PDE	RB	4	5101	$L^\infty$	1.47e-04	3.0
example #2c	PDE	RB	25	5101	$L^1$	2.48e-14	0.5

We detail hereafter these examples.

**4.1.1. Example #1: spline approximation of a closed-form score.** In a first numerical experiment, we design a one-dimensional analytic toy model. Cubic spline data interpolation is used to build the surrogate  $S(x)$  from the snapshots and

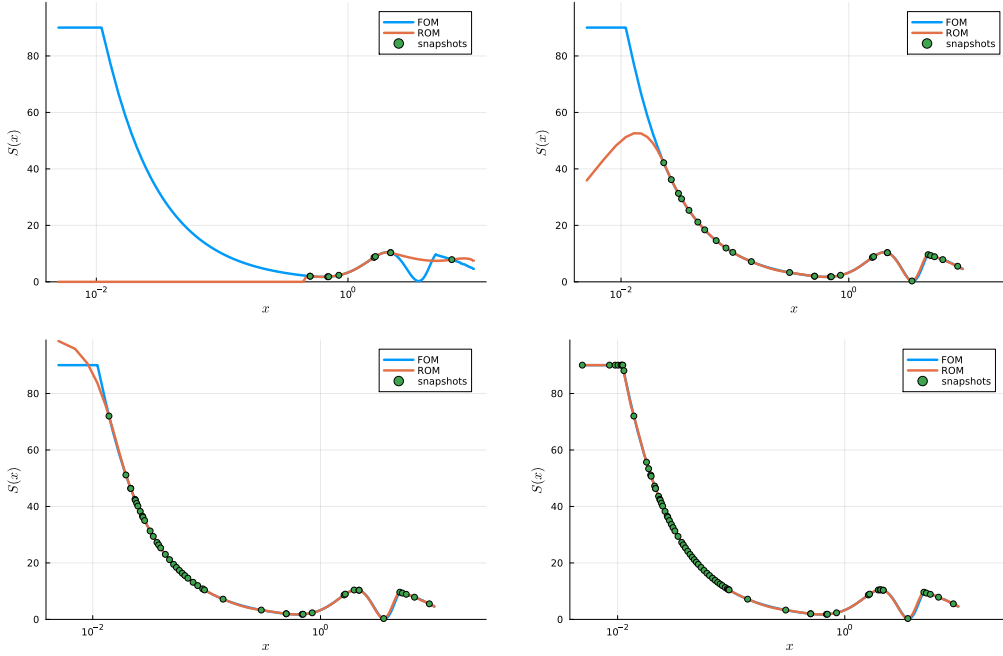


FIG. 4.1. **Illustration of the algorithm adaptive behavior (example #1).** The four plots illustrate the algorithm at iteration  $k = 1, 20, 40, 60$ . They display in blue the score function  $S^*$  called full order model (FOM) (the rare event being  $S^*(X) \geq 90$ ), in red the reduced score approximation  $S^{(k)}$  called reduced order model (ROM), in green the snapshots  $(X_1, \dots, X_k)$  used to build the reduced score.

the error bound is set to two times the absolute value of the true error  $E(x) = 2|S(x) - S^*(x)|$ .

Using spline interpolation is not recommended in practice, as it easily creates wide oscillations that can artificially create spurious high score areas that mislead the sampling routines. It is nonetheless an interesting benchmark for difficult situations where the reduced modeling produces very sub-optimal approximations in some areas.

The score function, depicted in Figure 4.1, is defined for  $x \in \mathcal{X} = \mathbb{R}_+$  by  $S^*(x) = |\Psi^*(x)|$  with

$$\Psi^*(x) = \begin{cases} l_{\max} & \text{if } x \leq 1/l_{\max} \\ 1/x + f(x) & \text{else} \end{cases},$$

where  $f(x) = \alpha (\mathbf{1}_{x_0 \leq x < x_1} \sin(x - x_0)^2 + \mathbf{1}_{x_1 \leq x} (\sin(x_1 - x_0)^2 - \beta(x - x_1)))$ , given parameters  $x_1 > x_0 > 1/l_{\max}$  and  $\alpha, \beta > 0$ . The reference distribution  $\pi$  is a log-normal distribution such that  $\ln(x) \sim \mathcal{N}(1.5, 1.5)$ . Note that here the rare event is  $\{S^* \geq l_{\max}\}$ . The parameter are set to  $x_0 = 0.5$ ,  $x_1 = 5$ ,  $\alpha = 15$ ,  $\beta = 0.1$  and  $l_{\max} = 90$ , corresponding to the rare event probability (given in closed-form) and given in the table above.

The initial surrogate is built using the 10 snapshots drawn according to the reference density.

Figure 4.1 shows a typical evolution of the spline-based score approximation as well as the selected snapshots over the iterations of the algorithm.

**4.1.2. Example #2: reduced basis approximation of a PDE score.** In our second numerical experiment, the score function is defined as  $S^*(x) = \|\Psi^*(x)\|_p$  with  $p = 1$  or  $p = \infty$ ,  $\Psi^*(x)$  being the high-fidelity numerical approximation of the solution of a bi-dimensional PDE parametrized by a  $q$ -dimensional vector  $x = (x_1 \cdots x_q)^\top \in \mathcal{X} = \mathbb{R}_+^q$ . The chosen PDE corresponds to the well-studied *thermal block problem*, which models heat diffusion in an heterogeneous media. The reference distribution for  $x$  will be a log-normal distribution with independent components with  $\ln(x_i) \sim \mathcal{N}(1.5, 1.5)$ , for  $i = 1, \dots, q$ . The probability of interest will correspond to the occurrence of the temperature exceeding the given threshold  $l_{\max}$  either in average (examples #1, #2a #2c) or locally (example #2b) over the spatial domain  $\Omega = [0, 1]^2$ . PDE solutions will be approximated using the popular reduced basis (RB) method.

We begin by describing the parametric PDE and the high-fidelity numerical approximation of its solution. The temperature function  $\Psi$  over  $\Omega$  is ruled by the elliptic problem

$$\begin{cases} -\nabla \cdot (\kappa(\cdot; x) \nabla \Psi) = 1 & \text{in } \Omega \\ \Psi = 0 & \text{on } \partial\Omega \end{cases},$$

where the function  $\kappa(\cdot; x) : \Omega \rightarrow \mathbb{R}_+$  represents a parametric diffusion coefficient function. This PDE models the stationary temperature field  $\Psi$  of an heterogeneous medium which is uniformly heated in the domain and thermostatted at the boundary.

The high-fidelity approximation  $\Psi^*(x)$  of the solution of this diffusion problem entails the solution of a variational problem. Precisely, let  $\mathcal{V}$  denote the Hilbert space  $\mathcal{H}_0^1(\Omega)$  and  $\|\nu\|_{\mathcal{V}} = \|\nabla \nu\|_{L^2}$  the norm induced by the inner product defined over  $\mathcal{V}$ . For any  $x \in \mathcal{X}$ , we define the solution  $\Psi^*(x) \in \mathcal{V}$  of the variational problem  $a(\Psi(x), \nu; x) = \int_{\Omega} \nu d\Omega$ ,  $\forall \nu \in \mathcal{V}$ , where  $a : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}$  is the bilinear form  $a(u, \nu; x) = \int_{\Omega} \kappa(\cdot; x) \nabla u \cdot \nabla \nu d\Omega$ . We consider the piecewise constant diffusion coefficient  $\kappa(\cdot; x) = \sum_{i=1}^q \mathbf{1}_{\Omega_i}(\cdot) x_i$  with the partition  $\Omega = \cup_{i=1}^q \Omega_i$  and  $\Omega_i$ 's being non-overlapping squares. Setting the parameter space  $\mathcal{X} = (\mathbb{R}_+^*)^q$  implies that the bilinear form is strongly coercive, and the Lax-Milgram lemma ensures the variational problem has a unique solution.

Consider next a finite  $h$ -dimensional subspace  $\mathcal{V}_h \subset \mathcal{V}$ . The high-fidelity approximation  $\Psi^*(x) \in \mathcal{V}_h$  of the solution  $\Psi^*(x)$  is defined for any  $x \in \mathcal{X}$  by

$$\Psi^*(x) = \arg \min_{\nu \in \mathcal{V}_h} \|\Psi^*(x) - \nu\|^2,$$

with the energy norm induced by the bilinear form  $\|\nu\| = \sqrt{a(\nu, \nu; x)}$ . The high-fidelity approximation  $\Psi^*(x)$  is obtained in practice by solving a large system of  $h$  linear equations. [22]. For the function *score* defined as the  $L^p$ -norm, the rare event  $\{S^*(X) = \|\Psi^*(X)\|_{L^p} > l_{\max}\}$  corresponds to reaching a critical temperature  $l_{\max}$ , e.g. a melting temperature, on average in the domain ( $p = 1$ ) or in some point of the domain ( $p = \infty$ ).

We are interested in the construction of an inexpensive approximation to the set of high-fidelity solutions

$$\mathcal{M} = \{\Psi^*(x) \in \mathcal{V}_h : x \in \mathcal{X}\},$$

by evaluating only some of its elements. By selecting a well-chosen set of  $K$  snapshots  $\{\Psi^*(x_i)\}_{i=1}^K$ , the RB framework approximates  $\mathcal{M}$  in a well-chosen linear subspace  $\mathcal{V}_K \subset \mathcal{V}_h$  of much lower dimension  $K \ll h$ . Using this framework yields *i*)



an approximation  $\Psi$  of the high-fidelity solution  $\Psi^*$  which is used to compute the approximation  $S(x) = \|\Psi(x)\|_{L^p}$  of the score  $S^*(x)$ , *ii*) a so-called *a posteriori* error bound on  $\|\Psi^*(x) - \Psi(x)\|_{\mathcal{V}_h}$  denoted by  $\Delta_\Psi(x)$ . The high-fidelity solution, the RB approximation and the *a posteriori* error bounds are computed via the *pymor* Python<sup>®</sup> toolbox [19] available at <https://github.com/pymor/pymor>.

To use the ARMS algorithm, we need to design an error bound  $E(x)$  majoring the absolute error on the score  $|S^*(x) - S(x)|$ , relying on the *a posteriori* error bound on  $\|\Psi^*(x) - \Psi(x)\|_{\mathcal{V}_h}$ . In the case of  $p = 1$ , the error bound satisfies

$$\begin{aligned} |S^*(x) - S(x)| &= \left| \|\Psi^*(x)\|_{L^1} - \|\Psi(x)\|_{L^1} \right| \\ &\leq \|\Psi^*(x) - \Psi(x)\|_{L^1} \leq \|\Psi^*(x) - \Psi(x)\|_{L^2} \leq \pi^{-1} \|\Psi^*(x) - \Psi(x)\|_{\mathcal{V}_h} \\ &\leq \pi^{-1} \Delta_\Psi(x) = E(x), \end{aligned}$$

where we have used the triangular inequality, the Cauchy-Schwarz, the Poincaré inequality with the constant  $\pi^{-1}$  on the unit square domain and the *a posteriori* error bound. As the finite element basis is an interpolating one, we will assume  $h$  sufficiently large so that we can compute the  $L^1$ -norm as the average of the finite element coefficients. In the case of  $p = \infty$ , except in the one-dimensional case, one can not in general rigorously upper bound  $\|\Psi^*(x) - \Psi(x)\|_{L^\infty}$  by  $\|\Psi^*(x) - \Psi(x)\|_{\mathcal{V}_h}$ . We will nevertheless use for this experiment the rough approximation  $|S^*(x) - S(x)| \approx \|\Psi^*(x) - \Psi(x)\|_{\mathcal{V}_h} \leq \Delta_\Psi(x) = E(x)$ , possibly violating assumption (1.2).

Let us comment on two important features of the RB approximation in regards to our rare event simulation framework. First, we emphasize that our algorithm enriches the RB approximation space using a generalization of the popular *weak greedy* algorithm. Indeed, the latter selects at each iteration the snapshot in the solution set  $\mathcal{M}$  which possess the worst *a posteriori* error bound for the current RB approximation. This setup corresponding in our algorithm to the degenerated case where the learning parameter is set to  $\tau = +\infty$ . At the opposite, the setup where  $\tau = 0$  will correspond to sampling according to the proposal, independently of the approximation error. In the general setup where  $0 < \tau < \infty$ , snapshot selection will balance these two cases, as the RB will be enriched with a sample drawn according to a weighted version of the proposal distribution, with weights proportional to the function  $e^{\tau E^{(k)}}$ .

Second, as in standard RB approaches, we here extensively rely on the affine parametric dependence of the RB approximation by splitting the assembly of the reduced matrices and vectors in an *off-line* and *on-line* phases. The former, which is performed once for all at each RB update, entails the computation of all the  $h$ -dependent and  $x$ -independent structures. It requires  $\mathcal{O}(Kh^2 + K^2h)$  operations. Besides, one must add the evaluation of a snapshot *per se*, which is  $\mathcal{O}(h^3)$ . In the latter, on-line phase, for a given parameter  $x \in \mathcal{X}$ , we assemble and solve the RB system and provide the error bound in  $\mathcal{O}(K^3)$ , with a cost depending only on  $K$ . For further details, we refer the reader to the introductory book on RB [22].

As there is no analytic value for the rare event probabilities sought, we perform standard AMS calculations with a score based on this high-fidelity approximation  $\Psi^*(x)$  and using  $N = 1e3$  particles. Averaging 10 runs, we obtain, after intensive computation, an empirical Monte Carlo approximation for  $p^* = \pi(\mathbf{1}_{S^* > l_{\max}})$  which is presented in the table above.

Figure 4.2 and 4.3 show typical evolutions along the ARMS algorithm iterations for example #2a and #2c. We observe that in the first iterations of the algorithm, the RB approximation of the high-fidelity solution of the PDE is too rough to char-

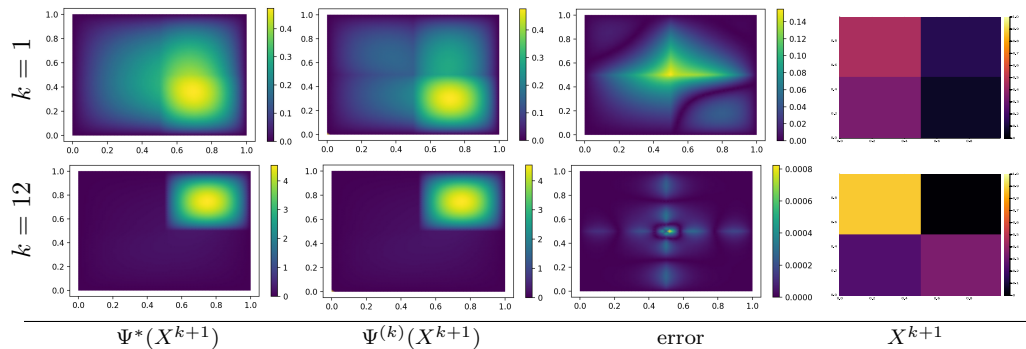


FIG. 4.2. **Illustration of the algorithm adaptive behavior (example #2a).** The figure displays the high-fidelity solution of the PDE (left), its reduced basis approximation (middle left) and the related absolute error (middle right), evaluated at the snapshot  $X^{k+1}$  (right) sampled at iteration  $k$  by our ARMS algorithm.

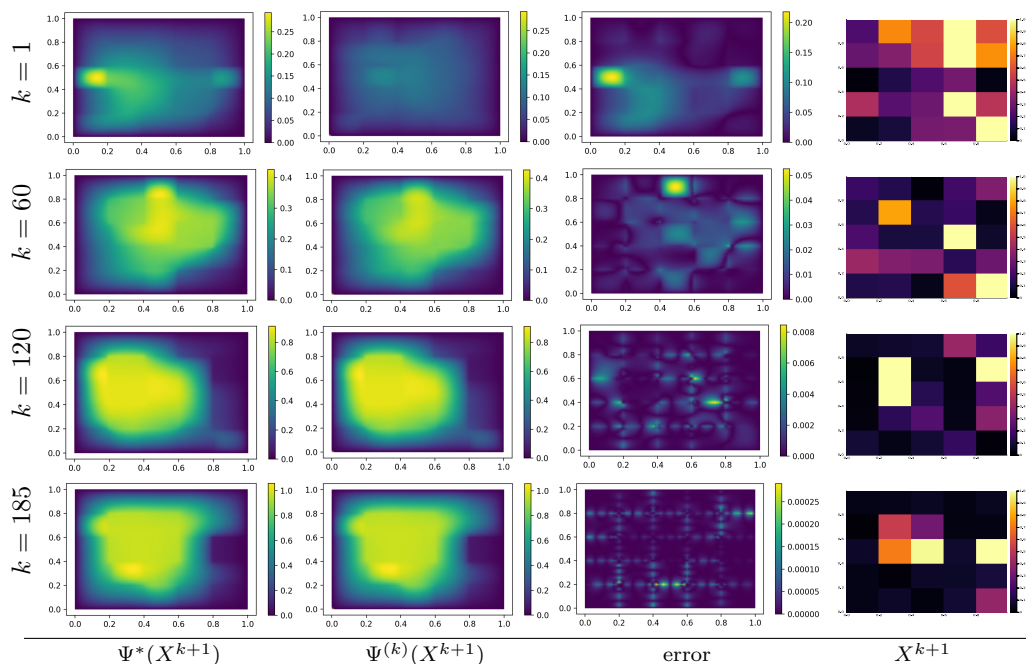


FIG. 4.3. **Illustration of the algorithm adaptive behavior (example #2c).** The figure displays the high-fidelity solution of the PDE (left), its reduced basis approximation (middle left) and the related absolute error (middle right), evaluated at the snapshot  $X^{k+1}$  (right) sampled at iteration  $k$  by our ARMS algorithm.

acterize the rare event of interest. While the number of iterations increase, RB gain in accuracy, enabling importance sampling of the rare event.

In the context of example #2a, a typical rare event sampled by ARMS corresponds to a diffusion coefficient one decade lower than other diffusion coefficients. To provide a physical interpretation of this rare event, note that each of the 4 subdomains touches the zero border condition, so that no coefficient configuration can isolate a subdomain. Consequently, increasing the average temperature can only be achieved by slowing

down the heat (contributed by the homogeneous source) lost by the system, which corresponds to configurations with low diffusion coefficients. As these configurations lie in the tail of the reference distribution, such events are rare, and of these events, the least rare corresponds to a decrease in one of the 4 parameters.

In the context of example #2c, a typical rare event sampled by ARMS corresponds to setting the value of most diffusion coefficients touching the zero boundary condition, one decade lower than the value of other diffusion coefficients. The physical interpretation of this rare event is clear: isolating the domain boundary slows the loss of heat contributed by the source, and therefore induces an increase in the average temperature within the domain. Note that such a configuration is extremely rare (on the order of  $1e-14$ ) according to the reference distribution, as 11 coefficients out of 25 must be sampled in the tail of the reference distribution.

**4.2. ARMS setup & evaluation criteria.** We now detail the configuration of the ARMS algorithm. The mutation Markov Chain Monte Carlo transition is a Metropolis accept/reject applied on an Ornstein-Uhlenbeck proposal, with an adaptive variance keeping the average acceptance rate over particles in a reasonable range (see *e.g.* [15]).

We use a simple scheme to initialize and update the learning parameter  $\tau$ : in the early iterations of our algorithm ( $k \leq K_{j_0}$ ) we set the learning parameter to  $\tau^{(k)} = +\infty$  (which corresponds to the weak greedy algorithm setting in the RB framework); once the sampled snapshots have hit  $j_0$  times the rare event set, we set this parameter to  $\tau^{(k)} = 0$  for  $k > K_{j_0}$  (which corresponds to sampling the next snapshots according to the current proposals).

If not explicitly specified in the text, ARMS parameters is configured in accordance with the table below.

	$N$	$M/N$	$K$	$T$	$R$	$j_0$	$c$
example #1	[2.0e2, 8.0e4]	0.3	300	1e2	100	5	5.0e-2
example #2a	[2.5e2, 4.0e2]	0.1	300	3e1	50	2	3.0
example #2b	[1.0e2, 4.0e3]	0.3	200	3e1	100	5	5.0e-2
example #2c	[1.0e2, 2.0e2]	0.1	300	3e1	50	2	3.0

In the table above,  $N$  is the number of particles,  $M$  is the number of particles killed at each selection step,  $K$  is the snapshot budget,  $T$  is the number of Markov transitions at each mutation step,  $R$  is the number of independent ARMS runs for each problem,  $j_0$  is the number of hits before the estimator is triggered and  $c$  is the log worst-cost threshold.

We plot the median and quantiles related to the IS estimate  $\hat{p}_{IS}$  and to the AMS estimate  $\hat{p}_{AMS}$  as a function of the number of snapshots  $k$ . ARMS performance is quantified by evaluating the expected cost as a function of the relative root expected square error. More specifically, after  $k$  snapshots, the expected relative square error is approximated over runs by the empirical average

$$\frac{\mathbb{E}(\hat{p} - p^*)^2}{p^{*2}} \approx \frac{1}{R} \sum_{r=1}^R \frac{(\hat{p}_r - p^*)^2}{p^{*2}}, \quad (4.1)$$

where  $\hat{p}_r$  denotes, at the  $r$ -th run and after  $k$  snapshots, either AMS estimate  $\hat{p}_{AMS}$  given in (3.12), or the importance sampling estimate  $\hat{p}_{IS}$  given in (3.11) (with the convention that  $\hat{p}_{IS} = 0$  if  $H = 0$ ). After  $k$  snapshots, the expected cost (in units of

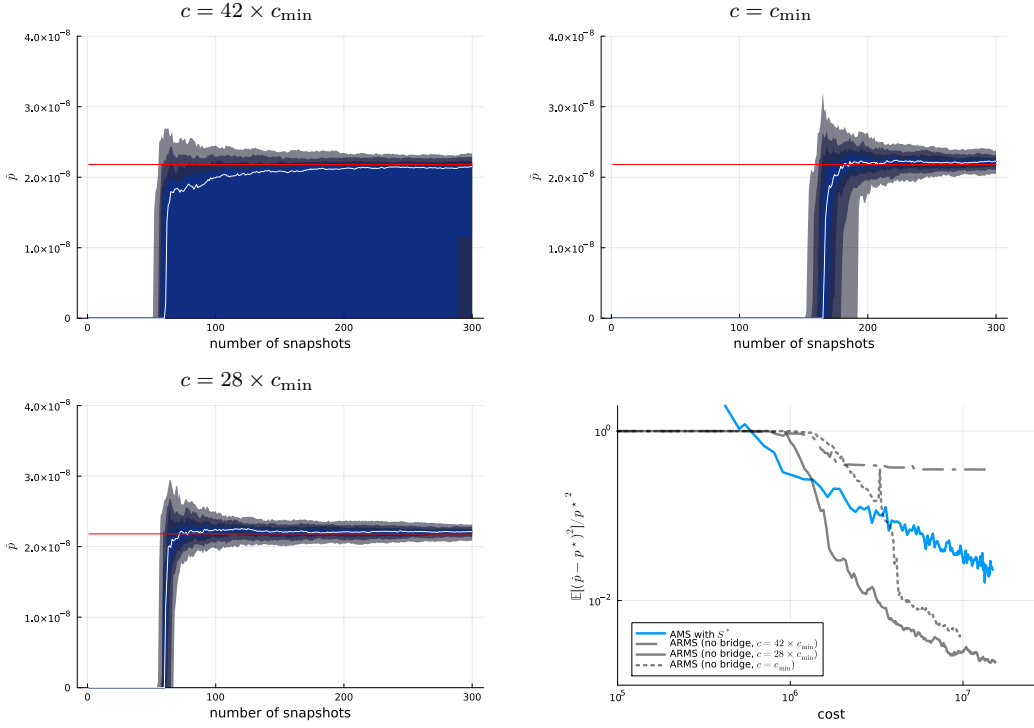


FIG. 4.4. **Influence of the entropic criterion threshold  $c$  (example #1).** Upper and lower left plots: median (in white) and quantiles (in shades of blue-grey) related to the IS estimate (3.11) as a function of the number of snapshots (equal to the number of iteration  $k$  of the algorithm) for  $N = 1e3$ , a reduced model gain of  $\times 25$ , and 3 different values of the worst log cost threshold  $c$ . The red line denotes the true rare event probability. Lower right plot: expected relative square error (4.1) as a function of the expected cost (4.2) (grey curves) for the IS estimate (3.11). Comparison to the performance of standard AMS with  $S^*$  and  $N \in [1e2, 1e4]$  (blue curve).

the cost of evaluating  $S^*$ ) is approximated over the runs by the empirical average

$$\text{cost}(k) \simeq k + \frac{1}{R} \sum_{r=1, k'=1}^{R, k} N_r^{(k')} \text{gain}^{(k')} \quad (4.2)$$

with  $\text{gain}^{(k')}$  representing the ratio of the cost of evaluating the reduced score  $S^{(k')}$  to the cost of the true score  $S^*$ , and where  $N_r^{(k')}$  represents the number of evaluations of the reduced score at the  $k'$ -th iteration and  $r$ -th execution of the ARMS algorithm.

Please note that for the first example,  $\text{gain}^{(k')}$  will be set artificially to emulate realistic cases in which the cost of evaluating the surrogate is indeed cheaper than that of the true model (see Figure 4.5).

We will compare the performance of ARMS with the empirical average of the estimates obtained by performing  $R$  standard AMS with the true score  $S^*$ .

**4.3. Results.** We begin by analyzing the results obtained with the toy model of example #1. Figure 4.4 evaluates the influence of the threshold  $c$  on the entropic criterion of Section 2.2. The three quantiles show the distribution of the IS estimate as a function of the number of snapshots for different values of the entropy criterion threshold  $c$ . As expected, we note (upper left) that a too large threshold produces over

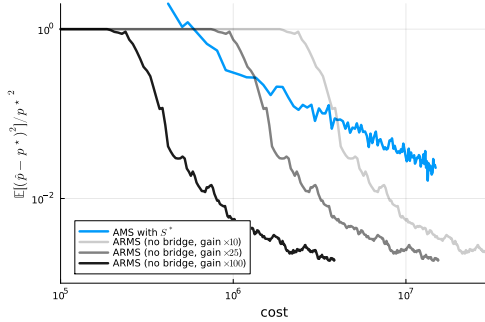


FIG. 4.5. **Influence of the reduced score gain (example #1).** Expected cost (4.2) as a function of the relative expected squared error (4.1) for the IS estimate (3.11), for different values of the ratio between the cost of evaluating  $S^{(k)}$  and  $S^*$  (grey curves) and for  $N = 1e3$  and the threshold  $c = 28 \times c_{\min}$ . Comparison to the performance of standard AMS with  $S^*$  and  $N \in [1e2, 1e4]$  (blue curve).

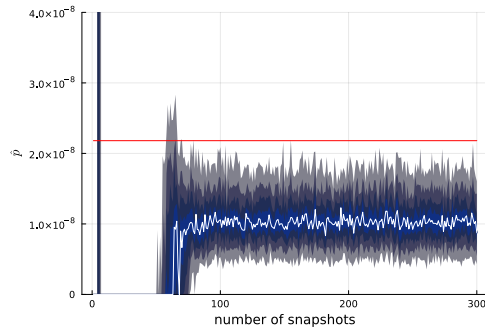


FIG. 4.6. **Bias of the reduced score based estimate (example #1).** Median (in white) and quantiles (in shades of blue-grey) related to the AMS estimate (3.12) as a function of the number of snapshots; the red line denotes the true rare event probability. Number of samples  $N = 1e3$ , entropic criterion set to  $c = 28 \times c_{\min}$ , no bridging, never stopping the reduced score update ( $\epsilon < 0$ ) and gain between surrogate and exact is assumed to be 25.

40% of erroneous IS estimates, associated with zero probability regardless of the number of reduced score updates. Although the remaining 60% of estimates nevertheless become accurate after around 60 updates of the reduced score, this demonstrates that increasing the critical level too quickly tends to concentrate particles around unimportant local minima. This results in sampling snapshots and refining the reduced score exclusively in regions of state space unrelated to the rare event of interest. To avoid this undesirable scenario, it is possible to set the threshold to its minimum value  $c_{\min}$  (upper right), which has the disadvantage of delaying the convergence of the IS estimate. Indeed, the plot shows convergence of all estimates towards a non-zero value close to  $p^*$ , but with a considerable delay, since in this case about a hundred more updates of the reduced score are required. A well-tuned value between these extrema (lower left) achieves convergence after about 70 snapshots, with a characterization of the distribution of IS estimates comparable to the safe distribution obtained with the minimum value.

This trend is confirmed by plotting (lower right) the relative root expected square error (4.1) as a function of the expected cost (4.2). In the case of a not-too-high threshold  $c_{\min}$  or  $28c_{\min}$ , given a certain precision (*i.e.*, after sampling enough snapshots),

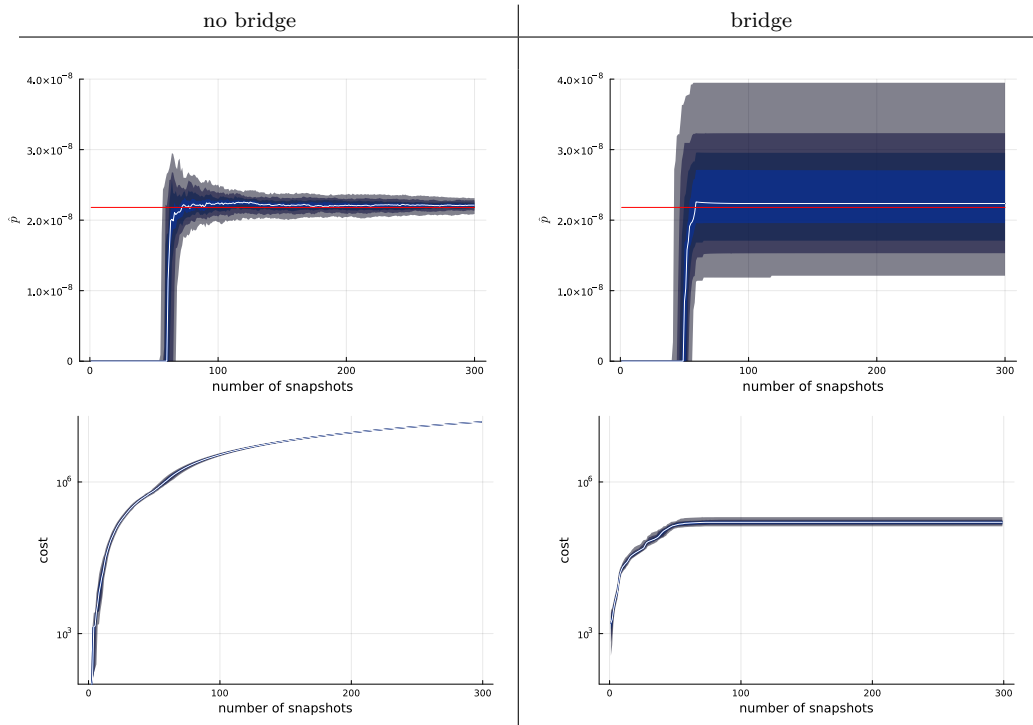


FIG. 4.7. **Influence of bridging (example #1).** Median (in white) and quantiles (in shades of blue-grey) related to the IS estimate (3.11) (above) and of the cost (4.2) (above), as a function of the number of snapshots when bridging simulation or not; the red line denotes the true rare event probability. Number of samples  $N = 1e3$ , entropic criterion set to  $c = 28 \times \epsilon_{\min}$ , stopping updates of the reduced score after convergence only in the case of bridging and gain between surrogate and exact is assumed to be 25. Note that, although imperceptible in the log-log plot at bottom right, the cost increases strictly with the number of snapshots.

we observe a gain in expected cost from around a decade onwards, which increases more and more. Clearly, a value of  $42c_{\min}$  is too high and do not achieve comparable accuracy, regardless of the number of snapshots.

Figure 4.5 illustrates how the gain in terms of expected cost over accuracy can be dramatic, as the cost of evaluating  $S^{(k)}$  becomes negligible compared to  $S^*$ .

Figure 4.6 shows that the AMS estimate (3.12) significantly underestimates the true probability, demonstrating that relying solely on reduced scores  $S^{(k)}$  (here spline approximations) of the true score  $S^*$  induces a significant bias. Indeed, in our specific example, the non-differentiability of  $S^*$  at the rare event boundary leads to high-frequency oscillations (which tend to decrease with the number of interpolation points, but never disappear completely) in the spline model. The consequence is that a significant proportion of particles in the rare event set  $\{x \in \mathcal{X} : S^*(x) \geq l_{\max}\}$  are to be excluded from the set  $\{x \in \mathcal{X} : S^{(k)}(x) \geq l_{\max}\}$ , leading to a negative bias. On the contrary, we observe that the IS estimate (3.11) is not biased, independently of the problem of convergence of  $S^{(k)}$  to  $S^*$ .

The evaluation above focused on the ARMS algorithm in which the simulation is restarted after each reduced score update and in which the stop criterion introduced

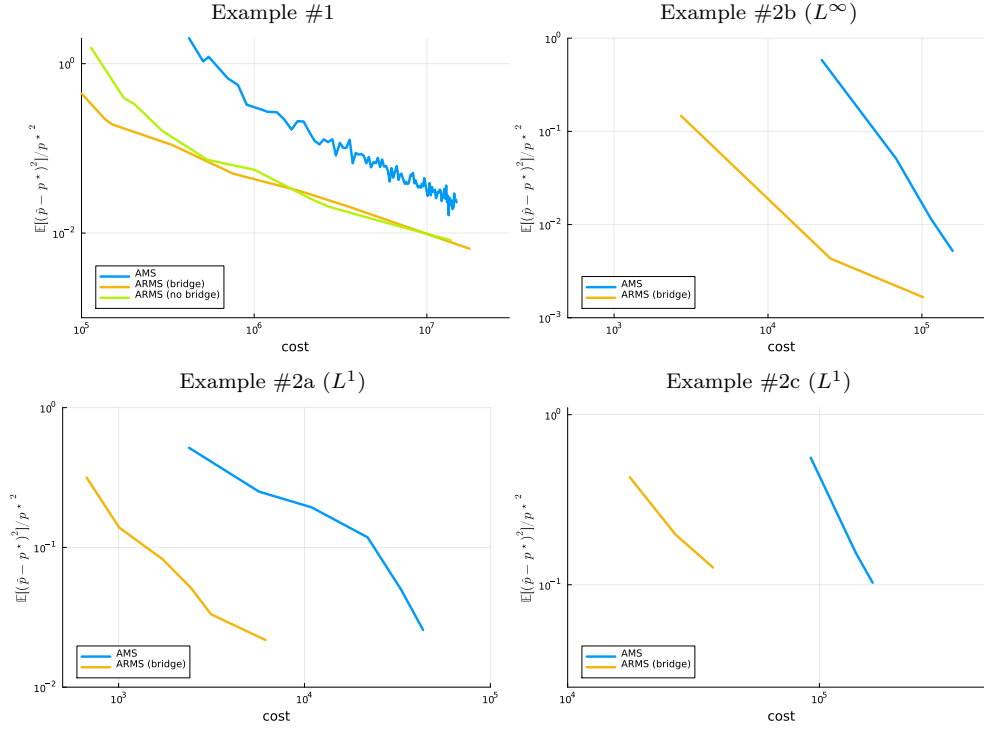


FIG. 4.8. **Performance of ARMS (example #1 & 2).** Relative expected square error (4.1) as a function of the expected cost (4.2) of the IS estimate (3.11) at the final iteration (once the snapshots budget  $K$  is exhausted), for a varying sample size  $N$ . Curves can be compared to standard AMS estimation using the score  $S^*$ , for a varying sample size  $N$ .

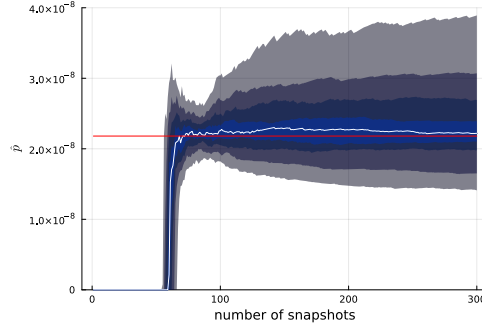


FIG. 4.9. **A bad idea: using the stopping criterion without bridging (example #1).** Median (in white) and quantiles (in shades of blue-grey) related to the IS estimate (3.11) as a function of the number of snapshots when stopping the reduced score updates ( $\epsilon = 0$ ) without bridging; the red line denotes the true rare event probability. Number of samples  $N = 1e3$ , entropic criterion set to  $c = 28 \times c_{\min}$  and gain between surrogate and exact is assumed to be 25.

in Section 2.5 is never triggered ( $\epsilon < 0$ ). Figure 4.7 extends the analysis by evaluating the influence of bridging the simulations between two iterations, as presented in Section 3.3. The entropic criterion is set to  $c = 28 \times c_{\min}$ . As shown in figure 4.7, the gain in computational cost brought about by bridging is offset by an increase in variance. One can remark that, when using the bridging procedure, and for a given

sample size, we cannot expect to increase the accuracy for the IS estimate (3.11) indefinitely, contrary to when the simulation is restarted after each reduced score update. Indeed, the estimate depends crucially on the variance of the normalization constant  $Z_{l^{(k)}}^{N,(k)}$  which generates a residual error (upper right box of Figure 4.7). In the case of bridging, this residual error will indeed propagate across the iterations according to (3.9) instead of being re-estimated by restarting a fresh simulation at each reduced score update. Nevertheless, the plot of expected squared error versus predicted cost in the upper left-hand panel of figure 4.8 shows that, in the case of small sample sizes yielding a relative square error above 20%, some saving of order of 2 is achieved by a bridging procedure compared to restarting the simulation, while both procedures are equivalent for larger sample sizes.

On the other hand, we notice that enabling the stopping criterion without bridging, is a bad idea as it induces a spread around the mean of the distribution of the IS estimate as illustrated in Figure 4.9. Indeed, similarly to the above explanation, the IS estimate will be affected by the apparent bias of the estimates  $Z_{l^{(k)}}^{N,(k)}$  in (3.11), all identical for indices  $k$  above some random time (determined by the stopping criterion). The phenomenon does not appear when the stopping criterion is off (Figure 4.7) and we restart the simulation as before and calculate a new normalization constant at each iteration. The stopping criterion is, in fact, simply used to marginally reduce the cost in the bridging procedure, without changing the outcome (not shown in plot).

We then move on to example #2, the PDE-based scores and their RB approximations. Overall, the trends of the algorithm for the RB approximation of a PDE score are very similar to those of the previous example. We thus omit the precise evaluation made for the previous example, and jump directly to the analysis of performances in Figure 4.8. We use an initial reduced score built on a 5-dimensional RB approximation. In the case of examples #2a and #2b, we observe that about 10 to 20 more snapshots (*i.e.*, ARMS iterations) are needed to reach convergence, with a median estimate around  $p^*$ . In the case of example #2c, about 170 to 200 additional snapshots are needed to reach a good accuracy of the reduced model and converge towards  $p^*$ . Indeed, in the latter situation, the parameter space has a much higher dimensionality (25 dimensions versus 4), and learning the reduced model for parameters in the (initially unknown) region where rare events occur is a difficult task performed progressively over the course of ARMS iterations.

In all examples, Figure 4.8 shows that the proposed ARMS method enables important computational savings as compared to the use of the true model alone. Savings are consistently around a factor 10 in all cases. As mentioned in section 4.1.2, in the case of the #2b example, ARMS uses error estimates that are not hard bounds for the  $L^\infty$  score. Nevertheless, it is surprising that ARMS manages to accurately estimate the desired probability and achieve performances comparable to those obtained with the  $L^1$  score.

**5. Conclusions and Perspectives.** This work proposes an adaptive algorithm called ARMS that reduces the cost of AMS simulation when the score function  $S^*$  is very expensive to compute.

The general idea of ARMS is to perform iterative importance sampling of a sequence of target distributions of the form  $\propto \mathbf{1}_{S^* > l^{(k)}} d\pi$  parametrized at iteration  $k$  by the level  $l^{(k)} \in \mathbb{R}$ , using a sequences of proposals of the form  $\mu_{l^{(k)}}^{(k)} \propto \mathbf{1}_{S^{(k)} > l^{(k)}} d\pi$ , much cheaper to simulate with AMS. As the algorithm is iterated, the proposals are adapted by adjusting the level  $l^{(k)}$  and refining the approximation of the reduced



score  $S^{(k)}$  using the history of score evaluations previously calculated for importance sampling.

As  $S^{(k)}$  gains in accuracy and  $l^{(k)}$  increases, the reduced score approximation should be refined using samples of the state space in regions closer and closer to the rare event of interest. However, an essential ingredient in achieving this desirable convergence is to determine a criteria for setting the first non-achievable (in the sense of Importance Sampling cost) level  $l^{(k)}$  by the accuracy of the current reduced score approximation  $S^{(k)}$ . In addition, one must ensure that the support of the proposal always includes the final rare event target. Justified by theoretical arguments, these two problems are addressed in ARMS using entropy constraints:  $l^{(k)}$  is chosen such that *i*) the relative entropy between a pessimistic rare event target at level  $l^{(k)}$  and the proposal corresponds to a given logarithmic cost of importance sampling; *ii*) the relative entropy between an optimistic final rare event target and the proposal remains finite. Pessimistic and optimistic rare event targets are constructed using a quantification of the worst-case error associated with the reduced score approximation.

A final point addressed by ARMS is the sampling of the proposal given by a new approximation of the reduced score  $S^{(k+1)}$ . In particular, the algorithm finds a (usually lower) level  $l_b(k+1)$  at which the updated proposal is dominated (inclusion in terms of support) by a well-chosen previous (carefully recorded) proposal, so that the updated proposal can be sampled by 'bridging' quickly from the recorded proposal in a single AMS simulation step. In addition, the updated proposal at level  $l_b(k+1)$  must satisfy entropy constraints *i*) and *ii*) of the previous paragraph.

In addition to demonstrating the unbiasedness of our estimator in an idealized setting, we evaluate in our numerical experiments the empirical convergence of the ARMS algorithm on a toy model and a rare-event problem based on a more realistic PDE. Our study confirms that entropy criteria play a crucial role in ARMS. In particular, as expected, too high a value of the logarithm of the cost in the entropy criterion *i*) can take ARMS into spurious regions of the state-space remote from the region of interest, resulting in a large variance of the IS estimator. Too small a value gives a more desirable variance but a slower rate of convergence. We show empirically that the approximation procedure using the entropy criterion *iii*) significantly reduces simulation cost (around a factor 10) while achieving a comparable squared error. In the same way in both experiments, we observe a squared error gain of the order of a decade for a given simulation cost. Our numerical experiments show that bridging is a good option for small computational budgets, with a saving of around 2 compared with restarting the simulation at each iteration. However, for larger computational budgets, both procedures appear to be equivalent in terms of cost multiplied by squared error.

The prospects arising from this work are numerous. These include answering open theoretical questions, including the consistency of ARMS, which is proven in the present paper only in a formal sense. No straightforward answers seem available, as we are confronted with a non-standard framework in which interacting particles in AMS simulations have distributions with varying supports which depend on random approximations of the reduced score. Future research directions also include extending ARMS to smooth proposal distributions, with the advantage of removing nesting and domination constraints and the possibility of being more flexible in the choice of error estimates that need not have strict error bounds.

## Appendix A. Sample Size for Importance Sampling.

### A.1. Generic quantification rules. Importance sampling of some generic dis-

tribution  $\eta \propto \gamma$  known up to a normalizing constant  $p = \gamma(\mathbf{1})$  using a proposal distribution  $\mu$  is based on the following identity

$$\gamma(\varphi) = p\eta(\varphi) = \mathbb{E}_\mu [f(X)\varphi(X)] \quad \forall \varphi : \mathcal{X} \rightarrow \mathbb{R}$$

where  $f(x) = \frac{d\gamma}{d\mu}(x)$  is computable in closed-form.

Unfortunately, importance sampling becomes quickly useless if the target distribution diverges too widely from the proposal, a typical phenomenon in high dimension. This can be readily seen by computing the relative variance of the importance sampling estimator for  $\varphi = \mathbf{1}$ :

$$\text{Var}_\mu(f(X))/p^2 = \eta\left(\frac{d\eta}{d\mu}\right) - 1 = e^{\text{Ent}_2(\eta|\mu)} - 1, \quad (\text{A.1})$$

where we have denoted  $\text{Ent}_2(\eta | \mu) = \ln \eta\left(\frac{d\eta}{d\mu}\right)$  the order 2 Rényi entropy (also known as  $\chi^2$ -divergence). The latter scales linearly in  $d$  if  $\mu$  and  $\eta$  are a  $d$ -fold product measure, generating a variance exponential with  $d$ .

As is thoroughly argued in [8] (see also [7]), the variance  $(e^{\text{Ent}_2(\eta|\mu)} - 1)$  is in general a pessimistic quantification of importance sampling, whose cost is better estimated using a  $\mathbb{L}^1$ -error, or a success probability within a tolerance precision. For a target  $\eta$  with proposal  $\mu$ , it has been proven in [8] that the cost of importance sampling – in terms of sample size  $K$  – is roughly given (up to a reasonable tail condition on the log-density  $\ln \frac{d\eta}{d\mu}$ ), by the exponential of the relative entropy between  $\eta$  and  $\mu$ :  $K \simeq e^{\text{Ent}(\eta|\mu)}$ .

REMARK A.1. *Jensen inequality*

$$\text{Ent}_2(\eta | \mu) = \ln \eta\left(\frac{\eta}{\mu}\right) \geq \eta\left(\ln \frac{\eta}{\mu}\right) = \text{Ent}(\eta | \mu),$$

shows that indeed, variance is a more pessimistic estimator of the required sample size for importance sampling.

**A.2. Particular case of interest.** We will particularize the above discussion to the family of distributions studied in our paper: distributions of propositions of the form  $\mu \propto \mathbf{1}_B d\pi$  and a target  $\eta \propto \mathbf{1}_A d\pi$  such that  $A \subseteq B$ .

In accordance with the importance sampling identity, the normalization  $p$  is estimated without bias using the random variable  $f(X) = \pi(B) \mathbf{1}_A(X)$  where  $X \sim \mu$ .

In this remarkable case, it turns out that the relative variance and the exponential of the relative entropy reduces to the same quantity. In other words, we reach the limit of Jensen’s inequality in the Remark A.1

$$\text{Ent}_2(\eta | \mu) = \text{Ent}(\eta | \mu) = \ln \frac{\pi(B)}{\pi(A)}.$$

Therefore, in our study, we will refer indifferently to the exponential of the relative entropy  $e^{\text{Ent}(\eta|\mu)}$  or to the relative variance  $e^{\text{Ent}_2(\eta|\mu)}$  to designate the *importance sampling cost* of the target  $\eta$  by the proposal  $\mu$ .

**Appendix B. Consistency of the Estimator.** This section analyses the idealized algorithm obtained with  $N = +\infty$  many particles (clones) in ARMS;

For simplicity (but the general case is similar), we study the estimator (2.9) in the case of a test function  $\varphi = \mathbf{1}$ . We thus consider the rare event probability estimator:  $\hat{p}_l^H = \hat{\gamma}_l^H(\mathbf{1})$ .

PROPOSITION B.1. *The estimator  $\hat{p}_l^H$  is unbiased, i.e.,  $\mathbb{E} \hat{p}_l^H = p_l^*$ , and its relative variance is given by*

$$\text{Var} [\hat{p}_l^H] / p_l^{*2} = \frac{1}{H^2} \sum_{k=K_{j_0}+1}^{K_{j_0}+H} \left( \frac{\mathbb{E}(Z_{l^{(k)}}^{(k)})}{p_l^*} - 1 \right), \quad (\text{B.1})$$

where the random variable  $Z_{l^{(k)}}^{(k)}$  is a function of the past snapshots.

*Proof.* We first show that for any test function  $\varphi$  the random sequence  $H \mapsto M_H^\varphi \stackrel{\text{def}}{=} H(\hat{\gamma}_l^H(\varphi) - \gamma_l^*(\varphi))$ , is a martingale with respect to the filtration generated by the snapshots. Indeed, defining this filtration by  $\mathcal{F}_H = \sigma(X_1, \dots, X_{K_{j_0}+1+H})$ , it follows from the definition (2.9) that

$$\begin{aligned} \mathbb{E}[M_H^\varphi | \mathcal{F}_{H-1}] &= \mathbb{E} \left[ \sum_{k=K_{j_0}+1}^{K_{j_0}+H} f_l^{(k)}(X_{k+1}) \varphi(X_{k+1}) - H \gamma_l^*(\varphi) | \mathcal{F}_{H-1} \right], \\ &= M_{H-1}^\varphi + \mathbb{E}[f_l^{(K_{j_0}+H)}(X_{K_{j_0}+H+1}) \varphi(X_{K_{j_0}+H+1}) - \gamma_l^*(\varphi) | \mathcal{F}_{H-1}], \\ &= M_{H-1}^\varphi + \pi(\varphi \mathbf{1}_{S^* > l}) - \gamma_l^*(\varphi) = M_{H-1}^\varphi. \end{aligned}$$

Now, since  $M_H^\varphi$  is a martingale, the estimator is unbiased  $\mathbb{E} \hat{\gamma}_l^H(\varphi) = \gamma_l^*(\varphi)$ , and its variance is given by its quadratic variations given by the sum over  $k$  of the local variances. In particular, using (A.1), the variance of the rare event probability estimator is then

$$\text{Var} [\hat{p}_l^H] / p_l^{*2} = \frac{1}{H^2} \sum_{k=K_{j_0}+1}^{K_{j_0}+H} \left( \mathbb{E}(e^{\text{Ent}_2(\eta_l^* | \mu_{l^{(k)}}^{(k)})}) - 1 \right),$$

which for  $\eta_l^* \propto \mathbf{1}_{S^* > l} d\pi$  and  $\mu_{l^{(k)}}^{(k)} \propto \mathbf{1}_{S^{(k)} > l^{(k)}} d\pi$  yields (B.1).  $\square$

## REFERENCES

- [1] Siu-Kui Au and James L Beck. Estimation of small failure probabilities in high dimensions by subset simulation. *Probabilistic engineering mechanics*, 16(4):263–277, 2001.
- [2] Mathieu Balesdent, Jerome Morio, and Julien Marzat. Kriging-based adaptive importance sampling algorithms for rare event estimation. *Structural Safety*, 44:1–10, 2013.
- [3] Julien Bect, Ling Li, and Emmanuel Vazquez. Bayesian subset simulation. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):762–786, 2017.
- [4] Claire Cannamela, Josselin Garnier, and Bertrand Iooss. Controlled stratification for quantile estimation. *The Annals of Applied Statistics*, 2(4):1554–1580, 2008.
- [5] F. Cérou, P. Del Moral, T. Furon, and A. Guyader. Sequential Monte Carlo for rare event estimation. *Stat. Comput.*, 22(3):795–808, 2012.
- [6] Frédéric Cérou and Arnaud Guyader. Fluctuation analysis of adaptive multilevel splitting. *Annals of Applied Probability*, 26(6):3319–3380, December 2016.
- [7] Frédéric Cérou, Patrick Héas, and Mathias Rousset. Entropy minimizing distributions are worst-case optimal importance proposals. *arXiv preprint arXiv:2212.04292*, 2022.
- [8] Sourav Chatterjee and Persi Diaconis. The sample size required in importance sampling. *The Annals of Applied Probability*, 28(2):1099–1135, 2018.
- [9] Peng Chen and Alfio Quarteroni. Accurate and efficient evaluation of failure probability for partial differential equations with random input data. *Computer Methods in Applied Mechanics and Engineering*, 267:233–260, 2013.

- [10] Pieter-Tjerk De Boer, Dirk P Kroese, Shie Mannor, and Reuven Y Rubinstein. A tutorial on the cross-entropy method. Annals of operations research, 134(1):19–67, 2005.
- [11] Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential monte carlo samplers. Journal of the Royal Statistical Society Series B: Statistical Methodology, 68(3):411–436, 2006.
- [12] Benjamin Echard, Nicolas Gayton, and Maurice Lemaire. AK-MCS: an active learning reliability method combining kriging and monte carlo simulation. Structural Safety, 33(2):145–154, 2011.
- [13] Daniel Elfverson, Robert Scheichl, Simon Weissmann, and F Alejandro DiazDelaO. Adaptive multilevel subset simulation with selective refinement. arXiv preprint arXiv:2208.05392, 2022.
- [14] Laurent Gallimard. Adaptive reduced basis strategy for rare-event simulations. International Journal for Numerical Methods in Engineering, 120(3):283–302, 2019.
- [15] Paul H Garthwaite, Yanan Fan, and Scott A Sisson. Adaptive optimal scaling of metropolis–hastings algorithms using the robbins–monro process. Communications in Statistics-Theory and Methods, 45(17):5098–5111, 2016.
- [16] Patrick Héas. Selecting reduced models in the cross-entropy method. SIAM/ASA Journal on Uncertainty Quantification, 8(2):511–538, 2020.
- [17] Matthias Heinkenschloss, Boris Kramer, and Timur Takhtaganov. Adaptive reduced-order model construction for conditional value-at-risk estimation. SIAM/ASA Journal on Uncertainty Quantification, 8(2):668–692, 2020.
- [18] Chunyan Ling, Zhenzhou Lu, Kaixuan Feng, and Xiaobo Zhang. A coupled subset simulation and active learning kriging reliability analysis method for rare failure events. Structural and Multidisciplinary Optimization, 60(6):2325–2341, 2019.
- [19] René Milk, Stephan Rave, and Felix Schindler. pymor–generic algorithms and interfaces for model order reduction. SIAM Journal on Scientific Computing, 38(5):S194–S216, 2016.
- [20] Maliki Moustapha, Stefano Marelli, and Bruno Sudret. Active learning for structural reliability: Survey, general framework and benchmark. Structural Safety, 96:102174, 2022.
- [21] Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. Siam Review, 60(3):550–591, 2018.
- [22] Alfio Quarteroni, Andrea Manzoni, and Federico Negri. Reduced basis methods for partial differential equations: an introduction, volume 92. Springer, 2015.
- [23] Nassim Razaaly, Daan Crommelin, and Pietro Marco Congedo. Efficient estimation of extreme quantiles using adaptive kriging and importance sampling. International Journal for Numerical Methods in Engineering, 121(9):2086–2105, 2020.
- [24] Elisabeth Ullmann and Iason Papaioannou. Multilevel estimation of rare events. SIAM/ASA Journal on Uncertainty Quantification, 3(1):922–953, 2015.
- [25] Felipe Uribe, Iason Papaioannou, Youssef M Marzouk, and Daniel Straub. Cross-entropy-based importance sampling with failure-informed dimension reduction for rare event simulation. SIAM/ASA Journal on Uncertainty Quantification, 9(2):818–847, 2021.
- [26] Fabian Wagner, Jonas Latz, Iason Papaioannou, and Elisabeth Ullmann. Multilevel sequential importance sampling for rare event estimation. SIAM Journal on Scientific Computing, 42(4):A2062–A2087, 2020.
- [27] Jize Zhang and Alexandros A Taflanidis. Adaptive kriging stochastic sampling and density approximation and its application to rare-event estimation. ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering, 4(3):04018021, 2018.