



HAL
open science

Adaptive reduced tempering for bayesian inverse problems and rare event simulation

Frédéric Cérou, Patrick Héas, Mathias Rousset

► **To cite this version:**

Frédéric Cérou, Patrick Héas, Mathias Rousset. Adaptive reduced tempering for bayesian inverse problems and rare event simulation. 2024. hal-04764656

HAL Id: hal-04764656

<https://inria.hal.science/hal-04764656v1>

Preprint submitted on 4 Nov 2024

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 TEMPERING FOR BAYESIAN INVERSE PROBLEMS AND RARE EVENT SIMULATION

FRÉDÉRIC CÉROU*, PATRICK HÉAS*, MATHIAS ROUSSET*

Abstract. This work proposes an adaptive sequential Monte Carlo sampling algorithm for solving inverse Bayesian problems in a context where a (costly) likelihood evaluation can be approximated by a surrogate, constructed from previous evaluations of the true likelihood. A rough error estimation of the obtained surrogates is required. The method is based on an adaptive sequential Monte-Carlo (SMC) simulation that jointly adapts the likelihood approximations and a standard tempering scheme of the target posterior distribution. This algorithm is well-suited to cases where the posterior is concentrated in a rare and unknown region of the prior. It is also suitable for solving low-temperature and rare-event simulation problems. The main contribution is to propose an entropy criteria that associates to the accuracy of the current surrogate a maximum inverse temperature for the likelihood approximation. The latter is used to sample a so-called snapshot, perform an exact likelihood evaluation, and update the surrogate and its error quantification. Some consistency results are presented in an idealized framework of the proposed algorithm. Our numerical experiments use in particular a reduced basis approach to construct approximate parametric solutions of a partially observed solution of an elliptic Partial Differential Equation. They demonstrate the convergence of the algorithm and show a significant cost reduction (close to a factor 10) for comparable accuracy.

Key words. Bayesian sampling, reduced modeling, importance sampling, sequential Monte Carlo, adaptive tempering, inverse problems, low-temperature simulation, rare events, relative entropy.

1. Introduction.

1.1. Context and problem. We consider the classical problem of designing an algorithmic Monte Carlo procedure that samples according to a target probability distribution $\eta_{\beta_\infty}^*$ (for one specific $\beta_\infty \in \mathbb{R}$) defined explicitly up to a normalization as an element of the family

$$\eta_\beta^* \stackrel{\text{def}}{=} \frac{1}{Z_\beta^*} e^{\beta S^*} d\pi, \quad \beta \in \mathbb{R}. \quad (1.1)$$

In the above, the quantity $Z_\beta^* \stackrel{\text{def}}{=} \pi(e^{\beta S^*})$ denotes the associated unknown normalization which typically must also be computed by the Monte Carlo procedure, π denotes a given, easily simulable, reference probability distribution over a state space \mathcal{X} , and $S^* : \mathcal{X} \rightarrow \mathbb{R}$, denotes a given yet expensive to compute real-valued *score* function.

The following two contexts will be used as the main motivations. In both cases, the score function is of the form

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

where *score* is an easily computable function of physical interest depending on some additional parameters (like observed data), and $\Psi^*(x)$ is the outcome of a numerical computation of a complex physical system parametrized by the variable $x \in \mathcal{X}$.

- In the Bayesian formulation of inverse problems, $\beta_\infty = 1$ and $\eta_{\beta_\infty}^*$ is the posterior distribution of some uncertainty parameters $x \in \mathcal{X}$ of a partially observed physical system $\Psi^*(x)$. π models the prior distribution of the uncertainty parameters and the weight $\propto e^{S^*}$ is the likelihood of those parameters given some partial observations. As usual, one wants to Monte Carlo sample and then approximate expectations according to this posterior.

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

- In rare events problems, one wants to estimate the small probability $p^* \stackrel{\text{def}}{=} \pi(\{S^* \geq 1\})$ corresponding to a large prescribed level $\ell \in \mathbb{R}$, with a score defined as $S^*(x) = 1 - \max(\ell - \Psi^*(x), 0)/\ell$. As $\lim_{\beta_\infty \rightarrow \infty} \pi(e^{\beta_\infty S^*}) = p^*$, the estimation of the rare event probability turns out to be the estimation of the normalization constant $Z_{\beta_\infty}^*$ as $\beta_\infty \rightarrow \infty$, *i.e.*, as the temperature $1/\beta_\infty$ tends to zero.

In this work, we will be interested in scenarios in which the practitioner has to face two types of difficulty. The first difficulty is classical and arises when β increases (or when the measurement noise becomes small in a Bayesian context), yielding a low-temperature problem, similar to global optimization. It confronts the practitioner with a computational challenge: the density defining η_β^* concentrates the majority of the probability mass in specific but possibly unknown areas of \mathcal{X} , which are roughly described by the maxima of S^* . A direct sampling with a Monte Carlo simulation of π is often time-consuming or even infeasible, as it can require a very large sample size. For instance to obtain a reasonable estimation variance, the sample size should be of order of $1/p^*$ in the rare event simulation context, which is obviously prohibitive. A very popular general strategy to simulate a sample according to η_β^* is to resort to a *Sequential Monte Carlo* (SMC) strategy (see [16]) that starts with a Monte Carlo sample of particles distributed according to π and then sequentially samples η_β^* for increasing values of β by resorting to a combination of Importance Sampling (IS), resampling (selection) of weighted samples and mutations of particles based on suitable Markov Chain Monte Carlo transitions. More specifically, we will be interested in the adaptive variant of tempering [3].

The second difficulty appears when the numerical evaluation of Ψ^* is extremely expensive; so expensive that the number of evaluations required in, say, a usual SMC approach becomes prohibitive. In order to circumvent this issue, we assume that the physical model Ψ^* can be approximated by a *reduced model* denoted Ψ ; the cost of evaluating Ψ being small as compared to one evaluation of the true score with Ψ^* . The reduced model Ψ 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 function $E(x)$ satisfying an approximate pointwise estimate of the form

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

We underline that the pointwise estimate (1.2) need not to be a very precise estimate on the error, but only a rough estimate revealing the trend of the error.

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 Ψ^* evaluated for each of those snapshots. The output is the reduced model Ψ . A reduced model can be constructed by various, potentially very different, means. In this work, we will consider any reduced modeling procedure and any approximate error quantification of the form (1.2).

When $\Psi^*(x)$ is the solution of a Partial Differential Equation (PDE) parametrized by the uncertainty variable x , a model reduction procedure particularly suited to this context is *reduced basis* [29].

The associated error estimate in (1.2) is updated simultaneously, and for reasons that will be explained in Section 2.2, we additionally require that the error $E^{(k)}$ vanishes on snapshot points $X^{(k)}$ already evaluated at iteration k :

$$E^{(X_1, \dots, X_k)}(X^{(j)}) = 0, \quad \forall j \leq k. \quad (1.3)$$

In order to improve the readability, and throughout the paper, we may use the superscript (k) instead of (X_1, \dots, X_k) , or even drop it for any quantity that depends on the sample of snapshots (X_1, \dots, X_k) . For instance the reduced score constructed out of k snapshots may be denoted:

$$S(x) \stackrel{\text{def}}{=} S^{(k)}(x) \stackrel{\text{def}}{=} S^{(X_1, \dots, X_k)}(x).$$

1.2. Previous works. Advanced methods for Bayesian sampling of (1.1) are gradient-based *Markov chain Monte Carlo* (MCMC) methods such as *Hamiltonian Monte Carlo* [6] and *preconditioning* (of gradients) techniques [14]. Estimation with such advanced MCMC techniques can still exhibit slow convergence rates, as the simulation may need a large time to explore the target distribution, especially in high-dimension [22]. Sequential Monte Carlo (SMC) methods or other methods based on interacting clones/particles are an interesting alternative for Bayesian sampling, and have shown to be stable in certain high-dimensional context [2], in particular *adaptive SMC* simulation [3]. These methods have been employed to solve Bayesian inverse problems [5, 23]. Other sampling methods such as *parallel tempering* have been proposed to accelerate convergence [24]. Nevertheless, sampling with any of these methods become difficult in the present context, where only a too small budget of evaluations of the score function S^* is affordable.

In recent years, numerous MCMC and SMC methods based on reduced score S approximating S^* have been studied to reduce the simulation cost. Among MCMC approaches for inverse Bayesian problems, as expected, convergence results have been obtained for a sequence of reduced scores converging to S^* [13]. Refinements of this approach involve integrating an iterative update of the $S^* - S$ approximation error into a reduced MCMC sampling scheme using S , for a given a predetermined estimate of the error (based on the evaluation of the error for a set of samples drawn from the prior) [7]. It should be noted, however, that a priori sampling to build the reduced model, characterize the approximation error or do both is probably not a relevant approach in our context. Indeed, in Bayesian inverse problems with informative measurements or in low-temperature problems, the target distribution usually concentrates its mass in a small and specific but unknown region of the prior support. More adapted to this context is the strategy to adapt the reduced model during the simulation, using samples whose distribution gradually matches the final target distribution.

Adaptive construction of the reduced score embedded in a tempering scheme was suggested in [26]. As mentioned above, such an adaptive construction of the approximation is relevant to our context, but this work remains quite specific as it relies on parametric Gaussian proposals designed by a cross-entropy method. Moreover, the adaptive construction of reduced score does not use possible error estimates.

Regarding SMC approaches, *multilevel SMC* has been proposed to optimize the product cost \times variance of the estimates; in the latter 'multilevel' refers to a sequence

of approximations of the score function at different discretization levels [4]. However, multilevel SMC requires precise characterization of the model-dependent variance bounds of the estimates, which is a tedious task. In addition, although it relies on a given hierarchy of reduced models or approximations, these are not updated online using snapshots.

An adaptive SMC approach using a predetermined hierarchy of reduced scores has also been proposed in [25]. This work tackles the tricky problem of selecting the right reduced model at the current temperature during simulation. Indeed, choosing a proper criterion to decide along the simulation when to refine the reduced score is quite an open question in general [20], although various heuristics have been proposed in the context of rare event simulation relying on some error quantification [1, 21, 11]. In their work, the authors of [25] propose a heuristic consisting on minimizing the variation of the particles weights testing the case of a reduced score refined or not. In this work, we go beyond these ideas and generally answer the question of when to update the reduced score and with which snapshots.

The authors of [25] have also proposed an analogous method [30] in the context of rare event simulation, in which a predetermined hierarchy of reduced model is given. The authors of the present work have also proposed a method similar to the one studied in the present paper in a rare event context: the *adaptive reduced multilevel splitting* (ARMS) algorithm [11] (*Nota Bene*: 'multilevel' here has historically a completely different meaning as the one in 'multilevel SMC' mentioned before). ARMS is based on a sequential simulation scheme very similar to SMC, which adapts a reduced score approximation while getting closer to the target rare event distribution. This method is nevertheless specific to the rare event simulation context in the sense it is dedicated to sample targets being the reference distribution conditionally to be in some rare event set. The method proposed in this paper (adaptive reduced tempering, ART) can be seen as a generalization of the ARMS algorithm to the case of general target distributions of the form of (1.1), which are relevant for Bayesian problems. However, as explained *e.g.* in Section 1.1, the ART method is also suited to the context of rare event simulation, by considering the latter as a zero temperature limit of a Gibbs family.

As compared to the pieces of work mentioned above, the main novelty of the proposed ART algorithm is to propose an adaptive rule that triggers, within a SMC sampling method with tempered surrogate targets, when to sample a new snapshot of the true model (serving to update of the reduced model). Details are given in the next section.

1.3. Contributions. This work proposes a methodology for sampling the flow of targets (1.1) combining an adaptive Sequential Monte Carlo approach and the availability of reduced modeling with error quantification. To do so, we consider the following exponential family of importance distributions

$$\mu_\beta \stackrel{\text{def}}{=} \frac{1}{Z_\beta} e^{\beta S} d\pi, \quad \beta \in \mathbb{R} \quad (1.4)$$

where the normalization $Z_\beta \stackrel{\text{def}}{=} \pi(e^{\beta S})$ represents an approximation of Z_β^* . An important idea is to interpret the latter as *proposal distributions* (also called *importance distributions*), in the sense of importance sampling, of the target family (1.1). In the context of rare event simulation, using proposal distributions in this exponential family rather than *e.g.* in the form of indicator of events $\propto \mathbf{1}_{S \geq \ell} d\pi$ with $\ell \in \mathbb{R}$ en-

ables to get rid of the possible nesting conditions between successive proposals in the algorithm (see [11]).

The idea behind this work is to design an algorithm iterating on k where k is the number of newly sampled snapshots, and giving at iteration k a well-chosen inverse temperature $\beta^{(k)}$ able to perform a reasonable importance sampling of the target distribution $\eta_{\beta^{(k)}}^*$ defined by (1.1). This adaptive importance sampling is performed with the sequence of proposals

$$\mu_{\beta^{(k)}} = \frac{1}{Z_{\beta^{(k)}}^{(k)}} e^{\beta^{(k)} S^{(k)}} d\pi, \quad k \geq 1$$

where the reduced score $S^{(k)}$ is much cheaper to compute, and only the evaluation of the importance weight is expensive. Very importantly, as the iterations progress, the proposal will adapt by adjusting the inverse temperature $\beta^{(k)}$ and refining the approximation of the reduced score $S^{(k)}$ using the history of the previous computations of the true model. In this framework, three issues emerge.

A first important one is that one has to address the definition of a criterion for setting the lowest inverse temperature $\beta^{(k)}$ suitable for the accuracy of the current reduced score approximation $S^{(k)}$. Justified by theoretical arguments, we propose to constrain the relative entropy¹ between a *surrogate target* denoted $\check{\eta}_\beta$ and the proposal as follows (for a given threshold c_1):

$$\beta^{(k)} \text{ verifies } \text{Ent}(\check{\eta}_{\beta^{(k)}}^{(k)} | \mu_{\beta^{(k)}}^{(k)}) \leq c_1. \quad (1.5)$$

The surrogate target is constructed using the approximate error bound (1.2) of the reduced score approximation as follows:

$$\check{\eta}_\beta \stackrel{\text{def}}{=} \frac{1}{Z_\beta} e^{\beta(S-E)} d\pi \quad \text{with} \quad \check{Z}_\beta \stackrel{\text{def}}{=} \pi(e^{\beta(S-E)}). \quad (1.6)$$

The choice is motivated by the fact that $\check{\eta}_\beta$ is roughly speaking a version of η_β where typical states with large expected errors (hence 'far away' from the already sampled snapshots) are penalized.

A second issue that we will address in this work is the tuning of the expensive snapshot sampling and its associated reduced modeling update. At iteration $k+1$ the snapshot X_{k+1} will be sampled from the proposal $\mu_{\beta^{(k)}}^{(k)}$ which ensures thanks to (1.5) a sampling in areas consistent with the true target distribution $\eta_{\beta^{(k)}}^*$ at inverse temperature $\beta^{(k)}$. In a first phase of the algorithm, we sample snapshots with an additional strong penalty weight favoring candidate states with a larger error (hence requiring priority correction). In a final phase of the algorithm, we remove that penalty to obtain snapshot-based optimal importance sampling estimation.

A last issue that has been addressed in this work is the simulation of a new proposal when at the end of iteration k the reduced score is updated $S^{(k)} \rightarrow S^{(k+1)}$. We will propose to *bridge* proposals by searching among past simulations for an iteration index k_{bridge} with associated approximate score $S^{(k_{\text{bridge}})}$, and by finding an inverse

¹or Kullback-Leibler divergence between probabilities defined by $\text{Ent}(\eta | \mu) \stackrel{\text{def}}{=} \int \log \frac{d\eta}{d\mu} d\eta$ when η is dominated by μ and extended to $+\infty$ otherwise.

temperature β_{bridge} such that the relative entropy between the new proposal $\mu_{\beta_{\text{bridge}}}^{(k+1)}$ and the past proposal $\mu_{\beta^{(k_{\text{bridge}})}}^{(k_{\text{bridge}})}$ is kept below a certain standard threshold:

$$\beta_{\text{bridge}}, k_{\text{bridge}} \text{ verifies } \text{Ent}(\mu_{\beta_{\text{bridge}}}^{(k+1)} \mid \mu_{\beta^{(k_{\text{bridge}})}}^{(k_{\text{bridge}})}) \leq c_2. \quad (1.7)$$

This work will demonstrate the following facts about the proposed algorithmic methodology. First, at a mathematical level, the proposed importance sampling estimator is shown to be consistent in an idealized setting. Second, we exhibit in our numerical experiments its practical convergence. Finally and most importantly, by evaluating the computational complexity in relation to the estimation error, we show that the use of reduced modeling (as compared to similar methods with only true model evaluation) significantly reduces the simulation cost while achieving a comparable accuracy.

The paper is organized as follows. An idealized version of the proposed algorithm, in which we assume the exact evaluation of expectations and the exact sampling of distributions, is presented in Section 2 together with the underlying methodological concepts. We then provide in Section 3 a practical implementation of this idealized algorithm using SMC simulations. In Section 4, we numerically evaluate the proposed algorithm in the case of rare event simulation and inverse Bayesian problems. A final section deals with conclusions.

2. Idealized Adaptive Reduced Tempering. We describe in this section the ingredients of an idealized adaptive reduced tempering (ART) algorithm that will lead to the consistent estimation of the expectation $\eta_{\beta_\infty}^*(\varphi)$ for any generic test function φ .

2.1. Iterative importance sampling. We consider a vanishing initial inverse temperature $\beta^{(0)} = 0$, so that the initial proposal distribution matches the reference distribution $\mu_{\beta^{(0)}}^{(0)} = \pi$. Using this initialization, we will design in Section 2.2 an adaptive procedure which determines at each iteration $k \geq 1$ (k being the number of snapshot evaluations using the true score S^*) some inverse temperature $\beta^{(k)} \in [0, \beta_\infty]$ “suited” to the current reduced score $S^{(k)}$. To this temperature is associated the target $\eta_{\beta^{(k)}}^*$ as defined by (1.1) and the proposal distribution $\mu_{\beta^{(k)}}^{(k)}$ as defined by (1.4). Given this adaptive procedure determining the sequence of inverse temperatures $\{\beta^{(k)}\}_k$, assume we can compute the normalizations $\{Z_{\beta^{(k)}}^{(k)}\}_k$ and that we can draw exactly according to the proposals $X_{k+1} \sim \mu_{\beta^{(k)}}^{(k)}$ in order to compute the sample of snapshots $\{(X_1, S^*(X_1)), (X_2, S^*(X_2)), \dots\}$. Then, an *iterative importance sampling estimator* of the unnormalized target $\gamma_{\beta_\infty}^* \stackrel{\text{def}}{=} Z_{\beta_\infty} \eta_{\beta_\infty}^*$ is given by

$$\hat{\gamma}_{\beta_\infty}^H(\varphi) = \frac{1}{H} \sum_{k=K_{j_0}+1}^{K_{j_0}+H} Z_{\beta^{(k)}}^{(k)} e^{\beta_\infty S^*(X_{k+1}) - \beta^{(k)} S^{(k)}(X_{k+1})} \varphi(X_{k+1}), \quad (2.1)$$

with H denoting the size of the sample used for estimation and where K_{j_0} denotes a random iteration index whose definition will be given in due time. The self-normalized estimate of the target $\eta_{\beta_\infty}^*$ is

$$\hat{\eta}_{\beta_\infty}^H(\varphi) = \frac{\hat{\gamma}_{\beta_\infty}^H(\varphi)}{\hat{\gamma}_{\beta_\infty}^H(\mathbf{1})}. \quad (2.2)$$

Therefore, to enable the evaluation of these estimates, we need at iteration k of the algorithm to *i*) adapt the inverse temperature $\beta^{(k)}$, so that it is “suited” to the current reduced score function $S^{(k)}$, *ii*) compute efficiently the normalization $Z_{\beta^{(k)}}$. We discuss these two issues in the next sections.

2.2. Critical inverse temperature. The idea is to tune at iteration k the inverse temperature $\beta^{(k)}$ with a criterion evaluating that surpassing this inverse temperature should require, in some sense, a better surrogate model, and thus new snapshots. The adjusted proposal $\mu_{\beta^{(k)}}^{(k)}$ 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 to detect when and where a new snapshot is needed.

The general insight is the following. Consider the following simpler problem: one is given the family of target distributions in the form of (1.1), a given computational budget, and the family of importance sampling proposal distributions in the form of (1.4). One wants to perform some sequential algorithm to sample the distributions $\mu_{\beta}^{(k)}$ along an increasing sequence of inverse temperatures β as high as possible, and then use importance sampling. However, we do not want to waste resources for temperature that are “unreasonably” low. As an idealized criterion giving achievable temperatures, we propose the *cost of importance sampling* of the true target. As is thoroughly argued in [12] (see also [10, 11]), the variance is in general a pessimistic quantification of importance sampling. For a target η with proposal μ , it has been proven in [12] that the cost of importance sampling – in terms of sample size – 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 η and μ , *i.e.*, $e^{\text{Ent}(\eta | \mu)}$, where $\text{Ent}(\eta | \mu) \stackrel{\text{def}}{=} \eta(\log \eta) - \eta(\log \mu)$. We will thus call $\text{Ent}(\eta | \mu)$ the log cost of importance sampling of the target η by the proposal μ . The first non-achievable temperature, hence a kind of critical temperature, is thus given when the latter matches a certain given numerical threshold c_1 . By doing so, we ensure that the snapshot sampled with respect to $\mu_{\beta}^{(k)}$ will not diverge too much from the target distribution η_{β}^* at an appropriate inverse temperature β .

In a practical context, the log cost $\text{Ent}(\eta_{\beta}^* | \mu_{\beta}^{(k)})$ has to be estimated using some prior information on η_{β}^* , for instance using some form of error quantification. We propose to consider a surrogate of the target η_{β}^* , denoted by $\check{\eta}_{\beta}^{(k)}$, constructed according to (1.6) with the reduced score $S^{(k)}$ and its error quantification $E^{(k)}$ of the form of (1.2). We stress that this choice of surrogate is merely conventional and other cases could be studied in future work. Our choice is simply a limit case² of possible true score imposed by the approximate error bound (1.2). The critical temperature will be estimated using the surrogate target by searching the closest temperature such that the log cost $\text{Ent}(\check{\eta}_{\beta}^{(k)} | \mu_{\beta}^{(k)})$ matches the threshold c_1 . We emphasize that the irrelevant situation where relative entropy is uniformly zero is excluded by the constraint (1.3), which requires that the error $E^{(k)}$ vanishes on snapshot points $X^{(k)}$ already evaluated at iteration k , prohibits constant error quantification functions unless the error estimate E is identically zero. We’d like to point out that this condition is naturally satisfied by many reduced models, and in particular by reduced

²Taking $S + E$ instead of $S - E$ in the definition of the surrogate $\check{\eta}_{\beta}$ is a possibility but it seems to lead to less robust calculations because it favors states with high errors. We did not studied in detail different options for surrogate targets in the present work.

bases. However, this requirement is not a hard limit on the method, as the choice of the surrogate $\check{\eta}_\beta$ can be adapted in other cases.

Given an initial inverse temperature β_0 , this yields the following rigorous definition of the inverse of the critical temperature:

$$\beta^{(k)}(c_1, \beta_0) \stackrel{\text{def}}{=} \min\{\beta : \text{Ent}(\check{\eta}_\beta^{(k)} | \mu_\beta^{(k)}) > c_1, \beta_0 \leq \beta \leq \beta_\infty\}, \quad (2.3)$$

with the initial temperature at $k = 0$ set to zero, *i.e.*, $\beta^{(0)}(c_1, \beta_0) \stackrel{\text{def}}{=} 0$. For convenience, the dependency on c_1 and β_0 of $\beta^{(k)}$ will be removed if the context makes this clear. $\beta^{(k)}(c_1, \beta_0)$ is an inverse temperature such that, in the case where the true target is given by the surrogate $\check{\eta}_l^{(k)}$, a log cost of c_1 is necessary and sufficient in order to perform importance sampling. We will define β_0 later in Section 2.5.

2.3. Adaptive tempering down to the critical temperature. In the previous section, we have defined $\beta^{(k)}$ as the critical level at the k -th iteration of the algorithm. However, the computation of the normalization $Z_{\beta^{(k)}}^{(k)}$ related to the distribution $\mu_{\beta^{(k)}}^{(k)}$ is often unaffordable by direct Monte Carlo simulation for a large $\beta^{(k)}$, as it corresponds to a difficult low-temperature problem. A frequently invoked procedure for these problems involves forming a sequence of tempered probability distributions $\mu_\beta^{(k)}$ among the family (1.4), defined with the reduced score $S^{(k)}$ and parameters $\beta_0 \leq \beta_1 \leq \dots \leq \beta_n = \beta^{(k)}$ [3]. Assume given an initial normalization constant $Z_{\beta_0}^{(k)}$ related to an initial distribution $\mu_{\beta_0}^{(k)}$. It is straightforward to show that, using this sequence, the sought normalization can be decomposed into the product

$$Z_{\beta^{(k)}}^{(k)} = Z_{\beta_0}^{(k)} \prod_{i=1}^{n-1} \mu_{\beta_i}^{(k)}(G_i), \quad G_i(x) = e^{(\beta_{i+1} - \beta_i)S^{(k)}(x)}, \quad (2.4)$$

where, as we will see in Section 3, the factor $\mu_{\beta_i}^{(k)}(G_i)$ can be recursively approximated using a SMC simulation. An adaptive version of this procedure, known as *adaptive tempering*, corresponds to a flow of distributions with a sequence of temperature parameters chosen adaptively. In this work, the inverse temperature β_{i+1} will be adapted using the rule

$$\beta_{i+1} := \max\{\beta : \text{Ent}(\mu_\beta^{(k)} | \mu_{\beta_i}^{(k)}) \leq c_2, \beta_i \leq \beta \leq \beta^{(k)}\}, \quad (2.5)$$

where $c_2 > 0$ is some given prescribed level. This one dimensional problem is easily solved since $\beta \mapsto \text{Ent}(\mu_\beta^{(k)} | \mu_{\beta_i}^{(k)})$ is increasing.³ The entropic criterion used in (2.5) may be viewed as an extension to smooth distributions of the criterion used in adaptive multilevel splitting [11]. However, other criteria are possible, see for example [3].

2.4. Snapshot sampling. As defined Section 2.1, only the snapshots sampled at and after iteration $K_{j_0} + 1$ are used to build the importance sampling estimator $\hat{\gamma}_{\beta_\infty}^H$ and $\hat{\eta}_{\beta_\infty}^H$. The idea here is that the snapshots should be first used to refine the reduced score $S^{(k)}$ around the regime of interest. The snapshots can be used for importance sampling once the reduced score is sufficiently accurate in the sense of a limited variance for importance sampling with distribution $\mu_{\beta^{(k)}}^{(k)}$ of the target

³Some simple calculation shows that $\partial_\beta \text{Ent}(\mu_\beta^{(k)} | \mu_{\beta_i}^{(k)}) = (\beta - \beta_i) \text{var}_{\mu_\beta^{(k)}}(S^{(k)})$

$\eta_{\beta^{(k)}}^*$. We choose to define the *trigger index* K_{j_0} (random index independent of the future snapshots) as the first iteration index k for which among the set of already computed critical inverse temperatures $\{\beta^{(k')}, k' \leq k\}$, at least j_0 of them are equal to the targeted inverse temperature β_∞ , for some prescribed $j_0 \geq 1$. This snapshot sampling approach corresponds to some extent to a particular setting of the so-called *learning function* in the active learning literature. [28].

In summary, in the first iterations new snapshots are sampled using a large penalty favoring high error bounds:

$$X_{k+1} \sim \frac{1}{Z_\tau^{(k)}} e^{\tau E^{(k)}(x)} \mu_{\beta^{(k)}}^{(k)}(dx) \quad \text{if } k < K_{j_0}.$$

In practice in the SMC implementation, for not too large sample size N , we may pick τ large enough to be virtually equivalent to $+\infty$. Then, once the reduced score is related to a temperature close to the target one, new snapshots are sampled according to the importance distribution:

$$X_{k+1} \sim \mu_{\beta^{(k)}}^{(k)} \quad \text{if } k \geq K_{j_0}.$$

2.5. Bridging reduced scores. Given an inverse temperature $\beta^{(k)}$ and an associated importance distribution $\mu_{\beta^{(k)}}^{(k)}$ related to the reduced score $S^{(k)}$, the problem is to determine a new initial inverse temperature denoted by $\beta_{\text{bridge}}(k+1)$, “suited” to the updated reduced score $S^{(k+1)}$ related to the updated importance distribution $\mu_{\beta_{\text{bridge}}(k+1)}^{(k+1)}$. In order to set this temperature in consistency with the critical temperature constraint (2.3) and the adaptive tempering constraint (2.5), it is natural to consider the problem

$$\text{Find: } \beta \in \mathcal{C}(k), \quad (2.6)$$

where

$$\mathcal{C}(k) \stackrel{\text{def}}{=} \left\{ \beta : \text{Ent} \left(\check{\eta}_\beta^{(k+1)} \mid \mu_\beta^{(k+1)} \right) \leq c_1, \text{Ent} \left(\mu_\beta^{(k+1)} \mid \mu_{\beta^{(k)}}^{(k)} \right) \leq c_2, \beta^{(k)} \leq \beta \leq \beta_\infty \right\}. \quad (2.7)$$

However the feasible set $\mathcal{C}(k)$ might be an empty set. The idea is thus to extend problem (2.6) to indices $k' \in \llbracket 0, k \rrbracket$, and search for the largest index k' such that the set $\mathcal{C}(k') \neq \emptyset$, by keeping in memory the past proposal distributions $\{\mu_{\beta^{(k')}}^{(k')}\}_{k' \leq k}$. In order to set a feasible bridging inverse temperature $\beta_{\text{bridge}}(k+1)$, it is then natural to consider the problem

$$\text{Find: } \beta \in \mathcal{C}(k_{\text{bridge}}) \quad \text{where } k_{\text{bridge}} \stackrel{\text{def}}{=} \text{largest } k' : k' \in \llbracket 0, k \rrbracket \text{ s.t. } \mathcal{C}(k') \neq \emptyset. \quad (2.8)$$

In practice, the set $\mathcal{C}(k_{\text{bridge}})$ is not a singleton and we need to define $\beta_{\text{bridge}}(k+1)$ as a particular element of this set. One possible choice for $\beta_{\text{bridge}}(k+1)$, which we will use, is to select the largest element of $\mathcal{C}(k_{\text{bridge}})$. Note that for $k' = 0$, the initial proposal distribution is related to a vanishing inverse temperature $\beta^{(0)} = 0$. Therefore $\mathcal{C}(0) = \{0\}$, which shows that a solution of (2.8) always exists.

It is straightforward to show that, at the bridging temperature solution $\beta_{\text{bridge}}(k+1)$, the normalization constant $Z_{\beta_{\text{bridge}}(k+1)}^{(k+1)}$ related to the distribution $\mu_{\beta_{\text{bridge}}(k+1)}^{(k+1)}$ satisfies

$$Z_{\beta_{\text{bridge}}(k+1)}^{(k+1)} = Z_{\beta^{(k_{\text{bridge}})}}^{(k_{\text{bridge}})} \mu_{\beta^{(k_{\text{bridge}})}}^{(k_{\text{bridge}})} \left(e^{\beta_{\text{bridge}}(k+1) S^{(k+1)}(x) - \beta^{(k_{\text{bridge}})} S^{(k_{\text{bridge}})}(x)} \right). \quad (2.9)$$

The adaptive tempering procedure discussed in Section 2.3 can then be deployed: the recursion is initialized at iteration $k + 1$ using the reduced score for $S^{(k+1)}$ as $\beta_0 := \beta_{\text{bridge}}(k + 1)$ (see definition (2.5)), and is run up to the critical parameter $\beta_n := \beta^{(k+1)}(c_1, \beta_{\text{bridge}}(k + 1))$ (as given by definition (2.3)).

2.6. Stopping criterion. In principle, the algorithm can continue as long as one wishes, for instance until the budget of snapshot samples is exhausted. However, in practice, we stop updating the reduced score of the algorithm if the divergence between the surrogate target and the reduced proposal is very small and if the associated critical inverse temperature equals β_∞ . More precisely, no update of the reduced model is performed if the following condition is satisfied:

$$\text{Ent} \left(\tilde{\eta}_{\beta^{(k)}}^{(k)} \mid \mu_{\beta^{(k)}}^{(k)} \right) \leq \epsilon \quad \& \quad \beta^{(k)} = \beta_\infty, \quad (2.10)$$

where $\epsilon > 0$ is a given threshold parameter close to the machine precision. In that case the algorithm amounts to a standard importance sampling procedure based on the proposal μ_{β_∞} computed with the same, final, reduced model.

2.7. The algorithm. The algorithm presented below gathers the ingredients presented in the previous sections. It is an idealized pseudo-code and, in practice, it is not directly implementable. Indeed, it assumes exact evaluation of expectations and exact sampling according to the proposals. In Section 3, we will propose a practical implementation of this idealized algorithm using a sample of N clones (or *particles*) SMC methodology. The latter will simulate the flow of proposals for diminishing temperatures using Monte Carlo samples with approximately correct distributions for large sample sizes $N \rightarrow +\infty$.

As shown in Appendix A with a martingale argument, the unnormalized estimator (2.1) computed by the algorithm below is unbiased

$$\mathbb{E} \hat{\gamma}_{\beta_\infty}^H = \gamma_{\beta_\infty}^*,$$

and its variance scales in $\mathcal{O}(H^{-1})$. Note that in this idealized setup, samples are assumed to be drawn exactly from the sequence of importance distributions. In practice, as detailed in Section 3, this assumption will not hold exactly because of the adaptive features used in the sampling of the proposals by the SMC routine, yielding a small bias of order the inverse of the sample size.

Algorithm (ART)

Require: inverse temperature target β_∞ , initial reduced score $S^{(1)}$ and error $E^{(1)}$, budget of snapshots K , worst-case log cost threshold c_1 , adaptive tempering parameter c_2 , reduced score precision $\epsilon \ll c_1$, trigger parameter j_0 , learning intensity τ .

▷ (initialization)

- 1: Set $\beta_{\text{bridge}}(1) = 0$, $H = 0$, $j = 0$, $k = 0$, $Z_\beta^{(1)} = 1$
- 2: **while** $K > 0$ **do**
- 3: $k \leftarrow k + 1$
 ▷ (critical temperature and adaptive tempering, see Section 2.2 & 2.3)
- 4: Simulate the flow of proposals $\mu_{\beta_i}^{(k)}$ for $i = 0, \dots, n$ using the recursion

$$\begin{cases} \beta_0 = \beta_{\text{bridge}}(k) \\ \beta_{i+1} = \min\{\beta : \text{Ent}(\mu_\beta^{(k)} \mid \mu_{\beta_i}^{(k)}) > c_2, \beta_i \leq \beta \leq \beta_\infty\}, \end{cases} \quad (2.11)$$

where the index of the final inverse temperature β_n is

$$n = \min\{i : \text{Ent}(\tilde{\eta}_{\beta_{i+1}}^{(k)} \mid \mu_{\beta_{i+1}}^{(k)}) > c_1 \quad \text{or} \quad \beta_i = \beta_\infty\}. \quad (2.12)$$

5: Set $\beta^{(k)} = \beta_n$ and compute the associated normalisation $Z_{\beta^{(k)}}^{(k)}$ with (2.4).
 \triangleright (snapshot sampling, see Section 2.4)

6: **if** $\beta^{(k)} = \beta_\infty$ **then** $j \leftarrow j + 1$ **end if**

7: **if** $j \geq j_0$ **then** draw $X \sim \mu_{\beta^{(k)}}^{(k)}$, **else** $X \sim e^{\tau E^{(k)}} \mu_{\beta^{(k)}}^{(k)} / Z_\tau^{(k)}$.

8: Compute the snapshot $S^*(X) = \text{score}(\Psi(X))$
 \triangleright (stopping criterion, see Section 2.6)

9:

10: **if** $\text{Ent}(\tilde{\eta}_{\beta^{(k)}}^{(k)} \mid \mu_{\beta^{(k)}}^{(k)}) > \epsilon$ or $\beta^{(k)} < \beta_\infty$ **then**

11: Update reduced model: build $(S^{(k+1)}, E^{(k+1)})$ by enriching $(S^{(k)}, E^{(k)})$ with $(X, S^*(X))$.

12: **else**

13: Don't update the reduced model: set $(S^{(k+1)}, E^{(k+1)}) \leftarrow (S^{(k)}, E^{(k)})$.

14: **end if**
 \triangleright (importance sampling, see Section 2.1)

15: **if** $j > j_0$ **then**

16: $H \leftarrow H + 1$

17:
$$\hat{\gamma}_{\beta_\infty}^H \leftarrow \frac{(H-1)}{H} \hat{\gamma}_{\beta_\infty}^{H-1} + \frac{1}{H} Z_{\beta^{(k)}}^{(k)} e^{\beta_\infty S^*(X) - \beta^{(k)} S^{(k)}(X)} \delta_X,$$

18: **end if**

19: Set $K \leftarrow K - 1$
 \triangleright (bridging, see Section 2.5)

20: Find the largest $k' \in \llbracket 0, k \rrbracket$ such that exists $\beta' \in \mathcal{C}(k')$ where:

$$\mathcal{C}(k') = \{\beta' : \text{Ent}(\tilde{\eta}_{\beta'}^{(k+1)} \mid \mu_{\beta'}^{(k+1)}) \leq c_1, \text{Ent}(\mu_{\beta'}^{(k+1)} \mid \mu_{\beta^{(k')}}^{(k')}) \leq c_2, \beta^{(k')} \leq \beta' \leq \beta_\infty\}$$

and denote a solution pair by $(k_{\text{bridge}}(k+1), \beta_{\text{bridge}}(k+1))$.

21: Simulate the proposal $\mu_{\beta_{\text{bridge}}(k+1)}^{(k+1)}$ starting from $\mu_{\beta_{\text{bridge}}(k+1)}^{(k_{\text{bridge}}(k+1))}$

22: Compute using (2.9) the normalization $Z_{\beta_{\text{bridge}}(k+1)}^{(k+1)}$.

23: **end while**

24: **Return** $\hat{\gamma}_{\beta_\infty}^H / \hat{\gamma}_{\beta_\infty}^H(\mathbf{1})$

3. Algorithm in Practice: SMC Simulation. We describe now how an iteration of the idealized ART algorithm presented above may be implemented in practice. Suppose we start the k -th iteration of the ART algorithm with a given sample of size N approximately distributed according to the law $\mu_{\beta_{\text{bridge}}(k)}^{(k)}$, as well as with a given estimate of the associated normalization $Z_{\beta_{\text{bridge}}(k)}^{(k)}$. We will use in practice a SMC sampling methodology [3] to approximate the adaptive tempering step introduced in Section 2.3 and the bridging procedure proposed in Section (2.5). The following sections describe these approximations, with the other steps in the algorithm approximated using SMC flowing directly from them.

3.1. Approximation of a flow of tempered distribution. The SMC simulation generates a finite sequence of N -samples approximately distributed according to a sequence of distributions $\{\mu_{\beta_i}^{(k)}\}_{1 \leq i \leq n}$, as well as estimators of the related normalizations $\{Z_{\beta_i}^{(k)}\}_{1 \leq i \leq n}$. The i -th empirical distribution related to the N -sample $(\Xi_1^{(i)}, \dots, \Xi_N^{(i)})$ (we omit in notation the dependence to the current iteration k) generated by SMC will be denoted

$$\mu_{\beta_i}^{N, (k)} \stackrel{\text{def}}{=} \frac{1}{N} \sum_{\ell=1}^N \delta_{\Xi_\ell^{(i)}} \simeq \mu_{\beta_i}^{(k)},$$

11

and the estimator of the normalization constant $Z_{\beta_n}^{(k)}$ will be denoted $Z_{\beta_n}^{N,(k)}$.

The sequence of inverse temperatures $\beta_0 = \beta_{\text{bridge}}(k) < \beta_1 < \dots < \beta_n = \beta^{(k)}$ will be redefined in this SMC simulation framework as an approximation targeting the sequence defined in (2.11), with the final element being an approximation of the critical inverse temperature (2.12). However, before going into the latter approximations, let us begin by describing how starting from $\mu_{\beta_i}^{N,(k)}$, SMC simulates at the i -th iteration the empirical distribution $\mu_{\beta_{i+1}}^{N,(k)}$. SMC performs successively: *i*) a weighting and selection step of the N particles, and *ii*) a mutation step of each individual particle. The weighting step *i*) assigns weights to the empirical distribution $\mu_{\beta_i}^{N,(k)}$ so that it approximates $\mu_{\beta_{i+1}}^{(k)}$

$$\mu_{\beta_{i+1}}^{N,(k)} = \sum_{\ell=1}^N w_{\ell}^{(i)} \delta_{\Xi_{\ell}^{(i)}} \quad \text{with} \quad w_{\ell}^{(i)} = \frac{G_i(\Xi_{\ell}^{(i)})}{\sum_{\ell'=1}^N G_i(\Xi_{\ell'}^{(i)})},$$

where $G_i(x)$ is given by (2.4). It is straightforward to show that the associated normalizations follow the recursion

$$Z_{\beta_{i+1}}^{N,(k)} = Z_{\beta_i}^{N,(k)} \mu_{\beta_i}^{N,(k)}(G_i).$$

The above weighting step is followed by a selection step, where the particles are re-sampled according to their weights $w_{\ell}^{(i)}$. For this purpose, we choose the efficient systematic resampling scheme among the multiple available resampling techniques [18]. The mutation step *ii*) then consists of a Markovian transition that leaves $\mu_{\beta_{i+1}}^{(k)}$ invariant and is applied individually and independently a certain number of times to the N particles. In practice, the Markovian transitions are performed using a Metropolis-Hasting algorithm, with an adaptive Markovian kernel in order to maintain a given Metropolis acceptance rate [19].

We emphasize that tempering with SMC 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. Approximation of critical temperature and adaptive tempering.

The simulation starts from an i.i.d. sample with the reference distribution $(\beta_{\text{bridge}}(1) = 0$ so that $\mu_{\beta_{\text{bridge}}(1)}^{(1)} = \pi$). We recall that after k iterations, the algorithm provides a current reduced score $S^{(k)}$, an associated pointwise error estimate $E^{(k)}$, an initial inverse temperature $\beta_{\text{bridge}}(k)$ and a sample of size N approximately distributed according to $\mu_{\beta_{\text{bridge}}(k)}^{(k)}$, denoted using a sample-size superscript by $\mu_{\beta_{\text{bridge}}(k)}^{N,(k)}$.

The ART algorithm needs in practice an approximation of the critical inverse temperature $\beta^{(k)}$ given by the final inverse temperature of index determined by the criterion (2.12). We approximate this criterion by substituting the continuous distributions by sample-based empirical approximations. More precisely, at the i -th iteration of the tempering scheme, we use the empirical distribution $\mu_{\beta_i}^{N,(k)}$ to approximate

the worst-case log cost for an inverse temperature β close to β_i as

$$\begin{aligned} \text{Ent} \left(\tilde{\eta}_\beta^{(k)} \mid \mu_\beta^{(k)} \right) &= \log \frac{Z_\beta^{(k)}}{\tilde{Z}_\beta^{(k)}} - \beta \tilde{\eta}_\beta^{(k)}(E^{(k)}), \\ &\approx \log \frac{\mu_{\beta_i}^{N,(k)}(e^{(\beta-\beta_i)S^{(k)}})}{\mu_{\beta_i}^{N,(k)}(e^{(\beta-\beta_i)S^{(k)}-\beta E^{(k)}})} - \beta \frac{\mu_{\beta_i}^{N,(k)}(e^{(\beta-\beta_i)S^{(k)}-\beta E^{(k)}}) E^{(k)}}{\mu_{\beta_i}^{N,(k)}(e^{(\beta-\beta_i)S^{(k)}-\beta E^{(k)}})}, \end{aligned} \quad (3.1)$$

where we have used the definitions of the proposal (1.4) and the surrogate target (1.6).

The ART algorithm also needs in practice a Monte-Carlo approximation of the sequence of temperature given in (2.11). Analogously to what was done previously, at the i -th iteration of the tempering scheme, we use the empirical distribution $\mu_{\beta_i}^{N,(k)}$ to approximate the adaptive criterion for β close to β_i as

$$\begin{aligned} \text{Ent} \left(\mu_\beta^{(k)} \mid \mu_{\beta_i}^{(k)} \right) &= \log \frac{Z_{\beta_i}^{(k)}}{Z_\beta^{(k)}} + (\beta - \beta_i) \mu_\beta^{(k)}(S^{(k)}), \\ &\approx -\log \mu_{\beta_i}^{N,(k)}(e^{(\beta-\beta_i)S^{(k)}}) + (\beta - \beta_i) \frac{\mu_{\beta_i}^{N,(k)}(e^{(\beta-\beta_i)S^{(k)}}) S^{(k)}}{\mu_{\beta_i}^{N,(k)}(e^{(\beta-\beta_i)S^{(k)}})}, \end{aligned} \quad (3.2)$$

Now, using these two approximations, the SMC simulation proceeds as explained in Section 3.1: it creates a random sequence $(\beta_i)_{i \geq 0}$ of increasing inverse temperature starting from $\beta_0 = \beta_{\text{bridge}}(k)$, related to a sequence of empirical distributions $\mu_{\beta_i}^{N,(k)}$ for $i \geq 0$, until it reaches the critical inverse temperature $\beta^{(k)}(c_1, \beta_{\text{bridge}}(k))$. More precisely, at the i -th iteration of the tempering scheme, we propose an increase for β_i by $\delta\beta$ given as the solution of an optimization problem defined with the approximation (3.2)

$$\delta\beta = \min \left\{ \delta\beta' : \delta\beta' \frac{\mu_{\beta_i}^{N,(k)}(e^{\delta\beta' S^{(k)}}) S^{(k)}}{\mu_{\beta_i}^{N,(k)}(e^{\delta\beta' S^{(k)}})} - \log \mu_{\beta_i}^{N,(k)}(e^{\delta\beta' S^{(k)}}) > c_2, 0 \leq \delta\beta' \leq \beta_\infty - \beta_i \right\}.$$

This one-dimensional problem can be efficiently solved using a standard bisection method. The increase $\delta\beta$ is accepted if it is below or equal the critical inverse temperature, *i.e.*, using approximation (3.1) if

$$\log \frac{\mu_{\beta_i}^{N,(k)}(e^{\delta\beta S^{(k)}})}{\mu_{\beta_i}^{N,(k)}(e^{\delta\beta S^{(k)}-(\beta_i+\delta\beta)E^{(k)}})} - (\beta_i + \delta\beta) \frac{\mu_{\beta_i}^{N,(k)}(e^{\delta\beta S^{(k)}-(\beta_i+\delta\beta)E^{(k)}}) E^{(k)}}{\mu_{\beta_i}^{N,(k)}(e^{\delta\beta S^{(k)}-(\beta_i+\delta\beta)E^{(k)}})} \leq c_1. \quad (3.3)$$

Therefore if (3.3) holds then $\beta_{i+1} = \beta_i + \delta\beta$, else the critical inverse temperature is reached, *i.e.*, $\beta^{(k)}(c_1, \beta_{\text{bridge}}(k)) = \beta_i$.

3.3. Approximation for bridging. As exposed in Section 2.5, the bridging procedure at the end of iteration k consists in selecting a past index $= k_{\text{bridge}} \leq k$ and a new initial inverse temperature $\beta_{\text{bridge}}(k+1)$ which will be used to initialize the SMC simulation with the new score approximation $S^{(k+1)}$ starting from the past, recorded, empirical distribution $\mu_{\beta^{(k_{\text{bridge}})}}^{N,(k_{\text{bridge}})}$. In line with the definition of problem (2.8), we select

for this purpose the largest lexicographic pair $(k', \beta) \in \llbracket 0, k \rrbracket \times \mathcal{C}(k')$. However, in order to provide an implementation of this step of the ART algorithm, it is necessary to define an empirical counterpart to this optimization problem. As done previously, using $\mu_{\beta^{(k')}}^{N, (k')}$, the worst log-cost is approximated for β close to $\beta^{(k')}$ by

$$\text{Ent} \left(\tilde{\eta}_{\beta}^{(k+1)} \mid \mu_{\beta}^{(k+1)} \right) \approx \log \frac{\mu_{\beta^{(k')}}^{N, (k')} (e^{\phi_{\beta, k'}})}{\mu_{\beta^{(k')}}^{N, (k')} (e^{\phi_{\beta, k'} - \beta E^{(k+1)}})} - \beta \frac{\mu_{\beta^{(k')}}^{N, (k')} (e^{\phi_{\beta, k'} - \beta E^{(k+1)}}) E^{(k+1)}}{\mu_{\beta^{(k')}}^{N, (k')} (e^{\phi_{\beta, k'} - \beta E^{(k+1)}})}, \quad (3.4)$$

while the entropic adaptive criterion is approximated for β close to $\beta^{(k')}$ by

$$\text{Ent} \left(\mu_{\beta}^{(k+1)} \mid \mu_{\beta^{(k')}}^{(k')} \right) \approx -\log \mu_{\beta^{(k')}}^{N, (k')} (e^{\phi_{\beta, k'}}) + \frac{\mu_{\beta^{(k')}}^{N, (k')} (e^{\phi_{\beta, k'}} \phi_{\beta, k'})}{\mu_{\beta^{(k')}}^{N, (k')} (e^{\phi_{\beta, k'}})}, \quad (3.5)$$

with

$$\phi_{\beta, k'} = \beta S^{(k+1)} - \beta^{(k')} S^{(k')}.$$

A SMC approximation of the bridging procedure (2.8) consists then in solving a sequence of one-dimensional constrained optimization problem of the form

$$\beta' = \max\{\beta : (3.4) \leq c_1, (3.5) \leq c_2, \beta^{(k')} \leq \beta \leq \beta_{\infty}\},$$

starting from index $k' = k$ and decreasing the index until the problem above admits a solution, in which case

$$(k_{\text{bridge}}, \beta_{\text{bridge}}(k+1)) = (k', \beta').$$

We remark that the set of constraints related to this problem is not necessarily convex. However, we can efficiently solved this one-dimensional problem with standard optimization techniques. We choose to use an interior point method. Analogously to the procedure detailed in Section 3.1, the normalization estimate is then updated as

$$Z_{\beta_{\text{bridge}}(k+1)}^{N, (k+1)} = Z_{\beta^{(k_{\text{bridge}})}}^{N, (k_{\text{bridge}})} \mu_{\beta^{(k_{\text{bridge}})}}^{N, (k_{\text{bridge}})} (G_{\text{bridge}}), \quad (3.6)$$

where

$$G_{\text{bridge}}(x) = e^{\beta_{\text{bridge}}(k+1)S^{(k+1)}(x) - \beta^{(k_{\text{bridge}})}S^{(k_{\text{bridge}})}(x)}.$$

The N -sample (Ξ_1, \dots, Ξ_N) related to the empirical distribution $\mu_{\beta^{(k_{\text{bridge}})}}^{N, (k_{\text{bridge}})}$ is resampled according to the weights $\frac{G_{\text{bridge}}(\Xi_i)}{\sum_{i'=1}^N G_{\text{bridge}}(\Xi_{i'})}$, and mutated to target the new initial distribution $\mu_{\beta_{\text{bridge}}(k+1)}^{(k+1)}$ serving at the next iteration.

We omit to comment the SMC approximation for snapshot sampling or computing the stopping index in the ART algorithm as they follow naturally from the above discussion or from the simple substitution of distributions by their empirical equivalents. This concludes the k -th iteration of the algorithm.

3.4. SMC-based estimators. Using the sequence of empirical distributions $\{\mu_{\beta^{(k)}}^{N,(k)}\}_{k=1}^K$ generated by SMC simulations along the ART algorithm iterates, we can write a closed-form approximation of the importance sampling estimator given in (2.1) - (2.2). The SMC approximation of the IS estimator $\hat{\gamma}_{\beta_\infty}^H$ of the the unnormalized distribution $\gamma_{\beta_\infty}^*$ is given by

$$\hat{\gamma}_{\beta_\infty,IS}^{N,H} = \frac{1}{H} \sum_{k=K_{j_0}+1}^{K_{j_0}+H} Z_{\beta^{(k)}}^{N,(k)} e^{\beta_\infty S^*(X_{k+1}) - \beta^{(k)} S^{(k)}(X_{k+1})} \delta_{X_{k+1}}. \quad (3.7)$$

Moreover, SMC simulation also provides as by-products an estimator of $\gamma_{\beta_\infty}^*$, which we will call reduced SMC (RSMC) estimator because based solely on the reduced score, given by

$$\hat{\gamma}_{\beta_\infty,RSMC}^{N,H} = \frac{1}{HN} \sum_{k=K_{j_0}+1}^{K_{j_0}+H} Z_{\beta^{(k)}}^{N,(k)} \sum_{i=1}^N e^{(\beta_\infty - \beta^{(k)}) S^{(k)}} \delta_{\Xi_i^{(k)}}, \quad (3.8)$$

where $(\Xi_1^{(k)}, \dots, \Xi_N^{(k)})$ denotes the N -sample related to the empirical distribution $\mu_{\beta^{(k)}}^{N,(k)}$. $\hat{\gamma}_{\beta_\infty,RSMC}^{N,H}$ is expected to be increasingly accurate in the case where the reduced score $S^{(k)}$ converges to S^* as k increases. However, we will see in our numerical experiments that the RSMC estimator can nevertheless retain a significant bias in some cases, even if the reduced score converges to the true one. The SMC-based estimates of the normalized distribution $\eta_{\beta_\infty}^*$ are

$$\hat{\eta}_{\beta_\infty,IS}^{N,H} = \hat{\gamma}_{\beta_\infty,IS}^{N,H} / \hat{\gamma}_{\beta_\infty,IS}^{N,H}(\mathbf{1}) \quad \text{and} \quad \hat{\eta}_{\beta_\infty,RSMC}^{N,H} = \hat{\gamma}_{\beta_\infty,RSMC}^{N,H} / \hat{\gamma}_{\beta_\infty,RSMC}^{N,H}(\mathbf{1}). \quad (3.9)$$

4. Numerical Simulations. Our numerical simulations focus on target distributions of the form (1.1), where S^* implies the high-dimensional solution of an elliptic PDE parametrized by some random input vector. We propose to evaluate numerically the ART algorithm for sampling these distributions, where the PDE solution is approximated by a popular reduced basis (RB) approach. The evaluation is completed by a study of the ART's performance in the context of a one-dimensional yet challenging toy problem. Our numerical simulations are carried out using an implementation of the ART algorithm in the Julia language, available at <https://gitlab.inria.fr/pheas/adaptive-reduced-tempering>. After introducing the models, namely the PDE and its high-fidelity and reduced solutions, we will study two related problems: a rare event estimation problem and a Bayesian inverse problem.

4.1. Model #RB: heat diffusion and reduced basis approximations. Let $\Psi^*(x)$ be the high-fidelity numerical approximation of the solution of a bi-dimensional PDE on the spatial domain $\Omega = [0, 1]^2$ parametrized by a 4-dimensional random vector $x = (x_1 \cdots x_4)^\top \in \mathcal{X} = \mathbb{R}_+^4$. The chosen PDE corresponds to the well-studied *thermal block problem*, which models heat diffusion in an heterogeneous media.

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

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

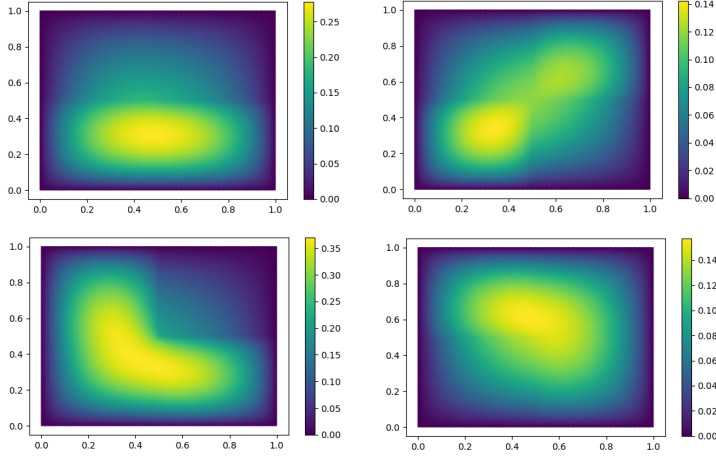


FIG. 4.1. **Illustration of the thermal block problem (Model #RB).** High-fidelity solution $\Psi^*(X)$ of the PDE for 4 different samples X of the reference log-normal distribution (images taken from [11]).

where the function $\kappa(\cdot; x) : \Omega \rightarrow \mathbb{R}_+$ represents a parametric diffusion coefficient function.

The high-fidelity solution $\Psi^*(x)$ 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 $\kappa(\cdot; x) = \sum_{q=1}^4 \mathbf{1}_{\Omega_q}(\cdot) x_q$ with the partition $\Omega = \cup_{q=1}^4 \Omega_q$ and Ω_q 's being non-overlapping squares. Setting the parameter space $\mathcal{X} = (\mathbb{R}_+)^4$ 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. [29] Figure 4.1 shows several high-fidelity solutions $\Psi^*(x)$ of the PDE for samples drawn according to the reference distribution.

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 Ψ of the high-fidelity solution Ψ^* , *ii*) a so-called *a posteriori* error upper bound on the finite element version of the reference $\mathcal{H}_0^1(\Omega)$ norm $\|\Psi^*(x) - \Psi(x)\|_{\mathcal{V}_h}$ denoted by

$$\|\Psi^*(x) - \Psi(x)\|_{\mathcal{V}_h} \leq \Delta_{\Psi}(x).$$

The high-fidelity solution, the RB approximation and the a posteriori error bounds are computed via the “pymor” Python[®] toolbox [27] available at <https://github.com/pymor/pymor>.

The computation of the RB approximation and error bound relies 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 based on the snapshots⁴. It requires $\mathcal{O}(Kh^2 + K^2h)$ operations. Besides, one must add the evaluation of the snapshot, 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 [29].

For sampling problems involving this PDE model, the reference distribution π for x will be a log-normal distribution of independent components with $\ln(x_q) \sim \mathcal{N}(0.6, 0.8^2)$, for $q = 1, \dots, 4$.

4.2. Model #S: closed-form score. From an evaluation perspective, it is difficult to obtain closed-form reference posterior distributions in the context of the previous PDE model. In order to further evaluate the proposed algorithm with a closed-form characterizable reference probability distribution, we rely on the one-dimensional closed-form function proposed in [11]. In this toy model, the previous high-fidelity PDE solution Ψ^* is substituted by a one-dimensional function $\Psi^* : \mathbb{R}_+ \rightarrow \mathbb{R}$, the RB approximation Ψ is replaced by a cubic spline approximation $\Psi : \mathbb{R}_+ \rightarrow \mathbb{R}$ and we assume that the error $|\Psi(x) - \Psi^*(x)|$ is directly available. For $x \in \mathcal{X} = \mathbb{R}_+$, the one-dimensional function is

$$\Psi^*(x) = \begin{cases} \ell & \text{if } x \leq 1/\ell \\ 1/x + f(x) & \text{else} \end{cases}, \quad (4.1)$$

where $f(x) = 15(\mathbf{1}_{x_0 \leq x < x_1} \sin(x - x_0)^2 + \mathbf{1}_{x_1 \leq x} (\sin(x_1 - x_0)^2 - 0.1(x - x_1)))$ with parameters $x_0 = 0.5$ and $x_1 = 5$.

For sampling problems involving this toy model, the reference distribution π will be a log-normal distribution such that $\ln(x) \sim \mathcal{N}(1.5, 1.5^2)$.

4.3. Sampling problems.

4.3.1. Problem #RE: rare event simulation. In these numerical experiments, we consider a target distribution $\eta_{\beta_\infty}^*$ of the form (1.1) with $\text{score}(\cdot) = 1 - \max(\ell - \|\cdot\|_p, 0)/\ell$ and $p > 0$ yielding the score function

$$S^*(x) = 1 - \max(\ell - \|\Psi^*(x)\|_p, 0)/\ell, \quad \ell > 0. \quad (4.2)$$

The rare event problem is the estimation of the probability $\lim_{\beta_\infty \rightarrow \infty} \pi(e^{\beta_\infty S^*}) = \pi(\{S^* \geq 1\})$. In practice, we set $\beta_\infty = 50$. The function $\Psi^*(x)$ is either

- the high-fidelity approximation of the PDE solution for parameter x defined in Section 4.1 (**model #RB**),
- the closed-form function (4.1) defined in Section 4.2 (**model #S**).

⁴With regard to the snapshots used to enrich the RB approximation space, it should be noted that the sampling of snapshots in the first iterations of ART corresponds to the popular *weak greedy* algorithm. However, once the target temperature has been reached (j_0 times), the weak greedy algorithm is replaced by drawing snapshots according to the importance distribution.

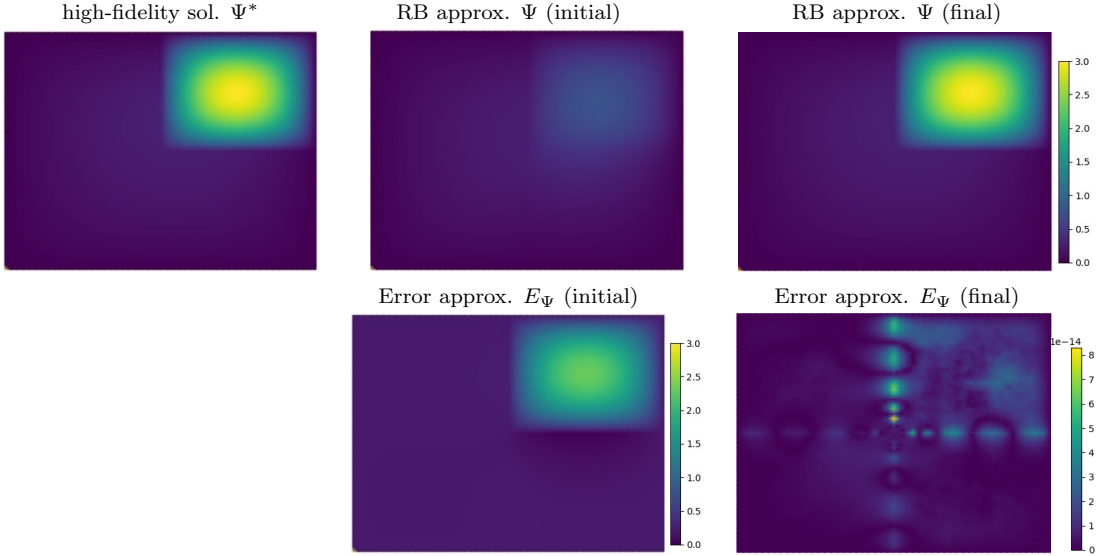


FIG. 4.2. **Problems #RE with model #RB.** Illustration of the need for an adaptive approach to construct relevant reduced-score approximations of the PDE solution Ψ^* . The RB approximation Ψ and its error $|\Psi - \Psi^*|$ are displayed for a sample X in the rare event $\{x : \|\Psi^*(x)\|_\infty \geq 3\}$, at the start and end of the algorithm (notice the different error scales in the colorbar).

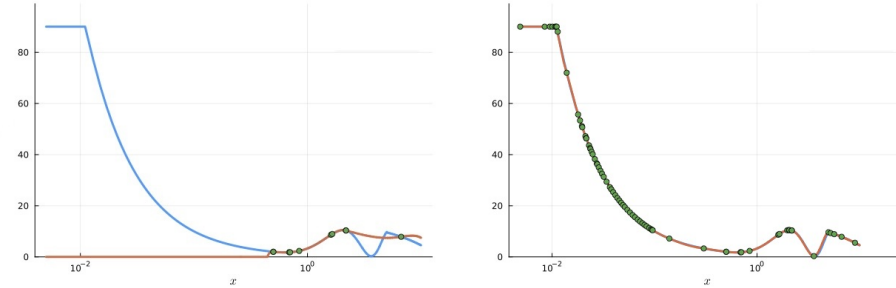


FIG. 4.3. **Problem #RE with model #S.** Illustration of the need for an adaptive approach to construct relevant approximations of Ψ^* as defined in (4.1). Function Ψ^* (blue curve), its spline approximation Ψ (red curve) and the snapshots used to build it (green dots) are displayed at the start and end of the algorithm. The rare event is $\{x : \Psi^*(x) \geq 90\}$ (the “plateau” around $x = 1e - 2$).

Let $p^* = \pi(e^{\beta_\infty S^*})$ denote the target rare event probability. In the case of model #RB, this probability corresponds to the occurrence of reaching the critical temperature ℓ , e.g. a melting temperature, on average in the domain ($p = 1$) or in some point of the domain ($p = \infty$). For this model, we perform standard adaptive SMC simulation with a score based on the high-fidelity solution $\Psi^*(x)$ and using 1000 particles. By averaging over 10 runs, we obtain after an intensive calculation empirical approximations of the target rare event probabilities: $p^* \approx 6.07e - 04$ (for $p = 1$ and $\ell = 0.5$) and $p^* \approx 1.47e - 04$ (for $p = \infty$ and $\ell = 3$). In the case of model #S, the target probability $p^* \approx 2.18e - 8$ (for $\ell = 90$) is closed-form and analytically computable. We use the ART algorithm to estimate probability p^* , evaluating the unnormalized estimates $\hat{\gamma}_{\beta_\infty, IS}^{N, H}(\mathbf{1})$ and $\hat{\gamma}_{\beta_\infty, RSMC}^{N, H}(\mathbf{1})$ defined in (3.7) and (3.8).

As needed by the ART algorithm, we design the error bound $E(x)$ on the absolute error on the score $|S^*(x) - S(x)|$, relying on the available error quantification. More

precisely, in the case of model #S, as $\Psi(x) \in \mathbb{R}$ and as we have assumed the error is directly available, we simply have that $|S^*(x) - S(x)| \leq |\Psi(x) - \Psi^*(x)| = E(x)$. In the case of model #RB with $p = 1$, an error function defined with the a posteriori error $\Delta_\Psi(x)$ as $E(x) = \pi^{-1}\Delta_\Psi(x)$ satisfies

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

where we have used the definition of the score function, the triangular inequality, the Cauchy-Schwarz inequality, the Poincaré inequality with the constant π^{-1} on the unit square domain and finally the a posteriori error inequality. 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 $p = \infty$, one can not in general rigorously derive an upper bound on the error relying on the a posteriori estimate $\Delta_\Psi(x)$. However, as the ART algorithm has the advantage of not requiring a rigorous error bound, we simply define the approximate error bound as

$$|S^*(x) - S(x)| \approx \Delta_\Psi(x) = E(x). \quad (4.3)$$

Figure 4.2 illustrates a high-fidelity approximation related to the rare event (*i.e.*, $\{x : S^*(x) \geq 1\} = \{x : \|\Psi^*(x)\|_\infty \geq 3\}$) obtained by our algorithm in the context of model #RB. We use an initial RB approximation built with 5 snapshots drawn according to the reference density. The same figure shows that this initial RB approximation is insufficiently accurate to discriminate this rare event since $\{x : \|\Psi(x)\|_\infty \geq \ell\} = \emptyset$. Figure 4.3 shows that similarly, in the case of model #S, the initial spline approximation built with 10 snapshots is insufficiently accurate to discriminate the rare event.

For these low-temperature problems, ART performance is quantified by evaluating the relative expected square error as a function of the expected cost over R independent runs. More specifically, let \hat{p}_r be a short-hand notation for either the estimate $\hat{\gamma}_{\beta_\infty, IS}^{N, H}(\mathbf{1})$ or $\hat{\gamma}_{\beta_\infty, RSMC}^{N, H}(\mathbf{1})$ at the r -th run. After exhausting the budget K of snapshots, the relative expected square error of the IS or RSMC estimator is approximated over runs by the empirical average

$$\frac{1}{R} \sum_{r=1}^R \frac{(\hat{p}_r - p^*)^2}{p^{*2}}, \quad (4.4)$$

and their expected cost (in units of the cost of evaluating S^*) is approximated by

$$K + \frac{1}{R} \sum_{r=1, k=1}^{R, K} N_r^{(k)} \text{gain}^{(k)} \quad (4.5)$$

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 ART algorithm.




FIG. 4.4. **Problem #BI with model #RB.** Illustration of the observations of the inverse problem. The high-fidelity solution Ψ^* for a sample X in the rare event set (obtained at the end of the algorithm for the sampling problem #RE) is observed on a random set of points up to Gaussian noise.

4.3.2. Problem #BI: Bayesian inverse problem. In this numerical experiment, we want to sample a Bayesian posterior distribution $\eta_{\beta_\infty}^*$ of the form (1.1) at the inverse temperature $\beta_\infty = 1$, for π being the log-normal prior distribution of model #RB, and given a set of observations following a Gaussian likelihood model. Let the supremum norm of the PDE solution and of its RB approximation evaluated at the diffusion parameter vector x be denoted respectively by the functions $\varphi^*(\cdot) = \|\Psi^*(\cdot)\|_\infty$ and $\varphi(\cdot) = \|\Psi(\cdot)\|_\infty$. The objective of this Bayesian sampling problem is to estimate the posterior distribution of X , given the PDE solution on a finite set of points corrupted by noise. For evaluation purposes, we will focus in this experiment in the distribution of the scalar observables $\varphi^*(X)$ and $\varphi(X)$. In order to generate the observations, we consider the high-fidelity PDE solution $\Psi^*(x^*)$ evaluated at x^* , where the latter parameter is sampled in the tail of the prior distribution π (using a sample of the low-temperature simulation experiment, see Section 4.3.1). Let $\Psi_z^* \in \mathbb{R}$ be the PDE solution Ψ^* taken at point $z \in \Omega$. Given a finite set of points $\mathcal{Z} \subset \Omega$, the set $\{\Psi_z^*(x^*)\}_{z \in \mathcal{Z}}$ is corrupted by an independent Gaussian noise of variance σ_{obs}^2 to generate the set of observations $\{\Psi_z^{obs}\}_{z \in \mathcal{Z}}$. The log likelihood of the Gaussian observation model is given by the score function

$$S^*(x) = -\frac{1}{\sigma_{obs}^2} \sum_{z \in \Omega_{obs}} (\Psi_z^*(x) - \Psi_z^{obs})^2.$$

We perform standard adaptive SMC simulation with a score based on the high-fidelity solution $\Psi^*(x)$ and using 500 particles. By averaging over 10 runs, we obtain after an intensive calculation an empirical reference approximation of the distribution of $\varphi^*(X)$, whose cumulative distribution function will be denoted by $F_{\varphi^*}(x)$. We perform R independent runs of the ART algorithm to estimate the latter function, using the normalized distribution estimates defined in (3.9). More precisely, $F_{\varphi^*}(x)$ is approximated using the IS estimate $\hat{\eta}_{\beta_\infty, IS}^{N, H}$ where the sample X_{k+1} is replaced by $\varphi^*(X_{k+1})$, or using the RSMC estimate $\hat{\eta}_{\beta_\infty, RSMC}^{N, H}$ where the N -sample $(\Xi_1^{(k)}, \dots, \Xi_N^{(k)})$ is replaced by $(\varphi(\Xi_1^{(k)}), \dots, \varphi(\Xi_N^{(k)}))$. We will adopt in the following the short-hand notation $F_{\varphi^*, r}^{N, H}$ for the cumulative distribution function related to the IS or RSMC estimates of $F_{\varphi^*}(x)$, obtained at the r -th run.

In this experiment, we rely on the reduced score $S(x) = -\sum_{z \in \Omega_{obs}} (\Psi_z(x) - \Psi_z^{obs})^2 / \sigma_{obs}^2$ where $\Psi_z(x)$ is the value of the RB approximation $\Psi(x)$ at the point

$z \in \Omega$. Using (4.3), an approximate majoration is

$$|S^*(x) - S(x)| = \frac{1}{\sigma_{obs}^2} \left| \sum_{z \in \Omega_{obs}} (\Psi_z(x) - \Psi_z^*(x))(\Psi_z(x) + \Psi_z^*(x) - 2\Psi_z^{obs}) \right| \approx E(x),$$

with the bound

$$E(x) = \frac{1}{\sigma_{obs}^2} \sum_{z \in \Omega_{obs}} E_\psi(x)(E_\psi(x) + 2|\Psi_z(x) - \Psi_z^{obs}|).$$

In this Bayesian sampling experiment, ART performance is quantified by evaluating the expected Kolmogorov-Smirnov statistic for the cumulative distribution function $F_{\varphi^*}(x)$

$$\frac{1}{R} \sum_{r=1}^R \sup_x |F_{\varphi,r}^{N,H}(x) - F_{\varphi^*}(x)|, \quad (4.6)$$

as a function of the expected cost (4.5).

4.4. Numerical results.

4.4.1. Experimental setup for evaluation. ART is configured as follows for the different problems and models. The mutation proposal is based on an Ornstein-Uhlenbeck [19] adaptive state movement. The worst-case log cost threshold c_1 , the adaptive tempering parameter c_2 , the snapshot budget of K , the trigger parameter j_0 , the number M of Markov transitions at each mutation step, the number R of runs and the number N of particles are given, for each experiment, in the table below.

	c_1	c_2	K	j_0	M	R	$N \in$
problem #RE & model #RB	1e-2	1e-3	200	5	30	50	[1e2, 3e3]
problem #RE & model #S	1e-3	1e-3	200	5	100	100	[1e3, 1e4]
problem #BI & model #RB	1e-2	1e-3	200	5	30	100	[1e1, 5e2]

For problem #BI (which uses model #RB), we draw uniformly on the square Ω a set \mathcal{Z} of 20 random points z , where we observe the solution Ψ_z^* up to a Gaussian noise of variance $1e - 2$. The observations are displayed in Figure 4.4.

For problems #RB (resp. for problem #BI), we evaluate the empirical relative square error (4.4) (resp. the expected Kolmogorov-Smirnov statistic (4.6)) versus the average cost and compare to results obtained for several dozen of runs with the true score S^* and for a varying number of particles, using either the standard AMS algorithm [8] (except for problem #BI, which can not be solved by AMS) or adaptive SMC simulation [3]. In the case of problem #RE, we also compare to performances obtained for several dozen of runs of the ARMS algorithm [11] with a varying number of particles.

4.4.2. Rare event simulation (Problem #RE). Figure 4.5 and 4.6 compare the performances of the ART algorithm with those of the state-of-the-art methods.

We note that the reduced sampling methods (ARMS and ART's IS estimate) share similar performances in the case of our rare event simulations. We nevertheless remark that depending on the model, there exists a slight gain or loss of the order of a few units for ART in comparison to ARMS (a gain in the case of model #RB for $p = 1$ and a loss in the case of model #S). Discriminating the most performant reduced method (ART or ARMS) is not obvious as performances appear in our experiments to be model

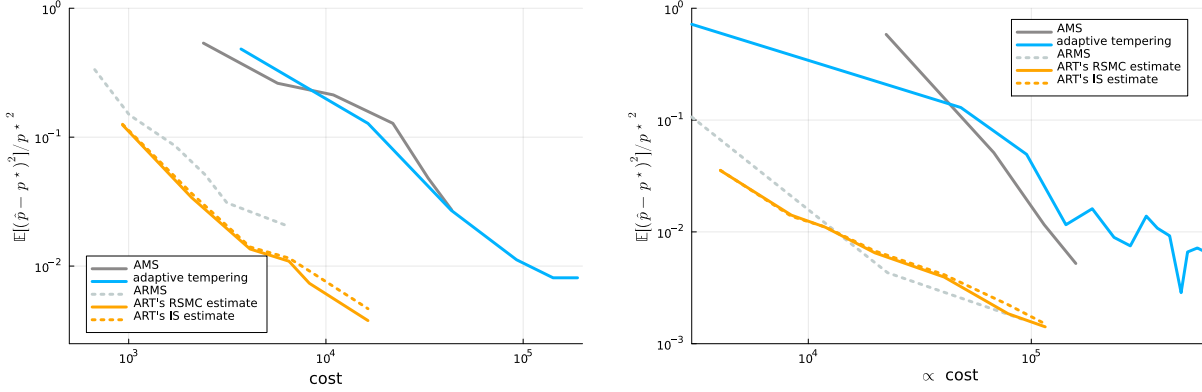


FIG. 4.5. *Problem #RE (rare event simulation) & model #RB (PDE-based score).* Relative expected square error as a function of the expected cost for ART's IS and RSMC estimates, for a score function (4.2) with $p = 1$ (left) and $p = \infty$ (right), and for a varying number of particles N . Comparison to the performance of adaptive tempering with S^* and to results of AMS and ARMS computed in [11].

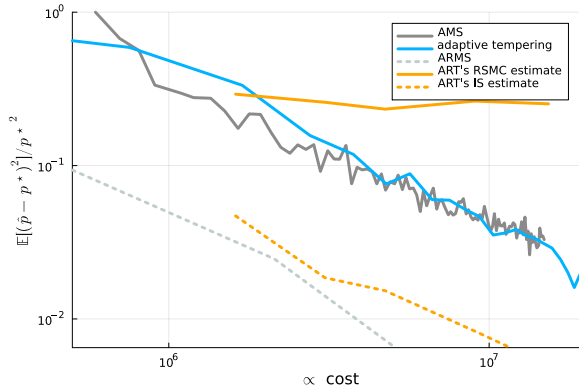


FIG. 4.6. *Problem #RE (rare event simulation) & model #S (closed-form score).* Relative expected square error as a function of the expected cost for ART's IS and RSMC estimates for a varying number of particles N . Comparison to the performance of adaptive tempering with S^* and to results of AMS and ARMS computed in [11].

dependent and sensitive to the setting of the free parameters (given in table above). However, we emphasize that the ART algorithm has the advantage of being simpler to implement as it does not require the use of the nesting and domination conditions of the ARMS algorithm. Furthermore, unlike ARMS, to maintain consistency (at least in an idealized case) ART only requires the use of an approximate error bound in cases where a rigorous one is not available (*e.g.*, in the case of model #RB with $p = \infty$). Last but not least, the performances of the two reduced sampling methods are far beyond the standard ones in terms of square error times cost: we observe a gain of the order of a decade or two of the expected square error for a given cost.

Although ART's RSMC and IS estimates are almost indistinguishable for model #RB, we note that ART's RSMC estimate (which relies solely on the reduced score) deteriorates in the case of model #S, the square error tending to be constant with respect to the cost. For this case, using the spline approximation S rather than the true

score S^* induces a significant bias. Indeed, 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 $\{x \in \mathcal{X} : S^*(x) \geq 1\}$ are to be excluded from the set $\{x \in \mathcal{X} : S(x) \geq 1\}$, leading to a negative bias. On the contrary, we observe that ART's IS estimate is not biased and retains its good performance in this case, independently of the problem of convergence of S to S^* .

For both standard and reduced estimates (ARMS and ART's IS estimate), we observe approximately a unitary negative slope (on a logarithmic scale) of the variance as a function of cost, implying that the product of variance and cost is constant. Such a behavior suggests asymptotic normality of the proposed IS estimate, as it has been demonstrated for AMS [9].

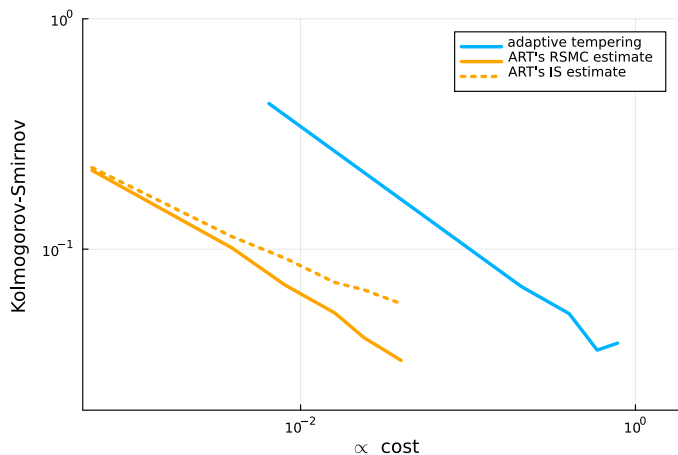


FIG. 4.7. **Problem #BI (Bayesian inversion) & model #RB (PDE-based score).** Kolmogorov-Smirnov distance as a function of the expected cost for ART's IS and RSMC estimates for a varying number of particles N . Comparison to the performance of adaptive tempering with S^* .

4.4.3. Bayesian inversion (Problem #BI). We next turn to the Bayesian inverse problem. Figure 4.7 shows that the ART algorithm decreases the cost of simulation in comparison to adaptive tempering, of about a decade to obtain a similar sampling accuracy according to the Kolmogorov-Smirnov distance (4.6).

We observe a sub-unitary negative slope (of minus one half on a logarithmic scale) of this distance as a function of cost for ART's RSMC estimate, implying that the product of the square of the Kolmogorov-Smirnov distance and the cost is constant, and suggesting that the product of the estimator's variance and the cost is constant for any test function. Again this behavior supports the conjecture of asymptotic normality for the proposed estimate.

In contrast to the case of problem #RB with model #S, we find that ART's IS estimate performs slightly worse than ART's RSMC estimate. Indeed, its slope is less than the slope minus one half observed for the RSMC estimate. In this case, we conjecture that for sufficiently large N and a sufficiently accurate approximation S , the bias of the ART's IS and RSMC normalized estimates are negligible in comparison to their variances, and furthermore, the variance of ART's IS estimate will be greater than that of ART's RSMC estimate. Indeed, let's assume that $S^{(k)} \simeq S^*$ for iterations

$k > K_{j_0}$ (which is a plausible assumption based on the illustrations in figures 4.2 and 4.3). The RSMC estimate should then have a variance of at best $1/NH$ [15, Prop. 9.5.6] and a negligible bias [17]. On the other hand, based on the Appendix A, the IS estimate should have a variance of the order of $1/H$ and a negligible bias. Consequently, considering N tending to infinity and H fixed, the mean square error of the RSMC should tend towards zero with a speed at best inversely proportional to the cost (4.5) (of the order of KN), while the mean square error of the IS estimate will tend at best towards a constant (of the order of $1/H$).

In summary, our numerical experiments have demonstrated that a computational saving of over a decade is achieved with the ART algorithm, compared with state-of-the-art standard approaches. In addition, the product of variance and cost appears to be constant for these estimates in some regimes, suggesting asymptotic normality. The results also show that, for a particular problem, the RSMC estimate can be biased even in the case where the reduced score tends towards the true score. Conversely, on the Bayesian inverse problem, the RSMC estimate seems to perform slightly better than the IS estimate for a fixed budget of snapshots.

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

The general idea of ART is to perform iterative importance sampling of a sequence of target distributions of the form $\propto e^{\beta^{(k)} S^*} d\pi$ parametrized at iteration k by the inverse temperature $\beta^{(k)} \in \mathbb{R}_+$, using a sequences of proposals of the form $\mu_{\beta^{(k)}}^{(k)} \propto e^{\beta^{(k)} S^{(k)}} d\pi$, much cheaper to simulate with SMC. As the algorithm is iterated, the proposals are adapted by adjusting the inverse temperature $\beta^{(k)}$ 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 $\beta^{(k)}$ increases, the reduced score approximation should be refined using samples of the state space in regions closer and closer to the low-temperature event of interest. However, an essential ingredient in achieving this desirable convergence is to determine the first non-achievable inverse temperature $\beta^{(k)}$ suitable for the accuracy of the current reduced score approximation $S^{(k)}$. Justified by theoretical arguments, this problem is addressed in ART using an entropy constraint: $\beta^{(k)}$ is chosen such that the relative entropy between a surrogate target and the proposal corresponds to a given logarithmic cost of importance sampling. The surrogate target is constructed using an approximation of the quantification of the worst-case error associated with the reduced score approximation.

A final point addressed by ART is the adaptation of the proposal to a new approximation of the reduced score $S^{(k+1)}$. To achieve this, ART searches among past simulation indexes $k_{bridge} < k+1$ for an approximation $S^{(k_{bridge})}$ and an initial parameter $\beta_{bridge}(k+1)$ such that the relative entropy between the new proposal $\mu_{\beta_{bridge}(k+1)}^{(k+1)}$ at its initial inverse temperature and the past proposal $\mu_{\beta^{(k_{bridge})}}^{(k_{bridge})}$ at its critical inverse temperature is kept below a certain standard threshold and, in addition, such that relative entropy between the surrogate target and the proposal is below the given logarithmic cost of importance sampling. The algorithm then bridges the two proposals in a single SMC simulation step.

In addition to demonstrating their consistency in an idealized setting, we evaluate in our numerical experiments the empirical convergence of the two proposed estima-

tors, an importance sampling estimator and a reduced SMC estimator, in the context of rare-event simulation and for a Bayesian inverse problem based on a PDE. We exhibit empirically that we significantly reduce simulation cost of order of a decade while achieving a comparable accuracy.

Appendix A. Consistency. We show the consistency of the unnormalized IS estimator $\hat{\gamma}_{\beta_\infty}^H$ in the idealized setting, see Section 2.1.

PROPOSITION A.1. *For any test function φ , the estimator $\hat{\gamma}_{\beta_\infty}^H$ is unbiased, i.e., $\mathbb{E} \hat{\gamma}_{\beta_\infty}^H(\varphi) = \gamma_{\beta_\infty}^*(\varphi)$, and its relative variance is*

$$\frac{\text{Var} \left[\hat{\gamma}_{\beta_\infty}^H(\varphi) \right]}{\gamma_{\beta_\infty}^*(\varphi)^2} = \frac{1}{H^2} \sum_{k=K_{j_0}+1}^{K_{j_0}+H} \frac{\mathbb{E}(Z_{\beta^{(k)}}^{(k)} \pi(e^{2\beta_\infty S^* - \beta^{(k)} S^{(k)}} \varphi^2))}{\gamma_{\beta_\infty}^*(\varphi)^2} - 1,$$

where the expectation is with respect to the random variable $\beta^{(k)} S^{(k)}$.

Proof. We first show that the random sequence

$$H \mapsto M_H^\varphi \stackrel{\text{def}}{=} H(\hat{\gamma}_{\beta_\infty}^H(\varphi) - \gamma_{\beta_\infty}^*(\varphi)),$$

is a martingale with respect to the filtration generated by the snapshots $\mathcal{F}_H = \sigma(X_1, \dots, X_{K_{j_0}+H})$. Indeed, defining the density ratio

$$f_{\beta_\infty}^{(k)}(x) = Z_{\beta^{(k)}}^{(k)} e^{\beta_\infty S^*(x) - \beta^{(k)} S^{(k)}(x)},$$

the IS estimator (2.1) is rewritten as

$$\hat{\gamma}_{\beta_\infty}^H(\varphi) = \frac{1}{H} \sum_{k=K_{j_0}+1}^{K_{j_0}+H} f_{\beta_\infty}^{(k)}(X_{k+1}) \varphi(X_{k+1}),$$

and we have

$$\begin{aligned} \mathbb{E}[M_H^\varphi | \mathcal{F}_{H-1}] &= \mathbb{E} \left[\sum_{k=K_{j_0}+1}^{K_{j_0}+H} f_{\beta_\infty}^{(k)}(X_{k+1}) \varphi(X_{k+1}) - H \gamma_{\beta_\infty}^*(\varphi) \middle| \mathcal{F}_{H-1} \right], \\ &= M_{H-1}^\varphi + \mathbb{E}[f_{\beta_\infty}^{(K_{j_0}+H)}(X_{K_{j_0}+H+1}) \varphi(X_{K_{j_0}+H+1}) | \mathcal{F}_{H-1}] - \gamma_{\beta_\infty}^*(\varphi), \\ &= M_{H-1}^\varphi, \end{aligned}$$

where the last equality is due to the fact that the snapshot $X_{K_{j_0}+H+1}$ is drawn according to the proposal

$$\mu_{\beta^{(K_{j_0}+H)}}^{(K_{j_0}+H)} = e^{\beta^{(K_{j_0}+H)} S^{(K_{j_0}+H)}} d\pi / Z_{\beta^{(K_{j_0}+H)}}^{(K_{j_0}+H)}.$$

Now, since M_H^φ is a martingale, the estimator is unbiased $\mathbb{E} \hat{\gamma}_{\beta_\infty}^H(\varphi) = \gamma_{\beta_\infty}^*(\varphi)$, and its variance is given by the average of the quadratic variations of the sequence. The $K_{j_0} + H$ -th element of the sequence, the expectation of the quadratic variation is (in this ideal setup the normalization constant can be computed exactly given some $\beta^{(k)}$ and $S^{(k)}$)

$$\text{Var}_{\mu_{\beta^{(K_{j_0}+H)}}^{(K_{j_0}+H)}} \left(\frac{d\gamma_{\beta_\infty}^*}{d\mu_{\beta^{(K_{j_0}+H)}}^{(K_{j_0}+H)}}(X) \varphi(X) \middle| \mathcal{F}_{H-1} \right) = Z_{\beta^{(k)}}^{(k)} \pi(e^{2\beta_\infty S^* - \beta^{(k)} S^{(k)}} \varphi^2) - \gamma_{\beta_\infty}^*(\varphi)^2,$$

where $\beta^{(k)} S^{(k)}$ is a variable depending on the randomness of the previous snapshots. Therefore, the variance is

$$\text{Var} [\hat{\gamma}_{\beta_{\infty}}^H(\varphi)] = \frac{1}{H^2} \sum_{k=K_{j_0}+1}^{K_{j_0}+H} \mathbb{E}(Z_{\beta^{(k)}}^{(k)} \pi(e^{2\beta_{\infty} S^* - \beta^{(k)} S^{(k)}} \varphi^2)) - \gamma_{\beta_{\infty}}^*(\varphi)^2. \quad \square$$

REFERENCES

- [1] Julien Bect, Ling Li, and Emmanuel Vazquez. Bayesian subset simulation. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):762–786, 2017.
- [2] Alexandros Beskos, Dan Crisan, and Ajay Jasra. On the stability of sequential Monte Carlo methods in high dimensions. *The Annals of Applied Probability*, 24(4):1396 – 1445, 2014.
- [3] Alexandros Beskos, Ajay Jasra, Nikolas Kantas, and Alexandre Thiery. On the convergence of adaptive sequential monte carlo methods. *The Annals of Applied Probability*, 26(2):1111–1146, 2016.
- [4] Alexandros Beskos, Ajay Jasra, Kody Law, Raul Tempone, and Yan Zhou. Multilevel sequential monte carlo samplers. *Stochastic Processes and their Applications*, 127(5):1417–1440, 2017.
- [5] Alexandros Beskos, Ajay Jasra, Ege A Muzaffer, and Andrew M Stuart. Sequential monte carlo methods for bayesian elliptic inverse problems. *Statistics and Computing*, 25:727–737, 2015.
- [6] Tan Bui-Thanh and Mark Girolami. Solving large-scale pde-constrained bayesian inverse problems with riemann manifold hamiltonian monte carlo. *Inverse Problems*, 30(11):114014, 2014.
- [7] Daniela Calvetti, Matthew Dunlop, Erkki Somersalo, and Andrew Stuart. Iterative updating of model error for bayesian inversion. *Inverse Problems*, 34(2):025008, 2018.
- [8] Frédéric Cérou, Pierre Del Moral, Teddy Furon, and Arnaud Guyader. Sequential Monte Carlo for rare event estimation. *Stat. Comput.*, 22(3):795–808, 2012.
- [9] Frédéric Cérou and Arnaud Guyader. Fluctuation analysis of adaptive multilevel splitting. *Annals of Applied Probability*, 26(6):3319–3380, December 2016.
- [10] Frédéric Cérou, Patrick Héas, and Mathias Rousset. Entropy minimizing distributions are worst-case optimal importance proposals. *arXiv preprint arXiv:2212.04292*, 2022.
- [11] Frédéric Cérou, Patrick Héas, and Mathias Rousset. Adaptive reduced multilevel splitting. *arXiv preprint arXiv:2312.15256*, 2024.
- [12] Sourav Chatterjee and Persi Diaconis. The sample size required in importance sampling. *The Annals of Applied Probability*, 28(2):1099–1135, 2018.
- [13] Patrick R Conrad, Andrew D Davis, Youssef M Marzouk, Natesh S Pillai, and Aaron Smith. Parallel local approximation mcmc for expensive models. *SIAM/ASA Journal on Uncertainty Quantification*, 6(1):339–373, 2018.
- [14] S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White. MCMC Methods for Functions: Modifying Old Algorithms to Make Them Faster. *Statistical Science*, 28(3):424 – 446, 2013.
- [15] Pierre Del Moral. Mean field simulation for monte carlo integration. *Monographs on Statistics and Applied Probability*, 126:26, 2013.
- [16] 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.
- [17] Pierre Del Moral, Arnaud Doucet, and Gareth W Peters. Sharp propagation of chaos estimates for feynman–kac particle models. *Theory of Probability & Its Applications*, 51(3):459–485, 2007.
- [18] Randal Douc and Olivier Cappé. Comparison of resampling schemes for particle filtering. In *ISPA 2005. Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis, 2005.*, pages 64–69. Ieee, 2005.
- [19] 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.
- [20] Andrew Gelman and Xiao-Li Meng. Simulating normalizing constants: from importance sampling to bridge sampling to path sampling. *Statistical Science*, 13(2):163 – 185, 1998.
- [21] Patrick Héas. Selecting reduced models in the cross-entropy method. *SIAM/ASA Journal on Uncertainty Quantification*, 8(2):511–538, 2020.
- [22] Patrick Héas, Frédéric Cérou, and Mathias Rousset. Chilled sampling for uncertainty quantification: a motivation from a meteorological inverse problem. *Inverse Problems*, 40(2):025004, 2023.

- [23] Nikolas Kantas, Alexandros Beskos, and Ajay Jasra. Sequential monte carlo methods for high-dimensional inverse problems: A case study for the navier–stokes equations. SIAM/ASA Journal on Uncertainty Quantification, 2(1):464–489, 2014.
- [24] Jonas Latz, Juan P Madrigal-Cianci, Fabio Nobile, and Raúl Tempone. Generalized parallel tempering on bayesian inverse problems. Statistics and Computing, 31(5):67, 2021.
- [25] Jonas Latz, Iason Papaioannou, and Elisabeth Ullmann. Multilevel sequential2 monte carlo for bayesian inverse problems. Journal of Computational Physics, 368:154–178, 2018.
- [26] Jinglai Li and Youssef M Marzouk. Adaptive construction of surrogates for the bayesian solution of inverse problems. SIAM Journal on Scientific Computing, 36(3):A1163–A1186, 2014.
- [27] 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.
- [28] Maliki Moustapha, Stefano Marelli, and Bruno Sudret. Active learning for structural reliability: Survey, general framework and benchmark. Structural Safety, 96:102174, 2022.
- [29] Alfio Quarteroni, Andrea Manzoni, and Federico Negri. Reduced basis methods for partial differential equations: an introduction, volume 92. Springer, 2015.
- [30] 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.