

Offline-Enhanced Reduced Basis Method Through Adaptive Construction of the Surrogate Training Set

Jiahua Jiang¹  · Yanlai Chen¹ · Akil Narayan²

Received: 16 March 2017 / Revised: 18 August 2017 / Accepted: 29 August 2017 /
Published online: 21 September 2017
© Springer Science+Business Media, LLC 2017

Abstract The reduced basis method (RBM) is a popular certified model reduction approach for solving parametrized partial differential equations. One critical stage of the *offline* portion of the algorithm is a greedy algorithm, requiring maximization of an error estimate over parameter space. In practice this maximization is usually performed by replacing the parameter domain continuum with a discrete “training” set. When the dimension of parameter space is large, it is necessary to significantly increase the size of this training set in order to effectively search parameter space. Large training sets diminish the attractiveness of RBM algorithms since this proportionally increases the cost of the offline phase. In this work we propose novel strategies for offline RBM algorithms that mitigate the computational difficulty of maximizing error estimates over a training set. The main idea is to identify a subset of the training set, a “surrogate training set” (STS), on which to perform greedy algorithms. The STS we construct is much smaller in size than the full training set, yet our examples suggest that it is accurate enough to induce the solution manifold of interest at the current offline RBM iteration. We propose two algorithms to construct the STS: our first algorithm, the successive maximization method, is inspired by inverse transform sampling for non-standard univariate probability distributions. The second constructs an STS by identifying

Y. Chen and J. Jiang were partially supported by National Science Foundation Grants DMS-1216928 and DMS-1719698. A. Narayan is partially supported by NSF DMS-1552238, AFOSR FA9550-15-1-0467, and DARPA N660011524053.

✉ Jiahua Jiang
jjjiang@umassd.edu

Yanlai Chen
yanlai.chen@umassd.edu

Akil Narayan
akil@sci.utah.edu

¹ Department of Mathematics, University of Massachusetts Dartmouth, 285 Old Westport Road, North Dartmouth, MA 02747, USA

² Scientific Computing and Imaging (SCI) Institute and Department of Mathematics, University of Utah, 72 S Campus Drive, Salt Lake City, UT 84112, USA

pivots in the Cholesky decomposition of an approximate error correlation matrix. We demonstrate the algorithm through numerical experiments, showing that it is capable of accelerating offline RBM procedures without degrading accuracy, assuming that the solution manifold has rapidly decaying Kolmogorov width.

Keywords Reduced basis method · Surrogate parameter domain · Cholesky decomposition · Adaptivity · Greedy algorithm

1 Introduction

Increasing computer power and computing availability makes the simulation of large-scale or high-dimensional numerical problems modeling complex phenomena more accessible, but any reduction in computational effort is still required for sufficiently onerous many-query computations and repeated output evaluations for different values of some inputs of interest. Much recent research has concentrated on schemes to accelerate computational tools for such many-query or parameterized problems; these schemes include Proper Orthogonal Decomposition (POD) [3,38], balanced truncation method [13,24], Krylov Subspace methods [10,34] and the Reduced Basis Method (RBM) [4,33]. The basic idea behind each of these model order reduction techniques is to iteratively project an associated large algebraic system to a small system that can effectively capture most of the information carried by the original model. In the context of parameterized partial differential equations (PDE), the fundamental reason that such model reduction approaches are accurate is that, for many PDE's of interest, the solution manifold induced by the parametric variation has rapidly decaying Kolmogorov width [28].

Among the strategies listed above, one of the most appealing methods is the Reduced Basis Method. RBM seeks to parametrize the random inputs and select the most representative points in the parameter space by means of a greedy algorithm that leverages an *a posteriori* error estimate. RBM algorithm are split into *offline* and *online* stages. During the offline stage, the parameter dependence is examined and the greedy algorithm is used to judiciously select a small number of parameter values on which the full, expensive PDE solver is solved. The solutions on this small parameter set are called *snapshots*. During the online stage, an approximate solution for any new parameter value is efficiently computed via a linear combination of the offline-computed snapshots. This linear combination can usually be accomplished with orders of magnitude less effort than a full PDE solve. Thus, RBM invests significant computational effort in an offline stage so that the online stage is efficient [6,25,29,30,32].

The Reduced Basis Method was initially introduced for structure analysis [26] and recently has been undergoing vast development in theory [11,12,23,27,33] and applied to many engineering problems [5,7,9,22,31,32]. If the input parameters to the PDE are random, then RBM is philosophically similar to stochastic collocation [1], but uses an adaptive sampling criterion dictated by the PDE's *a posteriori* error estimate. It has been shown that RBM can help delay the curse of dimensionality when solving parameterized problems in uncertainty quantification whenever the solution manifold can be well approximated by a low dimensional space [21].

RBM is motivated by the observation that the set of all solutions to the parametrized problem under variation of the parameter can be well approximated by linear combinations of a low number of basis functions. Early RBM research concentrated on problems with a

low-dimensional parameter due to the lack of effective tools to sample the *a posteriori* error estimate over high-dimensional spaces [33]. An effective procedure for greedily selecting parameter values for use in the RBM procedure must simultaneously leverage the structure of the error estimate along with efficient methods for searching over high-dimensional spaces. Some recent effort in the RBM framework has been devoted to inexpensive computation of the *a posteriori* error estimate along with effective sampling strategies [37].

The portion of the RBM algorithm most relevant in the context of this article is in the offline stage: Find a parameter value that maximizes a given (computable) error estimate. This maximization is usually accomplished in the computational setting by replacing the parameter domain continuum by a large, discrete set called the *training set*. Even this naïve procedure requires us to compute the value of the error estimate at every point in the training set. Thus, the work required scales proportionally to the training set size. When the parameter is high-dimensional, the size of the training set must be very large if one seeks to search over all regions of parameter space. This onerous cost of the offline stage debilitates RBM in this scenario. Thus, assuming the training set must be large, a more sophisticated scheme for maximizing the error estimate must be employed.

This is not the first attempt to enhance the offline phase of RBM. Existing approaches include multistage greedy algorithm [35], training set adaptivity and adaptive parameter domain partitioning [16], random greedy sampling [18], greedy sampling acceleration through nonlinear optimization [36]. Problems remain despite these developments. For example, the randomization in part of the greedy procedure reduces the stability of the whole algorithm. And the choice of the tuning parameter may make the improved methods excessively empirical. Our approach in this paper differentiates itself in that it is deterministic, more stable, more predictable with respect to the tuning parameters. Most importantly, it is the first attempt to utilize the intermediate results that were discarded (i.e. the error estimates for the whole training set) to achieve the enhancement. Indeed, we propose two novel strategies for mitigating the cost of searching over a training set of large size. The essential idea in both approaches is to perform some computational analysis on the *a posteriori* error estimate in order to construct a *surrogate training set*, a subset of the original training set with a much smaller size, that can effectively predict the general trend of the error estimate. The construction of this surrogate training set must be periodically repeated during the iterative phase of the offline RBM algorithm. Ideally, we want to decrease the computational burden of the offline algorithm without lowering the fidelity of the RBM procedure. The following qualitative characteristics are the guiding desiderata for construction of the surrogate training set:

1. the information used for the surrogate training set construction should be inexpensive to obtain
2. the parametric variation on the surrogate training set should be representative of that in the original training set
3. the size of the surrogate training set should be significantly smaller than that of the original training set

Our proposed offline-enhanced RBM strategies are as follows:

1. Successive Maximization Method for Reduced Basis Method (SMM-RBM)—We construct an empirical cumulative distribution function of the *a posteriori* error estimate on the training set, and deterministically subsample a surrogate training set according to this distribution.

Table 1 Notation used throughout this article

μ	Parameter in $\mathcal{D} \subseteq \mathbb{R}^p$
$u(\mu)$	Function-valued solution of a parameterized PDE
\mathcal{N}	Degrees of freedom (DoF) in PDE “truth” solver
$u^{\mathcal{N}}(\mu)$	Truth solution (finite-dimensional)
N	Number of reduced basis snapshots, $N \ll \mathcal{N}$
N_{max}	Terminated number of reduced bases
μ^j	“Snapshot” parameter values, $j = 1, \dots, N$
$X_N^{\mathcal{N}}$	Span of $u^{\mathcal{N}}(\mu^k)$ for $k = 1, \dots, N$
$u_N^{\mathcal{N}}(\mu)$	Reduced basis solution, $u_N^{\mathcal{N}} \in X_N^{\mathcal{N}}$
$e_N(\mu)$	Reduced basis solution error, equals $u^{\mathcal{N}}(\mu) - u_N^{\mathcal{N}}(\mu)$
Ξ_{train}	Parameter training set, a finite subset of \mathcal{D}
$\Delta_N(\mu)$	Error estimate (upper bound) for $\ e_N(\mu)\ $
ϵ_{tol}	Error estimate stopping tolerance in greedy sweep

2. Cholesky Decomposition Reduced Basis Method (CD-RBM)—An approximate correlation matrix (Gramian) of errors over the training set is computed, and the pivots in a pivoted Cholesky decomposition [17] identify the surrogate training set.

We note that both of our strategies are empirical in nature. In particular, it is relatively easy to manufacture error estimate data so that our construction of a surrogate training set does not accurately capture the parameter variation over the full training set. One of the main observations we make in our numerical results section is that such an adversarial situation does not occur for the parameterized PDEs that we investigate. Our procedure also features some robustness: a poorly constructed surrogate training set does not adversely affect either the efficiency or the accuracy of the RBM simulation.

The remainder of this paper is organized as follows. A parametrized PDE with random input data is set up with appropriate assumptions on the PDE operator in Sect. 2. The general framework and properties of the Reduced Basis Method are likewise introduced in Sect. 2. Section 3 is devoted to the development of our novel offline-enhanced RBM methods, consisting of SMM-RBM and CD-RBM. A rough complexity analysis of these methods is given in Sect. 3.3. Our numerical examples are shown in Sect. 4.

2 Background

In this section, we introduce the Reduced Basis Method (RBM) in its classical form; much of this is standard in the RBM literature. The reader familiar with RBM methodology may skip this section, using Table 1 as a reference for our notation.

2.1 Problem Setting

Let $\mathcal{D} \subset \mathbb{R}^p$ be the range of variation of a p -dimensional parameter and $\Omega \subset \mathbb{R}^d$ (for $d = 2$ or 3) a bounded spatial domain. We consider the following parametrized problem: Given $\mu \in \mathcal{D}$, the goal is to evaluate the output of interest

$$s(\mu) = \ell(u(\mu)), \tag{2.1}$$

where the function $u(\boldsymbol{\mu}) \in X$ satisfies

$$a(u(\boldsymbol{\mu}), v, \boldsymbol{\mu}) = f(v; \boldsymbol{\mu}), \quad v \in X, \tag{2.2}$$

which is a parametric partial differential equation (pPDE) written in a weak form. Here $X = X(\Omega)$ is a Hilbert space satisfying $H_0^1(\Omega) \subset X(\Omega) \subset H^1(\Omega)$. We denote by $(\cdot, \cdot)_X$ the inner product associated with the space X , whose induced norm $\|\cdot\|_X = \sqrt{(\cdot, \cdot)_X}$ is equivalent to the usual $H^1(\Omega)$ norm. We assume that $a(\cdot, \cdot; \boldsymbol{\mu}) : X \times X \rightarrow \mathbb{R}$ is continuous and coercive over X for all $\boldsymbol{\mu}$ in \mathcal{D} , that is,

$$\gamma(\boldsymbol{\mu}) := \sup_{w \in X} \sup_{v \in X} \frac{a(w, v; \boldsymbol{\mu})}{\|w\|_X \|v\|_X} < \infty, \quad \forall \boldsymbol{\mu} \in \mathcal{D}, \tag{2.3a}$$

$$\alpha(\boldsymbol{\mu}) := \inf_{w \in X} \frac{a(w, w; \boldsymbol{\mu})}{\|w\|_X^2} \geq \alpha_0 > 0, \quad \forall \boldsymbol{\mu} \in \mathcal{D}. \tag{2.3b}$$

$f(\cdot)$ and $\ell(\cdot)$ are linear continuous functionals over X , and for simplicity we assume that ℓ is independent of $\boldsymbol{\mu}$.

We assume that $a(\cdot, \cdot; \boldsymbol{\mu})$ is “affine” with respect to functions of the parameter $\boldsymbol{\mu}$: there exist $\boldsymbol{\mu}$ -dependent coefficient functions $\Theta_a^q : \mathcal{D} \rightarrow \mathbb{R}$ for $q = 1, \dots, Q_a$, and corresponding continuous $\boldsymbol{\mu}$ -dependent bilinear forms $a^q(\cdot, \cdot) : X \times X \rightarrow \mathbb{R}$ such that

$$a(w, v; \boldsymbol{\mu}) = \sum_{q=1}^{Q_a} \Theta_a^q(\boldsymbol{\mu}) a^q(w, v). \tag{2.4}$$

This assumption of affine parameter dependence is common in the reduced basis literature [33], and remedies are available [2] when it is not satisfied.

Finally, we assume that there is a finite-dimensional discretization for the model problem (2.2): The solution space X is discretized by an \mathcal{N} -dimensional subspace $X^{\mathcal{N}}$ (i.e., $\dim(X^{\mathcal{N}}) = \mathcal{N}$) and (2.1) and (2.2) are discretized as

$$\begin{cases} \text{For } \boldsymbol{\mu} \in \mathcal{D}, \text{ solve} \\ s^{\mathcal{N}} = \ell(u^{\mathcal{N}}(\boldsymbol{\mu})) \text{ where } u^{\mathcal{N}}(\boldsymbol{\mu}) \in X^{\mathcal{N}} \text{ satisfies} \\ a(u^{\mathcal{N}}, v; \boldsymbol{\mu}) = f(v; \boldsymbol{\mu}) \quad \forall v \in X^{\mathcal{N}}. \end{cases} \tag{2.5}$$

The relevant quantities such as the coercivity constant (2.3b) are defined according to the discretization,

$$\alpha^{\mathcal{N}}(\boldsymbol{\mu}) = \inf_{w \in X^{\mathcal{N}}} \frac{a(w, w; \boldsymbol{\mu})}{\|w\|_X^2}, \quad \forall \boldsymbol{\mu} \in \mathcal{D}.$$

In the RBM literature, any discretization associated to \mathcal{N} is called a “truth” discretization. E.g., $u^{\mathcal{N}}$ is called the “truth solution”.

2.2 RBM Framework

We assume \mathcal{N} is large enough so that solving (2.5) gives highly accurate approximations for $\boldsymbol{\mu} \in \mathcal{D}$. However, large \mathcal{N} also means that solving (2.5) is expensive, and in many-query contexts (e.g., optimization) a direct approach to solving (2.5) is computationally infeasible. The situation is exacerbated when the $\boldsymbol{\mu} \mapsto u^{\mathcal{N}}(\boldsymbol{\mu})$ response is sought in a real-time fashion. The reduced basis method is a reliable model reduction tool for these scenarios.

This section presents a brief overview of the standard RBM algorithm. Given a finite training set of parameter samples $\Xi_{\text{train}} \subset \mathcal{D}$ as well as a prescribed maximum dimension N_{max} (usually $\ll \mathcal{N}$), we approximate the solution set

$$\{u(\boldsymbol{\mu}) \mid \boldsymbol{\mu} \in \mathcal{D}\} \subset X^{\mathcal{N}}$$

via an N -dimensional subspace of $X^{\mathcal{N}}$, with $N \leq N_{\max}$. In RBM, this is accomplished via the N -dimensional cardinal Lagrange space

$$X_N^{\mathcal{N}} := \text{span}\{u^{\mathcal{N}}(\boldsymbol{\mu}^n), 1 \leq n \leq N\}, \quad N = 1, \dots, N_{\max} \tag{2.6}$$

in a hierarchical manner by iteratively choosing samples $S_N = \{\boldsymbol{\mu}^1, \dots, \boldsymbol{\mu}^N\}$ from Ξ_{train} until N is large enough so that a prescribed accuracy tolerance ϵ_{tol} is met. The $u^{\mathcal{N}}(\boldsymbol{\mu}^n)$ for $1 \leq n \leq N$ are the so-called ‘‘snapshots’’, and are obtained by solving (2.5) with $\boldsymbol{\mu} = \boldsymbol{\mu}^n$.

It is obvious that both S_N and $X_N^{\mathcal{N}}$ are nested; that is, $S_1 \subset S_2 \subset \dots \subset S_{N_{\max}}$ and $X_1^{\mathcal{N}} \subset X_2^{\mathcal{N}} \subset \dots \subset X_{N_{\max}}^{\mathcal{N}} \subset X^{\mathcal{N}}$. This condition is fundamental in ensuring efficiency of the resulting RB approximation. Given $\boldsymbol{\mu} \in \mathcal{D}$, we seek a surrogate RB solution $u_N^{\mathcal{N}}(\boldsymbol{\mu})$ in the reduced basis space $X_N^{\mathcal{N}}$ for the truth approximation $u^{\mathcal{N}}(\boldsymbol{\mu})$ by solving the following reduced system

$$\begin{cases} \text{For } \boldsymbol{\mu} \in \mathcal{D}, \text{ evaluate} \\ s_N^{\mathcal{N}} = \ell(u_N^{\mathcal{N}}(\boldsymbol{\mu})) \text{ s.t. } u_N^{\mathcal{N}}(\boldsymbol{\mu}) \in X_N^{\mathcal{N}} \subset X^{\mathcal{N}} \text{ satisfies} \\ a(u_N^{\mathcal{N}}, v; \boldsymbol{\mu}) = f(v) \quad \forall v \in X_N^{\mathcal{N}}. \end{cases} \tag{2.7}$$

In comparison to the \mathcal{N} -dimensional system (2.5), the reduced system (2.7) is N -dimensional; when $N \ll \mathcal{N}$, this results in a significant computational savings. The Galerkin procedure in (2.7) selects the best solution in $X_N^{\mathcal{N}}$ satisfying the pPDE.¹ Then, the RB solution $u_N^{\mathcal{N}}(\boldsymbol{\mu})$ for any parameter $\boldsymbol{\mu} \in \mathcal{D}$ can be expressed as

$$u_N^{\mathcal{N}}(\boldsymbol{\mu}) = \sum_{m=1}^N u_{Nm}^{\mathcal{N}}(\boldsymbol{\mu}) u^{\mathcal{N}}(\boldsymbol{\mu}^m) \tag{2.8}$$

Here $\{u_{Nm}^{\mathcal{N}}(\boldsymbol{\mu})\}_{m=1}^N$ are the unknown RB coefficients that can be obtained by solving (2.7). Upon replacing the reduced basis solution in (2.7) by (2.8) and taking the $X_N^{\mathcal{N}}$ basis functions $v_n = u^{\mathcal{N}}(\boldsymbol{\mu}^n)$, $1 \leq n \leq N$, as the test functions for Galerkin approximation, we obtain the RB ‘‘stiffness’’ equations

$$\sum_{m=1}^N a(v_m, v_n; \boldsymbol{\mu}) u_{Nm}^{\mathcal{N}}(\boldsymbol{\mu}) = f(v_n; \boldsymbol{\mu}) \quad 1 \leq n \leq N \tag{2.9}$$

Once this system is solved for the coefficients $u_{Nm}(\boldsymbol{\mu})$, the RB output $s_N^{\mathcal{N}}(\boldsymbol{\mu})$ can be subsequently evaluated as

$$s_N^{\mathcal{N}}(\boldsymbol{\mu}) = \ell(u_N^{\mathcal{N}}(\boldsymbol{\mu})). \tag{2.10}$$

It is not surprising that the accuracy of the RB solution $u_N^{\mathcal{N}}$ and of the corresponding computed output of interest $s_N^{\mathcal{N}}$ both depend crucially on the construction of the reduced basis approximation space. The procedure we use for efficiently selecting representative parameters $\boldsymbol{\mu}^1, \dots, \boldsymbol{\mu}^N$ and the corresponding snapshots defining the reduced basis space plays an essential role in the reduced basis method.

¹ In implementations, in order to ameliorate ill-conditioning issues that may arise in (2.7) we first apply the Gram–Schmidt process with respect to the $(\cdot, \cdot)_X$ inner product each time a new snapshot $u^{\mathcal{N}}(\boldsymbol{\mu}^N)$ is generated to obtain a $(\cdot, \cdot)_X$ -orthonormal basis $\{s_N^{\mathcal{N}}\}_{N=1}^{N_{\max}}$. We omit explicitly denoting or showing this orthogonalization procedure.

2.3 Selecting Snapshots: Enhancing Offline RBM Computations

This paper’s main contribution is the development of novel procedures for selecting the snapshot set $\{\boldsymbol{\mu}^N\}_{N=1}^{N_{max}}$. The main idea of our procedure is very similar to classical RBM methods, the latter of which is the greedy scheme

$$\boldsymbol{\mu}^{N+1} = \operatorname{argmax}_{\boldsymbol{\mu} \in \Xi_{\text{train}}} \Delta_N(\boldsymbol{\mu}), \quad \Delta_N(\boldsymbol{\mu}) \geq \|u_N^{\mathcal{N}}(\boldsymbol{\mu}) - u^{\mathcal{N}}(\boldsymbol{\mu})\|_{X^{\mathcal{N}}}. \tag{2.11}$$

Here $\Delta_N(\cdot)$ is an efficiently-computable *a posteriori* error bound for the RBM procedure [33], which is crucial for the reliability of the reduced basis space constructed according to the weak greedy algorithm discussed in ‘‘Appendix A.1’’. Toward that end, we reconsider the numerical schemes for the truth approximation (2.5) and for the RB solution (2.7). Defining the error $e_N(\boldsymbol{\mu}) := u^{\mathcal{N}}(\boldsymbol{\mu}) - u_N^{\mathcal{N}}(\boldsymbol{\mu}) \in X^{\mathcal{N}}$, linearity of a yields the following error equation:

$$a(e_N(\boldsymbol{\mu}), v; \boldsymbol{\mu}) = r_N(v; \boldsymbol{\mu}) \quad \forall v \in X^{\mathcal{N}}, \tag{2.12}$$

with the residual $r_N(v; \boldsymbol{\mu}) \in (X^{\mathcal{N}})'$ (the dual of $X^{\mathcal{N}}$) is defined as $f(v; \boldsymbol{\mu}) - a(u_N^{\mathcal{N}}(\boldsymbol{\mu}), v; \boldsymbol{\mu})$. The Riesz representation theorem and the Cauchy–Schwarz inequality implies that $\|e_N(\boldsymbol{\mu})\|_X \leq \frac{\|r_N(\cdot; \boldsymbol{\mu})\|_{(X^{\mathcal{N}})'}}{\alpha_{LB}^{\mathcal{N}}(\boldsymbol{\mu})}$, where $\alpha^{\mathcal{N}}(\boldsymbol{\mu}) = \inf_{w \in X^{\mathcal{N}}} \frac{a(w, w; \boldsymbol{\mu})}{\|w\|_X^2}$ is the stability (coercivity) constant for the elliptic bilinear form a . This implies that we can define the *a posteriori* error estimator for the solution as

$$\Delta_N(\boldsymbol{\mu}) = \frac{\|r_N(\cdot; \boldsymbol{\mu})\|_{(X^{\mathcal{N}})'}}{\alpha_{LB}^{\mathcal{N}}(\boldsymbol{\mu})} \geq \|e_N(\boldsymbol{\mu})\|_{X^{\mathcal{N}}} \tag{2.13}$$

The efficiency of computing the *a posteriori* error estimation relies on that of the lower bound of the coercivity constant $\alpha_{LB}^{\mathcal{N}}(\boldsymbol{\mu})$ as well as the value $\|r_N(\cdot; \boldsymbol{\mu})\|_{(X^{\mathcal{N}})'}$, for $\forall \boldsymbol{\mu} \in \mathcal{D}$. The coercivity constant $\alpha^{\mathcal{N}}$ can be nontrivial to compute, but there are constructive algorithms to address this [8, 19, 20]. The residual is typically computed by the RBM offline–online decomposition, which is presented in ‘‘Appendix A.2’’.

Once $\boldsymbol{\mu}^{N+1}$ is selected, standard RBM mechanics can be used to construct $u_{N+1}^{\mathcal{N}}$ so that the procedure above can be iterated to compute $\boldsymbol{\mu}^{N+2}$.

A classical RBM algorithm computes the maximum over Ξ_{train} above in a brute-force manner; since Ξ_{train} is large and this maximization must be done for every $N = 1, \dots, N_{max}$, this process of selecting snapshots is usually one of the more computationally expensive portions of RBM algorithms.

This manuscript is chiefly concerned with ameliorating the cost of selecting snapshots; we call this an ‘‘offline-enhanced’’ Reduced Basis Method.² We present two algorithms that are alternatives to the brute-force approach (2.11). Instead of maximizing Δ_N over Ξ_{train} , we instead maximize over a subset of Ξ_{train} that we call the ‘‘surrogate training set’’. The efficient computational determination of the surrogate training set, and subsequent empirical studies investigating the accuracy and efficiency of offline-enhanced methods compared to classical RBM, are the remaining topics of this manuscript.

This paper does not make novel contributions to any of the other important aspects of RBM algorithms (offline/online decompositions, error estimate computations, etc.), but in the interests of completeness we include in the ‘‘Appendix’’ a brief overview of the remaining portions of RBM algorithms.

² ‘‘Offline’’ is a standard descriptor for this general portion of the full RBM algorithm; see the ‘‘Appendix’’.

Table 2 Notation used for offline-enhanced RBM

<i>Algorithm template notation</i>	
ℓ	Number of “outer” loops in the offline enhancement procedure Algorithm 1
E_ℓ	The largest error estimator at the beginning of outer loop ℓ
N_ℓ	Number of offline-enhanced snapshots chosen at iteration ℓ
Ξ_{sur}	Surrogate Training Set (STS), a subset of Ξ_{train}
M_ℓ	STS parameter at outer loop iteration ℓ
K_{damp}	A constant integer controlling the damping ratio $\frac{1}{K_{\text{damp}} \times (\ell + 1)} \in (0, 1)$
<i>SMM notation</i>	
I_N^M	M equispaced samples on the interval $(\epsilon_{\text{tol}}, \max_{\mu \in \Xi_{\text{train}}} \Delta_N(\mu))$
<i>CD notation</i>	
G	$\Xi_{\text{train}} \times \Xi_{\text{train}}$ Gramian error matrix, with entries $G_{i,j} = (e_N(\mu^i), e_N(\mu^j))_{X_N}$
\tilde{G}	$\Xi_{\text{train}} \times \Xi_{\text{train}}$ approximate Gramian error matrix, with entries $\tilde{G}_{i,j} = (\tilde{e}_N(\mu^i), \tilde{e}_N(\mu^j))_{X_N}$

3 Offline-Enhanced RBM: Design and Analysis

The complexity of the offline stage, where the optimization (2.11) is performed, depends on N_{train} . Although this dependence is only linear, a large N_{train} can easily make the computation onerous; such a situation arises when the parameter domain \mathcal{D} has large dimension p . In this case standard constructions for N_{train} yield grid-based training sets that grow exponentially with p , even when the more parsimonious sparse grid constructions are involved [39]. Our goal in this project is to ameliorate the cost of sweeping over a very large training set in (2.11) without sacrificing the quality of the reduced basis solution. We call this approach an *Offline-enhanced Reduced Basis Method*.

The basic idea of our approach is to perform the standard RBM greedy algorithm on a surrogate training set (STS) constructed as subsets of the original Ξ_{train} . The STS is constructed adaptively, and construction is periodically repeated after a small batch of snapshots are selected. We let Ξ_{Sur} denote these constructed STS’s; they are small enough compared to Ξ_{train} to offer considerable acceleration of the greedy sweep (2.11), yet large enough to capture the general landscape of the solution manifold. We present in Algorithm 1 a general template for our Offline-Enhanced Reduced Basis Method. This algorithm can be implemented once we describe how Ξ_{Sur} are constructed; these descriptions are the topic of the next sections. In Table 2 we summarize the notation in Algorithm 1.

The first contribution of our paper resides in the unique structure of this template. Each global greedy sweep (i.e., over Ξ_{train} , and labeled “One-step greedy” in Algorithm 1) is followed by multiple targeted sweeps over the (smaller) STS Ξ_{Sur} (labeled “Multi-step greedy”). These latter sweeps produce a computational savings ratio of $1 - |\Xi_{\text{Sur}}|/|\Xi_{\text{train}}|$ because they operate on Ξ_{Sur} instead of on Ξ_{train} .

Note that we still require occasional global greedy sweeps, even though they are expensive. These global sweeps are necessary to retain reliability of the greedy algorithm.

Algorithm 1 The Offline-enhanced Reduced Basis Method template. Algorithms for constructing the STS Ξ_{Sur} are described in Sects. 3.1 and 3.2.

```

1: Input:  $\Xi_{\text{train}}$ , an accuracy tolerance  $\varepsilon_{\text{tol}}$ .
2: Randomly select the first sample  $\mu^1 \in \Xi_{\text{train}}$ , and set  $N = 1, \varepsilon = 2\varepsilon_{\text{tol}}$ , and  $\ell = 0$ .
3: Obtain truth solution  $u^{\mathcal{N}}(\mu^1)$ , and set  $X_1^{\mathcal{N}} = \text{span}\{u^{\mathcal{N}}(\mu^1)\}$ .
4: while ( $\varepsilon > \varepsilon_{\text{tol}}$ ) do
5:   Set  $\ell \leftarrow \ell + 1$ .
6:   One-step greedy scan on  $\Xi_{\text{train}}$ 
7:     for each  $\mu \in \Xi_{\text{train}}$  do
8:       Obtain RBM solution  $u_N^{\mathcal{N}}(\mu) \in X_N^{\mathcal{N}}$  and error estimate  $\Delta_N(\mu)$ 
9:     end for
10:     $\mu^{N+1} = \underset{\mu \in \Xi_{\text{train}}}{\text{argmax}} \Delta_N(\mu), \varepsilon = \Delta_N(\mu^{N+1}), E_\ell = \varepsilon$ .
11:    Augment RB space  $X_{N+1}^{\mathcal{N}} = X_N^{\mathcal{N}} \oplus \{u^{\mathcal{N}}(\mu^{N+1})\}$ .
12:    Set  $N \leftarrow N + 1, N_\ell \leftarrow 0$ 
13:   Construct STS  $\Xi_{\text{Sur}}$  based on  $\{(u_{N-1}^{\mathcal{N}}(\mu), \Delta_{N-1}(\mu)) : \mu \in \Xi_{\text{train}}\}$ .
14:   Multi-step greedy scan on  $\Xi_{\text{Sur}}$ 
15:     while ( $\varepsilon > \varepsilon_{\text{tol}}$ ) and ( $\varepsilon > E_\ell \frac{1}{K_{\text{damp}} \times (\ell+1)}$ ) do
16:       for each  $\mu \in \Xi_{\text{Sur}}$  do
17:         Obtain RBM solution  $u_N^{\mathcal{N}}(\mu) \in X_N^{\mathcal{N}}$  and error estimate  $\Delta_N(\mu)$ 
18:       end for
19:        $\mu^{N+1} = \underset{\mu \in \Xi_{\text{Sur}}}{\text{argmax}} \Delta_N(\mu), \varepsilon = \Delta_N(\mu^{N+1})$ .
20:       Augment RB space  $X_{N+1}^{\mathcal{N}} = X_N^{\mathcal{N}} \oplus \{u^{\mathcal{N}}(\mu^{N+1})\}$ 
21:       Set  $N \leftarrow N + 1, N_\ell \leftarrow N_\ell + 1$ 
22:     end while
23:   end while

```

The other main contribution of our paper is the creation of two strategies for constructing the surrogate training set Ξ_{Sur} , which is the topic of the next two subsections. Our two procedures are the Successive Maximization Method (SMM) and the Cholesky Decomposition Method (CDM). Once they are described, we may use them in the algorithmic template that fully describes Offline-enhanced RB methods. SMM and CDM are intrinsically different in their construction, yet our numerical experiments show that they both work very well, accelerating the offline portion of the RBM algorithm significantly without sacrificing accuracy for the examples we have tested.

The details of SMM are described in Sect. 3.1, and those of CDM in Sect. 3.2. Before describing these details, we make three general remarks concerning the Algorithm template:

- Motivation for constructing the STS** In a standard sweep of (2.11) to identify μ^* from Ξ_{train} that maximizes the error estimate $\Delta_N(\mu)$, we actually must compute $\Delta_N(\mu)$ for *all* $\mu \in \mathcal{D}$. Standard RB algorithms discard this information upon identifying μ^* . However, this is valuable, quantitative information about $\|e_N(\mu)\| = \|u^{\mathcal{N}}(\mu) - u_N^{\mathcal{N}}(\mu)\|$ for *any* $\mu \in \mathcal{D}$. Construction of Ξ_{Sur} attempts to utilize this information that was otherwise discarded to identify not just μ^* , but a collection of parameters that can describe the landscape of $\Delta(\cdot)$. In other words, we gauge the accuracy of the reduced solution in $X_{N+1}^{\mathcal{N}}$ for all μ , and trim from Ξ_{train} those parameters whose corresponding solutions are deemed good enough. Roughly speaking, we set

$$\Xi_{\text{Sur}} := \{\mu : u_{N+1}^{\mathcal{N}}(\mu) \text{ is predicted to be "inaccurate"}\}. \tag{3.1}$$

Note that one can mathematically devise adversarial scenarios where such a procedure can discard values in Ξ_{train} that later will be important. However, the outer loop of the template is designed so that we reconsider any parameter values that may have been discarded at one point. The goal is to construct Ξ_{Sur} in a balanced way: A strict definition of “inaccurate” in (3.1) makes Ξ_{Sur} too large and no computational savings are realized; a lax definition chooses too few values for Ξ_{Sur} and the RB surrogate will not be accurate.

- **Stopping criteria for the STS** On outer loop round ℓ , we repeatedly sweep the current STS Ξ_{Sur} after it is constructed until

$$\max_{\mu \in \Xi_{\text{Sur}}} \Delta_N(\mu) \leq E_\ell \frac{1}{((\ell + 1) \times K_{\text{damp}})},$$

where E_ℓ is the starting (global) maximum error estimate for this outer loop iteration. The damping ratio $\frac{1}{((\ell + 1) \times K_{\text{damp}})}$, enforces that the maximum error estimate over the STS decreases by a controllable factor K_{damp} ; in this paper we take K_{damp} to be constant in ℓ . However, this damping ratio should be determined by the practitioner and the problem at hand. Taking K_{damp} as a constant works well in our test problems.

- **Cost of constructing Ξ_{Sur}** The cost of constructing Ξ_{Sur} is an overhead cost for each outer loop of Algorithm 1. Therefore, we must formulate this construction so that the overhead cost is worth the effort. For example, if the cost of evaluating $\Delta(\cdot)$ at one value is C , and we select N_ℓ snapshots from Ξ_{Sur} at outer iteration ℓ , then we attain cost savings when

$$\frac{\text{Cost of constructing } \Xi_{\text{Sur}}}{CN_\ell} < |\Xi_{\text{train}}| - |\Xi_{\text{Sur}}|.$$

This yields qualitative information about the efficiency of the method: when the cost of constructing Ξ_{Sur} is negligible, we may take a large Ξ_{Sur} , but when this cost is large, we require a significant size reduction in order to amortize the initial investment.

3.1 Successive Maximization Method

Our first approach for constructing the surrogate training set is the Successive Maximization Method (SMM). This procedure is motivated by the notion that the difference between the norm of the errors $\|e(\mu_1)\|_X - \|e(\mu_2)\|_X$ is partially indicative of the difference between the solutions. Computation of the true error norms is impractical, so like standard offline RBM procedures we leverage the *a posteriori* error estimate $\Delta_N(\mu)$ defined in (2.13).

Suppose we have already selected N snapshots; when selecting parameter value $N + 1$ via (2.11), we must compile the values $\Delta_N(\Xi_{\text{train}}) = \{\Delta_N(\mu) \mid \mu \in \Xi_{\text{train}}\}$. We use this collection to identify the surrogate training set. From our argument that the values $\|e(\mu)\|$ give us some indication about the actual solution, we equidistantly sample values from $\Delta_N(\Xi_{\text{train}})$ to construct the surrogate training set.

With ϵ_{tol} the stopping tolerance for the RB sweep, let $\Delta_N^{\text{max}} = \max_{\mu \in \Xi_{\text{train}}} \Delta_N(\mu)$. We define $I_N^{M_\ell}$ as an equi-spaced set between ϵ_{tol} and Δ_N^{max} :

$$I_N^{M_\ell} = \left\{ \nu_{N,m} := \epsilon_{\text{tol}} + (\Delta_N^{\text{max}} - \epsilon_{\text{tol}}) \frac{m}{M_\ell} : m = 0, \dots, M_\ell - 1 \right\}. \tag{3.2}$$

Roughly speaking, we attempt to construct Ξ_{Sur} as $\Xi_{\text{Sur}} = \Delta_N^{-1} \left(I_N^{M_\ell} \right) \cap \Xi_{\text{train}}$. Rigorously, we use

$$\Xi_{\text{Sur}} = \left\{ \mu_{N,m} : \mu_{N,m} = \operatorname{argmin}_{\mu \in \Xi_{\text{train}}} \left\{ \Delta_N(\mu) - v_{N,m} \text{ such that } \Delta_N(\mu) \geq v_{N,m} \right\} \right\}.$$

Note that we have $|\Xi_{\text{Sur}}| \leq M_\ell$ by this construction. M_ℓ can be chosen as any monotonically increasing function with respect to ℓ such as $10(\ell + 1)$. But in order to avoid excessively large Ξ_{sur} , we set $M_\ell = C_M(\ell + 1)$ in this paper, where C_M is a constant.

3.2 Cholesky Decomposition Method

The intuition of our second approach is that the angle between the error vectors is somewhat symptomatic of the difference between the solutions. Thus, we consider the (scaled) Gramian matrix G comprised of pairwise inner products of error vectors $e_N(\mu)$. I.e.,

$$G_{i,j} = \frac{(e(\mu_i), e(\mu_j))_X}{\|e(\mu_i)\|_X \|e(\mu_j)\|_X}, \quad \mu_i, \mu_j \in \Xi_{\text{train}}.$$

The matrix G is positive semi-definite, and thus admits a (pivoted) Cholesky decomposition. We suppress notation indicating that G depends on the current number of snapshots N .

Our approach here is to apply the *pivoted Cholesky decomposition* [17] of the matrix G with maximum times of pivoting M_ℓ , where $M_\ell = C_M(\ell + 1)$ and C_M is a constant. This decomposition of G orders the elements of Ξ_{train} according to the pivots. We identify the surrogate training set Ξ_{Sur} as the first M pivots (parameter values) selected by this procedure.

We don't literally evaluate the error vectors $e_N(\mu)$ or their Gramian G . Indeed, obtaining the error vectors $e(\mu)$ is as expensive as solving for the truth approximation, we have to approximate these vectors. A linear algebraic way to write the Galerkin system (2.5) is

$$\mathbb{A}(\mu)u^N(\mu) = f(\mu), \quad \mathbb{A} \in \mathbb{R}^{N \times N}$$

where \mathbb{A} , u^N , and f are discretization vectors associated to $a(\cdot, \cdot; \mu)$, u^N , and $f(\cdot; \mu)$, respectively. With this notation, we have

$$e(\mu) = u^N(\mu) - u_N^N(\mu) = \mathbb{A}_N^{-1}(\mu)r(u_N^N(\mu); \mu),$$

where the residual vector $r(\cdot; \mu)$ is defined as $r(v; \mu) = f - \mathbb{A}_N(\mu)v$. We propose to approximate the unknown $\mathbb{A}_N^{-1}(\mu)$ by

$$\tilde{\mathbb{A}}_N^{-1}(\mu) := \sum_{m=1}^Q u_{Q_m}^N(\mu)\mathbb{A}_N^{-1}(\mu^m), \tag{3.3}$$

where $\{\mu^m\}_{m=1}^Q$ are the key parameters, $\{u_{Q_m}^N(\mu)\}_{m=1}^Q$ are the RB coefficients for $u_Q^N(\mu)$ defined in (2.8) and $\{\mathbb{A}_N^{-1}(\mu^m)\}_{m=1}^Q$ are already computed when we solve $u^N(\mu^m)$. Since $(\mathbb{A}_N^{-1}(\mu) - \tilde{\mathbb{A}}_N^{-1}(\mu))f = u^N(\mu) - u_N^N(\mu)$, we argue that this approximation is reasonable.

The approximation of $e(\mu)$ can be expressed as $\tilde{e}(\mu) = \left(\sum_{m=1}^Q u_{Q_m}^N(\mu)\mathbb{A}_N^{-1}(\mu^m) \right) r(u_N^N(\mu); \mu)$ which admits an affine decomposition, and the Gramian matrix G is approximated by

$$\tilde{G}_{ij} = (\tilde{e}(\mu_i), \tilde{e}(\mu_j)) / (\|\tilde{e}(\mu_i)\|_X \times \|\tilde{e}(\mu_j)\|_X), \tag{3.4}$$

where $1 \leq i, j \leq N_{\text{train}}$.

Note that we can take Q to be smaller than the number of reduced basis N . For our experiments, we take $Q = \sqrt{N}$. This practice generates savings when forming the approximate inverse (3.3) since the matrices \mathbb{A} are of size \mathcal{N} . Note also that, in (3.3), for practical implementations we don't need to form any inverse matrices $\mathbb{A}_{\mathcal{N}}^{-1}(\boldsymbol{\mu}^m)$. In fact, these inverse matrices are multiplied with components of the residual $r(\mathbf{v}; \boldsymbol{\mu}) = f - \mathbb{A}_{\mathcal{N}}(\boldsymbol{\mu})\mathbf{v}$. Therefore, we only need to solve a number of linear systems to form the Gramian (3.4). Moreover, these linear solves are independent of the parameter hence can be performed offline. As a result, the Gramian (3.4) is obtained in an offline-online decomposed, thus highly efficient, fashion. We address the related complexity count in the next section and the ‘‘Appendix’’.

3.3 Complexity Analysis

In ‘‘Appendix A.2’’ we see that the computational complexity for the offline portion of the classical algorithm has order

$$\underbrace{\mathcal{N}^2 N^2 Q_a}_{\text{Reduced solve preparation}} + \underbrace{N_{\text{train}} W_\alpha + N_{\text{train}} (Q_a^2 N^3 + N^4) + N W_s}_{\text{Greedy sweeping}} + \underbrace{Q_a^2 N^3 W_m}_{\text{Estimator preparation}}.$$

This cost is dominated by the boldface term in the middle, especially when Ξ_{train} is large. We denote this cost C_{orig} ,

$$C_{\text{orig}} := N_{\text{train}} (Q_a^2 N^3 + N^4).$$

This is the portion of the offline cost that our Offline-Enhanced RBM is aiming to reduce.

Suppose we have computed n snapshots. Then the cost for assembling and solving the RB system for one given parameter value $\boldsymbol{\mu}$ is of order $n^2 Q_a + n^3$, while the cost for calculating the error certificate is of order $n^2 Q_a^2$. Therefore the total cost for one instance, denoted by $c(n)$, is of order $n^2 Q_a^2 + n^3$. This means that the complexity for the classical RBM to sweep over $N_{\text{train}} := |\Xi_{\text{train}}|$ parameter values is

$$N_{\text{train}} \sum_{n=1}^N c(n) = \mathcal{O} (N_{\text{train}} (Q_a^2 N^3 + N^4)).$$

To better analyze the cost of our Offline-enhanced approaches, we denote the *cumulative* number of chosen parameter values after the j -th outer loop iteration by

$$t_j = 1 + \sum_{k=1}^j N_k$$

We note that $t_0 = 1$ because in standard RB algorithms the first parameter value is randomly chosen before starting the greedy algorithm. If the RB procedure given by Algorithm 1 terminates after ℓ outer loop iterations with a total of N snapshots, then we have $t_\ell = N$. The cost of the Offline-enhanced approaches corresponding to the dominating cost of the classical approach C_{orig} is

$$\begin{aligned}
 \overbrace{C_{oe}}^{\text{Total cost}} &= \overbrace{N_{\text{train}} \sum_{j=0}^{\ell} c(t_j)}^{\text{One-step greedy scans on } \Xi_{\text{train}}} + \overbrace{\sum_{j=1}^{\ell} M_j \left(\sum_{n=t_{j-1}+1}^{t_j-1} c(n) \right)}^{\text{Multi-step greedy scans on } \Xi_{\text{Sur}}} \\
 &\lesssim N_{\text{train}} \ell c(N) + M_{\text{max}} \sum_{n=1}^N c(n), \quad \text{with } M_{\text{max}} := \max_{j=1, \dots, \ell} M_j \\
 &\lesssim N_{\text{train}} \ell (Q_a^2 N^2 + N^3) + M_{\text{max}} (Q_a^2 N^3 + N^4).
 \end{aligned}$$

We conclude that the dominating parameter sweeping costs between the classical and offline-enhanced approaches satisfy

$$C_{oe} < C_{orig} \left(\frac{\ell}{N} + \frac{M_{\text{max}}}{N_{\text{train}}} \right). \tag{3.5}$$

We make some remarks concerning this cost analysis:

- **Potential savings**—Since $\frac{M_{\text{max}}}{N_{\text{train}}}$ is negligible especially for the cases of our concern when the parameter dimension is high, (3.5) demonstrates that the savings is roughly $\frac{\ell}{N}$.
- **Surrogate Training Set construction for SMM-RBM**—The additional cost for the Offline-enhanced approaches is the construction of the surrogate training set. For SMM-RBM, this cost is essentially negligible. This SMM surrogate training set construction cost is mainly dependent on the cost of evaluating the error certificate Δ_N , but this cost has already been included in the analysis above. In practice, the surrogate training set construction amounts to a quick sorting of these certificates which results in a cost of $\mathcal{O}(N_{\text{train}} \log(N_{\text{train}}))$. While this cost does depend on N_{train} , it is much smaller than any of the terms in, e.g., C_{orig} .
- **Surrogate Training Set for CD-RBM**—CD-RBM entails a sequence of (pivoted) Cholesky Decomposition steps applied to the approximate error Gramian matrix \tilde{G} , (3.4). For the decomposition algorithm, it suffices to just supply the approximate errors $\tilde{z}(\mu)$ without constructing the full matrix \tilde{G} . Therefore, the cost is primarily devoted to computing these approximate errors for all μ . Evaluating these error functions can be accomplished in an offline-online way detailed in “Appendix A.3”. We summarize here the total cost for the STS construction.

$$\underbrace{\mathcal{N}^2 Q N_{\text{RB}} Q_a}_{\text{Offline Preparation}} + \underbrace{N_{\text{train}} \mathcal{N} Q N_{\text{RB}} Q_a \ell}_{\text{Approximate error calculation}} + \underbrace{N_{\text{train}} \sum_{j=1}^{\ell} \left(n_{\text{cd}}^j \right)^2}_{\ell \text{ runs of Pivoted CD algorithm}}$$

where n_{cd}^j is the number of steps of the pivoted Cholesky decomposition for the j -th iteration. We see that this algorithm can be more costly than SMM-RBM because of the factor $\mathcal{N} Q N_{\text{RB}} Q_a$. However, we observe that it is still much faster than the classical version since this factor is notably smaller than $Q_a^2 N_{\text{RB}}^3 + N_{\text{RB}}^4$. This is confirmed by our numerical examples presented in the next section.

4 Offline-Enhanced RBM: Numerical Results

In this section, we present numerical examples to illustrate the accuracy and efficiency enhancement of the proposed approaches compared to the conventional reduced basis method.

4.1 Test Problems

We test the two algorithms, SMM-RBM and CD-RBM, on two standard diffusion-type problems. They vary substantially in terms of parameter dimension and the truth solver. Our results show that our offline-enhanced procedures work well in both of these cases, and suggest that the offline-enhanced strategies may be beneficial for a wider class of parameter spaces and truth solvers.

Diffusion problem with two-dimensional parameter domain

$$(1 + \mu_1 x)u_{xx} + (1 + \mu_2 y)u_{yy} = e^{4xy} \quad \text{on } \Omega. \tag{4.1}$$

Here $\Omega = [-1, 1] \times [-1, 1]$ and we impose homogeneous Dirichlet boundary conditions on $\partial\Omega$. The truth approximation is a spectral Chebyshev collocation method based on $\mathcal{N}_x = 35$ degrees of freedom in each direction, with $\mathcal{N}_x^2 = \mathcal{N}$. The parameter domain \mathcal{D} for (μ_1, μ_2) is taken to be $[-0.99, 0.99]^2$. For the Ξ_{train} we discretize \mathcal{D} using a tensorial 160×160 Cartesian grid with 160 equi-spaced points in each dimension.

Thermal Block problem with nine-dimensional parameter domain

$$\begin{cases} -\nabla \cdot (a(x, \boldsymbol{\mu}) \nabla u(x, \boldsymbol{\mu})) = f & \text{on } \Omega, \\ u(x, \boldsymbol{\mu}) = g_D & \text{on } \Gamma_D, \\ \frac{\partial u}{\partial n} = g_N & \text{on } \Gamma_N. \end{cases} \tag{4.2}$$

Here $\Omega = [0, 1] \times [0, 1]$ which is partitioned into 9 blocks $\bigcup_{i=1}^9 B_i = \Omega$, Γ_D is the top boundary, and $\Gamma_N = \partial\Omega \setminus \Gamma_D$. The parameters μ_i , $1 \leq i \leq 9$ denote the heat conductivities:

Γ_D		
$\mu_7(B_7)$	$\mu_8(B_8)$	$\mu_9(B_9)$
$\mu_4(B_4)$	$\mu_5(B_5)$	$\mu_6(B_6)$
$\mu_1(B_1)$	$\mu_2(B_2)$	$\mu_3(B_3)$
Γ_{base}		

The diffusion coefficient $a(x, \boldsymbol{\mu}) = \mu_i$ if $x \in B_i$. The parameter vector is thus given by $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_9)$ in $\mathcal{D} = [0.1, 10]^9$. We take as the right hand side $f = 0$, $g_D = 0$, $g_N = 1$ on the bottom boundary Γ_{base} and $g_N = 0$ otherwise. The output of interest is defined as the integral of the solution over Γ_{base}

$$s(\boldsymbol{\mu}) = \int_{\Gamma_{\text{base}}} u(x, \boldsymbol{\mu}) dx \tag{4.3}$$

The truth approximation is obtained by FEM with $\mathcal{N} = 361$. A sufficient number of $N_{\text{train}} = 20,000$ samples are taken from randomly sampling within the parameter domain \mathcal{D} . The classical RB solver for this problem is provided by, and our Offline-enhanced approach is compared against, the RBmatlab package [14, 15].³

4.2 Results

We investigate the performance of the two offline-enhanced RB algorithms on the two test problems. The tuning parameters for Algorithm 1 for both examples are shown in Table 3.

For the first test case with a two-dimensional parameter domain that is easy to visualize, we display the location of the selected parameter values in Fig. 1. On the left is that for the

³ Available for download at <http://www.ians.uni-stuttgart.de/MoRePaS/software/index.html>.

Table 3 Offline-enhanced RBM parameters for the numerical results

Parameter	SMM-RBM	CD-RBM
M_ℓ	$2 \times (\ell + 1)$	$20 \times (\ell + 1)$
K_{damp}	1	10




Fig. 1 Selected parameter values by a classical offline RBM algorithm (left), by the SMM-RBM algorithm (center), and by the CDM-RBM algorithm (right). Points with larger radius are chosen earlier in the sequence, and for the SMM and CDM plots points earlier in the sequence have greater transparency. For the SMM and CDM plots, all chosen parameters within a batch (an outer loop in Algorithm 1) have the same radius

classical RBM. The larger the marker, the earlier the parameter is picked. At the middle and on the right are the parameter sets selected by SMM-RBM and CD-RBM respectively. For these two, the more transparent the marker, the earlier it is picked. A group of parameter values chosen at the same step have the same radius. This figure shows that the enhanced algorithm does “preserve” the usual property when it comes to the distribution of the key parameter values determined by RBM. That is, they tend to distribute along the boundary of the parameter domain. We observe similar phenomenon for the 9-dimensional point selection for the second test case, but omit the visualization for brevity and simplicity.

The accuracy and efficiency of the new algorithms are shown in Fig. 2. We see clearly that the *a posteriori* error estimate is converging exponentially for both SMM-RBM and CD-RBM, in the same fashion as the classical version of RBM. This shows that our accelerated algorithm does not appear to suffer accuracy degradation for these examples. In addition, we see a factor of 3-to-6 times runtime speedup.

The influence of the tuning parameters M_ℓ and K_{damp} are shown in Tables 4 and 5. We test two cases: $\epsilon_{tol} = 0.05$ and $\epsilon_{tol} = 0.01$. For each case, we choose 6 different tuning parameters and check the performance of our algorithm by measuring the relative time and N_{max} . The relative time for each scenario is defined as the corresponding running time scaled by the running time of the case with $K_{damp} = 1/2$ and $M_\ell = 4(\ell + 1)$. We observe that the performance of our approaches, especially the SMM-RBM, is rather monotonically improving as expected, and the size of the resulting RB space is rather stable. As a comparison, we implemented the random greedy sampling strategy proposed in [18]. The results are not included in this paper, but they appear to show non-negligible variations between the runs for the same tuning parameter. This variation is zero for our case since our algorithm is deterministic. More importantly, as we change the tuning parameter, the algorithm appears to be less stable and less monotonic comparing with ours.

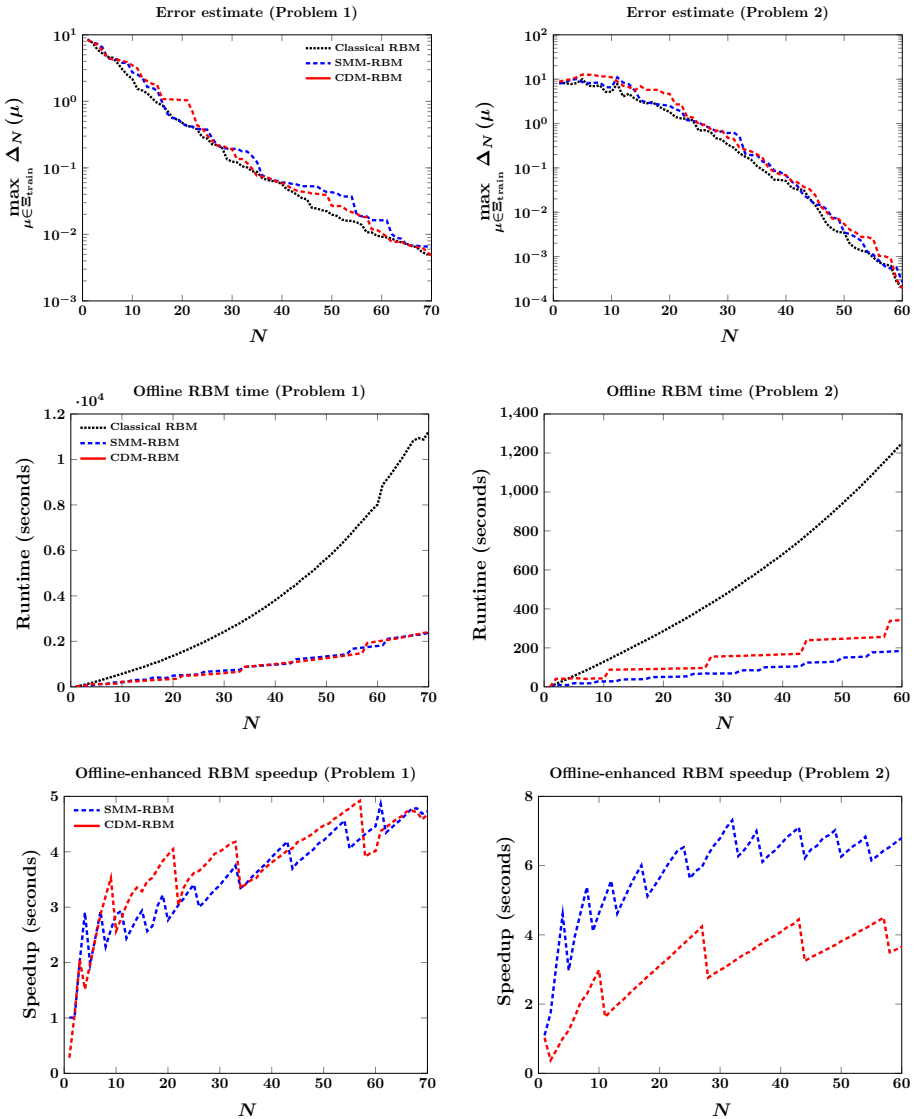


Fig. 2 Convergence and speedup of offline-enhanced RBM algorithms as a function of the total number of snapshots n . Error estimate convergence (top), computational runtime (middle), and speedup factor (bottom). Left-hand plots correspond to the test problem 1, and right-hand plots for test problem 2

Finally, to reveal the effectivity of the construction of the surrogate training set, we plot in Fig. 3 the *Surrogate Acceptance Ratio*

$$\text{SAR}(\ell) := \frac{N_\ell}{M_\ell}$$

as a function of the outer loop iteration index ℓ . This ratio quantifies how much of the surrogate training set is added to the snapshot parameter set at each outer loop iteration. Large ratios suggest that our construction of the surrogate training set effectively emulates

Table 4 Tuning parameter testing for SMM-RBM

K_{damp}	M_ℓ	$\varepsilon_{tol} = 0.05$		$\varepsilon_{tol} = 0.01$	
		Relative time	N_{max}	Relative time	N_{max}
1/2	$4(\ell + 1)$	1.0000	42	1.0000	62
1/5	$10(\ell + 1)$	0.8795	43	0.9885	62
1/10	$20(\ell + 1)$	0.8571	43	0.9043	60
1/20	$40(\ell + 1)$	0.8651	43	0.9490	63
1/30	$60(\ell + 1)$	0.8422	42	0.9026	62
1/40	$80(\ell + 1)$	0.7550	41	0.8400	61

Table 5 Tuning parameter testing for CD-RBM

K_{damp}	M_ℓ	$\varepsilon_{tol} = 0.05$		$\varepsilon_{tol} = 0.01$	
		Relative time	N_{max}	Relative time	N_{max}
1/10	$20(\ell + 1)$	1.0000	44	1.0000	63
1/20	$40(\ell + 1)$	0.9654	43	0.8769	61
1/30	$60(\ell + 1)$	0.9621	43	0.8816	60
1/40	$80(\ell + 1)$	0.9379	42	0.9016	62
1/50	$100(\ell + 1)$	0.9357	42	0.9018	62
1/60	$120(\ell + 1)$	0.9372	42	0.9082	63

the entire training set. In Fig. 3 we see that a significant portion (on average approximately 40% for the first case and 30% for the second case) of the surrogate training set is chosen by the greedy algorithm before continuing into another outer loop. We recall that the second case has a 9-dimensional parameter domain, and so our offline-enhanced RBM procedure can effectively choose surrogate training set even when the parameter dimension is large. The relatively large values of the SAR result in the computational speedup observed in Fig. 2.

5 Concluding Remarks

We proposed offline-enhanced reduced basis methods for building reduced-order models; RBM algorithms invest significant resources in an offline stage by studying a finite training set and judiciously choosing snapshots from this training set. Our novel approach substitutes the original training set with an adaptively constructed the surrogate training set that is much smaller in size, and thus reduces the computational time spent in the offline portion of the RBM algorithm. (Our algorithm leaves the online portion of RBM algorithms unchanged).

We provide two approaches to identify and construct the surrogate training set using two different perspectives: the SMM-RBM strategy constructs a surrogate training set by uniformly sampling parameters on the range of the *a posteriori* error estimate; the CD-RBM strategy uses the angle between two approximate error vectors at different locations in parameter space to identify the surrogate training set. Like RBM in general, our approaches are particularly useful in computing many-query reliable solutions parametrized PDE having a large number of random inputs. We have demonstrated the computational efficiency of our

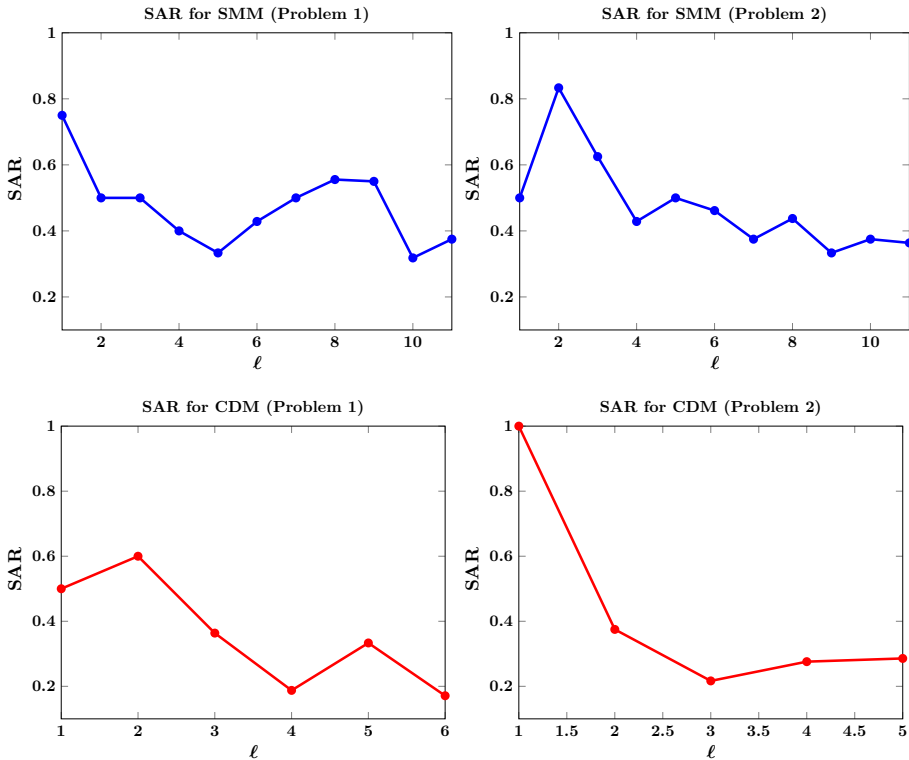


Fig. 3 The surrogate acceptance ratio (SAR) for the offline-enhanced RBM methods, SMM (top) and CDM (bottom). Test problem 1 is shown on the left, and test problem 2 on the right

proposed methods compared against the standard reduced basis method for two steady-state diffusion problems. The application of the offline-enhanced reduced basis method to more general problems with high dimensional parameter domains is ongoing research.

A Classical RBM Specifics: Greedy Algorithms, Efficiency, and Operational Count

This “Appendix” contains the mathematical and algorithmic portions of RBM algorithms that are not directly the subject of this manuscript. These specifics are well-known in the RBM literature and community, and we include this “Appendix” mainly for completeness of this manuscript. Section A.1 discusses the mathematical justification for why the greedy procedure (2.11) is a good selection of parameter snapshots. Section A.2 gives an overview of the RBM procedure, and quantifies the computational complexity of the RBM algorithm. Careful scrutiny of this operational count illustrates why RBM algorithms can simulate parameterized problems with \mathcal{N} -independent complexity in the online phase of the algorithm.

Finally, Sect. A.3 discusses an efficient methodology to compute entries of the approximate Gramian \tilde{G} used by (3.4) in the CDM algorithm. This procedure is a relatively straightforward application of the offline–online decomposition already employed by RBM algorithms.

A.1 Greedy and Weak Greedy Algorithms

The best N -dimensional RB space $X_N^{\mathcal{N}}$ in X among all possible N -dimensional subspaces of the solution manifold $u(\cdot; \mathcal{D})$ is in theory the one with the smallest Kolmogorov N -width d_N [28]:

$$d_N [u(\cdot; \mathcal{D})] := \inf_{\substack{X_N \subset X^{\mathcal{N}} \\ \dim X_N = N}} \sup_{\mu \in \mathcal{D}} \inf_{v \in X_N} \|u(\cdot, \mu) - v\|_X \tag{A.1}$$

The identification of an exact-infimizer for the outer “inf” is usually infeasible, but a prominent approach is to employ a greedy strategy which locates this N -dimensional space hierarchically. A first sample set $S_1 = \{\mu^1\}$ is identified by randomly selecting μ^1 from Ξ_{train} ; its associated reduced basis space $X_1^{\mathcal{N}} = \text{span}\{u^{\mathcal{N}}(\mu^1)\}$ is likewise computed. Subsequently parameter values are greedily chosen as sub-optimal solutions to an $L^2(\Xi_{\text{train}}; X)$ optimization problem [33]: for $N = 2, \dots, N_{\text{max}}$, we find

$$\mu^N = \operatorname{argmax}_{\mu \in \Xi_{\text{train}}} \|u^{\mathcal{N}}(\mu) - u_{N-1}^{\mathcal{N}}(\mu)\|_{X^{\mathcal{N}}} \tag{A.2}$$

where $u_{N-1}^{\mathcal{N}}(\mu)$ is the RB solution (2.8) in the current $(N - 1)$ -dimensional subspace. Direct calculation of $u^{\mathcal{N}}(\mu)$ to solve this optimization problem over all μ is impractical. Therefore, an even weaker greedy algorithm is usually employed where we replace the error $\|u^{\mathcal{N}}(\mu) - u_{N-1}^{\mathcal{N}}(\mu)\|_X$ by an inexpensive and computable *a posteriori* bound Δ_{N-1} (see the next section). After identifying μ^N , the parameter snapshot set and the reduced basis space are augmented, $S_N = S_{N-1} \cup \{\mu^N\}$ and $X_N^{\mathcal{N}} = X_{N-1}^{\mathcal{N}} \oplus \{u(\mu^N)\}$, respectively.

A.2 Offline–Online Decomposition

The last component of RBM that we plan to review in this section is the Offline–Online decomposition procedure [33]. The complexity of the offline stage depends on \mathcal{N} which is performed only once in preparation for the subsequent online computation, whose complexity is independent of \mathcal{N} . It is in the \mathcal{N} -independent online stage where RBM achieves certifiable orders-of-magnitude speedup compared with other many-query approaches. The topic of this paper addresses acceleration of the offline portion of the RBM algorithm. In order to put this contribution of this paper in context, in this section we perform a detailed complexity analysis of the decomposition.

We let $N_{\text{train}} = |\Xi_{\text{train}}|$ denote the cardinality (size) of Ξ_{train} ; $N \leq N_{\text{max}}$ is the dimension of the reduced basis approximation computed in the offline stage. Computation of the lower bound $\alpha_{LB}^{\mathcal{N}}(\mu)$ is accomplished via the Successive Constraint Method [20].

During the online stage and for any new μ , the online cost of evaluating $\alpha_{LB}^{\mathcal{N}}(\mu)$ is negligible, but we use W_α to denote the *average* cost for evaluating these values over Ξ_{train} (this includes the offline cost). W_s is the operational complexity of solving problem (2.5) once by the chosen numerical method. For most discretizations, $\mathcal{N}^2 \lesssim W_s \leq \mathcal{N}^3$. Finally, W_m is the work to evaluate the $X^{\mathcal{N}}$ -inner product $(f, g)_{X^{\mathcal{N}}}$ which usually satisfies $\mathcal{N} \lesssim W_m \lesssim \mathcal{N}^2$. Using these notations we can present a rough operation count for the three components of the algorithm.

A.2.1 Online Solve and its Preparation

The system (2.9) is usually of small size: a set of N linear algebraic equations for N unknowns, with $N \ll \mathcal{N}$. However, the formation of the stiffness matrix involves $u^{\mathcal{N}}(\mu^n)$ for $1 \leq n \leq$

N ; direct computation with these quantities requires \mathcal{N} -dependent complexity. It is the affine parameter assumption (2.4) that allows us to circumvent complexity in the online stage. By (2.4), the stiffness matrix for (2.9) can be expressed as

$$\sum_{m=1}^N \sum_{q=1}^{Q_a} \Theta^q(\boldsymbol{\mu}) a^q(u^{\mathcal{N}}(\boldsymbol{\mu}^m), u^{\mathcal{N}}(\boldsymbol{\mu}^n)) u_{Nm}^{\mathcal{N}}(\boldsymbol{\mu}) = f(u^{\mathcal{N}}(\boldsymbol{\mu}^n)), \quad n = 1, \dots, N \quad (\text{A.3})$$

During the offline stage, we can precompute the Q_a matrices $a^q(u^{\mathcal{N}}(\boldsymbol{\mu}^m), u^{\mathcal{N}}(\boldsymbol{\mu}^n)) \in \mathbb{R}^{N \times N}$ for $q = 1, \dots, Q_a$ with a cost of order $\mathcal{N}^2 N^2 Q_a$. During the online phase, we need only assemble the reduced stiffness matrix according to (A.3), and solve the reduced $N \times N$ system. The total online operation count is thus of order $Q_a N^3 + N^4$.

A.2.2 Error Estimator Calculations

With a cost of order $Q_a N \mathcal{N}$ in the offline stage, we can calculate functions C and \mathcal{L}_m^q , $1 \leq m \leq N, 1 \leq q \leq Q_a$ both defined by

$$\begin{cases} (C, v) = f(v)_{X^{\mathcal{N}}} \quad \forall v \in X^{\mathcal{N}} \\ (\mathcal{L}_m^q, v)_{X^{\mathcal{N}}} = -a^q(u^{\mathcal{N}}(\boldsymbol{\mu}^m), v) \quad \forall v \in X^{\mathcal{N}}. \end{cases} \quad (\text{A.4})$$

Here, we assume that the X -inner product can be “inverted” with cost of order \mathcal{N} , i.e. that the mass matrix is block diagonal. The availability of C and \mathcal{L}_m^q facilitates an Offline–Online decomposition of the term $\|r_N(\cdot; \boldsymbol{\mu})\|_{(X^{\mathcal{N}})}$ in the error estimate (2.13) due to that its square can be written as

$$\begin{aligned} (C, C)_{X^{\mathcal{N}}} + 2 \sum_{q=1}^{Q_a} \sum_{m=1}^N \Theta^q(\boldsymbol{\mu}) u_{Nm}^{\mathcal{N}}(\boldsymbol{\mu}) (C, \mathcal{L}_m^q)_X \\ + \sum_{q=1}^{Q_a} \sum_{m=1}^N \Theta^q(\boldsymbol{\mu}) u_{Nm}^{\mathcal{N}} \left\{ \sum_{q'=1}^{Q_a} \sum_{m'=1}^N \Theta^{q'}(\boldsymbol{\mu}) u_{Nm'}^{\mathcal{N}} (\mathcal{L}_m^q, \mathcal{L}_{m'}^{q'})_{X^{\mathcal{N}}} \right\}. \end{aligned} \quad (\text{A.5})$$

Therefore, in the offline stage we should calculate and store $(C, C)_{X^{\mathcal{N}}}, (C, \mathcal{L}_m^q)_{X^{\mathcal{N}}}, (\mathcal{L}_m^q, \mathcal{L}_{m'}^{q'})_{X^{\mathcal{N}}}, 1 \leq m, m' \leq N_{RB}, 1 \leq q, q' \leq Q_a$. This cost is of the order $Q_a^2 N^3 W_m$. During the online stage, given any parameter $\boldsymbol{\mu}$, we only need to evaluate $\Theta^q(\boldsymbol{\mu}), 1 \leq q \leq Q_a, u_{Nm}^{\mathcal{N}}(\boldsymbol{\mu}), 1 \leq m \leq N$, and compute the sum (A.5). Thus, the online operation count for each $\boldsymbol{\mu}$ is $O(Q_a^2 N^3)$.

A.2.3 Greedy Sweeping

In the offline phase of the algorithm, we repeatedly sweep Ξ_{train} for maximization of the error estimator $\Delta_n(\boldsymbol{\mu}), 1 \leq n \leq N$. The offline cost includes:

- computing the lower bound $\alpha_{LB}^{\mathcal{N}}(\boldsymbol{\mu})$. The operation count is $O(N_{\text{train}} W_\alpha)$,
- sweeping the training set by calculating the reduced basis solution and the *a posteriori* error estimate at each location. The operation count of the former one is $O(N_{\text{train}} (Q_a N_{RB}^3 + N_{RB}^4))$. The operation count of the latter one is $O(N_{\text{train}} Q_a^2 N_{RB}^3)$.
- solving system (2.5) N times. The total operation count is $O(N W_s)$.

A.2.4 Summary

The total offline portion of the algorithm has complexity of the order

$$\underbrace{\text{Reduced solve preparation}}_{\mathcal{N}^2 N^2 Q_a} \quad \underbrace{\text{Greedy sweeping}}_{+ N_{\text{train}} W_\alpha + N_{\text{train}} (Q_a^2 N^3 + N^4) + N W_s} + \underbrace{\text{Estimator preparation}}_{Q_a^2 N^3 W_m} .$$

The total online cost including the error certification is of order $Q_a^2 N^3$.

A.3 Offline–Online Decomposition for the Approximate CDM-RBM Gramian \tilde{G}

The entries of the matrix \tilde{G} defined in (3.4) can be efficiently computed assuming that we can compute the approximate errors $\{\tilde{\epsilon}(\boldsymbol{\mu}) : \boldsymbol{\mu} \in \Xi_{\text{train}}\}$ in an offline–online fashion. To accomplish this, note that $\mathbb{A}_{\mathcal{N}}(\boldsymbol{\mu})(\mathbf{v}) = \sum_{k=1}^{Q_a} \theta_k^a(\boldsymbol{\mu}) A_k(\mathbf{v})$ by (2.4), where $A_k(\cdot)$ is a nonparametric matrix operator, so that

$$\begin{aligned} \tilde{\epsilon}(\boldsymbol{\mu}) &= \left(\sum_{m=1}^Q u_{N_m}^{\mathcal{N}}(\boldsymbol{\mu}) \mathbb{A}_{\mathcal{N}}^{-1}(\boldsymbol{\mu}^m) \right) \left(f^{\mathcal{N}} - \mathbb{A}_{\mathcal{N}}(\boldsymbol{\mu}) \left(\sum_{m=1}^N u_{N_m}^{\mathcal{N}}(\boldsymbol{\mu}) u^{\mathcal{N}}(\boldsymbol{\mu}^m) \right) \right) \\ &= \sum_{m=1}^Q u_{N_m}^{\mathcal{N}}(\boldsymbol{\mu}) \left(\mathbb{A}_{\mathcal{N}}^{-1}(\boldsymbol{\mu}^m) f^{\mathcal{N}} \right) - \sum_{m=1}^Q \sum_{m'=1}^N u_{N_m}^{\mathcal{N}}(\boldsymbol{\mu}) u_{N_{m'}}^{\mathcal{N}}(\boldsymbol{\mu}) \left(\mathbb{A}_{\mathcal{N}}^{-1}(\boldsymbol{\mu}^m) \mathbb{A}_{\mathcal{N}}(\boldsymbol{\mu}) \left(u^{\mathcal{N}}(\boldsymbol{\mu}^{m'}) \right) \right) \\ &= \sum_{m=1}^Q u_{N_m}^{\mathcal{N}}(\boldsymbol{\mu}) \left(\mathbb{A}_{\mathcal{N}}^{-1}(\boldsymbol{\mu}^m) f^{\mathcal{N}} \right) - \sum_{m=1}^Q \sum_{m'=1}^N \sum_{k=1}^{Q_a} \theta_k^a(\boldsymbol{\mu}) u_{N_m}^{\mathcal{N}}(\boldsymbol{\mu}) u_{N_{m'}}^{\mathcal{N}}(\boldsymbol{\mu}) \left(\mathbb{A}_{\mathcal{N}}^{-1}(\boldsymbol{\mu}^m) A_k \left(u^{\mathcal{N}}(\boldsymbol{\mu}^{m'}) \right) \right), \end{aligned}$$

Therefore, we can split this computation into offline and online components as follows:

- **Offline** Calculate $\mathbb{A}_{\mathcal{N}}^{-1}(\boldsymbol{\mu}^m) f^{\mathcal{N}}$ and $\mathbb{A}_{\mathcal{N}}^{-1}(\boldsymbol{\mu}^m) A_k \left(u^{\mathcal{N}}(\boldsymbol{\mu}^{m'}) \right)$ for $1 \leq m' \leq N$, $1 \leq m \leq Q$, $1 \leq k \leq Q_a$, $Q \leq N$, with complexity $O(\mathcal{N}^2 Q N Q_a)$.
- **Online** Evaluate the coefficients $u_{N_m}^{\mathcal{N}}(\boldsymbol{\mu})$ and $\theta_k^a(\boldsymbol{\mu}) u_{N_m}^{\mathcal{N}}(\boldsymbol{\mu}) u_{N_{m'}}^{\mathcal{N}}(\boldsymbol{\mu})$ and form $\tilde{\epsilon}(\boldsymbol{\mu})$. The online computation has complexity $O(\mathcal{N} Q N Q_a)$.

References

1. Babuška, I., Nobile, F., Tempone, R.: A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM Rev.* **52**(2), 317–355 (2010)
2. Barrault, M., Maday, Y., Nguyen, N.C., Patera, A.T.: An ‘empirical interpolation’ method: application to efficient reduced-basis discretization of partial differential equations. *C. R. Math. Acad. Sci. Paris* **339**(9), 667–672 (2004)
3. Berkooz, G., Holmes, P., Lumley, J.L.: The proper orthogonal decomposition in the analysis of turbulent flows. In *Ann. Rev. Fluid Mech.* **25**, 539–575. Annual Reviews, Palo Alto, CA, (1993)
4. Binev, P., Cohen, A., Dahmen, W., DeVore, R., Petrova, G., Wojtaszczyk, P.: Convergence rates for greedy algorithms in reduced basis methods. *SIAM J. Math. Anal.* **43**(3), 1457–1472 (2011)
5. Chen, P., Quarteroni, A.: Accurate and efficient evaluation of failure probability for partial different equations with random input data. *Comput. Methods Appl. Mech. Eng.* **267**, 233–260 (2013)
6. Chen, P., Quarteroni, A., Rozza, G.: A weighted reduced basis method for elliptic partial differential equations with random input data. *SIAM J. Numer. Anal.* **51**(6), 3163–3185 (2013)
7. Chen, P., Quarteroni, A., Rozza, G.: Multilevel and weighted reduced basis method for stochastic optimal control problems constrained by Stokes equations. *Numer. Math.* **133**(1), 67–102 (2016)
8. Chen, Y.: A certified natural-norm successive constraint method for parametric inf–sup lower bounds. *Appl. Numer. Math.* **99**, 98–108 (2016)

9. Chen, Y., Hesthaven, J.S., Maday, Y., Rodríguez, J.: Certified reduced basis methods and output bounds for the harmonic Maxwell's equations. *SIAM J. Sci. Comput.* **32**(2), 970–996 (2010)
10. Feldmann, P., Freund, R.W.: Efficient linear circuit analysis by padé approximation via the lanczos process. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.* **14**(5), 639–649 (1995)
11. Grepl, M.A., Maday, Y., Nguyen, N.C., Patera, A.T.: Efficient reduced-basis treatment of nonaffine and nonlinear partial differential equations. *M2AN. Math. Model. Numer. Anal.* **41**(3), 575–605 (2007)
12. Grepl, M.A., Patera, A.T.: A posteriori error bounds for reduced-basis approximations of parametrized parabolic partial differential equations. *ESAIM Math. Model. Numer. Anal.* **39**(1), 157–181 (2005)
13. Gugercin, S., Antoulas, A.C.: A survey of model reduction by balanced truncation and some new results. *Int. J. Control* **77**(8), 748–766 (2004)
14. Haasdonk, B.: Reduced basis methods for parametrized pdes—a tutorial introduction for stationary and instationary problems. *Reduced Order Modelling, Luminy Book Series* (2014)
15. Haasdonk, B.: *RBmatlab* 2016)
16. Haasdonk, B., Dihlmann, M., Ohlberger, M.: A training set and multiple bases generation approach for parameterized model reduction based on adaptive grids in parameter space. *Math. Comput. Model. Dyn. Syst.* **17**(4), 423–442 (2011)
17. Harbrecht, H., Peters, M., Schneider, R.: On the low-rank approximation by the pivoted Cholesky decomposition. *Appl. Numer. Math.* **62**(4), 428–440 (2012)
18. Hesthaven, J.S., Stamm, B., Zhang, S.: Efficient greedy algorithms for high-dimensional parameter spaces with applications to empirical interpolation and reduced basis methods? *ESAIM Math. Model. Numer. Anal.* **48**(1), 259–283 (2014)
19. Huynh, D.B.P., Knezevic, D.J., Chen, Y., Hesthaven, J.S., Patera, A.T.: A natural-norm successive constraint method for inf–sup lower bounds. *Comput. Methods Appl. Mech. Eng.* **199**(29–32), 1963–1975 (2010)
20. Huynh, D.B.P., Rozza, G., Sen, S., Patera, A.T.: A successive constraint linear optimization method for lower bounds of parametric coercivity and inf–sup stability constants. *C. R. Math. Acad. Sci. Paris* **345**(8), 473–478 (2007)
21. Jiang, J., Chen, Y., Narayan, A.: A goal-oriented reduced basis methods-accelerated generalized polynomial Chaos algorithm. *SIAM/ASA J. Uncertain. Quant.* **4**(1), 1398–1420 (2016)
22. Lassila, T., Quarteroni, A., Rozza, G.: A reduced basis model with parametric coupling for fluid-structure interaction problems. *SIAM J. Sci. Comput.* **34**(2), A1187–A1213 (2012)
23. Maday, Y., Patera, A.T., Turinici, G.: Global a priori convergence theory for reduced-basis approximations of single-parameter symmetric coercive elliptic partial differential equations. *C. R. Math. Acad. Sci. Paris* **335**(3), 289–294 (2002)
24. Moore, B.C.: Principal component analysis in linear systems: controllability, observability, and model reduction. *IEEE Trans. Autom. Control* **26**(1), 17–32 (1981)
25. Noor, A.K.: Recent advances in reduction methods for nonlinear problems. *Comput. Struct.* **13**(1–3), 31–44 (1981)
26. Noor, A.K., Peters, J.M.: Reduced basis technique for nonlinear analysis of structures. *AIAA J.* **18**(4), 455–462 (1980)
27. Patera, A., Rozza, G.: Reduced basis approximation and a posteriori error estimation for parametrized partial differential equations. Copyright MIT (2007)
28. Pinkus, A.: *n*-widths in approximation theory, volume 7 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin (1985)
29. Porsching, T.A.: Estimation of the error in the reduced basis method solution of nonlinear equations. *Math. Comput.* **45**(172), 487–496 (1985)
30. Porsching, T.A., Lee, M.-Y.L.: The reduced basis method for initial value problems. *SIAM J. Numer. Anal.* **24**(6), 1277–1287 (1987)
31. Quarteroni, A., Rozza, G.: Numerical solution of parametrized Navier–Stokes equations by reduced basis methods. *Numer. Methods Partial Differ. Equ.* **23**(4), 923–948 (2007)
32. Quarteroni, A., Rozza, G., Manzoni, A.: Certified reduced basis approximation for parametrized partial differential equations and applications. *J. Math. Ind.* **1**: Art. 3, 44 (2011)
33. Rozza, G., Huynh, D.B.P., Patera, A.T.: Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations: application to transport and continuum mechanics. *Arch. Comput. Methods Eng.* **15**(3), 229–275 (2008)
34. Saad, Y.: *Iterative methods for sparse linear systems*. Society for Industrial and Applied Mathematics, 2nd edn. (2003)
35. Sen, S.: Reduced-basis approximation and a posteriori error estimation for many-parameter heat conduction problems. *Numer. Heat Transf. Part B Fundam.* **54**(5), 369–389 (2008)

36. Urban, K., Volkwein, S., Zeeb, O.: Greedy sampling using nonlinear optimization. In: *Reduced Order Methods for Modeling and Computational Reduction*, pp. 137–157. Springer (2014)
37. Veroy, K., Patera, A.T.: Certified real-time solution of the parametrized steady incompressible Navier–Stokes equations: rigorous reduced-basis a posteriori error bounds. *Internat. J. Numer. Methods Fluids* **47**(8–9), 773–788 (2005)
38. Willcox, K., Peraire, J.: Balanced model reduction via the proper orthogonal decomposition. *AIAA J.* **40**(11), 2323–2330 (2002)
39. Xiu, D., Hesthaven, J.S.: High-order collocation methods for differential equations with random inputs. *SIAM J. Sci. Comput.* **27**(3), 1118–1139 (2005)