

An improved discrete least-squares/reduced-basis method for parameterized elliptic PDEs

Max Gunzburger · Michael Schneier
Clayton Webster · Guannan Zhang

Received: date / Accepted: date

Abstract It is shown that the computational efficiency of the discrete least-squares (DLS) approximation of solutions of stochastic elliptic PDEs is improved by incorporating a reduced-basis method into the DLS framework. In particular, we consider stochastic elliptic PDEs with an affine dependence on the random variable. The goal is to recover the entire solution map from the parameter space to the finite element space. To this end, first, a reduced-basis solution using a weak greedy algorithm is constructed, then a DLS approximation is determined by evaluating the reduced-basis approximation instead of the full finite element approximation. The main advantage of the new approach is that one only need apply the DLS operator to the coefficients of the reduced-basis expansion, resulting in huge savings in both the storage of the DLS coefficients and the online cost of evaluating the DLS approximation. In addition, the recently developed quasi-optimal polynomial space is also adopted in the new approach, resulting in superior convergence rates for a wider class of problems than previous analyzed. Numerical experiments are provided that illustrate the theoretical results.

Keywords discrete least squares · reduced basis · quasi-optimal polynomials · random coefficients · partial differential equations

Mathematics Subject Classification (2000) 11A22 · 11A22 · 11A22

This material is based upon work supported in part by the U.S. Air Force of Scientific Research under grants 1854-V521-12 and FA9550-15-1-0001; the U.S. Defense Advanced Research Projects Agency, Defense Sciences Office under contract and award HR0011619523 and 1868-A017-15; the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under contracts and awards ERKJ259, ERKJ320, de-sc0009324, and de-sc0010678; by the National Science Foundation under the award numbers 1620027 and 1620280.

M. Gunzburger and M. Schneier
Department of Scientific Computing, Florida State University, Tallahassee FL 32306-4120, USA. E-mail: mgunzburger@fsu.edu, mhs13c@my.fsu.edu

C. Webster and G. Zhang
Computational and Applied Mathematics Division, Oak Ridge National Laboratory, Oak Ridge TN 37831-6164, USA. E-mail: webstercg@ornl.gov, zhangg@ornl.gov

1 Introduction

Mathematical models are used to understand and predict the behavior of complex systems arising in applications. Common input data for these type of models include forcing terms, boundary conditions, model coefficients, and the computational domain itself. Often, for any number of reasons there is a degree of uncertainty involved with these inputs. In order to obtain an accurate model one must incorporate such uncertainties into the governing equations and quantify their effect on model outputs of interest. In this paper, we focus on systems that can be modeled by elliptic partial differential equations (PDEs) with random input data that has an affine dependence on the random variables. In particular, we consider cases for which that data is parameterized, i.e., the random coefficients/fields in the PDEs are functions of a *finite* number of random parameters. This means that the PDE solution, denoted by $u(\mathbf{x}, \mathbf{y})$, can be viewed as a function of an N -dimensional vector random parameters, denoted by $\mathbf{y} = (y_1, \dots, y_N)^\top$. Our goal is to recover the entire solution map $\mathbf{y} \rightarrow u(\mathbf{x}, \mathbf{y})$ from the parameter space to the solution space.

It is well known that commonly used Monte Carlo methods are not feasible options for this task because they can only be used to compute limited types of statistics. Sparse polynomial approximations [12–14, 18, 19, 24], including stochastic Galerkin, stochastic collocation, discrete least-squares (DLS), and compressive sensing methods, etc. These methods take advantages of the smoothness of the solution map $\mathbf{y} \rightarrow u(\mathbf{x}, \mathbf{y})$ to reduce the complexity of approximating that map in high-dimensional parameter space. The purpose of building sparse approximations (surrogates) is to enable fast evaluations of the surrogates, e.g., when conducting uncertainty quantification (UQ) tasks. However, existing sparse approximation techniques only focus on complexity reduction with respect to the parameter dependence, and largely ignore the huge cost (in evaluating the surrogates) arising from the finite element discretization. Specifically, denote by J the degrees of freedom of the finite element discretization and by M the dimension of the sparse polynomial space. To approximate the entire solution map $\mathbf{y} \rightarrow u(\mathbf{x}, \mathbf{y})$, we have to build a polynomial approximation for each of the J finite element coefficients so that the final sparse approximation requires an $M \times J$ *dense* matrix to store all the coefficients. When J is large, as it is in practical applications, the required storage may well not be affordable. Moreover, the complexity of each evaluation of the sparse approximation will be roughly of $\mathcal{O}(JM)$, which is considered as a perhaps prohibitive cost, given that computing accurate statistical information may require a very large number of such evaluations.

To overcome the above challenges, we propose to incorporate the well-studied reduced-basis technique [6, 7, 20, 25] into the sparse polynomial approximation. Although we focus on improving discrete least-squares methods, our approach can also be potentially generalized to the other aforementioned approaches. The main idea is to construct a reduced-basis solution using a greedy algorithm [6] to reduce the number of coefficients that are dependent on the parameter vector \mathbf{y} and then apply the DLS operator only to the coefficients of the reduced-basis approximation. For example, if the dimension K of the reduced-basis space is such that $K \ll J$, then the DLS coefficients can be stored in an $M \times K$ matrix that is much smaller than the straightforward DLS case. Moreover, with respect to computational complexity,

our approach requires only $\mathcal{O}(K(M+J))$ operations to perform matrix-vector productions for each evaluation of the DLS approximation, which is also significantly smaller to the $\mathcal{O}(MJ)$ operations required for the straightforward DLS approximation.

The plan for the rest of the paper is as follows. In Section 2, we introduce the mathematical setting and assumptions needed throughout the rest of the paper. In Section 3.1, we introduce quasi-optimal polynomial spaces and, in Section 3.2, we discuss the formulation of the least-squares problem in Hilbert spaces using the quasi-optimal polynomial space introduced in Section 3.1. Then, in Section 4, we introduce the reduced-basis method and discuss its incorporation into the least-squares framework. In Section 5, we discuss the computational complexity of the discrete least-squares method as well as the reduced-basis method and provide the results of numerical experiments that illustrate our findings.

2 Problem setting

Let D denote a bounded Lipschitz domain in \mathbb{R}^d , $d \in \{1, 2, 3\}$, with boundary ∂D . We consider the parameterized elliptic partial differential problem

$$\begin{cases} -\nabla \cdot (a(\mathbf{x}, \mathbf{y}) \nabla u(\mathbf{x}, \mathbf{y})) = f(\mathbf{x}) & \forall (\mathbf{x}, \mathbf{y}) \in D \times \Gamma \\ u(\mathbf{x}, \mathbf{y}) = 0 & \forall (\mathbf{x}, \mathbf{y}) \in \partial D \times \Gamma \end{cases} \quad (1)$$

for the unknown function $u(\mathbf{x}, \mathbf{y})$, where $f(\mathbf{x})$ and $a(\mathbf{x}, \mathbf{y})$ are given functions and \mathbf{y} denotes a vector of parameters. In the stochastic setting, we have that \mathbf{y} is a random variable distributed according to a joint probability density function (PDF) $\rho(\mathbf{y})$. Although the case of a countably infinite number of random variables is of interest in some applications, here we assume, as is often the case, that the randomness present in a stochastic PDE can be approximated well in terms of a finite number of random variables. Therefore, we assume that the parameterized diffusion coefficient $a(\mathbf{x}, \mathbf{y})$ depends on a finite number N of random variables denoted by the vector $\mathbf{y} = (y_1, \dots, y_N)^\top \in \Gamma \subset \mathbb{R}^N$, where Γ denotes a parameter domain.

Here, we specialize to the case for which the components of \mathbf{y} are independent and identically distributed random variables so that Γ is a hyper-rectangle in \mathbb{R}^N and $\rho(\mathbf{y}) = \prod_{n=1}^N \rho_n(y_n)$ is a product of N one-dimensional PDFs. Without loss of generality, we can then assume that $\Gamma = [-1, 1]^N$. We note that although the assumption that the random variables are i.i.d. may appear restrictive, in practice, a wide range of problems can still be addressed. As discussed in [3], problems with non-independent random variables can be addressed via the introduction of auxiliary density functions.

We make several assumptions. First, we assume that there exist constants $0 < a_{\min} < a_{\max} < \infty$ such that

$$a_{\min} < a(\mathbf{x}, \mathbf{y}) < a_{\max} \quad \forall \mathbf{x} \in D, \text{ a.s. for } \mathbf{y} \in \Gamma. \quad (2)$$

Let $Y = L^2_\rho(\Gamma)$ denote the space of square integrable functions on Γ with respect to the weight $\rho(\mathbf{y})$. We also have the standard Sobolev space $X = H_0^1(D)$ equipped with

the norm $\|v\|_X = (\int_D \nabla v \cdot \nabla v d\mathbf{x})^{1/2}$; $X' = H^{-1}(D)$ denotes the corresponding dual space. Then, a weak formulation of (1) is, given $f \in X'$, to find $u \in X \otimes Y$ such that

$$\begin{aligned} & \int_{\Gamma} \int_D a(\mathbf{x}, \mathbf{y}) \nabla u(\mathbf{x}, \mathbf{y}) \cdot \nabla v(\mathbf{x}, \mathbf{y}) \rho(\mathbf{y}) d\mathbf{x} d\mathbf{y} \\ & = \int_{\Gamma} \int_D f(\mathbf{x}) v(\mathbf{x}, \mathbf{y}) \rho(\mathbf{y}) d\mathbf{x} d\mathbf{y} \quad \forall v \in X \otimes Y, \end{aligned} \quad (3)$$

where $X \otimes Y := L^2_{\rho}(\Gamma; X) := \{u : \int_{\Gamma} \|u(\cdot, \mathbf{y})\|_X^2 \rho(\mathbf{y}) d\mathbf{y} < \infty\}$. By (2) and the Lax-Milgram theorem, there exists a unique solution of (3) for any $f \in X'$ and that solution satisfies the bound

$$\|u\|_{X \otimes Y} \leq \frac{1}{a_{\min}} \|f\|_{X'}. \quad (4)$$

Our final assumptions about the coefficient $a(\mathbf{x}, \mathbf{y})$ is affine dependence on the random variables, i.e., it can be written in the form

$$a(\mathbf{x}, \mathbf{y}) = a_0(\mathbf{x}) + \sum_{n=1}^N a_n(\mathbf{x}) y_n \quad (5)$$

for some $a_n(\mathbf{x})$, $n \in \{0, \dots, N\}$. A coefficient of this form could be a truncated Karhunen-Loève (KL) expansion. The affine dependence is necessary to achieve satisfactory efficiency in constructing a reduced basis using greedy algorithms. We also note that this assumption implies the complex continuation of $a(\mathbf{x}, \mathbf{y})$, represented as the map $a(\mathbf{x}, \mathbf{y}) : \mathbb{C}^N \rightarrow L^{\infty}(D)$, is an $L^{\infty}(D)$ -valued *holomorphic function* on \mathbb{C}^N . This allows for the use of the quasi-optimal polynomial space introduced in Section 3.1.

For spatial discretization, we use standard finite element methods. Let $X_h^{\text{fe}} \subset X$ denote a standard finite element space of dimension J and let $\{\phi_j(\mathbf{x})\}_{j=1}^J$ denote a basis for X_h^{fe} that consists of piecewise-continuous polynomials defined with respect to a regular triangulation \mathcal{T}_h of D , where $h > 0$ denotes the maximum mesh spacing. For any $\mathbf{y} \in \Gamma$, the finite element approximation $u_h(\mathbf{x}, \mathbf{y}) \in X_h^{\text{fe}}$ is determined by solving

$$A(u_h(\mathbf{y}), v; \mathbf{y}) = (f, v) \quad \forall v \in X_h^{\text{fe}}, \quad (6)$$

where $A(u, v) := \int_D a(\mathbf{x}, \mathbf{y}) \nabla u(\mathbf{x}, \mathbf{y}) \cdot \nabla v(\mathbf{x}, \mathbf{y}) d\mathbf{x}$ is the bilinear form corresponding to the operator in (3) and (\cdot, \cdot) denotes the $L^2(D)$ inner product. Then, for any $\mathbf{y} \in \Gamma$, we have, for sufficient small h and for a constant C_S whose value is independent of h and \mathbf{y} , the error estimate

$$\|u - u_h\|_X \leq C_S h^{\alpha}. \quad (7)$$

The convergence rate α depends on the spatial regularity of u and the degree of the polynomial used. For a detailed treatment of finite element error analyses, see, e.g., [8].

3 Discrete least-squares approximation

In this section, we recall the formulation and theoretical results about random discrete L^2 approximations of solutions of the parameterized PDE problem (1). This presentation is brief and only discusses the DLS approximation for Hilbert-valued functions. For a more general and comprehensive analysis, see, e.g., [13, 22]. For the sake of further simplifying the exposition, we assume that the random variables are uniformly distributed on $[-1, 1]$ so that $\rho_n = \frac{1}{2}$ for $n = 1, \dots, N$.

3.1 Quasi-optimal polynomial spaces

The first step towards building a DLS approximation is to choose an appropriate polynomial space in $L_\rho(\Gamma)$. Because we assume a uniform measure, we use Legendre polynomials that are orthogonal with respect to this measure. Letting $\mathbf{v} = (v_1, v_2, \dots, v_N) \in \mathbb{N}_0^N$ denote a multi-index, the multidimensional Legendre polynomials are denoted by $L_{\mathbf{v}}(\mathbf{y}) = \prod_{i=1}^N L_{v_i}(y_i)$, where $L_{v_n}(y_n)$ denotes the L^2 -normalized one-dimensional Legendre polynomials [5].

In the construction of polynomial approximations with respect to parameter dependences, one wishes to select a multi-index set $\Lambda_M \subset \{\mathbf{v} = (v_1, v_2, \dots, v_N) : v_n \in \mathbb{N}_0\}$ such that the corresponding polynomial space $\text{span}\{L_{\mathbf{v}}(\mathbf{y}), \mathbf{v} \in \Lambda_M\}$ yields maximal accuracy for a given dimension M . To achieve this, there are two approaches that have been extensively studied [14, 17, 18]. The first approach is known as *best M -term approximation*. Using a truncated Legendre expansion and using the triangle inequality, we can express the error of the approximation in the form

$$\left\| u(\mathbf{y}) - \sum_{\mathbf{v} \in \Lambda_M} c_{\mathbf{v}} L_{\mathbf{v}}(\mathbf{y}) \right\|_{X \otimes Y} \leq \sum_{\mathbf{v} \in \Lambda_M} \|c_{\mathbf{v}}\|_X, \quad (8)$$

where Λ_M is chosen such that the error (8) is minimized. This means that the indices $\mathbf{v} \in \Lambda_M$ correspond to the M largest values of $\|c_{\mathbf{v}}\|_X$. However, in practice, finding the best index set and polynomial space is an infeasible task because it requires computation of *all* the coefficients $c_{\mathbf{v}}$.

An alternative approach that tends to be less computationally intensive is referred to as *quasi-optimal polynomial approximation* [3, 26]. Rather than explicitly computing the coefficients in order to evaluate $\|c_{\mathbf{v}}\|_X$, we instead compute sharp estimates for $\|c_{\mathbf{v}}\|_X$ and use these to determine a quasi-optimal index set Λ_M . It has been shown that this method can achieve convergence rates similar to those of the best M -term approximation.

In order to establish a bound on the coefficients of the Legendre expansion, we need the following definition concerning uniform ellipticity in polyellipses.

Definition 1 For $0 < \delta < a_{\min}$ and $\boldsymbol{\varphi}$ denoting the sequence $\{\varphi_i\}_{i=1}^N$ with $\varphi_i > 1 \forall i$, we say the random field $a(\mathbf{x}, \cdot)$ satisfies the $(\delta, \boldsymbol{\varphi})$ -polyellipse uniform ellipticity assumption if it holds that

$$\mathcal{R}(a(\mathbf{x}, \mathbf{z})) \geq \delta$$

for all $\mathbf{x} \in D$ and $\mathbf{z} = \{z_i\}_{i=1}^N$ contained in the polyellipse

$$\mathcal{E} = \bigotimes_{1 \leq i \leq N} \left\{ z_i \in \mathbb{C} : \Re(z_i) = \frac{\varphi_i + \varphi_i^{-1}}{2} \cos(\theta), \Im(z_i) = \frac{\varphi_i - \varphi_i^{-1}}{2} \sin(\theta), \theta \in [0, 2\pi) \right\}.$$

It has been shown [26], for any diffusion coefficient $a(\mathbf{x}, \mathbf{y})$ satisfying the coercivity assumption in (2) and having the holomorphic parameter dependence, there always exists one $\boldsymbol{\varphi}$ for which this property is satisfied. Now, using this regularity condition, the holomorphy of the solution with respect to the random parameters follows and the bound on the coefficients of the L^2 -normalized Legendre expansion

$$\|c_{\mathbf{v}}\|_X \leq C_{\boldsymbol{\varphi}, \delta} \boldsymbol{\varphi}^{-\mathbf{v}} \prod_{i=1}^N \sqrt{2v_i + 1} \quad (9)$$

holds, where $C_{\boldsymbol{\varphi}, \delta} = \frac{\|f\|_{Y'}}{\delta} \prod_{i=1}^N \frac{\ell(\mathcal{E}_{\varphi_i})}{4(\varphi_i - 1)}$ with $\ell(\mathcal{E}_{\varphi_i})$ denoting the perimeter of the ellipse \mathcal{E}_{φ_i} . Note that Definition 1 holds for an infinite combination of $(\delta, \boldsymbol{\varphi})$ that we denoted by \mathbf{Ad} . For a given \mathbf{v} , the best coefficient bound will then be given by

$$\|c_{\mathbf{v}}\|_X \leq \inf_{(\delta, \boldsymbol{\varphi}) \in \mathbf{Ad}} C_{\boldsymbol{\varphi}, \delta} \boldsymbol{\varphi}^{-\mathbf{v}} \prod_{i=1}^N \sqrt{2v_i + 1}.$$

Solving this minimization problem is in general computationally infeasible. However, in case the basis functions a_k have non-overlapping supports, $\boldsymbol{\varphi}$ can be determined easily [3]. Problems with both overlapping support and nonoverlapping support are explored further in Section 5. We can now state an asymptotic bound for the quasi-optimal M -term approximation as follows:

Proposition 1 *Consider the Legendre series $\sum_{\mathbf{v} \in \Lambda} c_{\mathbf{v}} L_{\mathbf{v}}$ for u . Assume that (9) holds for all $\mathbf{v} \in \Lambda$. Let $\log(\varphi_n) = \lambda_n$ and Λ_M denote the set of indices corresponding to the M largest bounds in (9) determined by*

$$\Lambda_M = \left\{ \mathbf{v} \in \Lambda : \sum_{n=1}^N (2\lambda_n v_n - \log(2v_n + 1)) \leq j \right\} \quad (10)$$

for a given $j \in \mathbb{N}$. Then, for any $0 < \mu < 1$, it follows that

$$\left\| u - \sum_{\mathbf{v} \in \Lambda_M} c_{\mathbf{v}} L_{\mathbf{v}} \right\|_{X \otimes Y}^2 \leq C_{\boldsymbol{\varphi}, \delta}^2 C_u(\mu) M \exp\left(-2 \left(\frac{MN! \prod_{n=1}^N \lambda_n}{(1 + \mu)}\right)^{\frac{1}{N}}\right) \quad (11)$$

as j (and therefore M) goes to ∞ . Here, $C_u(\mu) = (4e + 4\mu e - 2) \frac{e}{e-1}$.

The index set should be chosen such that its size M allows for a specified level of accuracy to be reached using estimate (11). The value $0 < \mu < 1$ is related to the cardinality of our polynomial approximation, and decreases towards 0 as the cardinality of our polynomial approximation increases. A sharp mathematical formula for μ , given M , is currently an open problem, though it has been shown that even for a moderate value of μ one can still obtain a strong rate of convergence. For full details see [26, Section 4].

3.2 Discrete least-squares approximation in quasi-optimal spaces

Here we introduce the DLS method for approximating solutions of parametric PDE in (1) in the quasi-optimal polynomial space discussed above. This presentation is brief and only discusses the least-squares approximation for Hilbert-valued functions. For a more general and deeper analysis, see, e.g., [13, 22].

Let Y_{Λ_M} denote an M -dimensional quasi-optimal subspace in Y . We intend to build a DLS approximation in $X \otimes Y_{\Lambda_M}$ of the solution $u(\mathbf{x}, \mathbf{y}) \in X \otimes Y$ by the orthogonal projection, i.e.,

$$P_M[u] := \arg \min_{v \in X \otimes Y_{\Lambda_M}} \|u - v\|_{X \otimes Y}.$$

Letting $\boldsymbol{\ell}_M(\mathbf{y}) := (\ell_1(\mathbf{y}), \dots, \ell_M(\mathbf{y}))^\top$ denote the vector of re-indexed Legendre basis functions $\{\ell_{\mathbf{v}}(\mathbf{y}) : \mathbf{v} \in \Lambda_M\}$ for the subspace Y_{Λ_M} , we have that

$$P_M[u] = \sum_{m=1}^M c_m(\mathbf{x}) \ell_m(\mathbf{y}) \quad \text{with} \quad c_m(\mathbf{x}) = \langle u(\mathbf{x}, \cdot), \ell_m(\cdot) \rangle_Y \quad \text{for } m = 1, \dots, M,$$

where $\langle \cdot, \cdot \rangle_Y$ denotes the inner product on Y .

In general, we do not have available the solution of the PDE for all $\mathbf{y} \in \Gamma$, but only at a set of points $\{\mathbf{y}_i\}_{i=1}^S$, where $\mathbf{y}_i \in \Gamma$ are i.i.d. random variables distributed according to some distribution. We then consider the discrete (with respect to the \mathbf{y} dependence) least-squares problem

$$P_{M,S}[u] := \arg \min_{v \in X \otimes Y_{\Lambda_M}} \sum_{i=1}^S \|u(\mathbf{x}, \mathbf{y}_i) - v(\mathbf{x}, \mathbf{y}_i)\|_X^2 \quad (12)$$

that has a unique solution as long as $M \leq S$.

In practice, we do not have access to the exact solution $u(\mathbf{x}, \mathbf{y}_i)$ for $\mathbf{y}_i \in \Gamma$, so that we apply the DLS operator to the finite element solution $u_h \in X_h^{\text{fe}} \otimes Y$ and obtain the L^2 projection $P_{M,S}[u_h]$ in the subspace $X_h^{\text{fe}} \otimes Y_{\Lambda_M}$ which has the form

$$P_{M,S}[u_h] = \sum_{m=1}^M \sum_{j=1}^J c_{mj} \phi_j(\mathbf{x}) \ell_m(\mathbf{y}), \quad (13)$$

recalling that $\phi_j(\mathbf{x})$ are the finite element basis functions introduced in section 2. Letting $[\boldsymbol{\Phi}]_{ij} = \ell_j(\mathbf{y}_i)$, the coefficients $\{c_{mj}\}_{m=1, j=1}^{M, J}$ are the solution of the following linear system:

$$(\boldsymbol{\Phi}^\top \boldsymbol{\Phi}) \mathbf{C} = \boldsymbol{\Phi}^\top \mathbf{U}, \quad (14)$$

where $[\mathbf{C}]_{mj} = c_{mj}$ for $m = 1, \dots, M$, $j = 1, \dots, J$ and $[\mathbf{U}]_{ij} = u_h(\mathbf{x}_j, \mathbf{y}_i)$ for $i = 1, \dots, S$, $j = 1, \dots, J$.

4 Improved DLS methods based on reduced-basis solutions

The main purpose of building DLS approximations is to reduce the cost of obtaining approximate solutions of the PDE problem (1) at a large set of samples in Γ , i.e., reducing the online cost. We observe that the need to reduce costs is only necessary when the finite element degrees of freedom J is extremely large. Otherwise, for a small J , a classic finite element solver will be efficient enough to be used as an online solver. However, for a very large J , we can see from (14) that the coefficient \mathbf{C} , which is an $M \times J$ dense matrix, may require an unaffordable amount of storage. Moreover, the complexity of each evaluation of the DLS approximation would be of $\mathcal{O}(JM)$. To avoid such inefficiencies in both storage and computation, we propose to develop a new DLS method based on reduced-basis approximations of solutions of (1). A brief overview of a reduced-basis method is given in Section 4.1 and our approach is introduced in Section 4.2.

4.1 Reduced-basis methods

We briefly recall the reduced-basis technique; for a more in depth discussions about reduced-basis methods, see [7, 25]. The main idea of reduced-basis methods is to collect a set of deterministic solutions of the stochastic problem in (1) at a subset of the most representative samples in Γ , then uses these solutions as a basis to approximate solutions at other points in Γ through Galerkin projection. Specifically, when we have a subset, denoted by $\Xi_K := \{\mathbf{y}_i\}_{i=0}^{K-1}$, consisting of K representative samples, we can then solve the finite element system in (6) K times to obtain the set of K solutions $\{u_h^i(\mathbf{x}) = u_h(\mathbf{x}, \mathbf{y}_i)\}_{i=0}^{K-1}$. Using these solutions (snapshots), we can define a K -dimensional reduced space

$$X_K^{\text{rb}} = \text{span}\{u_h^i(\mathbf{x})\}_{i=0}^{K-1} \subset X_h^{\text{fe}}.$$

Then, we can construct a reduced-basis approximation $u_{h,K}$ by projecting u_h into X_K^{rb} , i.e., seeking

$$u_{h,K}(\mathbf{x}, \mathbf{y}) = \sum_{i=0}^{K-1} w_i^k(\mathbf{y}) \xi_i(\mathbf{x}) \in X_K^{\text{rb}}, \quad (15)$$

satisfying

$$A(u_{h,K}(\mathbf{y}), v; \mathbf{y}) = (f, v) \quad \forall v \in X_K^{\text{rb}}, \quad (16)$$

where $A(u, v; \mathbf{y})$ is the bilinear form defined in (6) and $\{\xi_i\}_{i=0}^{K-1}$ is the orthogonalized reduced basis of X_K^{rb} . Note that, for each $\mathbf{y} \in \Gamma$, the equation in (16) is equivalent to a linear system of K algebraic equations for the coefficients $\{w_i^k(\mathbf{y})\}_{i=0}^{K-1}$ in (15). In this way, when $K \ll J$, the computational cost of approximating $u(\mathbf{x}, \mathbf{y})$ for each $\mathbf{y} \in \Gamma$ is significantly reduced from solving a $J \times J$ linear system to solving a $K \times K$ linear system.

Now the question is how does one determine a good set of K samples $\Xi_K := \{\mathbf{y}_i\}_{i=0}^{K-1}$? Suppose one has in hand a set of samples Ξ_k ; one could start with $k = 0$, i.e., a single sample chosen at random or at the center of Γ . Then given the current set Ξ_k of samples, how does one find the next sample \mathbf{y}_k in an effective and efficient way

so as to improve the accuracy of the reduced-basis solution. The ideal choice is to use the *greedy algorithm* [6], i.e., find the next sample \mathbf{y}_k by solving the optimization problem

$$\mathbf{y}_k = \arg \sup_{\mathbf{y} \in \Gamma} \|u_{h,k}(\cdot, \mathbf{y}) - u_h(\cdot, \mathbf{y})\|_X, \quad (17)$$

i.e., locating the point $\mathbf{y}_k \in \Gamma$ at which the error between the current reduced-basis approximation and the finite element approximation is the largest. Unfortunately, solving the optimization problem (17) is not practical because it requires full information about the exact finite element solution $u_h(\mathbf{x}, \mathbf{y})$. To circumvent this issue, a variant of the greedy strategy (17), i.e., the *weak greedy algorithm*, has been shown to be computationally feasible in the context of solving parameterized PDEs [25]. The key idea of the weak greedy strategy is to find an accurate and computationally efficient surrogate of the error $u_{h,k}(\cdot, \mathbf{y}) - u_h(\cdot, \mathbf{y})$, and replace the true error in (17) with the surrogate to solve the optimization problem. To this end, we use the Galerkin residual as the surrogate to implement the weak greedy algorithm. Letting $e_{h,k}(\mathbf{x}, \mathbf{y}) := u_h(\mathbf{x}, \mathbf{y}) - u_{h,k}(\mathbf{x}, \mathbf{y}) \in X_h^{\text{fe}}$, we then have that, for any $\mathbf{y} \in \Gamma$,

$$R(v; \mathbf{y}) := A(e_{h,k}(\cdot, \mathbf{y}), v; \mathbf{y}) = (f, v) - A(u_{h,k}(\cdot, \mathbf{y}), v; \mathbf{y}) \quad \forall v \in X_h^{\text{fe}}. \quad (18)$$

Thanks to Riesz representation, we have $(\hat{e}_{h,k}, e_{h,k})_X = R(e_{h,k}; \mathbf{y})$, such that

$$\|e_{h,k}(\mathbf{x}, \mathbf{y})\|_X = \|u_h(\mathbf{x}, \mathbf{y}) - u_{h,k}(\mathbf{x}, \mathbf{y})\|_X \leq \frac{1}{\alpha_{\text{LB}}(\mathbf{y})} \|\hat{e}_{h,k}\|_X,$$

where $\alpha_{\text{LB}}(\mathbf{y}) = \min_{\mathbf{x} \in D} a(\mathbf{x}, \mathbf{y})$. Thus, we can replace $e_{h,k}$ with $\hat{e}_{h,k}$. In this effort, we also have to replace the search over all $\mathbf{y} \in \Gamma$ by a search over a discrete training set; for solutions manifolds that are sufficiently smooth, this step does not introduce unmanageable errors. Specifically, the construction of the reduced-basis method begins by choosing a training set Ξ_{train} of S_{train} points in Γ ; these points could be chosen randomly according to the joint PDF $\rho(\mathbf{y})$ associated with the random parameters $\mathbf{y} \in \Gamma$ or could be chosen deterministically. Then, the optimization problem in (17) is solved within the training set, i.e., \mathbf{y}_k is generated by

$$\mathbf{y}_k = \arg \sup_{\mathbf{y} \in \Xi_{\text{train}}} \left\{ \frac{1}{\alpha_{\text{LB}}(\mathbf{y})} \|\hat{e}_{h,k}(\cdot, \mathbf{y})\|_X \right\}. \quad (19)$$

The term $\alpha_{\text{LB}}(\mathbf{y})$ will be calculated over the training set using the Successive Constraint Method outlined in [20]

Due to the affine property of the coefficient in (5), the residual $R(v; \mathbf{y})$ in (18) can be decomposed as

$$\begin{aligned} R(v; \mathbf{y}) &= (f, v) - A_0(u_{h,k}(\cdot, \mathbf{y}), v) + \sum_{n=1}^N A_n(u_{h,k}(\cdot, \mathbf{y}), v; \mathbf{y}) y_n \\ &= (f, v) - \sum_{i=0}^{k-1} \left(A_0(\xi_i, v) + \sum_{n=1}^N A_n(\xi_i, v) y_n \right) w_i^k(\mathbf{y}), \end{aligned}$$

for all $v \in X_h^{\text{fe}}$. Due to the linearity of the above representation, we can determine $\hat{e}_{h,k}(\mathbf{x}, \mathbf{y}) \in X_h^{\text{fe}}$ efficiently by solving the following set of problems offline:

$$\begin{cases} (\hat{e}^f, v)_X = (f, v) & \forall v \in X_h^{\text{fe}}, \\ (\hat{e}_{n,i}, v)_X = A_n(\xi_i, v) & \forall v \in X_h^{\text{fe}} \text{ for } i = 0, \dots, k-1 \text{ and } n = 0, \dots, N, \end{cases} \quad (20)$$

such that $\hat{e}_{h,k}$ can be computed very efficiently by

$$\hat{e}_{h,k}(\mathbf{x}, \mathbf{y}) = \hat{e}^f(\mathbf{x}) - \sum_{i=0}^{k-1} w_i^k(\mathbf{y}) \left[\hat{e}_{0,i}(\mathbf{x}) + \sum_{n=1}^N y_n \hat{e}_{n,i}(\mathbf{x}) \right] \in X_h^{\text{fe}}.$$

We note that the quantities $\hat{e}^f(\mathbf{x})$, $\hat{e}_{0,i}(\mathbf{x})$, and $\hat{e}_{n,i}(\mathbf{x})$ are independent of \mathbf{y} and can therefore be stored in an offline phase.

To terminate the greedy procedure, we can preset some error tolerance ε_{tol} and end the algorithm when the approximation error is judged to be sufficiently small, i.e., $\{\|\hat{e}_{h,k}/\alpha_{\text{LB}}(\mathbf{y})\|_X \leq \varepsilon_{\text{rb}}, \forall \mathbf{y} \in \Xi_{\text{train}}\}$. In addition, how the points should be selected and the size of the training set is extremely problem dependent. In practice there are two approaches commonly used to construct Ξ_{train} . The first is an adaptive approach which starts with a small number of sample points and then greedily enriches the sample space based on these initial points; see [21]. The other method is to randomly sample the parameter space Γ according to the probability distribution associated with the problem.

4.2 Reduced-basis discrete least-squares (RB-DLS) approximation

As already mentioned, the goal of using reduced-basis approximations in the least-squares setting is to reduce the online cost, i.e., the cost of evaluating the final DLS approximation. Letting K denote the final value of k upon termination of the greedy algorithm, we observe that the parameter dependence of the reduced-basis solution in (15) only appears in the coefficients $\mathbf{w}_K(\mathbf{y}) := (w_1^K(\mathbf{y}), \dots, w_K^K(\mathbf{y}))^\top$ which is vector of size $K \ll J$. Thus, instead of applying the DLS operator $P_{M,S}[\cdot]$ to u_h , we apply it to the reduced-basis solution $u_{h,K}$, i.e.,

$$P_{M,S}[u_{h,K}] = \sum_{k=0}^K \left(\sum_{m=1}^M c_{mk}^{\text{rb}} \ell_m(\mathbf{y}) \right) \xi_k(\mathbf{x}) \quad (21)$$

which is equivalent to approximating the coefficient vector $\mathbf{w}_K(\mathbf{y})$ using the DLS method. The algebraic formulation for solving the coefficients c_{mk}^{rb} is

$$(\Phi^\top \Phi) \mathbf{C}^{\text{rb}} = \Phi^\top \mathbf{W}, \quad (22)$$

where $[\mathbf{C}^{\text{rb}}]_{mk} = c_{mk}^{\text{rb}}$ for $m = 1, \dots, M$, $k = 0, \dots, K$ and $[\mathbf{W}]_{ik} = w_i^K(\mathbf{y}_k)$ for $i = 1, \dots, S$, $k = 0, \dots, K$. Then, for each new sample $\mathbf{y} \in \Gamma$, the evaluation of the RB-DLS approximation in $P_{M,S}[u_{h,K}]$ can be conducted by

$$\mathbf{u}^{\text{RB-DLS}} = \mathbf{V}(\mathbf{C}^{\text{rb}})^\top \ell_M(\mathbf{y}), \quad (23)$$

where $\mathbf{u}^{\text{RB-DLS}} := (P_{M,S}[u_{h,K}](\mathbf{x}_1), \dots, P_{M,S}[u_{h,K}](\mathbf{x}_J))^\top$ and $\mathbf{V} = (\xi_0, \dots, \xi_K)$ is the reduced-basis matrix. In comparison, evaluating the classic DLS approximation $P_{M,S}[u_h]$ in (13) has to be done by

$$\mathbf{u}^{\text{DLS}} = \mathbf{C}^\top \boldsymbol{\ell}_M(\mathbf{y}), \quad (24)$$

where $\mathbf{u}^{\text{DLS}} := (P_{M,S}[u_h](\mathbf{x}_1), \dots, P_{M,S}[u_h](\mathbf{x}_J))^\top$ and \mathbf{C} is given in (14). The advantages of (23) over (24) can be seen from two aspects. In terms of storage, (23) only requires storage for a $J \times K$ matrix \mathbf{V} and an $M \times K$ matrix \mathbf{C}^{rb} , whereas (24) requires storage for an $M \times J$ matrix \mathbf{C} . Thus, when K is small, (23) requires much less storage than (24). In other words, the matrix $\mathbf{C}^{\text{rb}}\mathbf{V}^\top$ can be viewed as a *low-rank* (i.e., rank K) approximation of the matrix \mathbf{C} . In terms of computation, for each $\mathbf{y} \in \Gamma$, (23) requires $\mathcal{O}(K(M+J))$ operations to perform matrix-vector products, whereas (24) requires $\mathcal{O}(MJ)$ operations. This is another significant savings achieved by using our approach.

The total error of the RB-DLS approximation $P_{M,S}[u_{h,K}]$ can be split into the sums of the finite element discretization error, the reduced-basis error, and the DLS projection error, i.e.,

$$\begin{aligned} & \mathbb{E} [\|u - P_{M,S}[u_{h,K}]\|_X^2] \\ & \leq \underbrace{\mathbb{E} [\|u - u_h\|_X^2]}_{e_I} + \underbrace{\mathbb{E} [\|u_h - u_{h,K}\|_X^2]}_{e_{II}} + \underbrace{\mathbb{E} [\|u_{h,K} - P_{M,S}[u_{h,K}]\|_X^2]}_{e_{III}}. \end{aligned} \quad (25)$$

The first error e_I is easy to control/balance based on the classic finite element error analysis. The second error e_{II} is essentially controlled by the Kolmogorov width associated with the solution manifold in the finite element space X_h^{fe} . In the recent works [2, 16], it has been shown for the class of problems dealt with in this paper that the Kolmogorov width will decay at least algebraically; this gives us hope that our reduced-basis method will be successful. In practice it is difficult to determine the decay of the RB error a-priori, hence we use the *a-posteriori* estimate to balance the second error e_{II} by adjusting the threshold ε_{tol} .

Defining the term

$$Z(\Lambda_M) := \sum_{\mathbf{v} \in \Lambda_M} \|L_{\mathbf{v}}\|_{L^\infty}^2,$$

the third error e_{III} can be bounded for any $r > 0$ by [15]

$$e_{III} \leq (1 + \beta(S))e_M(u)^2 + 8H^2S^{-r}$$

as long as the number of sample points satisfies

$$\frac{S}{\ln(S)} \geq \frac{Z(\Lambda_M)}{\kappa}, \quad \kappa := \frac{1 - \ln(2)}{2 + 2r},$$

where $\beta(S) \rightarrow 0$ as $S \rightarrow +\infty$, H is the uniform upper bound of u , and $e_M(u)$ is the error in the best M -term approximation of $u_{h,K}$. The term r is related to the stability of the least squares system, full details can be found in [15]. As shown in [13], when using Legendre polynomials, the quantity $Z(\Lambda_M)$ can be bounded by $M \leq Z(\Lambda_M) \leq M^2$, when Λ_M is a lower set. When using Chebyshev polynomials, a better bound can be obtained for lower sets, i.e., $Z(\Lambda_M) \leq \min(M^{\log 3 / \log 2}, 2^N M)$. Even though the

polynomial space used in this work is not a lower set, the quasi-optimal polynomial space can be covered with a lower set which is only slightly larger allowing us to effectively use these bounds. As will be illustrated in section 5 to maintain stability of the DLS method using the quasi-optimal polynomial space sample points only on the order of $3M$ are required. Therefore the slight theoretical oversampling due to the lack of the quasi-optimal polynomial space not being a lower set does not have an impact in practice. With the use of the quasi-optimal error estimate in Proposition 1, the error e_{III} can then be bounded by

$$e_{\text{III}} \leq (S^{\frac{N}{N+1}} C_N + 8H^2) \exp(-(C_e S)^{\frac{1}{N+1}}), \quad (26)$$

where the constants C_e and C_N are given by

$$\begin{aligned} C_e &= \left(\frac{\tau 2^{N-1} N! \prod_{i=1}^N \lambda_i}{3^N (1+\mu)} \right) \\ C_N &= ((1+\beta(S)) C_{\varphi, \delta} C_u(\mu) \frac{\tau^{\frac{N}{N+1}}}{3 * 2^{2N}} \left(\frac{(1+\mu)}{N! \prod_{i=1}^N \lambda_i} \right)^{\frac{1}{N+1}} \end{aligned} \quad (27)$$

with $\tau = \frac{1-\ln 2}{2}$ and the constants $\mu, C_{\varphi, \delta}, C_u$ are defined in Section 3.1.

The error e_{III} can be balanced by constructing an appropriate quasi-optimal subspace. To do so, we will be required to determine the weights $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_N)$ in (11). We note that it is the ratio of the weights which will determine our polynomial space, the magnitude will simply dictate the pace at which the error decays. It is only possible to analytically construct the weights in the case where the functions $\{a_n(\mathbf{x})\}_{n=0}^N$ in (5) do not have overlapping supports, e.g., the inclusion problem investigated in [3]; otherwise the weights must be determined numerically. On the other hand, the optimal weights require the solution of a nonlinear optimization problem in the N -dimensional parameter space, which is also not feasible in practice. Hence, we instead follow a procedure of one-dimensional analyses as done in [4, 23]. We consider the subset $\mathbb{U} = \{\mathbf{v} \in \mathbb{N}_0^N : v_i = 0 \text{ if } i \neq n, v_n = 0, 1, 2, \dots\}$. Then, according to the decay rates established in the previous section, $|c_{\mathbf{v}}| \sim e^{-\boldsymbol{\lambda}_n}$, so the rate $\boldsymbol{\lambda}_n$ can then be estimated through linear regression of the quantities $\ln |c_{\mathbf{v}}|$. Now recalling definition 1 for any $0 \leq \delta \leq 1$ it must hold that $\mathcal{R}(a(\mathbf{x}, \mathbf{z})) \geq \delta$ for all $\mathbf{x} \in D$ and $\mathbf{z} = \{z_i\}_{i=1}^N$ in the polyellipse \mathcal{E} which is determined by the weights $\boldsymbol{\lambda}_n$. In order to ensure this holds we scale our weights by an appropriate constant. Even though this may not result in an optimal estimate, it will still manage to capture any anisotropic behavior present in the problem.

5 Numerical experiments

In this section, we illustrate the convergence as well as the computational efficiency of the DLS-RB method. All calculations in this section are effected using the FEniCS [1] (<http://fenicsproject.org/>) and Rbnics [20] <http://mathlab.sissa.it/rbnics> software suites. We will use the same problem formulation utilized in [9]. Consider the stochastic elliptic problem (3) with $D = [0, 1]^2$, the forcing term $f = 1$, and the

finite element discretization with fixed $h = \frac{1}{256}$. We take the coefficient $a(\mathbf{x}, \mathbf{y})$ to be a random field with expectation and correlation given as

$$\mathbb{E}[a(\mathbf{x})] = c \text{ for a fixed } c > 0 \text{ and } \text{Cov}[a](\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{(\mathbf{x} - \mathbf{x}')^2}{L^2}\right), \quad (28)$$

where L is the correlation length. This field can be represented by the following Fourier-type expansion

$$a(\mathbf{x}, \mathbf{y}) = \frac{1}{100} \left\{ c + \left(\frac{\sqrt{\pi L}}{2}\right)^{\frac{1}{2}} y_1 + \sum_{n=1}^{\infty} \sqrt{\xi_n} \left(\sin(n\pi x_1) y_{2n} + \cos(n\pi x_1) y_{2n+1} \right) \right\}, \quad (29)$$

where the uncorrelated random variables y_n have zero mean and unit variance, and the eigenvalues are equal to

$$\sqrt{\xi_n} = (\sqrt{\pi L})^{\frac{1}{2}} \exp\left(-\frac{(n\pi L)^2}{8}\right) \text{ for } n \geq 1. \quad (30)$$

Here, we take $c = 4$, $L = \frac{1}{8}$ and only retain the first 5 random variables $\mathbf{y} = (y_1, \dots, y_5)$ in the expansion (29). Even though the independence of the five random variables is only valid in the case of Gaussian distribution, we assume (y_1, \dots, y_5) are independent uniformly distributed random variables in $\Gamma = [-1, 1]^5$. The weights for the quasi-optimal subspace are found to be $\boldsymbol{\lambda} \approx (0.68, 0.66, 0.98, 1.37, 0.49)$ after rescaling. In order to measure the error in our examples we will consider the quantity of interest

$$Q(u) = \frac{1}{|D|} \int_D u d\mathbf{x}, \quad (31)$$

and examine the behavior of our algorithm in the norm

$$\mathbb{E}(\|Q(u_h) - Q(P_{M,S}[u_{h,k}](\mathbf{y}))\|_{\infty}) \approx \max_{\mathbf{y} \in \Xi_{\text{test}}} |Q(u_h(\mathbf{y})) - Q(P_{M,S}[u_{h,k}](\mathbf{y}))| \quad (32)$$

where Ξ_{test} is 10,000 uniformly distributed points and u_h is some reference finite element solution. In order to generate a reduced basis in our examples we will use a training set Ξ_{train} of 1,000 uniformly distributed points.

We note for this particular quantity of interest it can be shown using an Aubin-Nitsche duality argument from [10, 11] that the convergence rates will be twice as great as those in estimate (25).

5.1 Example 1

For our first example, we examine the convergence and stability of the DLS method independent of any reduced basis. Specifically, we are interested in illustrating that a linear rule maintains sufficient stability for the least-squares problem (in a moderately sized dimension) utilizing a quasi-optimal polynomial space. This is motivated by computational necessity in larger dimensions for which the cardinality of the quasi-optimal polynomial space can grow very quickly. We choose the number of sample

points $S = M$, $3M$, and M^2 . As can be seen in Figure 5.1, we obtain similar levels of accuracy using the linear rule as we do when using a quadratic rule in agreement with the numerical findings in [13, 22]. We also see that taking $S = M$ leads to an unstable approximation indicating that some scaling constant is required.

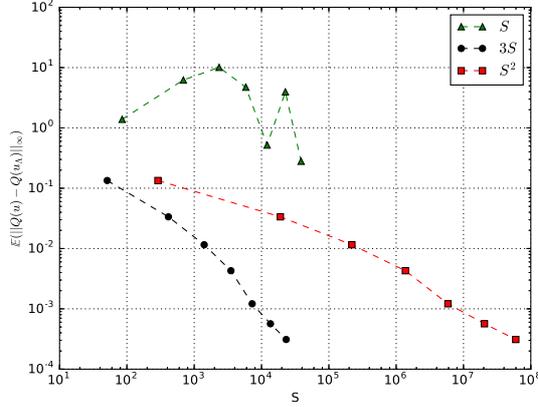


Fig. 1 Comparing the convergence of the DLS method utilizing different numbers of sample points with respect to the cardinality of the polynomial basis.

5.2 Example 2

Next, we are interested in the offline and online computational cost of the DLS method compared to that of the RB-DLS method. Beginning with the offline complexity of the DLS method we see, that the majority of the cost is incurred from setting up the right-hand side \mathbf{U} and then solving the system (14). To solve (14) we can use any number of methods, two of the more popular being the LU factorization or QR factorization, both of which have the same order of computational complexity. It then follows that the complexity for the DLS algorithm will scale as

$$\text{DLS}_{\text{cost}} = S \times \mathcal{O}(J^\alpha) + \mathcal{O}(M^3) + \mathcal{O}(M^2) \times J, \quad (33)$$

where $\mathcal{O}(J^\alpha)$ is the cost for solving the finite element system where α depends on both the solver and spatial dimension, $\mathcal{O}(M^3)$ is the cost associated with the LU or QR decomposition, and $\mathcal{O}(M^2) \times J$ is the cost for solving the system (14).

Next, we analyze the algorithm with the reduced basis incorporated into it. The construction of the reduced basis scales as

$$\text{RB}_{\text{cost}} = \mathcal{O}(S_{\text{train}}) \times \left(\sum_{\ell=1}^{K-1} w_{\text{online}}(\ell) \right) + K \times \mathcal{O}(J^\alpha), \quad (34)$$

where $\mathcal{O}(S_{\text{train}})$ is the cost of a max search in our training set, and $w_{\text{online}}(\ell) = \mathcal{O}(\ell^3)$ is the cost for calculating $\hat{e}_{h,K}$ and $u_{h,k}(\mathbf{y})$ for a value $\mathbf{y} \in \Xi_{\text{train}}$. The total cost for our algorithm, assuming no online enrichment of the reduced basis is necessary, will thus scale as

$$\text{RB-DLS}_{\text{cost}} = \text{RB}_{\text{cost}} + S \times \mathcal{O}(K^3) + \mathcal{O}(M^3) + \mathcal{O}(M^2) \times K + S \times N^2 \times K^2, \quad (35)$$

where $S \times \mathcal{O}(K^3)$ is the cost for solving the reduced basis system for $\{\mathbf{y}_i\}_{i=1}^S$ to form \mathbf{W} in (22), $\mathcal{O}(M^3)$ is the cost associated with the LU or QR decomposition, $\mathcal{O}(M^2) \times K$ is the cost for solving the system (22) and $S \times N^2 \times K^2$ is the cost for evaluating the error bound $\hat{e}_{h,K}$.

We see that the complexity of RB-DLS is dominated by the term $\mathcal{O}(M^2) \times K$ when K is large. On the other hand, when K is small the complexity of both algorithms is dominated by the term $\mathcal{O}(M^3)$. The key to the computational savings witnessed in the reduced-basis method is that the cost of the reduced-basis algorithm will be independent of J except in the offline portion. As seen in the above discussion, for large values of J the computational cost of the algorithm is dominated by the cost of finite element solves; therefore, we measure the offline computational cost in terms of the total number of full finite element solves necessary for the construction of the DLS and RB-DLS approximations.

Turning to the online computational cost, as described in Section 4, the cost of evaluating DLS for a given $\mathbf{y} \in \Gamma$ is of $\mathcal{O}(MJ)$ versus $\mathcal{O}(k(M+J))$ for RB-DLS. To illustrate the significant cost savings of RB-DLS, we compare the total CPU time (in seconds) it takes to compute all RB-DLS and DLS approximations for all $\mathbf{y} \in \Xi_{\text{train}}$. As shown in Figure 2, we observe significant offline and online computational cost savings while still being able to achieve similar levels of accuracy from the RB-DLS method.

6 Conclusions

We integrated a reduced-basis method into the discrete least-squares framework utilizing a new quasi-optimal polynomial space. Through our numerical results, we demonstrated significant cost savings in both the offline and online portions of the discrete least-squares-reduce basis method compared to that for the original discrete least-squares algorithm. Again, we would like to emphasize that reduced basis plays a critical role in solving large-scale UQ problems involving expensive finite element discretization (e.g., with a very fine mesh), especially in the online phase. We note that this method is not without drawbacks. For the case where the Kolmogorov width of the PDE solution does not decay quickly we will need to use a large number of reduced basis functions in order to obtain an accurate reduced basis approximation. This could potentially make the DLS-RB method more expensive than the standalone DLS method. Additionally the quasi-optimal polynomial basis used in this work only applies to the parametrized diffusion equation. Thus for more complicated PDEs a different polynomial basis would have to be used. This new basis may have significantly worse stability and convergence properties when coupled with the DLS method possibly rendering it ineffective.

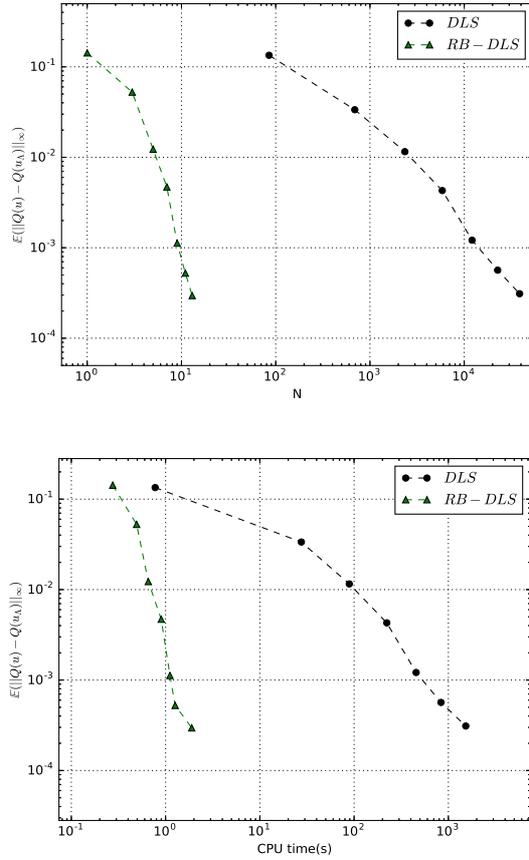


Fig. 2 A comparison of the offline (top) and online (bottom) error versus the cost for the DLS method versus the RB-DLS method measured in terms of the number of finite element solves in the offline phase (top) and CPU time in the online phase (bottom).

We note that while we paired the reduced basis method with the discrete least-squares algorithm, it is also possible to combine it with sparse-grid method as done in [9–11]. A detailed comparison of these approaches has not been done and will be a subject of future research.

As shown in this work, a polynomial approximation of the solution map $(\mathbf{x}, \mathbf{y}) \rightarrow u(\mathbf{x}, \mathbf{y})$ without using a reduced basis may lead to an unaffordable online cost in terms of storage requirement. Thus, model reduction should become a standard procedure in approximating/recovering the solution map $(\mathbf{x}, \mathbf{y}) \rightarrow u(\mathbf{x}, \mathbf{y})$.

References

1. Alnäs, M.S.: UFL: a Finite Element Form Language, chap. 17. Springer (2012)

2. Bachmayr, M., Cohen, A.: Kolmogorov widths and low-rank approximations of parametric elliptic pdes. arXiv preprint arXiv:1502.03117 (2015)
3. Beck, J., Nobile, F., Tamellini, L., Tempone, R.: Convergence of quasi-optimal stochastic galerkin methods for a class of pdes with random coefficients. *Computers & Mathematics with Applications* **67**(4), 732–751 (2014)
4. Beck, J., Tempone, R., Nobile, F., Tamellini, L.: On the optimal polynomial approximation of stochastic pdes by galerkin and collocation methods. *Mathematical Models and Methods in Applied Sciences* **22**(09), 1250,023 (2012)
5. Besselink, B., Tabak, U., Lutowska, A., van de Wouw, N., Nijmeijer, H., Rixen, D., Hochstenbach, M., Schilders, W.: A comparison of model reduction techniques from structural dynamics, numerical mathematics and systems and control. *Journal of Sound and Vibration* **332**(19), 4403 – 4422 (2013). DOI <http://doi.org/10.1016/j.jsv.2013.03.025>. URL <http://www.sciencedirect.com/science/article/pii/S0022460X1300285X>
6. Binev, P., Cohen, A., Dahmen, W., DeVore, R., Petrova, G., Wojtaszczyk, P.: Convergence rates for greedy algorithms in reduced basis methods. *SIAM Journal on Mathematical Analysis* **43**(3), 1457–1472 (2011)
7. Boyaval, S., Le Bris, C., Lelièvre, T., Maday, Y., Nguyen, N.C., Patera, A.T.: Reduced basis techniques for stochastic problems. *Archives of Computational methods in Engineering* **17**(4), 435–454 (2010)
8. Brenner, S.C., Scott, R.: *The mathematical theory of finite element methods*, vol. 15. Springer Science & Business Media (2008)
9. Chen, P., Quarteroni, A., Rozza, G.: Comparison between reduced basis and stochastic collocation methods for elliptic problems. *Journal of Scientific Computing* **59**(1), 187–216 (2014)
10. Chen, P., Schwab, C.: Sparse-grid, reduced-basis bayesian inversion. *Computer Methods in Applied Mechanics and Engineering* **297**(Supplement C), 84 – 115 (2015). DOI <https://doi.org/10.1016/j.cma.2015.08.006>. URL <http://www.sciencedirect.com/science/article/pii/S0045782515002601>
11. Chen, P., Schwab, C.: Sparse-grid, reduced-basis bayesian inversion: Nonaffine-parametric nonlinear equations. *Journal of Computational Physics* **316**(Supplement C), 470 – 503 (2016). DOI <https://doi.org/10.1016/j.jcp.2016.02.055>. URL <http://www.sciencedirect.com/science/article/pii/S0021999116001273>
12. Chkifa, A., Cohen, A., Christoph, S.: High-dimensional adaptive sparse polynomial interpolation and applications to parametric pdes. *Foundations of Computational Mathematics* **14**(4) (2014)
13. Chkifa, A., Cohen, A., Migliorati, G., Nobile, F., Tempone, R.: Discrete least squares polynomial approximation with random evaluations- application to parametric and stochastic elliptic pdes. *ESAIM: Mathematical Modelling and Numerical Analysis* (2015)
14. Chkifa, A., Cohen, A., Schwab, C.: Breaking the curse of dimensionality in sparse polynomial approximation of parametric pdes. *Journal de Mathématiques Pures et Appliquées* **103**(2), 400–428 (2015)
15. Cohen, A., Davenport, M.A., Leviatan, D.: On the stability and accuracy of least squares approximations. *Foundations of computational mathematics* **13**(5), 819–834 (2013)
16. Cohen, A., DeVore, R.: Kolmogorov widths under holomorphic mappings. *IMA Journal of Numerical Analysis* **36**(1), 1 (2016). DOI 10.1093/imanum/dru066. URL [+http://dx.doi.org/10.1093/imanum/dru066](http://dx.doi.org/10.1093/imanum/dru066)
17. Cohen, A., DeVore, R., Schwab, C.: Convergence rates of best n-term galerkin approximations for a class of elliptic pdes. *Foundations of Computational Mathematics* **10**(6), 615–646 (2010)
18. Cohen, A., DeVore, R., Schwab, C.: Analytic regularity and polynomial approximation of parametric and stochastic elliptic pde’s. *Analysis and Applications* **9**(01), 11–47 (2011)
19. Gunzburger, M.D., Webster, C.G., Zhang, G.: Stochastic finite element methods for partial differential equations with random input data. *Acta Numerica* **23**, 521–650 (2014). DOI 10.1017/S0962492914000075
20. Hesthaven, J.S., Rozza, G., Stamm, B.: *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*. SpringerBriefs in Mathematics. Springer International Publishing (2015)
21. 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: Mathematical Modelling and Numerical Analysis* **48**(01), 259–283 (2014)
22. Migliorati, G., Nobile, F., von Schwerin, E., Tempone, R.: Approximation of quantities of interest in stochastic pdes by the random discrete l^2 projection on polynomial spaces. *SIAM Journal on Scientific Computing* **35**(3), A1440–A1460 (2013)

23. Nobile, F., Tempone, R., Webster, C.G.: An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM Journal on Numerical Analysis* **46**(5), 2411–2442 (2008)
24. Nobile, F., Tempone, R., Webster, C.G.: A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM Journal on Numerical Analysis* **46**(5), 2309–2345 (2008)
25. Quarteroni, A., Rozza, G., Manzoni, A.: Certified reduced basis approximation for parametrized partial differential equations and applications. *Journal of Mathematics in Industry* **1**(1), 1–49 (2011)
26. Tran, H., Webster, C.G., Zhang, G.: Analysis of quasi-optimal polynomial approximations for parameterized pdes with deterministic and stochastic coefficients. *Numerische Mathematik* pp. 1–43 (2017). DOI 10.1007/s00211-017-0878-6. URL <http://dx.doi.org/10.1007/s00211-017-0878-6>