

Randomized residual-based error estimators for the Proper Generalized Decomposition approximation of parametrized problems

Kathrin Smetana*, Olivier Zahm†

October 28, 2019

Abstract

This paper introduces a novel error estimator for the Proper Generalized Decomposition (PGD) approximation of parametrized equations. The estimator is intrinsically random: It builds on concentration inequalities of Gaussian maps and an adjoint problem with random right-hand side, which we approximate using the PGD. The effectivity of this randomized error estimator can be arbitrarily close to unity with high probability, allowing the estimation of the error with respect to any user-defined norm as well as the error in some quantity of interest. The performance of the error estimator is demonstrated and compared with some existing error estimators for the PGD for a parametrized time-harmonic elastodynamics problem and the parametrized equations of linear elasticity with a high-dimensional parameter space.

Keywords: A posteriori error estimation, parametrized equations, Proper Generalized Decomposition, Monte-Carlo estimator, concentration phenomenon, goal-oriented error estimation

1 Introduction

Parametrized models are ubiquitous in engineering applications, life sciences, hazard event prediction and many other areas. Depending on the application, the parameters may, for instance, represent the permeability of soil, the Young’s modulus field, the geometry but also uncertainty in the data. In general, the goal is either to analyze the model response with respect to the parameter or to perform real-time simulations for embedded systems. In either case, the complexity of the model is the computational bottleneck, especially if computing the model response for one parameter value requires the solution to a complex system of partial differential equations (PDEs). Model order reduction methods such as the Reduced Basis (RB) method [26, 45, 24, 46, 49], the Proper Generalized Decomposition (PGD) [3, 4, 40], the hierarchical model reduction approach [51, 43, 42, 47, 8], or, more generally, tensor-based methods [25, 41] are commonly used to build a fast-to-evaluate approximation to the model response.

This paper is concerned with the PGD, which has been developed for the approximation of models defined in multidimensional spaces, such as problems in quantum chemistry, financial mathematics or (high-dimensional) parametrized models. For the variational version of the PGD, where an energy functional is minimized, convergence has been proved for the Poisson problem in [37]. For a review on the PGD method see for instance [14, 17, 16] and we refer to [15] for an overview of how to use PGD in different applications such as augmented learning and process optimization.

The goal of this paper is to introduce a novel error estimator for the PGD approximation. Error estimation is of crucial importance, in particular to stop the PGD enrichment process as soon as the desired precision is reached. To that end, we extend the randomized residual-based error estimator introduced in [48] to the PGD. The randomized error estimator features several desirable properties: first, it does not require the estimation of stability constants such as a coercivity or inf-sup constant and the effectivity index is close to unity with prescribed lower and upper bounds at high probability. The effectivity index can thus be selected by the user, balancing computational costs and desired sharpness of the estimator. Moreover, the presented framework yields error estimators with respect to user-defined norms, for instance the L^2 -norm or the H^1 -norm; the approach also permits error estimation of linear quantities of interest (QoI). Finally, depending on the desired effectivity the computation of the error estimator is

*Faculty of Electrical Engineering, Mathematics & Computer Science, University of Twente, Zilverling, P.O. Box 217, 7500 AE Enschede, The Netherlands, k.smetana@utwente.nl

†Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LJK, 38000 Grenoble, France, olivier.zahm@inria.fr

in general only as costly as the computation of the PGD approximation or even significantly less expensive, which makes our error estimator strategy attractive from a computational viewpoint. To build this estimator, we first estimate the norm of the error with a Monte Carlo estimator using Gaussian random vectors whose covariance is chosen according to the desired error measure, e.g., user-defined norms or quantity of interest. Then, we introduce a dual problem with random right-hand side the solution of which allows us to rewrite the error estimator in terms of the residual of the original equation. Approximating the random dual problems with the PGD yields a fast-to-evaluate error estimator that has low marginal cost.

Next, we relate our error estimator to other stopping criteria that have been proposed for the PGD solution. In [2] the authors suggest letting an error estimator in some linear QoI steer the enrichment of the PGD approximation. In order to estimate the error they employ an adjoint problem with the linear QoI as a right-hand side and approximate this adjoint problem with the PGD as well. It is important to note that, as observed in [2], this type of error estimator seems to require in general a more accurate solution of the dual than the primal problem. In contrast, the framework we present in this paper often allows using a dual PGD approximation with significantly less terms than the primal approximation. An adaptive strategy for the PGD approximation of the dual problem for the error estimator proposed in [2] was suggested in [23] in the context of computational homogenization. In [1] the approach proposed in [2] is extended to problems in nonlinear solids by considering a suitable linearization of the problem.

An error estimator for the error between the exact solution of a parabolic PDE and the PGD approximation in a suitable norm is derived in [34] based on the constitutive relation error and the Prager-Synge theorem. The discretization error is thus also assessed, and the computational costs of the error estimator depend on the number of elements in the underlying spatial mesh. Relying on the same techniques an error estimator for the error between the exact solution of a parabolic PDE and a PGD approximation in a linear QoI was proposed in [35], where the authors used the PGD to approximate the adjoint problem. In [38] the hypercircle theory introduced by Prager and Synge is invoked to derive an error between the exact solution of the equations of linear elasticity and the PGD approximation. While the estimators in [34, 35, 38] are used to certify the PGD approximation, in [12] the authors slightly modify the techniques used in [34, 35] to present an error estimator that is used to drive the construction of the PGD approximation within an adaptive scheme. To that end, the authors assume that also the equilibrated flux can be written in a variable separation form. In contrast to the error estimator we introduce in this paper, the estimators presented in [34, 35, 38, 12] also take into account the discretization error; for the a posteriori error estimator suggested in this paper this is the subject of a forthcoming work. Moreover, in [34, 35, 38, 12] the authors consider (parametrized) coercive elliptic and parabolic problems; it seems that the techniques just described have not been used yet to address problems that are merely inf-sup stable.

Error estimators based on local error indicators are proposed to drive mesh adaptation within the construction of the PGD approximation in [39]. In [7], an ideal minimal residual formulation is proposed in order to build optimal PGD approximations. In that work, an error estimator with controlled effectivity is derived from the solution of some adjoint problem. Finally, a PGD approximation of the adjoint problem is used to estimate a QoI in [33]. The latter is then employed to constrain the primal approximation, aiming at a (primal) PGD approximation more tailored to the considered QoI [33].

The randomized a posteriori error estimator proposed in [48], which we extend in this paper to the PGD, is inspired by the probabilistic error estimator for the approximation error in the solution of a system of ordinary differential equations introduced in [11, 27]. Here, the authors invoke the small statistical sample method from [32], which estimates the norm of a vector by its inner product with a random vector drawn uniformly at random on the unit sphere.

We note that randomized methods for error estimation are gaining interest in the reduced order modeling community in particular and in the context of parametrized PDEs generally. For instance in [5], random sketching techniques are used within the context of the RB method to speed-up the computation of the dual norm of the residual used as an error indicator. Here, the authors exploit that the residual lies in a low-dimensional subspace, allowing them to invoke arguments based on ε -nets in order to certify the randomized algorithm. In [30] for the RB method a probabilistic a posteriori error bound for linear scalar-valued QoIs is proposed, where randomization is done by assuming that the parameter is a random variable on the parameter set. In [53], an interpolation of the operator inverse is built via a Frobenius-norm projection and computed efficiently using randomized methods. An error estimator is obtained by measuring the norm of residual multiplied by the interpolation of the operator inverse, used here as a preconditioner. Next, in the context of localized model order reduction a reliable and efficient probabilistic a posteriori error estimator for the operator norm of the difference between a finite-dimensional linear operator and its orthogonal projection onto a reduced space is derived in [9]. Finally in the recent paper [21] the authors use concentration inequalities to bound the generalization error in the context of statistical learning for

high-dimensional parametrized PDEs in the context of Uncertainty Quantification.

The remainder of this article is organized as follows. After introducing the problem setting and recalling the PGD in Section 2, we introduce an unbiased randomized a posteriori error estimator in Section 3. As this error estimator still depends on the high-dimensional solutions of dual problems, we suggest in Section 4 how to construct dual PGD approximations for this setting efficiently and propose an algorithm where the primal and dual PGD approximations are constructed in an intertwined manner. In Section 5 we compare the effectivity of the randomized a posteriori error estimator with a stagnation-based error estimator and the dual norm of the residual for a time-harmonic elastodynamics problem and a linear elasticity problem with 20-dimensional parameter space. We close with some concluding remarks in Section 6.

2 Problem setting and the Proper Generalized Decomposition

We denote by $\boldsymbol{\mu} = (\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_p)$ a vector of p parameters which belongs to a set $\mathcal{P} \subset \mathbb{R}^p$ of admissible parameters. Here, we assume that \mathcal{P} admits a product structure $\mathcal{P} = \mathcal{P}_1 \times \dots \times \mathcal{P}_p$, where $\mathcal{P}_i \subset \mathbb{R}$ is the set of admissible parameters for $\boldsymbol{\mu}_i$, $i = 1, \dots, p$. For simplicity we assume that $\mathcal{P}_i = \{\boldsymbol{\mu}_i^{(1)}, \dots, \boldsymbol{\mu}_i^{(\#\mathcal{P}_i)}\}$ is a finite set of parameters (obtained for instance as an interpolation grid or by a Monte Carlo sample of some probability measure) so that \mathcal{P} is a grid containing $\#\mathcal{P} = \prod_{i=1}^p \#\mathcal{P}_i$ points.

Let $D \subset \mathbb{R}^d$ be a bounded spatial domain and let $\mathcal{A}(\boldsymbol{\mu}) : \mathcal{X} \rightarrow \mathcal{X}'$ be a parametrized linear differential operator, where \mathcal{X} is a suitable Sobolev space of functions defined on D ; \mathcal{X}' denotes the dual space of \mathcal{X} . We consider the following problem: for any $\boldsymbol{\mu} \in \mathcal{P}$ and $f(\boldsymbol{\mu}) \in \mathcal{X}'$, find $u(\boldsymbol{\mu}) \in \mathcal{X}$ such that

$$\mathcal{A}(\boldsymbol{\mu})u(\boldsymbol{\mu}) = f(\boldsymbol{\mu}) \quad \text{in } \mathcal{X}'. \quad (1)$$

We ensure well-posedness of problem (1) by requiring that

$$\beta(\boldsymbol{\mu}) := \inf_{v \in \mathcal{X}} \sup_{w \in \mathcal{X}} \frac{\langle \mathcal{A}(\boldsymbol{\mu})v, w \rangle_{\mathcal{X}', \mathcal{X}}}{\|v\|_{\mathcal{X}} \|w\|_{\mathcal{X}}} > 0 \quad \text{and} \quad \gamma(\boldsymbol{\mu}) := \sup_{v \in \mathcal{X}} \sup_{w \in \mathcal{X}} \frac{\langle \mathcal{A}(\boldsymbol{\mu})v, w \rangle_{\mathcal{X}', \mathcal{X}}}{\|v\|_{\mathcal{X}} \|w\|_{\mathcal{X}}} < \infty,$$

hold for all $\boldsymbol{\mu} \in \mathcal{P}$. Here $\langle \cdot, \cdot \rangle_{\mathcal{X}', \mathcal{X}}$ denotes the duality pairing of \mathcal{X}' and \mathcal{X} .

Next, we introduce a finite dimensional space $\mathcal{X}_N := \text{span}\{\psi_1, \dots, \psi_N\} \subset \mathcal{X}$ defined as the span of $N \gg 1$ basis functions ψ_1, \dots, ψ_N . For any $\boldsymbol{\mu} \in \mathcal{P}$, we define the high fidelity approximation of (1) as the function $u_N(\boldsymbol{\mu}) \in \mathcal{X}_N$ that solves the variational problem

$$\langle \mathcal{A}(\boldsymbol{\mu})u_N(\boldsymbol{\mu}), w_N \rangle_{\mathcal{X}', \mathcal{X}} = \langle f(\boldsymbol{\mu}), w_N \rangle_{\mathcal{X}', \mathcal{X}} \quad \text{for all } w_N \in \mathcal{X}_N. \quad (2)$$

Again, we ensure wellposed-ness of problem (2) by requiring the discrete inf-sup/sup-sup conditions

$$\beta_N(\boldsymbol{\mu}) := \inf_{v \in \mathcal{X}_N} \sup_{w \in \mathcal{X}_N} \frac{\langle \mathcal{A}(\boldsymbol{\mu})v, w \rangle_{\mathcal{X}', \mathcal{X}}}{\|v\|_{\mathcal{X}} \|w\|_{\mathcal{X}}} > 0 \quad \text{and} \quad \gamma_N(\boldsymbol{\mu}) := \sup_{v \in \mathcal{X}_N} \sup_{w \in \mathcal{X}_N} \frac{\langle \mathcal{A}(\boldsymbol{\mu})v, w \rangle_{\mathcal{X}', \mathcal{X}}}{\|v\|_{\mathcal{X}} \|w\|_{\mathcal{X}}} < \infty, \quad (3)$$

to hold for all $\boldsymbol{\mu} \in \mathcal{P}$. Expressing $u_N(\boldsymbol{\mu})$ in the basis $u_N(\boldsymbol{\mu}) = \sum_{i=1}^N u_i(\boldsymbol{\mu})\psi_i$ and defining the vector of coefficient $u(\boldsymbol{\mu}) := (u_1(\boldsymbol{\mu}), \dots, u_N(\boldsymbol{\mu})) \in \mathbb{R}^N$, we rewrite (2) as a parametrized algebraic system of linear equations

$$A(\boldsymbol{\mu})u(\boldsymbol{\mu}) = f(\boldsymbol{\mu}), \quad (4)$$

where the matrix $A(\boldsymbol{\mu}) \in \mathbb{R}^{N \times N}$ and the vector $f(\boldsymbol{\mu}) \in \mathbb{R}^N$ are defined as $A_{ij}(\boldsymbol{\mu}) = \langle \mathcal{A}(\boldsymbol{\mu})\psi_j, \psi_i \rangle_{\mathcal{X}', \mathcal{X}}$ and $f_i(\boldsymbol{\mu}) = \langle f(\boldsymbol{\mu}), \psi_i \rangle_{\mathcal{X}', \mathcal{X}}$ for any $1 \leq i, j \leq N$. Finally, with a slight abuse of notation, we denote by $\|\cdot\|_{\mathcal{X}_N}$ and $\|\cdot\|_{\mathcal{X}'_N}$ the norms defined on the algebraic space \mathbb{R}^N such that

$$\|v\|_{\mathcal{X}_N} := \left\| \sum_{i=1}^N v_i \psi_i \right\|_{\mathcal{X}}, \quad \text{and} \quad \|v\|_{\mathcal{X}'_N} := \sup_{w \in \mathcal{X}_N} \frac{\langle \sum_{i=1}^N v_i \psi_i, w \rangle_{\mathcal{X}', \mathcal{X}}}{\|w\|_{\mathcal{X}}} \quad \text{for any } v \in \mathbb{R}^N. \quad (5)$$

These norms can be equivalently defined by $\|v\|_{\mathcal{X}_N}^2 = v^T R_X v$ and $\|v\|_{\mathcal{X}'_N}^2 = v^T R_X^{-1} v$ for any $v \in \mathbb{R}^N$, where $R_X \in \mathbb{R}^{N \times N}$ is the gram matrix of the basis (ψ_1, \dots, ψ_N) such that $(R_X)_{ij} = (\psi_i, \psi_j)_{\mathcal{X}}$ where $(\cdot, \cdot)_{\mathcal{X}}$ is the scalar product in \mathcal{X} .

2.1 The Proper Generalized Decomposition

To avoid the curse of dimensionality we compute an approximation of $u(\boldsymbol{\mu}) = u(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_p) \in \mathbb{R}^N$ using the Proper Generalized Decomposition (PGD). The PGD relies on the separation of the variables $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_p$ to build an approximation of the form

$$\tilde{u}^M(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_p) = \sum_{m=1}^M u_{\mathcal{X}}^m u_1^m(\boldsymbol{\mu}_1) \cdots u_p^m(\boldsymbol{\mu}_p), \quad (6)$$

where the M vectors $u_{\mathcal{X}}^m \in \mathbb{R}^N$ and the $p \times M$ scalar univariate functions $u_i^m : \mathcal{D}_i \rightarrow \mathbb{R}$ need to be constructed. An approximation as in (6) can be interpreted as a low-rank tensor approximation in the canonical format, see [41]. With this interpretation, M is called the canonical rank of \tilde{u}^M seen as a tensor $\tilde{u}^M = \sum_{m=1}^M u_{\mathcal{X}}^m \otimes u_1^m \otimes \dots \otimes u_p^m$. Let us note that if the space of spatial functions \mathcal{X} was a tensor product space $\mathcal{X} = \mathcal{X}^1 \otimes \dots \otimes \mathcal{X}^d$, then one could have further separated the spatial variables and obtained approximations like $\tilde{u}^M = \sum_{m=1}^M u_{\mathcal{X}^1}^m \otimes \dots \otimes u_{\mathcal{X}^d}^m \otimes u_1^m \otimes \dots \otimes u_p^m$. This idea is however not pursued in this paper. The PGD is fundamentally an ℓ^2 -based model order reduction technique in the sense that it aims at constructing \tilde{u}^M such that the root mean squared (RMS) error

$$\|u - \tilde{u}^M\| := \sqrt{\frac{1}{\#\mathcal{D}} \sum_{\boldsymbol{\mu} \in \mathcal{D}} \|u(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})\|_{\mathcal{X}_N}^2}, \quad (7)$$

is controlled. However, minimizing this error over the functions \tilde{u}^M as in (6) is not feasible in practice. We now recall the basic ingredients of the PGD for the computation of $\tilde{u}^M(\boldsymbol{\mu})$, referring for instance to [14, 17, 16] for further details.

Progressive PGD

The progressive PGD relies on a greedy procedure where we seek the M terms in (6) one after another. After $M - 1$ iterations, we build a rank-one correction $u_{\mathcal{X}}^M u_1^M(\boldsymbol{\mu}_1) \cdots u_p^M(\boldsymbol{\mu}_p)$ of the current approximation $\tilde{u}^{M-1}(\boldsymbol{\mu})$ by minimizing some error functional described below. Once this correction is computed, we update the approximation

$$\tilde{u}^M(\boldsymbol{\mu}) = \tilde{u}^{M-1}(\boldsymbol{\mu}) + u_{\mathcal{X}}^M u_1^M(\boldsymbol{\mu}_1) \cdots u_p^M(\boldsymbol{\mu}_p), \quad (8)$$

and we increment $M \leftarrow M + 1$. Such a procedure is called *pure greedy* because it never updates the previous iterates.

We now describe how to construct the rank-one correction. A commonly used error functional is the mean squared residual norm, which yields the so-called *Minimal Residual PGD* [7, 22, 10]. The correction $u_{\mathcal{X}}^M u_1^M(\boldsymbol{\mu}_1) \cdots u_p^M(\boldsymbol{\mu}_p)$ is defined as a solution to

$$\min_{u_{\mathcal{X}} \in \mathbb{R}^N} \min_{u_1 : \mathcal{D}_1 \rightarrow \mathbb{R}} \cdots \min_{u_p : \mathcal{D}_p \rightarrow \mathbb{R}} \sum_{\boldsymbol{\mu} \in \mathcal{D}} \|A(\boldsymbol{\mu}) \left(\tilde{u}^{M-1}(\boldsymbol{\mu}) + u_{\mathcal{X}} u_1(\boldsymbol{\mu}_1) \cdots u_p(\boldsymbol{\mu}_p) \right) - f(\boldsymbol{\mu})\|_2^2, \quad (9)$$

where $\|\cdot\|_2$ denotes the canonical norm of \mathbb{R}^N . When the matrix $A(\boldsymbol{\mu})$ is symmetric definite positive (SPD), we can alternatively use the *Galerkin PGD* which consists in minimizing the mean squared energy

$$\min_{u_{\mathcal{X}} \in \mathbb{R}^N} \min_{u_1 : \mathcal{D}_1 \rightarrow \mathbb{R}} \cdots \min_{u_p : \mathcal{D}_p \rightarrow \mathbb{R}} \sum_{\boldsymbol{\mu} \in \mathcal{D}} \|A(\boldsymbol{\mu}) \left(\tilde{u}^{M-1}(\boldsymbol{\mu}) + u_{\mathcal{X}} u_1(\boldsymbol{\mu}_1) \cdots u_p(\boldsymbol{\mu}_p) \right) - f(\boldsymbol{\mu})\|_{A(\boldsymbol{\mu})^{-1}}^2, \quad (10)$$

where the residual is measured with the norm $\|\cdot\|_{A(\boldsymbol{\mu})^{-1}} = \sqrt{(\cdot)^T A(\boldsymbol{\mu})^{-1} (\cdot)}$. The Galerkin PGD is usually much more efficient compared to the Minimal Residual PGD, but the latter is much more stable for non SPD operators. Both for the Minimal Residual PGD (9) and the Galerkin PGD (10) Conditions (3) are sufficient to ensure the convergence of the pure greedy algorithm outlined above, see [7].

A simple and yet efficient approach to solve (9) or (10) is the *alternating least squares* (ALS) method. It relies on the fact that each of the $p + 1$ minimization problems is a least squares problem (the cost function being quadratic with respect to each of the $u_{\mathcal{X}}, u_1, u_2, \dots$) and thus is simple to solve. ALS consists in iteratively solving each of these least squares problems one after the other while keeping the other variables fixed. That is: starting with a random initialization $u_{\mathcal{X}}^{(0)}, u_1^{(0)}, \dots, u_p^{(0)}$, we first compute $u_{\mathcal{X}}^{(1)}$ solution to (9) (or to (10)) with $u_1 = u_1^{(0)}, \dots, u_p = u_p^{(0)}$ fixed. Then we compute $u_1^{(1)}$ solution to (9) (or to (10)) with $u_{\mathcal{X}} = u_{\mathcal{X}}^{(1)}, u_2 = u_2^{(0)}, \dots, u_p = u_p^{(0)}$ fixed etc. This yields a sequence of updates $u_{\mathcal{X}}^{(1)} \rightarrow u_1^{(1)} \rightarrow \dots \rightarrow u_p^{(1)} \rightarrow u_{\mathcal{X}}^{(2)} \rightarrow u_1^{(2)} \rightarrow \dots$ which we stop once the stagnation is reached. In practice, few iterations (three or four) suffice.

Stopping criterion for the greedy enrichment process

Ideally we want to stop the greedy iteration process in M as soon as the relative RMS error $\|u - \tilde{u}^M\|/\|u\|$ falls below a certain prescribed tolerance, which is why an error estimator for the relative RMS error is desirable. In *stagnation-based error estimators* the unknown solution u is replaced by a reference solution \tilde{u}^{M+k} for some $k \geq 1$

$$\Delta_{M,k}^{\text{stag}} := \frac{\|\tilde{u}^{M+k} - \tilde{u}^M\|}{\|\tilde{u}^{M+k}\|}, \quad (11)$$

where the norm $\|\cdot\|$ is defined as in (7). The estimator (11) is also often called a *hierarchical error estimator*, considering two different PGD approximations: a coarse one \tilde{u}^M and a fine one \tilde{u}^{M+k} . The quality of such an error estimator depends on the convergence rate of the PGD algorithm: a slow PGD convergence would require choosing a larger k to obtain an accurate estimate of the relative RMS error. Needless to say that the error estimator $\Delta_{M,k}^{\text{stag}}$ is more expensive to compute when $k \gg 1$. Proving that $\Delta_{M,k}^{\text{stag}}$ is an upper bound of the error and bounding the effectivity typically requires satisfying a so-called saturation assumption, i.e., having a uniform bound on the convergence rate of the PGD, a requirement which is hard to ensure in practice.

As one example of a *residual-based error estimator* we consider the relative RMS residual norm

$$\Delta_M^{\text{res}} := \sqrt{\frac{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|A(\boldsymbol{\mu})\tilde{u}^M(\boldsymbol{\mu}) - f(\boldsymbol{\mu})\|_{X'_N}^2}{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|f(\boldsymbol{\mu})\|_{X'_N}^2}}, \quad (12)$$

where the norm $\|\cdot\|_{X'_N}$ has been defined in (5). The advantage of this error estimator is that it comes with the guarantee that

$$\kappa_N^{-1} \frac{\|u - \tilde{u}^M\|}{\|u\|} \leq \Delta_M^{\text{res}} \leq \kappa_N \frac{\|u - \tilde{u}^M\|}{\|u\|}, \quad (13)$$

for any M , where $\kappa_N = \max_{\boldsymbol{\mu} \in \mathcal{P}} \gamma_N(\boldsymbol{\mu})/\beta_N(\boldsymbol{\mu}) \geq 1$ denotes the supremum of the condition number of $A(\boldsymbol{\mu})$ seen as an operator from $(\mathbb{R}^N, \|\cdot\|_{X_N})$ to $(\mathbb{R}^N, \|\cdot\|_{X'_N})$. This inequality can be easily proved using Assumption (3) and Definition (5). Inequality (13) ensures that the effectivity of Δ_M^{res} is contained in $[\kappa_N^{-1}, \kappa_N]$. Note however that estimating κ_N (or a sharp upper bound of it) is difficult in practice as it requires estimating the inf-sup/sup-sup constants $\beta_N(\boldsymbol{\mu})$ and $\gamma_N(\boldsymbol{\mu})$, a task that is general rather costly [29, 13, 28]. Moreover, for ill-conditioned problems as, say, the Helmholtz equation or a time-harmonic elastodynamics problem presented in section 5, the constant $\beta_N(\boldsymbol{\mu})$ can be arbitrarily close to zero, yielding a very large κ_N .

In [2] the authors suggest letting an estimator for the error in some linear QoI steer the enrichment of the PGD approximation. In order to estimate the error they employ an adjoint problem with the linear QoI as a right-hand side and approximate this adjoint problem with the PGD as well. In [34] the constitutive relation error and the Prager-Synge theorem is employed to derive an error estimator for the error between the exact solution of the PDE and the PGD approximation in a suitable norm. The error estimator requires the construction of equilibrated fluxes and has been proposed to use within an adaptive scheme to construct the PGD approximation in [12].

3 A randomized a posteriori error estimator

In this paper we aim at estimating the errors

$$\|u(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})\|_{\Sigma}, \quad \sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})\|_{\Sigma}^2}, \quad \text{or} \quad \sqrt{\frac{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})\|_{\Sigma}^2}{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu})\|_{\Sigma}^2}}, \quad (14)$$

where $\Sigma \in \mathbb{R}^{N \times N}$ is a symmetric positive semi-definite matrix which is chosen by the user according to the desired error measure, see subsection 3.1. This means that $\|\cdot\|_{\Sigma} = \sqrt{(\cdot)^T \Sigma (\cdot)}$ is a semi-norm unless Σ is SPD, in which case $\|\cdot\|_{\Sigma}$ is a norm. We show in subsection 3.2 that, at high probability, we can estimate accurately $\|v\|_{\Sigma}$ for any $v \in \mathbb{R}^N$ by means of *random Gaussian maps*. In subsection 3.3 we use union bound arguments to estimate the errors $\|u(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})\|_{\Sigma}$ and the norms $\|u(\boldsymbol{\mu})\|_{\Sigma}$ simultaneously for a finite number of parameters $\boldsymbol{\mu}$, which allows to estimate all quantities in (14).

3.1 Choice of the error measure

If we wish to estimate the error in the natural norm $\|\cdot\|_\Sigma = \|\cdot\|_{x_N}$ defined in (5), we simply choose $\Sigma = R_X$, where we recall that $(R_X)_{ij} = (\psi_i, \psi_j)_X$ is the gram matrix of the basis (ψ_1, \dots, ψ_N) . Alternatively, we can consider the error in the $L^2(D)$ -norm by letting $\Sigma = R_{L^2(D)}$ where the matrix $R_{L^2(D)}$ is defined by $(R_{L^2(D)})_{ij} = (\psi_i, \psi_j)_{L^2(D)}$. The presented framework also encompasses estimating the error in some vector-valued quantity of interest defined as a linear function of $u(\boldsymbol{\mu})$, say

$$s(\boldsymbol{\mu}) = Lu(\boldsymbol{\mu}) \in \mathbb{R}^m,$$

for some extractor $L \in \mathbb{R}^{m \times N}$. In this context we are interested in estimating $\|s(\boldsymbol{\mu}) - L\tilde{u}(\boldsymbol{\mu})\|_W$ where $\|\cdot\|_W$ is some norm associated with a SPD matrix $R_W \in \mathbb{R}^{m \times m}$. For example, one can be interested in computing the approximation error over a subdomain $D_{sub} \subsetneq D$, that is, estimating the error $u(\boldsymbol{\mu})|_{D_{sub}} - \tilde{u}(\boldsymbol{\mu})|_{D_{sub}}$ measured in the $L^2(D_{sub})$ -norm. In this situation, L would extract the relevant entries from the vector $\tilde{u}(\boldsymbol{\mu})$ and R_W would be defined by $(R_W)_{ij} = (\psi_i|_{D_{sub}}, \psi_j|_{D_{sub}})_{L^2(D_{sub})}$. With the choice $\Sigma = L^T R_W L$ we can write

$$\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_\Sigma^2 = (u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu}))^T (L^T R_W L) (u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})) = \|s(\boldsymbol{\mu}) - L\tilde{u}(\boldsymbol{\mu})\|_W^2,$$

so that measuring the error with respect to the norm $\|\cdot\|_\Sigma$ gives the error associated with the QoI. Notice that if $m < N$ the matrix Σ is necessarily singular and then $\|\cdot\|_\Sigma$ is a semi-norm. Finally, consider the scalar-valued QoI given by $s(\boldsymbol{\mu}) = l^T u(\boldsymbol{\mu})$ where $l \in \mathbb{R}^N$. This corresponds to the previous situation with $m = 1$ and $L = l^T$. The choice $\Sigma = ll^T$ yields $\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_\Sigma^2 = |s(\boldsymbol{\mu}) - L\tilde{u}(\boldsymbol{\mu})|^2$, where $|\cdot|$ denotes the absolute value.

3.2 Estimating norms using Gaussian maps

Random maps have proved themselves to be a very powerful tool to speed up computations in numerical linear algebra. They can also be exploited to estimate the norm of a vector very efficiently. In this subsection, we will present the required basics to first obtain an estimator for the norm of a fixed vector and then for a set of vectors. The latter is for instance of relevance if we want to compute the error (14). For a thorough introduction on how to estimate certain quantities in numerical linear algebra by random maps we refer to the recent monograph [50]; we follow subsection 2.2 of [48] here.

First, we will define a suitable random map $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}^K$ with $K \ll N$ that can be used to estimate the (semi-)norm $\|\cdot\|_\Sigma$ by $\|\Phi \cdot\|_2$, where $\|\cdot\|_2$ is the Euclidean norm in \mathbb{R}^K . To that end, let $Z \sim \mathcal{N}(0, \Sigma)$ be a zero mean Gaussian random vector in \mathbb{R}^N with covariance matrix $\Sigma \in \mathbb{R}^{N \times N}$. We emphasize that the user first selects the desired error measure which yields the (semi-)norm $\|\cdot\|_\Sigma$ and, subsequently, the covariance matrix is chosen as that very same matrix. Given a vector $v \in \mathbb{R}^N$ we can write

$$\|v\|_\Sigma^2 = v^T \Sigma v = v^T \mathbb{E}(ZZ^T) v = \mathbb{E}((Z^T v)^2),$$

where $\mathbb{E}(\cdot)$ denotes the expected value. This means that $(Z^T v)^2$ is an unbiased estimator of $\|v\|_\Sigma^2$. We may then use K independent copies of Z , which we denote by Z_1, \dots, Z_K , to define the random map $\Phi \in \mathbb{R}^{K \times N}$ as $\Phi_{(i,:)} = (1/\sqrt{K})Z_i^T$, where $(i, :)$ denotes the i -th row. Then, we can write

$$\|\Phi v\|_2^2 = \frac{1}{K} \sum_{i=1}^K (Z_i^T v)^2 \quad \text{for any } v \in \mathbb{R}^N. \quad (15)$$

We see that $\|\Phi v\|_2^2$ is a K -sample Monte-Carlo estimator of $\mathbb{E}((Z^T v)^2) = \|v\|_\Sigma^2$. By the independence of the Z_i 's, the variance of $\|\Phi v\|_2^2$ is $\text{Var}(\|\Phi v\|_2^2) = \frac{1}{K} \text{Var}((Z^T v)^2)$. Increasing K thus allows to reduce the variance of the estimate, and we get a more precise approximation of $\|v\|_\Sigma^2$. However, the variance is not always the best criterion to access the quality of the estimator.

Given a vector v and a realization of Φ , how much $\|\Phi v\|_2^2$ underestimates or overestimates $\|v\|_\Sigma^2$? Because Φ is a random variable, the answer to that question will hold only up to a certain probability which we need to quantify. To that end, we first note that the random variables $(Z_i^T v)/\|v\|_\Sigma$ for $i = 1, \dots, K$ are independent standard normal random variables so that we have

$$\|\Phi v\|_2^2 = \frac{\|v\|_\Sigma^2}{K} \sum_{i=1}^K \left(\frac{Z_i^T v}{\|v\|_\Sigma} \right)^2 \sim \frac{\|v\|_\Sigma^2}{K} Q, \quad (16)$$

| $\delta = 10^{-2}$ | $w = 2$ | $w = 4$ | $w = 10$ | $\delta = 10^{-4}$ | $w = 2$ | $w = 4$ | $w = 10$ |
|--------------------|---------|---------|----------|--------------------|---------|---------|----------|
| $\#m = 10^0$ | 24 | 6 | 3 | $\#m = 10^0$ | 48 | 11 | 6 |
| $\#m = 10^3$ | 60 | 13 | 7 | $\#m = 10^3$ | 84 | 19 | 9 |
| $\#m = 10^6$ | 96 | 21 | 11 | $\#m = 10^6$ | 120 | 26 | 13 |
| $\#m = 10^9$ | 132 | 29 | 15 | $\#m = 10^9$ | 155 | 34 | 17 |

Table 1: Minimal value of K for which Condition 18 is satisfied.

where $Q \sim \chi^2(K)$ is a random variable which follows a chi-squared distribution with K degrees of freedom. Denoting by $\mathbb{P}\{A\}$ the probability of an event A and by \overline{A} the complementary event of A , the previous relation yields

$$\mathbb{P}\left\{w^{-1}\|v\|_{\Sigma} \leq \|\Phi v\|_2 \leq w\|v\|_{\Sigma}\right\} = 1 - \mathbb{P}\left\{\overline{Kw^{-2} \leq Q \leq Kw^2}\right\},$$

for any $w \geq 1$. Then for any given (fixed) vector $v \in \mathbb{R}^N$, the probability that a realization of $\|\Phi v\|_2$ lies between $w^{-1}\|v\|_{\Sigma}$ and $w\|v\|_{\Sigma}$ is independent of v but also independent of the dimension N . The following proposition gives an upper bound for the failure probability $\mathbb{P}\{Kw^{-2} \leq Q \leq Kw^2\}$ in terms of w and K .

Proposition 3.1 (Proposition 2.1 in [48]). Let $Q \sim \chi^2(K)$ be a chi-squared random variable with $K \geq 3$ degrees of freedom. For any $w > \sqrt{e}$ we have

$$\mathbb{P}\left\{\overline{Kw^{-2} \leq Q \leq Kw^2}\right\} \leq \left(\frac{\sqrt{e}}{w}\right)^K.$$

Proof. The proof relies on the fact that we have closed form expressions for the law of $Q \sim \chi^2(K)$; for details see [48]. \square

Proposition 3.1 shows that the probability $\mathbb{P}\left\{\overline{Kw^{-2} \leq Q \leq Kw^2}\right\}$ decays at least exponentially with respect to K , provided $w \geq \sqrt{e}$ and $K \geq 3$. Then for any $v \in \mathbb{R}^N$, the relation

$$w^{-1}\|v\|_{\Sigma} \leq \|\Phi v\|_2 \leq w\|v\|_{\Sigma}, \quad (17)$$

holds with a probability greater than $1 - (\sqrt{e}/w)^K$. We highlight that this is a non-asymptotic result in the sense that it holds for finite K , not relying on the law of large numbers; for further details on those types of non-asymptotic concentration inequalities we refer to [50]. Moreover, note that by accepting a larger w and thus a larger deviation of the estimator from $\|v\|_{\Sigma}$ fewer samples are required to ensure that the probability of failure $(\sqrt{e}/w)^K$ is small. For instance with $w = 2$ and $K = 48$, relation (17) holds with a probability larger than 0.9998, while for $w = 4$ we only need $K = 11$ samples to obtain the same probability (see Table 1).

Next, by combining Proposition 3.1 with a union bound argument we can quantify the probability that relation (17) holds simultaneously for any vector in a finite set $\mathcal{M} \subset \mathbb{R}^N$.

Corollary 3.2. Given a finite collection of vectors $\mathcal{M} = \{v_1, v_2, \dots, v_{\#m}\} \subset \mathbb{R}^N$ and a failure probability $0 < \delta < 1$. Then, for any $w > \sqrt{e}$ and

$$K \geq \min \left\{ \frac{\log(\#\mathcal{M}) + \log(\delta^{-1})}{\log(w/\sqrt{e})}, 3 \right\} \quad (18)$$

we have

$$\mathbb{P}\left\{w^{-1}\|v\|_{\Sigma} \leq \|\Phi v\|_2 \leq w\|v\|_{\Sigma}, \forall v \in \mathcal{M}\right\} \geq 1 - \delta. \quad (19)$$

Table 1 gives numerical values of K that satisfy (18) depending on δ , w and $\#\mathcal{M}$. For example with $\delta = 10^{-4}$ and $w = 10$, estimating simultaneously the norm of 10^9 vectors requires only $K = 17$ samples. Again, we highlight that this result is independent on the dimension N of the vectors to be estimated. Moreover, Condition (18) allows reducing the number of samples K by increasing w^1 . Note that the computational complexity of the estimator $\|\Phi v\|_2$ and thus the a posteriori error estimator we propose in this paper depends on K (see section 4). By relying on Corollary 3.2 the user may thus choose between a less accurate but also less expensive estimator (higher w , lower K) or a more accurate but also more expensive estimator (lower w , higher K).

Remark 3.3 (Drawing Gaussian vectors). In actual practice we can draw efficiently from $Z \sim \mathcal{N}(0, \Sigma)$ using *any* factorization of the covariance matrix of the form of $\Sigma = U^T U$, for instance a Cholesky decomposition or a symmetric matrix square root. It is then sufficient to draw a standard Gaussian vector \hat{Z} and to compute the matrix-vector product $Z = U^T \hat{Z}$. As pointed-out in [5, Remark 2.9], one can take advantage of a potential block structure of Σ to build a (non-square) factorization U with a negligible computational cost.

¹Note that this is in contrast to the celebrated Johnson-Lindenstrauss lemma [18, 31]. The latter ensures that for any $0 < \varepsilon < 1$ and any finite set $\mathcal{M} \subset \mathbb{R}^N$ that contains the zero vector, the condition $K \geq 8\varepsilon^{-2} \log(\#\mathcal{M})$ ensures the existence of a linear map $\Phi: \mathbb{R}^N \rightarrow \mathbb{R}^K$ such that $\sqrt{1 - \varepsilon}\|v\|_2 \leq \|\Phi v\|_2 \leq \sqrt{1 + \varepsilon}\|v\|_2$ for all $v \in \mathcal{M}$. Increasing ε does therefore not allow reducing the number of samples K ; for a more detailed discussion see [48].

3.3 Randomized a posteriori error estimator

We apply the methodology described in the previous subsection to derive residual-based randomized a posteriori error estimators. Recall the definition of the random map $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}^K$, $\Phi_{(i,:)} = (1/\sqrt{K})Z_i^T$, $i = 1, \dots, K$, where $(i, :)$ denotes the i -th row and Z_1, \dots, Z_K are independent copies of $Z \sim \mathcal{N}(0, \Sigma)$. We may then define the estimators

$$\Delta(\boldsymbol{\mu}) = \|\Phi(u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu}))\|_2 = \sqrt{\frac{1}{K} \sum_{k=1}^K \left(Z_k^T (u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})) \right)^2} \quad \text{and} \quad \Delta = \sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \Delta(\boldsymbol{\mu})^2} \quad (20)$$

of the errors $\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_\Sigma$ and $\sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_\Sigma^2}$, respectively. As both $\Delta(\boldsymbol{\mu})$ and Δ require the high-dimensional solution $u(\boldsymbol{\mu})$ of problem (4), these estimators are computationally infeasible in practice. By introducing the residual

$$r(\boldsymbol{\mu}) = f(\boldsymbol{\mu}) - A(\boldsymbol{\mu})\tilde{u}(\boldsymbol{\mu}), \quad (21)$$

associated with Problem 4 and by exploiting the *error residual relationship* we rewrite the terms $Z_i^T(u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu}))$ in (20) as

$$Z_i^T(u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})) = Z_i^T A(\boldsymbol{\mu})^{-1} r(\boldsymbol{\mu}) = (A(\boldsymbol{\mu})^{-T} Z_i)^T r(\boldsymbol{\mu}) = Y_i(\boldsymbol{\mu})^T r(\boldsymbol{\mu}), \quad (22)$$

where $Y_i(\boldsymbol{\mu}) \in \mathbb{R}^N$ are the solutions to the *random dual problems*

$$A(\boldsymbol{\mu})^T Y_i(\boldsymbol{\mu}) = Z_i, \quad 1 \leq i \leq K. \quad (23)$$

Because the right hand side in (23) is random, the solutions $Y_1(\boldsymbol{\mu}), \dots, Y_K(\boldsymbol{\mu})$ are random vectors. Using relation (22) the error estimators $\Delta(\boldsymbol{\mu})$ and Δ (20) can be rewritten as

$$\Delta(\boldsymbol{\mu}) = \sqrt{\frac{1}{K} \sum_{k=1}^K \left(Y_k(\boldsymbol{\mu})^T r(\boldsymbol{\mu}) \right)^2} \quad \text{and} \quad \Delta = \sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \Delta(\boldsymbol{\mu})^2}. \quad (24)$$

This shows that $\Delta(\boldsymbol{\mu})$ can be computed by applying K linear forms to the residual $r(\boldsymbol{\mu})$. In that sense, $\Delta(\boldsymbol{\mu})$ can be interpreted as an a posteriori error estimator. However, computing the solutions to (23) is in general as expensive as solving the primal problem (4). In the next section, we show how to approximate the dual solutions $Y_1(\boldsymbol{\mu}), \dots, Y_K(\boldsymbol{\mu})$ via the PGD in order to obtain a fast-to-evaluate a posteriori error estimator. Before that, a direct application of Corollary 3.2 with $\mathcal{M} = \{u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu}) : \boldsymbol{\mu} \in \mathcal{P}\}$ gives the following control of the quality of $\Delta(\boldsymbol{\mu})$ and Δ .

Corollary 3.4. Given $0 < \delta < 1$ and $w > \sqrt{e}$, the condition

$$K \geq \min \left\{ \frac{\log(\#\mathcal{P}) + \log(\delta^{-1})}{\log(w/\sqrt{e})}, 3 \right\}, \quad (25)$$

is sufficient to ensure

$$\mathbb{P} \left\{ w^{-1} \Delta(\boldsymbol{\mu}) \leq \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_\Sigma \leq w \Delta(\boldsymbol{\mu}), \forall \boldsymbol{\mu} \in \mathcal{P} \right\} \geq 1 - \delta,$$

and therefore

$$\mathbb{P} \left\{ w^{-1} \Delta \leq \sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_\Sigma^2} \leq w \Delta \right\} \geq 1 - \delta.$$

It is important to note that Condition (25) depends only on the cardinality of \mathcal{P} . This means that K can be determined only knowing the number of parameters for which we need to estimate the error. Due to the logarithmic dependence of K on $\#\mathcal{P}$ Corollary 3.4 tolerates also (very) pessimistic estimates of $\#\mathcal{P}$.

We finish this section by noting that, using the same *error residual relationship* trick (20), the relative errors $\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_\Sigma / \|u(\boldsymbol{\mu})\|_\Sigma$ or $\sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_\Sigma^2} / \sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu})\|_\Sigma^2}$ can be estimated, respectively, using

$$\Delta^{\text{rel}}(\boldsymbol{\mu}) = \sqrt{\frac{\frac{1}{K} \sum_{i=1}^K \left(Y_i(\boldsymbol{\mu})^T r(\boldsymbol{\mu}) \right)^2}{\frac{1}{K} \sum_{i=1}^K \left(Y_i(\boldsymbol{\mu})^T f(\boldsymbol{\mu}) \right)^2}} \quad \text{and} \quad \Delta^{\text{rel}} = \sqrt{\frac{\frac{1}{K\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \sum_{k=1}^K \left(Y_k(\boldsymbol{\mu})^T r(\boldsymbol{\mu}) \right)^2}{\frac{1}{K\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \sum_{k=1}^K \left(Y_k(\boldsymbol{\mu})^T f(\boldsymbol{\mu}) \right)^2}}. \quad (26)$$

We then obtain a slightly modified version of Corollary 3.4; the proof can be found in the appendix.

Corollary 3.5. Given $0 < \delta < 1$ and $w > e$, the condition

$$K \geq \min \left\{ \frac{2 \log(2 \#\mathcal{P}) + 2 \log(\delta^{-1})}{\log(w/e)}, 3 \right\}, \quad (27)$$

is sufficient to ensure

$$\mathbb{P} \left\{ w^{-1} \Delta^{\text{rel}}(\boldsymbol{\mu}) \leq \frac{\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}}{\|u(\boldsymbol{\mu})\|_{\Sigma}} \leq w \Delta^{\text{rel}}(\boldsymbol{\mu}), \forall \boldsymbol{\mu} \in \mathcal{P} \right\} \geq 1 - \delta.$$

and also

$$\mathbb{P} \left\{ w^{-1} \Delta^{\text{rel}} \leq \frac{\sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}^2}}{\sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu})\|_{\Sigma}^2}} \leq w \Delta^{\text{rel}} \right\} \geq 1 - \delta.$$

Remark 3.6 (F-distribution). If we have $u(\boldsymbol{\mu})^T \Sigma (u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})) = 0$ the random variable $(\Delta^{\text{rel}}(\boldsymbol{\mu}) \|u(\boldsymbol{\mu})\|_{\Sigma})^2 / \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}^2$ follows an F -distribution with degrees of freedom K and K . We may then use the cumulative distribution function (cdf) of the F -distribution and an union bound argument to get an even more accurate lower bound for $\mathbb{P}\{w^{-1} \Delta^{\text{rel}}(\boldsymbol{\mu}) \leq \frac{\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}}{\|u(\boldsymbol{\mu})\|_{\Sigma}} \leq w \Delta^{\text{rel}}(\boldsymbol{\mu}), \forall \boldsymbol{\mu} \in \mathcal{P}\}$; for further details see section B. We remark that although in general we do not have $u(\boldsymbol{\mu})^T \Sigma (u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})) = 0$, often $u(\boldsymbol{\mu})$ and the error are nearly Σ -orthogonal. As a consequence we observe in the numerical experiments in Section 5 that the empirical probability density function (pdf) of $(\Delta^{\text{rel}}(\boldsymbol{\mu}) \|u(\boldsymbol{\mu})\|_{\Sigma}) / \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}$ follows the square root of the pdf of the F distribution pretty well.

Remark 3.7 (Scalar-valued QoI). When estimating the error in scalar-valued QoIs of the form of $s(\boldsymbol{\mu}) = l^T u(\boldsymbol{\mu})$, the covariance matrix is $\Sigma = ll^T$, see subsection 3.1. In that case the random vector $Z \sim \mathcal{N}(0, \Sigma)$ follows the same distribution as Xl where $X \sim \mathcal{N}(0, 1)$ is a standard normal random variable (scalar). The random dual problem (23) then becomes $A(\boldsymbol{\mu})^T Y_i(\boldsymbol{\mu}) = X_i l$ and the solution is $Y_i(\boldsymbol{\mu}) = X_i q(\boldsymbol{\mu})$ where $q(\boldsymbol{\mu})$ is the solution of the deterministic dual problem $A(\boldsymbol{\mu})^T q(\boldsymbol{\mu}) = l$. Dual problems of this form are commonly encountered for estimating linear quantities of interest, see [44] for a general presentation and [24, 46, 52, 2] for the application in reduced order modeling.

4 Approximation of the random dual solutions via the PGD

In order to obtain a fast-to-evaluate a posteriori error estimator we employ the PGD to compute approximations of the solutions of the K dual problems (23). Given K independent realizations Z_1, \dots, Z_K of $Z \sim \mathcal{N}(0, \Sigma)$, we denote by $Y(\boldsymbol{\mu}) = [Y_1(\boldsymbol{\mu}), \dots, Y_K(\boldsymbol{\mu})] \in \mathbb{R}^{N \times K}$ the horizontal concatenation of the K dual solutions (23) so that $Y(\boldsymbol{\mu})$ satisfies the parametrized matrix equation

$$A(\boldsymbol{\mu})^T Y(\boldsymbol{\mu}) = [Z_1, \dots, Z_K].$$

We use the PGD to approximate $Y(\boldsymbol{\mu})$ by

$$\tilde{Y}^L(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_p) = \sum_{m=1}^L Y_X^m Y_1^m(\boldsymbol{\mu}_1) \dots Y_p^m(\boldsymbol{\mu}_p), \quad (28)$$

where the matrices $Y_X^m \in \mathbb{R}^{N \times K}$ and the functions $Y_i^m : \mathcal{P}_i \rightarrow \mathbb{R}$ are built in a greedy way. The PGD algorithm we employ here is essentially the same as the one we used for the primal approximation, except with minor modifications regarding the fact that $Y(\boldsymbol{\mu})$ is a matrix and not a vector, see subsection 4.1. In subsection 4.2 we present several stopping criteria for the dual greedy enrichment, including an adaptive algorithm which builds the primal and dual PGD approximation simultaneously. By replacing $Y_i(\boldsymbol{\mu})$ by $\tilde{Y}_i(\boldsymbol{\mu})$, the i -th column of $\tilde{Y}(\boldsymbol{\mu})$, we define pointwise fast-to-evaluate a posteriori error estimators as

$$\tilde{\Delta}(\boldsymbol{\mu}) := \sqrt{\frac{1}{K} \sum_{i=1}^K (\tilde{Y}_i(\boldsymbol{\mu})^T r(\boldsymbol{\mu}))^2} \quad \text{and} \quad \tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu}) := \sqrt{\frac{\frac{1}{K} \sum_{i=1}^K (\tilde{Y}_i(\boldsymbol{\mu})^T r(\boldsymbol{\mu}))^2}{\frac{1}{K} \sum_{i=1}^K (\tilde{Y}_i(\boldsymbol{\mu})^T f(\boldsymbol{\mu}))^2}}, \quad (29)$$

and fast-to-evaluate error estimators for the (relative) RMS error as

$$\tilde{\Delta} = \sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \tilde{\Delta}(\boldsymbol{\mu})^2} \quad \text{and} \quad \tilde{\Delta}^{\text{rel}} := \sqrt{\frac{\frac{1}{K \#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \sum_{i=1}^K (\tilde{Y}_i(\boldsymbol{\mu})^T r(\boldsymbol{\mu}))^2}{\frac{1}{K \#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \sum_{i=1}^K (\tilde{Y}_i(\boldsymbol{\mu})^T f(\boldsymbol{\mu}))^2}}. \quad (30)$$

Similarly to Proposition 3.3 in [48], the following propositions bound the effectivity indices for the error estimators (29) and (30). This will be used in subsection 4.2 to steer the dual PGD approximation. Again, these results do not depend on the condition number of $A(\boldsymbol{\mu})$ as it is the case for the residual error estimator (12).

Proposition 4.1. Let $0 < \delta < 1$, $w > \sqrt{e}$ and assume

$$K \geq \min \left\{ \frac{\log(\#\mathcal{P}) + \log(\delta^{-1})}{\log(w/\sqrt{e})}, 3 \right\}. \quad (31)$$

Then the fast-to-evaluate estimators $\tilde{\Delta}(\boldsymbol{\mu})$ and $\tilde{\Delta}$ satisfy

$$\mathbb{P} \left\{ (\alpha_\infty w)^{-1} \tilde{\Delta}(\boldsymbol{\mu}) \leq \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_\Sigma \leq (\alpha_\infty w) \tilde{\Delta}(\boldsymbol{\mu}), \quad \boldsymbol{\mu} \in \mathcal{P} \right\} \geq 1 - \delta \quad (32)$$

and

$$\mathbb{P} \left\{ (\alpha_2 w)^{-1} \tilde{\Delta} \leq \sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_\Sigma^2} \leq (\alpha_2 w) \tilde{\Delta} \right\} \geq 1 - \delta, \quad (33)$$

where

$$\alpha_\infty := \max_{\boldsymbol{\mu} \in \mathcal{P}} \left(\max \left\{ \frac{\Delta(\boldsymbol{\mu})}{\tilde{\Delta}(\boldsymbol{\mu})}, \frac{\tilde{\Delta}(\boldsymbol{\mu})}{\Delta(\boldsymbol{\mu})} \right\} \right) \geq 1 \quad \text{and} \quad \alpha_2 := \max \left\{ \frac{\Delta}{\tilde{\Delta}}, \frac{\tilde{\Delta}}{\Delta} \right\} \geq 1. \quad (34)$$

Proof. The proof follows exactly the proof of Proposition 3.3 in [48] and we include it in the appendix for the sake of completeness. \square

Proposition 4.2. Let $0 < \delta < 1$, $w > e$ and assume

$$K \geq \min \left\{ \frac{2 \log(2 \#\mathcal{P}) + 2 \log(\delta^{-1})}{\log(w/e)}, 3 \right\}, \quad (35)$$

Then the fast-to-evaluate estimators $\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})$ and $\tilde{\Delta}^{\text{rel}}$ satisfy

$$\mathbb{P} \left\{ (\alpha_\infty^{\text{rel}} w)^{-1} \tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu}) \leq \frac{\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_\Sigma}{\|u(\boldsymbol{\mu})\|_\Sigma} \leq (\alpha_\infty^{\text{rel}} w) \tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu}), \quad \boldsymbol{\mu} \in \mathcal{P} \right\} \geq 1 - \delta \quad (36)$$

and

$$\mathbb{P} \left\{ (\alpha_2^{\text{rel}} w)^{-1} \tilde{\Delta}^{\text{rel}} \leq \frac{\sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_\Sigma^2}}{\sqrt{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu})\|_\Sigma^2}} \leq (\alpha_2^{\text{rel}} w) \tilde{\Delta}^{\text{rel}} \right\} \geq 1 - \delta, \quad (37)$$

where

$$\alpha_\infty^{\text{rel}} := \max_{\boldsymbol{\mu} \in \mathcal{P}} \left(\max \left\{ \frac{\Delta^{\text{rel}}(\boldsymbol{\mu})}{\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})}, \frac{\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})}{\Delta^{\text{rel}}(\boldsymbol{\mu})} \right\} \right) \geq 1 \quad \text{and} \quad \alpha_2^{\text{rel}} := \max \left\{ \frac{\Delta^{\text{rel}}}{\tilde{\Delta}^{\text{rel}}}, \frac{\tilde{\Delta}^{\text{rel}}}{\Delta^{\text{rel}}} \right\} \geq 1. \quad (38)$$

Proof. Can be proved completely analogously to Proposition 4.1. \square

4.1 PGD algorithm for the random dual solution

To construct the PGD approximation $\tilde{Y}^L(\boldsymbol{\mu})$ (28), we employ a pure greedy algorithm. All terms in (28) are constructed one after the other: after $L - 1$ iterations, the rank-one correction $Y_X^L Y_1^L(\boldsymbol{\mu}_1) \cdots Y_p^L(\boldsymbol{\mu}_p)$ is computed as the solution to

$$\min_{Y_X \in \mathbb{R}^{N \times K}} \min_{Y_1: \mathcal{P}_1 \rightarrow \mathbb{R}} \cdots \min_{Y_p: \mathcal{P}_p \rightarrow \mathbb{R}} \sum_{i=1}^K \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|A(\boldsymbol{\mu})^T (\tilde{Y}_i^{L-1}(\boldsymbol{\mu}) + Y_{X,i} Y_1(\boldsymbol{\mu}_1) \cdots Y_p(\boldsymbol{\mu}_p)) - Z_i\|_2^2, \quad (39)$$

where $\tilde{Y}_i^{L-1}(\boldsymbol{\mu})$ and $Y_{X,i}$ denote the i -th columns of the matrices $\tilde{Y}_i^{L-1}(\boldsymbol{\mu})$ and Y_X respectively. Formulation (39) corresponds to the *Minimal residual PGD* as it consists in minimizing the sum of the K residuals associated with the random dual problems (23). If the matrix $A(\boldsymbol{\mu})$ is SPD, one can also use the *Galerkin PGD* formulation

$$\min_{Y_X \in \mathbb{R}^{N \times K}} \min_{Y_1: \mathcal{P}_1 \rightarrow \mathbb{R}} \cdots \min_{Y_p: \mathcal{P}_p \rightarrow \mathbb{R}} \sum_{i=1}^K \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|A(\boldsymbol{\mu})^T (\tilde{Y}_i^{L-1}(\boldsymbol{\mu}) + Y_{X,i} Y_1(\boldsymbol{\mu}_1) \cdots Y_p(\boldsymbol{\mu}_p)) - Z_i\|_{A(\boldsymbol{\mu})^{-1}}^2, \quad (40)$$

where the canonical norm $\|\cdot\|_2$ has been replaced by the energy norm $\|\cdot\|_{A(\boldsymbol{\mu})^{-1}}$. As explained before, both of these formulations can be solved efficiently using the *alternating least squares*.

Comparing (39) or (40) with the formulations used for the primal PGD solution (9) or (10), the major difference is that the solution is matrix-valued for $K > 1$, so that the first minimization problem over $Y_{\mathcal{X}} \in \mathbb{R}^{N \times K}$ might be, at a first glance, more expensive to compute compared to the minimum over $u_{\mathcal{X}} \in \mathbb{R}^N$ in (9) or (10). However, the minimization problem over $Y_{\mathcal{X}} \in \mathbb{R}^{N \times K}$ possesses a particular structure which we may exploit for computational efficiency. This structure is that each column of $Y_{\mathcal{X}}$ is the solution to a least squares problem with the same operator but with different right-hand side. Indeed, for fixed $Y_1(\boldsymbol{\mu}_1), \dots, Y_p(\boldsymbol{\mu}_p)$, the i -th column of $Y_{\mathcal{X}}$ minimizes

$$Y_{\mathcal{X},i} \mapsto \sum_{\boldsymbol{\mu} \in \mathcal{D}} \|A(\boldsymbol{\mu})^T Y_{\mathcal{X},i} Y_1(\boldsymbol{\mu}_1) \cdots Y_p(\boldsymbol{\mu}_p) - (Z_i - A(\boldsymbol{\mu})^T \tilde{Y}_i^{L-1}(\boldsymbol{\mu}))\|_*^2,$$

with either $\|\cdot\|_* = \|\cdot\|_2$ or $\|\cdot\|_* = \|\cdot\|_{A(\boldsymbol{\mu})^{-1}}$. This means that $Y_{\mathcal{X},i}$ can be computed by solving a linear system of the form $\tilde{A}_{\mathcal{X}} Y_{\mathcal{X},i} = \tilde{Z}_i$ for some operator $\tilde{A}_{\mathcal{X}} \in \mathbb{R}^{N \times N}$ and right-hand side $\tilde{Z}_i \in \mathbb{R}^N$ which are assembled² using $A(\boldsymbol{\mu})^T$, Z_i and also $Y_1(\boldsymbol{\mu}_1), \dots, Y_p(\boldsymbol{\mu}_p)$ and $\tilde{Y}_i^{L-1}(\boldsymbol{\mu})$. Then, by concatenating the K right-hand sides, the matrix $Y_{\mathcal{X}}$ solution to (39) or (40) can be obtained by solving a matrix equation of the form

$$\tilde{A}_{\mathcal{X}} Y_{\mathcal{X}} = [\tilde{Z}_1, \dots, \tilde{Z}_K].$$

For certain solvers $Y_{\mathcal{X}}$ can be computed at a complexity which hardly depends on K . For instance, one can precompute say a sparse LU or Cholesky factorization of $\tilde{A}_{\mathcal{X}}$ and use this factorization to solve for the multiple right-hand sides. Thus, when $K \ll N$, the complexity is dominated by the factorization of $\tilde{A}_{\mathcal{X}}$ so that the costs for computing the dual PGD approximation is at least comparable to the one for the primal PGD approximation.

Remark 4.3 (Complexity comparison with the hierarchical error estimator). Thanks to the above remark, it is worth mentioning that the cost for computing the dual approximation $\tilde{Y}^L(\boldsymbol{\mu})$ is essentially the same as for computing a hierarchical estimate $\tilde{u}^{M+k}(\boldsymbol{\mu})$ with $k = L$. Indeed, both $\tilde{Y}^L(\boldsymbol{\mu})$ and $\tilde{u}^{M+k}(\boldsymbol{\mu})$ require the computation of $k = L$ PGD updates (primal or dual) which can be computed at comparable cost. Therefore, in the numerical experiments, we pay particular attention in comparing our randomized estimators (29) or (30) with the stagnation-based error estimator (11) with $k = L$ to ensure a fair comparison.

Remark 4.4. Instead of looking for $\tilde{Y}^L(\boldsymbol{\mu})$ as in (28), one could have build on the tensor product structure of $\mathbb{R}^{N \times K} = \mathbb{R}^N \otimes \mathbb{R}^K$ and look for approximation on the form of $Y_i(\boldsymbol{\mu}) \approx \sum_{m=1}^L Y_{\mathcal{X}}^m Y_K^m(i) Y_1^m(\boldsymbol{\mu}_1) \cdots Y_p^m(\boldsymbol{\mu}_p)$ where $Y_{\mathcal{X}}^m \in \mathbb{R}^N$ is a vector (not a matrix) and where $Y_K^m : \{1, \dots, K\} \rightarrow \mathbb{R}$. This amounts to consider the column index i as an additional parameter, and thus to increase the order of the problem by one. This option is however not pursued in this paper.

4.2 Primal-dual intertwined PGD algorithm

In this subsection we discuss how to choose the rank L of the dual approximation $\tilde{Y}^L(\boldsymbol{\mu})$.

First, one could choose the rank *based on a given (computational) budget*. Assuming for instance that the budget is tight one could only consider dual PGD approximations of say rank 1 to 5. The numerical experiments in section 5 show that this choice works remarkably well even for high dimensional examples. This choice, however, does not allow controlling the quality of the dual approximation so that neither $\alpha_{2/\infty}$ nor $\alpha_{2/\infty}^{\text{rel}}$ in Propositions 4.1 and 4.2 can be bounded.

In order to control the error in the dual approximation, we suggest building the primal and dual approximations simultaneously in an intertwined manner as described in Algorithm 1. We focus here on the relative error estimator Δ^{ref} to simplify the notation, but other errors such as Δ or $\max_{\boldsymbol{\mu} \in \mathcal{D}} \Delta(\boldsymbol{\mu})$ can be considered as well. First, we calculate K in line 2 for given effectivity αw and failure probability δ via (35). As $\tilde{\Delta}^{\text{ref}}$ invokes Proposition 3.1 $2\#\mathcal{D}$ -times in each greedy iteration and the maximal number of greedy iterations is M_{max} , we need to apply Proposition 3.1 at most $M_{\text{max}} 2\#\mathcal{D}$ -times during Algorithm 1 and thus have to replace $2\#\mathcal{D}$ by $M_{\text{max}} 2\#\mathcal{D}$ in (35) in the calculation of K . We then draw K independent Gaussian random vectors with mean zero and covariance matrix Σ .

Next, we describe how to build $\tilde{u}^M(\boldsymbol{\mu})$ and $\tilde{Y}^L(\boldsymbol{\mu})$ in an intertwined manner. Within each iteration of Algorithm 1 we first compute the *primal* rank-one correction using (9) or (10). Algorithm 1 terminates either if $\tilde{\Delta}^{\text{ref}} \leq \text{tol}$ or $M = M_{\text{max}}$.

²For example with the Galerkin PGD formulation (40), we have $\tilde{A}_{\mathcal{X}} = \sum_{\boldsymbol{\mu} \in \mathcal{D}} A(\boldsymbol{\mu})^T Y_1(\boldsymbol{\mu}_1)^2 \cdots Y_p(\boldsymbol{\mu}_p)^2$ and $\tilde{Z}_i = \sum_{\boldsymbol{\mu} \in \mathcal{D}} (Z_i - A(\boldsymbol{\mu})^T \tilde{Y}_i^{L-1}(\boldsymbol{\mu})) Y_1(\boldsymbol{\mu}_1) \cdots Y_p(\boldsymbol{\mu}_p)$.

Algorithm 1 Intertwined greedy construction of primal and dual PGD approximation

- 1: **INPUT:** tolerance tol , failure probability δ , maximal rank M_{max} , w , α , \mathcal{P} , number of increments k
- 2: Calculate K based on (35) replacing $2\#\mathcal{P}$ by $2M_{max}\#\mathcal{P}$
- 3: Draw K independent random Gaussian vectors with mean zero and covariance matrix Σ
- 4: Initialize $\tilde{u}^0(\boldsymbol{\mu}) = 0$, $M = 0$, $\tilde{\Delta}^{rel} = 2tol$, $L = 1$
- 5:
- 6: **while** $\tilde{\Delta}^{rel} > tol$ and $M < M_{max}$ **do**
- 7: *Enrich primal PGD approximation:*
- 8: $M \leftarrow M + 1$
- 9: Compute the primal rank-one correction using (9) or (10)
- 10: Update $\tilde{u}^M(\boldsymbol{\mu}) = \tilde{u}^{M-1}(\boldsymbol{\mu}) + u_{\mathcal{X}}^M u_1^M(\boldsymbol{\mu}_1) \dots u_p^M(\boldsymbol{\mu}_p)$
- 11:
- 12: *Enrich dual PGD approximation until tolerance reached:*
- 13: Compute $\alpha_{2,k}^{rel}$ as in (43)
- 14: **while** $\alpha_{2,k}^{rel} > \alpha$ **do**
- 15: $L \leftarrow L + 1$
- 16: Compute the primal rank-one correction using (39) or (40)
- 17: Update $\tilde{Y}^L(\boldsymbol{\mu}) = \tilde{Y}^{L-1}(\boldsymbol{\mu}) + Y_{\mathcal{X}}^L Y_1^L(\boldsymbol{\mu}_1) \dots Y_p^L(\boldsymbol{\mu}_p)$
- 18: Compute $\alpha_{2,k}^{rel}$ as in (43)
- 19: **end while**
- 20:
- 21: *Compute the error estimator:*
- 22: Update $\tilde{\Delta}^{rel}$ as in (30)
- 23: **end while**
- 24:
- 25: **OUTPUT:** $\tilde{u}^M(\boldsymbol{\mu})$ and $\tilde{\Delta}^{rel}$

However, due to the dual PGD approximation it is a priori not clear whether $\tilde{\Delta}^{rel}$ has an effectivity close to unity and thus estimates the relative error well. An effectivity that is well below one might result in an premature termination of the algorithm, which is clearly undesirable. This is why every time we need to estimate the error, we first assess the quality of $\tilde{\Delta}^{rel}$ before using it as a stopping criterion and then, if necessary, enrich the dual PGD approximation until the quality of $\tilde{\Delta}^{rel}$ is satisfactory.

Motivated by Proposition 4.2, the natural way to measure the quality of $\tilde{\Delta}^{rel}$ is $\alpha_2^{rel} = \max\{\Delta^{rel}/\tilde{\Delta}^{rel}; \tilde{\Delta}^{rel}/\Delta^{rel}\}$. Since the unbiased error estimator Δ^{rel} is not computationally feasible, as the solution $u(\boldsymbol{\mu})$ is not known in advance, one possibility is to replace Δ^{rel} with the computationally feasible error estimator

$$\Delta_{+k}^{rel} = \sqrt{\frac{\sum_{\boldsymbol{\mu} \in \mathcal{P}} \sum_{k=1}^K (Z_i^T (\tilde{u}^{M+k}(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})))^2}{\sum_{\boldsymbol{\mu} \in \mathcal{P}} \sum_{k=1}^K (Z_i^T \tilde{u}^{M+k}(\boldsymbol{\mu}))^2}}. \quad (41)$$

Here, the number of additional increments k are chosen depending on the available computational budget.

We also introduce

$$\tilde{\Delta}_{+k}^{rel} = \frac{\sqrt{\sum_{\boldsymbol{\mu} \in \mathcal{P}} \sum_{i=1}^K (\tilde{Y}_i^L(\boldsymbol{\mu})^T A(\boldsymbol{\mu}) (\tilde{u}^{M+k}(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})))^2}}{\sqrt{\sum_{\boldsymbol{\mu} \in \mathcal{P}} \sum_{i=1}^K (\tilde{Y}_i^L(\boldsymbol{\mu})^T A(\boldsymbol{\mu}) \tilde{u}^{M+k}(\boldsymbol{\mu}))^2}}, \quad (42)$$

which is the error estimator of the *known* increment $\tilde{u}^{M+k}(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})$. For larger k we expect $\tilde{u}^{M+k}(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})$ to be already a quite good error model justifying the choice $\tilde{\Delta}_{+k}^{rel} = \tilde{\Delta}^{rel}$. We recall and emphasize that the stagnation-based or hierarchical error estimator does neither ensure that the respective estimator bounds the error or allows to bound the effectivity unless we have for instance a uniform bound for the convergence rate of the PGD; see subsection 2.1. Therefore, we use the increments solely as an error model to train the dual PGD approximation such that we can then rely on Proposition 4.2. We then utilize in line 14

$$\alpha_{2,k}^{rel} := \max \left\{ \frac{\Delta_{\pm k}^{rel}}{\tilde{\Delta}_{\pm k}^{rel}}; \frac{\tilde{\Delta}_{\pm k}^{rel}}{\Delta_{\pm k}^{rel}} \right\}, \quad (43)$$




Figure 1: Left: Geometry of the wrench. Dirichlet condition $u = 0$ (black chopped lines), vertical unitary linear forcing f (red arrows). Right: Mesh for the finite element approximation.

where

$$\Delta_{-k}^{\text{rel}} = \sqrt{\frac{\sum_{\boldsymbol{\mu} \in \mathcal{D}} \sum_{k=1}^K (Z_i^T (\tilde{u}^M(\boldsymbol{\mu}) - \tilde{u}^{M-k}(\boldsymbol{\mu})))^2}{\sum_{\boldsymbol{\mu} \in \mathcal{D}} \sum_{k=1}^K (Z_i^T \tilde{u}^M(\boldsymbol{\mu}))^2}} \quad \text{and} \quad \tilde{\Delta}_{-k}^{\text{rel}} = \sqrt{\frac{\sum_{\boldsymbol{\mu} \in \mathcal{D}} \sum_{k=1}^K (\tilde{Y}_i^L(\boldsymbol{\mu})^T A(\boldsymbol{\mu}) (\tilde{u}^M(\boldsymbol{\mu}) - \tilde{u}^{M-k}(\boldsymbol{\mu})))^2}{\sum_{\boldsymbol{\mu} \in \mathcal{D}} \sum_{k=1}^K (\tilde{Y}_i^L(\boldsymbol{\mu})^T A(\boldsymbol{\mu}) \tilde{u}^M(\boldsymbol{\mu}))^2}}. \quad (44)$$

With this choice, one can compute $\alpha_{2,k}^{\text{rel}}$ with *no additional* computational effort. Apart from that, the motivation behind the definition of Δ_{-k}^{rel} and $\tilde{\Delta}_{-k}^{\text{rel}}$ is that we can estimate the norms of the increments $\tilde{u}^M(\boldsymbol{\mu}) - \tilde{u}^{M-k}(\boldsymbol{\mu})$ well. Assuming a relatively uniform convergence behavior of the primal PGD approximation, it can thus be conjectured that $\tilde{\Delta}^{\text{rel}}$ then also estimates the norm of $u(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})$ well; this is demonstrated in section 5. Algorithm 1 terminates the enrichment of the dual PGD approximation within each iteration if $\alpha_{2,k}^{\text{rel}}$ falls below some prescribed value $\alpha > 1$.

Remark 4.5 (Independence). We note that if the error in the primal PGD approximation is already close to the target tolerance tol the primal PGD approximation can become dependent on the respective realizations of Z_1, \dots, Z_K in the following sense: For some realizations the estimator $\tilde{\Delta}^{\text{rel}}$ lies below tol and thus terminates Algorithm 1 while for others we have $\tilde{\Delta}^{\text{rel}} > tol$ and thus obtain a primal PGD approximation of higher rank. A careful analysis of this effect on the failure probabilities in Proposition 4.2 is however not necessary as by replacing $2\#\mathcal{D}$ by $2M_{\max}\#\mathcal{D}$ in the calculation of K in line 2 we ensure that the number of samples K is chosen large enough such that we can estimate all possible outcomes.

5 Numerical experiments

We demonstrate the randomized error estimator for two test cases which are derived from the benchmark problem introduced in [36]. The source code is freely available in Matlab[®] at the address³ so that all numerical results presented in this section are entirely reproducible. We consider the system of linear PDEs whose solution $u : D \rightarrow \mathbb{R}^2$ is a vector field satisfying

$$-\text{div}(C : \varepsilon(u)) - k^2 u = f, \quad (45)$$

where $D \subset \mathbb{R}^2$ is a domain which has the shape of a wrench; see Fig. 1. Here, $\varepsilon(u) = \frac{1}{2}(\nabla u + \nabla u^T)$ is the strain tensor, k^2 the wave number, and C the fourth-order stiffness tensor derived from the plane stress assumption with Young's modulus E and fixed Poisson coefficient $\nu = 0.3$ such that

$$C : \varepsilon = \frac{E}{1 + \nu} \varepsilon + \frac{\nu E}{1 - \nu^2} \text{trace}(\varepsilon) I_2. \quad (46)$$

The right hand-side f and the boundary conditions are represented in Fig. 1. The finite element discretization using piecewise affine basis functions ψ_1, \dots, ψ_N yields $N = 1467$ degrees of freedom. In the remainder of the paper, we only consider estimating errors in the natural Hilbert norm $\|\cdot\|_X$ defined by

$$\|u\|_X^2 = \int_D \text{trace}(\nabla u(x)^T \nabla u(x)) + \|u(x)\|_2^2 dx,$$

so that $\Sigma = R_{X_N}$, as explained section 2 (see equation (5)).

³https://gitlab.inria.fr/ozahm/randomdual_pgd.git





Figure 2: Time-harmonic elastodynamics. Left: Norm of the solution on \mathcal{P} . Right: finite element solution $u(\boldsymbol{\mu})$ for $\boldsymbol{\mu} = 1$. The color represents the VonMises stress.





Figure 3: Time-harmonic elastodynamics. Comparison of several error estimators for the relative error of an PGD approximation with rank ten. Solid black line: Exact error (47). Dashed black line: $\Delta_{10}^{\text{res}}(\boldsymbol{\mu})$. Dotted red line: one realization of the exact relative random estimate $\Delta^{\text{rel}}(\boldsymbol{\mu})$ with $K = 6$, see Equation (26). Solid red line: approximation $\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})$ of $\Delta^{\text{rel}}(\boldsymbol{\mu})$ defined in (29) rank(\tilde{Y}) = 3 (left) and rank(\tilde{Y}) = 10 (right). Solid blue line: $\Delta_{10,k}^{\text{stag}}(\boldsymbol{\mu})$ with either $k = 3$ (left) or $k = 10$ (right).

5.1 Non-homogeneous time-harmonic elastodynamics problem

In this benchmark problem the Young's modulus is fixed to $E = 1$ uniformly over the domain D and the wave number k^2 varies between 0.5 and 1.2. The scalar parameter $\boldsymbol{\mu} = k^2$ varies on the parameter set \mathcal{P} defined as the uniform grid of $[0.5, 1.2]$ containing $\#\mathcal{P} = 500$ points. The parametrized operator derived from (45) is

$$\mathcal{A}(\boldsymbol{\mu}) = -\text{div}(C : \varepsilon(\cdot)) - \boldsymbol{\mu}(\cdot).$$

The parametrized stiffness matrix $A(\boldsymbol{\mu}) \in \mathbb{R}^{N \times N}$, obtained via a FE discretization, admits an affine decomposition $A(\boldsymbol{\mu}) = A_1 - \boldsymbol{\mu}A_2$, where $(A_1)_{ij} = \int_D \varepsilon(\psi_i) : C : \varepsilon(\psi_j) dx$ and $(A_2)_{ij} = \int_D \psi_i^T \psi_j dx$. As this benchmark problem contains only one scalar parameter, it is convenient to inspect how the method behaves within the parameter set. We observe in Fig. 2 that the norm of the solution blows up around three parameter values, which indicates resonances of the operator. This shows that for this parameter range, the matrix $A(\boldsymbol{\mu})$ is not SPD and therefore we use the Minimal Residual PGD formulations (9) and (39) for the primal and dual PGD approximation.

5.1.1 Estimating the relative error

We first compute a PGD approximation $\tilde{u}^M(\boldsymbol{\mu})$ with rank $M = 10$ and we aim at estimating the relative error

$$\frac{\|u(\boldsymbol{\mu}) - \tilde{u}^{10}(\boldsymbol{\mu})\|_{X_N}}{\|u(\boldsymbol{\mu})\|_{X_N}} \quad (47)$$

for every $\boldsymbol{\mu} \in \mathcal{P}$. In Fig. 3 we show a comparison of the performances of the randomized residual-based error estimator $\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})$ defined in (29), the error estimator based on the dual norm of the residual, and the stagnation-based error estimator defined, respectively, as

$$\Delta_M^{\text{res}}(\boldsymbol{\mu}) := \frac{\|A(\boldsymbol{\mu})\tilde{u}^M(\boldsymbol{\mu}) - f(\boldsymbol{\mu})\|_{\mathcal{X}'_N}}{\|f(\boldsymbol{\mu})\|_{\mathcal{X}'_N}} \quad \text{and} \quad \Delta_{M,k}^{\text{stag}}(\boldsymbol{\mu}) := \frac{\|\tilde{u}^{M+k}(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})\|_{\mathcal{X}_N}}{\|\tilde{u}^{M+k}(\boldsymbol{\mu})\|_{\mathcal{X}_N}}. \quad (48)$$

First, we observe in Fig. 3 that although the residual-based error estimator $\Delta_{10}^{\text{res}}(\boldsymbol{\mu})$ follows the overall variations in the error, it underestimates the error significantly over the whole considered parameter range. Following Remark 4.3, in order to have a fair comparison between the stagnation-based and the randomized error estimator, we set $k = L = \text{rank}(\tilde{Y}^L)$ so that the two estimators require the same computational effort in terms of number of PGD corrections that have to be computed. We observe that with either $k = L = 3$ (left) or $k = L = 10$ (right), the randomized residual-based error estimator is closer to the true relative error compared to the stagnation-based estimator $\Delta_{10,k}^{\text{stag}}(\boldsymbol{\mu})$. Notice however that the randomized error estimator $\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})$ is converging towards $\Delta^{\text{rel}}(\boldsymbol{\mu})$ with $\text{rank}(\tilde{Y})$ which is not the true relative error (see the dashed red lines on Fig. 3), whereas $\Delta_{10,k}^{\text{stag}}(\boldsymbol{\mu})$ will converge to the true relative error with k . Nevertheless, we do observe that for small $k = L$, our random dual is always competitive. In addition, increasing the number of samples K will reduce the variance of both $\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})$ and $\Delta^{\text{rel}}(\boldsymbol{\mu})$ such that the randomized error estimator will follow the error even more closely (see also Fig. 4 to that end). We finally note that Fig. 3 shows only one (representative) realization of $\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})$ and $\Delta^{\text{rel}}(\boldsymbol{\mu})$, and the behavior of the effectivity of $\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})$ over various realizations will be discussed next.

In detail, we first draw 100 different realizations of $Z = [Z_1, \dots, Z_K]$, $Z_i \sim \mathcal{N}(0, \Sigma)$, for $K = 1, 2, 6, 20$ and compute the corresponding dual solutions \tilde{Y}^L with either $L = \text{rank}(\tilde{Y}^L) = 1, 4, 8, 16$. Next, we compute the randomized error estimator $\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})$ and calculate the effectivity index

$$\eta(\boldsymbol{\mu}) = \frac{\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})}{\|u(\boldsymbol{\mu}) - \tilde{u}^{10}(\boldsymbol{\mu})\|_{\mathcal{X}_N} / \|u(\boldsymbol{\mu})\|_{\mathcal{X}_N}}, \quad (49)$$

for all $\boldsymbol{\mu} \in \mathcal{P}$. The resulting histograms are shown in Fig. 4. First, we emphasize that especially for $\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})$ based on a PGD approximation of the dual problems of rank 8, 16, and to some extent even 4, we clearly see a strong concentration of $\eta(\boldsymbol{\mu})$ around one as desired. If the number of samples K increases we observe, as expected, that the effectivity of $\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})$ is very close to one (see for instance picture with $K = 20$ and $\text{rank}(\tilde{Y}) = 8$ in Fig. 4). We recall that for the primal PGD approximation used for these numerical experiments we chose the rank $M = 10$. We thus highlight and infer that by using a dual PGD approximation of lower rank than those of the primal PGD approximation, we can obtain an error estimator with an effectivity index very close to one.

Using a dual PGD approximation of rank one and to some extent rank four, we observe in Fig. 4 that the mean over all samples of $\eta(\boldsymbol{\mu})$ is significantly smaller than one (see for instance picture with $K = 6$ and $\text{rank}(\tilde{Y}) = 1$ in Fig. 4). This behavior can be explained by the fact that in these cases the dual PGD approximation is not rich enough to characterize all the features of the error which results in an underestimation of the error.

Finally, as explained in Remark 3.6 and more detailed in section B, $\eta(\boldsymbol{\mu})$ is distributed as the square root of an F -distribution with degrees of freedom (K, K) , provided some orthogonality condition holds. The solid black line in Fig. 4 is the pdf of such a square root of an F -distribution with degrees of freedom (K, K) . In particular, we observe a nice accordance with the histograms showing that, for sufficiently large $L = \text{rank}(\tilde{Y})$, the effectivity indices concentrate around unity with K , as predicted by Proposition 4.2.

5.1.2 Testing the intertwined algorithm

We now illustrate the intertwined algorithm introduced in Subsection 4.2. Here, the goal is to monitor the convergence of the PGD algorithm during the iteration process $M = 1, 2, \dots$. We thus estimate the relative RMS error

$$\frac{\|u - \tilde{u}^M\|}{\|u\|} \stackrel{(\gamma)}{=} \sqrt{\frac{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})\|_{\mathcal{X}_N}^2}{\frac{1}{\#\mathcal{P}} \sum_{\boldsymbol{\mu} \in \mathcal{P}} \|u(\boldsymbol{\mu})\|_{\mathcal{X}_N}^2}},$$

with our relative randomized error estimator $\tilde{\Delta}^{\text{rel}}(\boldsymbol{\mu})$ introduced in (29) with the rank adaptation for the random dual approximation \tilde{Y}^L . We employ the economical error indicators within the intertwined algorithm 1 with $k = 6$. Fig. 5 shows the performance of our estimator (10 realizations) as well as the increase in $L = \text{rank}(\tilde{Y}^L)$ (averaged over the 10 realizations) during the PGD iteration process. We also compare to the relative RMS residual norm




Figure 4: Time-harmonic elastodynamics. Histograms of 100 realizations of the effectivity index $\{\eta(\boldsymbol{\mu}), \boldsymbol{\mu} \in \mathcal{D}\}$ defined by (49) for $K = 1, 2, 6, 20$ and $\text{rank}(\tilde{Y}) = 1, 4, 8, 16$. Solid line: pdf of the square root of the F-distribution with degrees of freedom K and K .

Δ_M^{res} (12) and to the relative RMS stagnation-based error estimator $\Delta_{M,k}^{\text{stag}}$ (11) with $k = 5$ for which we recall the expression

$$\Delta_M^{\text{res}} := \sqrt{\frac{\frac{1}{\#\mathcal{D}} \sum_{\boldsymbol{\mu} \in \mathcal{D}} \|A(\boldsymbol{\mu})\tilde{u}^M(\boldsymbol{\mu}) - f(\boldsymbol{\mu})\|_{\mathcal{X}'_N}^2}{\frac{1}{\#\mathcal{D}} \sum_{\boldsymbol{\mu} \in \mathcal{D}} \|f(\boldsymbol{\mu})\|_{\mathcal{X}'_N}^2}} \quad \text{and} \quad \Delta_{M,5}^{\text{stag}} := \sqrt{\frac{\frac{1}{\#\mathcal{D}} \sum_{\boldsymbol{\mu} \in \mathcal{D}} \|\tilde{u}^{M+5}(\boldsymbol{\mu}) - \tilde{u}^M(\boldsymbol{\mu})\|_{\mathcal{X}_N}^2}{\frac{1}{\#\mathcal{D}} \sum_{\boldsymbol{\mu} \in \mathcal{D}} \|\tilde{u}^{M+5}(\boldsymbol{\mu})\|_{\mathcal{X}_N}^2}} \quad (50)$$

We notice that our estimator is much more stable compared to the stagnation-based error estimator which tends to dramatically underestimate the error during the phases where the PGD is plateauing. Our estimator is also much more accurate compared to Δ_M^{res} . When comparing the results with $\alpha = 5$ (left) and $\alpha = 2$ (right), we observe that a large α results in smaller dual rank $L = \text{rank}(\tilde{Y}^L)$ but the quality of the error indicator is deteriorated. Reciprocally, a smaller α yields more demanding dual approximations so that $L = \text{rank}(\tilde{Y}^L)$ is, on average, larger. Let us notice that a rank explosion of the dual variable \tilde{Y}^L is not recommended because the complexity for evaluating $\tilde{\Delta}^{\text{rel}}$ will also explode. This will be further discussed in an upcoming paper.




Figure 5: Time-harmonic elastodynamics. Illustration of the intertwined algorithm 1. Red lines: error estimator $\tilde{\Delta}^{\text{rel}}$ produced by Algorithm 1 with $k = 6$ and with either $\alpha = 5$ (left) or $\alpha = 2$ (right). Each of the ten red lines corresponds to a different realization of the Z_1, \dots, Z_K with $K = 10$. Solid black line: exact relative RMS error $\|u - \tilde{u}^M(\boldsymbol{\mu})\|/\|u(\boldsymbol{\mu})\|$. Blue line: stagnation-based error estimator $\Delta_{M,k}^{\text{stag}}$ defined by (50) with $k = 5$. Dashed black lines: relative RMS residual norm Δ_M^{res} , see (50). Background color: value of $L = \text{rank}(\tilde{Y}^L)$ averaged over the ten realizations.




Figure 6: Linear elasticity. Left: One realization of the field $\log(E)$. Right: Corresponding solution $u(\boldsymbol{\mu})$. The color represents the VonMises stress.

5.2 Parametrized linear elasticity with high-dimensional parameter space

In this benchmark problem the wave number is set to $k^2 = 0$ so that (45) becomes an elliptic PDE. Young's modulus E is a random field on D that is log-normally distributed so that $\log(E)$ is a Gaussian random field on D with zero-mean and covariance function $\text{cov}(x, y) = \sigma_0 \exp(-(\|x - y\|_2/l_0)^2)$ with $\sigma = 0.4$ and $l_0 = 4$. A realization of the Young's modulus field is shown in Fig. 6. Since E is a random field, such a parametrized problem contains infinitely many parameters. We use here a truncated Karhunen Loève decomposition with $p = 20$ terms to approximate the random Young's modulus field by a finite number of parameter $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{20}$, that is

$$E \approx E(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{20}) = \exp\left(\sum_{i=1}^{20} \boldsymbol{\mu}_i \sqrt{\sigma_i} \phi_i\right), \quad (51)$$

where (σ_i, ϕ_i) is the i -th eigenpair of the covariance function cov and $\boldsymbol{\mu}_i$ are independent standard Gaussian random variables. This means that the parameter set is \mathbb{R}^{20} and that ρ is the density of a standard normal distribution in dimension $p = 20$. In actual practice, we replace the parameter set by a 20-dimensional grid $\mathcal{P} := (x_1, \dots, x_{50})^{20} \subset \mathbb{R}^{20}$, where x_1, \dots, x_{50} is an optimal quantization of the univariate standard normal distribution using 50 points which are obtained by Lloyd's algorithm [20]. Next, we use the Empirical Interpolation Method (EIM) [6] using 28 "magic points" in order to obtain an affine decomposition of the form

$$E(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{20}) \approx \sum_{j=1}^{28} E_j \zeta^j(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_{20}). \quad (52)$$




Figure 7: Linear elasticity. Monitoring the PGD convergence: Comparison of the residual-based error Δ_M^{res} defined in (12) (dashed black curves), the stagnation-based estimator $\Delta_{M,k}^{\text{stag}}$ defined in (11) (blue lines), and the estimator $\tilde{\Delta}^{\text{rel}}$ defined by (30) with $K = 3$ and with $L = \text{rank}(\tilde{Y}^L)$ either equal to 1 (left), 3 (middle), or 5 (right). The dashed red lines corresponds to the exact random dual estimator Δ^{rel} .

for some spatial functions $E_j : D \rightarrow \mathbb{R}$ and parametric functions $\zeta^j : \mathbb{R}^{20} \rightarrow \mathbb{R}$. With the EIM method, the functions ζ^j are defined by means of magic point $x^j \in D$ by $\zeta^j(\boldsymbol{\mu}) = \exp(\sum_{i=1}^{20} \boldsymbol{\mu}_i \sqrt{\sigma_i} \phi_i(x^j)) = \prod_{i=1}^{20} \zeta_i^j(\boldsymbol{\mu}_i)$, where $\zeta_i^j(\boldsymbol{\mu}_i) = \exp(\boldsymbol{\mu}_i \sqrt{\sigma_i} \phi_i(x^j))$. Together with the truncated Karhunen Loève decomposition (51) and with the EIM (52) the approximation of the Young's modulus fields does not exceed 1% of error, and the differential operator is given by

$$\mathcal{A}(\boldsymbol{\mu}) = -\text{div}(C(\boldsymbol{\mu}) : \varepsilon(\cdot)) \approx -\sum_{j=1}^{\tilde{L}} \text{div}(K_j : \varepsilon(\cdot)) \zeta_1^j(\boldsymbol{\mu}_1) \dots \zeta_{20}^j(\boldsymbol{\mu}_{20}),$$

where the forth-order stiffness tensors C_j are given as in (46) with E replaced by E_j . The parametrized stiffness matrix $A(\boldsymbol{\mu})$, obtained via a FE discretization, admits an affine decomposition $A(\boldsymbol{\mu}) = \sum_{i=1}^{28} A_i \zeta_1^i(\boldsymbol{\mu}_1) \dots \zeta_{20}^i(\boldsymbol{\mu}_{20})$, where $A_i \in \mathbb{R}^{N \times N}$ is the matrix resulting from the discretization of the continuous PDE operator $\text{div}(K_j : \varepsilon(\cdot))$. As $A(\boldsymbol{\mu})$ is a parametrized SPD matrix, we use the *Galerkin PGD* formulation (10) and (40) to compute the primal and dual PGD solutions $\tilde{u}^M(\boldsymbol{\mu})$ and $\tilde{Y}^L(\boldsymbol{\mu})$, respectively.

For this high-dimensional benchmark problem our goal is to monitor the convergence of the PGD. Again, we estimate the relative RMS error $\|u - \tilde{u}^M\|/\|u\|$ using $\tilde{\Delta}_M^{\text{rel}}$ as in (29) where we fix $K = 3$ and the rank of the approximate dual \tilde{Y}^L as $L = \text{rank}(\tilde{Y}^L) = 1, 3, 5$. The results are reported in Fig. 7. We highlight that, remarkably, a very small dual rank L is enough to obtain a very accurate estimation of the error. In particular, our approach outperforms the stagnation based error estimator, which is very unstable, and also the relative RMS residual norm, which is almost one order of magnitude off. As in the previous subsection, we mention that the evaluation of the estimator $\tilde{\Delta}^{\text{rel}}$ can be numerically expensive due to high primal and dual ranks (M and L).

6 Conclusions

In this paper we proposed a novel a posteriori error estimator for the PGD, extending the concept introduced in [48] to this setting. The proposed randomized error estimator does not require the estimation of stability constants, its effectivity is close to unity and lies within an interval chosen by the user at specified high probability. Moreover, the suggested framework allows estimating the error in any norm chosen by the user and the error in

some (vector-valued) QoI. To derive the error estimator we make use of the concentration phenomenon of Gaussian maps. Exploiting the error residual relationship and approximating the associated random dual problems via the PGD yields a fast-to-evaluate a posteriori error estimator. To ensure that also the effectivity of this fast-to-evaluate error estimator lies within a prescribed interval at high probability we suggested building the primal and dual PGD approximations in a intertwined manner, where we enrich the dual PGD approximation in each iteration until the error estimator reaches the desired quality.

The numerical experiments for a parametrized time-harmonic elastodynamics problem and a parametrized linear elasticity problem with 20-dimensional parameter space show that very often and even for a very tight effectivity interval it is sufficient to consider a dual PGD approximation that has a significantly lower rank than the primal PGD approximation. Thanks to its favorable computational complexity, the costs for the proposed a posteriori error estimator are about the same or much smaller than the costs for the primal PGD approximation. A numerical comparison with an error estimator defined as the dual norm of the residual and a stagnation-based error estimator (also called hierarchical error estimator) shows that the randomized error estimator provides a much more precise description of the behavior of the error at a comparable computational cost.

A Proofs

Proof of Corollary 3.5. Assuming $w > \sqrt{\epsilon}$ and $K \geq 3$ we can use Proposition 3.1 for any vector $v(\boldsymbol{\mu}) \in \mathbb{R}^N$ and a fixed parameter $\boldsymbol{\mu} \in \mathcal{P}$ and obtain

$$\mathbb{P}\left\{w^{-1}\|v(\boldsymbol{\mu})\|_{\Sigma} \leq \|\Phi v(\boldsymbol{\mu})\|_2 \leq w\|v(\boldsymbol{\mu})\|_{\Sigma}\right\} \geq 1 - \left(\frac{\sqrt{\epsilon}}{w}\right)^K. \quad (53)$$

Using the definition of $\Delta^{\text{rel}}(\boldsymbol{\mu})$, invoking (53) twice and using a union bound argument the following inequality

$$\frac{1}{w^2}\Delta^{\text{rel}}(\boldsymbol{\mu}) = \frac{1}{w^2} \frac{\|\Phi(u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu}))\|_2}{\|\Phi u(\boldsymbol{\mu})\|_2} \leq \frac{1}{w} \frac{\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}}{\|\Phi u(\boldsymbol{\mu})\|_2} \leq \frac{\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}}{\|u(\boldsymbol{\mu})\|_{\Sigma}} \leq w \frac{\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}}{\|\Phi u(\boldsymbol{\mu})\|_2} \leq w^2 \Delta^{\text{rel}}(\boldsymbol{\mu}) \quad (54)$$

holds at least with probability $1 - 2\left(\frac{\sqrt{\epsilon}}{w}\right)^K$. Using again a union bound argument yields

$$\mathbb{P}\left\{\frac{1}{w^2}\Delta^{\text{rel}}(\boldsymbol{\mu}) \leq \frac{\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}}{\|u(\boldsymbol{\mu})\|_{\Sigma}} \leq \frac{1}{w^2}\Delta^{\text{rel}}(\boldsymbol{\mu}), \forall \boldsymbol{\mu} \in \mathcal{P}\right\} \geq 1 - 2(\#\mathcal{P})\left(\frac{\sqrt{\epsilon}}{w}\right)^K.$$

For a given $0 < \delta < 1$ condition $K \geq (\log(2\#\mathcal{P}) + \log(\delta^{-1}))/\log(w/\sqrt{\epsilon})$, is equivalent to $1 - (2\#\mathcal{P})\left(\frac{\sqrt{\epsilon}}{w}\right)^K \geq 1 - \delta$ and ensures that (54) holds for all $\boldsymbol{\mu} \in \mathcal{P}$ with probability larger than $1 - \delta$. Setting $\hat{w} = w^2$ and renaming \hat{w} as w concludes the proof for $\Delta^{\text{rel}}(\boldsymbol{\mu})$; the proof for Δ^{rel} follows similar lines. \square

Proof of Proposition 4.1. We provide the proof for $\Delta(\boldsymbol{\mu})$ noting that the proof for Δ can be done all the same. By Corollary 3.4, it holds with probability larger than $1 - \delta$ that $w^{-1}\Delta(\boldsymbol{\mu}) \leq \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma} \leq w\Delta(\boldsymbol{\mu})$ for all $\boldsymbol{\mu} \in \mathcal{P}$. Then with the same probability we have

$$\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\| \leq w\Delta(\boldsymbol{\mu}) \leq w \left(\sup_{\boldsymbol{\mu}' \in \mathcal{P}} \frac{\Delta(\boldsymbol{\mu}')}{\tilde{\Delta}(\boldsymbol{\mu}')} \right) \tilde{\Delta}(\boldsymbol{\mu}) \stackrel{(34)}{\leq} (\alpha_{\infty} w) \tilde{\Delta}(\boldsymbol{\mu}),$$

and

$$\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\| \geq w^{-1}\Delta(\boldsymbol{\mu}) \geq w^{-1} \left(\inf_{\boldsymbol{\mu}' \in \mathcal{P}} \frac{\Delta(\boldsymbol{\mu}')}{\tilde{\Delta}(\boldsymbol{\mu}')} \right) \tilde{\Delta}(\boldsymbol{\mu}) \stackrel{(34)}{\geq} (\alpha_{\infty} w)^{-1} \tilde{\Delta}(\boldsymbol{\mu}),$$

for any $\boldsymbol{\mu} \in \mathcal{P}$, which yields (32) and concludes the proof. \square

B Connection of $\Delta^{\text{rel}}(\boldsymbol{\mu})$ with the F -distribution

If we have $u(\boldsymbol{\mu})^T \Sigma (u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})) = 0$ we can give an even more precise description of the random variable $(\Delta^{\text{rel}}(\boldsymbol{\mu})\|u(\boldsymbol{\mu})\|_{\Sigma})^2 / \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}^2$ and thus $\mathbb{P}\left\{w^{-1}\Delta^{\text{rel}}(\boldsymbol{\mu}) \leq \frac{\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}}{\|u(\boldsymbol{\mu})\|_{\Sigma}} \leq w\Delta^{\text{rel}}(\boldsymbol{\mu}), \forall \boldsymbol{\mu} \in \mathcal{P}\right\}$: In fact, under this assumption the random variable $(\Delta^{\text{rel}}(\boldsymbol{\mu})\|u(\boldsymbol{\mu})\|_{\Sigma})^2 / \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}^2$ follows an F -distribution with degrees of freedom K and K ; we also say $(\Delta^{\text{rel}}(\boldsymbol{\mu})\|u(\boldsymbol{\mu})\|_{\Sigma})^2 / \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}^2 \sim F(K, K)$ as stated in the following lemma.

Lemma B.1 (*F-distribution*). Assume that there holds $u(\boldsymbol{\mu})^T \Sigma(u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})) = 0$. Then the random variable $(\Delta^{\text{rel}}(\boldsymbol{\mu}) \|u(\boldsymbol{\mu})\|_{\Sigma}^2 / \|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}^2)$ follows an *F-distribution* with parameters K and K .

Proof. We will exploit the fact that in order to show that a random variable X follows an *F-distribution* it is sufficient to show that we have $X = \frac{Q_1/K_1}{Q_2/K_2}$, where Q_1 and Q_2 are independent random variables that follow a chi-square distribution with K_1 and K_2 degrees of freedom, respectively (see for instance [19, Definition 9.7.1]). Similar to (16) we have that

$$\Delta^{\text{rel}}(\boldsymbol{\mu})^2 \frac{\|u(\boldsymbol{\mu})\|_{\Sigma}^2}{\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}^2} = \frac{\frac{1}{K} \sum_{k=1}^K \left(Z_i^T (u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})) \right)^2}{\frac{1}{K} \sum_{k=1}^K \left(Z_i^T u(\boldsymbol{\mu}) \right)^2} \frac{\|u(\boldsymbol{\mu})\|_{\Sigma}^2}{\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}^2} = \frac{\frac{1}{K} \sum_{k=1}^K \left(Z_i^T \frac{u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})}{\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}} \right)^2}{\frac{1}{K} \sum_{k=1}^K \left(Z_i^T \frac{u(\boldsymbol{\mu})}{\|u(\boldsymbol{\mu})\|_{\Sigma}} \right)^2} \sim \frac{\frac{1}{K} Q_1}{\frac{1}{K} Q_2},$$

where Q_1 and Q_2 follow a chi-squared distribution with K degrees of freedom. It thus remains to show that under the above assumption Q_1 and Q_2 are independent. To that end, recall that Z_1, \dots, Z_K are independent copies of $Z \sim \mathcal{N}(0, \Sigma)$. Per our assumption that Σ is symmetric and positive semi-definite there exists a symmetric matrix $\Sigma^{1/2}$ that satisfies $\Sigma = \Sigma^{1/2} \Sigma^{1/2}$, exploiting for instance the Schur decomposition of Σ . Then we have $Z_i = \Sigma^{1/2} g_i$, where g_i are K independent copies of a standard Gaussian random vector. Thanks to the assumption $u(\boldsymbol{\mu})^T \Sigma(u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})) = 0$ the vectors $\hat{u}(\boldsymbol{\mu}) := \Sigma^{1/2} u(\boldsymbol{\mu})$ and $\hat{e}(\boldsymbol{\mu}) := \Sigma^{1/2} (u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu}))$ satisfy $\hat{u}(\boldsymbol{\mu})^T \hat{e}(\boldsymbol{\mu}) = 0$. As a consequence there exists an orthogonal matrix $V(\boldsymbol{\mu}) \in \mathbb{R}^{N \times N}$ that maps $\hat{u}(\boldsymbol{\mu})$ to e_1 and $\hat{e}(\boldsymbol{\mu})$ to e_2 , respectively, where $e_j, j = 1, \dots, N$ denote the standard basis vectors in \mathbb{R}^N . We can then write for $i = 1, \dots, K$

$$(Z_i^T u(\boldsymbol{\mu}))^2 = ([\Sigma^{1/2} g_i]^T u(\boldsymbol{\mu}))^2 = (g_i^T \Sigma^{1/2} u(\boldsymbol{\mu}))^2 = (g_i^T \hat{u}(\boldsymbol{\mu}))^2 = (g_i^T V(\boldsymbol{\mu})^T V(\boldsymbol{\mu}) \hat{u}(\boldsymbol{\mu}))^2 = ([V(\boldsymbol{\mu}) g_i]^T e_1)^2. \quad (55)$$

As there holds $V(\boldsymbol{\mu}) g_i \sim \mathcal{N}(0, I_N)$, where I_N is the identity matrix in \mathbb{R}^N , we have for $\bar{g}_i \sim \mathcal{N}(0, I_N)$

$$(Z_i^T u(\boldsymbol{\mu}))^2 = (\bar{g}_i^T e_1)^2 = (\bar{g}_i^1)^2 \quad \text{and} \quad [Z_i^T (u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu}))]^2 = (\bar{g}_i^T e_2)^2 = (\bar{g}_i^2)^2. \quad (56)$$

Here, \bar{g}_i^s denotes the s -th entry of the vector \bar{g}_i . As the entries of a standard Gaussian random vector are independent, we can thus conclude that the random variables Q_1 and Q_2 are independent as well. \square

We may then use the cumulative distribution function (cdf) of the *F-distribution* and an union bound argument to obtain the following result.

Corollary B.2. Assume that there holds $u(\boldsymbol{\mu})^T \Sigma(u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})) = 0$. Then, we have

$$\mathbb{P} \left\{ w^{-1} \Delta^{\text{rel}}(\boldsymbol{\mu}) \leq \frac{\|u(\boldsymbol{\mu}) - \tilde{u}(\boldsymbol{\mu})\|_{\Sigma}}{\|u(\boldsymbol{\mu})\|_{\Sigma}} \leq w \Delta^{\text{rel}}(\boldsymbol{\mu}), \forall \boldsymbol{\mu} \in \mathcal{P} \right\} \geq 1 - (\#\mathcal{P}) \left(I_{\left(\frac{K}{K+w^2K}\right)} \left(\frac{K}{2}, \frac{K}{2} \right) + 1 - I_{\left(\frac{Kw^2}{Kw^2+K}\right)} \left(\frac{K}{2}, \frac{K}{2} \right) \right), \quad (57)$$

where $I_x(a, b)$ is the regularized incomplete beta function defined as $(\int_0^x t^{a-1} (1-t)^{b-1} dt) / (\int_0^1 t^{a-1} (1-t)^{b-1} dt)$. Note that the cdf $F(x, K, K)$ of an *F-distribution* with K and K degrees of freedom is given as $F(x, K, K) = I_{\left(\frac{Kx}{Kx+K}\right)} \left(\frac{K}{2}, \frac{K}{2} \right)$.

References

- [1] I. ALFARO, D. GONZÁLEZ, S. ZLOTNIK, P. DÍEZ, E. CUETO, AND F. CHINESTA, *An error estimator for real-time simulators based on model order reduction*, Adv. Model. and Simul. in Eng. Sci., 2 (2015).
- [2] A. AMMAR, F. CHINESTA, P. DIEZ, AND A. HUERTA, *An error estimator for separated representations of highly multidimensional models*, Comput. Methods Appl. Mech. Engrg., 199 (2010), pp. 1872–1880.
- [3] A. AMMAR, B. MOKDAD, F. CHINESTA, AND R. KEUNINGS, *A new family of solvers for some classes of multidimensional partial differential equations encountered in kinetic theory modeling of complex fluids*, J. Non-Newtonian Fluid Mech., 139 (2006), pp. 153–176.
- [4] ———, *A new family of solvers for some classes of multidimensional partial differential equations encountered in kinetic theory modelling of complex fluids: Part ii: Transient simulation using space-time separated representations*, J. Non-Newtonian Fluid Mech., 144 (2007), pp. 98–121.
- [5] O. BALABANOV AND A. NOUY, *Randomized linear algebra for model reduction. part i: Galerkin methods and error estimation*, tech. rep., arXiv, 1803.02602, 2018.
- [6] M. BARRAULT, Y. MADAY, N. NGUYEN, AND A. PATERA, *An 'empirical interpolation' method: application to efficient reduced-basis discretization of partial differential equations*, C. R. Math. Acad. Sci. Paris Series I, 339 (2004), pp. 667–672.
- [7] M. BILLAUD-FRIESS, A. NOUY, AND O. ZAHM, *A tensor approximation method based on ideal minimal residual formulations for the solution of high-dimensional problems*, ESAIM Math. Model. Numer. Anal., 48 (2014), pp. 1777–1806.
- [8] J. BRUNKEN, T. LEIBNER, M. OHLBERGER, AND K. SMETANA, *Problem adapted Hierarchical Model Reduction for the Fokker-Planck equation*, in Proceedings of ALGORITMY 2016, 20th Conference on Scientific Computing, Vysoke Tatry, Podbanske, Slovakia, March 13-18, 2016, Proceedings of ALGORITMY 2016, 20th Conference on Scientific Computing, Vysoke Tatry, Podbanske, Slovakia, March 13-18, 2016, 2016, pp. 13–22.
- [9] A. BUHR AND K. SMETANA, *Randomized Local Model Order Reduction*, SIAM J. Sci. Comput., 40 (2018), pp. A2120–A2151.
- [10] E. CANCÉS, V. EHRLACHER, AND T. LELIÉVRE, *Greedy algorithms for high-dimensional non-symmetric linear problems*, in CANUM 2012, 41e Congrès National d'Analyse Numérique, vol. 41 of ESAIM Proc., EDP Sci., Les Ulis, 2013, pp. 95–131.
- [11] Y. CAO AND L. PETZOLD, *A posteriori error estimation and global error control for ordinary differential equations by the adjoint method*, SIAM J. Sci. Comput., 26 (2004), pp. 359–374.
- [12] L. CHAMOIN, F. PLED, P.-E. ALLIER, AND P. LADEVÈZE, *A posteriori error estimation and adaptive strategy for pgd model reduction applied to parametrized linear parabolic problems*, Comput. Methods Appl. Mech. Engrg., 327 (2017), pp. 118–146.
- [13] Y. CHEN, J. S. HESTHAVEN, Y. MADAY, AND J. RODRÍGUEZ, *Improved successive constraint method based a posteriori error estimate for reduced basis approximation of 2D Maxwell's problem*, M2AN Math. Model. Numer. Anal., 43 (2009), pp. 1099–1116.
- [14] F. CHINESTA AND E. AMMAR, A. AND CUETO, *Recent advances and new challenges in the use of the proper generalized decomposition for solving multidimensional models*, Arch. Comput. Methods Eng., 17 (2010), pp. 327–350.
- [15] F. CHINESTA AND E. CUETO, *PGD-Based Modeling of Materials, Structures and Processes*, Springer International Publishing, 2014.
- [16] F. CHINESTA, R. KEUNINGS, AND A. LEYGUE, *The proper generalized decomposition for advanced numerical simulations. A primer*, SpringerBriefs in Applied Sciences and Technology, Springer, Cham, 2014.
- [17] F. CHINESTA, P. LADEVÈZE, AND E. CUETO, *A short review on model order reduction based on proper generalized decomposition*, Arch. Comput. Methods Eng., 18 (2011), pp. 395–404.

- [18] S. DASGUPTA AND A. GUPTA, *An elementary proof of the Johnson-Lindenstrauss lemma*, International Computer Science Institute, Technical Report, (1999), pp. 99–006.
- [19] M. H. DEGROOT AND M. J. SCHERVISH, *Probability and statistics. Fourth Edition*, Pearson Education, 2012.
- [20] Q. DU, V. FABER, AND M. GUNZBURGER, *Centroidal voronoi tessellations: Applications and algorithms*, SIAM Rev., 41 (1999), pp. 637–676.
- [21] M. EIGEL, R. SCHNEIDER, P. TRUNSCHKE, AND S. WOLF, *Variational monte carlo – bridging concepts of machine learning and high-dimensional partial differential equations*, Adv. Comput. Math., (2019).
- [22] A. FALCÓ, L. HILARIO, N. MONTÉS, AND M. C. MORA, *Numerical strategies for the Galerkin-proper generalized decomposition method*, Math. Comput. Modelling, 57 (2013), pp. 1694–1702.
- [23] L. GIRALDI, A. NOUY, G. LEGRAIN, AND P. CARTRAUD, *Tensor-based methods for numerical homogenization from high-resolution images*, Comput. Methods Appl. Mech. Engrg., 254 (2013), pp. 154–169.
- [24] B. HAASDONK, *Reduced basis methods for parametrized PDEs – a tutorial*, in Model Reduction and Approximation, P. Benner, A. Cohen, M. Ohlberger, and K. Willcox, eds., SIAM, Philadelphia, PA, 2017, pp. 65–136.
- [25] W. HACKBUSCH, *Tensor spaces and numerical tensor calculus*, Springer, Heidelberg, 2012.
- [26] J. HESTHAVEN, G. ROZZA, AND B. STAMM, *Certified reduced basis methods for parametrized partial differential equations*, Springer Briefs in Mathematics, Springer, Cham, 2016.
- [27] C. HOMESCU, L. PETZOLD, AND R. SERBAN, *Error Estimation for Reduced-Order Models of Dynamical Systems*, SIAM Rev., 49 (2007), pp. 277–299.
- [28] D. HUYNH, D. KNEZEVIC, Y. CHEN, J. HESTHAVEN, AND A. PATERA, *A natural-norm Successive Constraint Method for inf-sup lower bounds*, Comput. Methods Appl. Mech. Engrg., 199 (2010), pp. 1963 – 1975.
- [29] D. B. P. HUYNH, G. ROZZA, S. SEN, AND A. T. PATERA, *A successive constraint linear optimization method for lower bounds of parametric coercivity and inf-sup stability constants*, C. R. Math. Acad. Sci. Paris, 345 (2007), pp. 473–478.
- [30] A. JANON, M. NODET, AND C. PRIEUR, *Goal-oriented error estimation for the reduced basis method, with application to sensitivity analysis*, J. Sci. Comput., 68 (2016), pp. 21–41.
- [31] W. B. JOHNSON AND J. LINDENSTRAUSS, *Extensions of Lipschitz mappings into a Hilbert space*, Contemp. Math., 26 (1984), p. 1.
- [32] C. S. KENNEY AND A. J. LAUB, *Small-sample statistical condition estimates for general matrix functions*, SIAM J. Sci. Comput., 15 (1994), pp. 36–61.
- [33] K. KERGRENE, L. CHAMOIN, M. LAFOREST, AND S. PRUDHOMME, *On a goal-oriented version of the proper generalized decomposition method*, J. Sci. Comput., (2019).
- [34] P. LADEVÈZE AND L. CHAMOIN, *On the verification of model reduction methods based on the proper generalized decomposition*, Comput. Methods Appl. Mech. Engrg., 200 (2011), pp. 2032–2047.
- [35] ———, *Toward guaranteed PGD-reduced models. Bytes and Science*, CIMNE: Barcelona, 2012, pp. 143–154.
- [36] R. LAM, O. ZAHM, Y. MARZOUK, AND K. WILLCOX, *Multifidelity dimension reduction via active subspaces*, arXiv preprint arXiv:1809.05567, (2018).
- [37] C. LE BRIS, T. LELIÈVRE, AND Y. MADAY, *Results and questions on a nonlinear approximation approach for solving high-dimensional partial differential equations*, Constr. Approx., 30 (2009), pp. 621–651.
- [38] J. P. MOITINHO DE ALMEIDA, *A basis for bounding the errors of proper generalised decomposition solutions in solid mechanics*, Internat. J. Numer. Methods Engrg., 94 (2013), pp. 961–984.
- [39] E. NADAL, A. LEYGUE, F. CHINESTA, M. BERINGHIER, J. J. RÓDENAS, AND F. J. FUENMAYOR, *A separated representation of an error indicator for the mesh refinement process under the proper generalized decomposition framework*, Comput. Mech., 55 (2015), pp. 251–266.

- [40] A. NOUY, *A generalized spectral decomposition technique to solve a class of linear stochastic partial differential equations*, Comput. Methods Appl. Mech. Engrg., 196 (2007), pp. 4521–4537.
- [41] A. NOUY, *Low-rank methods for high-dimensional approximation and model order reduction*, in Model reduction and approximation, P. Benner, A. Cohen, M. Ohlberger, and K. Willcox, eds., SIAM, Philadelphia, PA, 2017, pp. 171–226.
- [42] M. OHLBERGER AND K. SMETANA, *A dimensional reduction approach based on the application of reduced basis methods in the framework of hierarchical model reduction*, SIAM J. Sci. Comput., 36 (2014), pp. A714–A736.
- [43] S. PEROTTO, A. ERN, AND A. VENEZIANI, *Hierarchical local model reduction for elliptic problems: a domain decomposition approach*, Multiscale Model. Simul., 8 (2010), pp. 1102–1127.
- [44] N. A. PIERCE AND M. B. GILES, *Adjoint recovery of superconvergent functionals from PDE approximations*, SIAM Rev., 42 (2000), pp. 247–264.
- [45] A. QUARTERONI, A. MANZONI, AND F. NEGRI, *Reduced basis methods for partial differential equations*, vol. 92, Springer, Cham, 2016.
- [46] G. ROZZA, D. HUYNH, AND A. PATERA, *Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations: application to transport and continuum mechanics*, Arch. Comput. Meth. Eng., 15 (2008), pp. 229–275.
- [47] K. SMETANA AND M. OHLBERGER, *Hierarchical model reduction of nonlinear partial differential equations based on the adaptive empirical projection method and reduced basis techniques*, ESAIM Math. Model. Numer. Anal., 51 (2017), pp. 641–677.
- [48] K. SMETANA, O. ZAHM, AND A. T. PATERA, *Randomized residual-based error estimators for parametrized equations*, SIAM J. Sci. Comput., 41 (2019), pp. A900–A926.
- [49] K. VEROY, C. PRUD’HOMME, D. ROVAS, AND A. PATERA, *A posteriori error bounds for reduced-basis approximation of parametrized noncoercive and nonlinear elliptic partial differential equations*, 16th AIAA Computational Fluid Dynamics Conference, (2003).
- [50] R. VERSHYNIN, *High-dimensional probability. An introduction with applications in data science*, vol. 47 of Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, Cambridge, 2018.
- [51] M. VOGELIUS AND I. BABUŠKA, *On a dimensional reduction method. I. The optimal selection of basis functions*, Math. Comp., 37 (1981), pp. 31–46.
- [52] O. ZAHM, M. BILLAUD-FRIESS, AND A. NOUY, *Projection-based model order reduction methods for the estimation of vector-valued variables of interest*, SIAM J. Sci. Comput., 39 (2017), pp. A1647–A1674.
- [53] O. ZAHM AND A. NOUY, *Interpolation of inverse operators for preconditioning parameter-dependent equations*, SIAM J. Sci. Comput., 38 (2016), pp. A1044–A1074.