Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Seminar for Applied Mathematics SAM

Sparse-Grid, Reduced-Basis Bayesian Inversion

P. Chen and Ch. Schwab

Research Report No. 2014-36
December 2014

Seminar für Angewandte Mathematik
Eidgenössische Technische Hochschule
CH-8092 Zürich
Switzerland

Sparse-Grid, Reduced-Basis Bayesian Inversion

Peng Chen and Christoph Schwab

Seminar für Angewandte Mathematik
Eidgenössische Technische Hochschule
CH-8092 Zürich, Switzerland

{peng.chen,christoph.schwab}@sam.math.ethz.ch

last update December 2, 2014

Abstract

We analyze reduced basis acceleration of recently proposed deterministic Bayesian inversion algorithms for partial differential equations with uncertain distributed parameter, for observation data subject to additive, Gaussian observation noise. Specifically, Bayesian inversion of affine-parametric, linear operator families on possibly high-dimensional parameter spaces. We consider "high-fidelity " Petrov-Galerkin (PG) discretizations of these countably-parametric operator families: we allow general families of inf-sup stable, PG Finite-Element methods, covering most conforming primal and mixed Finite-Element discretizations of standard problems in mechanics. Reduced basis acceleration of the high-dimensional, parametric forward response maps which need to be numerically solved numerous times in Bayesian inversion is proposed and convergence rate bounds for the error in the Bayesian estimate incurred by the use of reduced bases are derived. As consequence of recent theoretical results on dimension-independent sparsity of parametric responses, and preservation of sparsity for holomorphic-parametric problems, we establish new convergence rates of greedy reduced basis approximations for both, the parametric forward maps as well as for the countably-parametric posterior densities which arise in Bayesian inversion. We show that the convergence rates for the reduced basis approximations of the parametric forward maps as well as of the countably-parametric, deterministic Bayesian posterior densities are free from the curse of dimensionality and depend only on the sparsity of the uncertain input data. In particular, we establish the quadratic convergence of the reduced basis approximation for the posterior densities with respect to that for the parametric forward maps.

Numerical experiments for model elliptic, affine-parametric problems in two space dimensions with hundreds of parameters are reported which confirm that the proposed adaptive, deterministic reduced basis algorithms indeed exploit sparsity of both, the parametric forward maps as well as the Bayesian posterior density.

Key words: Parametric Operator Equations, Bayesian Inversion, Reduced Basis, Sparse Grid, A Posteriori Error Estimate, A Priori Error Estimate, Best N-term Convergence, Curse of Dimensionality

Acknowledgement: Work supported by ERC AdG 247277

Contents

1 Introduction

The problem of prediction and estimation of responses from differential (or integral) equation models in engineering, subject to possibly large sets of noisy data, has found considerable interest in recent years. For models with a moderate, finite number of uncertain parameters, theoretical and numerical aspects are well-investigated. We refer to [34] and the references there for recent contributions regarding computational aspects. In applications involving systems governed by PDEs, however, Bayesian inversion is often subject to noisy data from a finite number K of observable responses (being continuous functionals on spaces of suitable weak solutions), and subject to uncertain *distributed parameters* in the model under consideration. Examples could be soil permeability in subsurface flow or, in solid mechanics, elastic constitutive parameters in heterogeneous solids, uncertain bottom topography in shallow water equations, to name but a few. In these cases, the parametric uncertainty is *distributed* in space (and, possibly, also in time) which mandates in parsimonius representations a rather large number of parameters. The problem of efficient computational treatment of high-dimensional, multi-parametric forward maps in computational Bayesian inversion in Partial Differential Equations (PDEs for short) attracted considerable attention recently; we mention only [6, 7, 5, 44]. In these references, several computational methods for "extracting" parsimonious, low-parametric "principal modes and components" of complex computational models have been proposed and found to be efficient in a host of problems governed by (linear and nonlinear) elliptic and parabolic PDEs. In each case, the feasibility of such extraction has been shown to afford large computational savings in parameter calibration and estimation. To unify and mathematically justify these obervations, for a general class of abstract, affine-parametric operator equations with uncertain inputs, is a purpose of the present paper. Based on previous results on sparsity of solution families of affine-parametric operator equations from [20, 28, 52, 31, 29, 30, 29, 30], Reduced Basis Methods (RBM), or more general model reduction techniques, have been applied to solve Bayesian inverse problems recently in [39, 35, 25, 33, 21]. These methods are either limited to low dimensional parametric forward models or lack efficient algorithms to construct the reduced bases without taking advantage of the sparsity of the parametric forward solution and the posterior density function.

In the present paper, we propose and analyze RBM acceleration of a class of affine-parametric, linear operator equations and its use in Bayesian inversion, in particular in Bayesian estimation algorithms in the infinite-dimensional setting proposed in [55], and in the context of dimension-adaptive, deterministic quadrature algorithms proposed in [49, 50]. RM methods have been largely developed during the past decade; we refer to [37, 45, 42, 48]. Recently, RB methods have also been successfully applied to PDEs with random inputs, cp. [3, 24, 27]. A detailed computational comparison of RB and (generalized) sparse grid (SG for short) stochastic collocation methods in terms of accuracy versus computational cost has been given in [14]. It was found that the former perform superior for linear elliptic problems. Reduced basis methods (RBM for short) have been proposed to deal with arbitrary prior probability distributions in [13, 15]. They have bee applied, in combination with sparse grid methods (SGM for short), to efficiently solve stochastic optimal control problems [11, 12]. In the present work, we investigate the RBM to solve "many-query" Bayesian estimation problems for parametric forward problems in high dimensions. While affine parameter dependence may be considered as rather specific, the presently considered class of affine-parametric operator equations comprises in fact a wide range of elliptic and parabolic PDEs with distributed parameter uncertainty, being parametrized by a Karhunen–Loève (or principal components) expansion (see, eg., [51] and the references there).

We exploit parametric holomorphy of solution families of affine-parametric operators on polydiscs. The holomorphic dependence on the parameters is inherited by the countably-parametric posterior density in Bayesian inversion (cp. [53, 50]). As shown in [53, 50], this holomorphic dependence implies dimension-independent convergence rates of N-term truncated Taylor expansions for both, the parametric forward maps and for the parametric Bayesian posterior density. A comparison argument implies the same rates for greedy approximations and RBM of the parametric forward maps *and* for the Bayesian posterior density. In particular, we establish the quasi-optimality of the reduced basis (RB) Petrov-Galerkin approximation. Based on this result, the N-width convergence rate for the PG-RB approximation of the parametric forward solution follows by comparison with best N-term polynomial approximation. Moreover, by exploiting the Taylor expansion of the posterior density on the RB solution, and using a dual approach for a corrected posterior density with a dual-weighted

residual, we obtain a superconvergence bound of the N-width the RB approximation.

In the present paper, we estimate errors due to different approximations on the Bayesian estimates; specifically, error due to parametric dimension truncation and due to high-fidelity PG discretization, SG sampling and Smolyak quadrature and, most importantly, due to RB approximation of the high-dimensional, parametric solution manifold of the forward problem. An analysis for combined/total error for the approximation of the parametric forward solution and for the posterior density as well as of quadrature errors with respect to the prior measure are also provided. These error bounds also lead to a new error bound for the Hellinger distance between the exact and approximate Bayesian posterior, and in particular to dimension-independent convergence rates for the errors of these quantities (forward solution, posterior density and measure distance) by the combined SG-RB approximation.

We report here numerical experiments with the proposed algorithms showing that they realize N-term approximation error bounds already for few, but also for high (hundreds of) parametric dimensions. The results indicate that the high-fidelity error estimate and SG interpolation error estimate are sharp, while the error estimates for the RB approximation, and in particular for the SG integration are likely not optimal. We also investigate the dependence of the number of RB approximation with respect to the high-fidelity PG discretization, the noise realization and scale, the locations and number of sensors/observations. In general, the number of reduced bases constructed from the dual-weighted residual based a posteriori error estimate are rather robust, remaining the same when refining the high-fidelity PG discretization when it is fine enough, displaying little sensitivity to the noise samples and small increase with decreasing noise scale, exhibiting small sensitivity with respect to the number K of sensors.

The structure of this paper is as follows. In Section 2, we present the problem of Bayesian inversion of parametric operator equations in a function space setting, based on [55, 22], and for the general class of affine-parametric operator equations considered here. Section 3 presents the results on adaptive SG interpolation and integration, from [10]. Section 4 summarizes the principal results and algorithms from the reduced basis method. Section 5 combines the basic error bounds, in particular the dimension-independent RB convergence rates of convergence for the RB approximation of the forward problems and for the Bayesian posterior density. Section 6 contains a suite of numerical experiments for a model parametric problem class which confirm the theoretical results. Section 7 contains some conclusions, and also indicate a number of generalizations, most notably to problems which do not allow an affine-parametric representation and nonlinear problems.

2 Bayesian Inversion

We define a class of operator equations which depend on an uncertain datum u taking values in a separable Banach space X (such as spaces for permeability), via a possibly countably infinite sequence $\boldsymbol{y} = (y_j)_{j\in\mathbb{J}}$ of parameters, and formulate theorems on Bayesian estimation for responses of these equations, with respect to a uniform, in a sense, prior measure π_0 on a suitable subset $X_+ \subset X$, conditional to noisy observation data $\delta \in Y$. Ie., the data δ is subject to additive, Gaussian observation noise $\eta \sim N(0,\Gamma)$ on Y; we assume that Y is finite-dimensional and that the covariance operator Γ is nondegenerate, ie. $Y \subseteq \mathbb{R}^K$ for some $K < \infty$, corresponding to, say, noisy data from K sensors.

2.1 Formulation of Bayesian Estimation

We review Bayesian Inversion, from [22, Sections 3 and 5] and [55], in the setting considered here. By $G : X \to \mathcal{X}$ we denote a "forward" response map from a separable Banach space X of uncertain data u in the operator A into a Banach space $\mathcal{X}$ of responses q which are accessible through a finite number K of bounded, linear functionals $\mathcal{O}() = (o_1, ..., o_K)^\top \in (\mathcal{X}')^K$.

The goal of Bayesian estimation is to predict a "most likely" value of a Quantity of Interest (QoI) $\Psi : \mathcal{X} \to \mathcal{Z}$ taking values in a Banach space $\mathcal{Z}$, conditional on given, noisy measurement data δ.

We consider forward models described by *linear operator equations*

$$\text{Given } u \in X,\ f \in \mathcal{Y}' \text{ find } q \in \mathcal{X} : \quad A(u)q = f \quad \text{in} \quad \mathcal{Y}' , \tag{2.1}$$

where the uncertain operator $A(u) \in \mathcal{L}(\mathcal{X}, \mathcal{Y}')$ is assumed to be boundedly invertible, uniformly w.r.

to uncertain input $u \in X_+ \subseteq X$ sufficiently close to a nominal input $\langle u \rangle \in X_+$, ie. for $\|u - \langle u \rangle\|_X$ small enough and a suitable, closed subset X_+ of X. Then, for such u, the response of (2.1) is

$$X_+ \ni u \mapsto q(u) := G(u; f) = (A(u))^{-1} f \in \mathcal{X} \ . \tag{2.2}$$

As we assume the forcing $f \in \mathcal{Y}'$ to be known, we omit the dependence of the response on f and simply write $q = G(u)$. We also assume given an observation functional $\mathcal{O}(\cdot) : \mathcal{X} \to \mathbb{R}^K$ denoting a *bounded linear observation operator* on the space $\mathcal{X}$ of system responses, i.e. $\mathcal{O} = (o_1, \ldots, o_K)^\top \in (\mathcal{X}')^K$, the dual space of the space $\mathcal{X}$ of system responses. We assume that the number of observations is finite so that $K < \infty$. Then $Y = \mathbb{R}^K$ equipped with the Euclidean norm, denoted by $|\cdot|$.

In this setting, we wish to predict *computationally* an expected (under the Bayesian posterior) system response of the QoI Ψ, conditional on measurement data $\delta \in Y$ corrupted by additive, gaussian observation noise η. Ie., we assume $\delta \in Y$ consists of exact system responses (for known, given $f \in \mathcal{Y}'$ and for some realization of the uncertain input $u \in X$) plus *additive, gaussian noise,*

$$\delta = \mathcal{O}(G(u)) + \eta \in Y = \mathbb{R}^K \tag{2.3}$$

where $\eta \sim \mathcal{N}(0, \Gamma)$, for a positive definite covariance operator Γ on $\mathbb{R}^K$ (ie., a symmetric, positive definite $K \times K$ covariance matrix Γ). With the *uncertainty-to-observation map* $\mathcal{G} : X \to \mathbb{R}^K : u \mapsto \mathcal{G}(u) = \mathcal{O}(G(u))$, it holds

$$\delta = \mathcal{G}(u) + \eta = (\mathcal{O} \circ G)(u) + \eta : X \mapsto L^2_\Gamma(\mathbb{R}^K)$$

where $L^2_\Gamma(\mathbb{R}^K)$ denotes random vectors taking values in $\mathbb{R}^K$ which are square integrable with respect to the Gaussian measure on $\mathbb{R}^K$. In view of Bayes' formula, we define the *observation noise covariance weighted least squares functional* (also referred to as "potential" in what follows) $\Phi : X \times Y \to \mathbb{R}$ by $\Phi(u; \delta) = \frac{1}{2}|\delta - \mathcal{G}(u)|^2_\Gamma$ where $|\cdot|_\Gamma = |\Gamma^{-\frac{1}{2}} \cdot|$. For given $u \in X$, the Bayesian potential takes the form

$$\Phi(u; \delta) = \frac{1}{2} \left((\delta - (\mathcal{O} \circ G)(u))^\top \Gamma^{-1} (\delta - (\mathcal{O} \circ G)(u)) \right) \ . \tag{2.4}$$

In [55] an infinite-dimensional version of the Bayes rule is shown to hold in the present setting. It states that, under appropriate continuity conditions on the uncertainty-to-observation map $\mathcal{G} = (\mathcal{O} \circ G)(\cdot)$ and given the prior measure π_0 on X, the posterior distribution π^δ of u, given data $\delta \in Y$, is absolutely continuous with respect to the prior π_0. In particular, then, the Radon-Nikodym derivative of the Bayesian posterior w.r. to the prior measure admits a bounded density $\Theta(u)$ w.r. to the prior π_0.

2.2 Uncertainty Parametrization

We parametrize the uncertain datum u in the forward equation (2.1) and in the posterior density Θ. In the context of PDEs, often the case where $u \in X$, a separable Banach space is of interest. We assume that there exists a countable, unconditional base $(\psi_j)_{j \in \mathbb{J}}$ of X such that, for some "nominal" value $\langle u \rangle \in X$ of the uncertain datum u, and for a random coefficient sequence $\boldsymbol{y} = (y_j)_{j \in \mathbb{J}}$ (depending on $u - \langle u \rangle \in X$) the uncertainty u is parametrized by $\boldsymbol{y}$ in the sense that

$$u = u(\boldsymbol{y}) := \langle u \rangle + \sum_{j \in \mathbb{J}} y_j \psi_j \in X \tag{2.5}$$

with unconditional convergence. We refer to $u - \langle u \rangle$ as "fluctuation" of u about the nominal value $\langle u \rangle \in X$; examples of (2.5) are Karhunen–Loève expansions.

The assumption (2.5) on affine parametrization of the distributed system uncertainty by the sequence $\boldsymbol{y} = (y_j)_{j \in \mathbb{J}}$ of (possibly countably many) parameters results in a parametric operator equation of the form

$$A(\boldsymbol{y}) = A_0 + \sum_{j \in \mathbb{J}} y_j A_j \in \mathcal{L}(\mathcal{X}, \mathcal{Y}') \ . \tag{2.6}$$

With the affine-parametric, linear operator family $A(\boldsymbol{y})$ we associate the parametric bilinear form

$\mathfrak{a}(\boldsymbol{y};\cdot,\cdot) : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ via

$$\mathfrak{a}(\boldsymbol{y}; w, v) :=_{\mathcal{Y}} \langle v, A(\boldsymbol{y}) w \rangle_{\mathcal{Y}'} , \qquad w \in \mathcal{X},\ v \in \mathcal{Y} . \tag{2.7}$$

Many choices for the functions ψ_j in (2.5) are conceivable; among them are standard spline bases, but also Karhunen–Loève eigenfunctions. With y_j denoting the coordinate variables, the parametrization (2.5) is deterministic. In order to place (2.3), (2.5) into the (probabilistic) Bayesian setting of [55], we introduce (after possibly rescaling the fluctuations) a "reference" parameter domain $U = [-1,1]^{\mathbb{J}} = \prod_{j\in\mathbb{J}}[-1,1]$, and equip this countable cartesian product of sets with the product sigma-algebra $\mathcal{B} = \bigotimes_{j\in\mathbb{J}} \mathcal{B}^1$, with $\mathcal{B}^1$ the sigma-algebra of Borel sets on $[-1,1]$. On the measurable space $(U,\mathcal{B})$ we introduce a probability measure π_0 (which will serve a Bayesian prior in what follows), and which we shall choose as $\pi_0 = \bigotimes_{j\in\mathbb{J}} \frac{1}{2}\lambda^1$ with λ^1 denoting the Lebesgue measure on $[-1,1]$. More general measures other than the uniform type can be efficiently treated by a weighted RB along the lines of [13]. Then, $(U,\mathcal{B},\pi_0)$ becomes (as countable product of probability spaces) a probability space on the set U of all sequences of coefficient vectors $\boldsymbol{y}$. Then the uncertain datum u in (2.5) becomes a random field, with π_0 charging the possible realizations of u. As indicated in [20, 49, 18], *analyticity of uncertainty parametrization* (2.5) *with respect to the parameter sequence* $\boldsymbol{y}$ can be used to derive sparsity results for this posterior.

2.3 Holomorphy of Parametric Forward Solutions

Analytic dependence of responses on the components y_j of the parameter $\boldsymbol{y} \in U$ plays an important role for polynomial approximation results, as well as for the sparsity of the Bayesian posterior. To state it, we denote by $\mathcal{D}_r = \{z \in \mathbb{C} : |z| \le r\}$ the closed disc in $\mathbb{C}$ of radius $r \ge 1$. All spaces in this section will be understood as Banach spaces over the coefficient field $\mathbb{C}$ without explicitly indicating so notationally.

Definition 2.1 *Given a summability exponent $0 < p < 1$ and a real number $\varepsilon > 0$, we say that the parametric family $\{q(\boldsymbol{y}) : \boldsymbol{y} \in U\} \subset \mathcal{X}$ is (p,ε)-analytic if*

$(p,\varepsilon):1$ *(well-posedness of the forward problem)*
for each $\boldsymbol{y} \in U$, there exists a unique realization $u(\boldsymbol{y}) \in X$ of the uncertainty and a unique solution $q(\boldsymbol{y}) \in \mathcal{X}$ of the forward problem (2.1)*. The forward problem is well-posed, uniformly w.r. to the parameter $\boldsymbol{y}$, i.e.*

$$\forall \boldsymbol{y} \in U : \quad \|q(\boldsymbol{y})\|_{\mathcal{X}} \le C_0 . \tag{2.8}$$

$(p,\varepsilon):2$ *(analyticity)*
There exists $0 < p < 1$ and a sequence $b = (b_j)_{j\in\mathbb{J}} \in \ell^p(\mathbb{J})$ such that for sufficiently small $0 < \varepsilon \le 1$ and for every $(\boldsymbol{b},\varepsilon)$-admissible sequence $\boldsymbol{\rho} = (\rho_j)_{j\in\mathbb{J}}$ of poly-radii $\rho_j > 1$, ie. such that

$$\sum_{j\in\mathbb{J}} (\rho_j - 1) b_j \le \varepsilon , \tag{2.9}$$

and such that solution map $U \ni \boldsymbol{y} \mapsto q(\boldsymbol{y}) \in \mathcal{X}$ admits an analytic continuation to the open polydisc $\mathcal{D}_\rho := \prod_{j\in\mathbb{J}} \mathcal{D}_{\rho_j} \subset \mathbb{C}^{\mathbb{J}}$ and satisfies the bound

$$\forall \boldsymbol{z} \in \mathcal{D}_\rho : \quad \|q(\boldsymbol{z})\|_{\mathcal{X}} \le C_\varepsilon \tag{2.10}$$

where $\boldsymbol{y} := \Re(\boldsymbol{z}) \in \bigotimes_{j\in\mathbb{J}}[-\rho_j,\rho_j] \subset \mathbb{R}^{\mathbb{J}}$.

We refer to ρ as *poly-radius*. In the case when $\rho_j = 1$ for all $j \in \mathbb{J}$, we simply write $\mathcal{D}$ in place of $\mathcal{D}_\rho$ to denote the unit disc in $\mathbb{C}^{\mathbb{J}}$. We observe $U \subset \mathcal{D}_\rho$ for $\rho_j > 1$. All assertions proved in the sequel hold in either case, and all constants are, in particular, independent of $J = \#(\mathbb{J})$.

In a stochastic interpretation of (2.6) in the context of Bayesian estimation, the parameter sequence $\boldsymbol{y} = (y_j)_{j\in\mathbb{J}}$ is assumed to be an i.i.d sequence of real-valued random variables $y_j \sim \mathcal{U}(-1,1)$, A_0 denotes a "nominal operator" (representing the non-perturbed system) and the sequence $(A_j)_{j\in\mathbb{J}} \subset \mathcal{L}(\mathcal{X},\mathcal{Y}')$ denotes a sequence of "fluctuations" about the "nominal operator" $A_0 = A(0)$. Affine parameter dependences (2.6) result for example when u in (2.5) is modelled as random field via its Karhunen–Loève expansion in X (or in a closed subspace of X).

In order for the sum in (2.6) to converge, we impose the following assumptions on the sequence $\{A_j\}_{j\geq 0} \subset \mathcal{L}(\mathcal{X}, \mathcal{Y}')$. In doing so, we associate with $A_j \in \mathcal{L}(\mathcal{X}, \mathcal{Y}')$ the bilinear forms $\mathfrak{a}_j(\cdot,\cdot) : \mathcal{X}\times\mathcal{Y} \to \mathbb{R}$ via

$$\forall w \in \mathcal{X},\ v \in \mathcal{Y}: \quad \mathfrak{a}_j(w,v) =_{\mathcal{Y}} \langle v, A_j w\rangle_{\mathcal{Y}'}\ , \quad j = 0,1,2....$$

Assumption 1 *The operator family $\{A_j\}_{j\geq 0} \in \mathcal{L}(\mathcal{X}, \mathcal{Y}')$ in (2.6) satisfies:*

1. *The "nominal" or "mean field" operator $A_0 \in \mathcal{L}(\mathcal{X}, \mathcal{Y}')$ is boundedly invertible, i.e. there exists $\beta_0 > 0$ such that*

$$\textbf{\textit{A1}} \quad \inf_{0\neq w\in\mathcal{X}} \sup_{0\neq v\in\mathcal{Y}} \frac{\mathfrak{a}_0(w,v)}{\|w\|_{\mathcal{X}}\|v\|_{\mathcal{Y}}} \geq \beta_0\ , \ \inf_{0\neq v\in\mathcal{Y}} \sup_{0\neq w\in\mathcal{X}} \frac{\mathfrak{a}_0(w,v)}{\|w\|_{\mathcal{X}}\|v\|_{\mathcal{Y}}} \geq \beta_0\ . \tag{2.11}$$

2. *The "fluctuation" operators $\{A_j\}_{j\geq 1}$ are small with respect to A_0 in the following sense: there exists a constant $0 < \kappa < 1$ such that*

$$\textbf{\textit{A2}} \quad \sum_{j\in\mathbb{J}} b_j \leq \kappa < 1\ , \quad \textit{where} \quad b_j := \|A_0^{-1}A_j\|_{\mathcal{L}(\mathcal{X},\mathcal{X})}\ , \quad j = 1,2,...\ . \tag{2.12}$$

3. *(p-summability) For some $0 < p < 1$, the operators $B_j := A_0^{-1}A_j$, $j = 1,2,...$, are p-summable, in the sense that with the sequence $\boldsymbol{b} := (b_j)_{j\in\mathbb{J}}$ as in (2.12) holds*

$$\textbf{\textit{A3}} \quad \|b\|^p_{\ell^p(\mathbb{J})} = \sum_{j\in\mathbb{J}} b_j^p < \infty\ . \tag{2.13}$$

Condition (2.12) (and, hence, Assumption 1) is sufficient for the bounded invertibility of $A(\boldsymbol{y})$, uniformly with respect to the parameter sequence $\boldsymbol{y} \in U = [-1,1]^{\mathbb{J}}$.

Theorem 2.1 *Under Assumption 1, for every realization $\boldsymbol{y} \in U$ of the parameters, the affine parametric operator family $A(\boldsymbol{y})$ is boundedly invertible, uniformly with respect to the parameter sequence $\boldsymbol{y} \in U$. In particular, for the parametric bilinear form $\mathfrak{a}(\boldsymbol{y};\cdot,\cdot) : \mathcal{X}\times\mathcal{Y} \to \mathbb{R}$ associated with $A(\boldsymbol{y}) \in \mathcal{L}(\mathcal{X}, \mathcal{Y}')$ via (2.7), there hold the uniform inf-sup conditions with $\beta = (1-\kappa)\beta_0 > 0$,*

$$\forall \boldsymbol{y} \in U: \quad \inf_{0\neq w\in\mathcal{X}} \sup_{0\neq v\in\mathcal{Y}} \frac{\mathfrak{a}(\boldsymbol{y};w,v)}{\|w\|_{\mathcal{X}}\|v\|_{\mathcal{Y}}} \geq \beta\ , \ \inf_{0\neq v\in\mathcal{Y}} \sup_{0\neq w\in\mathcal{X}} \frac{\mathfrak{a}(\boldsymbol{y};w,v)}{\|w\|_{\mathcal{X}}\|v\|_{\mathcal{Y}}} \geq \beta\ . \tag{2.14}$$

In particular, for every $f \in \mathcal{Y}'$ and for every $\boldsymbol{y} \in U$, the parametric operator equation

$$\textit{find } q(\boldsymbol{y}) \in \mathcal{X}: \quad \mathfrak{a}(\boldsymbol{y}; q(\boldsymbol{y}), v) = f(v) \quad \forall v \in \mathcal{Y} \tag{2.15}$$

admits a unique solution $q(\boldsymbol{y}) = (A(\boldsymbol{y}))^{-1}f$ which is uniformly bounded over U, ie.

$$\sup_{\boldsymbol{y}\in U} \|q(\boldsymbol{y})\|_{\mathcal{X}} \leq \frac{\|f\|_{\mathcal{Y}'}}{\beta}\ . \tag{2.16}$$

2.4 Parametric Bayesian posterior

The presently proposed, adaptive deterministic quadrature, and reduced basis approaches for Bayesian estimation via the computational realization of Bayes' formula is a *parametric, deterministic representation* of the derivative of the posterior measure with respect to the prior measure π_0. The prior measure π_0 being uniform in the present paper, we admit in (2.6) sequences $\boldsymbol{y}$ which take values in the parameter domain $U = [-1,1]^{\mathbb{J}}$. As explained in Section 2.2, we consider the parametric, deterministic forward problem (2.1), (2.5) in the probability space

$$(U, \mathcal{B}, \pi_0)\ . \tag{2.17}$$

We assume throughout what follows that the prior measure on the uncertain input data, parametrized in the form (2.6)*, is the uniform measure* $\pi_0(d\boldsymbol{y})$*.* With the parameter domain U as in (2.17) the parametric forward map is given by

$$\mathcal{G}(\boldsymbol{y}) := \mathcal{G}(u)\Big|_{u=\langle u\rangle+\sum_{j\in\mathbb{J}} y_j\psi_j} . \tag{2.18}$$

We abuse notation here by denoting this map by $\mathcal{G}()$ in either case, with the definition (2.18) depending implicitly on the choice of basis ψ_j. The mathematical foundation of Bayesian inversion in the present setting is a parametric version of Bayes' formula from [53] (see also [55]).

Theorem 2.2 *Under the above assumptions, in the parametrization* (2.5) *the posterior distribution* π^δ *of* $u \in X$*, given data* $\delta \in Y$*, is absolutely continuous with respect to the prior* π_0*, ie.*

$$\frac{d\pi^\delta}{d\pi_0}(\boldsymbol{y}) = \frac{1}{Z}\Theta(\boldsymbol{y}) \tag{2.19}$$

with the parametric Bayesian posterior $\Theta(\boldsymbol{y})$ *given by*

$$\Theta(\boldsymbol{y}) = \exp\big(-\Phi(u;\delta)\big)\Big|_{u=\langle u\rangle+\sum_{j\in\mathbb{J}} y_j\psi_j}, \tag{2.20}$$

where the Bayesian potential Φ *is as in* (2.4) *and the normalization constant* Z *is given by*

$$Z = \mathbb{E}^{\pi_0}\left[\Theta(\boldsymbol{y})\right] = \int_U \Theta(\boldsymbol{y}) d\pi_0(\boldsymbol{y}) . \tag{2.21}$$

Computational Bayesian inversion is concerned with approximation of a "most likely" *system response* $\phi : X \to \mathcal{Z}$ (sometimes also referred to as *Quantity of Interest (QoI)* which may take values in a Banach space $\mathcal{Z}$) for given (noisy) observation data δ of the QoI ϕ. In particular the choices $\phi(u) = G(u)$ (with $\mathcal{Z} = \mathcal{X}$) and $\phi(u) = G(u) \otimes G(u)$ (with $\mathcal{Z} = \mathcal{X} \otimes \mathcal{X}$) facilitate computation of the "most likely" (given the data δ) mean and covariance of the system's response.

With the QoI ϕ we associate the (infinite-dimensional) parametric map

$$\Psi(\boldsymbol{y}) \quad = \quad \Theta(\boldsymbol{y})\phi(u)\,|_{u=\langle u\rangle+\sum_{j\in\mathbb{J}} y_j\psi_j} = \exp\big(-\Phi(u;\delta)\big)\phi(u)\Big|_{u=\langle u\rangle+\sum_{j\in\mathbb{J}} y_j\psi_j} : U \to \mathcal{Z} . \tag{2.22}$$

Then the Bayesian estimate of the QoI ϕ, given noisy data δ (2.3), takes the form

$$\mathbb{E}^{\pi^\delta}[\phi] = \int_{\boldsymbol{y}\in U} \Psi(\boldsymbol{y})\pi^\delta(d\boldsymbol{y}) = \frac{1}{Z}\int_{\boldsymbol{y}\in U} \exp\big(-\Phi(u;\delta)\big)\phi(u)\Big|_{u=\langle u\rangle+\sum_{j\in\mathbb{J}} y_j\psi_j} \pi_0(d\boldsymbol{y}) . \tag{2.23}$$

Our aim is to approximate the expectations Z' and Z which, in the parametrization with respect to $\boldsymbol{y} \in U$, take the form of infinite-dimensional integrals with respect to the prior $\pi_0(d\boldsymbol{y})$.

Quadrature, or MC approximation of the Bayesian posterior distribution to estimate any other quantity of interest based on this posterior distribution, amounts to computing the pointwise densities $\Theta(\boldsymbol{y})$ and $\Psi(\boldsymbol{y})$ for $\boldsymbol{y} \in U$, and their expectations Z' and Z under the prior π_0. To estimate the computational accuracy, we proceed as in [53, 50].

Proposition 2.3 *Under Assumption 1, the parametric forward solution* $q(\boldsymbol{y})$ *with* $\boldsymbol{y} \in U$ *is* $(\boldsymbol{b}, p, \epsilon)$*-analytic in the polydisc* $\mathcal{D}_{\boldsymbol{\rho}}$ *for any* $(\boldsymbol{b}, \epsilon)$*-admissible poly-radius* $\boldsymbol{\rho}$*. In particular,* $q(\boldsymbol{y})$ *admits approximations* $q_{\Lambda_N}(\boldsymbol{y})$*, being* N*-term truncated series of tensorized Legendre polynomials such that*

$$\sup_{\boldsymbol{y}\in U} \|q(\boldsymbol{y}) - q_{\Lambda_N}(\boldsymbol{y})\|_{\mathcal{X}} \le C_q N^{-s} , \quad s = 1/p - 1 . \tag{2.24}$$

Here, the summation exponent p *is as in* (2.13) *and the sets* Λ_N *are collections of* N *finitely supported multi-indices indicating the active polynomial degrees in the approximation* $q_{\Lambda_N}(\boldsymbol{y})$*. The sets* Λ_N *can be chosen nested and 'downward closed ', as defined in [50]. The constant* $C_q > 0$ *is independent of the number of dimensions activated in* Λ_N*. Similarly,* $\Theta(\boldsymbol{y})$ *admits gpc approximations* $\Theta_{\Lambda_N}(\boldsymbol{y})$ *such*

that

$$\sup_{\boldsymbol{y}\in U} |\Theta(\boldsymbol{y}) - \Theta_{\Lambda_N}(\boldsymbol{y})| \leq C_\Theta N^{-s} , \quad s = 1/p - 1 , \tag{2.25}$$

where Λ_N could be different for q and Θ; the constant C_Θ does not depend on N but depends on the additive, gaussian observation noise covariance Γ as $\exp(b/\Gamma)$ for some $b > 0$. An analogous result holds also for the parametric QoI Ψ.

3 Dimension-adaptive sparse grid approximation

In this section, we introduce the dimension-adaptive SG techniques in order to evaluate efficiently the high-dimensional parametric response quantity; we draw upon recent work [10].

3.1 Univariate interpolation and integration

Let $\mathcal{I}$ denote a univariate interpolation operator associated with a sequence of interpolation nodes $y_1, \ldots, y_m \in U$ (in the univariate setting, ie., $U = [-1, 1]$) and corresponding Lagrange basis functions $l_1, \ldots, l_m : U \to \mathbb{R}$, $m \in \mathbb{N}_0$, i.e.,

$$\mathcal{I}g = \sum_{k=1}^{m} g(y_k) l_k(y), \tag{3.1}$$

where the function $g : U \to \mathcal{Z}$, represents, e.g. the parametric solution $q : U \to \mathcal{X}$ ($\mathcal{Z} = \mathcal{X}$), or the parametric Bayesian posterior $\Theta : U \to \mathbb{R}$ ($\mathcal{Z} = \mathbb{R}$). A particular choice for the basis functions and nodes are Lagrange polynomials based on the extrema of Chebyshev polynomials (see, eg., [47]):

$$l_k(y) = \prod_{l=1, l\neq k} \frac{y - y_l}{y_k - y_l}, \text{ where } y_k = \cos\left(\frac{k-1}{m-1}\pi\right), \ k = 1, \ldots, m . \tag{3.2}$$

By $i \in \mathbb{N}_+$ we denote the grid level, and by Ξ^i the corresponding set of nodes and by m^i the number of nodes on the grid of level i. In (3.2), we consider nested sets of nodes, i.e. $\Xi^i \subset \Xi^{i+1}$, $i = 1, \ldots, q$ for some $q \in \mathbb{N}_+$, and therefore do not indicate dependence on m in (3.2). Specifically, we choose Clenshaw–Curtis nodes, ie., $m^1 = 1$, $m^i = 2^{i-1} + 1$ for $i = 2, \ldots, q$. We point out, however, that all algorithms and error bounds which follow hold also for other nested sets of nodes, such as Leja nodes or symmetric Leja nodes [16], Gauss-Kronrod/Patterson nodes. Under such specification, the interpolation operator (3.1) can be rewritten in a hierarchical way indexed by q as

$$\mathcal{I}_q g = \sum_{i=1}^{q} \triangle^i g, \tag{3.3}$$

where $\triangle^i$ denotes an interpolation difference operator $\triangle^i = \mathcal{I}_i - \mathcal{I}_{i-1}$ for $i = 1, \ldots, q$, being $\mathcal{I}_0 g \equiv 0$. For ease of notation, we denote $\Xi^i_\triangle = \Xi^i \setminus \Xi^{i-1}$ for $i = 1, \ldots, q$ with $\Xi^0 = \emptyset$, and denote the nodes $y^i_k \in \Xi^i_\triangle$ and the corresponding basis functions l^i_k for $k \in m^i_\triangle$, where $m^i_\triangle = \{1, \ldots, m^i - m^{i-1}\}$ with $m^0 = 0$. Thanks to the nested structure $\Xi^{i-1} \subset \Xi^i$, we have $\mathcal{I}_{i-1} = \mathcal{I}_i \circ \mathcal{I}_{i-1}$. Moreover, $g(y^i_k) = \mathcal{I}_{i-1} g$ for $y^i_k \in \Xi^{i-1}$. Therefore, the interpolation operator (3.3) can be written more explicitly as

$$\mathcal{I}_q g = \sum_{i=1}^{q} (\mathcal{I}_i g - \mathcal{I}_i \circ \mathcal{I}_{i-1} g) = \sum_{i=1}^{q} \sum_{y^i_k \in \Xi^i_\triangle} \underbrace{(g(y^i_k) - \mathcal{I}_{i-1} g(y^i_k))}_{s^i_k} l^i_k(y), \tag{3.4}$$

where s^i_k denotes the so-called *hierarchical surplus* [8] that represents the interpolation error at the nodes in $\Xi^i_\triangle$. The univariate integration of g can be evaluated as

$$\mathbb{E}[g] \approx \mathbb{E}[\mathcal{I}_q g] = \sum_{i=1}^{q} \sum_{k \in m^i_\triangle} s^i_k w^i_k, \tag{3.5}$$

where the weights w_k^i can be computed using suitable quadrature rule, e.g. Clenshaw–Curtis, as

$$w_k^i = \int_U l_k^i(y)\pi_0(dy), \quad i = 1, \ldots, q,\ k \in m_\triangle^i. \tag{3.6}$$

Based on the hierarchical surplus s_k^i, we define the interpolation error indicator $\mathcal{E}_i$ and quadrature error indicator $\mathcal{E}_e$ for adaptive construction of the SG interpolation and Smolyak quadrature formula, respectively:

$$\mathcal{E}_i := \max_{k \in m_\triangle^i} s_k^q \text{ and } \mathcal{E}_i := \sum_{k \in m_\triangle^i} s_k^q w_k^q\ . \tag{3.7}$$

Remark 3.1 *Thanks to the nested structure of the interpolation/collocation nodes, the interpolation operator defined* (3.3) *and the associated integration operator defined in* (3.5) *can be efficiently evaluated in a hierarchical manner. Moreover, the error estimates* (3.7) *can be effectively used for the construction of the interpolation and integration operators. Besides the Clenshaw-Curtis nodes given in* (3.2)*, other feasible choices of nested nodes include real projection of Leja nodes [16] and Gauss–Kronrod nodes [43], etc.*

3.2 Smolyak sparse grid interpolation

In J dimensions, $J \in \mathbb{N}_+$, the Smolyak SG interpolation operator is defined as [54]

$$\mathcal{S}_q g = \sum_{|\boldsymbol{i}| \le q} \left(\triangle_1^{i_1} \otimes \cdots \otimes \triangle_J^{i_J}\right) g, \tag{3.8}$$

where the multi-index $\boldsymbol{i} = (i_1, \ldots, i_J) \in \mathbb{N}_+^J$ represents a multi-dimensional grid level; $\triangle_j^{i_j}$ is the interpolation difference operator at grid level i_j in dimension j; $q \ge J$ represents the total level of the sparse grid. The hierarchical structure of the sparse grid allows us to rewrite $\mathcal{S}_q$ as

$$\mathcal{S}_q = \mathcal{S}_{q-1} + \triangle\mathcal{S}_q \text{ where } \triangle\mathcal{S}_q = \mathcal{S}_q - \mathcal{S}_{q-1} = \sum_{|\boldsymbol{i}|=q} \left(\triangle_1^{i_1} \otimes \cdots \otimes \triangle_J^{i_J}\right)\ . \tag{3.9}$$

Thanks to the the relation $\mathcal{S}_{q-1} = \mathcal{S}_q \circ \mathcal{S}_{q-1}$ on $\Xi_\triangle^{\boldsymbol{i}} = \Xi_\triangle^{i_1} \times \cdots \times \Xi_\triangle^{i_J}$ with $|\boldsymbol{i}| = q$, we have

$$\triangle\mathcal{S}_q g(\boldsymbol{y}) = \sum_{|\boldsymbol{i}|=q} \sum_{\boldsymbol{k} \in m_\triangle^{\boldsymbol{i}}} \underbrace{\left(g(y_{k_1}^{i_1}, \ldots, y_{k_J}^{i_J}) - \mathcal{S}_{q-1} g(y_{k_1}^{i_1}, \ldots, y_{k_J}^{i_J})\right)}_{s_{\boldsymbol{k}}^{\boldsymbol{i}}} \underbrace{\left(l_{k_1}^{i_1}(y_1) \otimes \cdots \otimes l_{k_J}^{i_J}(y_J)\right)}_{l_{\boldsymbol{k}}^{\boldsymbol{i}}}, \tag{3.10}$$

where the set of index $m_\triangle^{\boldsymbol{i}} = m_\triangle^{i_1} \times \cdots \times m_\triangle^{i_J}$. Consequently, the interpolation operator $\mathcal{S}_q$ permits the following hierarchical representation

$$\mathcal{S}_q g = \sum_{|\boldsymbol{i}| \le q} \sum_{\boldsymbol{k} \in m_\triangle^{\boldsymbol{i}}} s_{\boldsymbol{k}}^{\boldsymbol{i}} l_{\boldsymbol{k}}^{\boldsymbol{i}}. \tag{3.11}$$

The integration of g can thus be approximated by the SG quadrature, the exact integral over the sparse interpolation

$$\mathbb{E}[g] \approx \mathbb{E}[\mathcal{S}_q g] = \sum_{|\boldsymbol{i}| \le q} \sum_{\boldsymbol{k} \in m_\triangle^{\boldsymbol{i}}} s_{\boldsymbol{k}}^{\boldsymbol{i}} w_{\boldsymbol{k}}^{\boldsymbol{i}}, \tag{3.12}$$

where the weight $w_{\boldsymbol{k}}^{\boldsymbol{i}} = \prod_{j=1}^J w_{k_j}^{i_j}$, with $w_{k_j}^{i_j} = \int_{U_j} l_{k_j}^{i_j}(y_j)\pi_0(dy_j)$ given in closed form.

3.3 Dimension-adaptive sparse grid approximation

As different y_j, $j = 1, \ldots, J$, may have very different influence on the output of interest, we relax the total degree structure $|\boldsymbol{i}| \le q$ of the Smolyak SG construction introduced in the last section and adopt a dimension-adaptive construction for a generalized sparse grid [26]. To start, we recall the notion

of downward closed index set: let $\boldsymbol{e}_j \in \{0,1\}^J$, whose j-th element is 1 and all the other elements are 0; a set Λ of indices is called downward closed if $\boldsymbol{i} \in \Lambda$ then $\boldsymbol{i} - \boldsymbol{e}_j \in \Lambda$ for each $j = 1, \ldots, J$ such that $i_j > 1$. For each downward closed set Λ_M of cardinality M, the generalized SG polynomial interpolation operator and the corresponding approximate quadrature are, resp.,

$$\mathcal{S}_{\Lambda_M} g(\boldsymbol{y}) = \sum_{\boldsymbol{i} \in \Lambda_M} \sum_{\boldsymbol{k} \in m^{\boldsymbol{i}}_{\triangle}} \underbrace{\left(g(\boldsymbol{y}^{\boldsymbol{i}}_{\boldsymbol{k}}) - \mathcal{S}_{\Lambda_M \setminus \{\boldsymbol{i}\}} g(\boldsymbol{y}^{\boldsymbol{i}}_{\boldsymbol{k}})\right)}_{s^{\boldsymbol{i}}_{\boldsymbol{k}}} l^{\boldsymbol{i}}_{\boldsymbol{k}}(\boldsymbol{y}) \text{ and } \mathbb{E}[\mathcal{S}_{\Lambda_M} g] = \sum_{\boldsymbol{i} \in \Lambda_M} \sum_{\boldsymbol{k} \in m^{\boldsymbol{i}}_{\triangle}} s^{\boldsymbol{i}}_{\boldsymbol{k}} w^{\boldsymbol{i}}_{\boldsymbol{k}} \,. \tag{3.13}$$

We next outline a heuristic algorithm for the construction of the dimension-adaptive, generalized SG: we initialize the set $\Lambda_1 = \{\mathbf{1}\}$, and compute $s^{\mathbf{1}}_{\mathbf{1}} = g(\boldsymbol{y}^{\mathbf{1}}_{\mathbf{1}})$; at the next level, we enrich Λ_1 with the indices of the "forward neighborhood" (also referred to as "reduced set of neighbors" in [50]) $\{\mathbf{1} + \boldsymbol{e}_j, j = 1, \ldots, J\}$ of the root index $\boldsymbol{i}$, i.e. $\Lambda_M = \{\mathbf{1}, \mathbf{1} + \boldsymbol{e}_j, j = 1, \ldots, J\}$ with $M = J + 1$, and compute the hierarchical surplus $s^{\boldsymbol{i}}_{\boldsymbol{k}}$ for $\boldsymbol{i} \in \Lambda_M \setminus \{\mathbf{1}\}$. Afterwards, we pick the new index $\boldsymbol{i}$ according to a suitable error indicator defined via the hierarchical surplus and enrich Λ_M with the indices from $\{\boldsymbol{i} + \boldsymbol{e}_j, 1 \le j \le J\}$ such that Λ_M remains downward closed . The interpolation error indicator is defined by incorporation (as a trade-off) of the work to evaluate the hierarchical surplus as

$$\boldsymbol{i} = \underset{\boldsymbol{i}' \in \mathscr{A}}{\operatorname{argmax}} \, \mathcal{E}_i(\boldsymbol{i}') \text{ with } \mathcal{E}_i(\boldsymbol{i}') = \frac{1}{|m^{\boldsymbol{i}'}_{\triangle}|} \sum_{\boldsymbol{k} \in m^{\boldsymbol{i}'}_{\triangle}} |s^{\boldsymbol{i}'}_{\boldsymbol{k}}|, \tag{3.14}$$

where $|m^{\boldsymbol{i}'}_{\triangle}|$ is the cardinality of the $m^{\boldsymbol{i}'}_{\triangle}$; here, $\mathscr{A}$ denotes the active index set collecting all those multiindices in Λ_M whose forward neighbors have not yet been processed. We refer to its complement as "old index set", and denote it as $\mathscr{O}$. Upon termination of the enrichment, the index $\boldsymbol{i}$ is moved from $\mathscr{A}$ to $\mathscr{O}$ and all elements in the reduced set of neighbors of $\boldsymbol{i}$ into $\mathscr{A}$ and Λ_M are added. Subsequently, we carry out the same procedure to construct the generalized SG until a stopping criterion is reached (typically, error tolerance or a maximum number of indices (or nodes). This is summarized in

Algorithm 1 Dimension-adaptive sparse grid construction

procedure INITIALIZATION:
  specify error tolerance ε_t and indicator $\mathcal{E}_{max} = 2\varepsilon_t$, set the maximum number of indices M_{max};
  set $M = 1$, $\boldsymbol{i} = \mathbf{1}$, initialize index sets $\mathscr{A} = \{\mathbf{1}\}$, $\mathscr{O} = \emptyset$, $\Lambda_M = \mathscr{A} \cup \mathscr{O}$, and evaluate $s^{\mathbf{1}}_{\mathbf{1}} = g(\boldsymbol{y}^{\mathbf{1}}_{\mathbf{1}})$;
end procedure
procedure CONSTRUCTION:
  while $\mathcal{E}_{max} > \varepsilon_t$ and $M \le M_{max}$ **do**
    set $\mathscr{O} = \mathscr{O} \cup \{\boldsymbol{i}\}$, $\mathscr{A} = \mathscr{A} \setminus \{\boldsymbol{i}\}$ and enrich $\mathscr{A}$ by the admissible forward neighbors of $\boldsymbol{i}$;
    compute the set of nodes $\Xi_{\triangle}$ different from old nodes at the newly added indices $m_{\triangle}$;
    for all $\boldsymbol{i} \in m_{\triangle}$ and $\boldsymbol{y}^{\boldsymbol{i}}_{\boldsymbol{k}} \in \Xi_{\triangle}$, evaluate $g(\boldsymbol{y}^{\boldsymbol{i}}_{\boldsymbol{k}})$ and the interpolation $\mathcal{S}_{\Lambda_M} g(\boldsymbol{y}^{\boldsymbol{i}}_{\boldsymbol{k}})$ by (3.13);
    compute the hierarchical surpluses $s^{\boldsymbol{i}}_{\boldsymbol{k}} = s(\boldsymbol{y}^{\boldsymbol{i}}_{\boldsymbol{k}}) - \mathcal{S}_{\Lambda_M} s(\boldsymbol{y}^{\boldsymbol{i}}_{\boldsymbol{k}})$ and error indicator $\mathcal{E}_i$ by (3.14);
    increase the number of indices $M = M + \#|m_{\triangle}|$, update the total index set $\Lambda_M = \mathscr{A} \cup \mathscr{O}$;
    pick the next index $\boldsymbol{i}$ such that $\boldsymbol{i} = \operatorname{argmax}_{\boldsymbol{i}' \in \mathscr{A}} \mathcal{E}_i(\boldsymbol{i}')$, set the indicator $\mathcal{E}_{max} = \mathcal{E}_i(\boldsymbol{i})$;
  end while
end procedure

The high-dimensional integration is performed by dimension-adaptive, sparse tensor product Smolyak quadrature with replacement of the interpolation error indicator $\mathcal{E}_i$ by the *(heuristic) integration error indicator*

$$\mathcal{E}_e(\boldsymbol{i}') = \frac{1}{|m^{\boldsymbol{i}'}_{\triangle}|} \left| \sum_{\boldsymbol{k} \in m^{\boldsymbol{i}'}_{\triangle}} s^{\boldsymbol{i}'}_{\boldsymbol{k}} w^{\boldsymbol{i}'}_{\boldsymbol{k}} \right|. \tag{3.15}$$

Reasonable error estimates for the interpolation and integration approximation are given by

$$\mathcal{E}_i(\mathscr{A}) = \max_{\boldsymbol{i} \in \mathscr{A}} \max_{\boldsymbol{k} \in m^{\boldsymbol{i}}_{\triangle}} |s^{\boldsymbol{i}}_{\boldsymbol{k}}| \text{ and } \mathcal{E}_e(\mathscr{A}) = \left| \sum_{\boldsymbol{i} \in \mathscr{A}} \sum_{\boldsymbol{k} \in m^{\boldsymbol{i}}_{\triangle}} s^{\boldsymbol{i}}_{\boldsymbol{k}} w^{\boldsymbol{i}}_{\boldsymbol{k}} \right|. \tag{3.16}$$

We remark that there might be situation where the error indicator becomes smaller than the error tolerance, the true approximation error still remains very large, which leads premature termination of the construction. This phenomenon is well-known in adaptive SG algorithms and has been referred to as *stagnation* (cp. [26]). To avoid stagnation, we propose to use the verified dimension-adaptive construction algorithm in [10].

4 Reduced basis approximation

This section is devoted to the development of efficient computational reduction techniques in order to tackle the computational challenge of the expensive forward emulator G in Bayesian estimation for large scale PDE. For this purpose, we rely on a reduced basis method [42, 46] in approximating a high-fidelity PG solution.

4.1 Petrov–Galerkin ("high-fidelity ") discretization

Our proposed RB acceleration will be based on so-called "high-fidelity " discretizations of the forward problems. As we do not assume the parametric operator $A(\boldsymbol{y})$ to be self-adjoint, we assume available *general, high-fidelity PG* (PG for short) discretization, which we formulate next: let $\{\mathcal{X}_h\}_{h>0} \subset \mathcal{X}$ and $\{\mathcal{Y}_h\}_{h>0} \subset \mathcal{Y}$ be two one-parameter families of subspaces of equal, finite dimension $N_h = \dim(\mathcal{X}_h) = \dim(\mathcal{Y}_h)$ which are dense in $\mathcal{X}$ and in $\mathcal{Y}$, respectively. Here, $h > 0$ is a discretization parameter such as the meswidth in Finite Element of Finite Volume methods, or the reciprocal of the spectral order in spectral methods.

The high-fidelity PG approximation of the parametric operator equation (2.15) reads:

$$\text{given any } \boldsymbol{y} \in U, \text{ find } q_h(\boldsymbol{y}) \in \mathcal{X}_h \text{ such that } \mathfrak{a}(\boldsymbol{y}; q_h(\boldsymbol{y}), v_h) = f(v_h) \quad \forall v_h \in \mathcal{Y}_h \ . \tag{4.1}$$

For the well-posedness of the discrete problem (4.1), we assume that the subspace sequences $\{\mathcal{X}_h\}_{h>0} \subset \mathcal{X}$ and $\{\mathcal{Y}_h\}_{h>0} \subset \mathcal{Y}$ are stable with respect to the nominal problem. Then, under Assumption (2.12), there exist $\beta_0^h > 0$ and $h_0 > 0$ such that for every $0 < h \leq h_0$, there hold the uniform (with respect to $\boldsymbol{y} \in U$) discrete inf-sup conditions

$$\inf_{0\neq w_h\in\mathcal{X}_h} \sup_{0\neq v_h\in\mathcal{Y}_h} \frac{\mathfrak{a}_0(w_h, v_h)}{\|w_h\|_{\mathcal{X}}\|v_h\|_{\mathcal{Y}}} \geq \beta_0^h > 0, \quad \inf_{0\neq v_h\in\mathcal{Y}_h} \sup_{0\neq w_h\in\mathcal{X}_h} \frac{\mathfrak{a}_0(w_h, v_h)}{\|w_h\|_{\mathcal{X}}\|v_h\|_{\mathcal{Y}}} \geq \beta_0^h > 0 \ , \tag{4.2}$$

$$\forall \boldsymbol{y} \in U: \inf_{0\neq w_h\in\mathcal{X}_h} \sup_{0\neq v_h\in\mathcal{Y}_h} \frac{\mathfrak{a}(\boldsymbol{y}; w_h, v_h)}{\|w_h\|_{\mathcal{X}}\|v_h\|_{\mathcal{Y}}} \geq \beta_h > 0, \inf_{0\neq v_h\in\mathcal{Y}_h} \sup_{0\neq w_h\in\mathcal{X}_h} \frac{\mathfrak{a}(\boldsymbol{y}; w_h, v_h)}{\|w_h\|_{\mathcal{X}}\|v_h\|_{\mathcal{Y}}} \geq \beta_h > 0 \ . \tag{4.3}$$

We expand the parametric high-fidelity PG solution $q_h(\boldsymbol{y})$ on the bases $(w_h^n)_{n=1}^{N_h}$ as

$$q_h(\boldsymbol{y}) = \sum_{n=1}^{N_h} q_h^n(\boldsymbol{y}) w_h^n \tag{4.4}$$

and denote the parametric coefficient vector $\mathbf{q}_h(\boldsymbol{y}) = (q_h^1(\boldsymbol{y}), \ldots, q_h^{N_h}(\boldsymbol{y}))^\top$. Due to the affine structure (2.6) we may express the bilinear form in (4.1) through the definition (2.7) as

$$\mathfrak{a}(\boldsymbol{y}; q_h(\boldsymbol{y}), v_h) = \mathfrak{a}_0(\boldsymbol{y}; q_h(\boldsymbol{y}), v_h) + \sum_{j\in\mathbb{J}} y_j \mathfrak{a}_j(q_h(\boldsymbol{y}), v_h) \ . \tag{4.5}$$

We introduce the high-fidelity PG matrices $\mathbb{A}_h^j$, $j \in \{0\} \cup \mathbb{J}$, through the bilinear forms in (4.1) as

$$(\mathbb{A}_h^j)_{nn'} := \mathfrak{a}_j(w_h^n, v_h^{n'}) \quad 1 \leq n, n' \leq N_h \tag{4.6}$$

and denote the parametric coefficient vector $\mathbf{f}_h = (f(v_h^1), \ldots, f(v_h^{N_h}))^\top$. With these notations, we

recast the high-fidelity PG problem (4.1) in algebraic form:

$$\text{given any } \boldsymbol{y} \in U, \text{ find } \mathbf{q}_h(\boldsymbol{y}) \in \mathbb{R}^{N_h} \text{ such that } \left(\mathbb{A}_h^0 + \sum_{j \in \mathbb{J}} y_j \mathbb{A}_h^j \right) \mathbf{q}_h(\boldsymbol{y}) = \mathbf{f}_h \ . \tag{4.7}$$

Observe that discrete inf-sup condition (4.3) implies nonsingularity of parametric stiffness matrix in (4.7). Once the parametric high-fidelity solution vector $\mathbf{q}_h(\boldsymbol{y})$ is available at given $\boldsymbol{y} \in U$, we can compute the high-fidelity approximation $\Theta_h(\boldsymbol{y})$ of the parametric Bayesian posterior density $\Theta(\boldsymbol{y})$ defined by

$$\Theta_h(\boldsymbol{y}) = \exp\left(-\frac{1}{2} (\delta - \mathbb{O}_h \mathbf{q}_h(\boldsymbol{y}))^\top \Gamma^{-1} (\delta - \mathbb{O}_h \mathbf{q}_h(\boldsymbol{y})) \right), \tag{4.8}$$

where $\mathbb{O}_h \in \mathbb{R}^{K \times N_h}$ denotes the high-fidelity observation matrix, the kth row vector of which is given by

$$\mathbf{o}_h^k = (o_k(w_h^1), \ldots, o_k(w_h^{N_h})), \tag{4.9}$$

being the linear functional o_k the kth element of the observation operator $\mathcal{O}$.

4.2 Reduced basis approximation

To achieve an accurate high-fidelity PG approximation of the foward problem often requires a prohibitive number $N_h = \dim \mathcal{X}_n = \dim \mathcal{Y}_h$ of degrees of freedom. The large-scale algebraic system (4.7) is then costly to solve, even at linear, ie., $O(N_h)$ complexity.

In order to further reduce this computational cost, particularly in the real-time or many-query context as in the Bayesian inverse problem, we propose acceleration by model order reduction techniques [9]. Specifically, a RBM [42] that only requires the solution of a small parametric algebraic system once RB approximation spaces which perform uniformly with respect to uncertainty u, resp. to the parameter $\boldsymbol{y} \in U$ are constructed.

To this end, let $\mathcal{X}_N \subset \mathcal{X}$ and $\mathcal{Y}_N \subset \mathcal{Y}$ denote two RB approximation spaces of equal, finite dimension N. In particular, we seek $\mathcal{X}_N \subset \mathcal{X}_h$ and $\mathcal{Y}_N \subset \mathcal{Y}_h$ such that $N \ll N_h$. In the context of the RBM [42], we choose the reduced bases such that $\mathcal{X}_N \subset \mathcal{M}_h$, being $\mathcal{M}_h$ the high-fidelity PG parametric solution manifold defined as $\mathcal{M}_h := \{q_h(\boldsymbol{y}) \in \mathcal{X}_h : q_h(\boldsymbol{y}) \text{ solves } (4.1), \forall \boldsymbol{y} \in U\}$. Moreover, we require that the bases in $\mathcal{X}_N$ are mutually orthonormal, which can be achieved by performing the Gram–Schmidt process in order to guarantee that the RB system is well-conditioned. The test space $\mathcal{Y}_N$ associated with $\mathcal{X}_N$ will be specified at the end of the next section. Analogous to the high-fidelity PG approximation, we denote by $(w_N^n)_{n=1}^N$ and $(v_N^n)_{n=1}^N$ the bases of the spaces $\mathcal{X}_N$ and $\mathcal{Y}_N$. Then the *reduced basis Petrov–Galerkin* (RB-PG) approximation of problem (2.15) reads:

$$\text{given any } \boldsymbol{y} \in U, \text{ find } q_N(\boldsymbol{y}) \in \mathcal{X}_N \text{ such that } \mathfrak{a}(q_N(\boldsymbol{y}), v_N; \boldsymbol{y}) = f(v_N) \quad \forall v_N \in \mathcal{Y}_N \ , \tag{4.10}$$

where the expansion of the reduced solution $q_N(\boldsymbol{y})$ on the bases in $\mathcal{X}_N$ is written as

$$q_N(\boldsymbol{y}) = \sum_{n=1}^N q_N^n(\boldsymbol{y}) w_N^n. \tag{4.11}$$

By $\mathbf{q}_N(\boldsymbol{y}) = (q_N^1, \ldots, q_N^N)^\top$, we denote the (parametric) RB coefficient vector. Again due to the affine structure of the diffusion coefficient in (2.5), the bilinear form in (4.10) can be written as

$$\mathfrak{a}(q_N(\boldsymbol{y}), v_N; \boldsymbol{y}) = \mathfrak{a}_0(q_N(\boldsymbol{y}), v_N) + \sum_{j \in \mathbb{J}} y_j \mathfrak{a}_j(q_N(\boldsymbol{y}), v_N). \tag{4.12}$$

In order to establish the algebraic formulation of the reduced problem (4.10), we first expand the reduced bases $(w_N^n)_{n=1}^N$ and $(v_N^n)_{n=1}^N$ on the high-fidelity bases $(w_h^n)_{n=1}^{N_h}$ and $(v_h^n)_{n=1}^{N_h}$, due to the fact

that $\mathcal{X}_N \subset \mathcal{X}_h$ and $\mathcal{Y}_N \subset \mathcal{Y}_h$, respectively, yielding

$$w_N^n = \sum_{m=1}^{N_h} w_N^{n,m} w_h^m, \text{ and } v_N^n = \sum_{m=1}^{N_h} v_N^{n,m} v_h^m, \ 1 \le n \le N \,. \tag{4.13}$$

Let $\mathbf{w}_N^n = (w_N^{n,1}, \ldots, w_N^{n,N_h})^\top$, and let $\mathbf{v}_N^n = (v_N^{n,1}, \ldots, v_N^{n,N_h})^\top$, $1 \le n \le N$, denote the coefficient vectors of the high-fidelity PG approximation. A key role is played by the "high-fidelity coefficient matrices", ie. $\mathbb{W} = (\mathbf{w}_N^1, \ldots, \mathbf{w}_N^N)$ and $\mathbb{V} = (\mathbf{v}_N^1, \ldots, \mathbf{v}_N^N)$. To this end, we construct reduced matrices $\mathbb{A}_N^j$, $j \in \{0\} \cup \mathbb{J}$, through the bilinear forms in (4.12) as

$$(\mathbb{A}_N^j)_{nn'} := \mathfrak{a}_j(w_N^n, v_N^{n'}) = \sum_{m=1}^{N_h} \sum_{m'=1}^{N_h} w_N^{n,m} \mathfrak{a}_j(w_h^m, v_h^{m'}) v_N^{n',m'} = {\mathbf{v}_N^n}^\top \mathbb{A}_h^j \mathbf{w}_N^{n'} \,, \tag{4.14}$$

so that the reduced matrices are given by

$$\mathbb{A}_N^j = \mathbb{V}^\top \mathbb{A}_h^j \mathbb{W}, \quad j \in \{0\} \cup \mathbb{J} \,. \tag{4.15}$$

We construct the reduced load vector $\mathbf{f}_N \in \mathbb{R}^N$ whose n-th ($1 \le n \le N$) entry is given by

$$(\mathbf{f}_N)_n = f(v_N^n) = \sum_{m=1}^{N_h} f(v_h^m) v_N^{n,m} = (\mathbf{v}_N^n)^\top \mathbf{f}_h, \tag{4.16}$$

yielding

$$\mathbf{f}_N = \mathbb{V}^\top \mathbf{f}_h \,. \tag{4.17}$$

The algebraic form of the RB Galerkin approximation (4.10) becomes:

$$\text{given any } \boldsymbol{y} \in U, \text{ find } \mathbf{q}_N(\boldsymbol{y}) \in \mathbb{R}^N \text{ such that } \left(\mathbb{A}_N^0 + \sum_{j \in \mathbb{J}} y_j \mathbb{A}_N^j \right) \mathbf{q}_N(\boldsymbol{y}) = \mathbf{f}_N \,. \tag{4.18}$$

Note that we only need to assemble $\mathbb{A}_N^j$, $j \in \{0\} \cup \mathbb{J}$, and $\mathbf{f}_N$ once with computational cost depending on the high-fidelity degree of freedom N_h. With the RB solution $\mathbf{q}_N(\boldsymbol{y})$ at some given $\boldsymbol{y} \in U$, we evaluate a RB approximation of the posterior density $\Theta_N(\boldsymbol{y})$ as an approximation of $\Theta_h(\boldsymbol{y})$ in (4.8) by (note that $\mathbb{W}\mathbf{q}_N(\boldsymbol{y})$ is an approximation of $\mathbf{q}_h(\boldsymbol{y})$)

$$\begin{aligned} \Theta_N(\boldsymbol{y}) &= \exp\left(-\frac{1}{2} (\delta - \mathbb{O}_N \mathbf{q}_N(\boldsymbol{y}))^\top \Gamma^{-1} (\delta - \mathbb{O}_N \mathbf{q}_N(\boldsymbol{y})) \right) \\ &= \exp\left(-\frac{1}{2} \left(\delta^\top \Gamma^{-1} \delta - 2\delta^\top \Gamma^{-1} \mathbb{O}_N \mathbf{q}_N(\boldsymbol{y}) + \mathbf{q}_N^\top \mathbb{O}_N^\top \Gamma^{-1} \mathbb{O}_N \mathbf{q}_N(\boldsymbol{y}) \right) \right) \end{aligned} \tag{4.19}$$

where $\delta^\top \Gamma^{-1} \delta$, $\delta^\top \Gamma^{-1} \mathbb{O}_N$ and $\mathbb{O}_N^\top \Gamma^{-1} \mathbb{O}_N$ can be assembled once, being the RB observation matrix $\mathbb{O}_N := \mathbb{O}_h \mathbb{W}$. Therefore, instead of solving the (generally large) linear system (4.7) from the high-fidelity PG discretization, which is of size $N_h \times N_h$, only the numerical solution of the much smaller, but generally dense reduced, linear system (4.18) of equations of size $N \times N$ is required, which is $O(N^3)$ due to the matrix being dense. The overall computational work is considerably smaller than (4.7) when $N \ll N_h$. The cost for evaluation of the density $\Theta_N(\boldsymbol{y})$ by formula (4.19) is only $O(N^2)$.

4.3 Construction of RB trial and test function spaces

Once the reduced bases are determined, the high-fidelity solution of a large system (4.7) can be approximated by the reduced solution of a small system (4.18) at any given $\boldsymbol{y} \in U$. The accuracy of RB approximation depends crucially on, roughly speaking, the ability of reduced bases to represent the high-fidelity solution manifold uniformly over the parameter space U. We present a greedy algorithm from [42] to adaptively construct the RB trial and test space sequences $\mathcal{X}_1 \subset \mathcal{X}_2 \subset \cdots \subset \mathcal{X}_N$ and $\mathcal{Y}_1 \subset \mathcal{Y}_2 \subset \cdots \subset \mathcal{Y}_N$ with favourable approximation and stability properties.

As mentioned in the last section, in the context of the RBM, reduced bases are constructed from the high-fidelity solutions at some parameter samples in U, known as "snapshots". To start, we pick the first sample $\boldsymbol{y}^{(1)} \in U$, e.g. by random sampling or by using the center of U, and subsequently solving the high-fidelity system (4.7) to get the first snapshot $q_h(\boldsymbol{y}^{(1)})$. The trial reduced space is initialized as $\mathcal{X}_N = \text{span}\{q_h(\boldsymbol{y}^{(1)})\}$ for $N = 1$ (the test reduced space $\mathcal{Y}_N$ with be defined later). In the following steps for $N = 1, 2, \ldots,$, we seek the sample $\boldsymbol{y}^{(N+1)} \in U$ by solving (approximately) a high-dimensional maximization problem [42], e.g. maximizing the RB error of the forward solution

$$\boldsymbol{y}^{(N+1)} = \underset{\boldsymbol{y} \in U}{\text{argsup}} \, ||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}}, \tag{4.20}$$

where $q_N(\boldsymbol{y})$ is the reduced solution of problem (4.10) in the reduced space given by

$$\mathcal{X}_N = \text{span}\{q_h(\boldsymbol{y}^{(n)}), 1 \leq n \leq N\}. \tag{4.21}$$

This is to say that $\boldsymbol{y}^{(N+1)}$ is a sample (not nec. unique) where the high-fidelity solution q_h is worst approximated by the RB solution q_N among all $\boldsymbol{y} \in U$, when measured in $||\cdot||_{\mathcal{X}}$ norm. However, as the maximization problem (4.20) is infinite dimensional (there is generally an infinite number of $\boldsymbol{y} \in U$) and requires the high-fidelity solution, it is computationally very expensive. In order to alleviate this cost, we replace U by a finite training set $\Xi_{train} \subset U$ and replace $||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}}$ for $N = 1, 2, \ldots,$ by some a posteriori error estimate $\triangle_N(\boldsymbol{y})$ for any $\boldsymbol{y} \in U$. We present the greedy algorithm in 2, also called *weak greedy algorithm* because of the replacement by the error estimate and the training set.

Algorithm 2 Greedy algorithm [42]

procedure INITIALIZATION:
  specify tolerance ϵ_t, a training set $\Xi_{train} \subset U$, and $N_{max} \in \mathbb{N}$, set $N = 1$, pick $\boldsymbol{y}^{(1)} \in \Xi_{train}$;
  solve (4.1) at $\boldsymbol{y}^{(1)}$ and construct $\mathcal{X}_1 = \text{span}\{q_h(\boldsymbol{y}^{(1)})\}$, evaluate $\triangle_N(\boldsymbol{y})$ at all $\boldsymbol{y} \in \Xi_{train}$;
end procedure
procedure CONSTRUCTION:
  while $\max_{\boldsymbol{y} \in \Xi_{train}} \triangle_N(\boldsymbol{y}) \geq \epsilon_t$ & $N \leq N_{max}$ **do**
    set $\boldsymbol{y}^{(N+1)} = \text{argmax}_{\boldsymbol{y} \in \Xi_{train}} \triangle_N(\boldsymbol{y})$;
    solve (4.1) at $\boldsymbol{y}^{(N+1)}$ to obtain $q_h(\boldsymbol{y}^{(N+1)})$;
    update $\mathcal{X}_{N+1} = \mathcal{X}_N \oplus \text{span}\{q_h(\boldsymbol{y}^{(N+1)})\}$;
    compute $\triangle_{N+1}(\boldsymbol{y})$ for all $\boldsymbol{y} \in \Xi_{train}$;
    set $N = N + 1$;
  end while
  set $N_{max} = N$;
end procedure

Some (empirical) finding to ensure computational efficiency and accuracy of the greedy algorithms from [42, 46]: (i) the training set Ξ_{train} ought to be sufficiently rich such that the solution manifold can be well explored over the training set, while it should not be too large to spoil the computational efficiency by exploring too many samples where the solution is well approximated; (ii) the error estimate should be inexpensive to evaluate in order to allow a rich training set; moreover, the error estimate must be reliable in order to provide a confident and sharp estimates of the RB approximation error for *all* samples, not only in Ξ_{train} but also in the entire parameter space U.

We leave the treatment of the second criterion to the next section and address the first criterion by developing an adaptive greedy algorithm in combination of the construction of the generalized SG in section 3.3. In fact, as we only need to evaluate the posterior density at the interpolation/quadrature nodes by the dimension-adaptive SG method, it is natural to take these nodes as the training samples. In Algorithm 3, we present an adaptive greedy algorithm that enable the construction of the reduced space and the evaluation of the reduced posterior density on the generalized SG simultaneously.

To this end, we specify the construction of the test space $\mathcal{Y}_N$ in two different cases. In the first case, we set $\mathcal{Y}_N = \mathcal{X}_N$, so that the PG approximation (4.1) is simply a Galerkin approximation with $\mathcal{Y}_h = \mathcal{X}_h$. In the second case, normally when $\mathcal{Y}_h \neq \mathcal{X}_h$ in the PG approximation problem (4.1), we adopt a minimum-residual (or least-squares) approach [36] to compute a set of parameter dependent

Algorithm 3 Adaptive greedy algorithm [10]

1: **procedure** INITIALIZATION:
2: specify tolerance ϵ_t, set $N = 1$, solve (4.1) at $\boldsymbol{y}_1^1$ and construct $\mathcal{X}_1 = \text{span}\{q_h(\boldsymbol{y}_1^1)\}$;
3: **end procedure**
4: **procedure** CONSTRUCTION:
5: at each step in line 9 of Algorithm 1, choose the training set $\Xi_{train} = \Xi_{\triangle}$;
6: solve problem (4.10) at each $\boldsymbol{y} \in \Xi_{train}$ and compute $\triangle_N(\boldsymbol{y})$ and $\Theta_N(\boldsymbol{y})$;
7: update Ξ_{train} as $\Xi_{train} = \Xi_{train} \setminus \{\boldsymbol{y} \in \Xi_{train} : \triangle_N(\boldsymbol{y}) < \epsilon_t\}$;
8: **while** $\max_{\boldsymbol{y}\in\Xi_{train}} \triangle_N(\boldsymbol{y}) \geq \epsilon_t$ **do**
9: set $\boldsymbol{y}^{(N+1)} = \text{argmax}_{\boldsymbol{y}\in\Xi_{train}} \triangle_N(\boldsymbol{y})$;
10: solve (4.1) at $\boldsymbol{y}^{(N+1)}$ to obtain $q_h(\boldsymbol{y}^{(N+1)})$;
11: update $\mathcal{X}_{N+1} = \mathcal{X}_N \oplus \text{span}\{q_h(\boldsymbol{y}^{(N+1)})\}$;
12: set $N = N + 1$ and repeat lines 6 – 7;
13: **end while**
14: **end procedure**

reduced bases $v_N^n(\boldsymbol{y})$, $1 \leq n \leq N$, by solving

$$(v_N^n(\boldsymbol{y}), v_h)_{\mathcal{Y}_h} = \mathfrak{a}(w_N^n, v_h; \boldsymbol{y}) \quad \forall v_h \in \mathcal{Y}_h, \text{ given } \boldsymbol{y} \in U, \tag{4.22}$$

where $(w_N^n)_{n=1}^N$ are the bases of $\mathcal{X}_N$ and the inner product $(\cdot,\cdot)_{\mathcal{Y}_h}$ is defined in $\mathcal{Y}_h$. In this case, the reduced algebraic system in problem (4.18) can be expressed more explicitly. Let $\mathbb{M}_h$ denote the mass matrix in $\mathcal{Y}_h$, whose entries are defined as $(\mathbb{M}_h)_{nn'} = (v_h^n, v_h^{n'})_h$, $1 \leq n, n' \leq N_h$. Then the matrix $\mathbb{A}_N^j$ in (4.15) is replaced by $\mathbb{A}_N^{j,j'} = (\mathbb{A}_h^j \mathbb{W})^\top \mathbb{M}_h^{-1} \mathbb{A}_h^{j'} \mathbb{W}$, and the vector $\mathbf{f}_N$ in (4.17) by $\mathbf{f}_N^j = (\mathbb{A}_h^j \mathbb{W})^\top \mathbb{M}_h^{-1} \mathbf{f}_h$ for $j, j' \in \{0\} \cup \mathbb{J}$. The RB algebraic system in problem (4.18) can be rewritten as

$$\left(\mathbb{A}_N^{0,0} + 2 \sum_{j\in\mathbb{J}} y_j \mathbb{A}_N^{0,j} + \sum_{j\in\mathbb{J}} \sum_{j'\in\mathbb{J}} y_j y_{j'} \mathbb{A}_N^{j,j'} \right) \mathbf{q}_N(\boldsymbol{y}) = \mathbf{f}_N^0 + \sum_{j\in\mathbb{J}} y_j \mathbf{f}_N^j, \tag{4.23}$$

where the matrices $\mathbb{A}_N^{j,j'}$ and vectors $\mathbf{f}_N^j$, $j, j' \in \{0\} \cup \mathbb{J}$, are assembled only once. For any given $\boldsymbol{y} \in U$, this requires $O((1 + |\mathbb{J}|)^2 N^2)$ and $O(N^3)$ operations for assembling and solving (4.23).

4.4 A posteriori error estimate

An inexpensive, reliable a-posteriori error estimator plays an crucial role in constructing suitable reduced basis trial- and test spaces as well as in quantifying the RB approximation error. As our primary interest is the posterior density Θ and the QoI Ψ, we propose a dual-weighted residual, goal-oriented a posterior error estimator to bound the error for the density $\Theta(\boldsymbol{y})$ between the high-fidelity approximation $\Theta_h(\boldsymbol{y})$ in (4.8) and the RB approximation $\Theta_N(\boldsymbol{y})$ in (4.19) for any $\boldsymbol{y} \in U$. A computable estimator for the the QoI Ψ can be obtained analogously.

For every $\boldsymbol{y} \in U$, the nonlinear (with respect to the high-fidelity solution $q_h(\boldsymbol{y})$) output $\Theta_h(\boldsymbol{y})$ is formally expanded about the reduced solution $q_N(\boldsymbol{y})$ as

$$\Theta_h(\boldsymbol{y}) \approx \Theta_N(\boldsymbol{y}) + \frac{\partial \Theta_h}{\partial q_h}\Big|_{q_N(\boldsymbol{y})} (q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})), \tag{4.24}$$

where the Fréchet derivative $\frac{\partial \Theta_h}{\partial q_h}\Big|_{q_N(\boldsymbol{y})}(w_h)$ ievaluated in the direction of w_h, $\forall w_h \in \mathcal{X}_h$, is given by

$$\frac{\partial \Theta_h}{\partial q_h}\Big|_{q_N(\boldsymbol{y})}(w_h) = \Theta_N(\boldsymbol{y}) \left((\delta - \mathcal{O}(q_N(\boldsymbol{y})))^\top \Gamma^{-1} \mathcal{O}(w_h) \right) . \tag{4.25}$$

As the posterior density $\Theta(\boldsymbol{y})$ is nonlinear with respect to the parametric forward solution map $q(\boldsymbol{y})$, we may not directly apply a primal-dual approach as in[36, 42]. Following [1], we associate a

dual problem with the density $\Theta(\boldsymbol{y})$, and construct a so-called dual-weighted-residual (DWR) [38] as an a posteriori error estimate. The parametric high-fidelity PG approximation of the dual problem corresponding to (4.1) is given by

$$\text{given any } \boldsymbol{y} \in U, \text{ find } \psi_h(\boldsymbol{y}) \in \mathcal{Y}_h \text{ such that } \mathfrak{a}(\boldsymbol{y}; w_h, \psi_h) = \frac{\partial \Theta_h}{\partial q_h}\Big|_{q_N(\boldsymbol{y})}(w_h) \quad \forall w_h \in \mathcal{X}_h \ . \tag{4.26}$$

We numerically approximate this high-fidelity PG solution with a RB solution by solving

$$\text{find } \psi_{N_{du}}(\boldsymbol{y}) \in \mathcal{Y}_{N_{du}} \text{ such that } \mathfrak{a}(\boldsymbol{y}; w^{du}_{N_{du}}, \psi_{N_{du}}) = \frac{\partial \Theta_h}{\partial q_h}\Big|_{q_N(\boldsymbol{y})}(w^{du}_{N_{du}}) \quad \forall w^{du}_{N_{du}} \in \mathcal{X}_{N_{du}}, \tag{4.27}$$

where N_{du} represents the number of reduced bases constructed for the dual problem (4.26). Note that the trial and test reduced spaces become $\mathcal{Y}_{N_{du}}$ and $\mathcal{X}_{N_{du}}$, whose construction is clear by following the same procedure as in Section 4.3. Once the parametric RB solution $\psi_{N_{du}}(\boldsymbol{y})$ has been computed, we evaluate the a-posteriori error estimate for $\Theta_h(\boldsymbol{y}) - \Theta_N(\boldsymbol{y})$ as

$$\triangle^{\Theta}_N(\boldsymbol{y}) := r(\psi_{N_{du}}(\boldsymbol{y}); \boldsymbol{y}) \ , \tag{4.28}$$

where the parametric, weak residual is defined for every $\boldsymbol{y} \in U$ as

$$r(v_h; \boldsymbol{y}) = f(v_h) - \mathfrak{a}(\boldsymbol{y}; q_N(\boldsymbol{y}), v_h) \quad \forall v_h \in \mathcal{Y}_h \ . \tag{4.29}$$

Due to (4.1) and (4.26), we have

$$\forall \boldsymbol{y} \in U: \quad r(\psi_h(\boldsymbol{y}); \boldsymbol{y}) = \mathfrak{a}(\boldsymbol{y}; e_h(\boldsymbol{y}), \psi_h(\boldsymbol{y})) = \frac{\partial \Theta_h}{\partial q_h}\Big|_{q_N(\boldsymbol{y})}(e_h(\boldsymbol{y})), \tag{4.30}$$

which is nothing but the second term in the expansion (4.24). Moreover, subtracting (4.28) from (4.30) with the notation $e^{du}_h(\boldsymbol{y}) = \psi_h(\boldsymbol{y}) - \psi_{N_{du}}(\boldsymbol{y})$ yields

$$\left|r(\psi_h(\boldsymbol{y}); \boldsymbol{y}) - \triangle^{\Theta}_N(\boldsymbol{y})\right| = |r(e^{du}_h(\boldsymbol{y}); \boldsymbol{y})| = |\mathfrak{a}(\boldsymbol{y}; e_h(\boldsymbol{y}), e^{du}_h(\boldsymbol{y}))| \leq \gamma_h(\boldsymbol{y})||e_h(\boldsymbol{y})||_{\mathcal{X}}||e^{du}_h(\boldsymbol{y})||_{\mathcal{Y}}, \tag{4.31}$$

where $\gamma_h(\boldsymbol{y})$ denotes the continuity constant of the bilinear form $\mathfrak{a}(\boldsymbol{y}; \cdot, \cdot)$ in the high-fidelity spaces $\mathcal{X}_h$ and $\mathcal{Y}_h$ at $\boldsymbol{y} \in U$, i.e. $\mathfrak{a}(\boldsymbol{y}; w_h, v_h) \leq \gamma_h(\boldsymbol{y})||w_h||_{\mathcal{X}}||v_h||_{\mathcal{Y}}$, $\forall w_h \in \mathcal{X}_h, \forall v_h \in \mathcal{Y}_h$. This demonstrates that as long as (4.30) provides an accurate error estimate for $\Theta_h(\boldsymbol{y}) - \Theta_N(\boldsymbol{y})$, the error estimate $\triangle^{\Theta}_N$ defined in (4.28) also does with additional error bounded by (4.31). Therefore, we may correct the RB approximation Θ_N by $\triangle^{\Theta}_N$ as

$$\tilde{\Theta}_N(\boldsymbol{y}) = \Theta_N(\boldsymbol{y}) + \triangle^{\Theta}_N(\boldsymbol{y}) \ . \tag{4.32}$$

In evaluation of $\triangle^{\Theta}_N$, the dual reduced matrices in (4.27) can be assembled as in the primal case:

$$\mathbb{A}^{du,j}_{N_{du}} = \mathbb{W}^{\top}_{du}\mathbb{A}^{j}_h\mathbb{V}_{du}, \quad j \in \{0\} \cup \mathbb{J}, \tag{4.33}$$

where $\mathbb{V}_{du}$ and $\mathbb{W}_{du}$ are the coefficient matrices corresponding to the dual reduced bases. The right hand side of the dual problem (4.27) can be evaluated as

$$\mathbf{f}^{du}_{N_{du}}(\boldsymbol{y}) := \Theta_N(\boldsymbol{y})\left(\mathbb{W}^{\top}_{du}\mathbb{O}^{\top}_h\Gamma^{-1}\delta - \mathbb{W}^{\top}_{du}\mathbb{O}^{\top}_h\Gamma^{-1}\mathbb{O}_h\mathbb{W}\mathbf{q}_N(\boldsymbol{y})\right), \tag{4.34}$$

where $\mathbb{W}^{\top}_{du}\mathbb{O}^{\top}_h\Gamma^{-1}\delta$ and $\mathbb{W}^{\top}_{du}\mathbb{O}^{\top}_h\Gamma^{-1}\mathbb{O}_h\mathbb{W}$ are assembled for just once. Once the reduced dual solution vector $\psi_{N_{du}}(\boldsymbol{y})$ is obtained by solving the following dual linear system

$$\left(\mathbb{A}^{du,0}_{N_{du}} + \sum_{j\in\mathbb{J}} y_j\mathbb{A}^{du,j}_{N_{du}}\right)\psi_{N_{du}}(\boldsymbol{y}) = \mathbf{f}^{du}_{N_{du}}(\boldsymbol{y}), \tag{4.35}$$

with $\mathbb{A}_{N_{du}}^{du,j}$, $j \in \{0\} \cup \mathbb{J}$, assembled for only once, we can evaluate the dual-weighted residual by

$$\triangle_N^{\Theta}(\boldsymbol{y}) = r(\psi_{N_{du}}(\boldsymbol{y}); \boldsymbol{y}) = \mathbf{f}_h^{\top} \mathbb{W}_{du} \psi_{N_{du}}(\boldsymbol{y}) - \sum_{j \in \{0\} \cup \mathbb{J}} y_j(\mathbf{q}_N(\boldsymbol{y}))^{\top} \mathbb{W}^{\top} \mathbb{A}_h^j \mathbb{W}_{du} \psi_{N_{du}}(\boldsymbol{y}), \tag{4.36}$$

where the quantities $\mathbf{f}_h^{\top} \mathbb{W}_{du}$ and $\mathbb{W}^{\top} \mathbb{A}_h^j \mathbb{W}_{du}$ are computed for just once. Note that the dual reduced problem (4.35) can be assembled and solved in $O((1 + |\mathbb{J}|)N^2)$ and $O(N^3)$ operations, respectively, and $\triangle_N^{\Theta}(\boldsymbol{y})$ is assembled in $O(NN_{du}(1 + |\mathbb{J}|))$ operations.

5 A priori error estimates

We provide a priori error estimates for the approximation of the parametric solution family $\{q(\boldsymbol{y}) \in \mathcal{X} : \boldsymbol{y} \in U\}$ with error contribution from parametric truncation, SG interpolation/integration, high-fidelity PG discretization and from RB acceleration, respectively. More importantly, we obtain error bounds and convergence rate estimates for the sparse approximation of the Bayesian posterior density Θ and for the integral of Θ with respect to the Bayesian prior distribution, which plays a key role in Bayesian inversion.

5.1 Dimension truncation

To obtain *computationally feasible* numerical approximation methods, we truncate the infinite sum in (2.6) to a finite number of J terms, so that the SG and the RB approximations introduced in section 3 and section 4 in J-dimensional parametric space become applicable. We remark that both approximations can be employed with adaptive truncation of the dimensions. We denote by $q_J(\boldsymbol{y})$ the solution of the corresponding parametric weak problem (2.15). Then Theorem 2.1 holds when $q(\boldsymbol{y})$ is replaced by $q_J(\boldsymbol{y})$. In addition to the assumption (2.13), which implies $\sum_{j \geq 1} b_j^p < \infty$ with b_j defined as in (2.12), we assume that the operators A_j are enumerated so that

$$b_1 \geq b_2 \geq \cdots \geq b_j \geq \cdots . \tag{5.1}$$

Proposition 5.1 *([32, Theorem 5.1]) Under Assumption 1, for every $f \in \mathcal{Y}'$, for every $\boldsymbol{y} \in U$ and for every $J \in \mathbb{N}$, the solution $q_J(\boldsymbol{y})$ of the J-term truncated parametric weak problem* (2.15) *satisfies*

$$\|q(\boldsymbol{y}) - q_J(\boldsymbol{y})\|_{\mathcal{X}} \leq C_p J^{-s}, \quad s = \frac{1}{p} - 1, \tag{5.2}$$

where the constant

$$C_p = \min\left(\frac{1}{1/p - 1}, 1\right) \left(\sum_{j \geq 1} b_j^p\right)^{1/p} \frac{C}{\beta} ||f||_{\mathcal{Y}'}, \tag{5.3}$$

being $0 < C < \infty$ independent of f, and b_j defined in (2.12)*, and β is the stability constant in* (2.14)*. Moreover, for every observational functional $\mathcal{O} \in \mathcal{X}'$, we have*

$$|\mathcal{O}(q(\boldsymbol{y})) - \mathcal{O}(q_J(\boldsymbol{y}))| \leq \tilde{C} C_p ||\mathcal{O}||_{\mathcal{X}'} J^{-s}, \quad s = \frac{1}{p} - 1, \tag{5.4}$$

where the constant $0 < \tilde{C} < \infty$ depends on K but not on f and $\mathcal{O}$.

5.2 Sparse grid approximation

To establish generalized SG interpolation and integration error estimates, we employ the results for the best M-term approximation by multidimensional polynomials [20, 19, 53, 17, 49].

We consider the generalized polynomial chaos (gPC) expansions

$$q(\boldsymbol{y}) = \sum_{\nu \in \mathcal{F}} q_\nu L_\nu(\boldsymbol{y}) \quad \text{and} \quad \Theta(\boldsymbol{y}) = \sum_{\nu \in \mathcal{F}} \Theta_\nu L_\nu(\boldsymbol{y}), \tag{5.5}$$

where the coefficient function $q_\nu \in \mathcal{X}$ and $\Theta_\nu \in \mathbb{R}$. Here, and throughout, $\mathcal{F}$ denotes the set of all "finitely supported" sequences of nonnegative integers, ie. sequences $\nu = (\nu_1, \ldots, \nu_j, ...) \in \mathbb{N}_0^{\mathbb{N}}$ for which $|\nu| = \nu_1 + \nu_2 + ... < \infty$. We remark that $\mathcal{F}$ is countable and that (5.5) hold with unconditional convergence in $\mathcal{X}$ resp. in $\mathbb{R}$. Furthermore, in (5.5), L_ν is a tensorized polynomial of degree ν, i.e. $L_\nu(\boldsymbol{y}) = \prod_{j \geq 1} L_{\nu_j}(y_j)$, where L_{ν_j} a polynomial of degree ν_j with respect to y_j and $||L_{\nu_j}||_{L^\infty([-1,1])} = 1$ for $\nu_j \in \mathbb{N}_0$, for instance, the normalized Legendre polynomial for Legendre polynomial chaos or simply y^{ν_j} in Taylor gPC representation, as in [17]. For a given downward closed index set $\Lambda_M \subset \mathcal{F}$ of finite cardinality $\#(\Lambda_M) = M < \infty$, we denote by q_{Λ_M} and Θ_{Λ_M} the M-term truncated gPC expansions

$$q_{\Lambda_M}(\boldsymbol{y}) = \sum_{\nu \in \Lambda_M} q_\nu L_\nu(\boldsymbol{y}) \quad \text{and} \quad \Theta_{\Lambda_M}(\boldsymbol{y}) = \sum_{\nu \in \Lambda_M} \Theta_\nu L_\nu(\boldsymbol{y}) \, . \tag{5.6}$$

As $\Lambda_M \subset \mathcal{F}$ is downward closed , the generalized SG interpolation formula (3.13) implies $\mathcal{S}_{\Lambda_M} v(\boldsymbol{y}) = v_{\Lambda_M}(\boldsymbol{y})$, $\forall v \equiv q \in \mathcal{X} \times P_{\Lambda_M}$ or $\forall v \equiv \Theta \in \mathbb{R} \times P_{\Lambda_M}$ with $P_{\Lambda_M} := \text{span}\{L_\nu, \nu \in \Lambda_M\}$, i.e. the interpolation is as accurate as a gPC approximation with the same downward closed set of active indices. Consequently, the generalized SG interpolation error can be bounded by the gPC approximation with proper consideration of the Lebesgue constant. By the holomorphy of the parametric forward solution $q(\boldsymbol{y})$ and of the parametric Bayesian posterior density $\Theta(\boldsymbol{y})$, the following approximation result can be established as in [17].

Proposition 5.2 *Under Assumption 1, let $(\Lambda_M)_{M \geq 1}$, $\#(\Lambda_M) = M$ denote a sequence of nested, downward closed sets of (indices of) M largest gpc coefficients $||q_\nu||_{\mathcal{X}}$, the error of the generalized SG interpolation* (3.13) *of the solution $q : U \to \mathcal{X}$ can be bounded by*

$$\sup_{\boldsymbol{y} \in U} ||q(\boldsymbol{y}) - \mathcal{S}_{\Lambda_M} q(\boldsymbol{y})||_{\mathcal{X}} \leq C_q^i M^{-s}, \quad s = \frac{1}{p} - 1, \tag{5.7}$$

where the constant $C_q^i = 2||(p_\theta(\nu)||q_\nu||_{\mathcal{X}})_{\nu \in \mathcal{F}}||_{\ell^p(\mathcal{F})} < \infty$, being $p_\theta(\nu) = \prod_{j \geq 1}(1 + \nu_j)^{\theta+1}$, and θ is such that the Lebesgue constant of the univariate interpolation (3.1) *$\lambda_m \leq m^\theta$. For instance, it has been shown that $0 < \theta < 1$ for Clenshaw–Curtis nodes [4] and $2 < \theta < 3$ for real-Leja nodes [16]. Analogously, for Λ_M corresponding to the M largest coefficients $|\Theta_\nu|$, we have*

$$\sup_{\boldsymbol{y} \in U} |\Theta(\boldsymbol{y}) - \mathcal{S}_{\Lambda_M} \Theta(\boldsymbol{y})| \leq C_\Theta^i M^{-s}, \quad s = \frac{1}{p} - 1, \tag{5.8}$$

where the constant $C_\Theta^i = 2||(p_\theta(\nu)|\Theta_\nu|)_{\nu \in \mathcal{F}}||_{\ell^p(\mathcal{F})} < \infty$. An analogous result holds also for the QoI Ψ.

Remark 5.1 *The interpolation error estimate is obtained with respect to the cardinality M of the index set Λ_M. For Leja and Gauss quadrature nodes, the same convergence rate with respect to the number of nodes can be obtained; the number of nodes is also M (or $2M$ for symmetric Leja nodes). As for Clenshaw–Curtis, though the number of nodes grows more quickly than the number of indices, the interpolation accuracy also grows faster than the number of indices, i.e. polynomials of degree $2^{\nu_j - 1} + 1$ (rather than ν_j) are exactly interpolated at level ν_j for $\nu_j > 1$, $1 \leq j \leq J$. Therefore, the same convergence rate with respect to the number of nodes is expected, based on the the numerical experiments in Section 6 ahead.*

Analogous results have been obtained for the generalized SG integration error in [49] by comparison to a Taylor expansion of the solution, which is summarized as follows in finite dimensions.

Proposition 5.3 *Under Assumption 1, we denote by $(\Lambda_M)_{M \geq 1} \subset \mathbb{N}_0^{\mathbb{N}}$ with $\#(\Lambda_M) = M < \infty$ a sequence of nested downward closed sets of (indices of) M largest coefficients $||q_\nu||_{\mathcal{X}}$ corresponding to a Taylor expansion of $q(\boldsymbol{y})$ at the nominal values $\boldsymbol{y} = 0$. Then, the error of the generalized SG*

integration (3.13) *of the solution* $q : U \to \mathcal{X}$ *with the SG corresponding to* Λ_M *can be bounded by*

$$||\mathbb{E}[q] - \mathbb{E}[\mathcal{S}_{\Lambda_M} q]||_{\mathcal{X}} \leq C_q^e M^{-s}, \quad s = \frac{1}{p} - 1, \tag{5.9}$$

where $C_q^e = 2||(p_e(\nu)||q_\nu||_{\mathcal{X}})_{\nu\in\mathcal{F}}||_{\ell^p(\mathcal{F})} < \infty$, *being* $p_e(\nu) = \prod_{j\geq 1}(1+\nu_j)^2$. *Similarly, the integration error of the Bayesian posterior density is bounded by*

$$|\mathbb{E}[\Theta] - \mathbb{E}[\mathcal{S}_{\Lambda_M}\Theta]| \leq C_\Theta^e M^{-s}, \quad s = \frac{1}{p} - 1, \tag{5.10}$$

where the constant $C_\Theta^e = 2||(p_e(\nu)|\Theta_\nu|)_{\nu\in\mathcal{F}}||_{\ell^p(\mathcal{F})} < \infty$. *An analogous result holds also for the QoI* Ψ.

Remark 5.2 *The algebraic error convergence rates for interpolation and integration by generalized SG in Theorem 5.2 and 5.3 have been obtained in infinite dimensional cases in [17] and [49], respectively, which also hold in finite dimension* $J < \infty$ *with constants which are independent of* J. *A sub-exponential error bound for the generalized SG interpolation (with constant depending on* J*) has been obtained in [41] in* $L^2(U,\mathcal{X})$*-norm instead of the* $L^\infty(U,\mathcal{X})$*-norm.*

Remark 5.3 *The results in Propositions 5.2 and 5.3 are based on the fact that the downward closed set* Λ_M *collects* M *largest coefficients of the gPC expansion. To date, however, there is no rigorous guarantee that the heuristic, dimension-adaptive algorithm 1 for the construction of the generalized SG gives rise to a quasi-optimal sequence* $\{\Lambda_M\}_{M\geq 1}$ *of such downward closed sets. Higher EOC*[1] *is observed for this algorithm in practice, particularly for the integration error, cp. the examples in Section 6 ahead.*

5.3 High-fidelity approximation

For the numerical analysis of the high-fidelity PG approximation error we work under *primal and dual approximation properties*: for $0 < t \leq \bar{t}$ and $0 < t' \leq \bar{t}'$, and for $0 < h \leq h_0$, there hold

$$\begin{aligned} \forall v \in \mathcal{X}_t \ : \quad & \inf_{v_h\in\mathcal{X}_h} \|v - v_h\|_{\mathcal{X}} \ \leq \ C_t\, h^t\, \|v\|_{\mathcal{X}_t}\ , \\ \forall w \in \mathcal{Y}_{t'} \ : \quad & \inf_{w_h\in\mathcal{Y}_h} \|w - w_h\|_{\mathcal{Y}} \ \leq \ C_{t'}\, h^{t'}\, \|w\|_{\mathcal{Y}_{t'}}\ . \end{aligned} \tag{5.11}$$

In (5.11), $\{\mathcal{X}_t\}_{t\geq 0}$, $\{\mathcal{Y}_t\}_{t\geq 0}$, are *scales of smoothness spaces* with

$$\begin{aligned} &\mathcal{X} = \mathcal{X}_0 \supset \mathcal{X}_1 \supset \mathcal{X}_2 \supset \cdots\ , \quad \mathcal{Y} = \mathcal{Y}_0 \supset \mathcal{Y}_1 \supset \mathcal{Y}_2 \supset \cdots\ , \quad \text{and} \\ &\mathcal{X}' = \mathcal{X}'_0 \supset \mathcal{X}'_1 \supset \mathcal{X}'_2 \supset \cdots\ , \quad \mathcal{Y}' = \mathcal{Y}'_0 \supset \mathcal{Y}'_1 \supset \mathcal{Y}'_2 \supset \cdots\ . \end{aligned} \tag{5.12}$$

The scales are assumed to be defined also for noninteger values of the smoothness parameter $t \geq 0$ by interpolation. For self-adjoint operators, usually $\mathcal{X}_t = \mathcal{Y}_t$. We assume *uniform parametric regularity in the scale* $\{\mathcal{X}_t\}_{t\geq 0}$, ie. that

$$\forall\, 0 \leq t \leq \bar{t}: \quad \sup_{\boldsymbol{y}\in U} \|A(\boldsymbol{y})^{-1}\|_{\mathcal{L}(\mathcal{Y}'_t,\mathcal{X}_t)} < \infty\ . \tag{5.13}$$

The maximum amount of smoothness in the scale $\mathcal{X}_t$, denoted by $\bar{t}$, depends on the problem class under consideration and on the Sobolev scale: e.g., for elliptic problems in polygonal domains, it is well known that choosing for $\mathcal{X}_t$ the usual Sobolev spaces will allow (5.13) with t only in a possibly small interval $0 < t \leq \bar{t}$, whereas choosing $\mathcal{X}_t$ as Sobolev spaces with weights allow large values of $\bar{t}$ (see, e.g., [40]).

Theorem 5.4 *Assume that the inf-sup condition* (4.2) *and* (4.3) *hold. Assume, moreover, that* (2.12) *holds. Then, there exist* $\beta_h > 0$ *and* $h_0 > 0$ *such that, for every* $0 < h \leq h_0$ *and for every* $\boldsymbol{y} \in U$, *the parametric Petrov-Galerkin approximation* $q_h(\boldsymbol{y}) \in \mathcal{X}_h$ *given by* (4.1), *admits a unique, parametric*

[1] EOC = Empirical Order of Convergence

high-fidelity solution $q_h(\boldsymbol{y})$ which satisfies the uniform a-priori estimate

$$\sup_{\boldsymbol{y}\in U} \|q_h(\boldsymbol{y})\|_{\mathcal{X}} \;\leq\; \frac{1}{\beta_h}\,\|f\|_{\mathcal{Y}'}\;, \tag{5.14}$$

and the (uniform w.r. to $\boldsymbol{y}\in U$) quasi-optimality estimate holds

$$\|q(\boldsymbol{y}) - q_h(\boldsymbol{y})\|_{\mathcal{X}} \;\leq\; \left(1 + \frac{||A(\boldsymbol{y})||_{\mathcal{L}(\mathcal{X},\mathcal{Y}')}}{\beta_h}\right) \inf_{0\neq w_h\in\mathcal{X}_h} \|q(\boldsymbol{y}) - w_h\|_{\mathcal{X}}\;. \tag{5.15}$$

Proof Under Assumption 1, (2.12), the validity of the discrete inf-sup conditions for the nominal bilinear form $\mathfrak{a}_0(\cdot,\cdot)$, see (2.11), with constant $\beta_0^h > 0$ independent of h, implies (4.2) and (4.3) for the bilinear form $\mathfrak{a}(\boldsymbol{y};\cdot,\cdot)$ with $\beta_h = (1-\kappa/2)\beta_0^h > 0$. This follows from the stability (2.11) for the nominal problem and a perturbation argument via Neumann series, from (2.12). □

Proposition 5.5 *Under Assumption 1 and condition* (5.13)*, for every $f\in\mathcal{Y}'$ and for every $\boldsymbol{y}\in U$, the approximations $q_h(\boldsymbol{y})$ are stable, i.e.,* (5.14) *holds. There exists a constant $C>0$ such that for every $\boldsymbol{y}\in U$, every $f\in\mathcal{Y}'_t$ with $0<t\leq\bar{t}$, as $h\to 0$ there holds*

$$\|q(\boldsymbol{y}) - q_h(\boldsymbol{y})\|_{\mathcal{X}} \;\leq\; C\,h^t\,\|f\|_{\mathcal{Y}'_t}\;. \tag{5.16}$$

Since we are interested in the expectations of functionals of the parametric solution, we will also impose a regularity assumption on the observation functional $\mathcal{O}(\cdot)\in\mathcal{X}'$:

$$\exists\, 0<t'\leq\bar{t}:\quad \mathcal{O}(\cdot)\in\mathcal{X}'_{t'}\;, \tag{5.17}$$

and the *adjoint regularity*: for t' as in (5.17) there exists $C_{t'}>0$ such that for every $\boldsymbol{y}\in U$,

$$v(\boldsymbol{y}) = (A^*(\boldsymbol{y}))^{-1}\mathcal{O}\in\mathcal{Y}_{t'}\;,\quad \|v(\boldsymbol{y})\|_{\mathcal{Y}_{t'}} \;\leq\; C_{t'}\,\|\mathcal{O}\|_{\mathcal{X}'_{t'}}\;. \tag{5.18}$$

Often, the discretization error of observation functionals $\mathcal{O}(q(\boldsymbol{y}))$ of the parametric solution is of interest. For sufficiently regular $\mathcal{O}(\cdot)$, the error $|\mathcal{O}(q(\boldsymbol{y})) - \mathcal{O}(q_h(\boldsymbol{y}))|$ may converge faster than $\|q(\boldsymbol{y}) - q_h(\boldsymbol{y})\|_{\mathcal{X}}$ by an Aubin-Nitsche duality argument. We state the result without proof.

Proposition 5.6 *Under Assumption 1 and the conditions* (5.13) *and* (5.18)*, for every $f\in\mathcal{Y}'_t$ with $0<t\leq\bar{t}$, for every $\mathcal{O}(\cdot)\in\mathcal{X}'_{t'}$ with $0<t'\leq\bar{t}$ and for every $\boldsymbol{y}\in U$, as $h\to 0$, there exists a constant $C>0$ independent of $h>0$ and of $\boldsymbol{y}\in U$ such that the Petrov-Galerkin approximations $\mathcal{O}(q_h(\boldsymbol{y}))$ satisfy, with $0<\tau := t+t'$,*

$$|\mathcal{O}(q(\boldsymbol{y})) - \mathcal{O}(q_h(\boldsymbol{y}))| \;\leq\; C\,h^\tau\,\|f\|_{\mathcal{Y}'_t}\,\|\mathcal{O}\|_{\mathcal{X}'_{t'}}\;. \tag{5.19}$$

5.4 Reduced basis approximation

First, we prove that both the reduced problem (4.10) and its adjoint (4.27) are well-posed under the construction of trial space $\mathcal{X}_N$ and test space $\mathcal{Y}_N$ in section 4.3.

Theorem 5.7 *Under the assumption of Theorem 5.4 for the high-fidelity problem, for every $N\in\mathbb{N}$, there exists a unique reduced, parametric solution $q_N(\boldsymbol{y})\in\mathcal{X}_N$ which satisfies the a-priori estimate*

$$||q_N(\boldsymbol{y})||_{\mathcal{X}} \leq \frac{1}{\beta_N}||f||_{\mathcal{Y}'} \tag{5.20}$$

where $\beta_N\geq\beta_h$. Moreover, for all $\boldsymbol{y}\in U$, there holds the quasi-optimality estimate

$$||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}} \leq \left(1 + \frac{||A(\boldsymbol{y})||_{\mathcal{L}(\mathcal{X},\mathcal{Y}')}}{\beta_N}\right) \inf_{0\neq w_N\in\mathcal{X}_N} ||q_h(\boldsymbol{y}) - w_N||_{\mathcal{X}}. \tag{5.21}$$

Analogously, for every $N \in \mathbb{N}$, there exists a unique reduced, adjoint solution $\psi_N(\boldsymbol{y}) \in \mathcal{Y}_{N_{du}}$ such that

$$||\psi_{N_{du}}||_{\mathcal{Y}} \leq \frac{1}{\beta_{N_{du}}} \left(|\delta| + \frac{||f||_{\mathcal{Y}'}}{\beta_h} ||\mathcal{O}||_{\mathcal{X}'} \right) \cdot |\Gamma^{-1}| \cdot ||\mathcal{O}||_{\mathcal{X}'}. \tag{5.22}$$

Moreover, for every $\boldsymbol{y} \in U$, there holds the quasi-optimality estimate

$$\|\psi_h(\boldsymbol{y}) - \psi_{N_{du}}(\boldsymbol{y})\|_{\mathcal{Y}} \leq \left(1 + \frac{||A^*(\boldsymbol{y})||_{\mathcal{L}(\mathcal{Y},\mathcal{X}')}}{\beta_{N_{du}}} \right) \inf_{0 \neq v_{N_{du}} \in \mathcal{Y}_{N_{du}}} ||\psi_h(\boldsymbol{y}) - v_{N_{du}}||_{\mathcal{Y}} \, . \tag{5.23}$$

Proof For any fixed $\boldsymbol{y} \in U$, we denote by $T_{\boldsymbol{y}}$ the linear map which, for any $w_h \in \mathcal{X}_h$, yields $T_{\boldsymbol{y}} w_h$, the "supremizer" which satisfies

$$(T_{\boldsymbol{y}} w_h, v_h)_{\mathcal{Y}_h} = \mathfrak{a}(\boldsymbol{y}; w_h, v_h) \quad \forall v_h \in \mathcal{Y}_h. \tag{5.24}$$

The construction of the test space $\mathcal{Y}_N$ in (4.22), implies $T_{\boldsymbol{y}} w_N \in \mathcal{Y}_N$ for any $w_N \in \mathcal{X}_N$. We point out that the test space $\mathcal{Y}_N$ depends on the high-dimensional parameter vector $\boldsymbol{y}$. However, there holds a uniform w.r. to $\boldsymbol{y}$ discrete inf-sup condition in the RB trial and test spaces $\mathcal{X}_N$ and $\mathcal{Y}_N$:

$$\begin{aligned}
\inf_{0 \neq w_N \in \mathcal{X}_N} \sup_{0 \neq v_N \in \mathcal{Y}_N} \frac{\mathfrak{a}(\boldsymbol{y}; w_N, v_N)}{||w_N||_{\mathcal{X}} ||v_N||_{\mathcal{Y}}} &= \inf_{0 \neq w_N \in \mathcal{X}_N} \frac{\mathfrak{a}(\boldsymbol{y}; w_N, T_{\boldsymbol{y}} w_N)}{||w_N||_{\mathcal{X}} ||T_{\boldsymbol{y}} w_N||_{\mathcal{Y}}} \\
&= \inf_{0 \neq w_N \in \mathcal{X}_N} \frac{(T_{\boldsymbol{y}} w_N, T_{\boldsymbol{y}} w_N)_{\mathcal{Y}_h}}{||w_N||_{\mathcal{X}} ||T_{\boldsymbol{y}} w_N||_{\mathcal{Y}}} \\
&= \inf_{0 \neq w_N \in \mathcal{X}_N} \frac{||T_{\boldsymbol{y}} w_N||_{\mathcal{Y}}}{||w_N||_{\mathcal{X}}} \\
&\geq \inf_{0 \neq w_h \in \mathcal{X}_h} \frac{||T_{\boldsymbol{y}} w_h||_{\mathcal{Y}}}{||w_h||_{\mathcal{X}}} \\
&= \inf_{0 \neq w_h \in \mathcal{X}_h} \sup_{0 \neq v_h \in \mathcal{Y}_h} \frac{\mathfrak{a}(\boldsymbol{y}; w_h, v_h)}{||w_h||_{\mathcal{X}} ||v_h||_{\mathcal{Y}}} \geq \beta_h,
\end{aligned} \tag{5.25}$$

where the first and second equalities are due to the definition of the linear, parametric operator $T_{\boldsymbol{y}}$, the first inequality is due to $\mathcal{X}_N \subset \mathcal{X}_h$. This implies the stability of the reduced problem (4.10) with inf-sup constant $\beta_N \geq \beta_h$. As a result, the existence of a unique solution $q_N(\boldsymbol{y}) \in \mathcal{X}_N$ for every $\boldsymbol{y} \in U$, the stability estimate (5.20) and the quasi-optimality estimate (5.21) are established. The lower bound on the stability constant $\beta_{N_{du}} > \beta_h$ can be proved along the lines used to establish (5.25). Therefore, we have the stability estimate for the adjoint solution

$$||\psi_{N_{du}}||_{\mathcal{Y}} \leq \frac{1}{\beta_{N_{du}}} \left|\left| \frac{\partial \Theta_h}{\partial q_h}\Big|_{q_N(\boldsymbol{y})} \right|\right|_{\mathcal{X}'} . \tag{5.26}$$

We consider the first order Fréchet derivative given in (4.25). Since each term of this derivative can be bounded, i.e. $\Theta_N(\boldsymbol{y}) \leq 1$, δ is finite almost surely, each entry of the matrix Γ^{-1} is finite, moreover the norm of the linear functional $||o_k||_{\mathcal{X}'} < \infty$ for each $1 \leq k \leq K$ and $||q_N(\boldsymbol{y})||_{\mathcal{X}}$ is bounded by the stability estimate (5.20), we have that the norm of the first order Fréchet derivative (4.25) (as a linear functional) can also be bounded. More explicitly, as $\Theta_N(\boldsymbol{y}) = \exp(-\Phi(u(\boldsymbol{y}); \delta)) \leq 1$, we have

$$\left|\left| \frac{\partial \Theta_h}{\partial q_h}\Big|_{q_N(\boldsymbol{y})} \right|\right|_{\mathcal{X}'} \leq \left(|\delta| + \frac{||f||_{\mathcal{Y}'}}{\beta_h} ||\mathcal{O}||_{\mathcal{X}'} \right) \cdot |\Gamma^{-1}| \cdot ||\mathcal{O}||_{\mathcal{X}'} < \infty \, . \tag{5.27}$$

The quasi optimality estimate (5.23) thus follows. □

In order to establish the best N-term convergence rate for the RB approximation, we take advantage of that for the SG interpolation approximation in the high-fidelity PG spaces.

Lemma 5.8 *Under Assumption 1,* (4.2) *and* (4.3)*, there exists a sequence of downward closed index sets $(\Lambda_N)_{N \geq 1}$, with $\#(\Lambda_N) = N$, corresponding to (indices of) N largest coefficients of the gPC*

expansion of the high-fidelity PG solution $q_h(\boldsymbol{y})$ of problem (4.1) *for every $\boldsymbol{y} \in U$, with the convergence rate estimate for the generalized SG error*

$$\sup_{y\in U} ||q_h(\boldsymbol{y}) - \mathcal{S}_{\Lambda_N} q_h(\boldsymbol{y})||_{\mathcal{X}} \leq C_q^i N^{-s}, \quad s = \frac{1}{p} - 1, \tag{5.28}$$

where C_q^i is a constant independent of N and of the PG meshwidth h. The sparse grid is unisolvent for $\mathbb{P}_{\Lambda_N} = \text{span}\{\boldsymbol{y}^\nu : \nu \in \Lambda_N\}$ and the number of interpolation points equals the number of indices in Λ_N. Similarly, we have

$$\sup_{y\in U} ||\psi_h(\boldsymbol{y}) - \mathcal{S}_{\Lambda^{du}_{N_{du}}} \psi_h(\boldsymbol{y})||_{\mathcal{Y}} \leq C_\psi^i N_{du}^{-s}, \quad s = \frac{1}{p} - 1, \tag{5.29}$$

where C_ψ^i is a constant independent of N_{du} and the meshwidth h.

Proof The existence of index sets Λ_N and the dimension-independent convergence rate follow from the Assumption 1. It remains to verify the relations (2.11), (2.12) and (2.13) in this assumption in the high-fidelity PG spaces $\mathcal{X}_h$ and $\mathcal{Y}_h$. The first one (2.11) is due to the assumption (4.2). As for the second term, we have

$$||A_0^{-1}A_j||_{\mathcal{L}(\mathcal{X}_h,\mathcal{X})} := \sup_{w_h\in\mathcal{X}_h} \frac{||A_0^{-1}A_j w_h||_{\mathcal{X}}}{||w_h||_{\mathcal{X}}} \leq \sup_{w\in\mathcal{X}} \frac{||A_0^{-1}A_j w||_{\mathcal{X}}}{||w||_{\mathcal{X}}} \leq b_j, \tag{5.30}$$

so that (2.12) and (2.13) also hold in $\mathcal{X}_h$ and $\mathcal{Y}_h$ provided they hold in $\mathcal{X}$ and $\mathcal{Y}$. For each index $\nu \in \Lambda_N$, we choose the set of interpolation points Ξ_ν such that $\Xi_\nu = \times_{j\geq 1}\Xi^j_{\nu_j}$ with $\#(\Xi^j_{\nu_j}) = \nu_j + 1$ and $\Xi^j_{\nu_j} \subset \Xi^j_{\nu'_j}$ if $\nu_j \leq \nu'_j$, and such that the Lebesgue constant corresponding to $\Xi^j_{\nu_j}$ is bounded by $(1+\nu_j)^\theta$ for a finite $\theta < \infty$. For instance, nested sequences of Leja points admit polynomial Lebesgue constants with $2 < \theta \leq 3$ [16]. The convergence rate for the high-fidelity dual solution can be proved by following the same lines. In particular,

$$||(A_0^*)^{-1}A_j^*||_{\mathcal{L}(\mathcal{Y}_h,\mathcal{Y})} := \sup_{v_h\in\mathcal{Y}_h} \frac{||(A_0^*)^{-1}A_j^* v_h||_{\mathcal{Y}}}{||v_h||_{\mathcal{Y}}} \leq \sup_{v\in\mathcal{Y}} \frac{||(A_0^*)^{-1}A_j^* v||_{\mathcal{Y}}}{||v||_{\mathcal{Y}}} \leq b_j\ . \tag{5.31}$$

Moreover, Assumption (1) can be verified for the reduced forward solution q_N in the reduced spaces $\mathcal{X}_N$ and $\mathcal{Y}_N$. More explicitly, (2.11) can be verified by setting $\boldsymbol{y} = \boldsymbol{0}$ in (5.25). Estimates (2.12) and (2.13) hold by

$$||A_0^{-1}A_j||_{\mathcal{L}(\mathcal{X}_N,\mathcal{X})} := \sup_{w_N\in\mathcal{X}_N} \frac{||A_0^{-1}A_j w_N||_{\mathcal{X}}}{||w_N||_{\mathcal{X}}} \leq \sup_{w\in\mathcal{X}} \frac{||A_0^{-1}A_j w||_{\mathcal{X}}}{||w||_{\mathcal{X}}} \leq b_j. \tag{5.32}$$

Therefore, the reduced forward solution $q_N \in \mathcal{X}_N$ is holomorphic in the holomorphic domain of the solution $q \in \mathcal{X}$. We note that the right hand side of the adjoint problem (4.26) is holomorphic with respect to the parameter $\boldsymbol{y}$ ($\boldsymbol{z}$ in complex domain) with the same holomorphic domain as the reduced forward solution q_N, since it is a composition of analytic functions of the forward solution q_N, which, together with Assumption (1), imply the sparsity of the adjoint solution and the convergence rate in (5.29). □

Thanks to this quasi-optimality result and the convergence rate for the SG interpolation error, we can prove the best N-term convergence rate for the RB approximation.

Theorem 5.9 *Under Assumption 1,* (4.2) *and* (4.3)*, the RB approximation error of the forward solution, ie., the solution of the reduced problem* (4.10) *can be bounded by*

$$\sup_{\boldsymbol{y}\in U} ||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}} \leq C_r N^{-s}, \quad s = \frac{1}{p} - 1\ , \tag{5.33}$$

where the positive constant C_r does not depend on N (but depends, in general, on the observation noise covariance Γ in (2.4)*). Analogously, the RB approximation error of the adjont solution of the*

reduced problem (4.27) *is bounded by*

$$\sup_{\boldsymbol{y}\in U} ||\psi_h(\boldsymbol{y}) - \psi_{N_{du}}(\boldsymbol{y})||_{\mathcal{Y}} \leq C_r^{du} N_{du}^{-s}, \quad s = \frac{1}{p} - 1 \,, \tag{5.34}$$

where the positive constant C_r^{du} *does not depend on the number of reduced dual bases* N_{du}*.*

Proof Lemma 5.8 implies that there exists a sequence of interpolation operators $(\mathcal{S}_{\Lambda_N})_{N\geq 1}$ satisfying the convergence rate for the interpolation error and that there are N interpolation points for the index set Λ_N. We construct the RB space $\mathcal{X}_N$ based on these N points. By the quasi-optimality of the RB approximation (5.21), we have

$$\sup_{\boldsymbol{y}\in U} ||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}} \leq C_a \sup_{\boldsymbol{y}\in U} \inf_{0\neq w_N\in\mathcal{X}_N} ||q_h(\boldsymbol{y}) - w_N|| \leq C_a \sup_{\boldsymbol{y}\in U} ||q_h(\boldsymbol{y}) - \mathcal{S}_{\Lambda_N} q_h(\boldsymbol{y})||_{\mathcal{X}} \,, \tag{5.35}$$

where the constant C_a is given by

$$C_a = \sup_{\boldsymbol{y}\in U} \left(1 + \frac{||A(\boldsymbol{y})||_{\mathcal{L}(\mathcal{X},\mathcal{Y}')}}{\beta_h}\right) < \infty, \tag{5.36}$$

and the second inequality is due to that the interpolation operator is linear, i.e. $\mathcal{S}_{\Lambda_N} q_h(\boldsymbol{y}) \in \mathcal{X}_N$. Therefore, a combination of (5.28) and (5.35) leads to the result (5.33) with $C_r = C_a C_q^i$. The result (5.34) can be obtained by following the same argument with $C_r^{du} = C_a C_\psi^i$.

□

Remark 5.4 *An alternative proof for this theorem in the case that* $\mathcal{X}$ *and* $\mathcal{Y}$ *are Hilbert spaces is to use the recent result in [18] for the bound of the RB best approximation error, ie., the Kolmogorov* N*-width estimate* $d_N \leq CN^{-s}$*, and the abstract result in [2] that the RB error achieves the same convergence rate as* d_N*. This comparison argument does not require the holomorphy. However, this argument of proof [23] for Banach spaces* $\mathcal{X}$ *and* $\mathcal{Y}$*, only a convergence rate estimate for the RB error of the form* $CN^{-s+\alpha}$*,* $\alpha \geq 1/2$*, which is inferior to* $d_N \leq CN^{-s}$*.*

Theorem 5.10 *Under Assumption 1,* (4.2) *and* (4.3)*, the RB approximation error of the posterior density* Θ_N *in* (4.19) *can be bounded by*

$$\sup_{\boldsymbol{y}\in U} |\Theta_h(\boldsymbol{y}) - \Theta_N(\boldsymbol{y})| \leq C_\Theta N^{-s}, \quad s = \frac{1}{p} - 1, \tag{5.37}$$

where C_Θ *does not depend on* N *(but on* Γ*) provided that* N *is sufficiently large. Moreover, the RB approximation error of the corrected posterior density* $\tilde{\Theta}_N$ *in* (4.32) *is bounded by*

$$\sup_{\boldsymbol{y}\in U} |\Theta_h(\boldsymbol{y}) - \tilde{\Theta}_N(\boldsymbol{y})| \leq C_{\tilde{\Theta}} N^{-2s}, \quad s = \frac{1}{p} - 1, \tag{5.38}$$

where the constant $C_{\tilde{\Theta}}$ *does not depend on* N *(but on* Γ*) provided that* N *is sufficiently large. An analogous result also holds for the QoI* Ψ *by following the same arguments for the RB error analysis for* Ψ_N *and* $\tilde{\Psi}_N$ *when the right hand side of the dual problem* (4.26) *is replaced by the Fréchet derivative of* Ψ*.*

Proof By definition, the RB error for the posterior density can be written as

$$\begin{aligned}
\Theta_h(\boldsymbol{y}) - \Theta_N(\boldsymbol{y}) &= \exp\left(-\frac{1}{2}(\delta - \mathcal{O}_h(q_h(\boldsymbol{y})))^\top \Gamma^{-1}(\delta - \mathcal{O}_h(q_h(\boldsymbol{y})))\right) \\
&\quad - \exp\left(-\frac{1}{2}(\delta - \mathcal{O}_h(q_N(\boldsymbol{y})))^\top \Gamma^{-1}(\delta - \mathcal{O}_h(q_N(\boldsymbol{y})))\right) \\
&= \Theta_N(\boldsymbol{y}) \left(\exp(\triangle) - 1\right) \\
&= \Theta_N(\boldsymbol{y}) \left(\triangle + \frac{1}{2!}\triangle^2 + \frac{1}{3!}\triangle^3 + \cdots\right),
\end{aligned} \tag{5.39}$$

where $\mathcal{O}_h$ is the high-fidelity approximation of $\mathcal{O}$ with $||\mathcal{O}_h||_{\mathcal{X}'} < \infty$ and

$$\begin{aligned}\triangle &= \frac{1}{2}\left((\delta - \mathcal{O}_h(q_N(\boldsymbol{y})))^\top \Gamma^{-1}(\delta - \mathcal{O}_h(q_N(\boldsymbol{y}))) - (\delta - \mathcal{O}_h(q_h(\boldsymbol{y})))^\top \Gamma^{-1}(\delta - \mathcal{O}_h(q_h(\boldsymbol{y})))\right) \\ &= \frac{1}{2}(2\delta - \mathcal{O}_h(q_h(\boldsymbol{y})) - \mathcal{O}_h(q_N(\boldsymbol{y})))^\top \Gamma^{-1}(\mathcal{O}_h(q_h(\boldsymbol{y})) - \mathcal{O}_h(q_N(\boldsymbol{y}))) \\ &\leq \left(|\delta| + ||\mathcal{O}_h||_{\mathcal{X}'}\frac{||f||_{\mathcal{Y}'}}{\beta_h}\right)\cdot|\Gamma^{-1}|\cdot||\mathcal{O}_h||_{\mathcal{X}'}||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}} =: \triangle_+ < \infty,\end{aligned} \tag{5.40}$$

where in the last inequality we have used the stability estimates (5.14) and (5.20) and the relation $\beta_N \geq \beta_h$ in (5.25). From (5.39), we have

$$|\Theta_h(\boldsymbol{y}) - \Theta_N(\boldsymbol{y})| \leq \Theta_N(\boldsymbol{y})\exp(\triangle_+)\triangle_+, \tag{5.41}$$

being $\triangle_+$ defined in (5.40), which yields (5.37) with the constant C_Θ given by

$$C_\Theta = \Theta_N(\boldsymbol{y})\exp(\triangle_+)\left(|\delta| + ||\mathcal{O}_h||_{\mathcal{X}'}\frac{||f||_{\mathcal{Y}'}}{\beta_h}\right)\cdot|\Gamma^{-1}|\cdot||\mathcal{O}_h||_{\mathcal{X}'}C_r < \infty. \tag{5.42}$$

Note that $\Theta_N(\boldsymbol{y}) \leq 1$ and $\exp(\triangle_+)$ is decreasing with respect to N, with $\triangle_+$ bounded by

$$\triangle_+ \leq 2\frac{||f||_{\mathcal{Y}'}}{\beta_h}\left(|\delta| + ||\mathcal{O}_h||_{\mathcal{X}'}\frac{||f||_{\mathcal{Y}'}}{\beta_h}\right)\cdot|\Gamma^{-1}|\cdot||\mathcal{O}_h||_{\mathcal{X}'}. \tag{5.43}$$

Therefore, C_Θ is uniformly bounded with respect to N and depends on $|\Gamma^{-1}|$, which could be big when small measurement noise is present, i.e. $|\Gamma^{-1}| \gg 1$. Moreover, $\Theta_N(\boldsymbol{y})\exp(\triangle_+)$ could be small even when $|\Gamma^{-1}|$ is very large, since $\Theta_N(\boldsymbol{y})$ roughly scales as $\exp(-\triangle_+)$, avoiding $C_\Theta \propto \exp(C|\Gamma^{-1}|)$.

In order to bound the error $|\Theta_h(\boldsymbol{y}) - \tilde{\Theta}_N(\boldsymbol{y})|$, we first split the error by

$$\begin{aligned}|\Theta_h(\boldsymbol{y}) - \tilde{\Theta}_N(\boldsymbol{y})| &= |\Theta_h(\boldsymbol{y}) - \Theta_N(\boldsymbol{y}) - \triangle_N^\Theta(\boldsymbol{y})| \\ &= |(\Theta_h(\boldsymbol{y}) - \Theta_N(\boldsymbol{y}) - \Theta_N(\boldsymbol{y})\triangle) + (\Theta_N(\boldsymbol{y})\triangle - r(\psi_h(\boldsymbol{y}))) + (r(\psi_h(\boldsymbol{y})) - \triangle_N^\Theta(\boldsymbol{y}))| \\ &\leq |\Theta_h(\boldsymbol{y}) - \Theta_N(\boldsymbol{y}) - \Theta_N(\boldsymbol{y})\triangle| + |\Theta_N(\boldsymbol{y})\triangle - r(\psi_h(\boldsymbol{y}))| + |r(\psi_h(\boldsymbol{y})) - \triangle_N^\Theta(\boldsymbol{y})|.\end{aligned} \tag{5.44}$$

The first term of (5.44) can be bounded by using (5.39) as

$$|\Theta_h(\boldsymbol{y}) - \Theta_N(\boldsymbol{y}) - \Theta_N(\boldsymbol{y})\triangle| \leq \frac{1}{2}\Theta_N(\boldsymbol{y})\exp(\triangle_+)\triangle_+^2 \leq C_I||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}}^2, \tag{5.45}$$

where the constant C_I is given by

$$C_I = \frac{1}{2}\Theta_N(\boldsymbol{y})\exp(\triangle_+)\left(|\delta| + ||\mathcal{O}_h||_{\mathcal{X}'}\frac{||f||_{\mathcal{Y}'}}{\beta_h}\right)^2\cdot|\Gamma^{-1}|^2\cdot||\mathcal{O}_h||_{\mathcal{X}'}^2, \tag{5.46}$$

which is uniformly bounded with respect to N as the analysis for (5.42) and $C_I \propto |\Gamma^{-1}|^2$. For the second term of (5.44), we observe that (being $r(\psi_h(\boldsymbol{y}); \boldsymbol{y})$ defined in (4.30))

$$\begin{aligned}|\Theta_N(\boldsymbol{y})\triangle - r(\psi_h(\boldsymbol{y}); \boldsymbol{y})| &= |(\mathcal{O}_h(q_N(\boldsymbol{y})) - \mathcal{O}_h(q_h(\boldsymbol{y})))^\top \Gamma^{-1}(\mathcal{O}_h(q_h(\boldsymbol{y})) - \mathcal{O}_h(q_N(\boldsymbol{y})))| \\ &\leq ||\mathcal{O}_h||_{\mathcal{X}'}^2\cdot|\Gamma^{-1}|\cdot||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}}^2,\end{aligned} \tag{5.47}$$

where we denote $C_{II} = ||\mathcal{O}_h||_{\mathcal{X}'}^2\cdot|\Gamma^{-1}|$. The third term of (5.44) is bounded as in (4.31). Therefore,

$$|\Theta_h(\boldsymbol{y}) - \tilde{\Theta}_N(\boldsymbol{y})| \leq (C_I + C_{II})||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}}^2 + \gamma_h(\boldsymbol{y})||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}}||\psi_h(\boldsymbol{y}) - \psi_N(\boldsymbol{y})||_{\mathcal{X}}, \tag{5.48}$$

which concludes the estimate (5.38) with the constant $C_{\tilde{\Theta}}$ given by

$$C_{\tilde{\Theta}} = (C_I + C_{II})C_r^2 + \gamma_h(\boldsymbol{y})C_r C_r^{du} < \infty, \tag{5.49}$$

which does not depend on N and which is uniformly bounded w.r. to $\boldsymbol{y}$ as $\gamma_h(\boldsymbol{y}) < \infty$, $\forall \boldsymbol{y} \in U$. □

5.5 Combined Error Bound

To estimate the combined effect of all approximations on the accuracy of the parametric forward map, we denote by $q_J, q_{J,h}, q_{J,h,N}$ approximate parametric forward maps of the solution $q \in \mathcal{X}$ which are obtained by dimension-truncation, high-fidelity (PG) discretization, and, finally, by RB approximation.

Theorem 5.11 *Given any $\boldsymbol{y} \in U$, there holds the following a priori error estimate*

$$||q(\boldsymbol{y}) - q_{J,h,N}(\boldsymbol{y})||_{\mathcal{X}} \leq C_p J^{-s} + C_t h^t + C_r N^{-s}, \quad s = \frac{1}{p} - 1, \tag{5.50}$$

where J is the number of truncated dimension, h is the high-fidelity discretization parameter (ie.spectral order or meshwidth), N denotes the number of reduced bases and the constants C_p, C_t and C_r are understood to be independent of h and of N, respectively and, in particular, independent of the number J of active parameters in the approximation. Then, the error in the Bayesian estimate due to generalized SG quadrature (3.13) *can be estimated by*

$$|||\mathbb{E}[q] - \mathbb{E}[\mathcal{S}_{\Lambda_M} q_{J,h,N}]|||_{\mathcal{X}} \leq C_p J^{-s} + C_t h^t + C_r N^{-s} + C_q^e M^{-s}, \quad s = \frac{1}{p} - 1, \tag{5.51}$$

where M is the cardinality of the downward closed set Λ_M and C_q^e is independent of M.

Proof Given any $\boldsymbol{y} \in U$, the global approximation error of the solution $q(\boldsymbol{y}) \in \mathcal{X}$ can be split into

$$||q(\boldsymbol{y}) - q_{J,h,N}(\boldsymbol{y})||_{\mathcal{X}} \leq \underbrace{||q(\boldsymbol{y}) - q_J(\boldsymbol{y})||_{\mathcal{X}}}_{\text{truncation}} + \underbrace{||q_J(\boldsymbol{y}) - q_{J,h}(\boldsymbol{y})||_{\mathcal{X}}}_{\text{high-fidelity}} + \underbrace{||q_{J,h}(\boldsymbol{y}) - q_{J,h,N}(\boldsymbol{y})||_{\mathcal{X}}}_{\text{reduced basis}}, \tag{5.52}$$

which represents the truncation error, the high-fidelity discretization error and the RB approximation error. The three terms can be bounded as in Proposition 5.1, Theorem 5.5, Theorem 5.9, respectively. Application of the triangular inequality to (5.52) leads to the global error estimate (5.50). As for the integration error, we have

$$|||\mathbb{E}[q] - \mathbb{E}[\mathcal{S}_{\Lambda_M} q_{J,h,N}]|||_{\mathcal{X}} \leq |||\mathbb{E}[q] - \mathbb{E}[q_{J,h,N}]|||_{\mathcal{X}} + |||\mathbb{E}[q_{J,h,N}] - \mathbb{E}[\mathcal{S}_{\Lambda_M} q_{J,h,N}]|||_{\mathcal{X}}, \tag{5.53}$$

where the first term can be bounded by (5.50). As the solution $q_{J,h,N}$ solves the reduced problem (4.10), which satisfies Assumption 1 in the reduced spaces $\mathcal{X}_N$ and $\mathcal{Y}_N$, so that the second term is bounded as in Theorem 5.3. Therefore, the estimate (5.51) follows. □

Let $\Theta_J(\boldsymbol{y})$, $\Theta_{J,h}(\boldsymbol{y})$, $\tilde{\Theta}_{J,h,N}(\boldsymbol{y})$ denote the approximate of the Bayesian posterior density $\Theta(\boldsymbol{y})$ by dimension truncation, high-fidelity discretization and RB approximation with correction (4.32), respectively. Then we have the following result.

Theorem 5.12 *Given any $y \in U$, there holds the following a priori error estimate*

$$|\Theta(\boldsymbol{y}) - \tilde{\Theta}_{J,h,N}(\boldsymbol{y})| \leq C_{\Theta,p} J^{-s} + C_{\Theta,\tau} h^\tau + C_{\tilde{\Theta}} N^{-2s}, \quad s = \frac{1}{p} - 1, \tag{5.54}$$

where the constant C_Θ is independent of J, h or N. Moreover, the integration error is bounded by

$$|\mathbb{E}[\Theta] - \mathbb{E}[\mathcal{S}_{\Lambda_M} \tilde{\Theta}_{J,h,N}]| \leq C_{\Theta,p} J^{-s} + C_{\Theta,\tau} h^\tau + C_{\tilde{\Theta}} N^{-2s} + C_{\Theta}^e M^{-s}, \quad s = \frac{1}{p} - 1\,. \tag{5.55}$$

Similar results hold for the approximations of the QoI Ψ and $\mathbb{E}[\Psi]$, with errors in the norm $\|\circ\|_{\mathcal{Z}}$.

Proof Given any $\boldsymbol{y} \in U$, the global approximation error of the density $\Theta(\boldsymbol{y})$ is split into

$$|\Theta(\boldsymbol{y}) - \tilde{\Theta}_{J,h,N}(\boldsymbol{y})| \leq \underbrace{|\Theta(\boldsymbol{y}) - \Theta_J(\boldsymbol{y})|}_{\text{truncation}} + \underbrace{|\Theta_J(\boldsymbol{y}) - \Theta_{J,h}(\boldsymbol{y})|}_{\text{high-fidelity}} + \underbrace{|\Theta_{J,h}(\boldsymbol{y}) - \tilde{\Theta}_{J,h,N}(\boldsymbol{y})|}_{\text{reduced basis}}. \tag{5.56}$$

The third term in the bound for the reduced basis error is estimated as in (5.38). To bound the first term, we write

$$|\Theta(\boldsymbol{y}) - \Theta_J(\boldsymbol{y})| = |\exp(-\Phi(\boldsymbol{y})) - \exp(-\Phi_J(\boldsymbol{y}))| \tag{5.57}$$

where the Bayesian potential $\Phi : U \to \mathbb{R}$ is defined in (2.4). The dimension-truncation of the Bayesian potential is given by

$$\Phi_J(\boldsymbol{y}) = \frac{1}{2}\left((\delta - \mathcal{O}(q_J(\boldsymbol{y})))^\top \Gamma^{-1}(\delta - \mathcal{O}(q_J(\boldsymbol{y})))\right) \geq 0\,. \tag{5.58}$$

For every $\boldsymbol{y} \in U$, the dimension-truncation error for Φ_J can be expressed as

$$\begin{aligned} |\Phi(\boldsymbol{y}) - \Phi_J(\boldsymbol{y})| &= \left|(-\delta + \mathcal{O}(q(\boldsymbol{y})) + \mathcal{O}(q_J(\boldsymbol{y})))^\top \Gamma^{-1}(\mathcal{O}(q(\boldsymbol{y}) - \mathcal{O}(q_J(\boldsymbol{y})))\right| \\ &\leq C_\Phi |\mathcal{O}(q(\boldsymbol{y}) - \mathcal{O}(q_J(\boldsymbol{y}))|\,, \end{aligned} \tag{5.59}$$

where the constant C_Φ is bounded by

$$C_\Phi \leq \left(|\delta| + 2\frac{||f||_{\mathcal{Y}'}}{\beta}||\mathcal{O}||_{\mathcal{X}'}\right) \cdot |\Gamma^{-1}|\,. \tag{5.60}$$

With the dimension truncation error bound (5.4), we obtain the first term in (5.54) with constant

$$C_{\Theta,p} = C_\Phi \tilde{C} C_p ||\mathcal{O}||_{\mathcal{X}'}. \tag{5.61}$$

The second term of high-fidelity error can be bounded similarly by using (5.19), with the constant $C_{\Theta,\tau}$ given by

$$C_{\Theta,\tau} = C_\Phi C ||f||_{\mathcal{Y}'_t} ||\mathcal{O}||_{\mathcal{X}'_{t'}}. \tag{5.62}$$

The integration error estimate (5.55) can be derived following the proof of Theorem 5.11. □

Remark 5.5 *The total computational cost of approximating $\Theta(\boldsymbol{y})$ for any $\boldsymbol{y} \in U$ and $\mathbb{E}[\Theta]$ by the high-fidelity–reduced basis–sparse grid scheme is dominated, when h is very small, by the solution of the N high-fidelity problems to construct the N-dimensional reduced basis spaces.*

We consider the approximation error of the posterior measure in Hellinger distance defined as

$$d_{\text{Hell}}(\pi,\pi') := \left(\frac{1}{2}\int_U \left(\sqrt{\frac{d\pi}{d\pi_0}} - \sqrt{\frac{d\pi'}{d\pi_0}}\right)^2 d\nu\right)^{1/2}. \tag{5.63}$$

Theorem 5.13 *The global approximation error for the posterior measure is bounded as*

$$d_{Hell}(\pi^\delta, \pi^\delta_{J,h,N,M}) \leq C(C_{\Theta,p}J^{-s} + C_{\Theta,\tau}h^\tau + C_{\tilde{\Theta}}N^{-2s} + C^e_{\tilde{\Theta}}M^{-s})\,, \tag{5.64}$$

where the constant C does not depend on J, h, N, M and is given explicitly in the proof; for sufficiently large values of J, N, M and sufficiently small h, $Z_{J,h,N,M} > 0$ and the approximate posterior measure $\pi^\delta_{J,h,N,M}$ is given by

$$\frac{d\pi^\delta_{J,h,N,M}}{d\pi_0}(\boldsymbol{y}) = \frac{\Theta_{J,h,N}(\boldsymbol{y})}{Z_{J,h,N,M}}\,. \tag{5.65}$$

Proof From the definition of the Hellinger distance we have

$$\left(d_{\text{Hell}}(\pi^\delta, \pi^\delta_{J,h,N,M})\right)^2 \leq I_1 + I_2\,, \tag{5.66}$$

where

$$I_1 = \frac{1}{Z}\int_U \left(\sqrt{\Theta(\boldsymbol{y})} - \sqrt{\Theta_{J,h,N}(\boldsymbol{y})}\right)^2 \pi_0(d\boldsymbol{y}) \; ; \tag{5.67}$$

and

$$I_2 = \left(Z^{-1/2} - Z_{J,h,N,M}^{-1/2}\right)^2 \int_U \Theta_{J,h,N}\pi_0(d\boldsymbol{y}) \; . \tag{5.68}$$

The first term can be bounded by

$$I_1 \leq \frac{1}{Z}\int_U \left(\sqrt{\Theta(\boldsymbol{y})} + \sqrt{\Theta_{J,h,N}(\boldsymbol{y})}\right)^{-2} \pi_0(d\boldsymbol{y})||\Theta - \Theta_{J,h,N}||^2_{L^\infty(U)}. \tag{5.69}$$

As the potential Φ is bounded from above for any δ by its definition (2.4), with δ finite $\mathbb{Q}_0$-almost surely, there exists $c_\delta > 0$ such that $\sqrt{\Theta(\boldsymbol{y})} \geq c_\delta$ and $\sqrt{\Theta_{J,h,N}(\boldsymbol{y})} \geq c_\delta$. More explicitly, we have

$$c_\delta = \exp\left(-\frac{1}{4}|\Gamma^{-1}|\max\left\{\left(|\delta| + \frac{||f||_{\mathcal{Y}'}}{\beta}||\mathcal{O}||_{\mathcal{X}'}\right)^2, \left(|\delta| + \frac{||f||_{\mathcal{Y}'}}{\beta_h}||\mathcal{O}_h||_{\mathcal{X}'}\right)^2\right\}\right) \tag{5.70}$$

As also the normalization constant $Z > 0$ (since its integrand $\Theta(\boldsymbol{y}) \geq c_\delta^2 > 0$ uniformly with respect to $\boldsymbol{y}$), we have

$$I_1 \leq C_1||\Theta - \Theta_{J,h,N}||^2_{L^\infty(U)} \leq C_1(C_{\Theta,p}J^{-s} + C_{\Theta,\tau}h^\tau + C_{\tilde{\Theta}}N^{-2s})^2, \quad s = \frac{1}{p} - 1 \; , \tag{5.71}$$

where $C_1 = 2c_\delta^{-2}/Z$. As for the second term I_2, we have

$$\begin{aligned} I_2 &\leq \max(Z^{-3}, Z_{J,h,N,M}^{-3})|Z - Z_{J,h,N,M}|^2 \\ &\leq C_2(C_{\Theta,p}J^{-s} + C_{\Theta,\tau}h^\tau + C_{\tilde{\Theta}}N^{-2s} + C_\Theta^e M^{-s})^2, \quad s = \frac{1}{p} - 1 \; , \end{aligned} \tag{5.72}$$

where $C_2 = \max(Z^{-3}, Z_{J,h,N,M}^{-3})$. This completes the proof. □

6 Numerical experiments

We consider a boundary value problem in the two dimensional physical domain $D = (0,1)^2$: given $\boldsymbol{y} \in U$, find $q(\boldsymbol{y}) \in H_0^1(D)$ such that

$$\begin{cases} -\text{div}(u(\boldsymbol{y})\nabla q(\boldsymbol{y})) &= f \quad \text{in } D, \\ q(\boldsymbol{y}) &= 0 \quad \text{on } \partial D, \end{cases} \tag{6.1}$$

where homogeneous Dirichlet boundary condition is prescribed over the entire boundary for simplicity. For the high-fidelity approximation, we use finite element method with piecewise linear polynomial basis on a uniform mesh of size $h = 2^{-n}$, $n \in \mathbb{N}_+$. The diffusion coefficient of the problem (6.1) is described by the affine parametric function (2.5). We assume that the K ($K = k^2$, $k \in \mathbb{N}_+$) observation data are given by Gaussian convolutions (signifying, for example, sensors such as transducers) of the forward solution at randomly sampled parameters $\boldsymbol{y} \in U$,

$$o_k(q) = \text{Gauss}(q; x_k, r_k) := \int_D \frac{1}{\sqrt{2\pi}r_k}\exp\left(-\frac{(x - x_k)^2}{2r_k^2}\right) q dx, \tag{6.2}$$

where the locations $x_1, \ldots, x_K$ are uniformly distributed inside the domain D and the width $r_1 = \cdots = r_K$ is such that an r_k ball about x_k is contained in D. The covariance operator of the gaussian observation noise η is chosen as $\Gamma = \sigma^2 I$, being I the $K \times K$ identity matrix.

In the numerical experiment, we select the bases in the affine diffusion coefficient field (3) to be

locally supported in the equal sized subdomains $D_j \subset D$, $j \in \mathbb{J}$; more explicitly,

$$\psi_0 = 1 \text{ and } \psi_j(x) = 0.95 j^{-\alpha} \chi_{D_j}(x), \quad j \in \mathbb{J},$$

where χ_{D_j} is the characteristic function supported in the subdomain D_j; $1 < \alpha \in \mathbb{R}_+$ is a scaling parameter.

6.1 Sparse grid approximation error

To test the SG approximation error for both the interpolation of the density $\Theta(\boldsymbol{y})$ and Z, we use both the interpolation and integration error indicators $\mathcal{E}_i$ and $\mathcal{E}_e$ defined in (3.14) and (3.15). The interpolation and integration errors are computationally estimated by evaluating $\mathcal{E}_i(\mathscr{A})$ and $\mathcal{E}_e(\mathscr{A})$ defined in (3.16). We specify the meshwidth $h = 2^{-5}$, the number of observations $K = 9$, the width $r_1 = 0.1$, and observation noise covariance $\Gamma = \sigma I$ with standard deviation $\sigma = 0.01$, the index set $\mathbb{J} = \{1, \ldots, 64\}$ and scaling parameter $\alpha = 2$. The realization of the parameters y_j, $j \in \mathbb{J}$, is randomly sampled from the uniform distribution $\mathcal{U}(-1, 1)$. The maximum number of interpolation points to construct the generalized SG is limited in all experiments by 10^4.

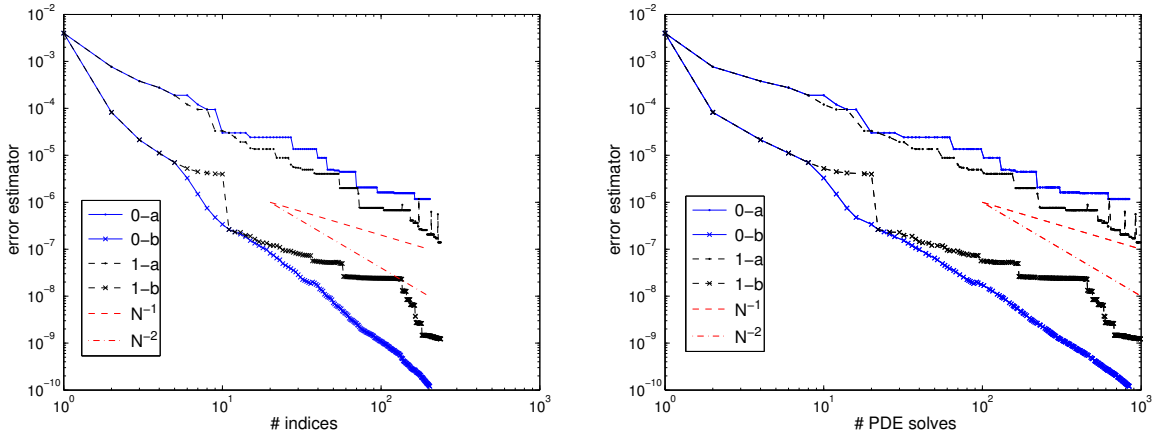

Figure 6.1: Interpolation error estimate (a) and integration (b) error estimate with respect to the number of indices (left) and number of PDE solves (right) with the generalized SG constructed by integration error indicator (0) and interpolation (1) error indicator.

Figure 6.1 displays the decay of the interpolation and integration error estimates $\mathcal{E}_i(\mathscr{A})$ with respect to the number of indices (left) and the number of interpolation nodes or PDE solves (right) with the SG constructed by two different error indicators. From the comparison, the interpolation error indicator appears to better steer the SG construction for interpolation in that it leads to faster convergence of the interpolation error, at least in the particular example considered here. In contrast, the comparison also shows that the integration error indicator affords better steering of the adaptive SG construction for integration. Moreover, the decay of interpolation error estimate with respect to the number of indices in this example appears to be better than the rate M^{-s} with $s = 1/p - 1 = \alpha - 1 = 1$ provided by the theoretical upper bound: we observe a convergence of approximately M^{-s} with respect to the number M of PDE solves, which supports Remark 5.1. The empirical integration error estimate indicates the convergence rate M^{-2}, which is however underestimated by the a priori error estimate M^{-1} in Theorem 5.3.

6.2 High-fidelity approximation error

The previous experiments were carried out with a high-fidelity Galerkin FEM based on a uniform quadrilateral mesh of meshwidth $h = 2^{-5}$. To verify the convergence of the high-fidelity PG discretization error and to verify that the number of reduced bases is substantially independent of the PG discretization parameter, we perform a sequence of runs for a fixed value of the parameter, and

with mesh sizes $h = 2^{-n}$, $n = 4, 5, 6, 7, 8$. We monitor the error $|\mathcal{O}_h(q_h(\bar{\boldsymbol{y}})) - \mathcal{O}_{\bar{h}}(q_{\bar{h}}(\bar{\boldsymbol{y}}))|$ with $\bar{h} = 2^{-8}$ and $\bar{\boldsymbol{y}}$ the parameter instance from the previous section in the SG approximation error test.

The left part of Figure 6.2 displays the convergence of the piecewise linear finite element approximation error with a least-squares fitted convergence rate $h^{-2.13}$. The tolerance is set to $\epsilon_t = 10^{-7}$ and the maximum number of nodes in the SG as 1000 to construct the reduced spaces. The number of reduced bases with respect to the mesh size is displayed on the right part of Figure 6.2, which increases by 7 from a very coarse mesh ($1/h = 16$) to a mesh of size $1/h = 32$, and only by 1 from $1/h = 32$ to $1/h = 64$. It does not change any more when refining the mesh from $1/h = 64$ to $1/h = 128$. This observation confirms the fact that once the high-fidelity PG approximation is sufficiently fine, the number of reduced bases depends essentially on the dimension of the *exact parametric* solution manifold rather than on the dimension of the high-fidelity space. Therefore, we expect to preserve the low-dimensional RB approximation when the high-fidelity degree of freedom becomes so large that solving high-fidelity problem at many (millions or more) samples becomes computationally unaffordable, whereas it is favorable to solve the high-fidelity problem at a limited number (tens or a few hundreds) of samples in order to construct reduced spaces and solve the corresponding reduced problems at all the other (millions or more) samples.

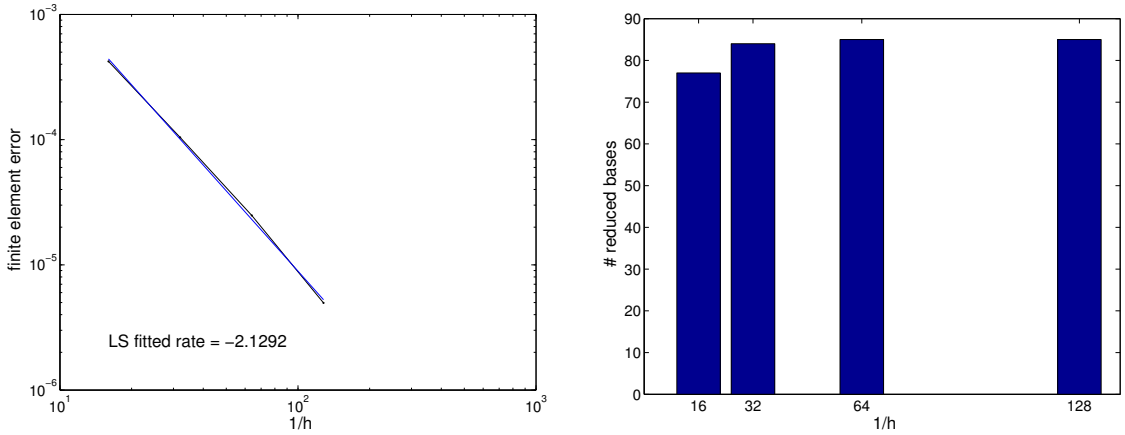

Figure 6.2: Left: decay of finite element error $|\mathcal{O}_h(q_h(\bar{\boldsymbol{y}})) - \mathcal{O}_{\bar{h}}(q_{\bar{h}}(\bar{\boldsymbol{y}}))|$ with respect to the mesh size $(1/h)$; right: the number of reduced bases (constructed with tolerance 10^{-7}) with respect to $1/h$.

6.3 Reduced basis approximation error

We first test the effectivity of the a posteriori error estimate (the dual-weighted residual) (4.28)

$$\text{effectivity} \ = \frac{\triangle_N^{\Theta}(\boldsymbol{y})}{|\Theta_h(\boldsymbol{y}) - \Theta_N(\boldsymbol{y})|}, \tag{6.3}$$

which is computed pointwise at all nodes in the sparse grid. The tolerance in the adaptive greedy algorithm 3 is set as $\epsilon_t = 10^{-8}$. The left part of Figure 6.3 shows the true error and the effectivity, from which we can observe that the effectivity is very close to one, i.e. $\triangle_N^{\Theta}$ is very close to the true error, which is uniformly smaller than the tolerance $\epsilon_t = 10^{-8}$. We also plot the ratio E_2/E_1 between the errors $E_2 = |\Theta_h(\boldsymbol{y}) - \tilde{\Theta}_N(\boldsymbol{y})|$ and $E_1 = |\Theta_h(\boldsymbol{y}) - \Theta_N(\boldsymbol{y})|$ for the approximates $\tilde{\Theta}_N$ and Θ_N, shown in the right part of Figure 6.3, which demonstrates that correction of the approximate posterior density Θ_N by the dual-weighted residual substantially improves the accuracy of the $\tilde{\Theta}_N$ in (4.32).

In another test, we use different tolerances ϵ_t for the construction of reduced spaces in order to demonstrate the accuracy of the RB approximate $\tilde{\Theta}_N$ and its integration. We set $\epsilon_t = 10^{-4}, 10^{-5}, 10^{-6}$ for the interpolation construction and $\epsilon_t = 10^{-5}, 10^{-7}, 10^{-9}$ for the integration construction. The convergence rate indicated by the interpolation error estimator and the integration error estimator are displayed in Figure 6.4. As the tolerance decreases, the RB approximation error likewise decreases such that the RB output converges to the high-fidelity output. Also, as the tolerance decreases, the

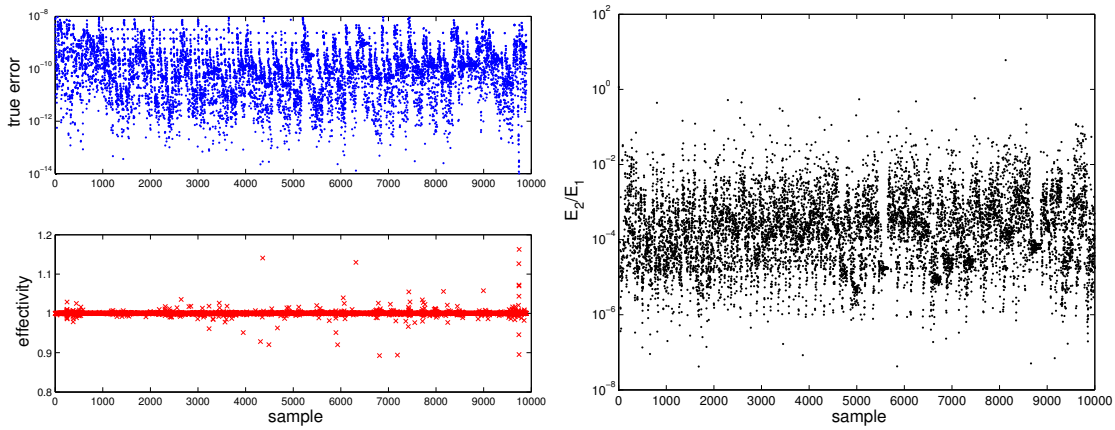

Figure 6.3: Left: true error (top) and effectivity (bottom); right: the effect of correction E_2/E_1.

number of reduced bases increases for both interpolation and integration (particularly for integration in this test case as the number of reduced bases becomes relatively large). To choose proper tolerances for the construction of the RB spaces, since the number of reduced bases remains essentially unchanged when increasing the number of high-fidelity degrees of freedom, we may first construct RB spaces based on a low-fidelity approximation (with fewer degrees of freedom), and subsequently choose a suitable tolerance for construction of the reduced spaces based on the high-fidelity model.

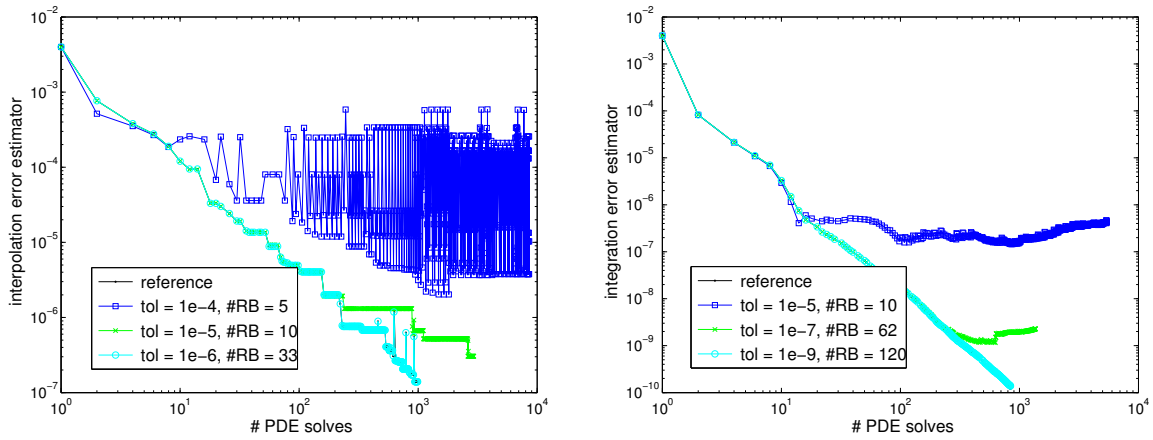

Figure 6.4: Decay of interpolation (left) and integration (right) error estimates at different tolerances.

In a new numerical experiment, we study the convergence of the RB approximation error with respect to the number of reduced bases. In the first part of the experiment, we set the parameter dimension as 64 and in the second part, we set it as 256. The reduced spaces are constructed by the greedy Algorithm 2 with the tolerance $\epsilon_t = 10^{-16}$, $N_{max} = 100$, and the training set Σ_{train} consisting of 1000 randomly samples. The reduced relative error for the output $\Theta(\boldsymbol{y})$ is taken as the maximum in a test set Σ_{test} consisting of another 100 randomly samples, i.e. $\max_{\boldsymbol{y}\in\Sigma_{test}} |\Theta_h(\boldsymbol{y}) - \Theta_N(\boldsymbol{y})|/\Theta_N(\boldsymbol{y})$. The decay of the RB approximation error is displayed in Figure 6.5, where on the left we can observe an asymptotic algebraic convergence rate $N^{-1.58}$ for the first 32 bases and $N^{-5.14}$ for the last 68 bases in the 64 dimensional approximation, while on the right we observe an asymptotic algebraic convergence rate $N^{-1.80}$ in the 256 dimensional approximation. The algebraic convergence behavior confirms the theoretical prediction given in section 5.4, however, the convergence rate $N^{-1.58}$ for the 64 dimensional case and $N^{-1.80}$ for 256 dimensional case demonstrates faster convergence than the theoretical prediction with convergence rate $N^{1-\alpha} = N^{-1}$, being the scaling parameter $\alpha = 2$. The

convergence rate $N^{-5.14}$ is higher than the theoretical prediction for the 64 dimensional approximation is due to all 64 dimensions having been sufficiently "resolved" by the 32 bases (with relative error around 10^{-5}) and the remaining 68 bases provide faster approximation.

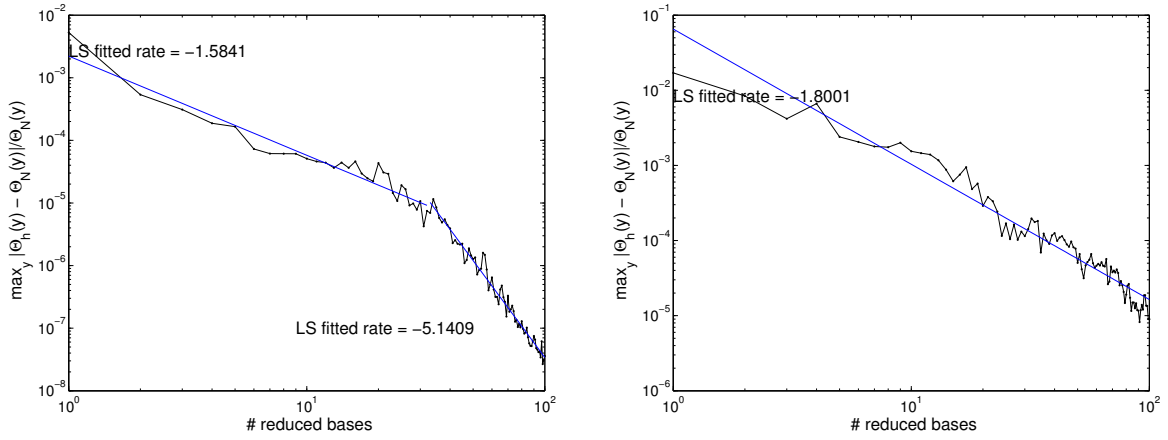

Figure 6.5: Decay of RB approximation error with respect to the number of reduced bases; left: 64 dimensions, EOC for the first 32 bases and the remaining 68 bases; right: 256 dimensions.

We run the same experiment by using the corrected output (4.32), which is expected to be more accurate than the value without correction. As expected, the convergence rate is larger, roughly twice the rate in Figure 6.5 (3.17 compared to 1.58, and 3.00 compared to 1.80), which demonstrates the quadratic effect (5.38). Therefore, by using the corrected output, we can achieve more accurate approximation with error decay of approximately N^{-3} compared to the decay of polynomial interpolation error of order N^{-1} in Figure 6.1. This also explains why a much smaller number of reduced bases is enough to accurately reconstruct the SG interpolation in Figure 6.4. The same experiment is

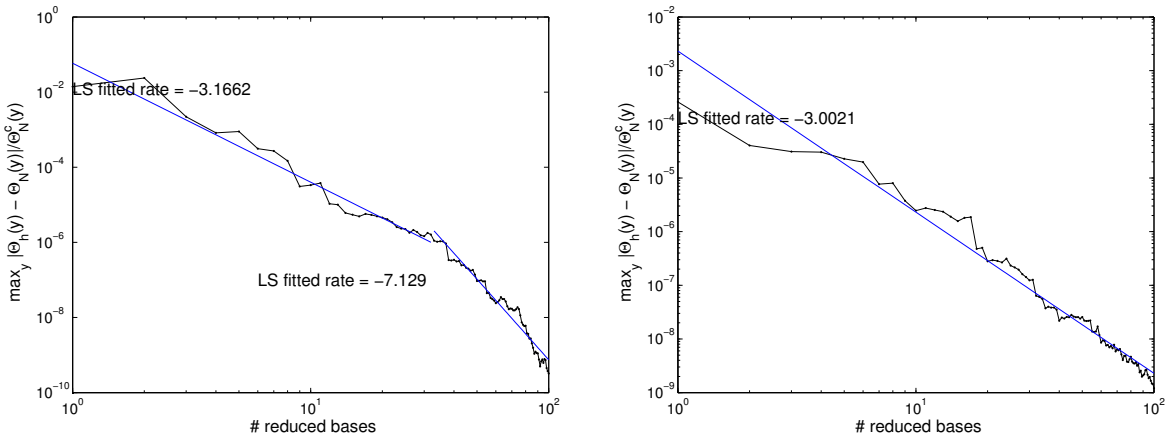

Figure 6.6: Decay of RB approximation error with the primal-dual corrected output (4.32); left: 64 dimensions, EOC for the first 32 bases and the other 68 bases; right: 256 dimensions.

also run to test the RB error of the QoI Ψ in (2.22) with $\phi(u(\boldsymbol{y})) = \mathcal{G}(\boldsymbol{y})$, ie. the observation of the forward solution, in 256 dimensions. Figure 6.7 shows the convergence rates of the RB errors without (left) and with (right) correction by the primal-dual error estimation approach, from which we can observe similar approximation property of RB for QoI to RB for the posterior density.

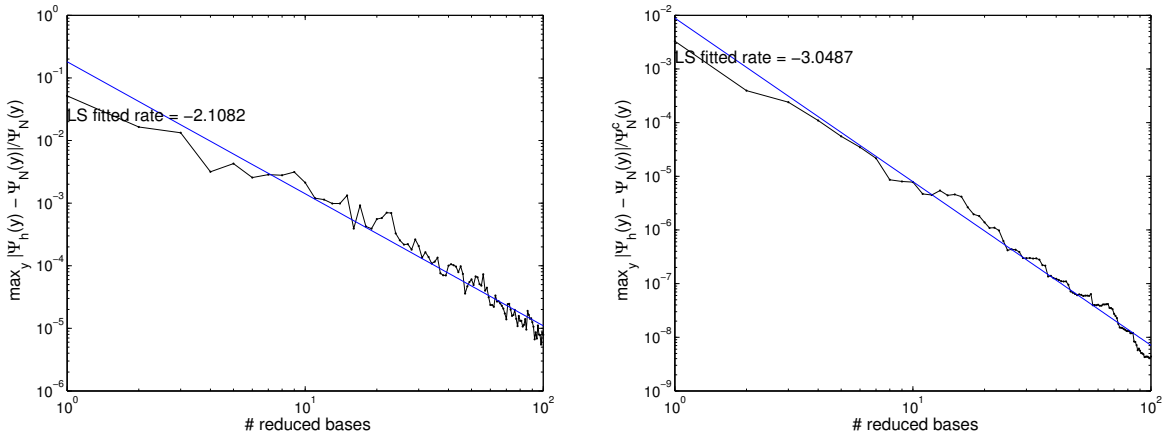

Figure 6.7: Decay of RB error for QoI without (left) and with (right) correction in 256 dimensions.

6.4 Sensitivity w.r. to the noise realizations η

We perform two numerical experiments in order to study the robustness of the computationally determined reduced bases with respect to observation noise realizations used during their construction. In both tests we construct the reduced spaces using a relative error estimate based on the dual-weighted residual, $\triangle_N^{(2)}(\boldsymbol{y})/\Theta_N(\boldsymbol{y})$, and set the tolerance as 10^{-5}. In the first test we use the same variance σ of the observation noise covariance $\Gamma = \sigma^2 I$ and ten random noise samples to numerically construct the reduced bases. Our numerical results on the number of reduced bases constructed individually at each of the noise samples are displayed on the left part of Figure 6.8. The results indicate that in this example, the variation of the number of reduced bases with respect to observation noise samples practically irrelevant, in the considered range of observations.

Moreover, when using reduced bases constructed from one noise sample (the last sample on the left part of Figure 6.8) to evaluate the density $\Theta_N(\boldsymbol{y})$ at a different observation noise sample (here the first sample), the relative error estimate remains mostly smaller than or very close to the preset tolerance at all of the 10000 parameter samples (nodes on the SG), as can be observed from the right part of the Figure 6.8. Therefore, the reduced bases are rather robust with respect to the realization of the random noise. Once constructed at some noise sample, the reduced bases can be directly used for other new observations.

In the second test, we consider the dependence of the number of reduced bases on the standard deviation of the observation noise. We specify $\sigma = \bar{\sigma} 2^n$, $\bar{\sigma} = 0.01$ (the value used in the preceding tests), $n = 2, 1, 0, -1, -2, -3, -4, -5, -6$, for the observation noise covariance operator Γ and construct reduced bases individually at each standard deviation (the $\mathcal{N}(0, \Gamma)$ gaussian observation noise being randomly sampled according to $\Gamma = \sigma I$, $\sigma = \bar{\sigma} 2^n$). From the left part of Figure 6.9 we can see that the number of reduced bases evidently increases when the standard deviation σ decreases and becomes stable as $\sigma \to 0$. The right part of Figure 6.9 shows the relative error estimate at the standard deviation $\sigma = \bar{\sigma} 2^{-6}$, evaluated by the reduced bases constructed from observation noise sampled with noise covariance $\Gamma = \sigma I$ with $\sigma = \bar{\sigma} 2^2$. In contrast to the previous test, the estimated error increases, but still remains close to the prescribed tolerance for many parameter samples. This observation suggests that when decreasing the observation noise covariance, it is necessary to enrich the reduced bases in order to preserve the same approximation accuracy.

6.5 Dependence on the location and the number of sensors in $\mathcal{O}$

We consider two cases of the influence of the sensors to the number of reduced bases: the locations $x_k \in D$ in (6.2) of the sensors and the number of the sensors. In the first case, we assume only one sensor ($K = 1$) located in nine different positions, $(x_1, x_2) = (m/4, n/4)$ where $m, n = 1, 2, 3$, the index of the sensor being given by $3(n-1) + m$. We construct the reduced bases with relative error estimate and error tolerance 10^{-5} as before. The number of reduced bases at different locations is

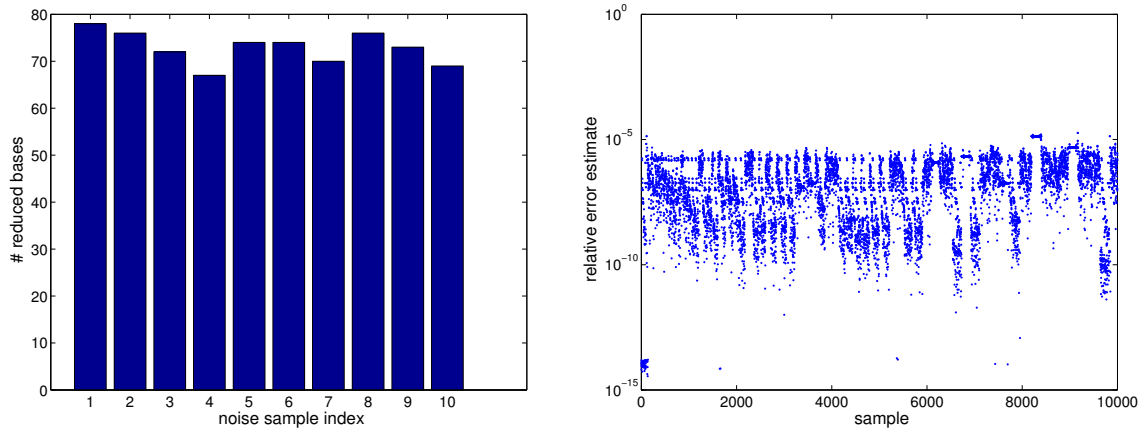

Figure 6.8: Left: the number of reduced bases at ten different noise samples; right: the relative error estimate at one noise sample evaluated by the reduced bases which were optimized from an independent noise sample

.

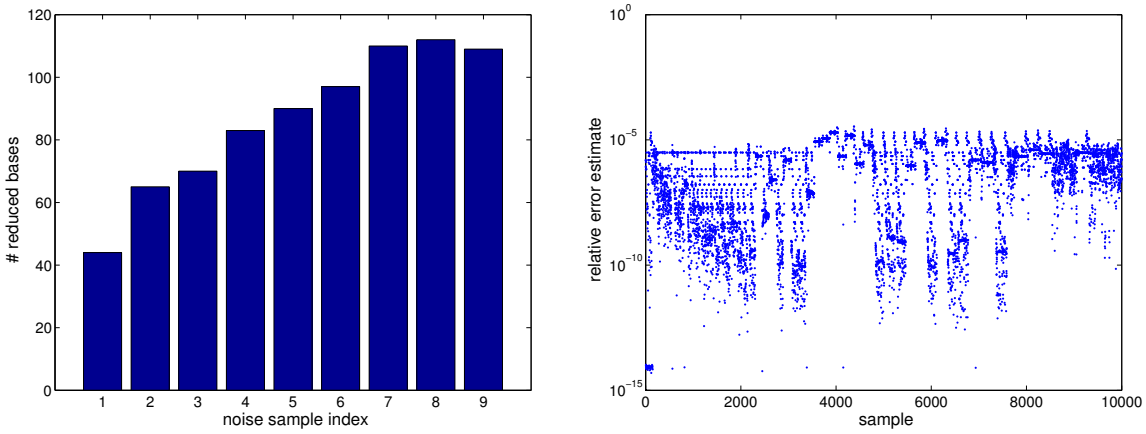

Figure 6.9: Left: the number of reduced bases at $\sigma = \bar{\sigma}2^n$, $n = 2, 1, 0, -1, -2, -3, -4, -5, -6$; right: the relative error estimate at $\sigma = \bar{\sigma}2^{-6}$ evaluated by the reduced bases constructed from $\sigma = \bar{\sigma}2^2$.

displayed in Figure 6.10, from which we can observe that the number of reduced bases varies in a large extent depending on the location of the sensor.

As n increases from 1 to 3 (corresponding to j increasing in $\mathbb{J}$), the number of reduced bases increases for all $m = 1, 2, 3$. In the second case, we set the number of sensors $K = n^2$, $n = 1, 3, 5, 7, 9, 11, 13, 15, 17$. When increasing the number of sensors, more reduced bases are constructed, particularly when the number of sensors, K, is small; the number of reduced bases appears to become stable once the number K of sensors is sufficiently large, as shown on the right part of Figure 6.10. This is evidence that the solution manifold is well approximated by the reduced space.

7 Concluding remarks

We proposed and analyzed reduced basis acceleration of the deterministic Bayesian inversion algorithms. For forward models given by linear and affine-parametric operator equations with uncertain distributed parameters and for general Petrov-Galerkin discretizations (covering, in particular, mixed variational formulations of elliptic and parabolic PDEs with uncertain coefficients), we established the best N-term convergence rate for the RB approximation of the parametric forward solution and the

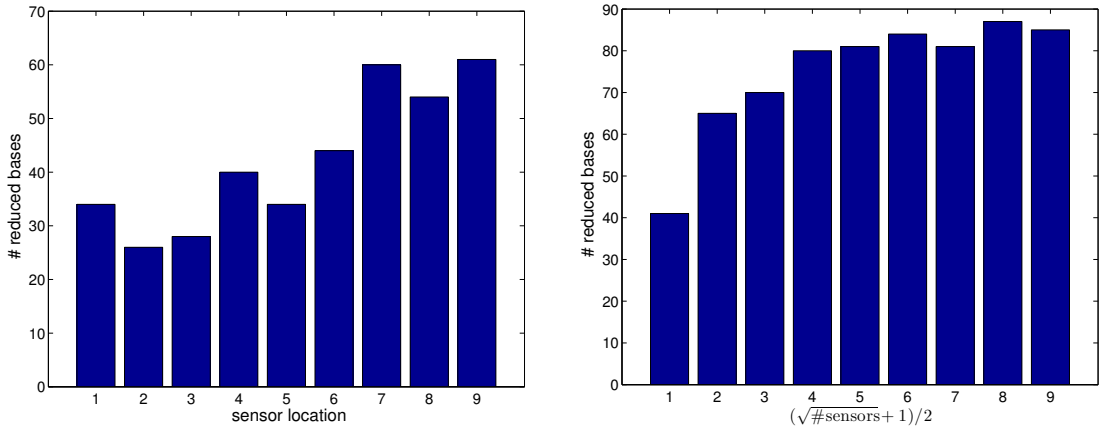

Figure 6.10: Dependence of the number of reduced bases on the location $(x_1, x_2) = (m/4, n/4)$, index $= 3(n-1)+m$, of one sensor (left) and on the number of sensors $K = n^2$, $n = 1, 3, \ldots, 17$ (right).

posterior density, independent of the dimension of the parameter space.

In particular, by numerically solving appropriate dual problems derived from the Fréchet-derivative of the parametric, deterministic Bayesian posterior density, and by correcting the posterior density evaluation with the corresponding dual-weighted residual, we obtained quadratic convergence of the RB approximation for the (nonaffine and nonlinear) posterior density (as compared to the N-width for the forward solution). We present numerical experiments which confirm the theory and indicate its practical relevance, in particular dramatic reduction of cost in computational Bayesian inversion and evaluation of the related QoIs.

We carried out an analysis of the a-priori error estimate for the combined error, ie., with dimension truncation, PG high-fidelity discretization, SG interpolation and integration, and RB approximation. The theoretical error bounds of each error contribution were tested individually in numerical experiments for a model diffusion problem with uncertain, distributed diffusion coefficient. Our theoretical a-priori error bounds were found to be sharp for the high-fidelity discretization and the SG interpolation. The experiments also revealed that the bounds for RB approximation and for the SG integration (one order underestimate from N^{-s-1} to N^{-s}) can likely be improved; To balance the error contribution from each approximation, heuristic strategies to estimate and control the various error components by sharp a posteriori error estimate are required. Particular difficulties arise from (1) the balance of the RB approximation error and of the high-fidelity PG discretization error (the former being measured with respect to the high-fidelity PG discretized quantities), (2) the balance of the RB approximation error with the SG integration error because the former is measured in the worst case while the latter is in average. Accurate and efficient algorithms for a combined adaptation will be addressed elsewhere.

Additional challenges towards the application of our algorithms for more complex engineering problems are due to nonaffine structure of the uncertainty/the random field with respect to the parameters/random variables, for instance for shape uncertainty or in the presence of log-normal random field; often, the forward models are, in addition, nonlinear. E.g. nonlinear, hyperbolic conservation laws, geometric or material nonlinearity in computational mechanics, nonlinear scattering theory, etc. The present theory will generalize, *provided the nonlinearities are analytic*. For problems with finite smoothness in parameters and/or solutions, all algorithms are immediately applicable. However, as the presently obtained error bounds rely on holomorphic parameter dependence of forward maps and Bayesian posterior density, convergence rates for finite parametric regularity are to be addressed.

References

[1] W. Bangerth and R. Rannacher. *Adaptive finite element methods for differential equations*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, 2003.

[2] P. Binev, A. Cohen, W. Dahmen, R. DeVore, G. Petrova, and P. Wojtaszczyk. Convergence rates for greedy algorithms in reduced basis methods. *SIAM Journal of Mathematical Analysis*, 43(3):1457–1472, 2011.

[3] S. Boyaval, C. Le Bris, T. Lelièvre, Y. Maday, N.C. Nguyen, and A.T. Patera. Reduced basis techniques for stochastic problems. *Archives of Computational Methods in Engineering*, 17:435–454, 2010.

[4] L. Brutman. Lebesgue functions for polynomial interpolation-a survey. *Annals of Numerical Mathematics*, 4:111–127, 1997.

[5] T. Bui-Thanh, O. Ghattas, J. Martin, and G. Stadler. A computational framework for infinite-dimensional Bayesian inverse problems Part I: The linearized case, with application to global seismic inversion. *SIAM J. Sci. Comput.*, 35(6):A2494–A2523, 2013.

[6] T. Bui-Thanh, K. Willcox, and O. Ghattas. Model reduction for large-scale systems with high-dimensional parametric input space. *SIAM Journal on Scientific Computing*, 30(6):3270–3288, 2008.

[7] T. Bui-Thanh, K. Willcox, O. Ghattas, and B. van Bloemen Waanders. Goal-oriented, model-constrained optimization for reduction of large-scale systems. *Journal of Computational Physics*, 224(2):880–896, 2007.

[8] H.J. Bungartz and M. Griebel. Sparse grids. *Acta Numerica*, 13(1):147–269, 2004.

[9] P. Chen. *Model order reduction techniques for uncertainty quantification problems*. PhD thesis, EPFL, 2014.

[10] P. Chen and A. Quarteroni. A new algorithm for high-dimensional uncertainty problems based on dimension-adaptive and reduced basis methods. *EPFL, MATHICSE Report 09, submitted*, 2014.

[11] P. Chen and A. Quarteroni. Weighted reduced basis method for stochastic optimal control problems with elliptic PDE constraints. *SIAM/ASA J. Uncertainty Quantification*, 2(1):364–396, 2014.

[12] P. Chen, A. Quarteroni, and G. Rozza. Multilevel and weighted reduced basis method for stochastic optimal control problems constrained by Stokes equations. *Submitted*, 2013.

[13] P. Chen, A. Quarteroni, and G. Rozza. A weighted reduced basis method for elliptic partial differential equations with random input data. *SIAM Journal on Numerical Analysis*, 51(6):3163 – 3185, 2013.

[14] P. Chen, A. Quarteroni, and G. Rozza. Comparison of reduced basis and stochastic collocation methods for elliptic problems. *Journal of Scientific Computing*, 59:187–216, 2014.

[15] P. Chen, A. Quarteroni, and G. Rozza. A weighted empirical interpolation method: a priori convergence analysis and applications. *ESAIM: Mathematical Modelling and Numerical Analysis*, 48:943–953, 7 2014.

[16] A. Chkifa. On the Lebesgue constant of Leja sequences for the complex unit disk and of their real projection. *Journal of Approximation Theory*, 166:176–200, 2013.

[17] A. Chkifa, A. Cohen, and Ch. Schwab. Breaking the curse of dimensionality in sparse polynomial approximation of parametric pdes. *Journal de Mathématiques Pures et Appliquées*, 2014.

[18] A. Cohen and R. DeVore. Kolmogorov widths under holomorphic mappings. *manuscript*, 2014.

[19] A. Cohen, R. DeVore, and Ch. Schwab. Convergence rates of best N-term Galerkin approximations for a class of elliptic SPDEs. *Foundations of Computational Mathematics*, 10(6):615–646, 2010.

[20] A. Cohen, R. Devore, and Ch. Schwab. Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDE's. *Analysis and Applications*, 9(01):11–47, 2011.

[21] T. Cui, Y.M. Marzouk, and K.E. Willcox. Data-driven model reduction for the Bayesian solution of inverse problems. *arXiv preprint arXiv:1403.4290*, 2014.

[22] M. Dashti and A.M. Stuart. The bayesian approach to inverse problems. 2014.

[23] R. DeVore, G. Petrova, and P. Wojtaszczyk. Greedy algorithms for reduced bases in banach spaces. *Constructive Approximation*, 37(3):455–466, 2013.

[24] H. Elman and Q. Liao. Reduced basis collocation methods for partial differential equations with random coefficients. *SIAM/ASA Journal on Uncertainty Quantification*, 1(1):192–217, 2013.

[25] D. Galbally, K. Fidkowski, K.E. Willcox, and O. Ghattas. Non-linear model reduction for uncertainty quantification in large-scale inverse problems. *International journal for numerical methods in engineering*, 81(12):1581–1608, 2010.

[26] T. Gerstner and M. Griebel. Dimension–adaptive tensor–product quadrature. *Computing*, 71(1):65–87, 2003.

[27] B. Haasdonk, K. Urban, and B. Wieland. Reduced basis methods for parameterized partial differential equations with stochastic influences using the Karhunen–Loève expansion. *SIAM/ASA J. Uncertainty Quantification*, 1(1):79–105, 2013.

[28] M. Hansen and Ch. Schwab. Sparse adaptive approximation of high dimensional parametric initial value problems. *Vietnam Journal of Mathematics*, 41(2):181–215, 2013.

[29] V.H. Hoang and Ch. Schwab. Analytic regularity and polynomial approximation of stochastic, parametric elliptic multiscale pdes. *Analysis and Applications (Singapore)*, 11(01), 2013.

[30] V.H. Hoang and Ch. Schwab. Sparse tensor galerkin discretizations for parametric and random parabolic pdes - analytic regularity and gpc approximation. *SIAM J. Mathematical Analysis*, 45(5):3050–3083, 2013.

[31] V.H. Hoang, Ch. Schwab, and A. Stuart. Complexity analysis of accelerated mcmc methods for bayesian inversion. *Inverse Problems*, 29(8), 2013.

[32] F. Kuo, Ch. Schwab, and I.H. Sloan. Quasi-monte carlo finite element methods for a class of elliptic partial differential equations with random coefficients. *SIAM J. Numerical Analysis*, 50(6):3351–3374, 2012.

[33] T. Lassila, A. Manzoni, A. Quarteroni, and G. Rozza. A reduced computational and geometrical framework for inverse problems in hemodynamics. *International journal for numerical methods in biomedical engineering*, 29(7):741–776, 2013.

[34] C. Lieberman and K. Willcox. Nonlinear goal-oriented Bayesian inference: application to carbon capture and storage. *SIAM J. Sci. Comput.*, 36(3):B427–B449, 2014.

[35] C. Lieberman, K. Willcox, and O. Ghattas. Parameter and state model reduction for large-scale statistical inverse problems. *SIAM Journal on Scientific Computing*, 32(5):2523–2542, 2010.

[36] Y. Maday, A.T. Patera, and D.V. Rovas. A blackbox reduced-basis output bound method for noncoercive linear problems. *Studies in Mathematics and its Applications*, 31:533–569, 2002.

[37] Y. Maday, A.T. Patera, and G. Turinici. A priori convergence theory for reduced-basis approximations of single-parameter elliptic partial differential equations. *Journal of Scientific Computing*, 17(1):437–446, 2002.

[38] M. Meyer and H.G. Matthies. Efficient model reduction in non-linear dynamics using the Karhunen–Loeve expansion and dual-weighted-residual methods. *Computational Mechanics*, 31(1-2):179–191, 2003.

[39] N.C. Nguyen, G. Rozza, D.B.P. Huynh, and A.T. Patera. Reduced basis approximation and a posteriori error estimation for parametrized parabolic PDEs; application to real-time Bayesian parameter estimation. *Biegler, Biros, Ghattas, Heinkenschloss, Keyes, Mallick, Tenorio, van Bloemen Waanders, and Willcox, editors, Computational Methods for Large Scale Inverse Problems and Uncertainty Quantification, John Wiley & Sons, UK*, 2009.

[40] V. Nistor and Ch. Schwab. High-order galerkin approximations for parametric second order elliptic partial differential equations. *Mathematical Methods and Models in Applied Sciences*, 23(9):1729–1760, 2013.

[41] F. Nobile, L. Tamellini, and R. Tempone. Convergence of quasi-optimal sparse grid approximation of Hilbert-valued functions: application to random elliptic PDEs. *EPFL MATHICSE report 12*, 2014.

[42] A.T. Patera and G. Rozza. Reduced basis approximation and a posteriori error estimation for parametrized partial differential equations. *Copyright MIT, http://augustine.mit.edu*, 2007.

[43] T.N.L. Patterson. The optimum addition of points to quadrature formulae. *Mathematics of Computation*, 22(104):847–856, 1968.

[44] N. Petra, J. Martin, G. Stadler, and O. Ghattas. A computational framework for infinite-dimensional Bayesian inverse problems, Part II: Stochastic Newton MCMC with application to ice sheet flow inverse problems. *SIAM J. Sci. Comput.*, 36(4):A1525–A1555, 2014.

[45] C. Prudhomme, Y. Maday, A.T. Patera, G. Turinici, D.V. Rovas, K. Veroy, and L. Machiels. Reliable real-time solution of parametrized partial differential equations: reduced-basis output bound methods. *Journal of Fluids Engineering*, 124(1):70–80, 2002.

[46] A. Quarteroni. *Numerical Models for Differential Problems*. Springer, Milano, 2nd ed., 2013.

[47] A. Quarteroni, R. Sacco, and F. Saleri. *Numerical Mathematics*. Springer, New York, 2007.

[48] G. Rozza, D.B.P. Huynh, and A.T. Patera. Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations. *Archives of Computational Methods in Engineering*, 15(3):229–275, 2008.

[49] C. Schillings and Ch. Schwab. Sparse, adaptive Smolyak quadratures for Bayesian inverse problems. *Inverse Problems*, 29(6), 2013.

[50] C. Schillings and Ch. Schwab. Sparsity in Bayesian inversion of parametric operator equations. *Inverse Problems*, 30(6), 2014.

[51] Ch. Schwab and C.J. Gittelson. Sparse tensor discretizations of high-dimensional parametric and stochastic pdes. *Acta Numerica*, 20:291–467, 2011.

[52] Ch. Schwab and V.H. Hoang. Regularity and generalized polynomial chaos approximation of parametric and random second-order hyperbolic partial differential equations. *Analysis and Applications*, 10(3), 2012.

[53] Ch. Schwab and A.M. Stuart. Sparse deterministic approximation of bayesian inverse problems. *Inverse Problems*, 28(4):045003, 2012.

[54] S.A. Smolyak. Quadrature and interpolation formulas for tensor products of certain classes of functions. In *Doklady Akademii Nauk SSSR*, volume 4, pages 240–243, 1963.

[55] A.M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numerica*, 19(1):451–559, 2010.

Recent Research Reports

Nr.	Authors/Title

2014-26 C. Schillings and Ch. Schwab
Scaling Limits in Computational Bayesian Inversion

2014-27 R. Hiptmair and A. Paganini
Shape optimization by pursuing diffeomorphisms

2014-28 D. Ray and P. Chandrashekar and U. Fjordholm and S. Mishra
Entropy stable schemes on two-dimensional unstructured grids

2014-29 H. Rauhut and Ch. Schwab
Compressive sensing Petrov-Galerkin approximation of high-dimensional parametric operator equations

2014-30 M. Hansen
A new embedding result for Kondratiev spaces and application to adaptive approximation of elliptic PDEs

2014-31 F. Mueller and Ch. Schwab
Finite elements with mesh refinement for elastic wave propagation in polygons

2014-32 R. Casagrande and C. Winkelmann and R. Hiptmair and J. Ostrowski
DG Treatment of Non-Conforming Interfaces in 3D Curl-Curl Problems

2014-33 U. Fjordholm and R. Kappeli and S. Mishra and E. Tadmor
Construction of approximate entropy measure valued solutions for systems of conservation laws.

2014-34 S. Lanthaler and S. Mishra
Computation of measure valued solutions for the incompressible Euler equations.

2014-35 P. Grohs and A. Obermeier
Ridgelet Methods for Linear Transport Equations