Model Order Reduction Methods in Computational Uncertainty Quantification

P. Chen and Ch. Schwab

Model Order Reduction Methods
in
Computational Uncertainty Quantification
in
Handbook of Uncertainty Quantification
R. Ghanem etal. (Eds). , Springer Publ. 2016

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 September 28, 2015

Abstract

This work surveys mathematical foundations of Model Order Reduction (MOR for short) techniques in accelerating computational forward and inverse UQ. Operator equations (comprising elliptic and parabolic Partial Differential Equations (PDEs for short) and Boundary Integral Equations (BIEs for short)) with distributed uncertain input, being an element of an infinite-dimensional, separable Banach space X, are admitted. Using an unconditional basis of X, computational UQ for these equations is reduced to numerical solution of countably parametric operator equations with smooth parameter dependence.

In computational forward UQ, efficiency of MOR is based on recent sparsity results for countably-parametric solutions which imply upper bounds on Kolmogorov N-widths of the manifold of (countably-)parametric solutions and Quantities of Interest (QoI for short) with dimension-independent convergence rates. Subspace sequences which realize the N-width convergence rates are obtained by greedy search algorithms in the solution manifold. Heuristic search strategies in parameter space based on finite searches over anisotropic sparse grids in parameter space render greedy searches in reduced basis construction feasible. Instances of the parametric forward problems which arise in the greedy searches are assumed to be discretized by abstract classes of Petrov-Galerkin (PG for short) discretizations of the parametric operator equation, covering most conforming primal, dual and mixed Finite Element Methods (FEM), as well as certain space-time Galerkin schemes for the application problem of interest. Based on the PG discretization, MOR for both linear and nonlinear, affine and nonaffine parametric problems are presented.

Computational inverse UQ for the mentioned operator equations is considered in the Bayesian setting of [M. Dashti and A.M. Stuart: Inverse problems a Bayesian perspective, arXiv:1302.6989v3, this Handbook]. The (countably-)parametric Bayesian posterior density inherits, in the absence of concentration effects for small observation noise covariance, the sparsity and N-width bounds of the (countably-)parametric manifolds of solution and QoI. This allows, in turn, for the deployment of MOR techniques for the parsimonious approximation of the parametric Bayesian posterior density, with convergence rates which are only limited by the sparsity of the uncertain inputs in the forward model.

Key words:

uncertainty quantification, sparse grid, reduced basis, empirical interpolation, greedy algorithm, high-fidelity, Petrov-Galerkin, a posteriori error estimate, a priori error estimate, Bayesian inversion

1 Introduction

A core task in computational forward and inverse UQ in computational science is the numerical solution of parametric numerical models for the system of interest. Uncertainty in numerical solutions obtained from the computational model can be crudely classified as follows:

i) *modeling error*: the mathematical model under consideration does not correctly describe the physical phenomena of interest: its exact solution does not predict the QoI properly. ii) *discretization error*: the discretization of the mathematical model (the substitution of the continous mathematical model by a discrete, finite-dimensional approximation which, in principle, is solvable to any prescribed numerical accuracy on computers in float point arithmetic) whose computational realization is used in forward UQ, introduces a mismatch between the "true response" (understood as exact solution of the mathematical model) and the "computed response" obtained from the computational model, *for a given (set of) input data*, and for a *prescribed quantity of interest (QoI)*. Discretization errors comprise replacing mathematical continuum models by Finite Difference, Finite Volume or Finite Element models, and inexact solution of the finite-dimensional problems which result from discretization, continuous time models by discrete timestepping, random variables by Monte-Carlo samples and their realization by random number generators. The classical paradigm of numerical analysis requires discretizations to be *stable* and *consistent*. Particular issues for discretizations in the context of UQ are *uniform stability and consistency*, with respect to all instances of the uncertain input u. iii) *computational error*: the discretized model for the numerical computation of the QoI which is obtained from a mathematical model and its subsequent discretization is not numerically solvable *for the given input data* with the compute ressources at hand. This could be, for example, due to CPU limitations, imprecise float point arithmetic (rounding), but also due to runtime failures of hardware components, partial loss of data at runtime, etc.

Under the (strong) assumptions that modeling error i) and computational error iii) be negligible, which is to say that epistemic uncertainty is absent, i.e., the mathematical model under consideration is well-posed and accounts in principle for all phenomena of interest, and the assumption that float point computations are reliable, a *first key task in computational forward UQ is the efficient computational prediction of the QoI of a mathematical model for any given instance of uncertain input u from a space of admissible input data X*. Computational challenges arise when the space X of uncertain inputs is infinite-dimensional. E.g. for *distributed uncertain input* (such as material properties in inhomogeneous solids or fluids, shapes obtained from noisy imaging, random forcing etc.), X is a (subset of a) function space. Computational UQ then involves *Uncertainty Parametrization* via a basis of X (such as, for example, a Fourier basis in $X = L^2$), resulting in a *parametric forward problem depending on sequences $\boldsymbol{y} = (y_j)_{j \geq 1}$ of uncertain input parameters*. Probability measures on X which encode information on aleatoric uncertainty can be introduced via countably products of probability measures on sequence space. Computational UQ procedures access the parametric response or forward solutions *nonintrusively*, i.e. by numerical solution of the forward model for instances of the parameter sequence $\boldsymbol{y}$. Due to the possibly infinite dimension of space of uncertain inputs $\boldsymbol{y}$, and due to the possibly high cost of each forward solve, two *key issues in computational UQ* are a) efficient computational sampling of uncertain inputs from high (possibly infinite) dimensional parameter spaces, and b) efficient numerical solution of parametric forward models.

Two key methodologies are reviewed, which address issue a) and issue b): *sparsity* of the parametric response map and *Model Order Reduction (MOR for short)*. Sparsity refers to the possibility to represent responses of the parametric forward model with user-specified accuracy $0 < \varepsilon << 1$ with responses at $O(\varepsilon^{-1/s})$ parameter instances where the sparsity parameter $s > 0$ and $O()$ are independent of the dimension of the parameter space. MOR refers to replacing the parametric response from the mathematical model for all admissible parameters by one parsimonious, so-called *reduced order model* respectively *surrogate model* which allows fast response evaluation with certified accuracy ε for any parameter instance $\boldsymbol{y}$. The accuracy which can be achieved by MOR with M parameters, for a given set of solutions, is determined by the so-called *n-width of the solution set in the space of all possible solutions*.

Recent results on sparsity of parametric forward models are reviewed, covering in particular uncertain coefficients, loadings and domains of definition in partial differential equation models which arise in engineering and in the sciences. *Sparsity adapted interpolation schemes in parameter space* allow to build *polynomial chaos surrogates of the parametric forward response maps*. The correspond-

ing point sets in parameter space are *(generalized) sparse grids*; they are based on dimension- and order-adaptive interpolation processes described here. As Monte-Carlo sampling, the performance of collocation in these deterministic point sets is independent of the dimension of the parameter space, and their convergence rates can exceed $1/2$, provided the parametric response exhibits sufficient sparsity.

Sampling of high-dimensional parameter spaces on generalized sparse grids can be used in so-called stochastic collocation for building polynomial surrogates of the parametric forward maps, in the "training" phase of MOR, and for deterministic Smolyak quadrature over all uncertainties to evaluate response statistics in forward UQ and Bayesian estimates in inverse UQ, for example.

The chapter's focus is therefore on model order reduction and sparse, adaptive collocation methods in high-dimensional parameter spaces in computational forward and inverse UQ. All concepts are developed on abstract classes of forward problems with parametric *distributed* uncertain input data, which is to say that the uncertain data may take values in an infinite-dimensional function space. Mathematical forward models which are considered here are specified in terms of (linear or nonlinear) PDEs, of elliptic or parabolic type, with smooth nonlinearities. Key steps in the approach are:

- Uncertainty parametrization: upon choosing a basis $\boldsymbol{\Psi}$ of the space X of uncertain input data, the forward problem turns into a parametric problem. The choice of basis $\boldsymbol{\Psi}$ is, in general, highly nonunique. To ensure stability in float point realizations of the parametric forward problem care must be taken to choose *well-conditioned* bases $\boldsymbol{\Psi}$ of X. For example, in Hilbert spaces X, ideally orthonormal bases $\boldsymbol{\Psi}$ should be chosen, or at least so-called *Riesz-bases* with good Riesz constants.

- Upon choosing a (well-conditioned) basis of X, a (countably-)parametric family of parametric, deterministic problems are obtained; these are referred to in the literature typically as "$\boldsymbol{y}$-PDE". For distributed uncertain inputs u, these are, usually, *infinite-dimensional parameter sequences* $\boldsymbol{y}$.

- The (minimal, for computational UQ) requirements of well-posedness of the forward problem for all admissible instances of the uncertain input u and for continuous dependence of the solution on the input imply that *the set of all parametric solutions* $\mathcal{M} := \{q(\boldsymbol{y}) : \boldsymbol{y} \in U\} \subset \mathcal{X}$ form a submanifold of the solution space $\mathcal{X}$. For smooth (linear or nonlinear) elliptic and parabolic problems, the manifold $\mathcal{M}$ is in fact an analytic manifold; see [25] and the references therein.

- "Smooth" (analytic) dependence of PDE on uncertain inputs, resp. parameters implies exponential smallness of the n-width $\mathcal{M} \subset \mathcal{X}$. This enables, in principle, *accelerated forward solves with work which scales logarithmic in accuracy ε.*

- Determining (near-)optimal subspaces which realize Kolmogorov n-width approximation bounds is feasible using so-called *reduced basis methods*. Section provides 3 key elements of the corresponding algorithms and the related theoretical results, with particular attention to the countably-parametric problems resulting from forward problems with distributed, uncertain input data. Detailed references to the literature on these techniques are provided.

- The availability of suitably compressed approximations of forward models with uncertain inputs implies, in particular, dramatic accelerations of any algorithm which involves numerous, repeated approximate evaluations of these forward models; for instance, optimization problems under uncertainty [11, 16], PDE constrained optimization algorithms, and the corresponding inverse problems. An application to *Inverse Uncertainty Quantification by Bayesian Inversion* is presented in Section 4. Bayes' theorem provides an expression for the "most likely", expected output in a Quantity of Interest (QoI), conditional on a set of given, noisy observations of (functionals of) the forward response. The numerical realization becomes, with the mentioned uncertainty parametrization, an *infinite-dimensional integration problem* against a probability measure which is only known up to a normalization constant. Current computational approaches to deal with this problem are various variants of Monte-Carlo Methods, such as Markov Chain Monte Carlo (MCMC), sequential Monte-Carlo, etc. Due to their generally slow convergence, numerous evaluations of the forward models are necessary, for a large number of proposals (generated by the corresponding samplers) of the uncertain input data. In this setting, running the

Markov chains on a reduced basis surrogate of the forward model can afford dramatic reductions in CPU time; as shown in Section 4, the resulting computational Bayesian estimates of the expected QoI will inherit an error which is of the order of the error incurred in the MOR.

The propagation of MOR error bounds translates one-to-one to other recently developed computational methods for Bayesian inversion which circumvent the use of MC sampling. These methods rather tackle the infinite-dimensional, parametric integrals obtained by inserting the uncertainty parametrization into Bayes' formula, for example by adaptive Smolyak or higher order Quasi Monte-Carlo integration.

2 Forward UQ

By forward uncertainty quantification ("forward UQ" for short) we denotes the efficient computational realization of the uncertainty-to-solution map. The present section specifies an abstract class of smooth, possibly nonlinear, operator equations with distributed, uncertain input data which allow for efficient computational forward UQ. The common feature of this class of problems is based on *holomorphic extension* of the parametric uncertainty-to-solution maps, as described in [25] and the references there.

2.1 A class of forward problems with uncertain input data

By $\mathcal{X}$ and $\mathcal{Y}$, we denote separable Hilbert spaces with duals $\mathcal{X}'$ and $\mathcal{Y}'$, respectively.

They are used for the formulation of the mathematical model of the forward problem: system responses q (such as temperature, concentration, displacements, electric fields etc.) take values in $\mathcal{X}$, whereas loads and source terms are understood as objects in $\mathcal{Y}'$, i.e. they are assumed to act on "test functions" $v \in \mathcal{Y}$.

Admissible mathematical models take form of a *residual map* $\mathcal{R}$ which associates, for a given uncertain input $u \in X$, to each state $q \in \mathcal{X}$ a response $\mathcal{R} : q \mapsto \mathcal{R}(u; q) \in \mathcal{Y}'$. With the space X of uncertain parameters $u \in X$ being infinite-dimensional, we speak about $u \in X$ as *distributed, uncertain parameters*. Three prototypical examples are presented: diffusion with uncertain diffusion coefficient, smooth nonlinear elliptic problem in a parametric domain and a parabolic problem.

Example 2.1 *(Linear diffusion problem with uncertain diffusion coefficient) The mathematical model is set in a bounded domain $D \subset \mathbb{R}^d$ (assumed certain), with diffusion coefficient $u(x) \in L^\infty(D)$ (assumed uncertain) and a source term $f(x) \in L^2(D)$ (assumed certain), find a concentration $q(x) \in H_0^1(D)$ such that*

$$\mathcal{R}(q; u) := \operatorname{div}(u(x)\operatorname{grad}q(x)) + f(x) = 0 \quad in \quad H^{-1}(D) . \tag{2.1}$$

Here the spaces for the system response q are $\mathcal{X} = \mathcal{Y} = H_0^1(D)$, the space for the uncertain input is $X = L^\infty(D)$, and $\mathcal{Y}' = (H_0^1(D))^* = H^{-1}(D)$. Note that (2.1) is to hold in the weak or variational sense of $\mathcal{Y}' = H^{-1}(D)$; this gives rise to the variational form of the residual equation: find $q \in \mathcal{X}$ such that

$$0 =_{\mathcal{Y}} \langle v, \mathcal{R}(q; u)\rangle_{\mathcal{Y}'} = \int_D \operatorname{grad}v \cdot u(x)\operatorname{grad}q(x)dx - \int_D v(x)f(x) \quad for\ all \quad v \in \mathcal{Y} = H_0^1(D) . \tag{2.2}$$

Here, and in what follows, $_{\mathcal{Y}}\langle \cdot, \cdot\rangle_{\mathcal{Y}'}$ denotes the $\mathcal{Y} \times \mathcal{Y}'$-duality.

Analogous formulations arise for any second order, linear elliptic PDE in divergence form, such as Helmholtz-equations for time-harmonic wave propagation in random media, or in (displacement formulations) of boundary value problems in computational mechanics in solids with uncertain material properties.

Example 2.2 *In a bounded domain $D \subset \mathbb{R}^n$, and in the time interval $I = (0, T)$ for a time-horizon $0 < T < \infty$, and for the affine-parametric, elliptic operator $A(\boldsymbol{y})q = \operatorname{div}(u(x)\operatorname{grad}q(x))$ as in (2.1), for given $f(x, t)$ and for given $u_0 \in L_2(D)$, the parametric, linear parabolic evolution problem is considered*

$$B(\boldsymbol{y})q := \partial_t q - A(\boldsymbol{y})q = f , \quad q(\cdot, t)|_{\partial D} = 0 \quad in \quad (0, T) \times D , \quad q(\cdot, 0) = q_0 . \tag{2.3}$$

The parabolic, parametric evolution operator $B(\boldsymbol{y})$ in (2.3) allows for a weak residual formulation analogous to (2.1) with the Bochner spaces $\mathcal{X} = L_2(I;V) \cap H^1(I;V')$, $\mathcal{Y} = L_2(I;V) \times H$, $V = H_0^1(D)$ and $H = L_2(D)$. Here, the parametric bilinear form $B(\boldsymbol{y}; w, v)$ is defined, for $v = (v_1, v_2) \in L_2(I;V) \times H$, by

$$B(\boldsymbol{y}; w, v) := \int_I \langle \tfrac{dw}{dt}(t), v_1(t) \rangle_H + \int_D u(x, \boldsymbol{y}) \nabla w \cdot \nabla v_1 \, dx dt + \langle w(0), v_2 \rangle_H \;,$$

where $u(x, \boldsymbol{y})$ denotes the linear-parametric, isotropic diffusion coefficient as in (2.13).

An abstract setting is considered which accommodates both examples (and more general models) in a unified fashion. For a distributed, uncertain parameter $u \in X$, one considers a "forward" operator $\mathcal{R}(q; u)$ depending on u and acting on $q \in \mathcal{X}$. Assuming at our disposal a "nominal parameter instance" $\langle u \rangle \in X$ (such as, for example, the expectation of an X-valued random field u), and for $u \in B_X(\langle u \rangle; R)$, an open ball of sufficiently small radius $R > 0$ in X centered at a nominal input instance $\langle u \rangle \in X$, the nonlinear operator equation is considered

$$\text{given } u \in B_X(\langle u \rangle; R) \,, \text{find } q \in \mathcal{X} \quad \text{s.t.} \quad {}_{\mathcal{Y}'}\langle \mathcal{R}(q; u), v \rangle_{\mathcal{Y}} = 0 \quad \forall v \in \mathcal{Y} \;. \tag{2.4}$$

Given $u \in B_X(\langle u \rangle; R)$, a solution q_0 of (2.4) is called *regular at u* if and only if $\mathcal{R}(\cdot; u)$ is differentiable with respect to q and if the differential $D_q\mathcal{R}(q_0; u) \in \mathcal{L}(\mathcal{X}; \mathcal{Y}')$ is an isomorphism. For the well-posedness of operator equations involving $\mathcal{R}(q; u)$, one assumes the map $\mathcal{R}(\cdot; u) : \mathcal{X} \mapsto \mathcal{Y}'$ admits a family of regular solutions *locally, in an open neighborhood of the nominal parameter instance $\langle u \rangle \in X$.*

Assumption 1 *The structural conditions*

$$\mathcal{R}(q; u) = A(q; u) - F(u) \quad in \quad \mathcal{Y}' \;, \tag{2.5}$$

hold, and for all u in a sufficiently small, closed neighborhood $\tilde{X} \subseteq X$ of $\langle u \rangle \in X$ the parametric forward problem: for every $u \in \tilde{X} \subseteq X$, given $F(u) \in \mathcal{Y}'$, find $q(u) \in \mathcal{X}$ such that the residual equation (2.4) is well-posed. Ie., for every fixed $u \in \tilde{X} \subset X$, and for every $F(u) \in \mathcal{Y}'$, there exists a unique solution $q(u)$ of (2.4) which depends continuously on u.

The set $\{(q(u), u) : u \in \tilde{X}\} \subset \mathcal{X} \times X$ is called a *regular branch of solutions* of (2.5) if

$$\begin{aligned} \tilde{X} \ni u \mapsto q(u) \text{ is continuous as mapping from } X \mapsto \mathcal{X} \;, \\ \mathcal{R}(q(u); u) = 0 \quad in \quad \mathcal{Y}' \;. \end{aligned} \tag{2.6}$$

The solutions is called in the regular branch (2.6) *nonsingular* if, in addition, the differential

$$(D_q\mathcal{R})(q(u); u) \in \mathcal{L}(\mathcal{X}, \mathcal{Y}') \text{ is an isomorphism from } \mathcal{X} \text{ onto } \mathcal{Y}', \text{ for all } u \in \tilde{X} \;. \tag{2.7}$$

The following proposition collects well-known sufficient conditions for well-posedness of (2.5). For regular branches of nonsingular solutions given by (2.5) - (2.7), the mathematical model is well-posed if the differential $D_q\mathcal{R}$ satisfies the so-called *inf-sup conditions*. In UQ, these classical (eg. [4, 56]) conditions are to hold *uniformly with respect to the uncertain input data $u \in \tilde{X} \subseteq X$.*

Proposition 2.1 *Assume that $\mathcal{Y}$ is reflexive and that, for some nominal value $\langle u \rangle \in X$ of the uncertain input data, the operator equation (2.5) admits a regular branch of solutions (2.6). Then the differential $D_q\mathcal{R}$ at $(\langle u \rangle, q_0)$ given by the bilinear map*

$$\mathcal{X} \times \mathcal{Y} \ni (\varphi, \psi) \mapsto {}_{\mathcal{Y}'}\langle D_q\mathcal{R}(q_0; \langle u \rangle)\varphi, \psi \rangle_{\mathcal{Y}}$$

is boundedly invertible, uniformly with respect to $u \in \tilde{X}$ where $\tilde{X} \subset X$ is an open neighborhood of the nominal instance $\langle u \rangle \in X$ of the uncertain parameter. In particular, there exists a constant $\beta > 0$

such that there holds

$$\forall u \in \tilde{X} : \quad \inf_{0 \neq \varphi \in \mathcal{X}} \sup_{0 \neq \psi \in \mathcal{Y}} \frac{{}_{\mathcal{Y}'}\langle (D_q\mathcal{R})(q_0;u)\varphi, \psi\rangle_{\mathcal{Y}}}{\|\varphi\|_{\mathcal{X}}\|\psi\|_{\mathcal{Y}}} \geq \beta > 0 \,,$$
$$\inf_{0 \neq \psi \in \mathcal{Y}} \sup_{0 \neq \varphi \in \mathcal{X}} \frac{{}_{\mathcal{Y}'}\langle (D_q\mathcal{R})(q_0;u)\varphi, \psi\rangle_{\mathcal{Y}}}{\|\varphi\|_{\mathcal{X}}\|\psi\|_{\mathcal{Y}}} \geq \beta > 0 \tag{2.8}$$

and

$$\forall u \in \tilde{X} : \quad \|(D_q\mathcal{R})(q_0, u)\|_{\mathcal{L}(\mathcal{X}, \mathcal{Y}')} = \sup_{0 \neq \varphi \in \mathcal{X}} \sup_{0 \neq \psi \in \mathcal{Y}} \frac{{}_{\mathcal{Y}'}\langle (D_q\mathcal{R})(q_0;u)\varphi, \psi\rangle_{\mathcal{Y}}}{\|\varphi\|_{\mathcal{X}}\|\psi\|_{\mathcal{Y}}} \leq \beta^{-1} \,. \tag{2.9}$$

The inf-sup conditions (2.8) and (2.9) are implied, for linear, self-adjoint PDEs such as the diffusion equation in Example 2.1, by the more familiar concept of *coercivity*.

Example 2.3 *In the context of Example 2.1, one verifies with $\mathcal{Y} = \mathcal{X} = H_0^1(D)$ that*

$$_{\mathcal{Y}'}\langle (D_q\mathcal{R})(q_0;u)\varphi, \psi\rangle_{\mathcal{Y}} = \int_D \nabla\psi \cdot u(x)\nabla\varphi(x)dx \,.$$

In particular, for linear operator equations the residual $\mathcal{R}(q;u)$ is linear with respect to q and the differential $(D_q\mathcal{R})(q_0;u)$ in (2.8) does not depend on q_0. Note that uniform validity of (2.8) for all realizations of the uncertain input u implies a constraint on the set $\tilde{X} \subseteq X$ of admissible data: one may choose, for example, $X = L^\infty(D)$ and require u to take values on

$$\tilde{X} = \{u \in X : \operatorname{essinf}(u) \geq c_0 > 0\} \,.$$

Then (2.8) (and thus also (2.7)) are implied by the coercivity *of $(D_q\mathcal{R})(q_0;u)$ in $\mathcal{X} = H_0^1(D)$*

$$_{\mathcal{Y}'}\langle (D_q\mathcal{R})(q_0;u)\varphi, \varphi\rangle_{\mathcal{Y}} \int_D \nabla\varphi \cdot u(x)\nabla\varphi(x)dx \geq c_0 \int_D |\nabla\varphi(x)|^2 dx \geq c_0 \frac{1}{2}(1 + C_P)\|\varphi\|_{\mathcal{X}}^2 \,,$$

where $C_P(D) > 0$ denotes the constant in the Poincaré-inequality $\|\nabla\varphi\|_{L^2(D)}^2 \geq C_P\|\varphi\|_{L^2(D)}^2$, *valid in $\mathcal{X}$ uniformly for all inputs $u \in \tilde{X}$. The saddle point stability conditions (2.8) and the possibility of different trial and testfunction spaces are not necessary here.*

Remark 2.1 *The saddle point stability conditions (2.8) are, however, indispensable for indefinite variational problems such as the space-time formulation of the parabolic evolution problem (2.3). For (2.3), the inf-sup conditions (2.8) have been verified in [64, Appendix]. Consider, for further illustration, the* Helmholtz Equations *for the propagation of time-harmonic pressure amplitude $\pi(x,t)$ in an uncertain, linearly elastic medium, in a bounded, certain domain $D \subset \mathbb{R}^d$, with homogeneous Dirichlet boundary conditions on ∂D.*

Separation of variables $\pi(x,t) = \exp(\iota\omega t)q(x)$ implies [Notation: ω also element. event] the Helmholtz equation,

$$-\operatorname{div}(u(x)\nabla q(x)) - \omega^2 q(x) = f(x) \quad in \quad D \,, \quad q|_{\partial D} = 0 \,. \tag{2.10}$$

One chooses again $\mathcal{X} = \mathcal{Y} = H_0^1(D)$ and $u \in \tilde{X}$ as in Example 2.3. Then, for large frequency $\omega > 0$, the saddle point stability conditions (2.8) hold only if $\omega \notin \bigcup_{u \in \tilde{X}} \Sigma(A(\cdot;u))$, where $\Sigma(A(\cdot;u)) \subset \mathbb{R}_{>0}$ denotes the (discrete and countable) spectrum of the second order elliptic operator $A(q;u) := -\operatorname{div}(u(x)\nabla q(x)) \in \mathcal{L}(\mathcal{X}; \mathcal{X}')$ as is revealed by an straightforward eigenfunction argument. The stability constant β in (2.8) is $\beta = \inf_{\lambda \in \Sigma(A(\cdot;u)), u \in \tilde{X}}\{|\lambda - \omega^2|\}$.

Under conditions (2.8) and (2.9), for every $u \in \tilde{X} \subseteq X$, there exists a unique, regular solution $q(u)$ of (2.5) which is uniformly bounded with respect to $u \in \tilde{X}$ in the sense that there exists a constant $C(F, \tilde{X}) > 0$ such that

$$\sup_{u \in \tilde{X}} \|q(u)\|_{\mathcal{X}} \leq C(F, \tilde{X}) \,. \tag{2.11}$$

For (2.8) - (2.11) being valid, we shall say that the set $\{(q(u), u) : u \in \tilde{X}\} \subset \mathcal{X} \times \tilde{X}$ forms a *regular branch of nonsingular solutions*.

If the data-to-solution map $\tilde{X} \ni u \mapsto q(u)$ is also Fréchet differentiable with respect to u at every point of the regular branch $\{(q(u); u) : u \in \tilde{X}\} \subset \mathcal{X} \times \tilde{X}$, the dependence of the "forward map", i.e. the mapping relating u to $q(u)$ with the branch of nonsingular solutions, is locally Lipschitz on $\tilde{X}$: there exists a Lipschitz constant $L(F, \tilde{X})$ such that

$$\forall u, v \in \tilde{X}: \quad \|q(u) - q(v)\|_{\mathcal{X}} \leq L(F, \tilde{X})\|u - v\|_X . \tag{2.12}$$

This follows from the identity $(D_u q)(u) = -(D_q \mathcal{R})^{-1}(D_u \mathcal{R})$, and from the isomorphism property $(D_u \mathcal{R}_q)(q_0; \langle u \rangle) \in \mathcal{L}_{iso}(\mathcal{X}, \mathcal{Y}')$ which is implied by (2.8) and (2.9), and from the continuity of the differential $D_q \mathcal{R}$ on the regular branch.

In what follows, the abstract setting (2.4) is considered with uniformly continuously differentiable mapping $\mathcal{R}(q; u)$ in a product of neighborhoods $B_X(\langle u \rangle; R) \times B_{\mathcal{X}}(q(\langle u \rangle); R) \subset X \times \mathcal{X}$ of sufficiently small radius $R > 0$, satisfying the structural assumption (2.5). In Proposition 2.1 and throughout what follows, $q(\langle u \rangle) \in \mathcal{X}$ denotes the unique regular solution of (2.5) at the nominal input $\langle u \rangle \in X$.

2.2 Uncertainty parametrization

As mentioned in the introduction, a key step in computational UQ is the *parametrization of the uncertain input data* $u \in X$ in terms of an (possibly infinite) sequence $\boldsymbol{y} = (y_j)_{j \geq 1}$ of parameters, taking values in a *parameter domain* U.

In the particular case where $u \in X$ is a random variable taking values in (a subset $\tilde{X}$ of) the Banach space X, probabilistic UQ involves probability measures on the space X of uncertain input data.

In *uncertainty parametrization*, X is assumed to be separable. This is guaranteed in particular when X is finite-dimensional, i.e., when *the uncertain input data consists of a finite number of parameters*.

In the case of *distributed uncertain input* $u \in X$, the data space X is infinite-dimensional; to facilitate uncertainty parametrization, X is assumed to admit an unconditional Schauder basis $\boldsymbol{\Psi}$: $X = \text{span}\{\psi_j : j \geq 1\}$. This is in particular the case for separable Hilbert spaces X. Then, every $u \in \tilde{X} \subset X$ can be parametrized *linearly* in this basis, i.e. it admits a *parametric representation with linear dependence on the parameters* y_j:

$$u = u(\boldsymbol{y}) := \langle u \rangle + \sum_{j \geq 1} y_j \psi_j \quad \text{for some } \boldsymbol{y} = (y_j)_{j \geq 1} \in U . \tag{2.13}$$

Some examples of linear uncertainty parametrizations (2.13) are (i) Karhunen–Loève expansions which arise, in particular, from a numerical PCA of a random field model of uncertain input u (see, e.g., [66, 63, 67, 29]) (ii) unconditional Schauder bases (see, e.g., [24]), (iii) wavelet or trigonometric bases, (iv) iso-geometric geometry parametrizations from computer aided design (see, e.g., [2] and the references there).

The representation (2.13) is not unique: rescaling y_j and ψ_j will not change u. One assumes, therefore, throughout what follows that *the basis sequence* $\{\psi_j\}_{j \geq 1}$ *is normalized such that the parameter domain is* $U = [-1, 1]^{\mathbb{N}}$. However, for *gaussian random field inputs*, the assumption of bounded parameter ranges is not satisfied: see [43, 63] for parametric formulations in this case.

Note that often, in applications, the dependence of u on the parameters y_j is not granted: one distinguishes

a) *linear parametrization* (2.13),

b) *affine or separable parametrization*: the uncertain/parametric heat conductivity is given by [55]

$$u(x, \boldsymbol{y}) = \sum_{j \geq 1} \theta_j(\boldsymbol{y})\psi_j(x) , \quad \boldsymbol{y} \in U . \tag{2.14}$$

where we emphasize that for every $j \geq 1$, each coordinate function $\theta_j(\boldsymbol{y})$ may depend on *all* coordinates $y_j \in \boldsymbol{y}$. A typical example of a separable uncertainty parametrization is the so-called

thermal fin problem (cp. [55])

$$u(x, \boldsymbol{y}) = \sum_{j=1}^{J} \chi_{D_j}(x) 10^{y_j} \, , \tag{2.15}$$

where $D := \bigcup_{j=1}^{J} D_j$ is decomposed into J non-overlapping subdomains D_j, $j = 1, \ldots, J$; χ_{D_j} is a characteristic function such that $\chi_{D_j}(x) = 1$ if $x \in D_j$ and 0 otherwise; $y_j \in [-1, 1]$, $j = 1, \ldots, J$.

c) *nonlinear transformation of an affine parametrization.* This case occurs for example in log-gaussian models with a positivity constraint, such as a linear diffusion equation with a log-gaussian permeability: for illustration, consider the Dirichlet problem

$$- \nabla \cdot (a(x, \omega) \nabla u(x, \omega)) = f(x) \quad \text{in} \quad D \, , \quad u|_{\partial D} = 0 \, . \tag{2.16}$$

Here,

$$a(x, \omega) = \exp(g(x, \omega)), \quad \text{where} \quad g(x, \omega) = \sum_{j \geq 1} Y_j(\omega) \psi_j(x) \, , \quad Y_j \sim N(0, 1) \, . \tag{2.17}$$

Due to Y_j taking values in all of $\mathbb{R}$ with positive probability, in (2.17) the normalization of the terms is effected by requiring that the standard deviation of Y_j be one.

d) *nonseparable, nonlinear parametric operator equations.* Examples of this class typically arise in problems of *domain uncertainty*: upon diffeomorphic transformation of the problem to a fixed nominal domain, we obtain a parametric problem with uncertain operator whose coefficients are rational functions of the parameters. We refer the reader to Example 2.4 ahead for illustration.

Bases in the uncertain input space X are, in general, not unique, even when fixing the scaling of the coordinates y_j in (2.13). Being bases of X, they are mathematically equivalent. *In the context of computational UQ, the concrete choice of basis can have a significant impact on the numerical stability of UQ algorithms.* For illustration we mention the (textbook) example $X = L^2(-1, 1)$, for which two bases are given by $\boldsymbol{\Psi}_1 = \{1, x, x^2, \ldots\}$ and $\boldsymbol{\Psi}_2 = \{P_j(x) : j = 0, 1, 2, \ldots\}$ denoting P_j the classical Legendre polynomial of degree $j \geq 0$, with normalization $P_j(1) = 1$. Both bases, $\boldsymbol{\Psi}_1$ and $\boldsymbol{\Psi}_2$, as well as the trigonometric functions constituting a Karhunen–Loève basis of $X = L^2(-1, 1)$, are *global*, meaning that their elements are supported in the entire domain $[-1, 1]$. Alternative bases of $X = L^2(-1, 1)$ with local supports are spline wavelet bases, such as the Haar wavelet basis. Most localized bases have an intrinsic limit on the approximation order which can be reached for sufficiently smooth uncertain input data. Uncertainty parametrization with localized bases can, however, substantially increase sparsity in the parametric forward map.

Norm-convergence of the series (2.13) in X is implied by the *summability condition*

$$\sum_{j \geq 1} \|\psi_j\|_X < \infty \, , \tag{2.18}$$

"Uncertain input data" can also signify *domain uncertainty*, ie. the shape of the physical domain D in which the boundary value problem is considered is uncertain. Upon suitable domain parametrization, such problems also are covered by the ensuing parametric, variational formulation; domain uncertainty in the physical domain can be reduced by domain mapping to a parametric problem in a fixed, nominal domain D_0. This was considered for (2.2) with parametric coefficient u depending on the shape of the [25] for a particular example, and the following, smooth and nonlinear elliptic problem from [20]; see also [48] for applications to artery variability in Hemodynamics, to [42] for application of reduced basis techniques in electromagnetic scattering.

Example 2.4 *(Domain uncertainty) A nonlinear operator equation in a random domain [25] is considered. The basic approach to dealing with domain uncertainty consists in* domain mapping *to a fixed reference domain, and in transforming the mathematical model from the physical domain to a fixed reference domain which is assumed known. Note that in certain classes of mathematical models, such*

as for example, elastic deformation of a continuous medium, such reference domains arise naturally; in mathematical elasticity, the reference domain would be referred to as reference configuration. *Note also that the reference domain may not necessarily be attained by concrete realizations of the uncertain input. Also, in stochastic domain modelling, the reference domain need not coincide with the nominal domain.*

In transforming to a fixed reference domain, parametric domain uncertainty is transferred to the differential operator on the reference domain. Due to the smooth, but highly nonlinear nature of the domain transformations, the resulting parametric differential operators on the reference domain exhibit, as a rule, highly nonlinear, rational dependence w.r. to the parameter sequence $\boldsymbol{y}$, even if the mathematical model is linear.

Here we consider in addition a mathematical model which is nonlinear, also w.r. to the state variable q taking values in the function space $\mathcal{X}$.

The mathematical model in the physical domain reads: given $\boldsymbol{y} \in U$, find $q(\boldsymbol{y}) : D_{u(\boldsymbol{y})} \to \mathbb{R}$ such that

$$- \triangle q(\boldsymbol{y}) + q^3(\boldsymbol{y}) = f \quad in \ D_{u(\boldsymbol{y})}, \quad q(\boldsymbol{y}) = 0 \quad on \ \partial D_{u(\boldsymbol{y})}, \tag{2.19}$$

where the random domain $D_{u(\boldsymbol{y})}$ is homeomorphic to the unit disc, and explicitly given by

$$D_{u(\boldsymbol{y})} := \{x = (r\cos(\theta), r\sin(\theta)) : 0 \leq r < u(\boldsymbol{y}; \theta), 0 \leq \theta < 2\pi\}, \ \boldsymbol{y} \in U \ . \tag{2.20}$$

Here, the random radius $u(\boldsymbol{y})$, as defined in (2.13), is given explicitly by

$$u(0) = \langle u \rangle = 1 \ and \ \psi_j = \frac{0.5}{j^\alpha} \sin(j\theta) \quad j \geq 1, \ where \ \alpha > 2 \ . \tag{2.21}$$

By T_u we denotes a transformation map from the nominal domain $D_{\langle u \rangle}$, the unit disk of $\mathbb{R}^2$ centered at the origin, to the parametric domain D_u, $T_u(r\cos(\theta), r\sin(\theta)) := (u(\boldsymbol{y})r\cos(\theta), u(\boldsymbol{y})r\sin(\theta))$. The nonlinear operator equation (2.19) in the parametric, uncertain physical domain becomes a parametric, nonlinear equation in the fixed nominal domain, which reads as: given a parameter sequence $\boldsymbol{y} \in U$, find a parametric response $q(\boldsymbol{y}) : D_{\langle u \rangle} \to \mathbb{R}$ such that

$$\begin{cases} -div(M(\boldsymbol{y})\nabla q(\boldsymbol{y})) + q^3(\boldsymbol{y})d(\boldsymbol{y}) &= f d(\boldsymbol{y}) \quad in \ D_{\langle u \rangle}, \\ q(\boldsymbol{y}) &= 0 \quad on \ \partial D_{\langle u \rangle}, \end{cases} \tag{2.22}$$

where $d(\boldsymbol{y})$ denotes the determinant of the Jacobian dT_u of the map T_u, given as $d(\boldsymbol{y}) = (u(\boldsymbol{y}))^2$;

$$M(\boldsymbol{y}) := d(\boldsymbol{y})dT_u^{-1}dT_u^{-\top} = \begin{pmatrix} 1 + (b(\boldsymbol{y}))^2 & -b(\boldsymbol{y}) \\ -b(\boldsymbol{y}) & 1 \end{pmatrix} \quad where \ b(\boldsymbol{y}) := \frac{\partial_\theta u(\boldsymbol{y})}{u(\boldsymbol{y})} \ . \tag{2.23}$$

This example fits into the abstract setting of Section 2.1 as follows: the uncertain datum $u = u(\boldsymbol{y}; \cdot) \in X_t = C_{per}^t([0, 2\pi])$ where the degree of smoothness $t = t(\alpha)$ depends on the exponent $\alpha > 2$ in (2.21). The spaces $\mathcal{X}$ and $\mathcal{Y}$ then are function spaces on the nominal domain, and chosen as $\mathcal{X} = \mathcal{Y} = H_0^1(D_{\langle u \rangle})$.

We note that uncertain inputs u with "higher regularity" (when measured in a smoothness scale $\{X_t\}_{t \geq 0}$ with $X = X_0 \supset X_1 \supset X_2 \supset ...$ on the admissible input data) correspond to stronger decay of ψ_j: for $u \in X_t \subset X$, in (2.13) the $\{\psi_j\}_{j \geq 1}$ are assumed scaled such that

$$\boldsymbol{b} := \{\|\psi_j\|_X\}_{j \geq 1} \in \ell^p(\mathbb{N}) \ for \ some \ \ 0 < p = p(t) < 1 \ , \tag{2.24}$$

where the sequence $\boldsymbol{b} = (b_j)_{j \geq 1}$ is given by $b_j := \|\psi_j\|_X$. We also introduce the subset

$$U = \{\boldsymbol{y} \in [-1, 1]^{\mathbb{N}} : u(\boldsymbol{y}) := \langle u \rangle + \sum_{j \geq 1} y_j \psi_j \in \tilde{X}\} \ . \tag{2.25}$$

See (2.21) in Example 2.4, where the exponent $\alpha > 2$ determines the (Hölder) smoothness of the domain transformation, and where the spaces X_t correspond to Hölder spaces of 2π-periodic functions.

Once an unconditional basis $\{\psi_j\}_{j \geq 1}$ of X has been chosen, every realization $u \in X$ can be

identified in a one-to-one fashion with the pair $(\langle u \rangle, \boldsymbol{y})$ where $\langle u \rangle$ denotes the *nominal instance* of the uncertain datum u and $\boldsymbol{y}$ is the coordinate vector in representation (2.13). Inserting (2.13) into (2.4), we obtain under Assumption 1 the *equivalent, countably-parametric form*: given $F : U \to \mathcal{Y}'$,

$$\text{find } q(\boldsymbol{y}; F) \in \mathcal{X} : \ \forall \boldsymbol{y} \in U : \ \mathcal{R}(q; \boldsymbol{y}) := A(q; \boldsymbol{y}) - F(\boldsymbol{y}) = 0 \ \text{ in } \ \mathcal{Y}' . \tag{2.26}$$

Remark 2.2 *In what follows, by a slight abuse of notation,* one identifies the subset U in (2.25) with the countable set of parameters from the infinite-dimensional parameter domain $U \subseteq \mathbb{R}^{\mathbb{N}}$ *without explicitly writing so. The operator $A(q; u)$ in (2.5) then becomes, via the parametric dependence $u = u(\boldsymbol{y})$, a parametric operator family $A(q; u(\boldsymbol{y}))$ which one denotes (with slight abuse of notation) by $\{A(q; \boldsymbol{y}) : \boldsymbol{y} \in U\}$, with the parameter set $U = [-1, 1]^{\mathbb{N}}$ (again, one uses in what follows this definition in place of the set U as defined in (2.25)). If $A(q; \boldsymbol{y})$ in (2.5) is linear, one has $A(q; \boldsymbol{y}) = A(\boldsymbol{y})q$ with $A(\boldsymbol{y}) \in \mathcal{L}(\mathcal{X}, \mathcal{Y}')$. One does not assume, however, that the maps $q \mapsto A(q; \boldsymbol{y})$ are linear in what follows, unless explicitly stated.*

With this understanding, and under the assumptions (2.11) and (2.12), the operator equation (2.5) will admit, for every $\boldsymbol{y} \in U$, a unique solution $q(\boldsymbol{y}; F)$ which is, due to (2.11) and (2.12), uniformly bounded and depends Lipschitz continuously on the parameter sequence $\boldsymbol{y} \in U$: there holds

$$\sup_{\boldsymbol{y} \in U} \|q(\boldsymbol{y}; F)\|_{\mathcal{X}} \leq C(F, U), \tag{2.27}$$

and, if the local Lipschitz condition (2.12) holds, there exists a Lipschitz constant $L > 0$ such that

$$\|q(\boldsymbol{y}; F) - q(\boldsymbol{y}'; F)\|_{\mathcal{X}} \leq L(F, U)\|u(\boldsymbol{y}) - u(\boldsymbol{y}')\|_X . \tag{2.28}$$

The Lipschitz constant in (2.28) is not, in general, equal to $L(F, \tilde{X})$ in (2.12): it depends on the nominal instance $\langle u \rangle \in X$ and on the choice of basis $\{\psi_j\}_{j \geq 1}$.

Unless explicitly stated otherwise, throughout what follows, we identify $q_0 = q(\boldsymbol{0}; F) \in \mathcal{X}$ in Proposition 2.1 with the solution of (2.4) at the nominal input $\langle u \rangle \in X$.

2.3 Parameter sparsity in Forward UQ

In forward UQ for problems with distributed, uncertain input data, upon uncertainty parametrizations such as (2.13), the solution of the forward problem become, as functions of the parameters y_j in the sequence $\boldsymbol{y}$, countably parametric maps from the parameter space U to the solutions' state space $\mathcal{X}$. Efficient computational UQ for such problems is crucially related to the approximation of such countably-parametric maps with *convergence rates which are independent of the dimension,* ie., independent of the number of coordinates which are active in the approximation. Mathematical results are reviewed which allow to establish such approximation results, with a key insight being that the attainable convergence rate is independent of the dimension of the space of active parameters, and depends only on the "sparsity" of the parametric map which is to be approximated. One technique to verify parametric sparsity for a broad class of parametric problems is to verify the existence of suitable *holomorphic extensions* of the parameter to solution map into the complex domain. The existence of such extensions is closely related to polynomial approximations of the parametric maps. The version presented here applies the holomorphic extension of countably-parametric maps for which polynomial approximations take the form of so-called *generalized polynomial chaos expansions* which will be described in detail below. The results presented in this section are recent; the reader is referred to [26, 25] and the references there.

2.3.1 $(\boldsymbol{b}, p)$-holomorphy

For $s > 1$, introduce the Bernstein ellipse in the complex plane

$$\mathcal{E}_s := \left\{ \frac{w + w^{-1}}{2} : 1 \leq |w| \leq s \right\} \subset \mathbb{C}, \tag{2.29}$$

which has semi axes of length $\frac{s+s^{-1}}{2} > 1$ and $\frac{s-s^{-1}}{2} > 0$ and denote

$$\mathcal{E}_{\boldsymbol{\rho}} := \bigotimes_{j \geq 1} \mathcal{E}_{\rho_j} \subset \mathbb{C}^{\mathbb{N}} , \qquad (2.30)$$

the tensorized poly-ellipse when $\boldsymbol{\rho} := (\rho_j)_{j \geq 1}$ is a sequence of semi-axis sums $\rho_j > 1$. With the convention $\mathcal{E}_1 = [-1, 1]$, one also admits $\rho_j = 1$ in (2.30), so that $U \subseteq \mathcal{E}_{\boldsymbol{\rho}}$.

Sparsity analysis of parametric maps $q : U \mapsto \mathcal{X} : \boldsymbol{y} \to q(\boldsymbol{y})$ as in [25] and the references there relies on holomorphic extensions of parametric solutions q from U to $\mathcal{E}_{\rho}$.

Definition 2.1 *For a positive sequence $\boldsymbol{b} = (b_j)_{j \geq 1} \in \ell^p(\mathbb{N})$ for some $0 < p < 1$, a parametric mapping $U \ni \boldsymbol{y} \mapsto q(\boldsymbol{y}) \in X$ satisfies the $(\boldsymbol{b}, p)$-holomorphy assumption in the Hilbert space $\mathcal{X}$ if and only if*

1. *For each $\boldsymbol{y} \in U$, there exists a unique $q(\boldsymbol{y}) \in X$ and the map $\boldsymbol{y} \mapsto q(\boldsymbol{y})$ from U to $\mathcal{X}$ is uniformly bounded, i.e.*

$$\sup_{\boldsymbol{y} \in U} \|q(\boldsymbol{y})\|_{\mathcal{X}} \leq C_0 , \qquad (2.31)$$

   for some finite constant $C_0 > 0$.

2. *For some $0 < \varepsilon < 1$ exists a constant $C_\varepsilon \geq C_0 > 0$ such that for any sequence $\boldsymbol{\rho} := (\rho_j)_{j \geq 1}$ of semiaxis sums ρ_j strictly larger than 1 that is $(\boldsymbol{b}, \varepsilon)$-admissible, ie.*

$$\sum_{j=1}^{\infty} (\rho_j - 1) b_j \leq \varepsilon, \qquad (2.32)$$

   the parametric map $\boldsymbol{y} \mapsto q(\boldsymbol{y}) \in X$ admits a complex extension $\boldsymbol{z} \mapsto q(\boldsymbol{z})$ (taking values in the complexification of the space $\mathcal{X}$) that is a holomorphic mapping with respect to each variable z_j on a set of the form $\mathcal{O}_\rho := \bigotimes_{j \geq 1} \mathcal{O}_{\rho_j}$, $\mathcal{O}_{\rho_j} \subset \mathbb{C}$ is an open set containing $\mathcal{E}_{\rho_j}$, and the modulus $\|q(\boldsymbol{z})\|_{\mathcal{X}}$ of this extension is bounded on $\mathcal{E}_\rho$ in (2.30) according to

$$\sup_{\boldsymbol{z} \in \mathcal{E}_\rho} \|q(\boldsymbol{z})\|_{\mathcal{X}} \leq C_\varepsilon . \qquad (2.33)$$

The significance of $(\boldsymbol{b}, p)$ holomorphy lies in the following facts: a) solution of well-posed, countably possibly nonlinear parametric operator equations with $(\boldsymbol{b}, p)$ holomorphic, parametric operator families are $(\boldsymbol{b}, p)$ holomorphic, b) $(\boldsymbol{b}, p)$ holomorphic parametric solution maps $\{q(\boldsymbol{y}) : \boldsymbol{y} \in U\} \subset \mathcal{X}$ allow for tensorized so-called "polynomial chaos" approximations with dimension-independent N-term approximation rates which depend only on the summability exponent p of the sequence $\boldsymbol{b}$, c) $(\boldsymbol{b}, p)$ holomorphic parametric solution maps $\{q(\boldsymbol{y}) : \boldsymbol{y} \in U\} \subset \mathcal{X}$ can also be constructively approximated by sparse, Smolyak type interpolation methods, see [23, 22], d) $(\boldsymbol{b}, p)$ holomorphy is preserved under composition with holomorphic maps, in particular for example in the context of *Bayesian inverse problems*; see [62, 61] for details. e) $(\boldsymbol{b}, p)$ holomorphic parametric solution maps $\{q(\boldsymbol{y}) : \boldsymbol{y} \in U\} \subset \mathcal{X}$ allow for *low-parametric, reduced basis surrogates*. Points b) and c) are explained next, detailing in particular computational approximation strategies for the efficient computation of sparse approximations of countably-parametric solution families.

2.3.2 Sparse Polynomial Approximation

$(\boldsymbol{b}, p)$-holomorphy ensures ℓ^p summability of gpc Legendre coefficients of $(\boldsymbol{b}, p)$ holomorphic parametric solution maps $\{q(\boldsymbol{y}) : \boldsymbol{y} \in U\} \subset \mathcal{X}$. To state the result, from [23], for any coefficient bound sequence $c := (c_\nu)_{\nu \in \mathcal{F}} \subset \mathbb{R}$, one associates its downward closed envelope $\mathbf{c} := (\mathbf{c}_\nu)_{\nu \in \mathcal{F}}$ defined by

$$\mathbf{c}_\nu := \sup_{\mu \geq \nu} |c_\nu|, \quad \nu \in \mathcal{F}, \qquad (2.34)$$

where $\mu \geq \nu$ means that $\mu_j \geq \nu_j$ for all j. An index set $\Lambda \subset \mathcal{F}$ is downward closed if and only if

$$\nu \in \Lambda \quad \text{and} \quad \mu \leq \nu \Rightarrow \mu \in \Lambda. \qquad (2.35)$$

For a summability exponent $p > 0$, one introduces the space $\ell^p_m(\mathcal{F})$ of sequences that have their downward closed envelope in $\ell^p(\mathcal{F})$. One approximates the parametric responses by truncating the tensorized Legendre ("generalized polynomial chaos") series

$$q(\boldsymbol{y}) = \sum_{\boldsymbol{\nu} \in \mathcal{F}} q_{\boldsymbol{\nu}} P_{\boldsymbol{\nu}}(\boldsymbol{y}), \tag{2.36}$$

where the convergence is understood to be unconditional (in particular, the limit exists and is independent of the particular enumeration of $\mathcal{F}$) and where the tensorized Legendre polynomials $P_{\boldsymbol{\nu}}(\boldsymbol{y})$ are given by $P_{\boldsymbol{\nu}}(\boldsymbol{y}) := \prod_{j \geq 1} P_{\nu_j}(y_j)$, with P_n denoting the univariate Legendre polynomial of degree n for the interval $[-1, 1]$ with the classical normalization $\|P_n\|_{L^\infty([-1,1])} = |P_n(1)| = 1$. The series (2.36) may be rewritten as

$$q(\boldsymbol{y}) = \sum_{\boldsymbol{\nu} \in \mathcal{F}} v_{\boldsymbol{\nu}} L_{\nu}(\boldsymbol{y}) , \tag{2.37}$$

where $L_{\boldsymbol{\nu}}(\boldsymbol{y}) := \prod_{j \geq 1} L_{\nu_j}(y_j)$, with L_n denoting the version of P_n normalized in $L^2([-1, 1], \frac{dt}{2})$, ie.

$$q_{\boldsymbol{\nu}} = \left(\prod_{j \geq 1} (1 + 2\nu_j) \right)^{1/2} v_{\boldsymbol{\nu}} . \tag{2.38}$$

Theorem 2.2 *([22]) For a $(\boldsymbol{b}, p)$-holomorphic, parametric map $U \ni \boldsymbol{y} \to q(\boldsymbol{y}) \in \mathcal{X}$ in a Hilbert space $\mathcal{X}$, the sequences $(\|q_\nu\|_\mathcal{X})_{\nu \in \mathcal{F}}$ and $(\|v_\nu\|_\mathcal{X})_{\nu \in \mathcal{F}}$ of (norms of) the tensorized Legendre coefficients belong to $\ell^p_m(\mathcal{F})$, and*

$$q(\boldsymbol{y}) = \sum_{\boldsymbol{\nu} \in \mathcal{F}} q_\nu P_{\boldsymbol{\nu}} = \sum_{\boldsymbol{\nu} \in \mathcal{F}} v_{\boldsymbol{\nu}} L_{\boldsymbol{\nu}}, \tag{2.39}$$

holds in the sense of unconditional convergence in $L^\infty(U, \mathcal{X})$.

There exists a sequence $(\Lambda_N)_{N \geq 1}$, with $\#(\Lambda_N) = N$ of nested downward closed sets such that

$$\inf_{w \in \mathcal{X}_{\Lambda_n}} \|q - w\|_{L^\infty(U,\mathcal{X})} \leq C(N+1)^{-s}, \quad s = \frac{1}{p} - 1, \tag{2.40}$$

where for any finite set $\Lambda \subset \mathcal{F}$ one defines

$$\mathcal{X}_\Lambda := \mathrm{span} \left\{ \sum_{\boldsymbol{\nu} \in \Lambda} w_{\boldsymbol{\nu}} \boldsymbol{y}^{\boldsymbol{\nu}} : w_{\boldsymbol{\nu}} \in \mathcal{X} \right\} . \tag{2.41}$$

2.3.3   Sparse Grid Interpolation

Polynomial interpolation processes on the spaces $\mathcal{X}_\Lambda$ for general downward closed sets Λ of multiindices have been introduced and studied in [23]. Given $z := (z_j)_{j \geq 1}$, a sequence of pairwise distinct points of $[-1, 1]$, one associates with any finite subset $\Lambda \subset \mathcal{F}$ the following *sparse interpolation grid in U*:

$$\Gamma_\Lambda := \{z_{\boldsymbol{\nu}} : \boldsymbol{\nu} \in \Lambda\} \quad \text{where} \quad z_{\boldsymbol{\nu}} := (z_{\nu_j})_{j \geq 1} . \tag{2.42}$$

If $\Lambda \subset \mathcal{F}$ is downward closed , then the sparse grid Γ_Λ is unisolvent for $\mathbb{P}_\Lambda$: for any function g defined in Γ_Λ and taking values in $\mathcal{X}$, there exists a unique sparse grid interpolation polynomial $I_\Lambda g$ in $\mathbb{P}_\Lambda$ that coincides with g on Γ_Λ. The interpolation polynomial $I_\Lambda g \in \mathbb{P}_\Lambda \otimes \mathcal{X}$ can be computed recursively: if $\Lambda := \{\nu^1, \cdots, \nu^N\}$ such that for any $k = 1 \cdots, N$, $\Lambda_k := \{\nu^1, \cdots, \nu^k\}$ is downward closed , then

$$I_\Lambda g = \sum_{i=1}^{N} g_{\nu^i} H_{\nu^i}, \tag{2.43}$$

where the polynomials $(H_{\boldsymbol{\nu}})_{\boldsymbol{\nu} \in \Lambda}$ are a hierarchical basis of $\mathbb{P}_\Lambda$ given by

$$H_{\boldsymbol{\nu}}(\boldsymbol{y}) := \prod_{j \geq 1} h_{\nu_j}(y_j) \quad where \quad h_0(t) = 1 \ \ and \ \ h_k(t) = \prod_{j=0}^{k-1} \frac{t - z_j}{z_k - z_j}, \ \ k \geq 1, \tag{2.44}$$

and where the coefficients g_{ν^k} are recursively defined by

$$g_{\nu^1} := g(z_0), \quad g_{\nu^{k+1}} := g(z_{\nu^{k+1}}) - I_{\Lambda_k} g(z_{\nu^{k+1}}) = g(z_{\nu^{k+1}}) - \sum_{i=1}^{k} g_{\nu^i} H_{\nu^i}(z_{\nu^{k+1}}) . \tag{2.45}$$

The sparse grid $\Gamma_\Lambda \subset U$ is unisolvent for the space $\mathcal{X}_\Lambda$ of multivariate polynomials with coefficients in $\mathcal{X}$. The interpolation operator that maps functions defined on U with values in $\mathcal{X}$ into $\mathcal{X}_\Lambda$ can be computed by the recursion (2.43) if one admits $g_\nu \in \mathcal{X}$. Naturally, in this case the coefficients g_ν being elements of a function space can not be exactly represented, and must be additionally approximated, e.g. by a Finite Element or a collocation approximation in a finite-dimensional subspace $\mathcal{X}_h \subset \mathcal{X}$.

The following result recovers the best N-term approximation rate $\mathcal{O}(N^{-s})$ in (2.40) for the interpolation in $\mathbb{P}_\Lambda$ different choice of downward closed sets Λ. See [22] for a proof.

Theorem 2.3 *For any $(\boldsymbol{b}, p)$-holomorphic, $\mathcal{X}$-valued parametric map $\boldsymbol{y} \mapsto q(\boldsymbol{y})$ there exists a constant $C > 0$ and a nested sequence of downward closed sets $(\Lambda_n)_{N \geq 1}$ with $\#(\Lambda_N) = N$ for which*

$$\|q - I_{\Lambda_N} q\|_{L^\infty(U,\mathcal{X})} \leq C(N + 1)^{-s}, \ \ s = \frac{1}{p} - 1 . \tag{2.46}$$

3 Model Order Reduction

Given any sample $\boldsymbol{y} \in U$, an accurate solution of the forward PDE model (2.26) relies on a stable and consistent numerical solver with high precision, which typically requires a high-fidelity discretization of the PDE model and a computationally expensive solving of the corresponding algebraic system. Such a large-scale computation for a large number of samples is the most critical challenge in UQ problems. This section outlines *model order reduction* (MOR for short) methods in order to effectively alleviate the computational burden while facilitating certified accuracy of the parametric solution as well as its related quantities of interest. The material in this section is related to developments during the past decade. Our presentation is therefore synoptic, and the reader is referred to the surveys [41, 47, 49] and the references there for more detailed elaboration, and further references.

3.1 High-Fidelity Approximation

At first, a stable and consistent high fidelity approximation of the solution of the parametric problem (2.26) following [20] is presented. To guarantee the stability of the HiFi approximation at any given $\boldsymbol{y} \in U$, one considers the Petrov-Galerkin (PG) discretization in the one-parameter family of pairs of *subspaces* $\mathcal{X}_h \subset \mathcal{X}$ and $\mathcal{Y}_h \subset \mathcal{Y}$ with equal dimensions, i.e. $N_h = \dim(\mathcal{X}_h) = \dim(\mathcal{Y}_h) < \infty$, where h represents a discretization parameter, for instance the meshwidth of a PG Finite Element discretization. To ensure the convergence of the HiFi PG solution $q_h \in \mathcal{X}_h$ to the exact solution $q \in \mathcal{X}$ as $h \to 0$, one assumes the subspace families $\mathcal{X}_h$ and $\mathcal{Y}_h$ to be dense in $\mathcal{X}$ and $\mathcal{Y}$ as the discretization parameter (being, for example, a meshwidth or an inverse spectral order) $h \to 0$, i.e.

$$\forall w \in \mathcal{X} : \ \lim_{h \to 0} \inf_{w_h \in \mathcal{X}_h} \|w - w_h\|_\mathcal{X} = 0, \quad and \quad \forall v \in \mathcal{Y} : \ \lim_{h \to 0} \inf_{v_h \in \mathcal{Y}_h} \|v - v_h\|_\mathcal{Y} = 0 . \tag{3.1}$$

Moreover, to quantify the convergence rate of the discrete approximation, one introduces a *scale of smoothness spaces* $\mathcal{X}^s \subset \mathcal{X} = \mathcal{X}^0$ and $\mathcal{Y}^s \subset \mathcal{Y} = \mathcal{Y}^0$ indexed by the smoothness parameter $s > 0$. Here, one has in mind for example spaces of functions with s extra derivatives in Sobolev or Besov spaces. Then, for appropriate choices of the subspaces $\mathcal{X}_h$ and $\mathcal{Y}_h$ hold the approximation properties: there exist constants $C_s > 0$ such that for all $0 < h \leq 1$ holds

$$\forall w \in \mathcal{X}^s : \ \inf_{w_h \in \mathcal{X}_h} \|w - w_h\|_\mathcal{X} \leq C_s h^s \|w\|_{\mathcal{X}^s} \quad and \quad \forall v \in \mathcal{Y}^s : \ \inf_{v_h \in \mathcal{Y}_h} \|v - v_h\|_\mathcal{Y} \leq C_s h^s \|v\|_{\mathcal{Y}^s} . \tag{3.2}$$

Here, the constant C_s is assumed independent of the discretization parameter h but may depend on the smoothness parameter s. For small values of h and/or if s is large, the PG discretization produces high-fidelity (HiFi) approximations $q_h \in \mathcal{X}_h$ of the true solution $q \in \mathcal{X}$ by solving

$$\text{given } \boldsymbol{y} \in U, \text{ find } q_h(\boldsymbol{y}) \in \mathcal{X}_h: \quad _{\mathcal{Y}'}\langle \mathcal{R}(q_h(\boldsymbol{y}); \boldsymbol{y}), v_h \rangle_{\mathcal{Y}} = 0 \quad \forall v_h \in \mathcal{Y}_h \ . \tag{3.3}$$

A globally convergent Newton iteration method can be applied to solve the nonlinear, parametric HiFi-PG approximation problem (3.3) numerically, see [30, 20] for details.

To establish the well-posedness of the HiFi-PG approximation problem (3.3) as well as the a-priori and a-posteriori error estimates for the approximate solution q_h, the following assumptions are imposed.

Assumption 2 *Let $a(\cdot, \cdot; \boldsymbol{y}) : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ denote the parametric bilinear form for each $\boldsymbol{y} \in U$ associated with the Fréchet derivative of $\mathcal{R}$ at q, i.e.*

$$a(w, v; \boldsymbol{y}) :=_{\mathcal{Y}'} \langle D_q \mathcal{R}(q(\boldsymbol{y}); \boldsymbol{y})(w), v \rangle_{\mathcal{Y}} \quad \forall w \in \mathcal{X}, \forall v \in \mathcal{Y} \ . \tag{3.4}$$

The following conditions are assumed to hold

A1 stability: *the parametric bilinear form a satisfies the discrete HiFi-PG inf-sup condition*

$$\forall \boldsymbol{y} \in U: \quad \inf_{0 \neq w_h \in \mathcal{X}_h} \sup_{0 \neq v_h \in \mathcal{Y}_h} \frac{a(w_h, v_h; \boldsymbol{y})}{||w_h||_{\mathcal{X}} ||v_h||_{\mathcal{Y}}} =: \beta_h(\boldsymbol{y}) \geq \beta_h > 0 \ , \tag{3.5}$$

where the inf-sup constant $\beta_h(\boldsymbol{y})$ depends on h and on $\boldsymbol{y}$ and may vanish $\beta_h(\boldsymbol{y}) \to 0$ as $h \to 0$.

A2 consistency: *the best approximation satisfies the consistent approximation property*

$$\forall \boldsymbol{y} \in U: \quad \lim_{h \to 0} \frac{1}{\beta_h^2(\boldsymbol{y})} \inf_{w_h \in \mathcal{X}_h} ||q(\boldsymbol{y}) - w_h||_{\mathcal{X}} = 0 \ . \tag{3.6}$$

In view of the convergence rate in (3.2), (3.6) amounts to require $h^s/\beta_h^2(\boldsymbol{y}) \to 0$ as $h \to 0$.

A3 local Lipschitz continuity: *there exists ϵ_0 and $L > 0$ such that for all $w \in \mathcal{X}$ with $||q(\boldsymbol{y}) - w||_{\mathcal{X}} \leq \epsilon_0$, there holds*

$$\forall \boldsymbol{y} \in U: \quad ||D_q \mathcal{R}(q(\boldsymbol{y}); \boldsymbol{y}) - D_q \mathcal{R}(w; \boldsymbol{y})||_{\mathcal{L}(\mathcal{X}, \mathcal{Y}')} \leq L||q(\boldsymbol{y}) - w||_{\mathcal{X}} \ . \tag{3.7}$$

Assumption 2 is sufficient to guarantee the existence of a solution $q_h(\boldsymbol{y}) \in \mathcal{X}_h$ of the HiFi-PG approximation problem (3.3) for any $\boldsymbol{y} \in U$, which is locally unique and satisfies a-priori error estimate. The results are presented in the following theorem, whose proof follows that in [56].

Theorem 3.1 *Under Assumption 2, there exists $h_0 > 0$ and $\eta_0 > 0$ such that for $0 < h \leq h_0$, there exists a solution $q_h(\boldsymbol{y}) \in \mathcal{X}_h$ of the HiFi-PG approximation problem (3.3), which is unique in $\mathcal{B}_{\mathcal{X}}(q(\boldsymbol{y}); \eta_0 \beta_h(\boldsymbol{y}))$. Moreover, for $0 < h \leq h_0$, there holds the a-priori error estimate*

$$||q(\boldsymbol{y}) - q_h(\boldsymbol{y})||_{\mathcal{X}} \leq 2\frac{||a(\boldsymbol{y})||}{\beta(\boldsymbol{y})} \left(1 + \frac{||a(\boldsymbol{y})||}{\beta_h(\boldsymbol{y})}\right) \inf_{w_h \in \mathcal{X}_h} ||q(\boldsymbol{y}) - w_h||_{\mathcal{X}} \ , \tag{3.8}$$

where $||a(\boldsymbol{y})|| := ||D_q \mathcal{R}(q(\boldsymbol{y}); \boldsymbol{y})||_{\mathcal{L}(\mathcal{X}, \mathcal{Y}')}$. Depending on the smoothness parameter $s > 0$ (see (3.2)) and the polynomial degree $r \geq 1$ of the Finite Element space, one has

$$\inf_{w_h \in \mathcal{X}_h} ||q(\boldsymbol{y}) - w_h||_{\mathcal{X}} \leq Ch^k ||q(\boldsymbol{y})||_{\mathcal{X}^s}, \quad k = \min\{s, r\} \ , \tag{3.9}$$

where C is independent of the mesh size h and uniformly bounded w.r.t. $\boldsymbol{y}$. Moreover, one has the a-posteriori error estimate

$$||q(\boldsymbol{y}) - q_h(\boldsymbol{y})||_{\mathcal{X}} \leq \frac{4}{\beta(\boldsymbol{y})} ||\mathcal{R}(q_h(\boldsymbol{y}); \boldsymbol{y})||_{\mathcal{Y}'} \ . \tag{3.10}$$

14

In many (but not all) practical applications in UQ, the stability constants $\beta(\boldsymbol{y})$ and $\beta_h(\boldsymbol{y})$ in Assumption 2 are independent of $\boldsymbol{y}$ and of h: consider specifically the parametric elliptic diffusion problem in Example 2.3. In this example, one has that for $\mathcal{X} = \mathcal{Y} = H_0^1(D)$ holds, for every $\boldsymbol{y} \in U$, that $\beta_h(\boldsymbol{y}) \geq \beta(\boldsymbol{y}) \geq c_0(1 + C_P)/2$.

3.2 Reduced Basis Compression

In order to avoid too many computationally expensive numerical solutions of the HiFi-PG problem (3.3) at a large number of required samples $\boldsymbol{y} \in U$, one computes surrogate solutions with certified accuracy and inexpensive cost by applying reduced basis (RB) compression techniques [57, 59, 3]. The rationality lies in the fact that the intrinsic dimension of the solution manifold $\mathcal{M}_h := \{q_h(\boldsymbol{y}), \boldsymbol{y} \in U\}$ could be low, even if the dimension of parameter space is high or infinite, so that the parametric solution can be compressed into a low-dimensional subspace of the HiFi space.

One assumes available a pair of N-dimensional subspaces $\mathcal{X}_N \subset \mathcal{X}_h$ and $\mathcal{Y}_N \subset \mathcal{Y}_h$ with $N \ll N_h$, which are known as RB (trail and test) spaces, whose construction are detailed in the next section. Then the RB-PG compression problem of the HiFi-PG approximation (2.26) is formulated as

$$\text{given } \boldsymbol{y} \in U, \ \ \text{find } q_N(\boldsymbol{y}) \in \mathcal{X}_N : \ \ \ {}_{\mathcal{Y}'}\langle \mathcal{R}(q_N(\boldsymbol{y}); \boldsymbol{y}), v_N \rangle_{\mathcal{Y}} = 0 \ \ \ \forall v_N \in \mathcal{Y}_N , \tag{3.11}$$

which can be solved by Newton iteration algorithm [20]. Note that the RB-PG problem (3.11) is nothing different from the HiFi-PG problem (3.3) except for the trial and test spaces, which indicate that the RB solution $q_N(\boldsymbol{y})$ is a PG compression/projection of the HiFi solution $q_h(\boldsymbol{y})$ from the HiFi space into the RB space. For the well-posedness of the RB solution, one makes the following assumptions.

Assumption 3 *Holding Assumption 2, with the same notation of the bilinear form $a(\cdot, \cdot; \boldsymbol{y}) : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ defined as in (3.4), one makes the further assumptions that*

A1 stability: *the parametric bilinear form a satisfies the discrete RB-PG inf-sup condition: there holds*

$$\forall \boldsymbol{y} \in U : \ \ \inf_{0 \neq w_N \in \mathcal{X}_N} \sup_{0 \neq v_N \in \mathcal{Y}_N} \frac{a(w_N, v_N; \boldsymbol{y})}{||w_N||_{\mathcal{X}} ||v_N||_{\mathcal{Y}}} =: \beta_N(\boldsymbol{y}) \geq \beta_N > 0 , \tag{3.12}$$

where β_N is a lower bound of the inf-sup constant $\beta_N(\boldsymbol{y})$, which depends on N and on $\boldsymbol{y}$ and may converge to the HiFi inf-sup constant $\beta_N(\boldsymbol{y}) \to \beta_h(\boldsymbol{y})$ as $N \to N_h$.

A2 consistency: *the best approximation satisfies the consistent approximation property*

$$\forall \boldsymbol{y} \in U : \ \ \lim_{N \to N_h} \frac{1}{\beta_N^2(\boldsymbol{y})} \inf_{w_N \in \mathcal{X}_N} ||q_h(\boldsymbol{y}) - w_N||_{\mathcal{X}} = 0 . \tag{3.13}$$

Proceeding as in [56], one can establish the following error estimates for the RB solution (see [20])

Theorem 3.2 *Under Assumption 3, there exist $N_0 > 0$ and $\eta_0' > 0$ such that for $N \geq N_0$, there exists a solution $q_N(\boldsymbol{y}) \in \mathcal{X}_N$ of the RB-PG compression problem (3.11), which is unique in $\mathcal{B}_{\mathcal{X}}(q_h(\boldsymbol{y}); \eta_0' \beta_N(\boldsymbol{y}))$. Moreover, for any $N \geq N_0$, there holds the a priori error estimate*

$$||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}} \leq 2 \frac{||a(\boldsymbol{y})||}{\beta_h(\boldsymbol{y})} \left(1 + \frac{||a(\boldsymbol{y})||}{\beta_N(\boldsymbol{y})}\right) \inf_{w_N \in \mathcal{X}_N} ||q_h(\boldsymbol{y}) - w_N||_{\mathcal{X}} . \tag{3.14}$$

Moreover, one has the a-posteriori error estimate

$$||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}} \leq \frac{4}{\beta_h(\boldsymbol{y})} ||\mathcal{R}(q_N(\boldsymbol{y}); \boldsymbol{y})||_{\mathcal{Y}'} . \tag{3.15}$$

Remark 3.1 *Note that both the a-priori and the a-posteriori error estimates of the RB solution turn out to be the same as those of the HiFi solution with different stability constants and different approximation spaces. These results are obtained as a consequence of the fact that the RB-PG problem (3.11) is nothing different from the HiFi-PG problem (3.3) except in different approximation spaces.*

3.3 Reduced Basis Construction

As the computational cost for the solution of the RB-PG problem (3.11) critically depends on the RB degrees of freedom (dof) N, one needs to construct the optimal RB space $\mathcal{X}_N$ that is most "representative" for all the parametric solutions with required approximation accuracy, such that N is as small as possible. However, it is computationally unfeasible to obtain such a optimal subspace $\mathcal{X}_N$ as it is an infinite dimensional optimization problem involving expensive HiFi solutions. In the following, two practical algorithms are presented that allow in practice quasi-optimal construction of RB trial spaces $\mathcal{X}_N$. Construction of the RB test space $\mathcal{Y}_N$ is deferred to the next sections.

3.3.1 Proper Orthogonal Decomposition

Proper orthogonal decomposition (POD) [9], also known as principle component analysis (PCA for short) in statistics or Karhunen–Loève (KL for short) decomposition in stochastic analysis, aims to extract the maximum information/energy/variance from a finite number of available solution "snapshots". Such solution snapshots could be, e.g., solutions at a finite set of parameter values in our context. In practice, the POD is determined from a finite training set $\Xi_t = \{\boldsymbol{y}^n \in U, n = 1, \ldots, N_t\}$ with N_t random samples, and the corresponding HiFi solutions $q_h(\boldsymbol{y})$, $\boldsymbol{y} \in \Xi_t$. The POD basis functions are defined as follows [58]: let $\mathbb{C}$ denote the correlation matrix with rank $N_r \leq N_t$, which is given by

$$\mathbb{C}_{mn} = (q_h(\boldsymbol{y}^m), q_h(\boldsymbol{y}^n))_{\mathcal{X}}, \quad m, n = 1, \ldots, N_t ; \tag{3.16}$$

let $(\lambda_n, \boldsymbol{\psi}_n)_{n=1}^{N_r}$ denote the eigenpairs of the correlation matrix $\mathbb{C}$, i.e.

$$\mathbb{C}\boldsymbol{\psi}_n = \lambda_n \boldsymbol{\psi}_n, \quad n = 1, \ldots, N_r . \tag{3.17}$$

Then the POD basis functions are given by

$$\zeta_h^n = \sum_{m=1}^{N_t} \frac{1}{\sqrt{\lambda_n}} \boldsymbol{\psi}_n^{(m)} q_h(\boldsymbol{y}^m), \quad n = 1, \ldots, N_r . \tag{3.18}$$

In common practice, instead of assembling the large correlation matrix $\mathbb{C}$ and compute its eigenpairs, one may apply singular value decomposition (SVD) method or its reduced version such as thin SVD [8, 58] in order to speed up the computation of the POD basis functions.

The POD basis functions are optimal in the "average" sense [58].

Proposition 3.3 *Let $W = \{w_h^1, \ldots, w_h^N\}$ denote any N-dimensional ($N \leq N_r$) orthonormal functions in $\mathcal{X}_h$, i.e. $(w_h^m, w_h^n)_{\mathcal{X}} = \delta_{mn}$, $m, n = 1, \ldots, N$; let P_N^W denote the $\mathcal{X}$-projection operator on W, i.e.*

$$P_N^W w_h = \sum_{n=1}^{N} (w_h, w_h^n)_{\mathcal{X}} w_h^n \quad \forall w_h \in \mathcal{X}_h . \tag{3.19}$$

Then POD basis functions $W_{pod} = \{\zeta_h^1, \ldots, \zeta_h^N\}$ given by (3.18) are orthonormal and satisfy

$$W_{pod} = \underset{W \subset \mathcal{X}_h}{\operatorname{argmin}} \sum_{n=1}^{N_t} ||q_h(\boldsymbol{y}^n) - P_N^W q_h(\boldsymbol{y}^n)||_{\mathcal{X}}^2 . \tag{3.20}$$

Moreover,

$$\sum_{n=1}^{N_t} ||q_h(\boldsymbol{y}^n) - P_N^{W_{pod}} q_h(\boldsymbol{y}^n)||_{\mathcal{X}}^2 = \sum_{n=N+1}^{N_r} \lambda_n . \tag{3.21}$$

Remark 3.2 *Proposition 3.3 implies that the POD basis functions achieve the optimal compression measured in the ensemble of square $\mathcal{X}$-norm of the orthogonal projection error. Moreover, the ensemble of the projection errors can be bounded explicitly according to (3.21), which can serve as an error indicator to choose the suitable number of POD basis functions given certain requirement of accuracy. Due to its optimality, POD has been widely used for reduced basis construction in Hilbert spaces [68, 6, 8].*

Remark 3.3 *However, to compute the POD basis functions, one needs to compute the HiFi solution at a sufficiently large number of properly chosen random samples. The possibly large training set could be prohibitive for the given computational budget, especially for high-dimensional problems that require numerous samples.*

3.3.2 Greedy Algorithm

In order to avoid solving too many HiFi-PG problems for the construction of the RB spaces with a relatively much smaller number of basis functions, one turns to a greedy algorithm [57, 54, 59, 3, 5], which only requires the same number of HiFi solutions as that of the RB basis functions. An abstract formulation of the greedy search algorithm reads: choose the first sample $\boldsymbol{y}^1$ such that

$$\boldsymbol{y}^1 := \operatorname*{argsup}_{\boldsymbol{y} \in U} ||q_h(\boldsymbol{y})||_{\mathcal{X}} , \tag{3.22}$$

at which one constructs the first RB space $\mathcal{X}_1 = \operatorname{span}\{q_h(\boldsymbol{y}^1)\}$. Then, for $N = 1, 2, \ldots$, one seeks the next sample $\boldsymbol{y}^{N+1}$ such that

$$\boldsymbol{y}^{N+1} := \operatorname*{argsup}_{\boldsymbol{y} \in U} ||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}} , \tag{3.23}$$

where $q_N(\boldsymbol{y})$ is the RB solution, and construct the new RB space $\mathcal{X}_{N+1} = \mathcal{X}_N \oplus \operatorname{span}\{q_h(\boldsymbol{y}^{N+1})\}$. However, both (3.22) and (3.23) are infinite dimensional optimization problems and necessitate many HiFi solutions for the evaluation of the RB errors. In order to tackle this challenge, the true error (3.23) is replaced ideally by a tight error bound $\triangle_N(\boldsymbol{y})$ [57, 59], i.e.

$$c\triangle_N(\boldsymbol{y}) \leq ||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}} \leq C\triangle_N(\boldsymbol{y}) \quad \forall \boldsymbol{y} \in U , \tag{3.24}$$

with constants $0 < c \leq C < \infty$ possibly depending on $\boldsymbol{y}$, and preferably $\gamma := c/C \approx 1$. Meanwhile, one can relax the first sample such that $||q_h(\boldsymbol{y}^1)||_{\mathcal{X}} \geq \gamma \sup_{\boldsymbol{y} \in U} ||q_h(\boldsymbol{y})||_{\mathcal{X}}$. It is crucial that the cost for the evaluation of the error bound $\triangle_N(\boldsymbol{y})$ should be so small that its evaluation at a large number of training samples remains feasible, i.e. the cost at each sample is effectively independent of the HiFi dof. The relaxation of the true error to an effective error bound leads to the development of the so-called *weak greedy algorithm* for which an a-priori error estimate of the error incurred by RB compression is established in the following theorem.

Theorem 3.4 *Let $d_N(\mathcal{M}_h, \mathcal{X}_h)$ denote the Kolmogorov N-width, ie., the worst-case scenario error of the $\mathcal{X}$-projection of the HiFi solution $q_h(\boldsymbol{y}) \in \mathcal{M}_h$ in the optimal among all possible N-dimensional subspaces $\mathcal{Z}_N \subset \mathcal{X}_h$. Specifically,*

$$d_N(\mathcal{M}_h, \mathcal{X}_h) := \inf_{\mathcal{Z}_N \subset \mathcal{X}_h} \sup_{\boldsymbol{y} \in U} \inf_{w_N \in \mathcal{Z}_N} ||q_h(\boldsymbol{y}) - w_N||_{\mathcal{X}} . \tag{3.25}$$

Let σ_N denote the worst-case scenario RB compression error, i.e.

$$\sigma_N := \sup_{\boldsymbol{y} \in U} ||q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})||_{\mathcal{X}} . \tag{3.26}$$

Then the following results hold for the convergence rates of the RB compression error [3, 31]:

- *If $d_N \leq C_0 N^{-\alpha}$ for some $C_0 > 0$ and $\alpha > 0$, and any $N = 1, 2, \ldots$,*
  then $\sigma_N \leq C_1 N^{-\alpha}$ for all $N = 1, 2, \ldots$, where $C_1 := 2^{5\alpha+1}\gamma^{-2}C_0$;

- *If $d_N \leq C_0 e^{-c_0 N^{\alpha}}$ for some $C_0 > 0$, $c_0 > 0$, $\alpha > 0$, and any $N = 1, 2, \ldots$,*
  then $\sigma_N \leq C_1 e^{-c_1 N^{\alpha}}$ for all $N = 1, 2, \ldots$, where $C_1 := \sqrt{2C_0}\gamma^{-1}$ and $c_1 := 2^{-1-2\alpha}c_0$.

Proof The proof of the results in the finite dimensional approximation spaces $\mathcal{X}_h$ and $\mathcal{M}_h$ follows those in [31] where $\mathcal{X}$ is a Hilbert space and the solution manifold $\mathcal{M}$ is a compact set in $\mathcal{X}$. $\qquad \square$

Remark 3.4 *The above results indicate that the RB compression by the (weak) greedy algorithm achieves optimal convergence rates in comparison with the Kolmogorov width, in the case of both algebraic rate and exponential rate. However, the Kolmogorov width is typically not available for general parametric problems. In our setting, ie. for smooth parameter dependence, it can be bounded from above by the sparse interpolation error estimate in (2.46), i.e. with algebraic convergence rate N^{-s}. Exponential convergence rates are shown in [13] for a one-dimensional parametric problem whose solution is analytic w.r.t. the parameter; and in [21] for elliptic, infinite dimensional parametric problems.*

Remark 3.5 *Construction of a N-dimensional RB space only requires N HiFi solutions by the greedy algorithm, which dramatically reduces the computational cost for the RB construction as long as evaluation of the error bound is inexpensive with operation independent of the HiFi dof N_h.*

3.4 Linear and Affine-parametric Problems

To illustrate the reduction in complexity which can be achieved by RB compression, linear and affine problems (e.g. Examples 2.1 and 2.3) with the uncertain parametrization given in Sec. 2.2 are first considered, for which one assumes the terms in (2.26) can be written more explicitly as

$$A(q; \boldsymbol{y}) = A(\boldsymbol{y})q = \sum_{j \geq 0} y_j A_j q \quad \text{and} \quad F(\boldsymbol{y}) = \sum_{j \geq 0} y_j F_j \, , \tag{3.27}$$

where one sets $y_0 = 1$ for notational simplicity and where $A_j \in \mathcal{L}(\mathcal{X}, \mathcal{Y}')$, and $F_j \in \mathcal{Y}'$, $j \geq 0$. The parametrization (3.27) is sometime also called *linear parametrization* uncertainty, whereas the term *affine parametrization* refers to separable expansions of the form (3.27) with y_j replaced by functions $\theta_j(\boldsymbol{y})$ where $\theta_0 = 1$ and where the θ_j, $j \geq 1$ depend on several or all parameters $y_j \in \boldsymbol{y}$, but are independent of the physical coordinates.

In computational practice, one truncates the affine expansion up to $J + 1$ terms with $J \in \mathbb{N}$, depending on the required accuracy of the truncation. For notational convenience, one defines $\mathbb{J} = \{0, 1, \ldots, J\}$. The ensuing development applies verbatim in the affine-parametric case, when the parameters y_j are replaced by $\theta_j(\boldsymbol{y})$ for $j \in \mathbb{J}$, with functions that are independent of the physical variable and where each $\theta_j(\boldsymbol{y})$ possibly depends on all coordinates $y_j \in \boldsymbol{y}$.

3.4.1 High Fidelity Approximation

Under the linear and affine assumptions, the parametric HiFi-PG approximation problem (3.3) becomes

$$\text{given } \boldsymbol{y} \in U, \quad \text{find } q_h(\boldsymbol{y}) \in \mathcal{X}_h : \quad \sum_{j \in \mathbb{J}} y_j \, _{\mathcal{Y}'} \langle A_j q_h(\boldsymbol{y}), v_h \rangle_{\mathcal{Y}} = \sum_{j \in \mathbb{J}} y_j \, _{\mathcal{Y}'} \langle F_j, v_h \rangle_{\mathcal{Y}} \quad \forall v_h \in \mathcal{Y}_h \, . \tag{3.28}$$

By $(w_h^n)_{n=1}^{N_h}$ and $(v_h^n)_{n=1}^{N_h}$ one denotes the basis functions of the HiFi trial and test spaces $\mathcal{X}_h$ and $\mathcal{Y}_h$. Then, the parametric solution $q_h(\boldsymbol{y})$ can be written as

$$q_h(\boldsymbol{y}) = \sum_{n=1}^{N_h} q_h^n(\boldsymbol{y}) w_h^n \, , \tag{3.29}$$

where $\boldsymbol{q}_h(\boldsymbol{y}) = (q_h^1(\boldsymbol{y}), \ldots, q_h^{N_h}(\boldsymbol{y}))^{\top}$ denotes the (parametric) coefficient vector of the HiFi PG solution $q_h(\boldsymbol{y})$. The algebraic formulation of (3.28) reads:

$$\text{given } \boldsymbol{y} \in U, \quad \text{find } \boldsymbol{q}_h(\boldsymbol{y}) \in \mathbb{R}^{N_h} : \quad \sum_{j \in \mathbb{J}} y_j \mathbb{A}_j^h \boldsymbol{q}_h(\boldsymbol{y}) = \sum_{j \in \mathbb{J}} y_j \boldsymbol{f}_j^h \, . \tag{3.30}$$

The HiFi matrix $\mathbb{A}_j^h \in \mathbb{R}^{N_h \times N_h}$ and the HiFi vector $\boldsymbol{f}_j^h \in \mathbb{R}^{N_h}$ can be assembled as

$$(\mathbb{A}_j^h)_{mn} = {}_{\mathcal{Y}'} \langle A_j w_h^n, v_h^m \rangle_{\mathcal{Y}} \text{ and } (\boldsymbol{f}_j^h)_m = {}_{\mathcal{Y}'} \langle F_j, v_h^m \rangle_{\mathcal{Y}} \quad m, n = 1, \ldots, N_h, j \in \mathbb{J} \, . \tag{3.31}$$

3.4.2 Reduced Basis Compression

Analogously, by $(w_N^n)_{n=1}^N$ and $(v_N^n)_{n=1}^N$ one denotes the basis functions of the RB trial and test spaces $\mathcal{X}_N$ and $\mathcal{Y}_N$, so that the RB solution $q_N(\boldsymbol{y})$ can be written as

$$q_N(\boldsymbol{y}) = \sum_{n=1}^N q_N^n(\boldsymbol{y}) w_N^n , \tag{3.32}$$

with the coefficient vector $\boldsymbol{q}_N(\boldsymbol{y}) = (q_N^1(\boldsymbol{y}), \dots, q_N^N(\boldsymbol{y}))^\top$. Then, the parametric RB-PG compression problem can be written in the algebraic formulation as

$$\text{given } \boldsymbol{y} \in U, \text{ find } \boldsymbol{q}_N(\boldsymbol{y}) \in \mathbb{R}^N : \quad \sum_{j \in \mathbb{J}} y_j \mathbb{A}_j^N \boldsymbol{q}_N(\boldsymbol{y}) = \sum_{j \in \mathbb{J}} y_j \boldsymbol{f}_j^N . \tag{3.33}$$

where the RB matrix $\mathbb{A}_j^N \in \mathbb{R}^{N \times N}$ and the RB vector $\boldsymbol{f}_j^N \in \mathbb{R}^N$ are obtained as

$$\mathbb{A}_j^N = \mathbb{V}^\top \mathbb{A}_j^h \mathbb{W} \text{ and } \boldsymbol{f}_j^N = \mathbb{V}^\top \boldsymbol{f}_j^h, \quad j \in \mathbb{J} , \tag{3.34}$$

where $\mathbb{W}$ and $\mathbb{V}$ are the transformation matrices between the HiFi and RB basis functions, i.e.

$$w_N^n = \sum_{m=1}^{N_h} \mathbb{W}_{mn} w_h^m \text{ and } v_N^n = \sum_{m=1}^{N_h} \mathbb{V}_{mn} v_h^m, \quad n = 1, \dots, N . \tag{3.35}$$

Thanks to the linear and affine structure of the parametric terms in (3.27), one can assemble the RB matrices $\mathbb{A}_j^N$ and the RB vectors $\boldsymbol{f}_j^N$, $j \in \mathbb{J}$, once and for all. For each given $\boldsymbol{y}$, one only needs to assemble and solve the RB algebraic system (3.33) with computational cost depends only on N as $O(N^2)$ for assembling and $O(N^3)$ for solving (3.33), which leads to considerable computational reduction as long as $N \ll N_h$.

3.4.3 Tight A-posteriori Error Bound

A tight and inexpensive error bound that facilitates the weak greedy algorithm for RB construction is designed based on Assumption 2, in particular **A1 stability**, where the bilinear form is defined as

$$\text{given any } \boldsymbol{y} \in U : \quad a(w, v; \boldsymbol{y}) :=_{\mathcal{Y}'} \langle A(\boldsymbol{y}) w, v \rangle_{\mathcal{Y}}, \quad \forall w \in \mathcal{X}, \forall v \in \mathcal{Y} , \tag{3.36}$$

which satisfies the stability condition in the HiFi spaces $\mathcal{X}_h$ and $\mathcal{Y}_h$ as in (3.5). Let the linear form be defined as

$$f(v; \boldsymbol{y}) =_{\mathcal{Y}'} \langle F(\boldsymbol{y}), v \rangle_{\mathcal{Y}}, \quad \forall v \in \mathcal{Y} , \tag{3.37}$$

then the RB residual in the HiFi space is defined as

$$r(v_h; \boldsymbol{y}) = f(v_h; \boldsymbol{y}) - a(q_N(\boldsymbol{y}), v_h; \boldsymbol{y}), \quad \forall v_h \in \mathcal{Y}_h . \tag{3.38}$$

Let $e_N(\boldsymbol{y}) = q_h(\boldsymbol{y}) - q_N(\boldsymbol{y})$ denote the RB error, then by the stability condition (3.5) one has

$$||e_N(\boldsymbol{y})||_{\mathcal{X}} \leq \frac{|a(e_N(\boldsymbol{y}), v_h; \boldsymbol{y})|}{\beta_h ||v_h||_{\mathcal{Y}}} = \frac{|r(v_h; \boldsymbol{y})|}{\beta_h ||v_h||_{\mathcal{Y}}} \leq \frac{||r(\cdot; \boldsymbol{y})||_{\mathcal{Y}'}}{\beta_h} =: \triangle_N(\boldsymbol{y}) , \tag{3.39}$$

which indicates that $\triangle_N(\boldsymbol{y})$ is a rigorous upper error bound for the RB error $e_N(\boldsymbol{y})$. On the other hand, to see the lower bound one defines the Riesz representation of the residual as $\hat{e}_N(\boldsymbol{y}) \in \mathcal{Y}_h$, i.e.

$$(\hat{e}_N(\boldsymbol{y}), v_h)_{\mathcal{Y}} = r(v_h; \boldsymbol{y}), \quad \forall v_h \in \mathcal{Y}_h , \tag{3.40}$$

so that $||\hat{e}_N(\boldsymbol{y})||_{\mathcal{Y}} = ||r(\cdot; \boldsymbol{y})||_{\mathcal{Y}'}$. By setting $v_h = \hat{e}_N(\boldsymbol{y})$ in (3.40), one has

$$||\hat{e}_N(\boldsymbol{y})||_{\mathcal{Y}}^2 = r(\hat{e}_N(\boldsymbol{y}); \boldsymbol{y}) = a(e_N(\boldsymbol{y}), \hat{e}_N(\boldsymbol{y}); \boldsymbol{y}) \leq \alpha_h ||e_N(\boldsymbol{y})||_{\mathcal{X}} ||\hat{e}_N(\boldsymbol{y})||_{\mathcal{Y}} , \tag{3.41}$$

where α_h is the continuity constant of the bilinear form a in $\mathcal{X}_h \times \mathcal{Y}_h$, which implies that

$$\frac{\beta_h}{\alpha_h} \triangle_N(\boldsymbol{y}) \leq ||e_N(\boldsymbol{y})||_{\mathcal{X}} . \tag{3.42}$$

Therefore, the error bound $\triangle_N(\boldsymbol{y})$ is tight with the constants in (3.23) as $c = \beta_h/\alpha_h$ and $C = 1$, and $\gamma = c/C = \beta_h/\alpha_h$.

For the evaluation of $\triangle_N(\boldsymbol{y})$, one makes use of the affine structure (3.27) by computing the Riesz representation $\mathcal{A}_j^n$ of the linear functional $a_j(w_N^n; \cdot) =_{\mathcal{Y}'} \langle A_j w_N^n, \cdot \rangle_{\mathcal{Y}} : \mathcal{Y}_h \to \mathbb{R}$ as the solution of

$$(\mathcal{A}_j^n, v_h)_{\mathcal{Y}} = a_j(w_N^n; v_h), \quad \forall v_h \in \mathcal{Y}_h, \quad j \in \mathbb{J}, n = 1, \dots, N , \tag{3.43}$$

where w_N^n is the nth RB basis function. Analogously, one computes the Riesz representation $\mathcal{F}_j$ of the linear functional $f_j(\cdot) =_{\mathcal{Y}'} \langle F_j, \cdot \rangle_{\mathcal{Y}} : \mathcal{Y}_h \to \mathbb{R}$ as the solution of

$$(\mathcal{F}_j, v_h) = f_j(v_h), \quad \forall v_h \in \mathcal{Y}_h, \quad j \in \mathbb{J} . \tag{3.44}$$

Finally, one can compute the dual norm of the residual in the error bound $\triangle_N(\boldsymbol{y})$ by

$$||\hat{e}_N(\boldsymbol{y})||_{\mathcal{Y}}^2 = \sum_{j,j' \in \mathbb{J}} y_j y_{j'} \left((\mathcal{F}_j, \mathcal{F}_{j'})_{\mathcal{Y}} - 2 \sum_{n=1}^{N} q_N^n(\boldsymbol{y})(\mathcal{F}_j, \mathcal{A}_{j'}^n)_{\mathcal{Y}} + \sum_{n,n'=1}^{N} q_N^n(\boldsymbol{y}) q_N^{n'}(\boldsymbol{y})(A_j^n, A_{j'}^{n'})_{\mathcal{Y}} \right) , \tag{3.45}$$

where $(\mathcal{F}_j, \mathcal{F}_{j'})_{\mathcal{Y}}$, $(\mathcal{F}_j, \mathcal{A}_{j'}^n)_{\mathcal{Y}}$, and $(A_j^n, A_{j'}^{n'})_{\mathcal{Y}}$, $j, j' \in \mathbb{J}$, $n, n' = 1, \dots, N$, can be computed once and for all. Given any $\boldsymbol{y}$, one only need to assemble (3.45) whose cost depends on N as $O(N^2)$, not on N_h, which results in effective computational reduction as long as $N \ll N_h$.

The lower bound of the stability constant β_h in $\triangle_N(\boldsymbol{y})$ can be computed for once based on the specific structure of the parametrization (e.g. at extreme points $\boldsymbol{y} = \{y_j = \pm 1 : j = 1, 2, ...\}$, or by a successive constraint method (SCM for short) [46, 45] for each $\boldsymbol{y}$, whose computational cost is independent of N_h.

3.4.4 Stable RB-PG Compression

Construction of the RB trial space $\mathcal{X}_N$ by both POD and greedy algorithm ensures the consistency of the RB-PG compression. For its stability, a suitable RB test space $\mathcal{Y}_N$ needs to be constructed depending on $\mathcal{X}_N$. In the case that $\mathcal{X} = \mathcal{Y}$, $\mathcal{X}_h = \mathcal{Y}_h$, and the linear problem with (3.27) is coercive, the choice $\mathcal{Y}_N := \mathcal{X}_N$ guarantees the coercivity (or stability) of the RB Galerkin compression.

In the case of saddle point variational formulations of the forward problem, such as time-harmonic acoustic or electromagnetic wave propagation, or when $\mathcal{X}_h \neq \mathcal{Y}_h$, MOR requires in addition to a reduction of the trial spaces also the numerical computation of a suitable inf-sup stable testfunction space. To this end, the so-called "supremizer" approach was proposed in [60], which is described as: denote by $T_{\boldsymbol{y}} : \mathcal{X}_h \to \mathcal{Y}_h$ a parameter dependent supremizer operator, which is defined by

$$(T_{\boldsymbol{y}} w_h, v_h)_{\mathcal{Y}} = a(w_h, v_h; \boldsymbol{y}) \quad \forall v_h \in \mathcal{Y}_h . \tag{3.46}$$

This definition implies $\sup_{v_h \in \mathcal{Y}_h} |a(w_h, v_h; \boldsymbol{y})| = |a(w_h, T_{\boldsymbol{y}} w_h; \boldsymbol{y})|$, i.e. $T_{\boldsymbol{y}} w_h$ is the supremizer of w_h in $\mathcal{Y}_h$ w.r.t. the bilinear form a. Then the $\boldsymbol{y}$-dependent RB test space $\mathcal{Y}_N^{\boldsymbol{y}}$ is constructed as

$$\mathcal{Y}_N^{\boldsymbol{y}} = \text{span}\{T_{\boldsymbol{y}} w_N, w_N \in \mathcal{X}_N\} . \tag{3.47}$$

For this construction, it holds that (see [17])

$$\beta_N(\boldsymbol{y}) := \inf_{w_N \in \mathcal{X}_N} \sup_{v_N \in \mathcal{Y}_N^{\boldsymbol{y}}} \frac{a(w_N, v_N; \boldsymbol{y})}{||w_N||_{\mathcal{X}} ||v_N||_{\mathcal{Y}}} \geq \beta_h(\boldsymbol{y}) . \tag{3.48}$$

This implies that *infsup-stability of the HiFi-PG discretization is inherited by the corresponding PG-RB trial- and test-spaces:* all RB-PG compression problems are infsup-stable under the stability assumption of the HiFi-PG approximation: **A1 stability** in Assumption 2. In particular, if the

HiFi-PG discretizations are infsup-stable uniformly with respect to the uncertain input parameter u (resp. its parametrization in terms of $\boldsymbol{y}$), so is any PG-RB method obtained with PG trial space $\mathcal{X}_N$ obtained by a greedy search, and the corresponding PG test space (3.47).

Due to the affine structure (3.27), one can compute $T_{\boldsymbol{y}} w_N$ for each $\boldsymbol{y} \in U$ as

$$T_{\boldsymbol{y}} w_N = \sum_{j \in \mathbb{J}} y_j T_j w_N, \quad \text{where } (T_j w_N, v_h)_{\mathcal{Y}} = a_j(w_N, v_h) \quad \forall v_h \in \mathcal{Y}_h \ , \tag{3.49}$$

where $T_j w_N$, $w_N \in \mathcal{X}_N$, $j \in \mathbb{J}$ needs to be computed only once; given any $\boldsymbol{y}$, $T_{\boldsymbol{y}} w_N$ can be assembled in $O(N)$ operations, which is independent of the N_h HiFi dof. The RB-PG compression problem (3.33) can be written more explicitly as

$$\text{given } \boldsymbol{y} \in U, \text{ find } \boldsymbol{q}_N(\boldsymbol{y}) \in \mathbb{R}^N : \quad \sum_{j,j' \in \mathbb{J}} y_j y_{j'} \mathbb{A}_{j,j'}^N \boldsymbol{q}_N(\boldsymbol{y}) = \sum_{j,j' \in \mathbb{J}} y_j y_{j'} \boldsymbol{f}_{j,j'}^N \ , \tag{3.50}$$

where the RB matrix $\mathbb{A}_{j,j'}^N \in \mathbb{R}^{N \times N}$ and the RB vector $\boldsymbol{f}_{j,j'}^N \in \mathbb{R}^N$ are given by

$$\mathbb{A}_{j,j'}^N = \mathbb{W}^\top (\mathbb{A}_{j'}^h)^\top \mathbb{M}_h^{-1} \mathbb{A}_j^h \mathbb{W} \text{ and } \boldsymbol{f}_{j,j'}^N = \mathbb{W}^\top (\mathbb{A}_{j'}^h)^\top \mathbb{M}_h^{-1} \boldsymbol{f}_j^h, \quad j,j' \in \mathbb{J} \ , \tag{3.51}$$

where $\mathbb{M}_h$ is the mass matrix with $(\mathbb{M}_h^{-1})_{n,n'} = (v_h^{n'}, v_h^n)_{\mathcal{Y}}$, $n, n' = 1, \ldots, N_h$. Since all these quantities are independent of $\boldsymbol{y}$, one can evaluate them once and for all.

Remark 3.6 *The stable RB-PG compression is equivalent to a least-squares RB-PG compression presented in [8], see also [58]. Alternatively, a minimum residual approach known as double greedy algorithm [28] can be applied for the construction of $\mathcal{Y}_N$ to ensure infsup stability.*

3.5 Nonlinear and Nonaffine-parametric Problems

The linearity and the affinity in the linear and affine problems play a crucial role in effective decomposition of the parameter dependent and independent quantities, which leads to the computational reduction of the RB-PG compression. For more general problems that involve nonlinear terms w.r.t. the state variable q and/or nonaffine terms w.r.t. the parameter $\boldsymbol{y}$, for instance Example 2.4, it is necessary to obtain an affine approximation of these terms in order to retain the effective decomposition and RB reduction. In this section, such an affine approximation based on empirical interpolation [1, 53, 10, 37, 15] is presented.

3.5.1 High Fidelity Approximation

To solve the nonlinear parametric HiFi-PG approximation problem (3.3), one applies a Newton iteration method based on the parametric tangent operator of the nonlinear residual [20]: given any $\boldsymbol{y} \in U$ and an initial guess of the solution $q_h^{(1)}(\boldsymbol{y}) \in \mathcal{X}_h$, for $k = 1, 2, \ldots$, one finds $\delta q_h^{(k)}(\boldsymbol{y}) \in \mathcal{X}_h$ such that

$$_{\mathcal{Y}'}\langle D_q \mathcal{R}(q_h^{(k)}(\boldsymbol{y}); \boldsymbol{y})(\delta q_h^{(k)}(\boldsymbol{y})), v_h \rangle_{\mathcal{Y}} = -_{\mathcal{Y}'}\langle \mathcal{R}(q_h^{(k)}(\boldsymbol{y}); \boldsymbol{y}), v_h \rangle_{\mathcal{Y}} \quad \forall v_h \in \mathcal{Y}_h \ ; \tag{3.52}$$

then the solution is updated according to

$$q_h^{(k+1)}(\boldsymbol{y}) = q_h^{(k)}(\boldsymbol{y}) + \eta^{(k)} \delta q_h^{(k)}(\boldsymbol{y}) \ , \tag{3.53}$$

where $\eta^{(k)}$ is a constant determined by a line search method [30]. The Newton iteration is stopped once

$$||\delta q_h^{(k)}(\boldsymbol{y})||_{\mathcal{X}} \leq \varepsilon_{tol} \quad \text{or} \quad ||\mathcal{R}(q_h^{(k)}(\boldsymbol{y}); \boldsymbol{y})||_{\mathcal{Y}'} \leq \varepsilon_{tol} \ , \tag{3.54}$$

being ε_{tol} a tolerance; then one sets $q_h(\boldsymbol{y}) = q_h^{(k+1)}(\boldsymbol{y})$, bearing in mind of the tolerance error.

With the notation of the basis $\{w_h^n\}_{n=1}^{N_h}$ and $\{v_h^n\}_{n=1}^{N_h}$ of $\mathcal{X}_h$ and $\mathcal{Y}_h$, respectively, one can write

$$q_h^{(k)}(\boldsymbol{y}) = \sum_{n=1}^{N_h} q_{h,n}^{(k)}(\boldsymbol{y}) w_h^n \text{ and } \delta q_h^{(k)}(\boldsymbol{y}) = \sum_{n=1}^{N_h} \delta q_{h,n}^{(k)}(\boldsymbol{y}) w_h^n \ , \tag{3.55}$$

with the coefficient vectors $\boldsymbol{q}_h^{(k)}(\boldsymbol{y}) = (q_{h,1}^{(k)}(\boldsymbol{y}), \ldots, q_{h,N_h}^{(k)}(\boldsymbol{y}))^\top$ and $\boldsymbol{\delta q}_h^{(k)}(\boldsymbol{y}) = (\delta q_{h,1}^{(k)}(\boldsymbol{y}), \ldots, \delta q_{h,N_h}^{(k)}(\boldsymbol{y}))^\top$ so that the algebraic formulation of the parametric HiFi-PG approximation problem (3.52) reads: find the coefficient vector $\boldsymbol{\delta q}_h^{(k)}(\boldsymbol{y}) := (\delta q_{h,1}^{(k)}(\boldsymbol{y}), \ldots, \delta q_{h,N_h}^{(k)}(\boldsymbol{y}))^\top \in \mathbb{R}^{N_h}$ such that

$$\mathbb{J}_h(\boldsymbol{q}_h^{(k)}(\boldsymbol{y}); \boldsymbol{y})\boldsymbol{\delta q}_h^{(k)}(\boldsymbol{y}) = -\boldsymbol{r}_h(\boldsymbol{q}_h^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \,, \tag{3.56}$$

where the Jacobian matrix $\mathbb{J}_h(\boldsymbol{q}_h^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \in \mathbb{R}^{N_h \times N_h}$ is given by

$$\left(\mathbb{J}_h(\boldsymbol{q}_h^{(k)}(\boldsymbol{y}); \boldsymbol{y})\right)_{nn'} =_{\mathcal{Y}'} \left\langle D_q \mathcal{R}(q_h^{(k)}(\boldsymbol{y}); \boldsymbol{y})(w_h^{n'}), v_h^n \right\rangle_{\mathcal{Y}}, \quad n, n' = 1, \ldots, N_h \,, \tag{3.57}$$

and the residual vector $\boldsymbol{r}_h(\boldsymbol{q}_h^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \in \mathbb{R}^{N_h}$ takes the form

$$\left(\boldsymbol{r}_h(\boldsymbol{q}_h^{(k)}(\boldsymbol{y}); \boldsymbol{y})\right)_n =_{\mathcal{Y}'} \left\langle \mathcal{R}\left(q_h^{(k)}(\boldsymbol{y}); \boldsymbol{y}\right), v_h^n \right\rangle_{\mathcal{Y}}, \quad n = 1, \ldots, N_h \,. \tag{3.58}$$

3.5.2 Reduced Basis Compression

To solve the nonlinear parametric RB-PG compression problem, one applies the same Newton iteration method as for solving the HiFi-PG approximation problem. More specifically, one sets an initial guess of the solution $q_N^{(1)}(\boldsymbol{y}) \in \mathcal{X}_N$ for any given $\boldsymbol{y} \in U$, for $k = 1, 2, \ldots$, one finds $\delta q_N^{(k)} \in \mathcal{X}_N$ such that

$$_{\mathcal{Y}'}\langle D_q \mathcal{R}(q_N^{(k)}(\boldsymbol{y}); \boldsymbol{y})(\delta q_N^{(k)}(\boldsymbol{y})), v_N \rangle_{\mathcal{Y}} = -_{\mathcal{Y}'}\langle \mathcal{R}(q_N^{(k)}(\boldsymbol{y}); \boldsymbol{y}), v_N \rangle_{\mathcal{Y}} \quad \forall v_N \in \mathcal{Y}_N \,; \tag{3.59}$$

then the RB solution is updated by

$$q_N^{(k+1)}(\boldsymbol{y}) = q_N^{(k)}(\boldsymbol{y}) + \eta^{(k)} \delta q_N^{(k)}(\boldsymbol{y}) \,, \tag{3.60}$$

where again $\eta^{(k)}$ is a constant determined by a line search method [30]. The stopping criterion is

$$||\delta q_N^{(k)}(\boldsymbol{y})||_{\mathcal{X}} \leq \varepsilon_{tol} \quad \text{or} \quad ||\mathcal{R}(q_N^{(k)}(\boldsymbol{y}); \boldsymbol{y})||_{\mathcal{Y}'} \leq \varepsilon_{tol} \,, \tag{3.61}$$

With the notation of the basis $(w_N^n)_{n=1}^N$ and $(v_N^n)_{n=1}^N$ for the RB trial and test spaces $\mathcal{X}_N$ and $\mathcal{Y}_N$, one can expand the RB solution $q_N^{(k)}(\boldsymbol{y})$ and its update $\delta q_N^{(k)}$ as

$$q_N^{(k)}(\boldsymbol{y}) = \sum_{n=1}^N q_N^{(n,k)}(\boldsymbol{y}) w_N^n \text{ and } \delta q_N^{(k)}(\boldsymbol{y}) = \sum_{n=1}^N \delta q_N^{(n,k)}(\boldsymbol{y}) w_N^n \tag{3.62}$$

with the coefficient vectors $\boldsymbol{q}_N^{(k)} = (q_N^{(1,k)}(\boldsymbol{y}), \ldots, q_N^{(N,k)}(\boldsymbol{y}))^\top$ and $\boldsymbol{\delta q}_N^{(k)} = (\delta q_N^{(1,k)}(\boldsymbol{y}), \ldots, \delta q_N^{(N,k)}(\boldsymbol{y}))^\top$. Then the algebraic formulation of the RB-PG compression (3.59) reads: find $\boldsymbol{\delta q}_N^{(k)} \in \mathbb{R}^N$ such that

$$\mathbb{J}_N(\boldsymbol{q}_N^{(k)}(\boldsymbol{y}); \boldsymbol{y})\boldsymbol{\delta q}_N^{(k)}(\boldsymbol{y}) = -\boldsymbol{r}_N(\boldsymbol{q}_N^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \,, \tag{3.63}$$

where the parametric RB Jacobian matrix $\mathbb{J}_N(\boldsymbol{q}_N^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \in \mathbb{R}^{N \times N}$ and the parametric RB residual vector $\boldsymbol{r}_N(\boldsymbol{q}_N^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \in \mathbb{R}^N$ are given (through the transformation matrix $\mathbb{W}$ and $\mathbb{V}$) by

$$\mathbb{J}_N(\boldsymbol{q}_N^{(k)}(\boldsymbol{y}); \boldsymbol{y}) = \mathbb{V}^\top \mathbb{J}_h(\mathbb{W}\boldsymbol{q}_N^{(k)}(\boldsymbol{y}); \boldsymbol{y})\mathbb{W} \text{ and } \boldsymbol{r}_N(\boldsymbol{q}_N^{(k)}(\boldsymbol{y}); \boldsymbol{y}) = \mathbb{V}^\top \boldsymbol{r}_h(\mathbb{W}\boldsymbol{q}_N^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \tag{3.64}$$

One can observe that, due to the nonlinearity (and/or nonaffinity) of the residual operator, neither the residual vector nor the Jacobian matrix allows affine decomposition of parameter dependent and independent terms which prevents effective RB-PG compression with computational reduction.

3.5.3 Empirical Interpolation

The Empirical Interpolation Method (EIM) was originally developed for affine decomposition of non-affine parametric functions [1]. It was later applied to decompose nonaffine parametric discrete func-

tion [10] (known as discrete EIM) and nonaffine-parametric operator [37, 51]. In this presentation, it is applied to decompose the residual vector $\boldsymbol{r}_h^{(k)}$ and its tangent derivative in (3.58) and (3.57).

For notational simplicity, suppose one has M_t training (residual) vectors

$$\boldsymbol{r}_m^t \in \mathbb{R}^{N_h}, \quad m = 1, \ldots, M_t, \tag{3.65}$$

for instance collected from the residual vectors $\boldsymbol{r}_h(\boldsymbol{q}_h^{(k)}; \boldsymbol{y})$ at different iteration steps k and different samples $\boldsymbol{y}$. A *greedy algorithm* is applied to construct the empirical interpolation for the approximation of any given $\boldsymbol{r} \in \mathbb{R}^{N_h}$: one picks the first EI basis $\boldsymbol{r}_1 \in \mathbb{R}^{N_h}$ as

$$\boldsymbol{r}_1 = \boldsymbol{r}_{m^*}^t, \quad \text{where } m^* = \operatorname*{argmax}_{1 \leq m \leq M_t} ||\boldsymbol{r}_m^t||_\infty , \tag{3.66}$$

where $|| \cdot ||_\infty$ can also be replaced by $|| \cdot ||_2$; then one chooses the first index $n_1 \in \{1, \ldots, N_h\}$, such that

$$n_1 = \operatorname*{argmax}_{1 \leq n \leq N_h} |(\boldsymbol{r}_1)_n| , \tag{3.67}$$

where $(\boldsymbol{r}_1)_n$ is the n-th entry of the vector $\boldsymbol{r}_1$. The number of EI basis is set as M, and $M = 1$ for the time being. For any $\boldsymbol{r} \in \mathbb{R}^{N_h}$, it is approximated by the (empirical) interpolation

$$\mathcal{I}_M \boldsymbol{r} = \sum_{m=1}^{M} c_m \boldsymbol{r}_m , \tag{3.68}$$

where the coefficient vector $\boldsymbol{c} = (c_1, \ldots, c_M)^\top$ is obtained by solving the interpolation problem

$$(\boldsymbol{r})_{m'} = \sum_{m=1}^{M} c_m (\boldsymbol{r}_m)_{m'}, \quad m' = n_1, \ldots, n_M . \tag{3.69}$$

More explicitly, let $\mathbb{P}_M \in \{0, 1\}^{M \times N_h}$ denote an index indicator matrix with nonzero entries $(\mathbb{P}_M)_{m,n_m} = 1$, $m = 1, \ldots, M$; let $\mathbb{R}_M \in \mathbb{R}^{N_h \times M}$ denote the EI basis matrix whose m-th column is $\boldsymbol{r}_m$, $m = 1, \ldots, M$. Then, the coefficient vector $\boldsymbol{c}$ can be written as

$$\boldsymbol{c} = (\mathbb{P}_M \mathbb{R}_M)^{-1} (\mathbb{P}_M \boldsymbol{r}) , \tag{3.70}$$

and the empirical interpolation becomes

$$\mathcal{I}_M \boldsymbol{r} = \mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1} \mathbb{P}_M \boldsymbol{r} . \tag{3.71}$$

For $M = 1, 2, \ldots$, the next EI basis $\boldsymbol{r}_{M+1} \in \mathbb{R}^{N_h}$ is constructed as

$$\boldsymbol{r}_{M+1} = \frac{\boldsymbol{r}_{m^*}^t - \mathcal{I}_M \boldsymbol{r}_{m^*}^t}{||\boldsymbol{r}_{m^*}^t - \mathcal{I}_M \boldsymbol{r}_{m^*}^t||_\infty}, \quad \text{where } m^* = \operatorname*{argmax}_{1 \leq m \leq M_t} ||\boldsymbol{r}_m^t - \mathcal{I}_M \boldsymbol{r}_m^t||_\infty , \tag{3.72}$$

and find the next index n_{M+1} as

$$n_{M+1} = \operatorname*{argmax}_{1 \leq n \leq N_h} |(\boldsymbol{r}_{M+1})_n| . \tag{3.73}$$

The greedy algorithm is terminated when $|(\boldsymbol{r}_{M+1})_{n_{M+1}}| \leq \varepsilon_{tol}$. The empirical interpolation is consistent in that when $M \to N_h$, one has $\boldsymbol{r}_M \to \boldsymbol{0}$ due to the interpolation property. Moreover, an priori error analysis shows that the greedy algorithm for EI construction leads to the same result for the convergence of the EI compression error in comparison with the Kolmogorov width as the bound stated in Theorem 3.4, except for a Lebesgue constant depending on M, see [52, 20] for more details.

For any $\boldsymbol{y} \in U$, let $\boldsymbol{q}_{N,M}(\boldsymbol{y}) \in \mathbb{R}^N$ denote the coefficient of the solution $q_{N,M}(\boldsymbol{y}) \in \mathcal{X}_N$ of the RB-PG compression problem with the empirical interpolation. By the empirical interpolation of the

HiFi residual vector in (3.64), one can approximate the RB residual vector as

$$\boldsymbol{r}_N(\boldsymbol{q}_N^{(k)}; \boldsymbol{y}) \approx \boldsymbol{r}_{N,M}(\boldsymbol{q}_{N,M}^{(k)}; \boldsymbol{y}) := \mathbb{V}^\top \mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1} \mathbb{P}_M \boldsymbol{r}_h(\mathbb{W}\boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y}) , \qquad (3.74)$$

where the $\boldsymbol{y}$-independent quantity $\mathbb{V}^\top \mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1} \in \mathbb{R}^{M \times M}$ can be computed once and for all, and the $\boldsymbol{y}$-dependent quantity $\mathbb{P}_M \boldsymbol{r}_h(\mathbb{W}\boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y})$ can be evaluated in $O(MN)$ operations as long as locally supported HiFi basis functions are used, e.g. Finite Element basis functions.

Similarly, one can approximate the HiFi Jacobian matrix in (3.64) as

$$\mathbb{J}_h(\mathbb{W}\boldsymbol{q}_N^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \approx \mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1} \mathbb{P}_M \mathbb{J}_h(\mathbb{W}\boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y}) , \qquad (3.75)$$

so that the RB Jacobian matrix in (3.64) can be approximated by

$$\mathbb{J}_N(\boldsymbol{q}_N^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \approx \mathbb{J}_{N,M}(\boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y}) := \mathbb{V}^\top \mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1} \mathbb{P}_M \mathbb{J}_h(\mathbb{W}\boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y})\mathbb{W} , \qquad (3.76)$$

where the $\boldsymbol{y}$-dependent quantity $\mathbb{P}_M \mathbb{J}_h(\mathbb{W}\boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y})\mathbb{W}$ can be computed efficiently with operations $O(M^2 N)$, as long as the Jacobian matrix $\mathbb{J}_h(\mathbb{W}\boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y})$ is sparse, which is the typical case for PDE approximation with locally supported basis functions. Direct approximation of the HiFi Jacobian matrix $\mathbb{J}_h$ by empirical interpolation has also been studied in [8].

By the above EI compression, $\boldsymbol{q}_{N,M}(\boldsymbol{y})$ is the solution of the problem

$$\mathbb{V}^\top \mathcal{I}_M \boldsymbol{r}_h(\mathbb{W}\boldsymbol{q}_{N,M}(\boldsymbol{y}); \boldsymbol{y}) = 0 . \qquad (3.77)$$

One observes that the RB solution (with EI) $\boldsymbol{q}_N(\boldsymbol{y}) \neq \boldsymbol{q}_{N,M}(\boldsymbol{y})$ due to the empirical interpolation error. Moreover, the RB-EI solution $q_{N,M}(\boldsymbol{y})$ converges to the RB solution $q_N(\boldsymbol{y})$ as $M \to N_h$.

3.5.4 A-posteriori Error Bound

For the derivation of a-posteriori error bound of the RB-EI solution $q_{N,M}(\boldsymbol{y})$ at any $\boldsymbol{y}$, recall the HiFi-PG problem in the algebraic formulation with slight abuse of notation: given any $\boldsymbol{y} \in U$, find $q_h(\boldsymbol{y}) \in \mathcal{X}_h$, such that

$$\boldsymbol{r}_h(q_h(\boldsymbol{y}); \boldsymbol{y}) = 0 . \qquad (3.78)$$

Analogously, recall the RB-EI-PG problem with slight abuse of notation: given any $\boldsymbol{y} \in U$, find $q_{N,M}(\boldsymbol{y}) \in \mathcal{X}_N$, such that

$$\mathbb{V}^\top \mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) = 0 . \qquad (3.79)$$

Subtracting (3.78) from (3.79), inserting two zero terms, one has by rearranging some terms that

$$\boldsymbol{r}_h(q_h(\boldsymbol{y}); \boldsymbol{y}) - \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) = - \left(\boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) - \mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) \right)$$
$$- \left(\mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) - \mathbb{V}^\top \mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) \right) . \qquad (3.80)$$

Taking $|| \cdot ||_2$ norm of both sides, one has for the left hand side that there exists a constant $\tilde{\beta}_h$ (due to the stability assumption in Assumption 2) and there exist $N_0 \in \mathbb{N}$ and $M_0 \in \mathbb{N}$ such that when $N > N_0$ and $M > M_0$ (due to the consistency of RB compression and EI compression), one has

$$||\boldsymbol{r}_h(q_h(\boldsymbol{y}); \boldsymbol{y}) - \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y})||_2 = ||D_q \boldsymbol{r}_h((q_h(\boldsymbol{y}); \boldsymbol{y}))(q_h(\boldsymbol{y}) - q_{N,M}(\boldsymbol{y}))||_2$$
$$+ o(||q_h(\boldsymbol{y}) - q_{N,M}(\boldsymbol{y})||_\mathcal{X}) \qquad (3.81)$$
$$\geq \tilde{\beta}_h ||q_h(\boldsymbol{y}) - q_{N,M}(\boldsymbol{y})||_\mathcal{X} .$$

Note that the constant $\tilde{\beta}_h$ can be computed empirically, e.g. at some extreme realization $\boldsymbol{y} = -\mathbf{1}$ or $\mathbf{1}$, or by SCM [46, 45]. On the other hand, the right hand side (RHS) of (3.81) can be bounded by

$$\text{RHS} \leq ||\boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) - \mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y})||_2$$
$$+ ||\mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) - \mathbb{V}^\top \mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y})||_2 , \qquad (3.82)$$

where the first term accounts for the empirical interpolation error, which can be approximated by

$$||\boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) - \mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y})||_2 \approx ||\mathcal{I}_{M+M'} \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) - \mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y})||_2$$

$$\text{(by (3.68))} = \left(\sum_{m,m'=M+1}^{M+M'} c_m c_{m'} \boldsymbol{r}_m^\top \boldsymbol{r}_{m'} \right)^{1/2}, \tag{3.83}$$

where one assumes that $\mathcal{I}_{M+M'}$ for some constant $M' \in \mathbb{N}$, e.g. $M' = 2$, is a more accurate EI compression operator for the residual, so that $(\mathcal{I} - \mathcal{I}_M)\boldsymbol{r}_h \approx (\mathcal{I}_{M+M'} - \mathcal{I}_M)\boldsymbol{r}_h$. The quantities $\boldsymbol{r}_m^\top \boldsymbol{r}_{m'}$, $m, m' = M+1, \dots, M+M'$, can be computed for only once, while for any $\boldsymbol{y} \in U$, the coefficients c_m, $m = M+1, \dots, M+M'$, can be evaluated by (3.70) with $O(M + M')$ operations.

The second term of (3.82) represents the RB compression error, which can be evaluated as (by noting that (3.79) holds)

$$||\mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) - \mathbb{V}^\top \mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y})||_2 = ||\mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y})||_2$$

$$= \left((\mathbb{P}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}))^\top \left((\mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1})^\top \mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1} \right) \mathbb{P}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) \right)^{1/2}, \tag{3.84}$$

where $(\mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1})^\top \mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1}$ can be evaluated for only once; given any $\boldsymbol{y} \in U$, evaluation of $\mathbb{P}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y})$ takes $O(MN)$ operations. Therefore, once the $\boldsymbol{y}$-independent quantities are (pre)computed, evaluation of the a-posteriori error bounds for both the EI compression error and the RB compression error can be achieved efficiently, with cost depending only on N and M for each given $\boldsymbol{y} \in U$.

Finally, one can define the a-posteriori error bound of the RB-EI compression error as

$$\triangle_N(\boldsymbol{y}) := \frac{1}{\tilde{\beta}_h} \left(||\boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y}) - \mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y})||_2 + ||\mathcal{I}_M \boldsymbol{r}_h(q_{N,M}(\boldsymbol{y}); \boldsymbol{y})||_2 \right), \tag{3.85}$$

which can be efficiently evaluated for each $\boldsymbol{y} \in U$ with cost independent of the HiFi dof N_h.

3.5.5 Stable RB-EI-PG Compression

By following the same procedure as in the linear and affine case in section 3.4.4, a stable RB-EI-PG compression problem can be obtained by least-squares formulation as: given $\boldsymbol{y} \in U$, with some initial solution $\boldsymbol{q}_{N,M}^{(1)} \in \mathbb{R}^N$, for $k = 1, 2$, find $\boldsymbol{\delta q}_{N,M}^{(k)}(\boldsymbol{y}) \in \mathbb{R}^N$, such that

$$\mathbb{J}_{N,M}^s(\boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \boldsymbol{\delta q}_{N,M} = -\boldsymbol{r}_h^s(\boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y}), \tag{3.86}$$

where the Jacobian matrix $\mathbb{J}_{N,M}^s(\boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y})$ with stabilization is given by

$$(\mathbb{P}_M \mathbb{J}_h(\mathbb{W} \boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \mathbb{W})^\top (\mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1})^\top \mathbb{M}_h^{-1} \mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1} \mathbb{P}_M \mathbb{J}_h(\mathbb{W} \boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \mathbb{W}, \tag{3.87}$$

with $(\mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1})^\top \mathbb{M}_h^{-1} \mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1}$ evaluated once and for all and with $\mathbb{P}_M \mathbb{J}_h(\mathbb{W} \boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \mathbb{W}$ evaluated in $O(M^2 N)$ operations. The stabilized RB-EI residual vector $\boldsymbol{r}_h^s(\boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y})$ is given by

$$(\mathbb{P}_M \mathbb{J}_h(\mathbb{W} \boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y}) \mathbb{W})^\top (\mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1})^\top \mathbb{M}_h^{-1} \mathbb{R}_M (\mathbb{P}_M \mathbb{R}_M)^{-1} \mathbb{P}_M \boldsymbol{r}_h(\mathbb{W} \boldsymbol{q}_{N,M}^{(k)}(\boldsymbol{y}); \boldsymbol{y}), \tag{3.88}$$

which can also be efficiently evaluated for each given $\boldsymbol{y}$ with an additional $O(MN)$ operations. Then until certain criterion is met, e.g. $||\boldsymbol{\delta q}_{N,M}^{(k)}||_2 \leq \varepsilon_{tol}$, the solution is updated as $\boldsymbol{q}_{N,M}^{(k+1)} = \boldsymbol{q}_{N,M}^{(k)} + \eta^{(k)} \boldsymbol{\delta q}_{N,M}^{(k)}$ with suitable constant $\eta^{(k)}$ obtained by a line search method.

3.6 Sparse Grid RB Construction

For the construction of the RB space $\mathcal{X}_N$, one can directly solve the optimization problem (3.23) with the true error replaced by suitable a-posteriori error estimate, for instance

$$\boldsymbol{y}^{N+1} := \operatorname*{argsup}_{\boldsymbol{y} \in U} \triangle_N(\boldsymbol{y}) \ . \tag{3.89}$$

This approach has been adopted in [7] by solving model-constrained optimization problems with Lagrangian formulation, which requires both full and reduced solution of adjoint problems, leading to possibly many more expensive solution of HiFi problems than the number of reduced basis functions. In very high or infinite dimensional parameter space, the optimization problem is typically very difficult to solve as there might be many local maximal points.

A more common approach is to replace the parameter space U by a training set Ξ_t, which consists of a finite number of samples that are rich enough to construct the most representative RB space, yet should be limited due to the constraint of computational cost. Hence, it remains to seek the next sample according to

$$\boldsymbol{y}^{N+1} := \operatorname*{argmax}_{\boldsymbol{y} \in \Xi_t} \triangle_N(\boldsymbol{y}) \ . \tag{3.90}$$

To choose the training samples, random sampling methods have been mostly used in practice [59]; adaptive sampling with certain saturation criteria [40] has also been developed recently to remove and add samples from the training set. In the present setting of uncertainty parametrization as introduced in section 2.2, one takes advantage of the sparsity of the parametric data-to-solution map which is implied by $(\boldsymbol{b}, p)$-holomorphy. This sparsity allows for dimension-independent convergence rates of adaptive sparse grid sampling based on an adaptive construction of a generalized/anisotropic sparse grid in the high-dimensional parameter space. The basic idea is to build the RB space $\mathcal{X}_N$ (and EI basis for nonlinear and nonaffine problems) in tandem with the adaptive construction of the sparse grid, see [14, 12] for more details. The advantages of this approach are threefold: the first is that the training samples as well as the sparse grid nodes for RB construction are "the most representative ones"; the second is that the computational cost for the sparse grid construction is reduced by replacing the HiFi solution at each sparse grid node by its RB surrogate solution. This provides a new algorithm for fast sparse grid construction with certificated accuracy; third, one can obtain an explicitly computable a priori error estimate for the RB compression error based on a computable bound of the sparse grid interpolation error, as stated in Theorem 2.3. However, these advantages are less pronounced if the parameter dependence of the parametric solution family of the forward UQ problem is less sparse; specifically, if the sparsity parameter p being $0 < p < 1$ in the $(\boldsymbol{b}, p)$-holomorphic property becomes large and close to 1.

4 Inverse UQ

The abstract, parametric problems which arise in forward UQ in Section 2 consisted in computing, for given, admissible uncertain input datum $u \in X$ (respective for any parameter sequence $\boldsymbol{y}$ in the parametrization (2.13) of u), an approximate response $q(u) \in \mathcal{X}$, respectively a *Quantity of Interest (QoI for short)* $\phi(q) \in \mathcal{Z}$ where $\phi(\cdot) : \mathcal{X} \to \mathcal{Z}$ is a continuous mapping, and $\mathcal{Z}$ denotes a suitable space containing realizations of the QoI. If, for example, solution values $q(u)$ are of interest, one chooses $\mathcal{Z} = \mathcal{X}$, if $\phi(\cdot) \in \mathcal{X}'$, one has $\mathcal{Z} = \mathbb{R}$. The sparsity results in Section 2, in particular the constructive interpolation approximation result Theorem 2.3 and the MOR results in

4.1 Bayesian Inverse Problems for Parametric Operator equations

Following [67, 29, 65, 61, 62], one equips the space of uncertain inputs X and the space of solutions $\mathcal{X}$ of the forward maps with norms $\| \cdot \|_X$ and with $\| \cdot \|_{\mathcal{X}}$, respectively. Consider the abstract (possibly nonlinear) operator equation (2.5) where the uncertain operator $A(\cdot; u) \in C^1(\mathcal{X}, \mathcal{Y}')$ is assumed to be boundedly invertible, at least locally for the uncertain input u sufficiently close to a nominal input $\langle u \rangle \in X$, i.e. for $\|u - \langle u \rangle\|_X$ sufficiently small so that, for such u, the response of the forward problem (2.5) is uniquely defined. Define the *forward response map*, which relates a given uncertain input u

and a given forcing F to the response q in (2.5) by

$$X \ni u \mapsto q(u) := G(u; F(u)) , \quad \text{where} \quad G(u, F) : X \times \mathcal{Y}' \mapsto \mathcal{X} . \tag{4.1}$$

To ease notation, one does not list the dependence of the response on F and simply denote the dependence of the forward solution on the uncertain input as $q(u) = G(u)$. Assume given an observation functional $\mathcal{O}(\cdot) : \mathcal{X} \to Y$, which denotes a *bounded linear observation operator* on the space $\mathcal{X}$ of observed system responses in Y. Throughout the remainder of this paper, one assumes that there is a finite number K of sensors, so that $Y = \mathbb{R}^K$ with $K < \infty$. Then $\mathcal{O} \in \mathcal{L}(\mathcal{X}; Y) \simeq (\mathcal{X}')^K$. One equips $Y = \mathbb{R}^K$ with the Euclidean norm, denoted by $|\cdot|$. For example, if $\mathcal{O}(\cdot)$ is a K-vector of observation functionals $\mathcal{O}(\cdot) = (o_k(\cdot))_{k=1}^K$.

In this setting, one wishes to predict *computationally* an expected (under the Bayesian posterior) system response of the QoI, conditional on given, noisy measurement data δ. Specifically, the data δ is assumed to consist of observations of system responses in the data space Y, corrupted by additive observation noise, e.g. by a realization of a random variable η taking values in Y with law $\mathbb{Q}_0$. One assumes additive, centered gaussian noise on the observed data $\delta \in Y$. Ie., the data δ is composed of the observed system response and the additive noise η according to $\delta = \mathcal{O}(G(u)) + \eta \in Y$. One assumes that $Y = \mathbb{R}^K$ and η is Gaussian, i.e. a random vector $\eta \sim \mathbb{Q}_0 \sim \mathcal{N}(0, \Gamma)$ with a positive definite covariance Γ on $Y = \mathbb{R}^K$ (i.e., a symmetric, positive definite covariance matrix $\Gamma \in \mathbb{R}^{K \times K}_{sym}$ which is assumed to be known. The *uncertainty-to-observation map* of the system $\mathcal{G} : X \to Y = \mathbb{R}^K$ is $\mathcal{G} = \mathcal{O} \circ G$, so that

$$\delta = \mathcal{G}(u) + \eta = (\mathcal{O} \circ G)(u) + \eta \in Y , \tag{4.2}$$

where $Y = L^2_\Gamma(\mathbb{R}^K)$ denotes random vectors taking values in $Y = \mathbb{R}^K$ which are square integrable with respect to the Gaussian measure on $Y = \mathbb{R}^K$. Bayes' formula [67, 29] yields a density of the Bayesian posterior with respect to the prior whose negative log-likelihood equals the observation noise covariance-weighted, least squares functional (also referred to as "potential" in what follows) $\Phi_\Gamma : X \times Y \to \mathbb{R}$ by $\Phi_\Gamma(u; \delta) = \frac{1}{2} |\delta - \mathcal{G}(u)|^2_\Gamma$, ie.

$$\Phi_\Gamma(u; \delta) = \frac{1}{2} |\delta - \mathcal{G}(u)|^2_\Gamma := \frac{1}{2} \left((\delta - \mathcal{G}(u))^\top \Gamma^{-1} (\delta - \mathcal{G}(u)) \right) . \tag{4.3}$$

In [67, 29], an infinite-dimensional version of Bayes' rule was shown to hold in the present setting. In particular, the local Lipschitz assumption (2.12) on the solutions' dependence on the data implies a corresponding Lipschitz dependence of the Bayesian Potential (4.3) on $u \in X$. Specifically, there holds the following version of Bayes' theorem. Bayes' Theorem states that, under appropriate continuity conditions on the uncertainty-to-observation map $\mathcal{G} = (\mathcal{O} \circ G)(\cdot)$ and on the prior measure $\boldsymbol{\pi}_0$ on $u \in X$, *for positive observation noise covariance* Γ in (4.3), the posterior $\boldsymbol{\pi}^\delta$ of $u \in X$ given data $\delta \in Y$ is absolutely continuous with respect to the prior $\boldsymbol{\pi}_0$.

Theorem 4.1 *([29, Thm. 3.3]) Assume that the potential $\Phi_\Gamma : X \times Y \mapsto \mathbb{R}$ is, for given data $\delta \in Y$, $\boldsymbol{\pi}_0$ measurable on $(X, \mathcal{B}(X))$ and that, for $\mathbb{Q}_0$-a.e. data $\delta \in Y$ there holds*

$$Z := \int_X \exp\left(-\Phi(u; \delta)\right) \boldsymbol{\pi}_0(\mathrm{d}u) > 0 .$$

Then the conditional distribution of $u|\delta$ exists and is denoted by $\boldsymbol{\pi}^\delta$. It is absolutely continuous with respect to $\boldsymbol{\pi}_0$ and there holds

$$\frac{d\boldsymbol{\pi}^\delta}{d\boldsymbol{\pi}_0}(u) = \frac{1}{Z} \exp\left(-\Phi(u; \delta)\right) . \tag{4.4}$$

In particular, then, the Radon-Nikodym derivative of the Bayesian posterior w.r.t. the prior measure admits a bounded density w.r.t. the prior $\boldsymbol{\pi}_0$ which is denoted by Θ, and which is given by (4.4).

4.2 Parametric Bayesian posterior

The uncertain datum u in the forward equation (2.5) is parametrized as in (2.13). Motivated by [61, 62], the basis for the presently proposed deterministic quadrature approaches for Bayesian estimation

via the computational realization of Bayes' formula is a *parametric, deterministic representation* of the derivative of the posterior measure $\boldsymbol{\pi}^\delta$ with respect to the *uniform prior measure $\boldsymbol{\pi}_0$ on the set U of coordinates in the uncertainty parametrization* (2.25). The prior measure $\boldsymbol{\pi}_0$ being uniform, one admits in (2.13) sequences $\boldsymbol{y}$ which take values in the parameter domain $U = [-1,1]^{\mathbb{J}}$, with an index set $\mathbb{J} \subset \mathbb{N}$. Consider the countably-parametric, deterministic forward problem in the probability space

$$(U, \mathcal{B}, \boldsymbol{\pi}_0) . \tag{4.5}$$

To ease notation, one assumes throughout what follows that the prior measure $\boldsymbol{\pi}_0$ on the uncertain input $u \in X$, parametrized in the form (2.13), is the uniform measure (the ensuing derivations are still applicable if $\boldsymbol{\pi}_0$ is absolutely continuous with respect to the uniform measure, with a smooth and bounded density). Being π_0 a countable product probability measure, this assumption implies the statistical independence of the coordinates y_j in the parametrization (2.13). With the parameter domain U as in (4.5) the parametric uncertainty-to-observation map $\Xi : U \to Y = \mathbb{R}^K$ is given by

$$\Xi(\boldsymbol{y}) = \mathcal{G}(u)\Big|_{u = \langle u \rangle + \sum_{j \in \mathbb{J}} y_j \psi_j} . \tag{4.6}$$

Our reduced basis approach is based on a parametric version of Bayes' Theorem 4.1, in terms of the uncertainty parametrization (2.13). To present it, one views U as the unit ball in $\ell^\infty(\mathbb{J})$, the Banach space of bounded sequences taking values in U.

Theorem 4.2 *Assume that $\Xi : \bar{U} \to Y = \mathbb{R}^K$ is bounded and continuous. Then $\boldsymbol{\pi}^\delta(d\boldsymbol{y})$, the distribution of $\boldsymbol{y} \in U$ given data $\delta \in Y$, is absolutely continuous with respect to $\boldsymbol{\pi}_0(d\boldsymbol{y})$, i.e. there exists a parametric density $\Theta(\boldsymbol{y})$ such that*

$$\frac{d\boldsymbol{\pi}^\delta}{d\boldsymbol{\pi}_0}(\boldsymbol{y}) = \frac{1}{Z}\Theta(\boldsymbol{y}) \tag{4.7}$$

with $\Theta(\boldsymbol{y})$ given by

$$\Theta(\boldsymbol{y}) = \exp\big(-\Phi_\Gamma(u; \delta)\big)\Big|_{u = \langle u \rangle + \sum_{j \in \mathbb{J}} y_j \psi_j} , \tag{4.8}$$

with Bayesian potential Φ_Γ as in (4.3) and with normalization constant Z given by

$$Z = \mathbb{E}^{\boldsymbol{\pi}_0}[\Theta] = \int_U \Theta(\boldsymbol{y}) \, d\boldsymbol{\pi}_0(\boldsymbol{y}) > 0 . \tag{4.9}$$

Bayesian inversion is concerned with the 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 $\mathcal{Z}$) of the QoI ϕ, conditional on given (noisy) observation data $\delta \in Y$. In particular the choice $\phi(u) = G(u)$ (with $\mathcal{Z} = \mathcal{X}$) facilitates computation of the "most likely" (as expectation under the posterior, given data δ) system response. With the QoI ϕ one associates the countably-parametric map

$$\Psi(\boldsymbol{y}) = \Theta(\boldsymbol{y})\phi(u) \,|_{u = \langle u \rangle + \sum_{j \in \mathbb{J}} y_j \psi_j} = \exp\big(-\Phi_\Gamma(u; \delta)\big)\phi(u)\Big|_{u = \langle u \rangle + \sum_{j \in \mathbb{J}} y_j \psi_j} : U \to \mathcal{Z} . \tag{4.10}$$

Then the Bayesian estimate of the QoI ϕ, given noisy observation data δ, reads

$$\mathbb{E}^{\boldsymbol{\pi}^\delta}[\phi] = Z'/Z, \ \ Z' := \int_{\boldsymbol{y} \in U} \Psi(\boldsymbol{y}) \, \boldsymbol{\pi}_0(d\boldsymbol{y}), \ \ Z := \int_{\boldsymbol{y} \in U} \Theta(\boldsymbol{y}) \, \boldsymbol{\pi}_0(d\boldsymbol{y}) . \tag{4.11}$$

The task in computational Bayesian estimation is therefore to approximate the ratio $Z'/Z \in \mathcal{Z}$ in (4.11). In the parametrization with respect to $\boldsymbol{y} \in U$, Z and Z' take the form of infinite-dimensional, iterated integrals with respect to the prior $\boldsymbol{\pi}_0(d\boldsymbol{y})$.

4.3 Well-posedness and approximation

For the computational viability of Bayesian inversion the quantity $\mathbb{E}^{\boldsymbol{\pi}^\delta}[\phi]$ should be stable under perturbations of the data δ and under changes in the forward problem stemming, for example, from discretizations as considered in Sections 3.1, 3.2.

Unlike deterministic inverse problems where the data-to-solution maps can be severely ill-posed, for $\Gamma > 0$ the expectations (4.11) are Lipschitz continuous with respect to the data δ, *provided that the potential Φ_Γ in (4.3) is locally Lipschitz with respect to the data δ* in the following sense.

Assumption 4 *Let $\tilde{X} \subseteq X$ and assume $\Phi_\Gamma \in C(\tilde{X} \times Y; \mathbb{R})$ is Lipschitz on bounded sets. Assume also that there exist functions $M_i : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_+$ (depending on $\Gamma > 0$) which are monotone, non-decreasing separately in each argument, such that for all $u \in \tilde{X}$, and for all $\delta, \delta_1, \delta_2 \in B_Y(0, r)$*

$$\Phi(u; \delta) \geq -M_1(r, \|u\|_X), \tag{4.12}$$

and

$$|\Phi_\Gamma(u; \delta_1) - \Phi_\Gamma(u; \delta_2)| \leq M_2(r, \|u\|_X)\|\delta_1 - \delta_2\|_Y . \tag{4.13}$$

Under Assumption 4, the expectation (4.11) depends Lipschitz on δ (see [29, Sec. 4.1] for a proof):

$$\forall \phi \in L^2(\boldsymbol{\pi}^{\delta_1}, X; \mathbb{R}) \cap L^2(\boldsymbol{\pi}^{\delta_2}, X; \mathbb{R}) \quad \|\mathbb{E}^{\boldsymbol{\pi}^{\delta_1}}[\phi] - \mathbb{E}^{\boldsymbol{\pi}^{\delta_2}}[\phi]\|_{\mathcal{Z}} \leq C(\Gamma, r)\|\delta_1 - \delta_2\|_Y . \tag{4.14}$$

Below, one shall be interested in the impact of approximation errors in the forward response of the system (e.g. due to discretization and approximate numerical solution of system responses) on the Bayesian predictions (4.11). For continuity of the expectations (4.11) w.r.t. changes in the potential, the following assumption is imposed.

Assumption 5 *Let $\tilde{X} \subseteq X$ and assume $\Phi \in C(\tilde{X} \times Y; \mathbb{R})$ is Lipschitz on bounded sets. Assume also that there exist functions $M_i : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_+$ which are monotonically non-decreasing separately in each argument, such that for all $u \in \tilde{X}$, and all $\delta \in B_Y(0, r)$, Equation (4.12) is satisfied and*

$$|\Phi_\Gamma(u; \delta) - \Phi_\Gamma^N(u; \delta)| \leq M_2(r, \|u\|_X)\|\delta\|_Y \psi(N) \tag{4.15}$$

where $\psi(N) \to 0$ as $N \to \infty$.

By $\boldsymbol{\pi}_N^\delta$ one denotes the Bayesian posterior, given data $\delta \in Y$, with respect to Φ_Γ^N.

Proposition 4.3 *Under Assumption 5, and the assumption that for $\tilde{X} \subseteq X$ and for some bounded $B \subset X$ one has $\boldsymbol{\pi}_0(\tilde{X} \cap B) > 0$ and*

$$X \ni u \mapsto \exp(M_1(\|u\|_X))(M_2(\|u\|_X))^2 \in L^1_{\boldsymbol{\pi}_0}(X; \mathbb{R}) ,$$

there holds, for every QoI $\phi : X \to \mathcal{Z}$ such that, although the convergence rate s can be substantially higher than the rate $1/2$ afforded by MCMC methods (cp. Section 4.4 and [29, Sections 5.1, 5.2]) that $\phi \in L^2_{\boldsymbol{\pi}^\delta}(X; \mathcal{Z}) \cap L^2_{\boldsymbol{\pi}_N^\delta}(X; \mathcal{Z})$ uniformly w.r.t. N, that $Z > 0$ in (4.9) and

$$\|\mathbb{E}^{\boldsymbol{\pi}^\delta}[\phi] - \mathbb{E}^{\boldsymbol{\pi}_N^\delta}[\phi]\|_{\mathcal{Z}} \leq C(\Gamma, r)\|\delta\|_Y \psi(N) . \tag{4.16}$$

For a proof of Proposition 4.3, see [29, Thm. 4.7, Rem. 4.8].

Below, concrete choices are presented for the convergence rate function $\psi(N)$ in estimates (4.15), (4.16) in terms of i) *"dimension truncation"* of the uncertainty parametrization (2.13), i.e. to a finite number of $s \geq 1$ terms in (2.13), and ii) *Petrov-Galerkin Discretization* of the dimensionally truncated problem, iii) *generalized polynomial chaos (gpc) approximation* of the dimensionally truncated problem for particular classes of forward problems. The verification of the consistency condition (4.15) in either of these cases will be based on (cf. [32])

Proposition 4.4 *Assume given a sequence $\{q^N\}_{N \geq 1}$ of approximations to the parametric forward response $X \ni u \mapsto q(u) \in \mathcal{X}$ such that, with the parametrization (2.13),*

$$\sup_{\boldsymbol{y} \in U} \|(q - q^N)(\boldsymbol{y})\|_{\mathcal{X}} \leq \psi(N) \tag{4.17}$$

with a consistency error bound $\psi \downarrow 0$ as $N \to \infty$ monotonically and uniformly w.r.t. $u \in \tilde{X}$ (resp. w.r.t. $\boldsymbol{y} \in U$). By G^N one denotes the corresponding (Galerkin) approximations of the parametric

forward maps. Then the approximate Bayesian potential

$$\Phi^N(u;\delta) = \frac{1}{2}(\delta - \mathcal{G}^N(u))^\top \Gamma^{-1}(\delta - \mathcal{G}^N(u)) : X \times Y \mapsto \mathbb{R}, \qquad (4.18)$$

where $\mathcal{G}^N := \mathcal{O} \circ G^N$, satisfies (4.15).

The preceding result shows that the consistency condition (4.17) for the approximate forward map q^N ensures corresponding consistency of the Bayesian estimate $\mathbb{E}^{\boldsymbol{\pi}_N^\delta}[\phi]$, due to (4.16). Note that so far, *no specific assumption on the nature of approximation of the forward map has been made.* Using a MOR surrogate of the parametric forward model, Theorem 3.4 allows to bound $\psi(N)$ in (4.17) by the corresponding worst case RB compression error σ_N in (3.26) which, is bound by the convergence rate of the corresponding N-width:

$$\psi(N) \le \sigma_N \lesssim d_N(\mathcal{M}_h; \mathcal{X}_h) \qquad (4.19)$$

in the various cases indicated in Theorem 3.4. Under Assumption 4, Proposition 4.4 ensures that Assumption 5 holds, with (4.19). One concludes in particular that replacing the forward model by a reduced basis surrogate will result in an error in the Bayesian estimate of the same asymptotic order of magnitdue, as $N \to \infty$. This justifies, for example, running Markov chains on the surrogate forward model obtained from MOR. In doing this, however, care must be taken to account for the constants implied by $\lesssim$ in (4.19): the constants do not depend on N, but large values of these constants can imply prohibitive errors for the (small) values N of the number of RB degrees of freedom employed in PG projections of MOR forward surrogates. In addition, it is pointed out that the estimate (4.16) depends on the observation noise covariance Γ, as well as on the size r of the observation data δ (measured in Y).

4.4 Reduced basis acceleration of MCMC

Markov Chain Monte Carlo (MCMC) methods compute the expectation $\mathbb{E}^{\boldsymbol{\pi}^\delta}[\phi]$ in (4.11) under the posterior by sampling from the posterior density. They proceed in approximation of the constant Z' in (4.11) by sample averages where, however, *the posterior distribution from which samples are to be drawn is itself to be determined during the course of the sampling process.* MCMC methods start by sampling from the (known) prior $\boldsymbol{\pi}_0$, and by updating, in the course of sampling, both numerator Z' as well as the normalizing constant Z in (4.11). Several variants exist; see [29, Sections 5.1,5.2] for a derivation of the *Metropolis-Hastings MCMC*, and to [29, Section 5.3] for *sequential Monte-Carlo (sMC) methods.* A convergence theory in term of certain *spectral gaps* is provided in [29, Thm. 5.13], resulting in the convergence rate $N^{-1/2}$ with N denoting the number of increments of the chain. In the context of the present paper, N denotes the number of (approximate) solves of the parametric forward problem (2.4), resp. of a discretization of it. Due to the low rate $1/2$ of the MCMC methods (and due to the high rejection rate of the samplers during burn-in), generally a very large number of samples is required. Moreover, successive updates of the MCMC samplers has an intrisically serial structure which, in turn, foils massive parallelism to compensate for the slow convergence rate. It is therefore of high interest to examine the possibility of accelerating MCMC methods. In [44], the impact of various discretization and acceleration techniques for MCMC methods were analyzed for the computation of the expectation $\mathbb{E}^{\boldsymbol{\pi}^\delta}[\phi]$ in (4.11); among them a generalized polynomial chaos (gpc) surrogate of the parametric forward may $U \ni \boldsymbol{y} \to q(\boldsymbol{y}) \in \mathcal{X}$. The theory in [44] can be extended using the consistency error bound Assumption 5 in the Bayesian potential, and Proposition 4.4 for the RB error in the forward map. Practical application and implementation of RB with MCMC for (Bayesian) inverse problems can be found, for instance, in [27, 49].

4.5 Dimension and order adaptive, deterministic quadrature

The parametric, deterministic infinite-dimensional integrals Z' and Z in (4.11) are, in principle, accessible to any quadrature strategy which is able to deal efficiently with the high dimension of the integration domain, and which is able to exploit $(\boldsymbol{b}, p)$-sparsity of the parametric integrand functions.

Following [39, 22], a greedy strategy based on reduced sets of indices which are neighboring the currently active set Λ, defined by

$$\mathcal{N}(\Lambda) := \{\nu \notin \Lambda : \nu - e_j \in \Lambda, \forall j \in \mathbb{I}_\nu \text{ and } \nu_j = 0, \forall j > j(\Lambda) + 1\}$$

for any downward closed index set $\Lambda \subset \mathcal{F}$ of currently active gpc modes, where $j(\Lambda) := \max\{j : \nu_j > 0 \text{ for some } \nu \in \Lambda\}$. This heuristic approach aims at controlling the global approximation error by locally collecting indices of the current set of neighbors with the largest estimates error contributions. In the following, the resulting algorithm to recursively build the downward closed index set Λ in the Smolyak quadrature which is adapted to the posterior density (and, due to the explicit expression (4.8) from Bayes' formula, also to the observation data δ) is summarized. The reader is referred to [39, 62, 61] for details and numerical results. Development and analysis of the combination of RB, MOR and ASG for Bayesian inversion are described in depth in [18, 19, 20], for both linear and nonlinear, both affine and nonaffine parametric problems.

1: **function** ASG
2:     Set $\Lambda_1 = \{0\}$, $k = 1$ and compute $\Delta_0(\Xi)$.
3:     Determine the reduced set (4.20) of neighbors $\mathcal{N}(\Lambda_1)$.
4:     Compute $\Delta_\nu(\Xi)$, $\forall \nu \in \mathcal{N}(\Lambda_1)$.
5:     **while** $\sum_{\nu \in \mathcal{N}(\Lambda_k)} \|\Delta_\nu(\Xi)\|_{\mathcal{S}} > tol$ **do**
6:         Select ν from $\mathcal{N}(\Lambda_k)$ with largest $\|\Delta_\nu\|_{\mathcal{S}}$ and set $\Lambda_{k+1} = \Lambda_k \cup \{\nu\}$.
7:         Determine the reduced set (4.20) of neighbors $\mathcal{N}(\Lambda_{k+1})$.
8:         Compute $\Delta_\nu(\Xi)$, $\forall \nu \in \mathcal{N}(\Lambda_{k+1})$.
9:         Set $k = k + 1$.
10:    **end while**
11: **end function**

4.6 Quasi Monte-Carlo Quadrature

The adaptive, deterministic quadrature Algorithm 4.5 realizes, theoretically and in practical experiments ([62, 61]), convergence rates $s = 1/p - 1$ which are determined only by the summability exponent p of the sequence $\boldsymbol{b}$ of (X-norms of) the basis $\boldsymbol{\Psi}$ adopted for the space X, in order to parametrize the uncertain input data u as in (2.13). The downside/ drawback of Algorithm 4.5 is that, although the (dimension-independent) convergence rate s can be substantially higher than the rate $1/2$ afforded by MCMC methods (cp. Section 4.4), provided the summability exponent p is sufficiently small, that it is intrinsically sequential in nature, due to the recursive construction of the active index sets; in this respect, it is analogous to MCMC methods which access the forward model (or a surrogate of it) through uncertainty instances produced by the sampler along the Markov chains. An alterative to these approaches which allows for dimension-independent convergence rate $s = 1/p$ in terms of the number of samples *and* which allows simultaneous, parallel access to the forward model in all instances of the uncertainty is the recently developed, higher order Quasi Monte-Carlo integration.

A remedy which affords dimension-independent, higher-order convergence rates limited only by uncertain input sparsity *and* which allows, at the same time, fully parallel evaluation of the integrals Z' and Z in the Bayesian estimate (4.11) is afforded by *Quasi Monte-Carlo Integration*. The reader is referred to [35] for a general survey, and numerous references. It has recently been shown that $(\boldsymbol{b}, p)$ sparsity implies, indeed, the dimension-independent convergence rate $s = 1/p$ for certain types of higher order QMC integration; see [34] for the theory for linear, affine parametric operator equations $q \mapsto A(\boldsymbol{y})q$ and to [33] for multilevel extensions, to [36] for the verification of the convergence conditions in [34] implied by $(\boldsymbol{b}, p)$ holomorphy. Computational construction of higher order QMC integration rules on the bounded parameter domain U is described in [38]. There exist also QMC integration methods for unbounded parameter regions. Such arise typically for *Gaussian random field (GRF for short) inputs u* taking values in X. Upon uncertainty parametrization with, for example, a Karhunen–Loève expansion into eigenfunctions of the covariance operator of the GRF (2.17), there result parametric deterministic problems with unbounded parameter ranges (consisting, for GRF's, of countable cartesian products of real lines, ie. $U = \mathbb{R}^{\mathbb{N}}$). In this case, the present theory still is

applicable, however, *all stability and equivalence constants will depend, generally, on the parameter* $\boldsymbol{y}$ *with the parametric dependence degenerating for "extremal events", i.e. realizations of u whose parameters in the tail of the prior* $\boldsymbol{\pi}_0$. *This is particularly relevant for uncertain input data which involve a gaussian random field* (2.17) *in some form.*

5 Software

- **rbMIT:** The general algorithms of RB based on Finite Element are implemented in the software package rbMIT ©MIT in MATLAB. It is implemented mainly for demonstration and education. However, it is also friendly to use for development and test of new algorithms. The code and an accompanying textbook [55] are available through the link:
  `http://augustine.mit.edu/methodology/methodology_rbMIT_System.htm`

- **RBniCS:** An RB extension of the Finite Element software package `FEniCS` ([50]) is under development and public domain through the link:`http://mathlab.sissa.it/rbnics`. Implementation includes POD and greedy algorithm for coercive problems, which is suited for an introductory course of RB together with the book [41].

- **Dune-RB:** It is a module for the Dune (Distributed and Unified Numerics Environment) library in C++. Template classes are available for RB construction based on several HiFi discretizations, including Finite Element and Finite Volume. Parallelization is available for RB construction too. Tutorials and code are available at `http://www.dune-project.org/`.

- **pyMOR:** pyMOR is implemented in Python for MOR for parameterized PDE. It has friendly interfaces and proper integration with external high-dimensional PDE solvers. Finite element and finite volume discretizations implemented based on the library of NumPy/SciPy are available. For more information, see `http://pymor.org`.

- **AKSELOS:** MOR remains the core technology for the startup company AKSELOS in several engineering fields, such as port infrastructure and industrial machinery. Different components are included such as FEA and CAD. The HiFi solution in AKSELOS is implemented with a HPC and cloud-based simulation platform and is available for commercial use. For more information see `http://www.akselos.com/`.

- For further libraries/software packages, we refer to `http://www.ians.uni-stuttgart.de/MoRePaS/software/`.

6 Conclusion

In this work, both mathematical and computational foundations of model order reduction techniques for UQ problems with distributed uncertainties are surveyed. Based on the recent development of sparse polynomial approximation in infinite dimensions, the convergence property of MOR constructed by greedy algorithm is established. In particular, under the sparsity of the uncertainties and the holomorphy of the forward solution maps w.r.t. the parameters, the dimension-independent convergence rate of the RB compression error can be achieved. Details of the construction and the compression of MOR are provided for both affine and nonaffine, linear and nonlinear problems modelled by parametric operator equations. Stability of the HiFi approximation and the RB compression is fulfilled by Petrov–Galerkin formulation with suitably constructed test spaces. Efficient MOR construction is realized by a greedy search algorithm with sparse sampling scheme, which further leads to a fast method for sparse grid construction. The MOR techniques are applied for both forward and inverse UQ problems, leading to considerable computational reduction in the *many-query* context for evaluation of statistics, and in the *real-time* context for fast Bayesian inversion.

7 Glossary

List of used abbreviations and definition of technical terms:

- UQ: Uncertainty Quantification

- MOR: Model Order Reduction

- PDE: Partial Differential Equations

- POD: Proper Orthogonal Decomposition

- RB: Reduced Basis

- EI: Empirical Interpolation

- SG: Sparse Grid

- FEM: Finite Element Method

- PG: Petrov–Galerkin

- QoI : Quantity of interest

- Fidelity (of a mathematical model): notion of quality of responses of a computational **surrogate** model for a given mathematical model.

- HiFi: High Fidelity

- SCM: Successive Constraint Method

- Surrogate Model: numerical model obtained by various numerical approximation of a mathematical model.

References

[1] M. Barrault, Y. Maday, N.C. Nguyen, and A.T. Patera. An empirical interpolation method: application to efficient reduced-basis discretization of partial differential equations. *Comptes Rendus Mathematique, Analyse Numérique*, 339(9):667–672, 2004.

[2] L. Beirão da Veiga, A. Buffa, G. Sangalli, and R. Vázquez. Mathematical analysis of variational isogeometric methods. *Acta Numer.*, 23:157–287, 2014.

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

[4] F. Brezzi, J. Rappaz, and P.-A. Raviart. Finite-dimensional approximation of nonlinear problems. I. Branches of nonsingular solutions. *Numer. Math.*, 36(1):1–25, 1980/81.

[5] A. Buffa, Y. Maday, A.T. Patera, C. Prudhomme, and G. Turinici. A priori convergence of the greedy algorithm for the parametrized reduced basis method. *ESAIM: Mathematical Modelling and Numerical Analysis*, 46(03):595–603, 2012.

[6] T. Bui-Thanh, M. Damodaran, and K.E. Willcox. Aerodynamic data reconstruction and inverse design using proper orthogonal decomposition. *AIAA journal*, 42(8):1505–1516, 2004.

[7] 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.

[8] K. Carlberg, C. Bou-Mosleh, and C. Farhat. Efficient non-linear model reduction via a least-squares Petrov–Galerkin projection and compressive tensor approximations. *International Journal for Numerical Methods in Engineering*, 86(2):155–181, 2011.

[9] A. Chatterjee. An introduction to the proper orthogonal decomposition. *Current science*, 78(7):808–817, 2000.

[10] S. Chaturantabut and D.C. Sorensen. Nonlinear model reduction via discrete empirical interpolation. *SIAM Journal on Scientific Computing*, 32(5):2737–2764, 2010.

[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 and A. Quarteroni. A new algorithm for high-dimensional uncertainty quantification based on dimension-adaptive sparse grid approximation and reduced basis methods. *Journal of Computational Physics*, 298:176 – 193, 2015.

[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] P. Chen, A. Quarteroni, and G. Rozza. Multilevel and weighted reduced basis method for stochastic optimal control problems constrained by Stokes equations. *Numerische Mathematik*, pages 1–36, 2015.

[17] P. Chen, A. Quarteroni, and G. Rozza. Reduced order methods for uncertainty quantification problems. *ETH Report 03, Submitted*, 2015.

[18] P. Chen and Ch. Schwab. Sparse-grid, reduced-basis Bayesian inversion. Technical Report 2014-36, Seminar for Applied Mathematics, ETH Zürich, Switzerland, 2014.

[19] P. Chen and Ch. Schwab. Adaptive sparse grid model order reduction for fast Bayesian estimation and inversion. Technical Report 2015-08, Seminar for Applied Mathematics, ETH Zürich, Switzerland, 2015.

[20] P. Chen and Ch. Schwab. Sparse-grid, reduced-basis bayesian inversion: Nonaffine-parametric nonlinear equations. Technical Report 2015-21, Seminar for Applied Mathematics, ETH Zürich, 2015.

[21] P. Chen and Ch. Schwab. Exponential Convergence of GEIM. Technical report, Seminar for Applied Mathematics, ETH Zürich, 2016. (to appear).

[22] A. Chkifa, A. Cohen, R. DeVore, and Ch. Schwab. Adaptive algorithms for sparse polynomial approximation of parametric and stochastic elliptic pdes. *M2AN Math. Mod. and Num. Anal.*, 47(1):253–280, 2013.

[23] A. Chkifa, A. Cohen, and Ch. Schwab. High-dimensional adaptive sparse polynomial interpolation and applications to parametric pdes. *Journ. Found. Comp. Math.*, 14(4):601–633, 2013.

[24] Z. Ciesielski and J. Domsta. Construction of an orthonormal basis in $C^m(I^d)$ and $W_p^m(I^d)$. *Studia Math.*, 41:211–224, 1972.

[25] A. Cohen, A. Chkifa, and Ch. Schwab. Breaking the curse of dimensionality in sparse polynomial approximation of parametric pdes. *Journ. Math. Pures et Appliquees*, 103(2):400–428, 2015.

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

[27] 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.

[28] W. Dahmen, C. Plesken, and G. Welper. Double greedy algorithms: reduced basis methods for transport dominated problems. *ESAIM: Mathematical Modelling and Numerical Analysis*, 48(03):623–663, 2014.

[29] M. Dashti and A.M. Stuart. The Bayesian approach to inverse problems. *Handbook of UQ (R. Ghanem etal., Eds.)*, 2016.

[30] P. Deuflhard. *Newton methods for nonlinear problems: affine invariance and adaptive algorithms*, volume 35. Springer Science & Business Media, 2011.

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

[32] J. Dick, R. Gantner, Q.T. LeGia, and Ch. Schwab. Higher order Quasi Monte Carlo integration for Bayesian inversion of holomorphic, parametric operator equations. Technical report, Seminar for Applied Mathematics, ETH Zürich, 2015.

[33] J. Dick, Frances Y. Kuo, Quoc T. LeGia, and Christoph Schwab. Multi-level higher order QMC Galerkin discretization for affine parametric operator equations. Technical Report 2014-14, Seminar for Applied Mathematics, ETH Zürich, 2014.

[34] J. Dick, F.Y. Kuo, Q.T. Le Gia, D. Nuyens, and Ch. Schwab. Higher order QMC Petrov-Galerkin discretization for affine parametric operator equations with random field inputs. *SIAM J. Numerical Analysis*, 52(6):2676–2702, 2014.

[35] J. Dick, F.Y. Kuo, and I.H. Sloan. High-dimensional integration: the quasi-Monte Carlo way. *Acta Numer.*, 22:133–288, 2013.

[36] J. Dick, Q.T. LeGia, and Ch. Schwab. Higher order Quasi Monte Carlo integration for holomorphic, parametric operator equations. Technical Report 2014-23, Seminar for Applied Mathematics, ETH Zürich, 2014.

[37] M. Drohmann, B. Haasdonk, and M. Ohlberger. Reduced basis approximation for nonlinear parametrized evolution equations based on empirical operator interpolation. *SIAM Journal on Scientific Computing*, 34(2):A937–A969, 2012.

[38] R.N. Gantner and Ch. Schwab. Computational higher order quasi-monte carlo integration. Technical Report 2014-25, Seminar for Applied Mathematics, ETH Zürich, 2014.

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

[40] J. Hesthaven, B. Stamm, and S. Zhang. Efficient greedy algorithms for high-dimensional parameter spaces with applications to empirical interpolation and reduced basis methods. *ESAIM: Mathematical Modelling and Numerical Analysis*, 48(1):259–283, 2011.

[41] J.S. Hesthaven, G. Rozza, and B. Stamm. *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*. SpringerBriefs in Mathematics. Springer, 2016.

[42] J.S. Hesthaven, B. Stamm, and S. Zhang. Certified reduced basis method for the electric field integral equation. *SIAM J. Sci. Comput.*, 34(3):A1777–A1799, 2012.

[43] V.H. Hoang and Ch. Schwab. n-term Wiener chaos approximation rates for elliptic PDEs with lognormal Gaussian random inputs. *Math. Mod. Meth. Appl. Sci.*, 24(4):797–826, 2014.

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

[45] D.B.P. Huynh, D.J. Knezevic, Y. Chen, J.S. Hesthaven, and A.T. Patera. A natural-norm successive constraint method for inf-sup lower bounds. *Computer Methods in Applied Mechanics and Engineering*, 199(29):1963–1975, 2010.

[46] D.B.P Huynh, G. Rozza, S. Sen, and A.T. Patera. A successive constraint linear optimization method for lower bounds of parametric coercivity and inf-sup stability constants. *Comptes Rendus Mathematique, Analyse Numérique*, 345(8):473–478, 2007.

[47] T. Lassila, A. Manzoni, A. Quarteroni, and G. Rozza. Generalized reduced basis methods and n-width estimates for the approximation of the solution manifold of parametric PDEs. *Analysis and Numerics of Partial Differential Equations Series: Springer INdAM Series, Vol. 4, Brezzi, F.; Colli Franzone, P.; Gianazza, U.; Gilardi, G. (Eds.)*, 2013.

[48] T. Lassila and G. Rozza. Parametric free-form shape design with PDE models and reduced basis method. *Computer Methods in Applied Mechanics and Engineering*, 199(23):1583–1592, 2010.

[49] Toni Lassila, Andrea Manzoni, Alfio Quarteroni, and Gianluigi Rozza. A reduced computational and geometrical framework for inverse problems in hemodynamics. *Int. J. Numer. Methods Biomed. Eng.*, 29(7):741–776, 2013.

[50] A. Logg, K.A. Mardal, and G. Wells. *Automated solution of differential equations by the finite element method: The FEniCS book*, volume 84. Springer Science & Business Media, 2012.

[51] Y. Maday, O. Mula, A.T. Patera, and M. Yano. The generalized empirical interpolation method: stability theory on Hilbert spaces with an application to the Stokes equation. *Computer Methods in Applied Mechanics and Engineering*, 287:310–334, 2015.

[52] Y. Maday, O. Mula, and G. Turinici. A priori convergence of the Generalized Empirical Interpolation Method. In *10th international conference on Sampling Theory and Applications (SampTA 2013)*, pages 168–171, 2013.

[53] Y. Maday, N.C. Nguyen, A.T. Patera, and G.S.H. Pau. A general, multipurpose interpolation procedure: the magic points. *Communications on Pure and Applied Analysis*, 8(1):383–404, 2009.

[54] 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.

[55] 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.

[56] J. Pousin and J. Rappaz. Consistency, stability, a priori and a posteriori errors for Petrov-Galerkin methods applied to nonlinear problems. *Numerische Mathematik*, 69(2):213–231, 1994.

[57] 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.

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

[59] 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.

[60] G. Rozza and K. Veroy. On the stability of the reduced basis method for stokes equations in parametrized domains. *Computer Methods in Applied Mechanics and Engineering*, 196(7):1244–1260, 2007.

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

[62] C. Schillings and Ch. Schwab. Sparsity in Bayesian inversion of parametric operator equations. *Inverse Problems*, 30(6):065007, 30, 2014.

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

[64] Ch. Schwab and R. Stevenson. Space-time adaptive wavelet methods for parabolic evolution problems. *Math. Comp.*, 78(267):1293–1318, 2009.

[65] Ch. Schwab and A. M. Stuart. Sparse deterministic approximation of Bayesian inverse problems. *Inverse Problems*, 28(4):045003, 32, 2012.

[66] Ch. Schwab and R. A. Todor. Karhunen–Loève approximation of random fields by generalized fast multipole methods. *Journal of Computational Physics*, 217(1):100–122, 2006.

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

[68] K. Willcox and J. Peraire. Balanced model reduction via the proper orthogonal decomposition. *AIAA journal*, 40(11):2323–2330, 2002.

Recent Research Reports

Nr.	Authors/Title

2015-18   R. Hiptmair
           Maxwell's Equations: Continuous and Discrete

2015-19   X. Claeys and R. Hiptmair and E. Spindler
           Second-Kind Boundary Integral Equations for Scattering at Composite Partly
           Impenetrable Objects

2015-20   R. Hiptmair and A. Moiola and I. Perugia
           A Survey of Trefftz Methods for the Helmholtz Equation

2015-21   P. Chen and Ch. Schwab
           Sparse-Grid, Reduced-Basis Bayesian Inversion: Nonaffine-Parametric Nonlinear
           Equations

2015-22   F. Kuo and R. Scheichl and Ch. Schwab and I. Sloan and E. Ullmann
           Multilevel Quasi-Monte Carlo Methods for Lognormal Diffusion Problems

2015-23   C. Jerez-Hanckes and Ch. Schwab
           Electromagnetic Wave Scattering by Random Surfaces: Uncertainty Quantification via
           Sparse Tensor Boundary Elements

2015-24   V. Kazeev and Ch. Schwab
           Quantized tensor-structured finite elements for second-order elliptic PDEs in two
           dimensions

2015-25   S. Etter
           Parallel ALS Algorithm for the Hierarchical Tucker Representation

2015-26   C. Hafner and R. Hiptmair and P. Souzangar
           Data Sparse Numerical Models for SNOM Tips

2015-27   F. Leonardi and S. Mishra and Ch. Schwab
           Numerical approximation of statistical solutions of incompressible flow