MATHICSE

Mathematics Institute of Computational Science and Engineering

School of Basic Sciences - Section of Mathematics

Analysis of the discrete L² projection on polynomial spaces with random evaluations

Giovanni Migliorati, Fabio Nobile, Erik von Schwerin, Raúl Tempone

Analysis of the discrete L^2 projection on polynomial spaces with random evaluations

Giovanni Migliorati[♯], Fabio Nobile[♯†], Erik von Schwerin[‡], Raul Tempone[‡]

December 6, 2011

[♯] MOX - Dipartimento di Matematica "Francesco Brioschi", Politecnico di Milano, Italy
[†] CSQI-MATHICSE, École Polytechnique Fédérale de Lausanne, Switzerland
[‡] Applied Mathematics and Computational Sciences, KAUST, Thuwal, Saudi Arabia
giovanni.migliorati@mail.polimi.it, raul.tempone@kaust.edu.sa

Abstract

We analyse the problem of approximating a multivariate function by discrete least-squares projection on a polynomial space starting from *random, noise-free* observations. An area of possible application of such technique is Uncertainty Quantification (UQ) for computational models.

We prove an optimal convergence estimate, up to a logarithmic factor, in the monovariate case, when the observation points are sampled in a bounded domain from a probability density function bounded away from zero, provided the number of samples scales quadratically with the dimension of the polynomial space.

Several numerical tests are presented both in the monovariate and multivariate case, confirming our theoretical estimates. The numerical tests also clarify how the convergence rate depends on the number of sampling points, on the polynomial degree, and on the smoothness of the target function.

1 Introduction

Given a smooth multivariate function $\phi = \phi(Y_1, \ldots, Y_N)$, where $Y_1, \ldots, Y_N$ are random variables, we consider in this work the classical problem of approximating ϕ in a multivariate polynomial space, starting from noise-free observations of ϕ on random samples of $Y_1, \ldots, Y_N$.

1

The motivation for such work come from the field of Uncertainty Quantification (UQ) in computational models [25, 30], where often uncertainty is present in many input parameters entering the mathematical model used to describe some problem in engineering, physics, biology, etc. and can be characterized in probabilistic terms by considering the input parameters as random variables. The goal of the analysis is typically to compute statistics of the solution to the mathematical model or some output quantities of interest.

It is assumed here that, for each value of the input parameters, the solution or output quantity can be accessed without errors. This is, of course, an idealization as deterministic approximation-type errors will typically be present whenever the model involves differential or integral operators. Also, round-off errors will be present as well. However, these sources of errors are quite different in nature from "measurement errors" appearing in an experimental setting, which are usually modeled as random and statistically independent. In the context of UQ in computational models, it is therefore reasonable to assume that the approximation errors can be kept under control by some careful a-posteriori error estimation and mesh refinement (see e.g. [1, 5] and references therein).

A technique that has received considerable attention in the last few years is the so called generalized Polynomial Chaos expansion (gPC); see e.g. [20, 29]. It consists in expanding the solution in polynomials of the input random variables. The use of global polynomial spaces is sound in many situations, where the input/output (parameters-to-solution) map is smooth. This is the case, for instance, in elliptic PDEs with random diffusion coefficient [2, 13, 14, 6].

Once a truncated gPC expansion has been computed by some means, it can be used later for inexpensive computations of solution statistics or as a reduced model of the input/output map for "global sensitivity analysis" [15, 28], or optimization under uncertainty [19].

As a tool to build such a gPC approximation, we consider in this work an L^2 projection, starting from a random sample of the input parameters. Such an idea has already been proposed in the framework of UQ and has been given several names: Point Collocation [22, 17, 26], non intrusive gPC [25, 18]regression [9, 10]. As a practical recipe, the number of samples drawn from the input distribution is typically taken to be 2 to 3 times the dimension of the polynomial space.

The proposed approach falls within the classical topic of polynomial regression estimation, i.e. minimization of the empirical L^2 risk within the given polynomial space. We insist, however, that unlike the traditional statistical approach, here we consider noise-free data evaluated in random points.

A relevant question is whether such a minimizer is "optimal" in the sense that it achieves an approximation of the (unknown) function that is equivalent to the "best approximation" in the polynomial space.

A lot of literature is available on regression estimations in the noisy case. We recall here the book [21] that provides a general framework for analysis of regression estimators with random design, as well as the works [8, 7] that

show optimality of the noisy regression when using piecewise constant or linear functions. The aforementioned works give estimates on the expected L^2 error under the assumption that the function is bounded in L^∞ by some a priori fixed constant. Other works in the field of distribution-free regression with noise have derived convergence rates for the L^2 risk which are optimal up to logarithmic factors, *e.g.* [23, 24, 4].

The L^2-error in expectation is bounded by two terms: the (purely deterministic) best approximation error of the (unknown) function in the approximating space, and the statistical error due to the random sampling and the noise in the observations. The latter scales as $1/\sqrt{M}$ if M is the number of samples. In the aforementioned works, the statistical error does not vanish in the noise-free setting.

Hence, the main question that we address in this paper: in the noise-free polynomial approximation based on random samples, does the randomness of the evaluation points introduce a statistical error $\mathcal{O}(1/\sqrt{M})$ or not?

We study theoretically the problem for a monovariate function $f(Y)$ where Y is a bounded random variable with probability density function bounded away from zero. Denoting by $\#\Lambda$ the dimension of the polynomial space, we prove that the L^2 projection on a polynomial space with random samples leads to quasi optimal convergence rates (up to logarithmic factors) provided that the number of samples scales as $M \sim \#\Lambda^2$.

Similar results have been derived very recently in [12]. However, the proof proposed therein uses different techniques than ours.

We also show, in the general multivariate setting, the relation between the optimality of the projection on random points and the condition number of the corresponding design matrix, when using an orthonormal basis.

We present several numerical tests, both on monovariate and multivariate functions, that clearly show that a choice $M \sim \#\Lambda^2$ leads to a stable regression problem and an optimal approximation, whereas $M \sim \#\Lambda$ leads to an ill conditioned problem when $\#\Lambda$ is large and, eventually, to a divergent approximation. Moreover, our numerical tests show some significant differences between the one-dimensional and the multidimensional case.

The outline of the paper is the following: Section 2 introduces the approximation problem as an L^2 projection on a space of polynomials in N underlying variables; some common choices of polynomial spaces are described in Section 2.1. The optimality of the random L^2 projection, in terms of a best approximation constant, is shown in Section 2.2; the asymptotic behaviour of this best approximation constant, as the number of random evaluation points goes to infinity, is analyzed in Section 2.3. Next, Section 3 restricts the study to polynomial spaces in one variable, with uniformly distributed random points; in this case, a bound on the best approximation constant is derived and used to prove Theorem 2, which given the maximal polynomial degree provides a rule for the number of random points that makes the discrete random L^2 projection nearly optimal (up to a logarithmic factor) with any prescribed confidence level. Section 4 gives the

algebraic formulation of the random projection problem, in view of its numerical discretization. In particular, Section 4.1 provides an analysis of how the condition number of the design matrix depends on the best approximation constant of the random L^2 projector. Finally, Section 5 complements the analysis in Sections 2–4 with numerical tests, both in the one-dimensional case and in higher dimensions.

2 Discrete L^2 projection with random points

Let $\mathbf{Y} = [Y_1, \ldots, Y_N]$ be a vector of N random variables, taking values in a bounded set $\Gamma \subset \mathbb{R}^N$. We assume that Γ has a tensor structure $\Gamma = \Gamma_1 \times \cdots \times \Gamma_N$ and that the random vector $\mathbf{Y}$ has a probability density function $\rho : \Gamma \to \mathbb{R}^+$.

We consider a random variable $Z = \phi(\mathbf{Y})$, where $\phi : \Gamma \to \mathbb{R}$ is assumed to be a smooth function. The goal of the analysis is to compute statistical moments of Z. This will be achieved by first constructing a reduced model (approximate response function); i.e. we approximate the function $\phi(Y_1, \ldots, Y_N)$ by a suitable multivariate polynomial $\phi_\Lambda(Y_1, \ldots, Y_N)$. We then compute statistical moments using the approximate function ϕ_Λ.

We denote by

$$\mathbb{E}[Z] := \int_\Gamma \phi(\mathbf{Y})\rho(\mathbf{Y})d\mathbf{Y}$$

the expected value of the random variable Z and by

$$P(A) := \int_A \rho(\mathbf{Y})d\mathbf{Y}$$

the probability of the event $A \in \mathcal{B}(\Gamma)$, where $\mathcal{B}(\Gamma)$ is the Borel σ-algebra with respect to the measure $\rho(\mathbf{Y})d\mathbf{Y}$. Throughout the paper we also assume that

Assumption 1. $0 < \rho_{min} \leq \rho(\mathbf{y}) \leq \rho_{max} < \infty, \forall \mathbf{y} \in \Gamma$.

We introduce the space L_ρ^2 of square integrable functions $f : \Gamma \to \mathbb{R}$, endowed with the norm

$$\|f\|_{L_\rho^2} = \left(\int_\Gamma f^2(\mathbf{Y})\rho(\mathbf{Y})d\mathbf{Y} \right)^{1/2}.$$

Observe that under Assumption 1, the norm $\| \cdot \|_{L_\rho^2}$ is equivalent to the standard L^2 norm (with Lebesgue measure), *i.e.*

$$\sqrt{\rho_{\min}} \leq \frac{\|v\|_{L_\rho^2(\Gamma)}}{\|v\|_{L^2(\Gamma)}} \leq \sqrt{\rho_{\max}}, \qquad \forall\, v \in L_\rho^2.$$

Let $\mathbf{p} = (p_1, \ldots, p_N)$ be a multi-index and $\Lambda \subset \mathbb{N}^N$ an index set which is monotonous in the following sense:

4

Property 1 (Monotonicity of Λ). *Consider two multi-indices $\mathbf{p}', \mathbf{p}'' \in \mathbb{N}^N$ such that $p_n'' \le p_n', \forall n = 1, \dots, N$. The multi-index set Λ is monotonous if the following holds:*

$$\mathbf{p}' \in \Lambda \Longrightarrow \mathbf{p}'' \in \Lambda.$$

We denote by $\mathbb{P}_\Lambda(\Gamma)$ the multivariate polynomial space

$$\mathbb{P}_\Lambda(\Gamma) = \mathrm{span}\left\{\prod_{n=1}^N y_n^{p_n}, \text{ with } \mathbf{p} \in \Lambda\right\}, \tag{2.1}$$

and by $\#\Lambda = dim(\mathbb{P}_\Lambda)$ the dimension of the polynomial space, which corresponds to the cardinality of the index set Λ. For convenience, the set Λ can be arranged in lexicographical order:

Property 2 (Lexicographical order). *Given $\mathbf{p}', \mathbf{p}'' \in \Lambda$,*

$$\mathbf{p}' \overset{L}{<} \mathbf{p}'' \iff \exists \widetilde{n} > 0 \quad : \quad (p_{\widetilde{n}}' < p_{\widetilde{n}}'') \wedge (p_n' = p_n'' \quad \forall \, n < \widetilde{n}).$$

According to this order, we can let $\mathbf{p}^j$ denote the jth multi-index of Λ. Sometimes we refer to the elements of Λ with the generic multi-index $\mathbf{p}$, rather than listing them by the lexicographical index.

Since the monomial basis in (2.1) is very ill-conditioned, in practice we use an orthonormal polynomial basis. A typical choice is to take orthogonal polynomials with respect to the measure $\rho(\mathbf{Y})d\mathbf{Y}$.

We introduce an N-dimensional orthonormal polynomial basis $\{l_j\}_{j=1}^{\#\Lambda}$ of $\mathbb{P}_\Lambda$ with respect to the weighted inner product

$$(u,v)_{L_\rho^2} = \int_\Gamma u(\mathbf{Y})v(\mathbf{Y})\rho(\mathbf{Y})\, d\mathbf{Y},$$

i.e. $(l_i, l_j)_{L_\rho^2} = \delta_{ij}$. Assumption 1 guarantees that the orthonormal basis is complete in L_ρ^2 when $\Lambda = \mathbb{N}^N$.

In the case where the density factorizes as $\rho(\mathbf{Y}) = \prod_{n=1}^N \rho_n(Y_n)$ the basis can be constructed by tensorizing 1D orthogonal polynomials with respect to each weight ρ_n separately. Given n, let $\{\varphi_j^n(\cdot)\}_j$ be the orthogonal polynomials with respect to ρ_n. Picking the jth (according to Property 2) multi-index $\mathbf{p}^j \in \Lambda$, the corresponding jth multidimensional basis function is given by

$$l_j(\mathbf{Y}) = \prod_{n=1}^N \varphi_{p_n^j}^n(Y_n). \tag{2.2}$$

Thus, using the basis function provided by (2.2), the definition (2.1) of $\mathbb{P}_\Lambda$ becomes

$$\mathbb{P}_\Lambda(\Gamma) = \mathrm{span}\{l_j, \, j = 1, \dots, \#\Lambda\}. \tag{2.3}$$

Observe that in general (2.1) and (2.3) are equivalent only if the index set Λ satisfies the Monotonicity Property 1.

To construct the polynomial approximation ϕ_Λ we first sample the exact function ϕ in $M > \#\Lambda$ independent points $\mathbf{y}_1, \ldots, \mathbf{y}_M$, i.e. $\mathbf{y}_i$ are independent identically distributed (i.i.d.) random vectors with probability density function ρ. Then, we compute a discrete least square approximation of the values $\phi(\mathbf{y}_i)$ in the polynomial space $\mathbb{P}_\Lambda$, i.e.,

$$\phi_\Lambda = \Pi_M^{\Lambda,\omega}\phi = \operatorname*{argmin}_{v \in \mathbb{P}_\Lambda(\Gamma)} \frac{1}{M} \sum_{i=1}^M (\phi(\mathbf{y}_i) - v(\mathbf{y}_i))^2. \tag{2.4}$$

We will use the superscript (or subscript) ω to denote a quantity that depends on the random sample $\mathbf{y}_1, \ldots, \mathbf{y}_M$ (and therefore is random itself).

We now introduce the *random* discrete inner product

$$(u,v)_{M,\omega} = \frac{1}{M} \sum_{i=1}^M u(\mathbf{y}_i)v(\mathbf{y}_i), \tag{2.5}$$

and the corresponding norm $\|u\|_{M,\omega} = (u,u)_{M,\omega}^{1/2}$. Observe that this is actually a norm in $\mathbb{P}_\Lambda$ if for all $v \in \mathbb{P}_\Lambda$,

$$v(\mathbf{y}_i) = 0, \text{ for } i = 1, \ldots, M \qquad \Leftrightarrow \qquad v = 0, \tag{2.6}$$

which by Assumption 1 happens with probability one provided that M is large enough. Then, the discrete least square problem can be written equivalently as

$$\text{find } \Pi_M^{\Lambda,\omega}\phi \in \mathbb{P}_\Lambda(\Gamma) \quad \text{s.t.} \quad (\Pi_M^{\Lambda,\omega}\phi, v)_{M,\omega} = (\phi, v)_{M,\omega}, \quad \forall v \in \mathbb{P}_\Lambda(\Gamma).$$

2.1 Common multivariate polynomial spaces

Some of the most common choices of function spaces are, Tensor Product, Total Degree, and Hyperbolic Cross, which are defined by the index sets below. We index the set Λ by the subscript w denoting the maximum polynomial degree used:

$$\textbf{Tensor Product (TP)}, \quad \Lambda_{\text{w}} = \left\{ \mathbf{p} \in \mathbb{N}^N : \max_{n=1,\ldots,N} p_n \leq \text{w} \right\}, \tag{2.7}$$

$$\textbf{Total Degree (TD)}, \quad \Lambda_{\text{w}} = \left\{ \mathbf{p} \in \mathbb{N}^N : \sum_{n=1}^N p_n \leq \text{w} \right\}, \tag{2.8}$$

$$\textbf{Hyperbolic Cross (HC)}, \quad \Lambda_{\text{w}} = \left\{ \mathbf{p} \in \mathbb{N}^N : \prod_{n=1}^N (p_n + 1) \leq \text{w} + 1 \right\}. \tag{2.9}$$

These spaces are isotropic and the maximum polynomial degree w is the same in all spatial dimensions. Anisotropic versions could be considered as well [3].

The dimensions of TP and TD spaces are

$$\#TP(\mathrm{w}, N) = (\mathrm{w} + 1)^N, \tag{2.10}$$

$$\#TD(\mathrm{w}, N) = \binom{N + \mathrm{w}}{N}. \tag{2.11}$$

The dimension of the HC space is harder to quantify, so we report its exact dimension $\#HC(\mathrm{w}, N)$ in Fig. 1, computed for some values of w and N. A sharp upper bound, that holds when $N = 2$, is given by

$$\#HC(\mathrm{w}, 2) \leq \left\lfloor (\mathrm{w} + 1) \cdot (1 + \log(\mathrm{w} + 1)) \right\rfloor. \tag{2.12}$$
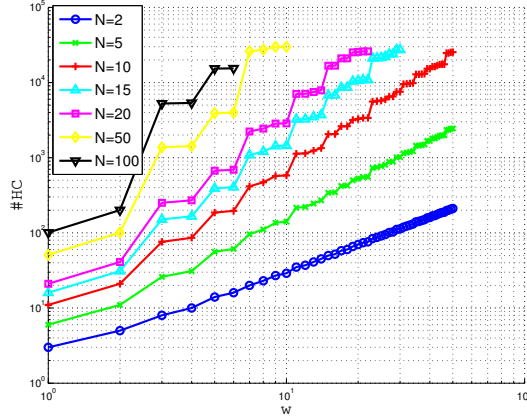
Fig. 1: Dimension of the HC space, $N = 2, 5, 10, 15, 20, 50, 100$.

2.2 L_ρ^2 optimality of the random discrete L^2 projection

Let us first introduce the following quantity

$$C^\omega(M, \Lambda) := \sup_{v \in \mathbb{P}_\Lambda \setminus \{v \equiv 0\}} \frac{\|v\|_{L_\rho^2}^2}{\|v\|_{M,\omega}^2}. \tag{2.13}$$

Notice that this quantity depends on the random sample and is therefore a random variable. The following result holds:

Proposition 1. *With $C^\omega(M, \Lambda)$ defined as in* (2.13)*, it holds that*

$$\|\phi - \Pi_M^{\Lambda,\omega} \phi\|_{L_\rho^2} \leq \left(1 + \sqrt{C^\omega(M, \Lambda)} \right) \inf_{v \in \mathbb{P}_\Lambda(\Gamma)} \|\phi - v\|_{L^\infty}. \tag{2.14}$$

7

Proof. Let v be an arbitrary polynomial in $\mathbb{P}_\Lambda$. Then

$$
\begin{aligned}
\|\phi - \Pi_M^{\Lambda,\omega}\phi\|_{L_\rho^2} &\leq \|\phi - v\|_{L_\rho^2} + \|v - \Pi_M^{\Lambda,\omega}\phi\|_{L_\rho^2} \\
&= \|\phi - v\|_{L_\rho^2} + \frac{\|v - \Pi_M^{\Lambda,\omega}\phi\|_{L_\rho^2}}{\|v - \Pi_M^{\Lambda,\omega}\phi\|_{M,\omega}}\|v - \Pi_M^{\Lambda,\omega}\phi\|_{M,\omega} \\
&\leq \|\phi - v\|_{L_\rho^2} + \sup_{v \in \mathbb{P}_\Lambda \backslash \{v \equiv 0\}} \frac{\|v\|_{L_\rho^2}}{\|v\|_{M,\omega}}\|v - \Pi_M^{\Lambda,\omega}\phi\|_{M,\omega} \\
&= \|\phi - v\|_{L_\rho^2} + \sqrt{C^\omega(M,\Lambda)}\left(\|\phi - v\|_{M,\omega}^2 - \|\phi - \Pi_M^{\Lambda,\omega}\phi\|_{M,\omega}^2\right)^{\frac{1}{2}} \\
&\leq \left(1 + \sqrt{C^\omega(M,\Lambda)}\right)\|\phi - v\|_{L^\infty}.
\end{aligned}
$$

The thesis follows from the arbitrariness of v. $\qquad\square$

Clearly, the convergence property of the random discrete projection is strictly related to the properties of the quantity $C^\omega(M,\Lambda)$. In Section 3 we will characterize this quantity in the particular case of a one dimensional problem ($N = 1$) and the uniform distribution. The extension to multivariate functions is an ongoing work.

2.3   Asymptotic limit of $C^\omega(M,\Lambda)$

Here we show that when $M \to \infty$, while Λ is kept fixed, the limit of $C^\omega(M,\Lambda)$ is almost surely 1. For the purpose of this analysis we also introduce the constant

$$
\widetilde{C}_\Lambda = \sup_{\varphi \in \mathbb{P}_\Lambda} \frac{\|\varphi\|_{L^\infty}^2}{\|\varphi\|_{L_\rho^2}^2} < +\infty. \tag{2.15}
$$

Remark 1. *Given any Λ, the constant $\widetilde{C}_\Lambda$ is always finite in any dimension N since the space $\mathbb{P}_\Lambda$ is finite dimensional and the domain is bounded. For example in 1D, any $v \in \mathbb{P}_{\mathrm{w}}(-1,1)$ can be expanded in Legendre series $v = \sum_{j=0}^{\mathrm{w}} \alpha_j \varphi_j$ with $\|v\|_{L^2}^2 = \sum_{j=0}^{\mathrm{w}} \alpha_j^2$. Moreover, assuming a uniform distribution ρ on the interval $[-1,1]$,*

$$
\begin{aligned}
\|v\|_{L^\infty} &\leq \sum_{j=0}^{\mathrm{w}} |\alpha_j| \|\varphi_j\|_{L^\infty} \leq \sum_{j=0}^{\mathrm{w}} |\alpha_j| \sqrt{j + \frac{1}{2}} \\
&\leq \sqrt{\sum_{j=0}^{\mathrm{w}} \alpha_j^2}\sqrt{\sum_{j=0}^{\mathrm{w}}\left(j + \frac{1}{2}\right)} = \frac{(\mathrm{w}+1)}{\sqrt{2}}\|v\|_{L^2} = (\mathrm{w}+1)\|v\|_{L_\rho^2},
\end{aligned}
$$

so $\widetilde{C}_\Lambda \leq \mathrm{w}+1$. One can tensorize the 1D case, to show that $\widetilde{C}_\Lambda$ is bounded with respect to the maximum polynomial degree w also in higher dimensions.

We recall that a sequence of functions $\{f_n\}$, where $f_n : S \to \mathbb{R}$ converges uniformly to f if and only if

$$\lim_{n \to \infty} \sup_{s \in S} \left| f_n(s) - f(s) \right| = 0.$$

We denote the uniform convergence by $f_n \overset{\text{unif.}}{\longrightarrow} f$. We also recall the definition of an equicontinuous sequence of functions:

Definition 1. *Consider the usual ε-δ definition of continuity for a real function f in the point x_0:*

$$\forall \varepsilon > 0 \quad \exists \delta = \delta(x_0, f) > 0 \quad \forall x \in dom(f) : |x - x_0| < \delta \Rightarrow |f(x) - f(x_0)| < \varepsilon. \tag{2.16}$$

We say that the family of continuous functions $F = \{f_i\}_i$ is

- *equicontinuous if δ in (2.16) is independent of f_i,*

- *uniformly equicontinuous if δ in (2.16) is independent of f_i and x_0.*

Now consider the set of functions $\widehat{\mathbb{P}}_\Lambda = \{v \in \mathbb{P}_\Lambda : \|v\|_{L^2_\rho} = 1\} \subset L^2_\rho$. For any outcome $\{\mathbf{y}_j\}_{j=1}^\infty$ of the random variables $\{\mathbf{Y}_j\}_{j=1}^\infty$ we define the sequence $\{\tau_{M,\omega}\}_M$ of functionals, whose elements are defined as

$$\tau_{M,\omega}(\cdot) = \frac{\| \cdot \|_{M,\omega}^2}{\| \cdot \|_{L^2_\rho}^2} : \widehat{\mathbb{P}}_\Lambda \to \mathbb{R}^+. \tag{2.17}$$

Proposition 2. *For any M and Λ the function $\tau_{M,\omega}$ in (2.17) is Lipschitz continuous over $\widehat{\mathbb{P}}_\Lambda$.*

Proof. Consider the constant $\widetilde{C}_\Lambda$ defined in (2.15). Clearly $\|v\|_{L^\infty} \leq \widetilde{C}_\Lambda$ for all $v \in \widehat{\mathbb{P}}_\Lambda$, so they are uniformly bounded. Taking arbitrary v_1 and v_2 in $\widehat{\mathbb{P}}_\Lambda$,

$$
\begin{aligned}
\left| \tau_{M,\omega}(v_1) - \tau_{M,\omega}(v_2) \right| &\leq \frac{1}{M} \sum_{j=1}^M \left| v_1(y_j)^2 - v_2(y_j)^2 \right| \\
&= \frac{1}{M} \sum_{j=1}^M \left| \left(v_1(y_j) + v_2(y_j) \right) \left(v_1(y_j) - v_2(y_j) \right) \right| \\
&\leq \frac{1}{M} \left(\|v_1\|_{L^\infty} + \|v_2\|_{L^\infty} \right) \sum_{j=1}^M \left| v_1(y_j) - v_2(y_j) \right| \\
&\leq 2\,\widetilde{C}_\Lambda \|v_1 - v_2\|_{L^\infty} \leq 2\,\widetilde{C}_\Lambda^2 \|v_1 - v_2\|_{L^2_\rho}. \qquad (2.18)
\end{aligned}
$$

$\square$

From Proposition 2 it follows immediately that the sequence $\{\tau_{M,\omega}\}$ is uniformly equicontinuous. Moreover, from the Strong Law of Large Numbers, since $\mathbb{E}[v(y_j)^2] < \infty$, it follows that

$$\|v\|_{M,\omega}^2 = \frac{1}{M}\sum_{j=1}^{M} v(x_j)^2 \xrightarrow{M\to+\infty} \mathbb{E}\Big[(v(y_j))^2\Big] = \|v\|_{L_\rho^2}^2,$$

almost surely; hence the sequence $\{\tau_{M,\omega}\}$ is also converging to 1 pointwise. We have then the following result:

Proposition 3. *Under the assumptions of Proposition 2 it holds*

$$\tau_{M,\omega}(v) \xrightarrow{unif.} 1, \quad \forall\, v \in \widehat{\mathbb{P}}_\Lambda, \quad a.s. \tag{2.19}$$

Proof. For any outcome ω, the sequence $\{\tau_{M,\omega}\}_{M=1}^{\infty}$ is (uniformly) equicontinuous on $\widehat{\mathbb{P}}_\Lambda$ and converges almost surely on $\widehat{\mathbb{P}}_\Lambda$, therefore it converges uniformly in $\widehat{\mathbb{P}}_\Lambda$ (see *e.g.* [27, Theorem 7.25]). $\square$

Theorem 1. *Let $C^\omega(M,\Lambda)$ be the constant defined in (2.13). Then*

$$\lim_{M\to\infty} C^\omega(M,\Lambda) = 1, \quad a.s.$$

Proof. We can extend the definition of $\tau_{M,\omega}$ in (2.17) to $\mathbb{P}_\Lambda \setminus \{v \equiv 0\}$ using the same formula. Then we rewrite the following supremum as

$$\lim_{M\to+\infty} \sup_{v\in\widehat{\mathbb{P}}_\Lambda} \left|\frac{\|v\|_{M,\omega}^2}{\|v\|_{L_\rho^2}^2} - 1\right| = \lim_{M\to+\infty} \sup_{v\in\mathbb{P}_\Lambda\setminus\{v\equiv 0\}} \left|\frac{\|v\|_{M,\omega}^2}{\|v\|_{L_\rho^2}^2} - 1\right|,$$

and, by the continuity of $\tau_{M,\omega}(v)$ in $\mathbb{P}_\Lambda$,

$$\tau_{M,\omega}(v) \xrightarrow{unif.} 1, \quad \forall\, v \in \mathbb{P}_\Lambda \setminus \{v \equiv 0\}, \tag{2.20}$$

from which we also deduce that

$$\Big(\tau_{M,\omega}(v)\Big)^{-1} \xrightarrow{unif.} 1, \quad \forall\, v \in \mathbb{P}_\Lambda \setminus \{v \equiv 0\}.$$

But this implies

$$\lim_{M\to+\infty} \sup_{v\in\mathbb{P}_\Lambda\setminus\{v\equiv 0\}} \left|\frac{\|v\|_{L_\rho^2}^2}{\|v\|_{M,\omega}^2} - 1\right| = 0,$$

which is the thesis. $\square$

3 Error Analysis in 1D for the Uniform Distribution

We restrict the analysis to the case $N = 1$ and consider the polynomial space $\mathbb{P}_\Lambda = \text{span}\{y^p, \ p = 0, \ldots, w\}$. For convenience, we rename the polynomial space as $\mathbb{P}_w$ and the random discrete projector as $\Pi_M^{w,\omega}$. The main result of this section is Theorem 2. Its proof is postponed until the end of the section as we will need several lemmas and intermediate results.

Theorem 2. *For any $\alpha \in (0, 1)$, under the condition*

$$\frac{M}{3 \log((M+1)/\alpha)} \geq 4\sqrt{3}\,w^2 \tag{3.1}$$

holds

$$\mathbb{P}\left(\|\phi - \Pi_M^{w,\omega}\phi\|_{L_\rho^2} \leq \left(1 + \sqrt{3\log\frac{M+1}{\alpha}}\right) \inf_{v \in \mathbb{P}_w} \|\phi - v\|_{L^\infty}\right) \geq 1 - \alpha. \tag{3.2}$$

The previous theorem states that, with confidence level $1 - \alpha$, the discrete projection with random points is (nearly) optimal up to a logarithmic factor in M, provided M is large enough and satisfies the condition (3.1).

Now we proceed to derive some probabilistic results on the distribution of order statistics for the standard uniform distribution.

3.1 Useful results on order statistics of the uniform distribution

To study order statistics of the uniform distribution it is more convenient to consider standard uniform random variables in $[0, 1]$ instead of $[-1, 1]$. Therefore, we introduce a linear transformation $\mathcal{T} : [-1, 1] \to [0, 1]$ and define a new set of i.i.d. random variables

$$U_i = \mathcal{T}(Y_i) = \frac{Y_i + 1}{2}, \qquad i = 1, \ldots, M.$$

Thus we continue working with a sample of M independent random variables from the standard uniform parent distribution,

$$U_1, \ldots, U_M \overset{\text{i.i.d.}}{\sim} \mathbb{U}(0, 1). \tag{3.3}$$

We know that the order statistics $U_{(1)} < U_{(2)} < \ldots < U_{(M)}$ are

$$U_{(k)} \sim \text{Beta}(k, M + 1 - k), \qquad k = 1, \ldots, M,$$

where we recall that a $\text{Beta}(k, M + 1 - k)$ random variable has distribution

$$f(y) = \frac{M!}{(k-1)!(M-k)!}y^{k-1}y^{M-k}, \qquad y \in [0, 1].$$

Let us define $\Xi_{(0)} = U_{(1)}$, $\Xi_{(M)} = 1 - U_{(M)}$, and $\Xi_{(k)} = U_{(k+1)} - U_{(k)}$ for $k = 1, \ldots, M-1$. It can be shown that $\Xi_{(k)} \sim \text{Beta}(1, M)$ for each $k = 0, \ldots, M$; see [16, page 14] where a more general result on the distribution of $U_j - U_i$ is proved, namely

$$U_j - U_i \sim \text{Beta}(j - i, M - j + i + 1), \quad 1 \le i < j \le M.$$

In particular, the distribution is independent of the values i and j, and depends only on their difference $j - i$. The following result gives a bound on the probability distribution of $\max_{k=0,\ldots,M} \Xi_{(k)}$.

Lemma 1. *For any $\alpha \in (0,1)$ and $M \in \mathbb{N}_+$*

$$\mathbb{P}\left(\max_{k=0,\ldots,M} \Xi_{(k)} > \frac{\log((M+1)/\alpha)}{M} \right) \le \alpha.$$

Proof. Trivially, if $\quad 0 < \alpha \le \dfrac{M+1}{\exp(M)}, \quad$ then $\quad \dfrac{\log\big((M+1)/\alpha\big)}{M} \ge 1 \quad$ and

$$\mathbb{P}\left(\max_{k=0,\ldots,M} \Xi_{(k)} > 1 \right) = 0 < \alpha.$$

Consider now $\quad \dfrac{M+1}{\exp(M)} < \alpha < 1$. Rewriting the random event

$$\left\{ \max_{k=0,\ldots,M} \Xi_{(k)} > \delta \right\} = \bigcup_{k=0}^{M} \left\{ \Xi_{(k)} > \delta \right\},$$

we have, for $0 < \delta < 1$,

$$\mathbb{P}\left(\max_{k=0,\ldots,M} \Xi_{(k)} > \delta \right) = \mathbb{P}\left(\bigcup_{k=0}^{M} \left\{ \Xi_{(k)} > \delta \right\} \right)$$

$$\le \sum_{k=0}^{M} \mathbb{P}\left(\Xi_{(k)} > \delta \right) = (M+1)(1-\delta)^M. \qquad (3.4)$$

Therefore

$$\mathbb{P}\left(\max_{k=0,\ldots,M} \Xi_{(k)} > \frac{\log((M+1)/\alpha)}{M} \right) \le (M+1)\left(1 - \frac{\log((M+1)/\alpha)}{M} \right)^M$$

$$= (M+1)\left(\left(1 - \frac{\log((M+1)/\alpha)}{M} \right)^{\frac{M}{\log((M+1)/\alpha)}} \right)^{\log((M+1)/\alpha)}$$

$$\le (M+1)\, e^{-\log((M+1)/\alpha)} = \alpha.$$

$\square$

3.2 Relation between $\|\cdot\|_{L_\rho^2}$ and $\|\cdot\|_{M,\omega}$ on $\mathbb{P}_w([-1,1])$

Here we go back to the uniform distribution in $\Gamma = [-1,1]$. We are interested in finding an inequality between the continuous norm $\|\cdot\|_{L_\rho^2}$ and the discrete norm $\|\cdot\|_{M,\omega}$, *i.e.* finding a suitable $C_M^{\mathrm{w},\omega}$ such that

$$\|v\|_{L_\rho^2}^2 \leq C_M^{\mathrm{w},\omega}\|v\|_{M,\omega}^2, \qquad \forall\, v \in \mathbb{P}_\mathrm{w}(\Gamma). \tag{3.5}$$

This will be the random constant appearing in Proposition 1. We will also need an inverse inequality for the derivative (see e.g. [11]),

$$\left\|\frac{\partial v}{\partial y}\right\|_{L^2(\Gamma)} \leq \widehat{C}\,\mathrm{w}^2\|v\|_{L^2(\Gamma)}, \qquad \forall\, v \in \mathbb{P}_\mathrm{w}(\Gamma). \tag{3.6}$$

In this case $\widehat{C} = \sqrt{3}$, while in general $\widehat{C}(\Gamma)$ depends on the length of Γ. The same inequality holds replacing $L^2(\Gamma)$ with L_ρ^2.

The sampled points $\{y_j\}_{j=1}^M$ are distinct almost surely. To each point y_j we associate an open interval I_j satisfying the requirement that

$$\left(\bigcup_{j=1}^M \overline{I_j}\right) \cap \Gamma = \Gamma. \tag{3.7}$$

In other words, the family of intervals $\{I_j\}_j$ is a (finite) covering of the domain Γ. We order the points in increasing order

$$-1 \leq y_1 < \ldots < y_M \leq 1,$$

keeping the notation $\{y_j\}$; it will be clear from the context if we refer to the ordered points or to the originally sampled points.

In one dimension, it is easy to build a finite covering of mutually disjoint intervals

$$I_j = \left(y_j - \Delta y_{j-1}, y_j + \Delta y_j\right), \qquad j = 1, \ldots, M, \tag{3.8}$$

taking

$$\Delta y_j = \begin{cases} |-1-y_1|, & j = 0, \\ \dfrac{|y_{j+1}-y_j|}{2}, & j = 1, \ldots, M-1, \\ |1-y_M|, & j = M. \end{cases} \tag{3.9}$$

In general, the sets I_j of this covering are not centered in the sample points. It is useful to split each interval I_j in its left and right part,

$$I_j^- = I_j \cap \{y \in \mathbb{R} : y \leq y_j\}, \qquad I_j^+ = I_j \cap \{y \in \mathbb{R} : y \geq y_j\},$$

13

with measures $|I_j^-| = \Delta y_{j-1}$ and $|I_j^+| = \Delta y_j$, respectively.

We also define the random variable ΔY whose realizations are

$$\Delta y = \max_{j=1,\ldots,M} |I_j|. \tag{3.10}$$

The link between the random variable ΔY and the result given in the previous section is the following:

Lemma 2 (Corollary of Lemma 1)**.** *For any $\alpha \in (0,1)$ and $M \in \mathbb{N}_+$*

$$\mathbb{P}\left(\Delta Y > \frac{3\log((M+1)/\alpha)}{M}\right) \leq \alpha.$$

Proof. It is enough to notice that, for each realization $(\xi_0, \ldots, \xi_M)$ of the random variables $(\Xi_{(0)}, \ldots, \Xi_{(M)})$, it holds

$$\max_{k=0,\ldots,M} \xi_k < \delta \quad \Longrightarrow \quad \Delta y < 3\delta,$$

and thus

$$\mathbb{P}(\Delta Y > 3\delta) \leq \mathbb{P}\left(\max_{k=0,\ldots,M} \Xi_{(k)} > \delta\right).$$

$\square$

We now define two events:

$$\Sigma_\mathrm{w} = \left\{\Delta Y \leq \frac{1}{4\widehat{C}\mathrm{w}^2}\right\} \quad \text{and} \quad \Sigma_M = \left\{\Delta Y \leq \frac{3\log((M+1)/\alpha)}{M}\right\}. \tag{3.11}$$

Observe that, taking M large enough, the probability of Σ_w can be made arbitrary close to 1. More precisely

Lemma 3. *For any $\alpha \in (0,1)$, under the condition*

$$\frac{M}{3\log\left((M+1)/\alpha\right)} \geq 4\,\widehat{C}\,\mathrm{w}^2, \tag{3.12}$$

the following inclusion holds

$$\Sigma_M \subset \Sigma_\mathrm{w}$$

and

$$\mathbb{P}(\Sigma_\mathrm{w}) \geq \mathbb{P}(\Sigma_M) \geq 1 - \alpha. \tag{3.13}$$

Proof. Clearly, under (3.12)

$$\Delta Y \leq \frac{3\log((M+1)/\alpha)}{M} \quad \Longrightarrow \quad \Delta Y \leq \frac{1}{4\widehat{C}\mathrm{w}^2}$$

and the inclusion $\Sigma_M \subset \Sigma_\mathrm{w}$ is proved. The bound from below on the corresponding probabilities is an immediate consequence of Lemma 2. $\square$

Now we are ready to formulate the main result of this subsection:

Theorem 3. *Define on Σ_{w} the random variable*

$$C_M^{\mathrm{w},\omega} = \frac{M\Delta Y/2}{1 - 2\Delta Y\,\widehat{C}\,\mathrm{w}^2}. \tag{3.14}$$

Then, in Σ_{w} it holds

$$\|v\|_{L_\rho^2}^2 \le C_M^{\mathrm{w},\omega}\|v\|_{M,\omega}^2, \qquad \forall\, v \in \mathbb{P}_{\mathrm{w}}(\Gamma). \tag{3.15}$$

Moreover, under condition (3.12), in $\Sigma_M \subset \Sigma_{\mathrm{w}}$ it holds

$$C_M^{\mathrm{w},\omega} \le 3\log((M+1)/\alpha).$$

Proof. For convenience the proof uses the standard $L^2(\Gamma)$ norm instead of the weighted norm L_ρ^2. Remember that in this case $\|\cdot\|_{L_\rho^2}^2 = \frac{1}{2}\|\cdot\|_{L^2(\Gamma)}^2$. To lighten the notation we also introduce on each interval $I_j^{\pm}$ the short notation

$$\|\cdot\|_{I_j^-} := \|\cdot\|_{L^2(I_j^-)},$$
$$\|\cdot\|_{I_j^+} := \|\cdot\|_{L^2(I_j^+)}.$$

Take $v \in \mathbb{P}_{\mathrm{w}}(\Gamma)$ arbitrarily. The following standard inequalities are used in the sequel:

$$\left| v(y)^2 - v(y_j)^2 \right| = \left| \int_{y_j}^{y} (v^2)'(\xi)d\xi \right| = \left| \int_{y_j}^{y} 2v(\xi)v(\xi)'d\xi \right|$$
$$\le \begin{cases} 2\|v\|_{I_j^+}\|v'\|_{I_j^+}, & \forall y \in I_j^+, \\ 2\|v\|_{I_j^-}\|v'\|_{I_j^-}, & \forall y \in I_j^-. \end{cases} \tag{3.16}$$

Integrating (3.16) over I_j^- and I_j^+ yields

$$\int_{I_j^-} v(y)^2 dy - v(y_j)^2|I_j^-| = \int_{I_j^-} \left(v(y)^2 - v(y_j)^2 \right) dy$$
$$\le \int_{I_j^-} \left| v(y)^2 - v(y_j)^2 \right| dy$$
$$\le 2\,|I_j^-|\,\|v\|_{I_j^-}\|v'\|_{I_j^-}, \tag{3.17}$$

and similarly for I_j^+

$$\int_{I_j^+} v(y)^2 dy - v(y_j)^2|I_j^+| \le 2\,|I_j^+|\,\|v\|_{I_j^+}\|v'\|_{I_j^+}. \tag{3.18}$$

15

Summing (3.17) and (3.18) we get

$$\int_{I_j} v(y)^2 dy - |I_j| v(y_j)^2 \leq 2\Big(|I_j^-| \, \|v\|_{I_j^-} \|v'\|_{I_j^-} + |I_j^+| \, \|v\|_{I_j^+} \|v'\|_{I_j^+} \Big)$$
$$\leq 2 \, |I_j| \, \|v\|_{I_j} \|v'\|_{I_j}, \tag{3.19}$$

which implies

$$\int_{I_j} v(y)^2 dy \leq |I_j| \Big(v(y_j)^2 + 2 \, \|v\|_{I_j} \|v'\|_{I_j} \Big). \tag{3.20}$$

Summing over all the intervals we have

$$\|v\|_{L^2}^2 = \sum_{j=1}^{M} \int_{I_j} v(y)^2 dy \qquad \Big[\text{substitute (3.20)} \Big]$$

$$\leq \sum_{j=1}^{M} |I_j| v(y_j)^2 + 2 \sum_{j=1}^{M} |I_j| \, \|v\|_{I_j} \|v'\|_{I_j} \qquad \Big[\text{using (3.10)} \Big]$$

$$\leq \Delta y \sum_{j=1}^{M} v(y_j)^2 + 2\Delta y \sum_{j=1}^{M} \|v\|_{I_j} \|v'\|_{I_j} \qquad \Big[\text{definition of } \| \cdot \|_{M,\omega} \Big]$$

$$= M\Delta y \|v\|_{M,\omega}^2 + 2\Delta y \sum_{j=1}^{M} \|v\|_{I_j} \|v'\|_{I_j} \qquad \Big[\text{Cauchy-Schwarz's ineq.} \Big]$$

$$\leq M\Delta y \|v\|_{M,\omega}^2 + 2\Delta y \, \|v\|_{L^2(\Gamma)} \|v'\|_{L^2(\Gamma)} \qquad \Big[\text{using (3.6)} \Big]$$

$$\leq M\Delta y \|v\|_{M,\omega}^2 + 2\Delta y \, \widehat{C} \, \mathrm{w}^2 \, \|v\|_{L^2(\Gamma)}^2. \tag{3.21}$$

Rearranging the terms in (3.21) we obtain

$$(1 - 2\Delta y \, \widehat{C} \, \mathrm{w}^2) \|v\|_{L^2(\Gamma)}^2 \leq M\Delta y \|v\|_{M,\omega}^2. \tag{3.22}$$

The coefficient in the left hand side is non zero on Σ_{w} so we have proved (3.14)-(3.15). If we now restrict to Σ_M under condition (3.12), then from Lemma 3 we have

$$\Delta Y \leq \frac{1}{4\widehat{C}\mathrm{w}^2} \quad \text{and} \quad \Delta Y \leq \frac{3\log((M+1)/\alpha)}{M}$$

so that

$$C_M^{\mathrm{w},\omega} = \frac{M\Delta Y/2}{1 - 2\Delta Y \, \widehat{C} \, \mathrm{w}^2} \leq 3\log((M+1)/\alpha)$$

and this concludes the proof. $\qquad \square$

Remark 2. *Theorem 3 states that on Σ_M, which is an event of probability at least $1-\alpha$, the discrete and continuous norms are equivalent up to a logarithmic factor if condition (3.12) is fulfilled, which, roughly speaking corresponds to $M \propto \mathrm{w}^2$.*

3.3 Proof of Theorem 2

The proof of this theorem is merely a collection of results from Proposition 1, Theorem 3, and Lemma 2.

Proof of Theorem 2. We consider the event Σ_M defined in (3.11). From Lemma 2 we know that under condition (3.1) the probability of this event is at least $1 - \alpha$. From Theorem 3 it holds

$$\|v\|^2_{L^2_\rho} \leq C^{\mathrm{w},\omega}_M \|v\|^2_{M,\omega} \quad \text{in } \Sigma_M$$

uniformly with respect to $v \in \mathbb{P}_{\mathrm{w}}(\Gamma)$ and the constant $C^{\mathrm{w},\omega}_M \leq 3 \log((M+1)/\alpha)$ *uniformly* in Σ_M. We then apply Proposition 1 in Σ_M to conclude that

$$\|\phi - \Pi^{\mathrm{w},\omega}_M \phi\|_{L^2_\rho} \leq \left(1 + \sqrt{C^{\mathrm{w},\omega}_M}\right) \inf_{v \in \mathbb{P}_{\mathrm{w}}} \|\phi - v\|_{L^\infty}, \quad \forall \phi \in L^\infty(\Gamma), \quad \forall \omega \in \Sigma_M.$$

Since under condition (3.1) the probability of Σ_M is at least $1 - \alpha$, this concludes the proof. $\qquad \square$

4 Algebraic Formulation

The value of $\#\Lambda$ depends on the particular polynomial space (TP, TD, HC, ...), the maximal polynomial degree used, w, and the dimension of the physical space, N. The number of points M must satisfy the constraint

$$M \geq \#\Lambda,$$

to have an overdetermined problem (more data than unknowns). We have shown in Section 2 that, for monovariate functions, M should scale as $M \propto \mathrm{w}^2$ to have a stable discrete projection. As a general rule to choose M for multivariate functions and arbitrary polynomial spaces we consider the formula

$$M = c \, (\#\Lambda)^\alpha, \tag{4.1}$$

where c is a positive constant, and $\alpha \geq 1$ a real number. We restrict our numerical tests in Section 5 to $1 \leq \alpha \leq 2$.

Given the polynomial space, we define the design matrix $D^\omega \in \mathbb{R}^{M \times \#\Lambda}$. The element $D^\omega_{i,j}$ contains the jth L^2_ρ-orthonormal basis function l_j evaluated in the ith sample point $\mathbf{y}_i$, that is

$$[D^\omega]_{i,j} = l_j(\mathbf{y}_i). \tag{4.2}$$

The discrete random projection $\Pi^{\Lambda,\omega}_M \phi$ can be expressed in terms of the orthonormal basis $\{l_j\}_j$ as

$$\Pi^{\Lambda,\omega}_M \phi(\mathbf{Y}) = \sum_{j=1}^{\#\Lambda} x^\omega_j l_j(\mathbf{Y}). \tag{4.3}$$

Then the algebraic problem to determine the unknown coefficients $\{x_j^\omega\}$ can be formulated as:

$$x^\omega = \underset{x \in \mathbb{R}^{\#\Lambda}}{\operatorname{argmin}} \|D^\omega x - b^\omega\|_2 \qquad (4.4)$$

where $b^\omega \in \mathbb{R}^{M \times 1}$ contains the evaluations of the target function ϕ in the M sample points: $b^\omega(i) = \phi(\mathbf{y}_i)$. The Normal Equations allow us to rewrite the rectangular system embedded in (4.4) as a square system:

$$(D^\omega)^T D^\omega x^\omega = (D^\omega)^T b^\omega. \qquad (4.5)$$

We use problem (4.4) to calculate the approximation error, computing its solution x^ω by QR factorization. On the other hand, formulation (4.5) will be useful to measure the ill-conditioning of the problem, through the condition number of the matrix $(D^\omega)^T D^\omega$. Alternatively, one can also solve (4.5) by Cholesky factorization.

To evaluate the approximation error, we have considered a cross-validation approach: a random set of 100 cross-validating points is chosen at the beginning, and the corresponding design matrix D_{cv} is computed. The evaluations of ϕ in these points are stored in b_{cv}. Then the cross-validated error in ∞-norm is defined as

$$\|\phi - \Pi_M^{\Lambda;\omega}\phi\|_{cv} = \|D_{cv}x^\omega - b_{cv}\|_\infty. \qquad (4.6)$$

Note that $\|\cdot\|_{cv}$ is not a norm on the function space of ϕ; we abuse the norm notation in the figures with cross-validation errors below to emphasize the dependence on ϕ. To estimate the variability of (4.6) due to the random sampling of the M collocation points, we have repeated the calculation over R independent sets of points $\{\mathbf{y}_j^{\omega_k}, \; j = 1, \ldots, M\}$, with $k = 1, \ldots, R$ and we have computed the average error

$$\overline{E}_{cv} = \frac{\sum_{k=1}^{R} \|D_{cv}x^{\omega_k} - b_{cv}\|_\infty}{R}, \qquad (4.7)$$

as well as the sample standard deviation by

$$s_E = \sqrt{\frac{1}{R-1} \sum_{k=1}^{R} \left(\|D_{cv}x^{\omega_k} - b_{cv}\|_\infty - \overline{E}_{cv} \right)^2}. \qquad (4.8)$$

We also aim to analyze the condition number of $(D^\omega)^T D^\omega$,

$$\operatorname{cond}\left((D^\omega)^T D^\omega\right) = \frac{\sigma_{\max}\left((D^\omega)^T D^\omega\right)}{\sigma_{\min}\left((D^\omega)^T D^\omega\right)}, \qquad (4.9)$$

where $\sigma_{\max}(\cdot)$ and $\sigma_{\min}(\cdot)$ are the maximum and minimum singular values. Again, denoting by D^{ω_k} the design matrix built with the k-th set $\{\mathbf{y}_j^{\omega_k}\}_j$, we

estimate the mean condition number $\overline{K}$ over the R repetitions as

$$\overline{K} = \frac{\sum_{k=1}^{R} \text{cond}\left((D^{\omega_k})^T D^{\omega_k} \right)}{R},$$ (4.10)

and the standard deviation as

$$s_K = \sqrt{\frac{1}{R-1} \sum_{k=1}^{R} \left(\text{cond}\left((D^{\omega_k})^T D^{\omega_k} \right) - \overline{K} \right)^2}.$$ (4.11)

4.1 The condition number of $(D^{\omega})^T D^{\omega}$

In the rest of the section we show how the condition number of problem (4.5) relates to some quantities that already appeared previously. All the contents of this section hold for a generic polynomial space, in any dimension. Accordingly, we refer to the polynomial space as $\mathbb{P}_\Lambda$.

In addition to the constant $C^{\omega}(M, \Lambda)$ already introduced in (2.13), we define the constant $c^{\omega}(M, \Lambda)$ as

$$c^{\omega}(M, \Lambda) = \sup_{\varphi \in \mathbb{P}_\Lambda \setminus \{\varphi \equiv 0\}} \frac{\|\varphi\|_{M,\omega}^2}{\|\varphi\|_{L_\rho^2}^2}.$$ (4.12)

Proposition 4. *The spectral condition number (2-norm) of the matrix $(D^{\omega})^T D^{\omega}$, as defined in (4.9) is equal to*

$$K\left((D^{\omega})^T D^{\omega} \right) = c^{\omega}(M, \Lambda)\, C^{\omega}(M, \Lambda).$$ (4.13)

Proof. Each realization of the random matrix $(D^{\omega})^T D^{\omega}$ is almost surely a symmetric and positive definite matrix under Assumption 1 on ρ.

The ith element of the vector $D^{\omega}x$ is

$$[D^{\omega}x^{\omega}]_i = \sum_{j=1}^{\#\Lambda} l_j(y_i) x_j^{\omega} = \Pi_M^{\Lambda,\omega} \phi(y_i).$$ (4.14)

By definition (2.5) of the random discrete inner product

$$\|\Pi_M^{\Lambda,\omega} \phi\|_{M,\omega}^2 = \frac{1}{M} \sum_{i=1}^{M} \left(\Pi_M^{\Lambda,\omega} \phi(y_i) \right)^2$$ (4.15)

and by (4.3) and the L_ρ^2-ortonormality of $\{l_j\}$

$$(x^{\omega})^T x^{\omega} = \sum_{j=1}^{\#\Lambda} (x_j^{\omega})^2 = \|\Pi_M^{\Lambda,\omega} \phi\|_{L_\rho^2}^2.$$ (4.16)

19

Using in sequence (4.16), (4.14) and (4.15), yields

$$\frac{(x^\omega)^T (D^\omega)^T D^\omega x^\omega}{(x^\omega)^T x^\omega} = \frac{\sum_{i=1}^{M} (\Pi_M^{\Lambda,\omega} \phi(y_i))^2}{\|\Pi_M^{\Lambda,\omega} \phi\|_{L_\rho^2}^2} = \frac{M \|\Pi_M^{\Lambda,\omega} \phi\|_{M,\omega}^2}{\|\Pi_M^{\Lambda,\omega} \phi\|_{L_\rho^2}^2}. \qquad (4.17)$$

This shows that (4.17) is the Rayleigh quotient of the matrix $(D^\omega)^T D^\omega$. So the largest and smallest eigenvalues of $(D^\omega)^T D^\omega$ correspond to $c^\omega(M, \Lambda)$ and $C^\omega(M, \Lambda)^{-1}$, respectively. Since $(D^\omega)^T D^\omega$ is symmetric positive definite, its singular values are also eigenvalues. The conclusion follows from the definition (4.9) of the condition number. $\square$

Remark 3. *In 1D and for the uniform distribution we can easily establish a norms equivalence between $\|\cdot\|_{M,\omega}$ and $\|\cdot\|_{L_\rho^2}$, collecting the results of Theorem 3 and Remark 1. Namely, under the condition (3.1), with probability $1 - \alpha$ we have*

$$\frac{1}{1 + \mathrm{w}} \|v\|_{M,\omega}^2 \le \|v\|_{L_\rho^2}^2 \le 3 \log\left(\frac{M+1}{\alpha}\right) \|v\|_{M,\omega}^2,$$

from which we get the bound on the condition number

$$cond\left((D^\omega)^T D^\omega\right) \le \frac{\log((M+1)/\alpha)}{\mathrm{w}+1}, \quad in \ \Sigma_M. \qquad (4.18)$$

where Σ_M is the event defined in (3.11) that has probability at least $1 - \alpha$.

However, we have observed numerically that the bound (4.18) is very pessimistic as under condition (3.1) the condition number seems to be uniformly bounded with respect to M and w.

A direct consequence of Theorem 1 is that

$$\mathrm{cond}\left((D^\omega)^T D^\omega\right) \overset{M \to +\infty}{\longrightarrow} 1, \quad \text{a.s.}$$

This is confirmed numerically. Fig. 2 shows the numerical results obtained for a 1D problem with an overkilling rule $M = 100 \, \#\Lambda^4$, to simulate the asymptotic case $M \uparrow \infty$.

5 Numerical Results

We present an illustrative selection of results from an extensive set of numerical tests. The aim is to seek the correct relation between the number of points to sample M and the dimension of the polynomial space to have a stable and accurate approximation. The following issues are addressed:

- how the condition number (4.9) depends on w, N, c, α and the choice of the polynomial space;
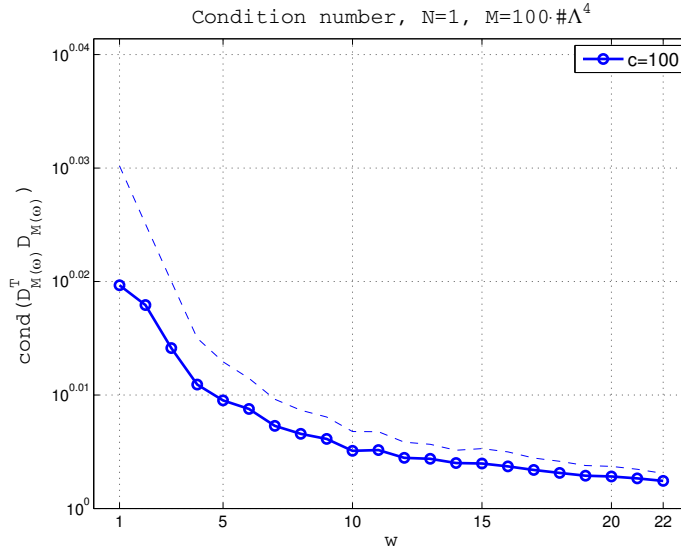
Fig. 2: Condition number (4.9), $N = 1$, $M = 100 \, \# \Lambda^4$. The continuous, marked, lines show the mean condition number (4.10) over 200 repetitions. The dashed lines show the mean (4.10) plus one standard deviation (4.11). The scale on the y-axis ranges from $10^0 = 1$ to $10^{0.04} = 1.0965$

- analogously, how the cross-validation error (4.6) behaves.

In the convergence plots presented in the rest of this section we show the average error (4.7) and condition number (4.10) as well as their average plus one standard deviation.

5.1 One dimensional case

We first investigate the dependence of the condition number (4.9) on the rule (4.1) used to select the number $M(\mathrm{w})$ of sampling points. Observe that in the one dimensional case $\# \Lambda = \mathrm{w} + 1$.

As seen in Fig. 3, the condition number behaves differently depending on the rule chosen. In Fig. 3-Left we report results obtained with the linear rule $M = c \, \# \Lambda$, corresponding to $\alpha = 1$ in (4.1). We tested several values for c ranging from 2 to 20. All cases show an exponential growth of the condition number with respect to w, with rates decreasing with increasing c (as one would expect). Using $R = 10000$ repetitions the observed average condition number still shows a large variability. This is due to the large standard deviations of the condition number, as indicated in the figure in dashed line.

Note that the range of w goes up to 25, so in this range the choice of the largest c yields a linear rule which uses more sample points than some of the quadratic rules (shown in Fig. 3-Right).

In contrast to the exponential growth observed when using the linear rule,

the results using the quadratic rule exhibit a condition number which is approximately constant for w ranging from 1 to 40. Fluctuations become smaller when c increases. This behavior is consistent with the theoretical results in Section 3.
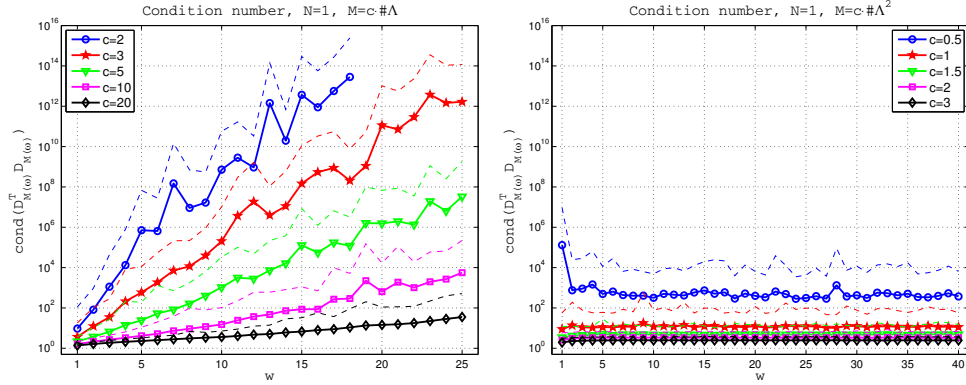
Fig. 3: Condition number (4.9), $N = 1$. The continuous, marked, lines show the mean condition number (4.10) over 10000 repetitions. The dashed lines show the mean (4.10) plus one standard deviation (4.11). Left: $M = c\#\Lambda$. Right: $M = c\#\Lambda^2$.

We now proceed to illustrate the convergence of the error for few functions of varying regularity in Figs. 4–7.

We focus on three target functions: an exponential function

$$\phi(Y) = \exp(Y), \qquad Y \in [-1, 1], \tag{5.1}$$

a meromorphic function

$$\phi(Y) = \frac{1}{1 + \beta Y}, \qquad Y \in [-1, 1], \tag{5.2}$$

that is a function which is analytic provided that $\beta \in (-1, 1)$, and a function with lower regularity

$$\phi(Y) = |Y|^3. \tag{5.3}$$

Fig. 4 shows the error computed as in (4.6), in approximating the exponential function (5.1) with different choices of c and α in the rule (4.1). The quadratic rule (to the right) displays the same exponential, optimal, convergence with respect to w independently of the constant c. The convergence is up to the machine precision.

In contrast, the linear rule (on the left) displays a deterioration of the convergence using small values of c. The deterioration is due to the ill-conditioning of the matrix $D_{M(\omega)}^T D_{M(\omega)}$ when w increases. As noted before, the largest value of c yields at least as many sample points as the quadratic rule with the smallest value of c in the shown range and the errors behave accordingly. Again the fluctuations in the average error decrease with increasing c.
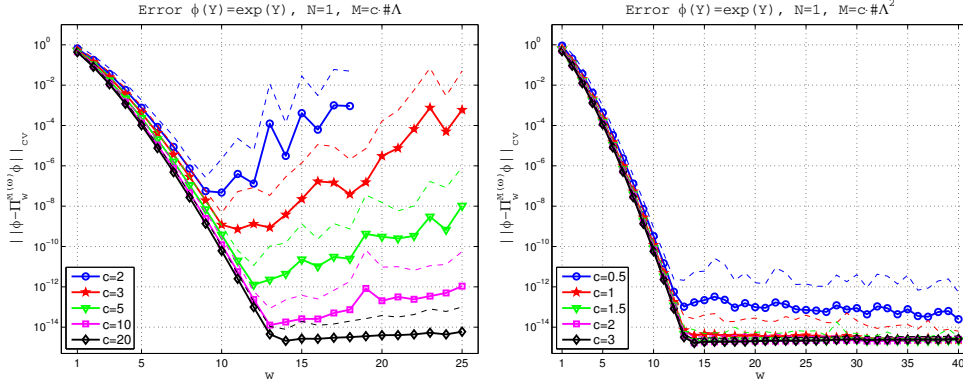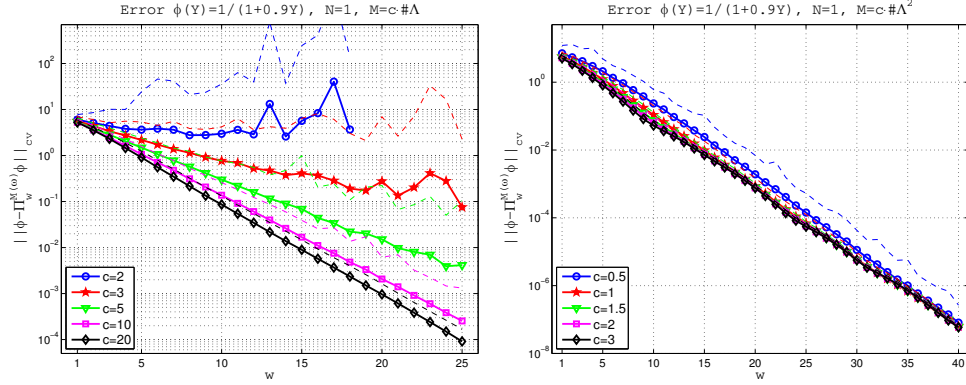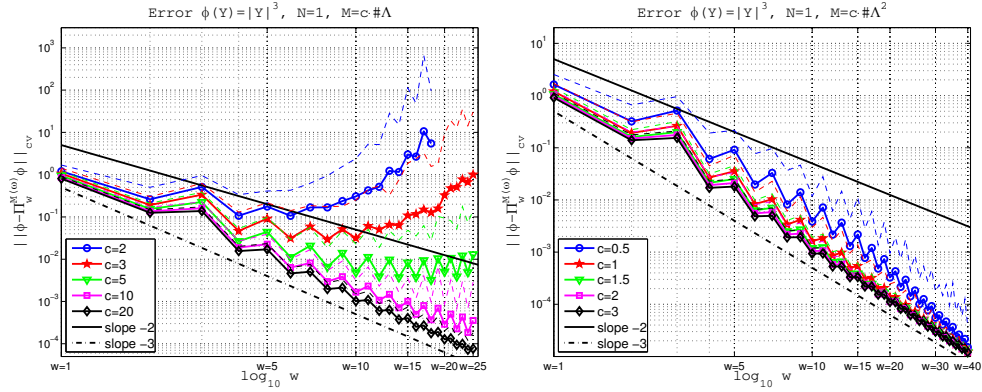
Fig. 4: Error (4.6) for the function (5.1). The continuous marked lines show the mean error (4.7) over 10000 repetitions. The dashed lines show the mean (4.7) plus one standard deviation (4.8). Left: $M = c \# \Lambda$. Right: $M = c \# \Lambda^2$.

The use of the meromorphic function (5.2) with $\beta = 0.5$ (Fig. 5) and $\beta = 0.9$ (Fig. 6) yields analogous error graphs, but with a slower convergence rate.
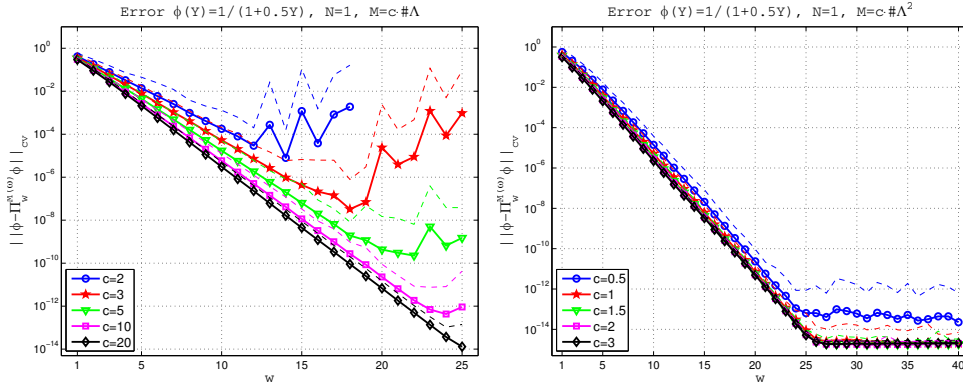
Fig. 5: Error (4.6) for the function (5.2) with $\beta = 0.5$. The continuous marked lines show the mean error (4.7) over 10000 repetitions. The dashed lines show the mean (4.7) plus one standard deviation (4.8). Left: $M = c \# \Lambda$. Right: $M = c \# \Lambda^2$.

Unlike the function (5.2), which is analytic in $[-1, 1]$, the function (5.3) is only in $C^2([-1, 1])$, but not in $C^3([-1, 1])$. This decreased regularity manifests in the slower decay of the approximation error in Fig. 7. Note that the dependence of the error on the polynomial degree w is displayed in log-log scale, so that the error no longer decreases exponentially with respect to w.

When taking the number of sample points according to the quadratic rule (Fig. 7-Right), the error decreases like w^{-3}, and in this range of w the error shows no tendency to blow up for any of the studied values of c.

On the other hand, using the linear rule (Fig. 7-Left) yields a deterioration:

Fig. 6: Error (4.6) for the function (5.2) with $\beta = 0.9$. The continuous marked lines show the mean error (4.7) over 10000 repetitions. The dashed lines show the mean (4.7) plus one standard deviation (4.8). Left: $M = c \# \Lambda$. Right: $M = c \# \Lambda^2$.

the critical w, above which the error starts to grow, increases with increasing c.

Note, in particular, that sooner or later the error starts to blow up for all shown constants. This is a clear indication that the linear rule does not lead to a stable and convergent approximation.

Fig. 7: Error (4.6) for the function (5.3). The continuous marked lines show the mean error (4.7) over 10000 repetitions. The dashed lines show the mean (4.7) plus one standard deviation (4.8). Left: $M = c \# \Lambda$. Right: $M = c \# \Lambda^2$.

From a practical point of view, we are mainly interested in the error as a function of the computational work, not the polynomial degree itself. Fig. 8 shows how the error depends on the total number of sampling points, when we consider the function (5.2) with $\beta = 0.5$. Note that Fig. 8 shows the same errors as Fig. 5 but with M instead of w on abscissas.

In Fig. 8-Left we show the linear case: the error decays exponentially with

increasing M in an initial phase, until the error starts to deteriorate. The convergence is faster for small values of c, but the deterioration also happens earlier, which prevents higher accuracies.

In Fig. 8-Right we show the quadratic case. In contrast to the linear case the convergence becomes subexponential with respect to M. On the other hand, all choices of $c \geq 1$ avoid the deterioration of the errors that we see using the linear rule, and the approximation remains stable and convergent. Fig. 9 compares the convergence of the error obtained with the linear and quadratic rules, with respect to M.

We remark that, even though the error deteriorates for high w when using the linear rule with a small c, we can still obtain an accuracy that suffices in many applications.
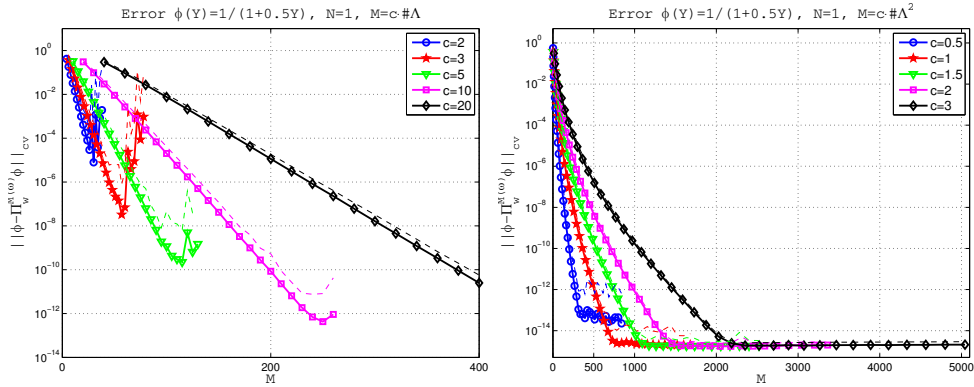
Fig. 8: Dependence of the error (4.6) on the number of sample points M. The function is (5.2) with $\beta = 0.5$. The continuous marked lines show the mean error (4.7) over 10000 repetitions. The dashed lines show the mean (4.7) plus one standard deviation (4.8). Left: $M = c \# \Lambda$. Right: $M = c \# \Lambda^2$.

The plots in Fig. 8 show how the convergence speed is affected by the ratio between M and w. The convergence is fastest with the lowest constant $c = M/\text{w}$ up to a certain point when blow up occurs.

A lower number of repetitions only turns out in an amplified variance in the results. In view of the multiD section, where we choose $R = 100$, we report also in Fig. 10 the same graph of Fig. 3-Right, but with $R = 200$.

5.2 Multidimensional case

We now proceed to the multidimensional case, where we have an increased freedom to choose the space $\mathbb{P}_\Lambda(\Gamma)$. We will restrict our examples to isotropic versions of the TP, TD and HC spaces mentioned above. In this section we choose $R = 100$ repetitions to estimate the variability of the error and condition number. In the linear case, the values assumed by c are $1.1, 1.25, 2, 5, 20$. In the multidimensional case a constant c slightly larger than 1 is enough to have a
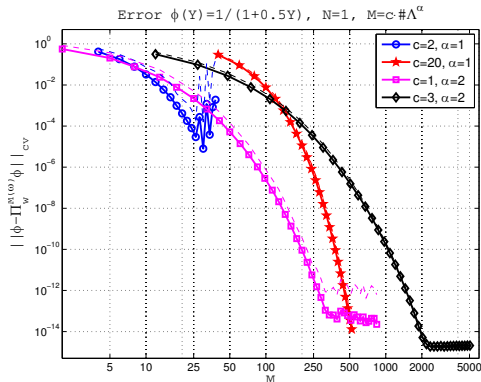
Fig. 9: Dependence of the error (4.6) on the number of sample points M. Selected data from Fig. 8 corresponding to the rules $M = 2 \# \Lambda$, $M = 20 \# \Lambda$, $M = 1 \# \Lambda^2$, and $M = 3 \# \Lambda^2$.
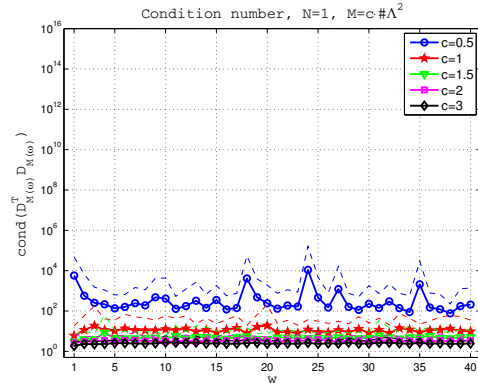
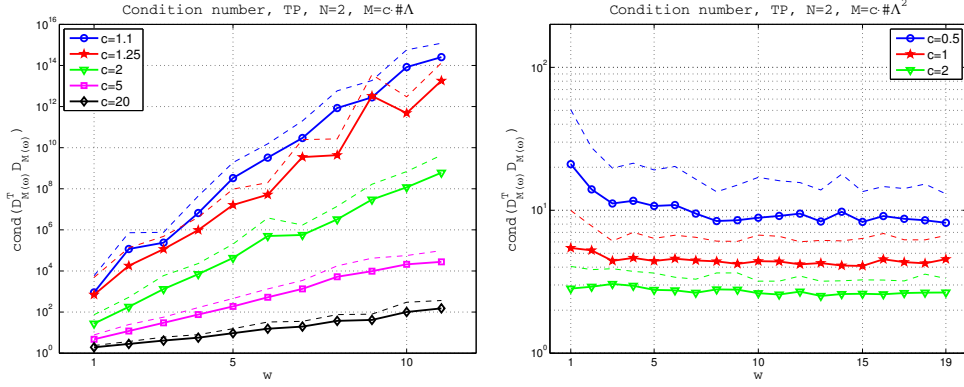Fig. 10: Condition number (4.9), $N = 1$. The continuous, marked, lines show the mean condition number (4.10) over 200 repetitions. The dashed lines show the mean (4.10) plus one standard deviation (4.11). $M = c \# \Lambda^2$.

good approximation. This is in contrast to the 1D case, where the linear rule with a constant $c = 2$ features a fast growth of the condition number and a large variability of the polynomial approximation.

5.2.1 Condition number

Fig. 11 shows the behavior of the condition number for the 2D TP space. We see again an exponential growth of the condition number when M is chosen according to the linear rule. Note that the dimension of the PC space is equal to $(\text{w} + 1)^2$ here.

As in the one dimensional case, choosing the number of sample points M like $M \propto \# \Lambda^2$ yields condition numbers that are approximately constant in the studied range of w. Compared to the one dimensional results of Fig. 3, the two dimensional results exhibit a lower variability.

Changing the polynomial space to TD we obtain Fig. 12, which looks similar to Fig. 11. The same holds with the HC space (Fig. 13). Therefore, the choice of the space does not seem to play a major role in the behavior of the condition number (4.9). Note that the lowest value of c on the left of Fig. 13 is 1, instead of 1.1.

The situation is similar in higher dimensions; see Figs. 14– 16. In addition, we observe a lower variability and a slower growth of the condition number with the linear rule, which, however, is still clearly exponential. Lastly, the HC space with the linear rule shows a very slow growth of the condition number also for very low values of c.

Fig. 11: Condition number (4.9), TP, $N = 2$. The continuous, marked, lines show the mean condition number (4.10) over 100 repetitions. The dashed lines show the mean (4.10) plus one standard deviation (4.11). Left: $M = c \# \Lambda$. Right: $M = c \# \Lambda^2$.
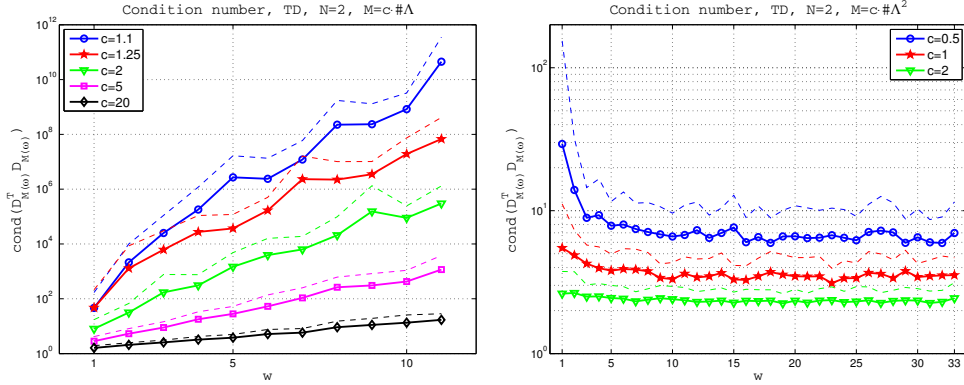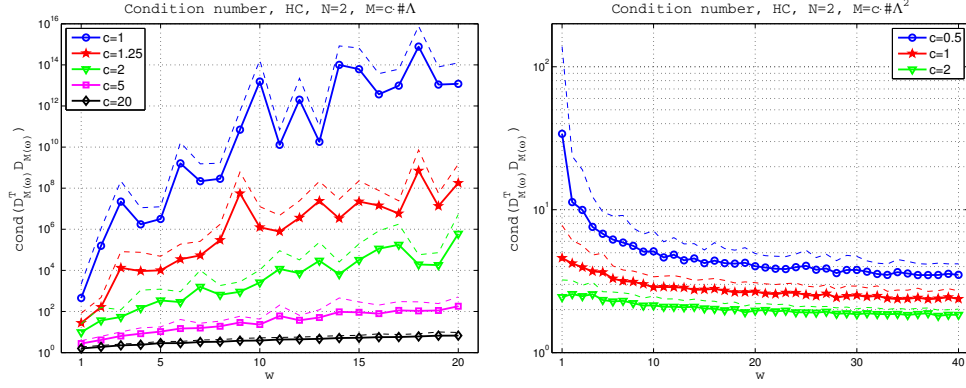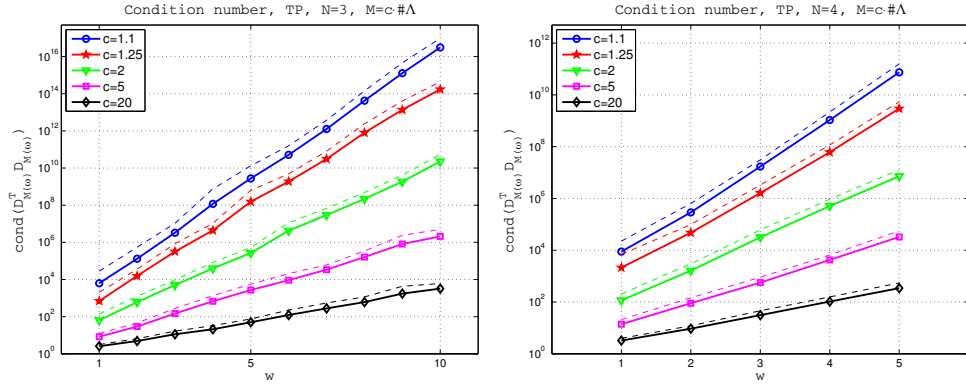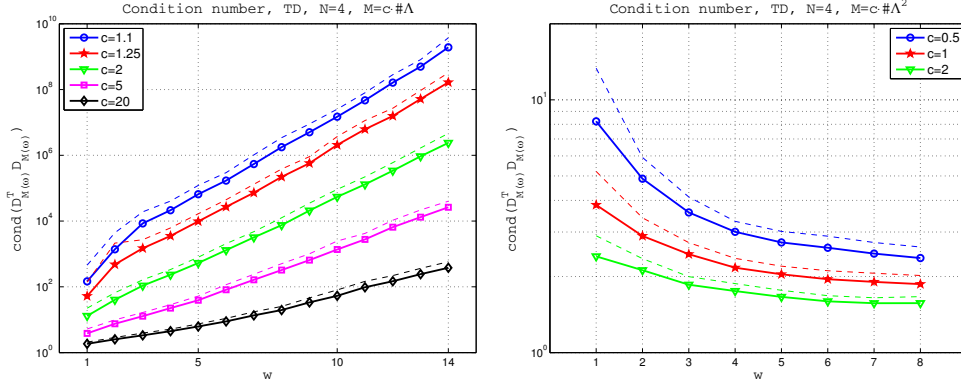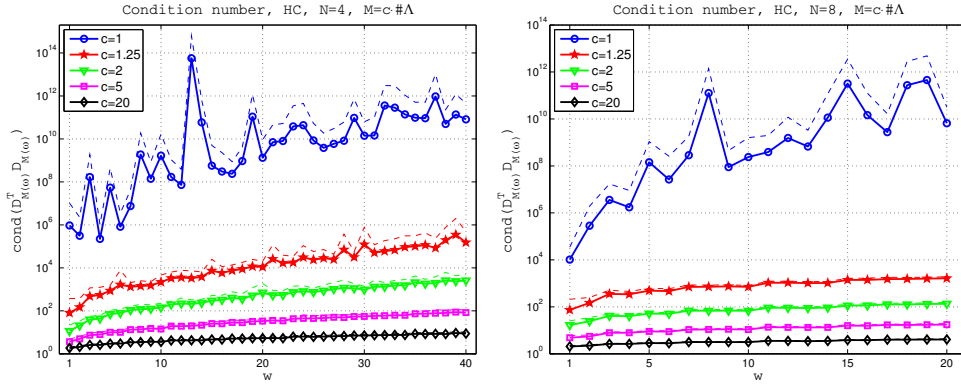

Fig. 12: Condition number (4.9), TD, $N = 2$. The continuous, marked, lines show the mean condition number (4.10) over 100 repetitions. The dashed lines show the mean (4.10) plus one standard deviation (4.11). Left: $M = c \# \Lambda$. Right: $M = c \# \Lambda^2$.

Fig. 13: Condition number (4.9), HC, $N = 2$. The continuous, marked, lines show the mean condition number (4.10) over 100 repetitions. The dashed lines show the mean (4.10) plus one standard deviation (4.11). Left: $M = c \# \Lambda$. Right: $M = c \# \Lambda^2$.

Fig. 14: Condition number (4.9), TP, $M = c \# \Lambda$. The continuous, marked, lines show the mean condition number (4.10) over 100 repetitions. The dashed lines show the mean (4.10) plus one standard deviation (4.11). Left: $N = 3$. Right: $N = 4$.

Fig. 15: Condition number (4.9), TD, $N = 4$. The continuous, marked, lines show the mean condition number (4.10) over 100 repetitions. The dashed lines show the mean (4.10) plus one standard deviation (4.11). Left: $M = c\,\#\Lambda$. Right: $M = c\,\#\Lambda^2$.

Fig. 16: Condition number (4.9), HC, $M = c\,\#\Lambda$. The continuous, marked, lines show the mean condition number (4.10) over 100 repetitions. The dashed lines show the mean (4.10) plus one standard deviation (4.11). Left: $N = 4$. Right: $N = 8$.

5.2.2 Approximation error

Let us consider the error in approximating the target function

$$\phi(\mathbf{Y}) = \frac{1}{1 + \dfrac{\beta}{N}\sum_{i=1}^{N} Y_i}, \qquad \mathbf{Y} \in [-1,1]^N, \tag{5.4}$$

which is a multidimensional generalization of (5.2), and inherits its regularity. We take $\beta = 0.5$ and start by considering the quadratic rule. Fig. 17 shows the optimal convergence rates. The TP space (Fig. 17-Left) seems to converge faster than the TD space (Fig. 17-Center), but looking at the dimension $\#\Lambda$ of the space we recognize that spaces of the same dimension (TP or TD) introduce similar approximation errors (see Tables 1 and 2). Instead, we see in Fig. 17-Right that the convergence of the HC space is slower, also when looking at the effective dimension of the space, in Table 3.
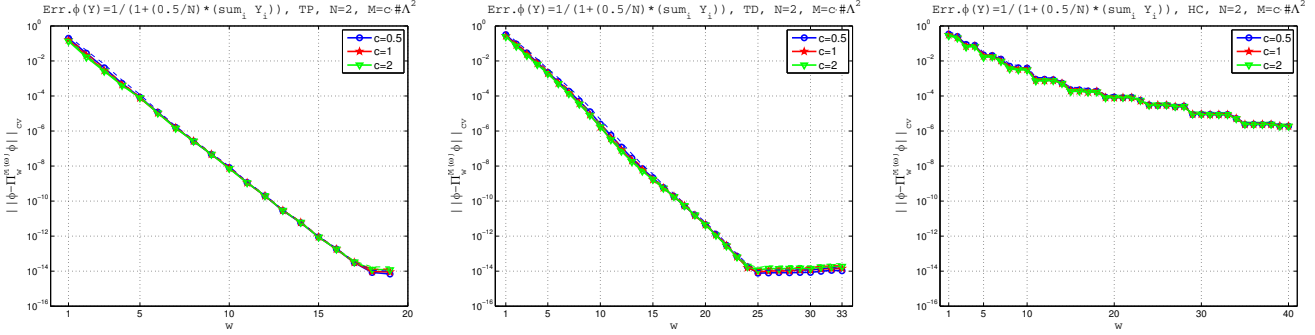
Fig. 17: Error (4.6) with the function (5.4), $\beta = 0.5$, $N = 2$, $M = c\,\#\Lambda^2$. The continuous marked lines show the mean error (4.7) over 100 repetitions. The dashed lines show the mean (4.7) plus one standard deviation (4.8). Spaces: TP (left), TD (center), HC (right).

w	# TP	$\overline{E}_{cv}$	w	# TD	$\overline{E}_{cv}$	w	# HC	$\overline{E}_{cv}$
5	36	$0.790 \cdot 10^{-4}$	7	36	$1.354 \cdot 10^{-4}$	12	37	$7.505 \cdot 10^{-4}$
7	64	$1.456 \cdot 10^{-6}$	10	66	$1.688 \cdot 10^{-6}$	19	66	$7.971 \cdot 10^{-5}$
10	121	$0.724 \cdot 10^{-8}$	14	120	$0.512 \cdot 10^{-8}$	31	119	$8.788 \cdot 10^{-6}$
11	144	$1.108 \cdot 10^{-9}$	16	153	$5.513 \cdot 10^{-10}$	40	160	$1.889 \cdot 10^{-6}$
15	256	$0.876 \cdot 10^{-12}$	21	253	$1.123 \cdot 10^{-12}$			

Table 1: Selected error values for TP and $M = 2\,\#$TP from Fig. 17-Left.

Table 2: Selected error values for TD and $M = 2\,\#$TD from Fig. 17-Center.

Table 3: Selected error values for HC and $M = 2\,\#$HC from Fig. 17-Right.

Fig. 18 shows the error in approximating the function (5.4) in TP (on the left) or TD (on the right) space, for $N = 2$. We observe a lower variability in

the error due to the reduced variability in the corresponding condition number. Despite the reduced variability we observe also in this case that the linear rule eventually leads to divergence when w increases if c is chosen too small. This effect will be much more dramatic for low regularity functions, as shown later on.
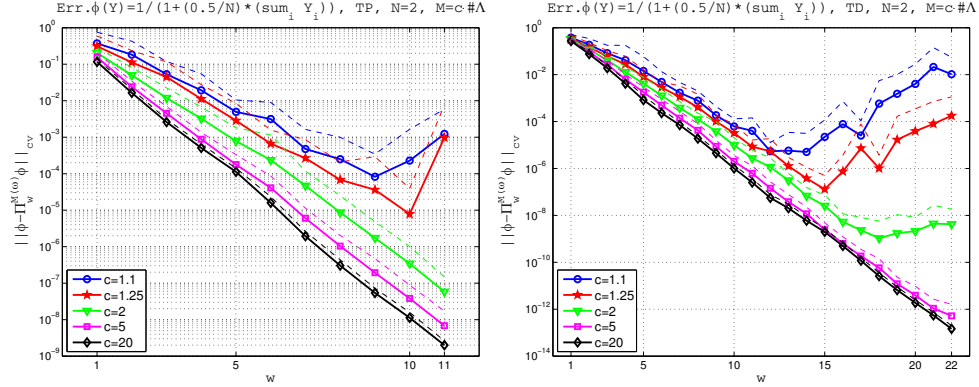

Fig. 18: Error (4.6) with the function (5.4), $N = 2$, $M = c \,\#\Lambda$. The continuous marked lines show the mean error (4.7) over 100 repetitions. The dashed lines show the mean (4.7) plus one standard deviation (4.8). Left: TP space. Right: TD space.

A further comparison can be made between Fig. 18-Left and Fig. 19-Left. In the latter case N is increased to 4 (TP space), and this yields a faster convergence with respect to w. Note however, that a larger amount of sample points is used, for corresponding w values.

Analogously, the comparison of Figs. 18-Right and Fig. 19-Right concerning the TD space reveals that in higher dimension the convergence is faster and more stable.

We also consider the function,

$$\phi(\mathbf{Y}) = \sum_{i=1}^{N} |Y_i|^3, \qquad \mathbf{Y} \in [-1, 1]^N, \tag{5.5}$$

which is a multidimensional extension of (5.3). Fig. 20 shows that the optimal convergence rate achieved by the quadratic rule does not depend on the choice of the space (*i.e.* TP, TD, or HC). In addition, for the linear rule Fig. 21 shows that the same convergence behavior is obtained when increasing the dimension N, but with a significantly decreased variability. Note that we even allowed c to take the value 1 in the linear rule.
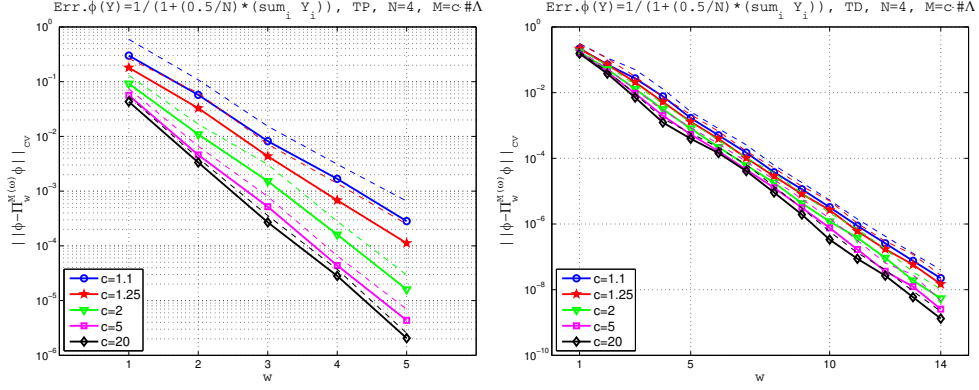
31

Fig. 19: Error (4.6) with the function (5.4), $\beta = 0.5$, $N = 4$, $M = c\,\#\Lambda$. The continuous marked lines show the mean error (4.7) over 100 repetitions. The dashed lines show the mean (4.7) plus one standard deviation (4.8). Spaces: TP (left), TD (right).
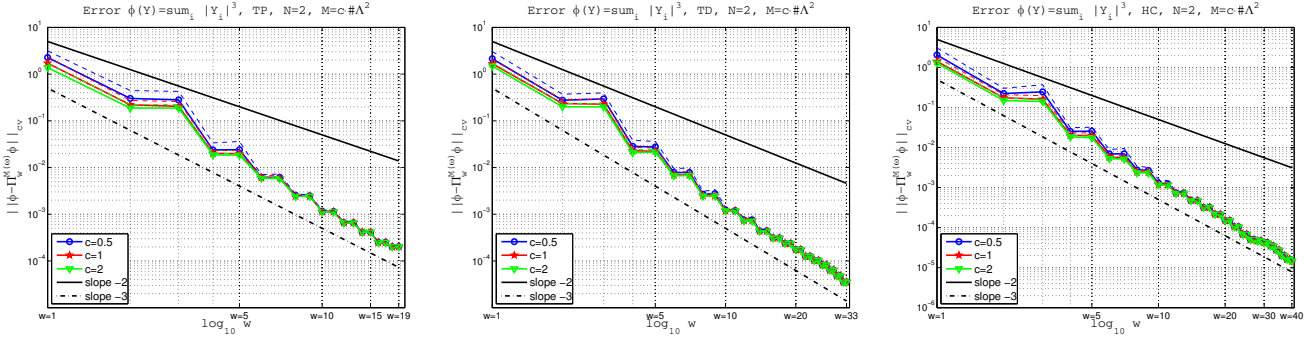

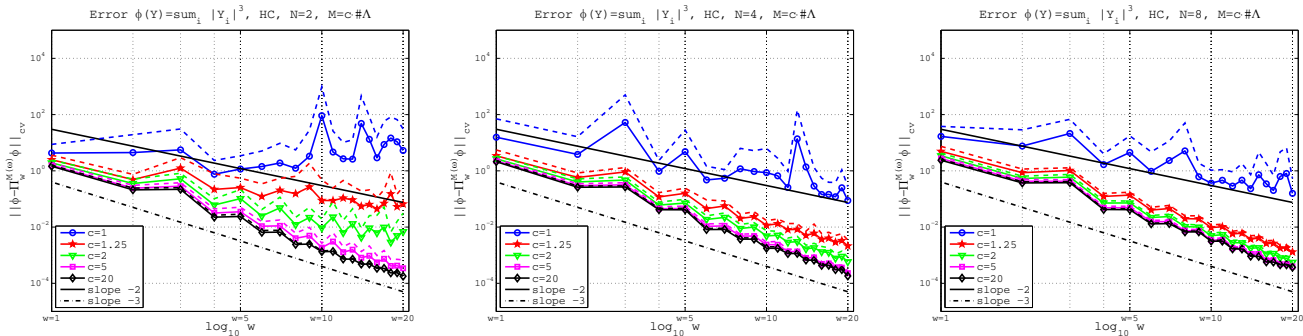Fig. 20: Error (4.6) with the function (5.5), $N = 2$, $M = c\,\#\Lambda^2$. The continuous marked lines show the mean error (4.7) over 100 repetitions. The dashed lines show the mean (4.7) plus one standard deviation (4.8). Spaces: TP (left), TD (center), HC (right).

32

Fig. 21: Error (4.6) with the function (5.5), HC space. The continuous marked lines show the mean error (4.7) over 100 repetitions. The dashed lines show the mean (4.7) plus one standard deviation (4.8). $M = c \,\#\Lambda$. $N = 2$ (left), $N = 4$ (center), $N = 8$ (right).

5.2.3 A function with lower regularity across a circle

Now we give an example of a function which is hard to approximate in the TD spaces. When $N = 2$, we consider the target function

$$\phi(\mathbf{Y}) = \left| \sum_{i=1}^{2} Y_i^2 - 0.5 \right|^3, \qquad \mathbf{Y} \in [-1, 1]^2, \tag{5.6}$$

which features a discontinuity in its derivatives over the circle with radius equal to $\sqrt{0.5}$ and centered in the origin. Note that (5.6) is a continuous function.

Choosing the quadratic rule leads to the expected theoretical convergence rates for both TP and TD spaces; see Fig. 22. However, the TD space exhibits a suboptimal convergence rate already when w ≤ 5.

When choosing the linear rule (Fig. 23), the results obtained with the TP space slightly differ from those obtained with the TD space. In particular, the convergence rate is slower than the theoretically predicted one (Fig. 23-Right).

6 Conclusions

In this work we have analysed the problem of approximating a multivariate function by discrete L^2 projection on a polynomial space starting from random, noise-free observations.

In the 1D case with sampling points drawn from a bounded domain and a probability density function bounded away from zero, we have shown that the discrete L^2 projection leads to optimal convergence rates, equivalent to the best approximation error in L^∞, up to a logarithmic factor, *provided the number of samples M scales quadratically with the dimension of the polynomial space $\#\Lambda$*. We have also shown how this result reflects on the condition number of the design matrix.
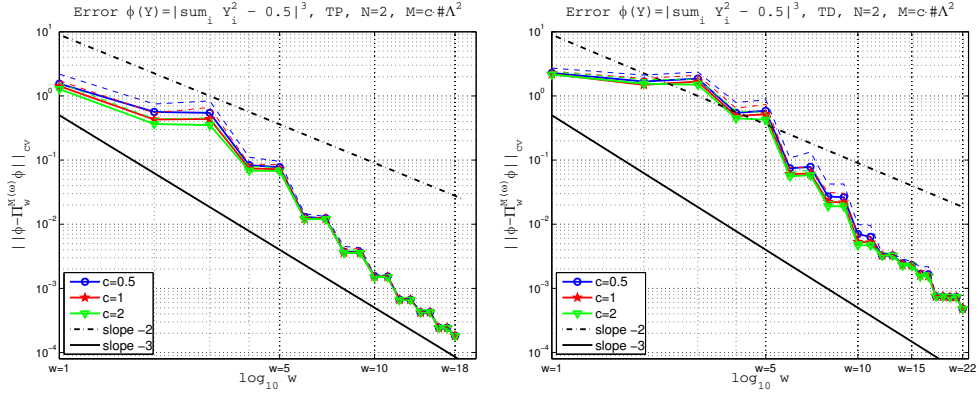
Fig. 22: Error (4.6) with the function (5.6), $N = 2$, $M = c \# \Lambda^2$. The continuous marked lines show the mean error (4.7) over 100 repetitions. The dashed lines show the mean (4.7) plus one standard deviation (4.8). Left: TP. Right: TD.
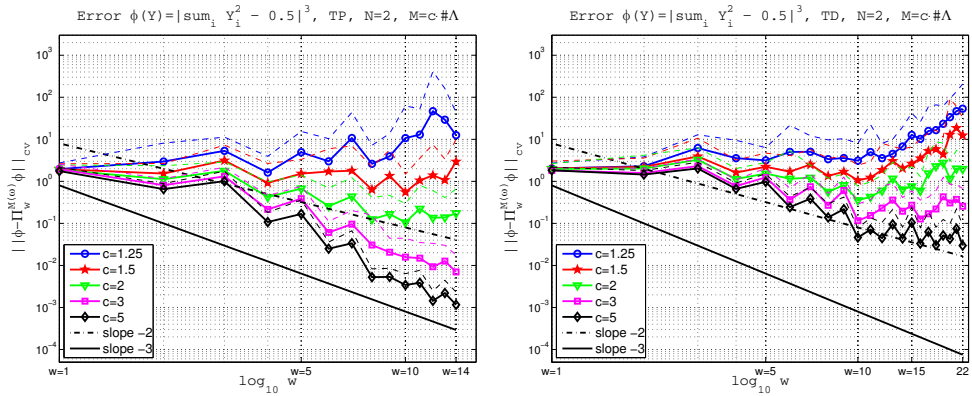

Fig. 23: Error (4.6) with the function (5.6), $N = 2$, $M = c \# \Lambda$. The continuous marked lines show the mean error (4.7) over 100 repetitions. The dashed lines show the mean (4.7) plus one standard deviation (4.8). Left: TP. Right: TD.

The numerical tests we have performed confirm the theoretical results corresponding to a uniform distribution of sample points. In our 1D tests, we clearly see that the condition $M \sim \#\Lambda^2$ guarantees a condition number of the design matrix bounded independently of the polynomial degree and an optimal convergence rate. On the other hand, the relation $M \sim \#\Lambda$ leads to a condition number growing exponentially fast with the polynomial degree and a convergece plot that features initially a suboptimal rate up to a critical polynomial degree beyond which divergence is observed. In addition, the sensitivity on the proportionality constant has been examined.

In high dimension we observe numerically in many cases that a choice $M \sim \#\Lambda$ does lead to optimal convergence rate within all reasonable tolerances (up to machine precision). Whether this is an indication that in high-D the relation $M \sim \#\Lambda$ is enough to have a stable and optimal approximation or just that the blow up of the error occurs much further (at tolerances below machine precision) is still an open question and a topic of current research.

In this work we have considered only functions with values in $\mathbb{R}$. In the field of UQ, one is often interested in functions with values in some Banach space, representing the solution of a (possibly non-linear) differential or intergral problem. Future research directions will include the extension of these results to Banach-valued functions.

Acknowledgements

References

[1] M. Ainsworth and J.T. Oden. *A posteriori error estimation in finite element analysis.* Pure and Applied Mathematics (New York). Wiley-Interscience [John Wiley & Sons], New York, 2000.

[2] I. Babuška, F. Nobile, and R. Tempone. A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM Rev.*, 52(2):317–355, 2010.

[3] J. Bäck, F. Nobile, L. Tamellini, and R. Tempone. Stochastic spectral Galerkin and collocation methods for PDEs with random coefficients: a numerical comparison. In J.S. Hesthaven and E.M. Ronquist, editors, *Spectral and High Order Methods for Partial Differential Equations*, volume 76, pages 43–62. Springer, 2011.

[4] A.M. Bagirov, C. Clausen, and M. Kohler. Estimation of a regression function by maxima of minima of linear functions. *IEEE Trans. Inform. Theory*, 55(2):833–845, 2009.

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

[6] J. Beck, F. Nobile, L. Tamellini, and R. Tempone. On the optimal polynomial approximation of stochastic PDEs by galerkin and collocation methods. MOX-Report 23-2011, Department of Mathematics, Politecnico di Milano, Italy, 2011. To appear in Math. Mod. Methods. Appl. Sci. (M3AS).

[7] P. Binev, A. Cohen, W. Dahmen, and R. DeVore. Universal algorithms for learning theory. II. Piecewise polynomial functions. *Constr. Approx.*, 26(2):127–152, 2007.

[8] P. Binev, A. Cohen, W. Dahmen, R. DeVore, and Vladimir Temlyakov. Universal algorithms for learning theory. I. Piecewise constant functions. *J. Mach. Learn. Res.*, 6:1297–1321, 2005.

[9] G. Blatman and B. Sudret. Sparse polynomial chaos expansions and adaptive stochastic finite elements using regression approach. *C.R.Mechanique*, 336:518–523, 2008.

[10] G. Blatman and B. Sudret. Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of Computational Physics*, 230(6):2345 – 2367, 2011.

[11] C. Canuto and A. Quarteroni. Approximation results for orthogonal polynomials in Sobolev spaces. *Math. Comp.*, 38(157):67–86, 1982.

[12] A. Cohen, M. Davenport, R. DeVore, and D. Leviatan. On the stability and accuracy of least squares approximations. Private communication.

[13] A. Cohen, R. DeVore, and C. Schwab. Convergence rates of best N-term Galerkin approximations for a class of elliptic sPDEs. *Found. Comput. Math.*, 10(6):615–646, 2010.

[14] A. Cohen, R. Devore, and C. Schwab. Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDE's. *Anal. Appl. (Singap.)*, 9(1):11–47, 2011.

[15] T. Crestaux, O. Le Maitre, and J.M. Maritinez. Polynomial chaos expansion for sensitivity analysis. *Reliability Engineering and System Safety*, 2009.

[16] H. A. David and H. N. Nagaraja. *Order statistics*. Wiley-Interscience, Third edition, 2003.

[17] B. Debusschere, H. Najm, P. Pebray, O. Knio, R. Ghanem, and O. Le Maitre. Numerical challenges in the use of polynomial chaos representations for stochastic processes. *SIAM Journal on Scientific Computing*, 26:698–719, 2004.

[18] M.S. Eldred. Recent advances in non-intrusive polynomial chaos and stochastic collocation methods for uncertainty analysis and design. In *AIAA*, 2009.

[19] M.S. Eldred. Design under uncertainty employing stochastic expansion methods. *International Journal for Uncertainty Quantification*, 1:119–146, 2011.

[20] G.R. Ghanem and P.D. Spanos. *Stochastic Finite Elements: a spectral approach*. Dover Publications, 2003.

[21] L. Gyørfi, M. Kohler, A. Krzyżak, and H. Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer Series in Statistics. Springer-Verlag, Berlin, first edition, 2002.

[22] S. Hosder, R. W. Walters, and M. Balch. Point-collocation nonintrusive polynomial chaos method for stochastic computational fluid dynamics. *AIAA Journal*, 48:2721–2730, 2010.

[23] M. Kohler. Inequalities for uniform deviations of averages from expectations with applications to nonparametric regression. *J. Statist. Plann. Inference*, 89(1-2):1–23, 2000.

[24] M. Kohler. Nonlinear orthogonal series estimates for random design regression. *J. Statist. Plann. Inference*, 115(2):491–520, 2003.

[25] O.P. Le Maître and O.M. Knio. *Spectral Methods for Uncertainty Quantification*. Springer, 2010.

[26] F. Nobile, R. Tempone, and C.G. Webster. A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J. Numer. Anal.*, 46(5):2309–2345, 2008.

[27] W. Rudin. *Principles of mathematical analysis.* McGraw-Hill Book Co., New York, third edition, 1976.

[28] Bruno Sudret. Global sensitivity analysis using polynomial chaos expansion. *Reliability Engineering and System Safety*, 2007.

[29] D. Xiu and G.E. Karniadakis. The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM J. Sci. Comput.*, 24(2):619–644 (electronic), 2002.

[30] Dongbin Xiu. Fast numerical methods for stochastic computations: a review. *Commun. Comput. Phys.*, 5(2-4):242–272, 2009.