SMAI-JCM
SMAI Journal of Computational Mathematics

Optimal weighted least-squares methods

Albert Cohen & Giovanni Migliorati

Volume 3 (2017), p. 181-203.
<http://smai-jcm.cedram.org/item?id=SMAI-JCM_2017__3__181_0>

SMAI Journal of Computational Mathematics
Vol. 3, 181-203 (2017)

Optimal weighted least-squares methods

Albert Cohen [1]
Giovanni Migliorati [2]

[1] Sorbonne Universités, UPMC Univ Paris 06, CNRS, UMR 7598, Laboratoire Jacques-Louis Lions, 4, place Jussieu 75005, Paris, France.
E-mail address: cohen@ljll.math.upmc.fr
[2] Sorbonne Universités, UPMC Univ Paris 06, CNRS, UMR 7598, Laboratoire Jacques-Louis Lions, 4, place Jussieu 75005, Paris, France.
E-mail address: migliorati@ljll.math.upmc.fr.

Abstract. We consider the problem of reconstructing an unknown bounded function u defined on a domain $X \subset \mathbb{R}^d$ from noiseless or noisy samples of u at n points $(x^i)_{i=1,\dots,n}$. We measure the reconstruction error in a norm $L^2(X, d\rho)$ for some given probability measure $d\rho$. Given a linear space V_m with $\dim(V_m) = m \leq n$, we study in general terms the weighted least-squares approximations from the spaces V_m based on independent random samples. It is well known that least-squares approximations can be inaccurate and unstable when m is too close to n, even in the noiseless case. Recent results from [6, 7] have shown the interest of using weighted least squares for reducing the number n of samples that is needed to achieve an accuracy comparable to that of best approximation in V_m, compared to standard least squares as studied in [4]. The contribution of the present paper is twofold. From the theoretical perspective, we establish results in expectation and in probability for weighted least squares in general approximation spaces V_m. These results show that for an optimal choice of sampling measure $d\mu$ and weight w, which depends on the space V_m and on the measure $d\rho$, stability and optimal accuracy are achieved under the mild condition that n scales linearly with m up to an additional logarithmic factor. In contrast to [4], the present analysis covers cases where the function u and its approximants from V_m are unbounded, which might occur for instance in the relevant case where $X = \mathbb{R}^d$ and $d\rho$ is the Gaussian measure. From the numerical perspective, we propose a sampling method which allows one to generate independent and identically distributed samples from the optimal measure $d\mu$. This method becomes of interest in the multivariate setting where $d\mu$ is generally not of tensor product type. We illustrate this for particular examples of approximation spaces V_m of polynomial type, where the domain X is allowed to be unbounded and high or even infinite dimensional, motivated by certain applications to parametric and stochastic PDEs.

Math. classification. 41A10, 41A25, 41A65, 62E17, 93E24.

Keywords. multivariate approximation, weighted least squares, error analysis, convergence rates, random matrices, conditional sampling, polynomial approximation.

1. Introduction

1.1. Reconstruction from pointwise data

Let X be a Borel set of $\mathbb{R}^d$. We consider the problem of estimating an unknown function $u : X \to \mathbb{R}$ from pointwise data $(y^i)_{i=1,\dots,n}$ which are either noiseless or noisy observations of u at points $(x^i)_{i=1,\dots,n}$ from X. In numerous applications of interest, some prior information is either established or assumed on the function u. Such information may take various forms such as:

(i) regularity properties of u, in the sense that it belongs to a given smoothness class;

(ii) decay or sparsity of the expansion of u in some given basis;

(iii) approximability of u with some prescribed error by given finite-dimensional spaces.

This research is supported by Institut Universitaire de France and the ERC AdV project BREAD.

Note that the above are often related to one another and sometimes equivalent, since many smoothness classes can be characterized by prescribed approximation rates when using certain finite-dimensional spaces or truncated expansions in certain bases.

This paper uses the third type of prior information, taking therefore the view that u can be "well approximated" in some space V_m of functions defined everywhere on X, such that $\dim(V_m) = m$. We work under the following mild assumption:

$$\text{for any } x \in X, \text{ there exists } v \in V_m \text{ such that } v(x) \neq 0. \tag{1.1}$$

This assumption holds, for example, when V_m contains the constant functions. Typically, the space V_m comes from a family $(V_j)_{j\geq 1}$ of nested spaces with increasing dimension, such as algebraic or trigonometric polynomials, or piecewise polynomial functions on a hierarchy of meshes.

We are interested in measuring the error in the $L^2(X, d\rho)$ norm

$$\|v\| := \left(\int_X |v|^2 d\rho\right)^{1/2},$$

where $d\rho$ is a given probability measure on X. We denote by $\langle \cdot, \cdot \rangle$ the associated inner product. One typical strategy is to pick the estimate from a finite-dimensional space V_m such that $\dim(V_m) = m$. The ideal estimator is given by the $L^2(X, d\rho)$ orthogonal projection of u onto V_m, namely

$$P_m u := \operatorname*{argmin}_{v\in V_m} \|u - v\|.$$

In general, this estimator is not computable from a finite number of observations. The best approximation error

$$e_m(u) := \min_{v\in V_m} \|u - v\| = \|u - P_m u\|,$$

thus serves as a benchmark for a numerical method based on a finite sample.

One objective of such numerical methods is therefore to approach the accuracy $e_m(u)$ by using a sample of minimal possible size. One of the main results of this paper shows that by using a well-chosen random sample of size n linearly proportional to m, up to a logarithmic factor, one can construct an estimator $\tilde{u} \in V_m$ such that $\|u - \tilde{u}\|$ is comparable to $e_m(u)$ in expectation, or to another best approximation error with high probability. The construction of this estimator is based on properly weighted least-squares methods.

1.2. **Discrete least-squares methods**

In the subsequent analysis, we make significant use of an arbitrary $L^2(X, d\rho)$ orthonormal basis $\{L_1, \ldots, L_m\}$ of the space V_m. We also introduce the notation

$$e_m(u)_\infty := \min_{v\in V_m} \|u - v\|_{L^\infty},$$

where L^∞ is meant with respect to $d\rho$, and observe that $e_m(u) \leq e_m(u)_\infty$ for any probability measure $d\rho$.

The *weighted least-squares* method consists in defining the estimator as

$$u_W := \operatorname*{argmin}_{v\in V_m} \frac{1}{n}\sum_{i=1}^n w^i |v(x^i) - y^i|^2, \tag{1.2}$$

where the weights $w^i > 0$ are given. In the noiseless case $y^i = u(x^i)$, this also writes

$$\operatorname*{argmin}_{v\in V_m} \|u - v\|_n, \tag{1.3}$$

where the discrete seminorm is defined by

$$\|v\|_n := \left(\frac{1}{n} \sum_{i=1}^{n} w^i |v(x^i)|^2 \right)^{1/2}. \tag{1.4}$$

This seminorm is associated with the semi-inner product $\langle \cdot, \cdot \rangle_n$. If we expand the solution to (1.3) as $\sum_{j=1}^{m} v_j L_j$, the vector $\mathbf{v} = (v_j)_{j=1,\dots,m}$ is the solution to the normal equations

$$\mathbf{G}\mathbf{v} = \mathbf{d}, \tag{1.5}$$

where the $m \times m$ Gramian matrix $\mathbf{G}$ has entries

$$\mathbf{G}_{j,k} := \langle L_j, L_k \rangle_n, \quad j,k = 1, \dots, m, \tag{1.6}$$

and where the data vector $\mathbf{d} = (d_j)_{j=1,\dots,m}$ is given by $d_j := \frac{1}{n} \sum_{i=1}^{n} w^i y^i L_j(x^i)$. This system always has at least one solution, which is unique when $\mathbf{G}$ is nonsingular. When $\mathbf{G}$ is singular, we may define u_W as the unique minimal ℓ^2 norm solution to (1.5).

Note that $\mathbf{G}$ is nonsingular if and only if $\|\cdot\|_n$ is a proper norm on the space V_m. Then, if the data are noisefree that is, when $y^i = u(x^i)$, we may also write

$$u_W = P_m^n u,$$

where P_m^n is the orthogonal projection onto V_m for the norm $\|\cdot\|_n$.

In practice, for the estimator (1.2) to be easily computable, it is important that the functions $L_1, \dots, L_m$ have explicit expressions that can be evaluated at any point in X so that the system (1.5) can be assembled. Let us note that computing this estimator by solving (1.5) only requires that $\{L_1, \dots, L_m\}$ be a basis of the space V_m, not necessarily orthonormal in $L^2(X, d\rho)$. Yet, since our subsequent analysis of this estimator makes use of an $L^2(X, d\rho)$ orthonormal basis, we simply assume that $\{L_1, \dots, L_m\}$ is of such type. In addition, this analysis shows that when using such an orthonormal basis, the condition number of the Gramian matrix $\mathbf{G}$ can be controlled, which warrants the stability of numerical computations.

In our subsequent analysis, we sometimes work under the assumption of a known uniform bound

$$\|u\|_{L^\infty} \leq \tau. \tag{1.7}$$

We introduce the truncation operator

$$z \mapsto T_\tau(z) := \operatorname{sign}(z) \min\{|z|, \tau\},$$

and we study the *truncated weighted least-squares approximation* defined by

$$u_T := T_\tau \circ u_W.$$

Note that, in view of (1.7), we have $|u - u_T| \leq |u - u_W|$ in the pointwise sense and therefore

$$\|u - u_T\| \leq \|u - u_W\|.$$

The truncation operator aims at avoiding instabilities which may occur when the matrix $\mathbf{G}$ is ill-conditioned. In this paper, we use randomly chosen points x^i, and corresponding weights $w^i = w(x^i)$, distributed in such a way that the resulting random matrix $\mathbf{G}$ concentrates towards the identity $\mathbf{I}$ as n increases. Therefore, if no L^∞ bound is known, an alternative strategy consists in setting to zero the estimator when $\mathbf{G}$ deviates from the identity by more than a given value in the spectral norm. We recall that for $m \times m$ matrices $\mathbf{X}$, this norm is defined as $\|\mathbf{X}\|_2 := \sup_{\|\mathbf{v}\|_2 = 1} \|\mathbf{X}\mathbf{v}\|_2$. More precisely, we introduce the *conditioned least-squares approximation*, defined by

$$u_C := \begin{cases} u_W, & \text{if } \|\mathbf{G} - \mathbf{I}\|_2 \leq \frac{1}{2}, \\ 0, & \text{otherwise.} \end{cases}$$

The choice of $\frac{1}{2}$ as a threshold for the distance between $\mathbf{G}$ and $\mathbf{I}$ in the spectral norm is related to our subsequent analysis. However, the value $\frac{1}{2}$ could be be replaced by any real number in $]0,1[$ up to some minor changes in the formulation of our results. Note that

$$\|\mathbf{G}-\mathbf{I}\|_2 \leq \frac{1}{2} \implies \text{cond}(\mathbf{G}) \leq 3. \tag{1.8}$$

In this paper, we consider least-squares approximations $\tilde{u}$ of the above types u_W, u_T and u_C.

It is well known that if $n \geq m$ is too much close to m, weighted least-squares methods may become unstable and inaccurate for most sampling distributions. For example, if $X = [-1,1]$ and $V_m = \mathbb{P}_{m-1}$ is the space of algebraic polynomials of degree $m-1$, then with $m = n$ the estimator coincides with the Lagrange polynomial interpolation which can be highly unstable and inaccurate, in particular for equispaced points. The question that we want to address here in general terms is therefore:

Given a space V_m and a measure $d\rho$, how to best choose the samples y^i and weights w^i in order to ensure that the $L^2(X,d\rho)$ error $\|u-\tilde{u}\|$ is comparable to $e_m(u)$, with n being as close as possible to m, for $\tilde{u} \in \{u_W, u_T, u_C\}$?

We address this question in the case where the x^i are randomly chosen. More precisely, we draw the independent x^i according to a certain probability measure $d\mu$ defined on X. A natural prescription for the success of the method is that $\|v\|_n$ approaches $\|v\|$ as n tends to $+\infty$.

1.3. **Previous results on unweighted least squares**

One first obvious choice is to use

$$d\mu = d\rho \quad \text{and} \quad w^i = 1, \quad i = 1,\dots,n, \tag{1.9}$$

that is, sample according to the measure in which we plan to evaluate the L^2 error and use equal weights. When using equal weights $w^i = 1$, the weighted least-squares estimator (1.2) becomes the *standard least-squares* estimator, as a particular case. The strategy (1.9) was analyzed in [4], through the introduction of the function

$$x \mapsto k_m(x) := \sum_{j=1}^{m} |L_j(x)|^2,$$

which is the diagonal of the integral kernel of the projector P_m. This function only depends on V_m and $d\rho$. It is strictly positive in X due to Assumption 1.1. Its reciprocal function is characterized by

$$\frac{1}{k_m(x)} = \min_{v \in V_m, v(x)=1} \|v\|^2,$$

and is called Christoffel function in the particular case where V_m is the space of algebraic polynomials of total degree $m-1$, see [12]. Obviously, the function k_m satisfies

$$\int_X k_m d\rho = m. \tag{1.10}$$

We define

$$K_m = K_m(V_m, d\rho) := \|k_m\|_{L^\infty},$$

and recall the following results from [4, 11] for the standard least-squares method with the weights and the sampling measure chosen as in (1.9).

Theorem 1.1. *For any $r > 0$, if m and n are such that the condition*

$$K_m \leq \kappa \frac{n}{\ln n}, \quad \text{with} \quad \kappa := \kappa(r) = \frac{1-\ln 2}{2+2r}, \tag{1.11}$$

is satisfied, then the following holds:

(i) *The matrix* $\mathbf{G}$ *satisfies the tail bound*

$$\Pr\left\{\|\mathbf{G}-\mathbf{I}\|_2 > \frac{1}{2}\right\} \leq 2n^{-r}. \tag{1.12}$$

(ii) *If* $u \in L^2(X, d\rho)$ *satisfies a uniform bound* (1.7)*, then the truncated least-squares estimator satisfies, in the noiseless case,*

$$\mathbb{E}(\|u-u_T\|^2) \leq (1+\varepsilon(n))e_m(u)^2 + 8\tau^2 n^{-r}, \tag{1.13}$$

where $\varepsilon(n) := \frac{4\kappa}{\ln(n)} \to 0$ *as* $n \to +\infty$*, and* κ *as in* (1.11)*.*

(iii) *With probability larger than* $1-2n^{-r}$*, the truncated and nontruncated least-squares estimators satisfy, in the noiseless case,*

$$\|u-u_T\| \leq \|u-u_W\| \leq (1+\sqrt{2})e_m(u)_\infty, \tag{1.14}$$

for all $u \in L^\infty(X, d\rho)$*.*

The second item in the above result shows that the optimal accuracy $e_m(u)$ is met in expectation, up to an additional term of order n^{-r}. When $e_m(u)$ has polynomial decay $\mathcal{O}(m^{-s})$, we are ensured that this additional term can be made negligible by taking r strictly larger than $s/2$, which amounts to taking $\kappa(r)$ small enough. Condition (1.11) imposes a minimal number of samples to ensure stability and accuracy of standard least squares. Since (1.10) implies that $K_m \geq m$, the fulfillment of this condition requires that n is at least of the order $m\ln(m)$. However simple examples show that the restriction can be more severe, for example if $V_m = \mathbb{P}_{m-1}$ on $X = [-1,1]$ and with $d\rho$ being the uniform probability measure. In this case, one choice for the L_j are the Legendre polynomials with proper normalization $\|L_j\|_{L^\infty} = |L_j(1)| = \sqrt{1+2j}$ so that $K_m = m^2$, and therefore condition (1.11) imposes that n is at least of order $m^2\ln(m)$. Other examples in the multivariate setting are discussed in [2, 3] which show that for many relevant approximation spaces V_m and probability measures $d\rho$, the behaviour of K_m is superlinear in m, leading to a very demanding regime in terms of the needed number n of samples. In the case of multivariate downward closed polynomial spaces, precise upper bounds for K_m have been proven in [3, 9] for measures associated to Jacobi polynomials. In addition, note that the above theory does not cover simple situations such as algebraic polynomials over unbounded domains, for example $X = \mathbb{R}$ equipped with the Gaussian measure, since the orthonormal polynomials L_j are unbounded for $j \geq 2$ and thus $K_m = \infty$ if $m \geq 2$.

The rest of our paper is organized as follows. In the next section §2, we present our main result which allows us to circumvent all the above limitations by using a proper weighted least-squares method and an optimal sampling measure $d\mu_m$. The proof of the main result is given in §3 in a concise form since it follows the same lines as the original results on standard least squares from [4, 11, 3]. We devote §4 to analogous results in the case of sample evaluations affected by additive noise, proving that the estimates are robust. The proposed method for sampling the optimal measure $d\mu_m$ is discussed in §5, and we illustrate its effectiveness in §6 by numerical examples.

2. Main results

2.1. Weighted least-squares approximation

In the present paper, we show that the above limitations can be overcome, by using a proper weighted least-squares method. We thus return to the general form of the discrete norm (1.4) used in the definition of the weighted least-squares estimator. We now use a sampling measure $d\mu$ which generally differs from $d\rho$ and is such that

$$w d\mu = d\rho,$$

where w is a positive function defined everywhere on X and such that $\int_X w^{-1} d\rho = 1$, and we then consider the weighted least-square method with weights given by

$$w^i = w(x^i).$$

With such a choice, the norm $\|v\|_n$ again approaches $\|v\|$ as n increases. The particular case $d\mu = d\rho$ and $w \equiv 1$ corresponds to the standard least-squares method analyzed by Theorem 1.1. Note that changing the sampling measure is a commonly used strategy for reducing the variance in Monte Carlo methods, where it is referred to as *importance sampling*.

With L_j again denoting the $L^2(X, d\rho)$ orthonormal basis of V_m, we now introduce the function

$$x \mapsto k_{m,w}(x) := \sum_{j=1}^{m} w(x)|L_j(x)|^2,$$

which only depends on V_m, $d\rho$ and w, as well as

$$K_{m,w} = K_{m,w}(V_m, d\rho, w) := \|k_{m,w}\|_{L^\infty}.$$

Note that, since the $\sqrt{w}L_j$ are an $L^2(X, d\mu)$ orthonormal basis of $\sqrt{w}V_m$, that is, the space consisting of the functions $\sqrt{w}g$ with $g \in V_m$, we find that $\int_X k_{m,w} d\mu = m$ and thus $K_{m,w} \geq m$. We prove in this paper the following generalization of Theorem 1.1.

Theorem 2.1. *For any $r > 0$, if m and n are such that the condition*

$$K_{m,w} \leq \kappa \frac{n}{\ln n}, \quad \text{with} \quad \kappa := \frac{1 - \ln 2}{2 + 2r}, \tag{2.1}$$

is satisfied, then the following holds:

(i) *The matrix $\mathbf{G}$ satisfies the tail bound*

$$\Pr\left\{\|\mathbf{G} - \mathbf{I}\|_2 > \frac{1}{2}\right\} \leq 2n^{-r}. \tag{2.2}$$

(ii) *If $u \in L^2(X, d\rho)$ satisfies a uniform bound (1.7), then the truncated weighted least-squares estimator satisfies, in the noiseless case,*

$$\mathbb{E}(\|u - u_T\|^2) \leq (1 + \varepsilon(n))e_m(u)^2 + 8\tau^2 n^{-r}, \tag{2.3}$$

where $\varepsilon(n) := \frac{4\kappa}{\ln(n)} \to 0$ as $n \to +\infty$, and κ as in (1.11).

(iii) *With probability larger than $1 - 2n^{-r}$, the nontruncated weighted least-squares estimator satisfies, in the noiseless case,*

$$\|u - u_W\| \leq (1 + \sqrt{2})e_m(u)_\infty, \tag{2.4}$$

for all $u \in L^\infty(X, d\rho)$.

(iv) *If $u \in L^2(X, d\rho)$, then the conditioned weighted least-squares estimator satisfies, in the noiseless case,*

$$\mathbb{E}(\|u - u_C\|^2) \leq (1 + \varepsilon(n))e_m(u)^2 + 2\|u\|^2 n^{-r}, \tag{2.5}$$

where $\varepsilon(n) := \frac{4\kappa}{\ln(n)} \to 0$ as $n \to +\infty$, and κ as in (1.11).

Notice that the probability and expectation are now taken with respect to the sampling measure $d\mu$, whereas in Theorem 1.1 they are taken in $d\rho$.

Let us mention that the quantity $K_{m,w}$ has been considered in [6], where similar stability and approximation results have been formulated in a slightly different form (see in particular Theorem 2.1 therein), in the specific framework of total degree polynomial spaces.

2.2. **Optimal sampling**

The interest of Theorem 2.1 is that it leads us in a natural way to an optimal sampling strategy for the weighted least-square method. We simply take

$$w := \frac{m}{k_m} = \frac{m}{\sum_{j=1}^m |L_j|^2}, \tag{2.6}$$

and with such a choice for w one readily checks that

$$d\mu := \frac{k_m}{m} d\rho, \tag{2.7}$$

is a probability measure on X since $\int_X k_m d\rho = m$.

In addition, we have for this particular choice that

$$k_{m,w} = w k_m = m,$$

and therefore

$$K_{m,w} = m.$$

We thus obtain the following result as a consequence of Theorem 2, which shows that the above choice of w and $d\mu$ allows us to obtain near-optimal estimates for the truncated weighted least-squares estimator, under the minimal condition that n is at least of the order $m \ln(m)$.

Corollary 2.2. *For any $r > 0$, if m and n are such that the condition*

$$m \leq \kappa \frac{n}{\ln n}, \quad \text{with} \quad \kappa := \frac{1 - \ln 2}{2 + 2r}, \tag{2.8}$$

is satisfied, then the conclusions (i), (ii), (iii) *and* (iv) *of Theorem 2.1 hold for weighted least squares with the choice of w and $d\mu$ given by* (2.6) *and* (2.7).

One of the interests of the above optimal sampling strategy is that it applies to polynomial approximation on unbounded domains that were not covered by Theorem 1.1, in particular $X = \mathbb{R}$ equipped with the Gaussian measure. In this case, the relevant target functions u are often nonuniformly bounded and therefore the results in items (ii) and (iii) of Theorem 2.1 do not apply. The result in item (iv) for the conditioned estimator u_C remains valid, since it does not require uniform boundedness of u.

Let us remark that all the above results are independent of the dimension d of the domain X. However, raising d has the unavoidable effect of restricting the classes of functions for which the best approximation error $e_m(u)$ or $e_m(u)_\infty$ have some prescribed decay, due to the well-known curse of dimensionality.

Note that the optimal pair $(d\mu, w)$ described by (2.6) and (2.7) depends on V_m, that is

$$w = w_m \quad \text{and} \quad d\mu = d\mu_m.$$

This raises a difficulty for properly choosing the samples in settings where the choice of V_m is not fixed a-priori, such as in adaptive methods. In certain particular cases, it is known that w_m and $d\mu_m$ admit limits w^* and $d\mu^*$ as $m \to \infty$ and are globally equivalent to these limits. One typical example is given by the univariate polynomial spaces $V_m = \mathbb{P}_{m-1}$, when $X = [-1, 1]$ and $d\rho = \rho dx$ where ρ is a Jacobi weight and dx is the Lebesgue measure on X. In this case $d\mu^*$ is the pluripotential equilibrium measure

$$d\mu^* = \frac{dx}{2\pi\sqrt{1 - x^2}},$$

see e.g. [8, 13], and one has

$$c d\mu^* \leq d\mu_m \leq C d\mu^*, \quad m \geq 1,$$

for some fixed constants $0 < c < C < \infty$. Thus, in such a case, the above corollary also holds for the choice $w = w^*$ and $d\mu = d\mu^*$ under the condition $m \leq \frac{c}{C}\kappa\frac{n}{\ln n}$. The development of sampling strategies in cases of varying values of m without such asymptotic equivalences is the object of current investigation.

A closely related weighted least-squares strategy was recently proposed and analyzed in [7], in the polynomial framework. There, the authors propose to use the renormalized Christoffel function (2.6) in the definition of the weights, however sampling from the fixed pluripotential equilibrium measure $d\mu^*$. Due to the fact that $d\mu_m$ differs from $d\mu^*$, the main estimate obtained in [7] (see p 3 therein) does not have the same simple form of a direct comparison between $\|u - u_T\|$ and $e_m(u)$ as in (ii) of Theorem 2.1. In particular, it involves an extra term $d(f)$ which does not vanish even as $n \to \infty$.

One intrinsic difficulty when using the optimal pair $(d\mu, w) = (d\mu_m, w_m)$ described by (2.6) and (2.7) is the effective sample generation, in particular in the multivariate framework since the measure $d\mu_m$ is generally not of tensor product type. One possible approach is to use Markov Chain Monte Carlo methods such as the Metropolis-Hastings algorithm, as explored in [6]. In such methods the samples are mutually correlated, and only asymptotically distributed according to the desired sampling measure. One contribution of the present paper is to propose a straightforward and effective sampling strategy for generating an arbitrary finite number n of independent samples identically distributed according to $d\mu_m$. This strategy requires that $d\rho$ has tensor product structure and that the spaces V_m are spanned by tensor product bases, such as for multivariate polynomial spaces, in which case $d\mu_m$ is generally not of tensor product type.

Another type of reconstruction in finite-dimensional spaces with arbitrary bases and using samples has been analysed in [1], known as "generalized sampling".

3. **Proof of Theorem 2.1**

The proof is structurally similar to that of Theorem 1.1 given in [4] for items (i) and (ii) and in [3] for item (iii), therefore we only sketch it. We observe that the matrix $\mathbf{G}$ in (1.6) can be written as $\mathbf{G} = \sum_{i=1}^n \mathbf{X}^i$ where the $\mathbf{X}^i$ are i.i.d. copies of the rank 1 random matrix

$$\mathbf{X} = \mathbf{X}(x) := \frac{1}{n}\left(w(x)L_j(x)L_k(x)\right)_{j,k=1,\dots,m},$$

with x a random variable distributed over X according to μ. One obviously has $\mathbb{E}(\mathbf{X}) = \mathbf{I}$. We then invoke the Chernov bound from [14] to obtain that if $\|\mathbf{X}\|_2 \leq R$ almost surely, then, for any $0 < \delta < 1$,

$$\Pr\left\{\|\mathbf{G} - \mathbf{I}\|_2 > \delta\right\} \leq 2m\left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^{1/R} = 2m\exp\left(-\frac{c_\delta}{R}\right), \tag{3.1}$$

with $c_\delta := \delta + (1-\delta)\ln(1-\delta) > 0$. Taking $\delta = \frac{1}{2}$, and observing that

$$\|\mathbf{X}(x)\|_2 = \frac{1}{n}w(x)\sum_{j=1}^m |L_j(x)|^2 = \frac{K_{m,w}(x)}{n},$$

we may thus take $R = \frac{K_{m,w}}{n}$ which yields (2.2) in item (i).

For the proof of (2.3) in item (ii), we first consider the event where $\|\mathbf{G} - \mathbf{I}\|_2 \leq \frac{1}{2}$. In this case we write

$$\|u - u_T\|^2 = \|T_\tau(u) - T_\tau(u_W)\|^2 \leq \|u - u_W\|^2 = \|u - P_m^n u\|^2 = \|g\|^2 + \|P_m^n g\|^2, \quad g := u - P_m u,$$

where we have used that $P_m^n P_m u = P_m u$ and that g is orthogonal to V_m, and thus

$$\|u - u_T\|^2 \leq e_m(u)^2 + \sum_{j=1}^m |a_j|^2,$$

where $\mathbf{a} = (a_j)_{j=1,\dots,m}$ is solution of the system

$$\mathbf{G}\mathbf{a} = \mathbf{b},$$

and $\mathbf{b} := (\langle g, L_k\rangle_n)_{k=1,\dots,m}$. Since $\|\mathbf{G}^{-1}\|_2 \leq 2$, it follows that

$$\|u - u_T\|^2 \leq e_m(u)^2 + 4\sum_{k=1}^{m} |\langle g, L_k\rangle_n|^2.$$

In the event where $\|\mathbf{G} - \mathbf{I}\|_2 > \frac{1}{2}$, we simply write $\|u - u_T\| \leq 2\tau$. It follows that

$$\mathbb{E}(\|u - u_T\|^2) \leq e_m(u)^2 + 4\sum_{k=1}^{m} \mathbb{E}(|\langle g, L_k\rangle_n|^2) + 8\tau^2 n^{-r}.$$

For the second term, we have

$$\begin{aligned}
\mathbb{E}(|\langle g, L_k\rangle_n|^2) &= \frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{n} \mathbb{E}(w(x^i)w(x^j)g(x^i)g(x^j)L_k(x^i)L_k(x^j)) \\
&= \frac{1}{n^2}\left(n(n-1)|\mathbb{E}(w(x)g(x)L_k(x))|^2 + n\mathbb{E}(|w(x)g(x)L_k(x)|^2)\right) \\
&= \left(1 - \frac{1}{n}\right)|\langle g, L_k\rangle|^2 + \frac{1}{n}\int_X |w(x)|^2|g(x)|^2|L_k(x)|^2 d\mu \\
&= \frac{1}{n}\int_X w(x)|g(x)|^2|L_k(x)|^2 d\rho,
\end{aligned}$$

where we have used the fact that g is $L^2(X,\rho)$-orthogonal to V_m and thus to L_k. Summing over k, we obtain

$$\sum_{k=1}^{m} \mathbb{E}(|\langle g, L_k\rangle_n|^2) \leq \frac{K_{m,w}}{n}\|g\|^2 \leq \frac{\kappa}{\ln(n)} e_m(u)^2,$$

and we therefore obtain (2.3).

For the proof of (2.4) in item (iii) we place ourselves in the event where $\|\mathbf{G} - \mathbf{I}\|_2 \leq \frac{1}{2}$. This property also means that

$$\frac{1}{2}\|\mathbf{v}\|_2^2 \leq \langle \mathbf{G}\mathbf{v}, \mathbf{v}\rangle_2 \leq \frac{3}{2}\|\mathbf{v}\|_2^2, \quad \mathbf{v} \in \mathbb{R}^m,$$

which can be expressed as a norm equivalence over V_m,

$$\frac{1}{2}\|v\|^2 \leq \|v\|_n^2 \leq \frac{3}{2}\|v\|^2, \quad v \in V_m. \tag{3.2}$$

We then write that for any $v \in V_m$,

$$\begin{aligned}
\|u - P_m^n u\| &\leq \|u - v\| + \|v - P_m^n u\| \\
&\leq \|u - v\| + \sqrt{2}\|v - P_m^n u\|_n \\
&\leq \|u - v\| + \sqrt{2}\|u - v\|_n \\
&\leq (1 + \sqrt{2})\|u - v\|_{L^\infty},
\end{aligned}$$

where we have used (3.2), the Pythagorean identity $\|u - v\|_n^2 = \|u - P_m^n u\|_n^2 + \|v - P_m^n u\|_n^2$, and the fact that both $\|\cdot\|$ and $\|\cdot\|_n$ are dominated by $\|\cdot\|_{L^\infty}$. Since v is arbitrary, we obtain (2.4).

Finally, (2.5) in item (iv) is proven in a very similar way as (2.3) in item (ii), by writing that in the event $\|\mathbf{G} - \mathbf{I}\|_2 > \frac{1}{2}$, we have $\|u - u_C\| = \|u\|$, so that

$$\mathbb{E}(\|u - u_C\|^2) \leq e_m(u)^2 + 4\sum_{k=1}^{m} \mathbb{E}(|\langle g, L_k\rangle_n|^2) + 2\|u\|^2 n^{-r},$$

and we conclude in the same way. ■

Remark 3.1. Let us stress that the probabilistic estimate (2.4) of item (iii) is stated in a *uniform* sense: once the draw falls in the event $\|\mathbf{G}-\mathbf{I}\|_2 \leq \frac{1}{2}$ which has high probability, the estimate (2.4) is valid for all functions u in L^∞. In this estimate, the L^2 error $\|u - P_m^n u\|$ is controlled by the best approximation error $e_m(u)_\infty$ in L^∞. One way to recover a control of the error by the more natural quantity $e_m(u)$ is to search for an estimate in a *nonuniform* sense: returning to the proof, in the estimate

$$\|u - P_m^n u\| \leq \|u - v\| + \sqrt{2}\|u - v\|_n,$$

we take $v := T_\tau(P_m u)$ as the truncated best L^2 approximation of u, where $\tau := \|u\|_{L^\infty}$, so that $\|u - v\| \leq e_m(u)$. Then, for the given function $u - v \in L^\infty$, one may apply a concentration inequality in order to bound $\|u - v\|_n$ by $\|u - v\| + \epsilon(n)$ with high probability where $\epsilon(n)$ decreases rapidly to 0 as n increases.

4. The noisy case

In a similar way as in [4, 10], we can analyze the case where the observations of u are affected by an additive noise. In practical situations the noise may come from different sources, such as a discretization error when u is evaluated by some numerical code, or a measurement error. The first one may be viewed as a perturbation of u by a deterministic funtion h, that is, we observe

$$y^i = u(x^i) + h(x^i).$$

The second one is typically modelled as a stochastic fluctuation, that is, we observe

$$y^i = u(x^i) + \eta^i.$$

where η^i are independent realizations of the centered random variable $\eta = y - u(x)$. Here, we do not necessarily assume η and x to be independent, however we typically assume that the noise is centered, that is,

$$\mathbb{E}(\eta|x) = 0,$$

and we also assume uniformly bounded conditional variance

$$\sigma^2 := \sup_{x\in X} \mathbb{E}(|\eta|^2|x) < \infty. \tag{4.1}$$

Note that we may also consider a noncentered noise, which amounts in adding the two contributions, that is,

$$y^i = u(x^i) + \beta^i, \quad \beta^i = h(x^i) + \eta^i, \tag{4.2}$$

with $h(x) = \mathbb{E}(\beta|x)$. The following result shows that the estimates in Theorem 2.1 are robust under the presence of such an additive noise.

Theorem 4.1. *For any $r > 0$, if m and n are such that condition* (2.1) *is satisfied, then the following holds for the noise model* (4.2)*:*

(i) *if $u \in L^2(X, d\rho)$ satisfies a uniform bound* (1.7)*, then the truncated weighted least-squares estimator satisfies*

$$\mathbb{E}(\|u - u_T\|^2) \leq (1 + 2\varepsilon(n))e_m(u)^2 + (8 + 2\varepsilon(n))\|h\|^2 + \frac{\overline{K}_{m,w}\sigma^2}{n} + 8\tau^2 n^{-r}, \tag{4.3}$$

(ii) *if $u \in L^2(X, d\rho)$, then the conditioned weighted least-squares estimator satisfies*

$$\mathbb{E}(\|u - u_C\|^2) \leq (1 + 2\varepsilon(n))e_m(u)^2 + (8 + 2\varepsilon(n))\|h\|^2 + \frac{\overline{K}_{m,w}\sigma^2}{n} + 2\|u\|^2 n^{-r}, \tag{4.4}$$

where in both cases $\varepsilon(n) := \frac{4\kappa}{\ln(n)} \to 0$ *as* $n \to +\infty$*, with* κ *as in* (1.11)*, and* $\overline{K}_{m,w} := \int_X k_{m,w} d\rho$.

Proof. We again first consider the event where $\|\mathbf{G} - \mathbf{I}\|_2 \leq \frac{1}{2}$. In this case we write

$$\|u - u_T\| \leq \|u - u_W\|,$$

and use the decomposition $u - u_W = g - P_m^n g - f$ where $g = u - P_m u$ as in the proof of Theorem 2.1 and f stands for the solution to the least-squares problem for the noise data $(\beta^i)_{i=1,\dots,n}$. Therefore

$$\|u - u_W\|^2 = \|g\|^2 + \|P_m^n g + f\|^2 \leq \|g\|^2 + 2\|P_m^n g\|^2 + 2\|f\|^2 = \|g\|^2 + 2\|P_m^n g\|^2 + 2\sum_{j=1}^m |n_j|^2,$$

where $\mathbf{n} = (n_j)_{j=1,\dots,m}$ is solution to

$$\mathbf{G}\mathbf{n} = \mathbf{b}, \quad \mathbf{b} := \left(\frac{1}{n}\sum_{i=1}^n \beta^i w(x^i) L_k(x^i)\right)_{k=1,\dots,m} = (b_k)_{k=1,\dots,m}.$$

Since $\|\mathbf{G}^{-1}\|_2 \leq 2$, it follows that

$$\|u - u_T\|^2 \leq e_m(u)^2 + 8\sum_{k=1}^m |\langle g, L_k\rangle_n|^2 + 8\sum_{k=1}^m |b_k|^2.$$

Compared to the proof of Theorem 2.1, we need to estimate the expectation of the third term on the right side. For this we simply write that

$$\mathbb{E}(|b_k|^2) = \frac{1}{n^2}\sum_{i=1}^n\sum_{j=1}^n \mathbb{E}(\beta^i w(x^i) L_k(x^i) \beta^j w(x^j) L_k(x^j)).$$

For $i \neq j$, we have

$$\mathbb{E}(\beta^i w(x^i) L_k(x^i) \beta^j w(x^j) L_k(x^j)) = \Big(\mathbb{E}(\beta w(x) L_k(x))\Big)^2 = \Big(\mathbb{E}(h(x) w(x) L_k(x))\Big)^2$$
$$= \left|\int_X h w L_k d\mu\right|^2 = |\langle h, L_k\rangle|^2.$$

Note that the first and second expectations are with respect to the joint density of (x, β) and the third one with respect to the density of x, that is, μ. For $i = j$, we have

$$\begin{aligned}
\mathbb{E}(|\beta^i w(x^i) L_k(x^i)|^2) &= \mathbb{E}(|\beta w(x) L_k(x)|^2) \\
&= \int_X \mathbb{E}(|\beta w(x) L_k(x)|^2 | x) d\mu \\
&= \int_X \mathbb{E}(|\beta|^2 | x) |w(x) L_k(x)|^2 d\mu \\
&= \int_X \mathbb{E}(|\beta|^2 | x) w(x) |L_k(x)|^2 d\rho \\
&= \int_X (|h(x)|^2 + \mathbb{E}(|\eta|^2 | x)) w(x) |L_k(x)|^2 d\rho \\
&\leq \int_X (|h(x)|^2 + \sigma^2) w(x) |L_k(x)|^2 d\rho.
\end{aligned}$$

Summing up on i, j and k, and using condition (2.1), we obtain that

$$\sum_{k=1}^m \mathbb{E}(|b_k|^2) \leq \left(1 - \frac{1}{n^2}\right)\|h\|^2 + \frac{K_{m,w}}{n}\|h\|^2 + \frac{\overline{K}_{m,w}}{n}\sigma^2 \leq \left(1 + \frac{\kappa}{\log n}\right)\|h\|^2 + \frac{\overline{K}_{m,w}\sigma^2}{n}. \tag{4.5}$$

For the rest we proceed as for item (ii) and (iv) in the proof of Theorem 2.1, using that in the event $\|\mathbf{G} - \mathbf{I}\|_2 > \frac{1}{2}$ we have $\|u - u_T\| \leq 2\tau$ and $\|u - u_C\| = \|u\|$. ∎

Remark 4.2. Note that for the standard least-squares method, corresponding to the case where $w \equiv 1$, we know that $\overline{K}_{m,w} = m$. The noise term thus takes the stardard form $\frac{m\sigma^2}{n}$, as seen for example in Theorem 3 of [4] or in Theorem 1 of [10]. Note that, in any case, condition (2.1) implies that this term is bounded by $\frac{\kappa\sigma^2}{\log n}$.

The conclusions of Theorem 4.1 do not include the estimate in probability similar to item (iii) in Theorem 2.1. We can obtain such an estimate in the case of a bounded noise, where we assume that $h \in L^\infty(X)$ and η is a bounded random variable, or equivalently, assuming that β is a bounded random variable, that is we use the noise model (4.2) with

$$|\beta| \leq D, \quad a.s. \tag{4.6}$$

For this bounded noise model we have the following result.

Theorem 4.3. *For any $r > 0$, if m and n are such that condition* (2.1) *is satisfied, then the following holds with probability larger than $1-2n^{-r}$, for the the noise model* (4.2) *under* (4.6)*: the nontruncated weighted least-squares estimator satisfies*

$$\|u - u_W\| \leq (1+\sqrt{2})e_m(u)_\infty + \sqrt{2}D, \tag{4.7}$$

for all $u \in L^\infty(X, d\rho)$.

Proof. Similar to the proof of (iii) in Theorem 2.1, we place ourselves in the event where $\|\mathbf{G}-\mathbf{I}\|_2 \leq \frac{1}{2}$ and use the norm equivalence (3.2). We then write that for any $v \in V_m$,

$$\|u - u_W\| \leq \|u - v\| + \|v - P_m^n u\| + \|P_m^n \beta\|.$$

The first two terms already appeared in the noiseless case and can be treated in the same way. The new term $P_m^n\beta$ corresponds to the weighted least-squares approximation from the noise vector, and satisfies

$$\|P_m^n\beta\| \leq \sqrt{2}\|P_m^n\beta\|_n \leq \sqrt{2}\|\beta\|_n \leq \sqrt{2}D.$$

This leads to (4.7). ∎

5. Random sampling from μ_m

The analysis in the previous sections prescribes the use of the optimal sampling measure $d\mu_m$ defined in (2.7) for drawing the samples $x^1, \ldots, x^n$ in the weighted least-squares method. In this section we discuss numerical methods for generating independent random samples according to this measure, in a specific relevant multivariate setting.

Here, we make the assumption that $X = \times_{i=1}^d X_i$ is a Cartesian product of univariate real domains X_i, and that $d\rho$ is a product measure, that is,

$$d\rho = \bigotimes_{i=1}^d d\rho_i,$$

where each $d\rho_i$ is a measure defined on X_i. We assume that each $d\rho_i$ is of the form

$$d\rho_i(t) = \rho_i(t)dt,$$

for some nonnegative continuous function ρ_i, and therefore

$$d\rho(x) = \rho(x)\,dx, \quad \rho(x) = \prod_{i=1}^d \rho_i(x_i), \quad x = (x_1, \ldots, x_d) \in X.$$

In particular $d\rho$ is absolutely continuous with respect to the Lebesgue measure.

We consider the following general setting: for each $i = 1, \dots, d$, we choose a univariate basis $(\phi_j^i)_{j\geq 0}$ orthonormal in $L^2(X_i, d\rho_i)$. We then define the tensorized basis

$$L_\nu(x) := \prod_{i=1}^d \phi_{\nu_i}^i(x_i), \quad \nu \in \mathbb{N}_0^d,$$

which is orthonormal in $L^2(X, d\rho)$. We consider general subspaces of the form

$$V_m := \operatorname{span}\{L_\nu \ : \ \nu \in \Lambda\},$$

for some multi-index set $\Lambda \subset \mathbb{N}_0^d$ such that $\#(\Lambda) = m$. Thus we may rename the $(L_\nu)_{\nu\in\Lambda}$ as $(L_j)_{j=1,\dots,m}$ after a proper ordering has been chosen, for example in the lexicographical sense.

The measure $d\mu_m$ is thus given by $d\mu_m(x) = \mu_m(x)dx$, where

$$\mu_m(x) := \frac{1}{m}\sum_{i=1}^m |L_i(x)|^2\rho(x) = \frac{1}{\#(\Lambda)}\sum_{\nu\in\Lambda}|L_\nu(x)|^2\rho(x), \quad x \in X. \tag{5.1}$$

We now discuss our sampling methods for generating n independent random samples $x^1, \dots, x^n$ identically distributed according to the multivariate density (5.1). Note that this density does not have a product structure, despite ρ is a product density. There exist many methods for sampling from multivariate densities. In contrast to Markov Chain Monte Carlo methods mentioned in the introduction, the methods that we next propose exploit the particular structure of the multivariate density (5.1), in order to generate independent samples in a straightforward manner, and sampling only from univariate densities. The first method is *sequential conditional sampling*, discussed in §5.1, and the second method is *mixture sampling*, discussed in §5.3.

Given the vector $x = (x_1, \dots, x_d)$ of all the coordinates, for any $A \subseteq \{1, \dots, d\}$, we introduce the notation

$$x_A := (x_i)_{i\in A}, \quad \bar{A} := \{1, \dots, d\} \setminus A, \quad x_{\bar{A}} := (x_i)_{i\in\bar{A}},$$

and

$$dx_A := \bigotimes_{i\in A} dx_i, \quad d\rho_A := \bigotimes_{i\in A} d\rho_i, \quad \rho_A(x_A) := \prod_{i\in A}\rho_i(x_i), \quad X_A := \bigtimes_{i\in A} X_i.$$

In the following, we mainly use the particular sets

$$A^q := \{1, \dots, q\} \quad \text{and} \quad \bar{A}^q := \{q+1, \dots, d\},$$

so that any $x \in X$ may be written as $x = (x_{A^q}, x_{\bar{A}^q})$.

Using such a notation, for any $q = 1, \dots, d$, we associate to the joint density μ_m its marginal density ψ_q of the first q variables, namely

$$\psi_q(x_{A^q}) := \int_{X_{\bar{A}^q}} \mu_m(x_{A^q}, x_{\bar{A}^q})\, dx_{\bar{A}^q}. \tag{5.2}$$

Since $(\phi_j^i)_{j\geq 0}$ is an orthonormal basis of $L^2(X_i, d\rho_i)$, for any $q = 1, \dots, d$ and any $\nu \in \mathbb{N}_0^d$, we obtain that

$$\int_{X_{\bar{A}^q}} |L_\nu(x_{A^q}, x_{\bar{A}^q})|^2\rho(x_{A^q}, x_{\bar{A}^q})dx_{\bar{A}^q} = \rho_{A^q}(x_{A^q})\prod_{i=1}^q |\phi_{\nu_i}^i(x_i)|^2, \quad x_{A^q} \in X_{A^q}.$$

Therefore, the marginal density (5.2) can be written in simple form as

$$\psi_q(x_{A^q}) = \frac{1}{\#(\Lambda)}\rho_{A^q}(x_{A^q})\sum_{\nu\in\Lambda}\prod_{i=1}^q |\phi_{\nu_i}^i(x_i)|^2. \tag{5.3}$$

In the next sections, for the given set Λ of interest, we use the notation

$$\lambda_j := \max_{\nu\in\Lambda}\nu_j \quad \text{and} \quad \lambda_\Lambda := \max_{j=1,\dots,d}\lambda_j.$$

5.1. **Sequential conditional sampling**

Based on the previous notation and remarks, we propose an algorithm which generates n samples $x^k = (x_1^k, \dots, x_d^k) \in X$ with $k = 1, \dots, n$, that are independent and identically distributed realizations from the density μ_m in (5.1).

The general principle of sequential conditional sampling is that any multivariate probability density $p(x_1, \dots, x_d)$ can be written as

$$p(x_1, \dots, x_d) = p_1(x_1)p_2(x_2|x_1) \cdots p_d(x_d|x_1, \dots, x_{d-1}), \tag{5.4}$$

using the conditional densities $p_k(x_k|x_1, \dots, x_{k-1})$ of x_k given $x_1, \dots, x_{k-1}$, that are defined as

$$p_k(x_k|x_1, \dots, x_{k-1}) := \frac{p_{1,\dots,k}(x_1, \dots, x_k)}{p_{1,\dots,k-1}(x_1, \dots, x_{k-1})}$$

using the marginal densities $p_{1,\dots,k}(x_1, \dots, x_k)$ and $p_{1,\dots,k-1}(x_1, \dots, x_{k-1})$ of p with respect to the first k and $k-1$ coordinates, respectively. Therefore, a random sample $x = (x_1, \dots, x_d)$ of p can be generated by first drawing x_1 according to $p_1(\cdot)$, and then x_k according to $p_k(\cdot|x_1, \dots, x_{k-1})$ for any $k = 2, \dots, d$. The method requires the explicit knowledge of all the marginals $p_{1,\dots,k}(\cdot|x_1, \dots, x_k)$ for any $k = 1, \dots, d-1$. This happens to be the case for the density μ_m in (5.1), whose marginals have been explicitly calculated in (5.3). We next discuss this method in more detail for our particular setting.

In the multivariate case the coordinates can be arbitrarily reordered. Start with the first coordinate x_1 and sample n points $x_1^1, \dots, x_1^n \in X_1$ from the univariate density

$$\varphi_1 : X_1 \to \mathbb{R} : t \mapsto \varphi_1(t) := \psi_1(t) = \frac{\rho_1(t)}{\#(\Lambda)} \sum_{\nu \in \Lambda} |\phi_{\nu_1}^1(t)|^2, \tag{5.5}$$

which coincides with the marginal ψ_1 of x_1 calculated in (5.3). In the univariate case $d = 1$ the algorithm terminates. In the multivariate case $d \geq 2$, by iterating q from 2 to d, consider the qth coordinate x_q, and sample n points $x_q^1, \dots, x_q^n \in X_q$ in the following way: for any $k = 1, \dots, n$, given the values $x_{A^{q-1}}^k = (x_1^k, \dots, x_{q-1}^k) \in X_{A^{q-1}}$ that have been calculated at the previous $q-1$ steps, sample the point $x_q^k \in X_q$ from the univariate density

$$\varphi_q : X_q \to \mathbb{R} : t \mapsto \varphi_q(t|x_{A^{q-1}}^k) := \rho_q(t) \frac{\sum_{\nu \in \Lambda} |\phi_{\nu_q}^q(t)|^2 \prod_{j=1}^{q-1} |\phi_{\nu_j}^j(x_j^k)|^2}{\sum_{\nu \in \Lambda} \prod_{j=1}^{q-1} |\phi_{\nu_j}^j(x_j^k)|^2}. \tag{5.6}$$

The expression on the right-hand side of (5.6) is continuous at any $t \in X_q$ and at any $x_{A^{q-1}}^k \in X_{A^{q-1}}$. Assumption 1.1 ensures that the denominator of (5.6) is strictly positive for any possible choice of $x_{A^{q-1}}^k = (x_1^k, \dots, x_{q-1}^k) \in X_{A^{q-1}}$, and also ensures that the marginal ψ_{q-1} is strictly positive at any point $x_{A^{q-1}}^k \in X_{A^{q-1}}$ such that $\rho_{A^{q-1}}(x_{A^{q-1}}^k) \neq 0$. For any $t \in X_q$ and any $x_{A^{q-1}}^k \in X_{A^{q-1}}$ such that $\rho_{A^{q-1}}(x_{A^{q-1}}^k) \neq 0$, the density φ_q satisfies

$$\varphi_q(t|x_{A^{q-1}}^k) = \frac{\psi_q(x_{A^{q-1}}^k, t)}{\psi_{q-1}(x_{A^{q-1}}^k)}, \tag{5.7}$$

where the densities ψ_q and ψ_{q-1} are the marginals defined in (5.2) and evaluated at the points $(x_{A^{q-1}}^k, t) \in X_{A^q}$ and $x_{A^{q-1}}^k \in X_{A^{q-1}}$, respectively. From (5.7), using (5.3) and simplifying the term $\rho_{A^{q-1}}(x_{A^{q-1}}^k) = \prod_{j=1}^{q-1} \rho_j(x_j^k) \neq 0$, one obtains the right-hand side of (5.6). The right-hand side of equation (5.7) is well defined for any $t \in X_q$ and any $x_{A^{q-1}}^k \in X_{A^{q-1}}$ such that $\rho_{A^{q-1}}(x_{A^{q-1}}^k) \neq 0$, and it is not defined at the points $x_{A^{q-1}}^k \in X_{A^{q-1}}$ such that $\rho_{A^{q-1}}(x_{A^{q-1}}^k) = 0$ where $\psi_{q-1}(x_{A^{q-1}}^k)$ vanishes. Nonetheless, (5.7) has finite limits at any point $(x_{A^{q-1}}^k, t) \in X_{A^q}$, and these limits equal expression (5.6).

According to technical terminology, the right-hand side of equation (5.7) is the conditional density of x_q given $x_1, \dots, x_{q-1}$ with respect to the density ψ_q, and φ_q is the continuous extension to X_{A^q} of this conditional density.

The densities $\varphi_1, \dots, \varphi_d$ defined in (5.5)–(5.6) can be concisely rewritten for any $q = 1, \dots, d$ as

$$\varphi_q(t|x^k_{A^{q-1}}) = \rho_q(t) \sum_{\nu \in \Lambda} \alpha_\nu(x^k_{A^{q-1}}) |\phi^q_{\nu_q}(t)|^2, \tag{5.8}$$

where the nonnegative weights $(\alpha_\nu)_{\nu \in \Lambda}$ are defined as

$$\alpha_\nu = \alpha_\nu(z_{A^{q-1}}) := \begin{cases} \dfrac{1}{\#(\Lambda)}, & \text{if } q = 1, \\ \dfrac{\prod_{j=1}^{q-1} |\phi^j_{\nu_j}(z_j)|^2}{\sum_{\widetilde{\nu} \in \Lambda} \prod_{j=1}^{q-1} |\phi^j_{\widetilde{\nu}_j}(z_j)|^2}, & \text{if } 2 \le q \le d, \end{cases}$$

for any $z_{A^{q-1}} = (z_1, \dots, z_{q-1}) \in X_{A^{q-1}}$. Since $\sum_{\nu \in \Lambda} \alpha_\nu = 1$, each density φ_q in (5.8) is a convex combination of the densities $\rho_q |\phi^q_1|^2, \dots, \rho_q |\phi^q_{\lambda_q}|^2$. Note that if the orthonormal basis $(\phi^q_j)_{j \ge 0}$ have explicit expressions and can be evaluated at any point in X_q, then the same holds for the univariate densities (5.8). In particular, in the polynomial case, for standards univariate densities ρ_i such as uniform, Chebyshev or Gaussian, the orthonormal polynomials $(\phi^i_j)_{j \ge 1}$ have expressions which are explicitly computable, for example by recursion formulas.

In Algorithm 1 we summarize our sampling method, that sequentially samples the univariate densities (5.8) to generate independent samples from the multivariate density (5.1). In the univariate case $d = 1$ the algorithm does not run the innermost loop, and only samples from φ_1. In the multivariate case $d \ge 2$ the algorithm runs also the innermost loop, and conditionally samples also from $\varphi_2, \dots, \varphi_d$. Our algorithm therefore relies on accurate sampling methods for the relevant univariate densities (5.8).

Algorithm 1 Sequential conditional sampling for μ_m.

INPUT: n, d, Λ, ρ_i, $(\phi^i_j)_{j \ge 0}$ for $i = 1, \dots, d$.

OUTPUT: $x^1, \dots, x^n \overset{\text{i.i.d.}}{\sim} \mu_m$.

for $k = 1$ to n **do**

  $\alpha_\nu \leftarrow (\#(\Lambda))^{-1}$, for any $\nu \in \Lambda$.

  Sample x^k_1 from $t \mapsto \varphi_1(t) = \rho_1(t) \sum_{\nu \in \Lambda} \alpha_\nu |\phi^1_{\nu_1}(t)|^2$.

  for $q = 2$ to d **do**

    $$\alpha_\nu \leftarrow \frac{\prod_{j=1}^{q-1} |\phi^j_{\nu_j}(x^k_j)|^2}{\sum_{\widetilde{\nu} \in \Lambda} \prod_{j=1}^{q-1} |\phi^j_{\widetilde{\nu}_j}(x^k_j)|^2}, \text{ for any } \nu \in \Lambda.$$

    Sample x^k_q from $t \mapsto \varphi_q(t) = \rho_q(t) \sum_{\nu \in \Lambda} \alpha_\nu |\phi^q_{\nu_q}(t)|^2$.

  end for

  $x^k \leftarrow (x^k_1, \dots, x^k_d)$.

end for

5.2. **Sampling the univariate densities**

We next discuss two possible methods for sampling from such densities: *rejection sampling* and *inversion transform sampling*, see *e.g.* [5]. Both methods equally apply to any univariate density φ_q, and therefore we present them for any q arbitrarily chosen from 1 to d.

Rejection sampling (RS). For applying this method, one needs to find a suitable univariate density Θ_q, whose support contains the support of φ_q, and a suitable real $M_q > 1$ such that

$$\varphi_q(t) \leq M_q \Theta_q(t), \quad t \in \mathrm{supp}(\varphi_q).$$

The density Θ_q should be easier to sample than φ_q, *i.e.* efficient pseudorandom number generators for sampling from Θ_q are available. The value of M_q should be the smallest possible. For sampling one point from φ_q using RS: sample a point z from Θ_q, and sample u from the standard uniform $\mathcal{U}(0,1)$. Then check if $u < \varphi_q(z)/M_q\Theta_q(z)$: if this is the case then accept z as a realization from φ_q, otherwise reject z and restart sampling z and u from beginning. On average, acceptance occurs once every M_q trials. Therefore, for a given q, sampling one point from φ_q by RS requires on average M_q evaluations of the function

$$t \mapsto \frac{\varphi_q(t)}{M_q\Theta_q(t)} = \frac{\rho_q(t)}{M_q\Theta_q(t)} \sum_{\nu \in \Lambda} \alpha_\nu |\phi^q_{\nu_q}(t)|^2.$$

This amounts to evaluating M_q times a subset of the terms $\phi^q_0, \ldots, \phi^q_{\lambda_q}$. The coefficients α_ν depend on the terms $\phi^j_0, \ldots, \phi^j_{\lambda_j}$ for $j = 1, \ldots, q-1$, which have been already evaluated when sampling the previous coordinates $1, \ldots, q-1$. Thus, if we use RS for sampling the univariate densities, the overall computational cost of Algorithm 1 for sampling n points $x^1, \ldots, x^n \in X$ is on average proportional to $n \sum_{q=1}^d M_q(\lambda_q + 1)$.

When the basis functions $(\phi^q_j)_{j \geq 0}$ form a bounded orthonormal system, an immediate and simple choice of the parameters in the algorithm is

$$M_q = \max_{\nu \in \Lambda} \|\phi^q_{\nu_q}\|^2_{L^\infty}, \quad \text{and} \quad \Theta_q(t) = \rho_q(t). \tag{5.9}$$

With such a choice, we can quantify more precisely the average computational cost of sampling n points in dimension d. When $(\phi^q_j)_{j \geq 0}$ are the Chebyshev polynomials, whose L^∞ norms satisfy $\|\phi^q_j\|_{L^\infty} \leq \sqrt{2}$, we obtain the bound $2n \sum_{q=1}^d (\lambda_q + 1) \leq 2nd(\lambda_\Lambda + 1) \leq 2ndm$. When $(\phi^q_j)_{j \geq 0}$ are the Legendre polynomials, whose L^∞ norms satisfy $\|\phi^q_j\|_{L^\infty} \leq \sqrt{2j+1}$, we have the crude estimate $2n \sum_{q=1}^d (\lambda_q + 1)^2 \leq 2nd(\lambda_\Lambda + 1)^2 \leq 2ndm^2$. In general, when $(\phi^q_j)_{j \geq 0}$ are Jacobi polynomials, similar upper bounds can be derived, and the dependence of these bounds on n and d is linear.

Inversion transform sampling (ITS). Let $\Phi_q : X_q \to [0,1]$ be the cumulative distribution function associated to the univariate density φ_q. In the following, only when using the ITS method, we make the further assumption that ρ_q vanishes at most a finite number of times in X_q. Such an assumption is fulfilled in many relevant situations, *e.g.* when ρ_q is the density associated to Jacobi or Hermite polynomials orthonormal in $L^2(X_q, d\rho_q)$. Together with Assumption 1.1, this ensures that the function $t \mapsto \Phi_q(t)$ is continuous and strictly increasing on X_q. Hence Φ_q is a bijection between X_q and $[0,1]$, and it has a unique inverse $\Phi_q^{-1} : [0,1] \to X_q$ which is continuous and strictly increasing on $[0,1]$. Sampling from φ_q using ITS can therefore be performed as follows: sample n independent realizations $u^1, \ldots, u^n$ identically distributed according to the standard uniform $\mathcal{U}(0,1)$, and obtain the n independent samples from φ_q as $(\Phi_q^{-1}(u^1), \ldots, \Phi_q^{-1}(u^n))$.

For any $u \in [0,1]$, computing $z = \Phi_q^{-1}(u) \in X_q$ is equivalent to find the unique solution $z \in X_q$ to $\Phi_q(z) = u$. This can be executed by elementary root-finding numerical methods, *e.g.* the

bisection method or Newton's method. As an alternative to using root-finding methods, one can build an interpolant operator $\mathcal{I}_q$ of Φ_q^{-1}, and then approximate $\Phi_q^{-1}(u) \approx \mathcal{I}_q(u)$ for any $u \in [0,1]$. Such an interpolant $\mathcal{I}_q$ can be constructed for example by piecewise linear interpolation, from the data $(\Phi_q(t_1^q), t_1^q), \ldots, (\Phi_q(t_{s_q}^q), t_{s_q}^q)$ at s_q suitable points $t_1^q < \ldots < t_{s_q}^q$ in X_q.

Both root-finding methods and the interpolation method require evaluating the function Φ_q pointwise in X_q. In general these evaluations can be computed using standard univariate quadrature formulas. When $(\phi_j^q)_{j\geq 0}$ are orthogonal polynomials, the explicit expression of the primitive of φ_q can be used for directly evaluating the function Φ_q.

Finally we discuss the overall computational cost of Algorithm 1 for sampling n points $x^1, \ldots, x^n \in X$ when using ITS for sampling the univariate densities. With the bisection method, this overall cost amounts to $n \sum_{q=1}^d \gamma_q W_q$, where γ_q is the maximum number of iterations for locating the zero in X_q up to some desired tolerance, and W_q is the computational cost of each iteration. With the interpolation of Φ_q^{-1}, the overall cost amounts to n evaluations of each interpolant $\mathcal{I}_q$, in addition to the cost of building the interpolants which does not depend on n.

5.3. **Mixture sampling**

An alternative to sequential conditional sampling in Algorithm 1 can be developed by exploiting again the structure of μ_m in (5.1). The density μ_m is a convex combination with equal weights $\frac{1}{m} = \frac{1}{\#(\Lambda)}$ of the $\#(\Lambda)$ densities $\rho|L_\nu|^2 = \prod_{q=1}^d \rho_q |\phi_{\nu_q}^q|^2$ for $\nu \in \Lambda$. The following Algorithm 2 exploits this additive mixture structure. It starts by drawing n independent random indices $\nu^1, \ldots, \nu^n \overset{\text{i.i.d.}}{\sim} \mathcal{U}(\Lambda)$, that is, we randomly choose each ν^j among the elements of Λ with equal probability $\frac{1}{\#(\Lambda)}$. These random variables select the densities $\rho|L_{\nu^1}|^2, \ldots, \rho|L_{\nu^n}|^2$ in the mixture, and each of these densities has product form $\rho|L_{\nu^k}|^2 = \prod_{q=1}^d \rho_q |\phi_{\nu_q^k}^q|^2$. For any q, sampling from the univariate density $\rho_q|\phi_{\nu_q^k}^q|^2$ can be done for example with the rejection sampling method described in section §5.2, by choosing a density Θ_q for which an efficient pseudorandom generator is available and a suitable $M_q > 1$ such that

$$\rho_q(t)|\phi_{\nu_q^k}^q(t)|^2 \leq M_q \Theta_q(t), \quad t \in X_q.$$

In general, sampling from the univariate densities $\rho_q|\phi_j^q|^2$ might not be an easy task: for example, when $(\phi_j^q)_{j\geq 0}$ are orthogonal polynomials, the density $\rho_q|\phi_j^q|^2$ vanishes (at least) at $2j$ points giving rise to large oscillations. This is the main drawback of Algorithm 2, in contrast to Algorithm 1 that prevents this effect by sampling from the density φ_q where the densities $\rho_q|\phi_1^q|^2, \ldots, \rho_q|\phi_{\lambda_q}^q|^2$ are piled up. The zeros of φ_q are drastically reduced in number (or even completely removed), thanks to the interlacing property of the zeros of orthogonal polynomials. An immediate consequence of reducing the number of zeros of φ_q is that the oscillations are also reduced in number and amplitude.

Remark 5.1. The computational cost of the various algorithms that we have proposed in this section for sampling μ_m is generally higher than the cost of sampling the product density ρ. However, this difference in the cost is negligible in frequently encountered situations where the computational cost of the pointwise evaluations of the function is several orders of magnitude larger than the computational cost of sampling the densities. For example, in the context of numerical approximation of parametric PDEs, each pointwise evaluation amounts to running a potentially costly finite element/difference/volume solver. In such cases, weighted least squares provide a substantial saving on the overall number of samples, at the price of a negligibly more costly generation of the samples.

Algorithm 2 Mixture sampling from μ_m.

INPUT: n, d, Λ, ρ_i, $(\phi_j^i)_{j\geq 0}$ for $i = 1, \dots, d$.

OUTPUT: $x^1, \dots, x^n \overset{\text{i.i.d.}}{\sim} \mu_m$.

```
for k = 1 to n do
    Sample ν^k from U(Λ)
    for q = 1 to d do
        Sample x_q^k from t ↦ ρ_q(t)|φ^q_{ν_q^k}(t)|^2.
    end for
    x^k ← (x_1^k, ..., x_d^k).
end for
```

6. Examples and numerical illustrations

This section presents the numerical performances of the weighted least-squares method compared to the standard least-squares method, in three relevant situations where $d\rho$ can be either the uniform measure, the Chebyshev measure, or the Gaussian measure. In each of these three cases, we choose w and $d\mu$ in the weighted least-squares method from (2.6) and (2.7), as prescribed by our analysis in Corollary 2.2. For standard least squares we choose w and $d\mu$ as in (1.9). Our tests focus on the condition number of the Gramian matrix, that quantifies the stability of the linear system (1.5) and the stability of the weighted and standard least-squares estimators. A meaningful quantity is therefore the probability

$$\Pr\{\text{cond}(\mathbf{G}) \leq 3\}, \tag{6.1}$$

where, through (1.8), the value three of the threshold is related to the parameter $\delta = 1/2$ in the previous analysis. For any n and m, from (1.8) the probability (6.1) is larger than $\Pr\{\|\mathbf{G} - \mathbf{I}\|_2 \leq \frac{1}{2}\}$. From Corollary 2.2, under condition (2.8) between n, m and r, the Gramian matrix of weighted least squares satisfies (2.2) and therefore the probability (6.1) is larger than $1 - 2n^{-r}$. For standard least squares, from Theorem 1.1 the Gramian matrix satisfies (6.1) with probability larger than $1 - 2n^{-r}$, but under condition (1.11).

In all the presented numerical tests the probability (6.1) is numerically approximated by its empirical counterpart, obtained by counting how many times the event $\text{cond}(\mathbf{G}) \leq 3$ occurs when repeating the random sampling one hundred times.

All the examples presented in this section are confined to multivariate approximation spaces of polynomial type. One natural assumption in this case is to require that the set Λ is *downward closed*, that is, satisfies

$$\nu \in \Lambda \quad \text{and} \quad \tilde{\nu} \leq \nu \implies \tilde{\nu} \in \Lambda,$$

where $\tilde{\nu} \leq \nu$ means that $\tilde{\nu}_j \leq \nu_j$ for all $i = 1, \dots, d$. Then V_m is the polynomial space spanned by the monomials

$$z \mapsto z^\nu := \prod_{j=1}^{d} z_j^{\nu_j},$$

and the orthonormal basis L_ν is provided by taking each $(\phi_j^i)_{j\geq 0}$ to be a sequence of univariate orthonormal polynomials of $L^2(X_i, d\rho_i)$.

In both the univariate and multivariate forthcoming examples, the random samples from the measure $d\mu_m$ are generated using Algorithm 1. The univariate densities $\varphi_1, \dots, \varphi_d$ are sampled using the

inversion transform sampling method. The inverse of the cumulative distribution function is approximated using the interpolation technique.

6.1. **Univariate examples**

In the univariate case $d = 1$, let the index set be $\Lambda = \{0, \dots, m-1\}$ and $V_m = \mathbb{P}_\Lambda = \text{span}\{z^k : k = 0, \dots, m-1\}$. We report in Fig. 6.1 the probability (6.1), when $\mathbf{G}$ is the Gramian matrix of the weighted least-squares method. Different combinations of values for m and n are tested, with three choices of the measure $d\rho$: uniform, Gaussian and Chebyshev. As in further figures, the empirically approximated probability is represented by the color level from black (0) to white (1). The results do not show perceivable differences among the performances of weighted least squares with the three different measures. In any of the three cases, $n/\ln(n) \geq 4m$ is enough to obtain a probability equal to one that $\text{cond}(\mathbf{G}) \leq 3$. This confirms that condition (2.8) with any choice of $r > 0$ ensures (6.1), since it demands for a larger number of samples.

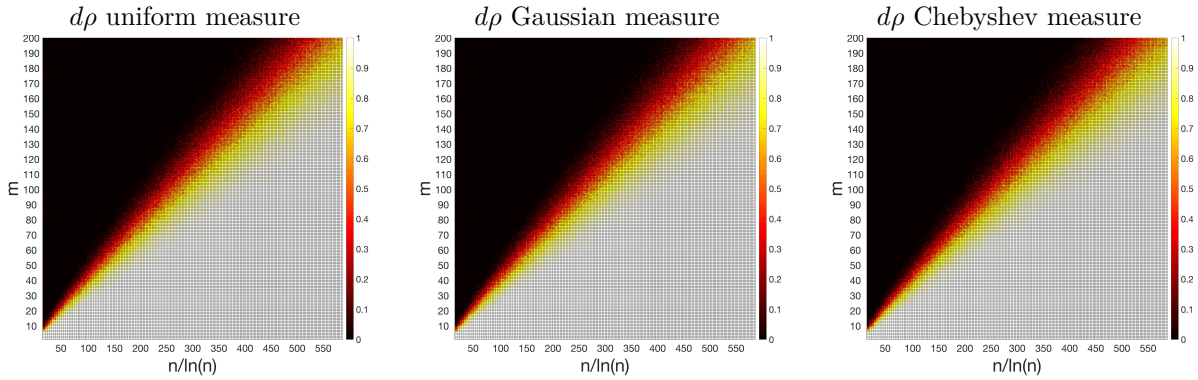

Figure 6.1. Weighted least squares, $Pr\{\text{cond}(G) \leq 3\}$, $d = 1$. Left: $d\rho$ uniform measure. Center: $d\rho$ Gaussian measure. Right: $d\rho$ Chebyshev measure.

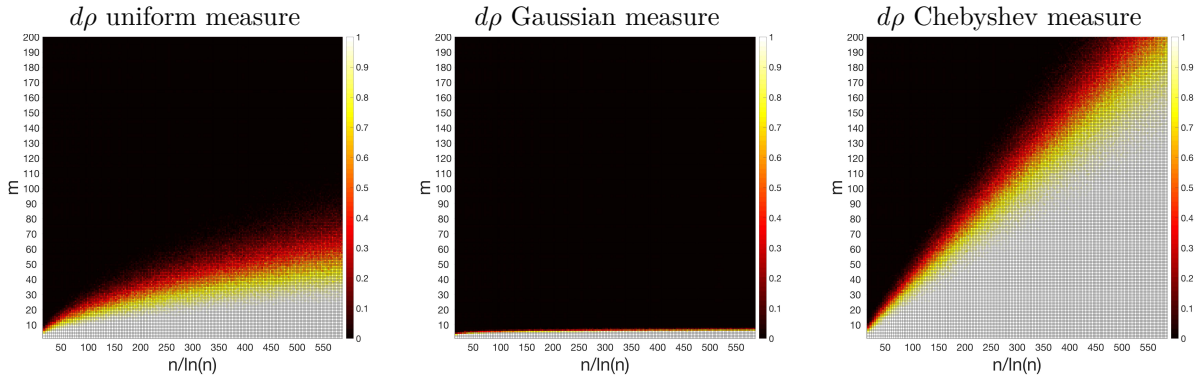

Figure 6.2. Standard least squares, $Pr\{\text{cond}(G) \leq 3\}$, $d = 1$. Left: $d\rho$ uniform measure. Center: $d\rho$ Gaussian measure. Right: $d\rho$ Chebyshev measure.

Fig. 6.2 shows the probability (6.1) when $\mathbf{G}$ is the Gramian matrix of standard least squares. With the uniform measure, the condition $n/\ln(n) \geq m^2$ is enough to have (6.1) with empirical probability larger than 0.95. When $d\rho$ is the Gaussian measure, stability requires a very large number

of evaluations, roughly $n/\ln(n)$ linearly proportional to $\exp(m/3)$. For the univariate Chebyshev measure, it is proven that standard least squares are stable under the same minimal condition (2.8) as for weighted least squares. In accordance with the theory, the numerical results obtained in this case with weighted and standard least squares are indistinguishable, see Fig. 6.1-right and Fig. 6.2-right.

6.2. **Multivariate examples**

Afterwards we present some numerical tests in the multivariate setting. In dimension d larger than one there are many possible ways to enrich the polynomial space $\mathbb{P}_\Lambda$. The number of different downward closed sets whose cardinality equals m gets very large already for moderate values of m and d. Therefore, in our numerical results, for a given dimension d, we first randomly generate a particular sequence $\Lambda_1 \subset \cdots \subset \Lambda_m$, where each $\Lambda_j \subset \mathbb{N}_0^d$ is downward closed, $\#(\Lambda_j) = \dim(\mathbb{P}_{\Lambda_j}) = j$ and the starting set Λ_1 contains only the null multi-index. More specifically, given Λ_k, we take $\Lambda_{k+1} = \Lambda_k \cup \{\nu^k\}$, where ν^k is randomly selected among the finitely many elements $\nu \notin \Lambda_k$ such that $\Lambda_k \cup \{\nu\}$ remains downward closed. Once such a sequence is fixed, the tests in Fig. 6.3 and Fig. 6.4 are performed using the embedded polynomial spaces $\mathbb{P}_{\Lambda_1} \subset \ldots \subset \mathbb{P}_{\Lambda_m}$, for both weighted and standard least squares and for the three choices of the measures $d\rho$. Such a choice allows us to establish a fair comparison between the two methods and among different measures, without the additional variability arising from modifications to the polynomial space. We comment further on the influence of the chosen sequence $\Lambda_1 \subset \cdots \subset \Lambda_m$.

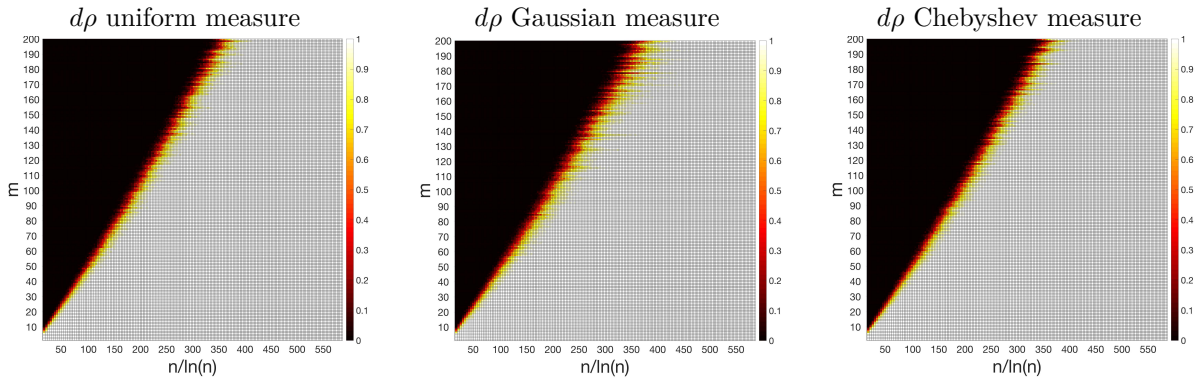

FIGURE 6.3. Weighted least squares, $Pr\{\text{cond}(G) \leq 3\}$, $d = 10$. Left: $d\rho$ uniform measure. Center: $d\rho$ Gaussian measure. Right: $d\rho$ Chebyshev measure.

We first report the results obtained for the tests in dimension $d = 10$. The results in Fig. 6.3 confirm that weighted least squares always yield an empirical probability equal to one that $\text{cond}(\mathbf{G}) \leq 3$, provided that $n/\log(n) \geq 2m$. This condition ensures that (2.8) with any choice of $r > 0$ implies (6.1), thus verifying Corollary 2.2. Again, the results do not show significant differences among the three choices of the measure $d\rho$: a straight line, with the same slope for all the three cases uniform, Chebyshev and Gaussian, separates the two regimes corresponding to empirical probabilities equal to zero and one. Compared to the univariate case in Fig. 6.1, the results in Fig. 6.3 exhibit a sharper transition between the two extreme regimes, and an overall lower variability in the transition regime.

The results for standard least squares with $d = 10$ are shown in Fig. 6.4. In the case of the uniform measure, in Fig. 6.4-right, stability is ensured if $n/\ln(n) \geq 3.5m$, which is more demanding than the condition $n/\ln(n) \geq 2m$ needed for the stability of weighted least squares in Fig. 6.3-right, but much less strict than the condition required with standard least squares in the univariate case, where $n/\ln(n)$ scales like m^2. These phenomena have already been observed and described in [11]. Similar results

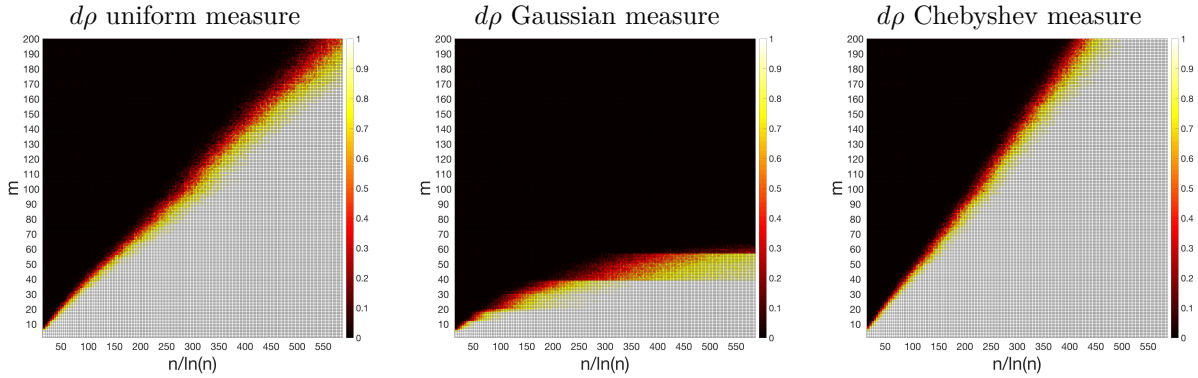

Figure 6.4. Standard least squares, $Pr\{\text{cond}(G) \leq 3\}$, $d = 10$. Left: $d\rho$ uniform measure. Center: $d\rho$ Gaussian measure. Right: $d\rho$ Chebyshev measure.

as those with the uniform measure are obtained with the Chebyshev measure in Fig. 6.4-left, where again standard least squares achieve stability using more evaluations than weighted least squares in Fig. 6.3-left. The case of the Gaussian measure drastically differs from the uniform and Chebyshev cases: the results in Fig. 6.4-center clearly indicate that a very large number of evaluations n compared to m is required to achieve stability of standard least squares.

Analogous results as those presented in Figs. 6.1 and 6.3 for weighted least squares have been obtained also in other dimensions, and with many other sequences of increasingly embedded polynomial spaces. In the next tables we report some of these results for selected values of $d = 1, 2, 5, 10, 50, 100$. We choose $n = 26599$ and $m = 200$ that satisfy condition (2.8) with $r = 1$, and report in Table 6.1 the empirical probabilities that approximate (6.1), again calculated over one hundred repetitions. This table provides multiple comparisons: weighted least squares versus standard least squares, for the three choices of the measure $d\rho$ (uniform, Gaussian and Chebyshev) and with d varying between 1 and 100.

method	$d\rho$	$d = 1$	$d = 2$	$d = 5$	$d = 10$	$d = 50$	$d = 100$
weighted LS	uniform	1	1	1	1	1	1
weighted LS	Gaussian	1	1	1	1	1	1
weighted LS	Chebyshev	1	1	1	1	1	1
standard LS	uniform	0	0	0.54	1	1	1
standard LS	Gaussian	0	0	0	0	0	0
standard LS	Chebyshev	1	1	1	1	1	1

Table 6.1. $Pr\{\text{cond}(G) \leq 3\}$, with $n = 26559$ and $m = 200$: weighted least squares versus standard least squares, $d\rho$ uniform versus $d\rho$ Gaussian versus $d\rho$ Chebyshev, $d = 1, 2, 5, 10, 50, 100$.

In Table 6.1, all the empirical probabilities related to results for weighted least squares are equal to one, and confirm the theory since, for the chosen values of n, m and r, the probability (6.1) is larger than $1 - 5.67 \times 10^{-7}$. This value is computed using estimate (3.1) from the proof of Theorem 2.1. In contrast to weighted least squares, whose empirical probability equal one independently of $d\rho$ and d, the empirical probability of standard least squares does depend on the chosen measure, and to some extent on the dimension d as well. With the uniform measure, the empirical probability that approximates (6.1) equals zero when $d = 1$ or $d = 2$, equals 0.54 when $d = 5$, and equals one when $d = 10$, $d = 50$ or $d = 100$. In the Gaussian case, standard least squares always feature null empirical

method	$d\rho$	$d=1$	$d=2$	$d=5$	$d=10$	$d=50$	$d=100$
weighted LS	uniform	1.5593	1.4989	1.4407	1.4320	1.4535	1.4179
weighted LS	Gaussian	1.5994	1.5698	1.4743	1.4643	1.4676	1.4237
weighted LS	Chebyshev	1.5364	1.4894	1.4694	1.4105	1.4143	1.4216
standard LS	uniform	19.9584	29.8920	3.0847	1.9555	1.7228	1.5862
standard LS	Gaussian	$\sim 10^{19}$	$\sim 10^{19}$	$\sim 10^{19}$	$\sim 10^{16}$	$\sim 10^{9}$	$\sim 10^{3}$
standard LS	Chebyshev	1.5574	1.5367	1.5357	1.4752	1.4499	1.4625

TABLE 6.2. Average of cond(G), with $n = 26559$ and $m = 200$: weighted least squares versus standard least squares, $d\rho$ uniform versus $d\rho$ Gaussian versus $d\rho$ Chebyshev, $d = 1, 2, 5, 10, 50, 100$.

probabilities. With the Chebyshev measure, the condition number of $\mathbf{G}$ for standard least squares is always lower than three for any tested value of d.

In addition to the results in Table 6.1, further information is needed for assessing how severe is the lack of stability when obtaining null empirical probabilities. To this aim, in Table 6.2 we also report the average value of cond($\mathbf{G}$), obtained when averaging the condition number of $\mathbf{G}$ over the same repetitions used to estimate the empirical probabilities in Table 6.1. The information in Table 6.2 is complementary to that in Table 6.1. On the one hand they point out the stability and robustness of weighted least squares, showing a tamed condition number with any measure $d\rho$ and any dimension d. On the other hand they provide further insights on stability issues of standard least squares and their dependence on $d\rho$ and d.

For standard least squares with the uniform measure, the average condition number reduces as the dimension d increases, in agreement with the conclusion drawn from Table 6.1. One possible explaination of this phenomenon is the following: while $K_m(P_\Lambda, d\rho)$ is known to satisfy the bound

$$K_m(\mathbb{P}_\Lambda, d\rho) \leq m^2, \tag{6.2}$$

for all downward closed sets Λ of cardinality m and in any dimension d, equality in this bound is only attained for certain sets Λ. In particular, it is attained for the sets Λ of rectangular shape, that is

$$\Lambda := \{\nu \,:\, \nu \leq \mu\}, \tag{6.3}$$

for some $\mu = (\mu_1, \ldots, \mu_d) \in \mathbb{N}_0^d$ such that $\prod_{j=1}^d (1+\mu_j) = m$. However, as d gets larger, the typical value of $K_m(\mathbb{P}_\Lambda, d\rho)$ may be significantly smaller for a general downward closed set, which is the case for our randomly generated sequence $\Lambda_1 \subset \cdots \subset \Lambda_m$. The Gramian matrix of standard least squares with the Gaussian measure is very ill-conditioned for all tested values of d, with again a reduction as d gets large. For standard least squares with the Chebyshev measure, the averaged condition number of $\mathbf{G}$ is only slightly larger than the one for weighted least squares.

As explained above, the results for standard least squares in Fig. 6.4, Table 6.1 and Table 6.2 are sensitive to the chosen sequence of polynomial spaces. Testing different sequences might produce different results, that however necessarily obey to the estimates proven in Theorem 1.1 with uniform and Chebyshev measures, when n, m and r satisfy condition (1.11). Many other examples with standard least squares have been extensively discussed in previous works *e.g.* [11, 3], also in situations where n, m and r do not satisfy condition (1.11) and therefore Theorem 1.1 does not apply. In general, when n, m and r do not satisfy (1.11), there exist multivariate polynomial spaces of dimension m such that the Gramian matrix of standard least squares with the uniform and Chebyshev measures does not satisfy (1.12). Examples of such spaces are discussed in [11, 3]. Using these spaces would yield null empirical probabilities in Table 6.1 for standard least squares with the uniform and Chebyshev measures.

For weighted least squares, when n, m and r satisfy condition (2.8), any sequence of polynomial spaces yields empirical probabilities close to one, according to Corollary 2.2. Indeed such a robustness with respect to the choices of $d\rho$, of the polynomial space and of the dimension d represents one of the main advantages of the weighted approach.

References

[1] B. Adcock and A. C. Hansen. A generalized sampling theorem for stable reconstructions in arbitrary bases. *J. Fourier Anal. Appl.*, 18:685–716, 2012.

[2] G. Chardon, A. Cohen, and L. Daudet. Sampling and reconstruction of solutions to the helmholtz equation. *Sampl. Theory Signal Image Process.*, 13:67–89, 2014.

[3] A. Chkifa, A. Cohen, G. Migliorati, F. Nobile, and R. Tempone. Discrete least squares polynomial approximation with random evaluations - application to parametric and stochastic elliptic PDEs. *M2AN*, 49:815–837, 2015.

[4] A. Cohen, M. A. Davenport, and D. Leviatan. On the stability and accuracy of least squares approximations. *Found. Comput. Math.*, 13:819–834, 2013.

[5] L. Devroye. *Non-Uniform Random Variate Generation*. Springer, 1985.

[6] A. Doostan and J. Hampton. Coherence motivated sampling and convergence analysis of least squares polynomial chaos regression. *Comput. Methods Appl. Mech. Engrg.*, 290:73–97, 2015.

[7] J. D. Jakeman, A. Narayan, and T. Zhou. A Christoffel function weighted least squares algorithm for collocation approximations. *Math. Comp.*, 86:1913–1947, 2017.

[8] A. Máté, P. Nevai, and V. Totik. Szegö's extremum problem on the unit circle. *Annals of Mathematics,*, 134:433–453, 1991.

[9] G. Migliorati. Multivariate Markov-type and Nikolskii-type inequalities for polynomials associated with downward closed multi-index sets. *J. Approx. Theory*, 189:137–159, 2015.

[10] G. Migliorati, F. Nobile, and R. Tempone. Convergence estimates in probability and in expectation for discrete least squares with noisy evaluations at random points. *J. Multivar. Analysis*, 142:167–182, 2015.

[11] G. Migliorati, F. Nobile, E. von Schwerin, and R. Tempone. Analysis of discrete L^2 projection on polynomial spaces with random evaluations. *Found. Comput. Math.*, 14:419–456, 2014.

[12] P. Nevai. Géza Freud, orthogonal polynomials and Christoffel Functions. A case study. *J. Approx. theory*, 48:3–167, 1986.

[13] E. B. Saff and V. Totik. *Logarithmic Potentials with External Fields*. Springer, 1997.

[14] J. Tropp. User friendly tail bounds for sums of random matrices. *Found. Comput. Math.*, 12:389–434, 2012.