

Multilevel weighted least squares polynomial approximation

Abdul-Lateef Haji-Ali^a, Fabio Nobile^b, Raúl Tempone^c, Sören Wolfers^{c*}

^a Oxford University

^b École Polytechnique Fédérale de Lausanne (EPFL), CSQI-MATH

^c King Abdullah University of Science and Technology (KAUST), Applied Mathematics and Computational Science

September 11, 2017

Abstract

Weighted least squares polynomial approximation uses random samples to determine projections of functions onto spaces of polynomials. It has been shown that, using an optimal distribution of sample locations, the number of samples required to achieve quasi-optimal approximation in a given polynomial subspace scales, up to a logarithmic factor, linearly in the dimension of this space. However, in many applications, the computation of samples includes a numerical discretization error. Thus, obtaining polynomial approximations with a single level method can become prohibitively expensive, as it requires a sufficiently large number of samples, each computed with a sufficiently small discretization error. As a solution to this problem, we propose a multilevel method that utilizes samples computed with different accuracies and is able to match the accuracy of single-level approximations with reduced computational cost. We derive complexity bounds under certain assumptions about polynomial approximability and sample work. Furthermore, we propose an adaptive algorithm for situations where such assumptions cannot be verified a priori. Finally, we provide an efficient algorithm for the sampling from optimal distributions and an analysis of computationally favorable alternative distributions. Numerical experiments underscore the practical applicability of our method.

Keywords Multilevel methods, least squares approximation, multivariate approximation, polynomial approximation, convergence rates, error analysis

Mathematics Subject Classification (2010) 41A10, 41A25, 41A63, 65B99, 65D10, 65N12, 65N15, 65N22, 65N35

1 Introduction

A common goal in uncertainty quantification [21] is the approximation of response surfaces

$$\gamma \mapsto f(\gamma) := Q(u_\gamma) \in \mathbb{R},$$

which describe how a quantity of interest Q of the solution u_γ to some partial differential equation (PDE) depends on parameters $\gamma \in \Gamma \subset \mathbb{R}^d$ of the PDE. The non-intrusive approach to this problem is to evaluate the response surface for finitely many values of γ and then to use an interpolation method, such as (tensor-)spline interpolation [8], kernel-based approximation (kriging) [29, 13], or (global) polynomial approximation [21].

In this work, we study a variant of polynomial approximation in which least squares projections onto finite-dimensional polynomial subspaces are computed using values of f at finitely many

*Corresponding author. Email address: `soeren.wolfers@kaust.edu.sa`

random locations. More specifically, given a probability measure μ on the parameter space Γ and a polynomial subspace $V \subset L^2_\mu(\Gamma)$, the approximating polynomial is determined as

$$\Pi_V f := \arg \min_{v \in V} \|f - v\|_N, \quad (1)$$

where $\|\cdot\|_N$ is a discrete approximation of the $L^2_\mu(\Gamma)$ norm that is based on evaluations in finitely many randomly chosen sample locations $\gamma_j \in \Gamma$, $j \in \{1, \dots, N\}$ and a weight function $w: \Gamma \rightarrow \mathbb{R}$.

The case where equally weighted samples are drawn independently and identically distributed from the underlying probability measure itself, $\gamma_j \sim \mu$, has been popular among practitioners for a long time and has been given a thorough theoretical foundation in the past decade [3, 5, 26]. More recently, the use of alternative sampling distributions and non-constant weights was studied in [27, 7, 16]. In particular, [16] presented a sampling distribution ν_V^* and a corresponding weight function for which the number of samples required to determine quasi-optimal approximations within V is bounded by $\dim V$ up to a logarithmic factor. (This result was proved in [16] for total degree polynomial spaces and generalized in [7] to more general function spaces.) Since this distribution depends on V , it is natural to ask how samples can be efficiently obtained from it and whether there is an alternative that works equally well for all polynomial subspaces V . To address the first question, we present and analyze an efficient algorithm to generate samples from ν_V^* in the case where Γ is a product domain and μ is a product measure. For more general cases, we also study Markov chain methods for sample generation and analyze the effect of small perturbations of the sampling distribution on the convergence estimates of [16, 7]. To address the second question, we provide upper and lower bounds on ν_V^* in the case where Γ is a hypercube. The lower bound allows us to make the error estimates obtained in [7] more explicit. The upper bound shows that the arcsine distribution, which was proposed in [27], performs just as well as ν_V^* up to a constant that is independent of V . This is advantageous for adaptive algorithms in which the polynomial subspace and the corresponding optimal sampling distribution vary during the iterations. However, the constant mentioned above does depend exponentially on the dimension d of Γ .

To motivate the main contribution of this work, namely the multilevel weighted least squares polynomial approximation method, we note that the response surface f from the beginning of this introduction cannot be evaluated exactly. Indeed, in most cases, the computation of $Q(u_\gamma)$ requires the numerical solution of a PDE. Thus, we can only compute approximations of f whose accuracy and computational work are determined by the PDE discretization. If we simply applied polynomial least squares approximation using a sufficiently fine discretization of the PDE for all evaluations, then we would quickly face prohibitively long runtimes. For this reason, we introduce a multilevel method that combines lots of cheap samples using coarse discretizations with relatively few more expensive samples using fine discretizations of the PDE. In the recent decade, such multilevel algorithms have been studied intensely for the approximation of expectations [17, 19, 20, 14]. The goal of this paper is to extend this earlier work to the reconstruction of the full response surface, using global polynomial approximation and estimating the resulting error in the L^2_μ norm.

To describe the multilevel method, assume that we want to approximate a function f . Assume furthermore that we can only evaluate functions f_l with $f_l \rightarrow f$ as $l \rightarrow \infty$ in a suitable sense and that the cost per evaluation increases as $l \rightarrow \infty$. A straightforward approach to this situation is to apply least squares approximation to some f_L that is sufficiently close to f . The theory of (weighted) polynomial least squares approximation then provides conditions on the number of samples required to achieve quasi-optimal approximation of f_L within a given space of polynomials V_L . However, this approach can be computationally expensive, as each evaluation of f_L requires the numerical solution of a PDE using a fine discretization. As an alternative, our proposed multilevel algorithm starts out with a least squares approximation of f_0 using a relatively large polynomial subspace V_0 and correspondingly many samples. To correct for the committed error $f - f_0$, the algorithm then adds polynomial approximations of $f_l - f_{l-1}$ that lie in subspaces V_l , $l \in \{1, \dots, L\}$. Since we assume that $f_l \rightarrow f$ in an appropriate sense, the differences $f_l - f_{l-1}$ may be approximated using smaller polynomial subspaces for $l \rightarrow \infty$. Exploiting this fact, it is possible to obtain approximations with significantly reduced computational work. Indeed, we show that under certain conditions the work that the multilevel method requires to attain an accuracy of $\epsilon > 0$ is the same as the work that

regular least squares polynomial approximation would require if f could be evaluated exactly. It is clear that such a result is not always possible. For example, if f were constant, then polynomial least squares approximations in any fixed polynomial subspace would yield the exact solution given a sufficiently large sample size. This means that the work required to achieve an accuracy $\epsilon > 0$ would be bounded as $\epsilon \rightarrow 0$, which can clearly not be true for an algorithm that uses evaluations from approximate functions f_l that become more expensive to evaluate as $l \rightarrow \infty$. Instead, the optimal computational work required for an accuracy of $\epsilon > 0$ by such an algorithm would be determined by the convergence of $f_l \rightarrow f$ and by the work that is required for evaluations of f_l . Our results show that for many problems, the two cases described above are dichotomous. This means that the total computational work is determined either by the convergence and work associated with f_l or by the convergence of polynomial least squares approximation using exact evaluations (see Theorem 4.1 for a more formal statement).

The remainder of this paper is structured as follows. In Section 2, we review the theoretical analysis of weighted least squares approximation. In Section 3, we discuss different sampling strategies. We propose algorithms to sample the optimal distribution and we discuss the consequences of using perturbed distributions. In Section 4, we introduce a novel multilevel algorithm and prove our main results concerning the work and convergence of this algorithm. For situations in which the regularity of f and the convergence of f_l are not known, we propose an adaptive algorithm in Section 5. We discuss the applicability of our method to problems in uncertainty quantification in Section 6. Finally, we present numerical experiments in Section 7.

2 Weighted least squares polynomial approximation

In this section, we provide a short summary of the theory of weighted discrete least squares polynomial approximation, closely following [7].

Assume that we want to approximate a function $f \in L^2_\mu(\Gamma)$, where $\Gamma \subset \mathbb{R}^d$ and μ is a probability measure on Γ . The strategy of weighted discrete least squares polynomial approximation is to

- choose a finite-dimensional space $V \subset L^2_\mu(\Gamma)$ of polynomials on Γ
- choose a function $\rho: \Gamma \rightarrow \mathbb{R}$ that satisfies $\int_\Gamma \rho(\gamma)\mu(d\gamma) = 1$ and $\rho > 0$
- generate $N > 0$ independent random samples from the *sampling distribution* ν defined by $\frac{d\nu}{d\mu} := \rho$,

$$\gamma_j \sim \nu, \quad j \in \{1, \dots, N\}.$$

Here, $\frac{d\nu}{d\mu}$ denotes the density, or Radon-Nikodym derivative, of the probability measure ν with respect to the reference measure μ .

- evaluate f at γ_j , $j \in \{1, \dots, N\}$
- define the *weight function* $w := \frac{1}{\rho}: \Gamma \rightarrow \mathbb{R}$
- and finally define the *weighted discrete least squares approximation*

$$\Pi_V f := \arg \min_{v \in V} \|f - v\|_N, \quad (2)$$

where

$$\|g\|_N^2 := \langle g, g \rangle_N := \frac{1}{N} \sum_{j=1}^N w(\gamma_j) |g(\gamma_j)|^2 \quad \forall g: \Gamma \rightarrow \mathbb{R}. \quad (3)$$

It is straightforward to show that the coefficients \mathbf{v} of $\Pi_V f$ with respect to any basis $(B_j)_{j=1}^m$ of V are given by

$$\mathbf{G}\mathbf{v} = \mathbf{c}, \quad (4)$$

with $\mathbf{G}_{ij} := \langle B_i, B_j \rangle_N$, and $c_j := \langle f, B_j \rangle_N$, $i, j \in \{1, \dots, m\}$, assuming that \mathbf{G} is invertible. If \mathbf{G} is not invertible, then Equation (2) has multiple solutions and we define $\Pi_V f$ as the one with the minimal $L^2_\mu(\Gamma)$ norm.

Remark 1. Assembling the matrix \mathbf{G} requires $\mathcal{O}(m^2N)$ operations. However, using the fact that $\mathbf{G} = \mathbf{M}^\top \mathbf{M}$ for $\mathbf{M}_{ij} := N^{-1/2} \sqrt{w(\gamma_i)} B_j(\gamma_i)$, matrix vector products with \mathbf{G} can be computed at the lower cost $\mathcal{O}(mN)$ as $\mathbf{G}\mathbf{x} = \mathbf{M}^\top(\mathbf{M}\mathbf{x})$.

Since $w\rho = 1$, the semi-norm defined in Equation (3) is a Monte Carlo approximation of the $L_\mu^2(\Gamma)$ norm. Therefore, we may expect that the error $\|f - \Pi_V f\|_{L_\mu^2(\Gamma)}$ is close to the optimal one,

$$e_{V,2}(f) := \min_{v \in V} \|f - v\|_{L_\mu^2(\Gamma)}. \quad (5)$$

Part (iii) of Theorem 2.1 below shows that this is true in expectation, provided that the number of samples N is coupled appropriately to the dimension $m = \dim V$ of the approximating polynomial subspace and provided that we ignore outcomes where \mathbf{G} is ill-conditioned. For results in probability, we need to replace the best $L_\mu^2(\Gamma)$ approximation by the best approximation in a weighted supremum norm,

$$e_{V,w,\infty}(f) := \inf_{v \in V} \sup_{\gamma \in \Gamma} |f(\gamma) - v(\gamma)| \sqrt{w(\gamma)}. \quad (6)$$

Theorem 2.1 (Convergence of weighted least squares, [7, Theorem 2]). For arbitrary $r > 0$, define

$$\kappa := \frac{1/2 - 1/2 \log 2}{1 + r}.$$

Assume that for all $\gamma \in \Gamma$ there exists $v \in V$ such that $v(\gamma) \neq 0$ and denote by $(B_j)_{j=1}^m$ an L_μ^2 -orthonormal basis of V . Finally, assume that

$$K_{V,w} := \|w \sum_{j=1}^m B_j^2\|_{L^\infty(\Gamma)} \leq \kappa \frac{N}{\log N}. \quad (7)$$

(i) With probability larger than $1 - 2N^{-r}$, we have

$$\|\mathbf{G} - \mathbf{I}\| \leq \frac{1}{2}, \quad (8)$$

where \mathbf{G} is the matrix from Equation (4), \mathbf{I} is the identity matrix, and $\|\cdot\|$ denotes the spectral matrix norm.

(ii) If $\|\mathbf{G} - \mathbf{I}\| \leq 1/2$, then for all f with $\sup_{\gamma \in \Gamma} |f(\gamma)| \sqrt{w(\gamma)} < \infty$, we have

$$\|f - \Pi_V f\|_{L_\mu^2(\Gamma)} \leq (1 + \sqrt{2}) e_{V,w,\infty}(f).$$

(iii) If $f \in L_\mu^2(\Gamma)$, then

$$\mathbb{E} \|f - \Pi_V^c f\|_{L_\mu^2(\Gamma)}^2 \leq \left(1 + \frac{4\kappa}{\log N}\right) e_{V,2}^2(f) + 2\|f\|_{L_\mu^2(\Gamma)}^2 N^{-r},$$

where \mathbb{E} denotes the expectation with respect to the N -fold draw from the sampling distribution ν and

$$\Pi_V^c f := \begin{cases} \Pi_V f & \text{if } \|\mathbf{G} - \mathbf{I}\| \leq \frac{1}{2}, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. It is proved in [7, Theorem 2] that the bound in part (ii) holds for a fixed f with probability larger than $1 - 2N^{-r}$. A look at the proof reveals that the bound only depends on the event $\|\mathbf{G} - \mathbf{I}\| \leq 1/2$ and not on the specific choice of f . The remaining claims are exactly as in [7]. \square

3 Sampling strategies

It was observed in [7] that the constant $K_{V,w}$ in Equation (7) satisfies

$$\begin{aligned} m &= \int \sum_{j=1}^m B_j(\gamma)^2 \mu(d\gamma) \\ &\leq \left(\int w^{-1}(\gamma) \mu(d\gamma) \right) \left\| \sum_{j=1}^m w B_j^2 \right\|_{L^\infty(\Gamma)} \\ &= K_{V,w} \end{aligned} \tag{9}$$

and that the inequality becomes an equality for the weight $w_V^* = \rho_V^{*-1}$ that is associated with the density

$$\rho_V^*(\gamma) := \frac{1}{m} \sum_{j=1}^m |B_j(\gamma)|^2. \tag{10}$$

For this choice, Theorem 2.1 roughly asserts that the number of samples required to determine a near-optimal approximation of f in an m -dimensional space V is smaller than $Cm \log m$ for some $C > 0$. In the remainder of this work, we refer to w_V^* , ρ_V^* , and

$$\nu_V^* : \frac{d\nu_V^*}{d\mu} := \rho_V^* \tag{11}$$

as the *optimal* weight, density, and distribution, respectively. Since the optimal distribution ν_V^* depends on V , practical implementations need to address the question how to obtain samples from ν_V^* for general subspaces V . Furthermore, since ρ_V^* depends on V , the weight in $e_{V,\infty}(f)$ in part (ii) of Theorem 2.1 does as well. To address these issues, we present two types of results in this section.

First, we discuss how to obtain samples from ν_V^* . For the case where Γ is a d -dimensional hypercube and $\mu = \bigotimes_{j=1}^d \mu_j$ with μ_j satisfying certain assumptions, we propose a method for the generation of N samples whose computational work is bounded in expectation by the product KdN with a constant K that depends only on the measures μ_j . For non-product domains or measures, we briefly discuss how to use Markov chain Monte Carlo (MCMC) sampling for the generation of samples from approximate distributions and how perturbations of the sampling distributions affect the error estimates.

Second, we prove for the case $\Gamma = [0, 1]^d$ and under a rather permissive assumption on μ (in particular, μ must be absolutely continuous with respect to the Lebesgue measure) that for any polynomial subspace V the density of the optimal distribution ν_V^* with respect to the Lebesgue measure λ satisfies

$$C^{-d} < \frac{d\nu_V^*}{d\lambda} \leq C^d p_d^\infty, \tag{12}$$

where $0 < C < \infty$ is independent of V , and p_d^∞ is the Lebesgue density of the d -dimensional arcsine distribution,

$$p_d^\infty(\gamma) := \prod_{j=1}^d \frac{1}{\pi \sqrt{\gamma_j(1-\gamma_j)}}. \tag{13}$$

The lower bound in Equation (12) implies that the optimal weight w_V^* is bounded above by $C^d \frac{d\mu}{d\lambda}$, which can be used to make the error estimate in part (ii) of Theorem 2.1 more explicit.

By the upper bound, we may use samples from the d -dimensional arcsine distribution instead of the optimal distribution. Indeed, the upper bound implies that the weight function w associated with the arcsine distribution satisfies $K_{V,w} \leq C^d m$. Thus, the required number of samples is increased by at most a factor that is independent of V . The advantages are that samples from the arcsine distribution can be generated efficiently, that we can use samples from the same distribution for all polynomial subspaces, and that the weight w is easy to analyze and independent of V .

3.1 Sampling from the optimal distribution

We now describe an efficient algorithm to obtain samples from ν_V^* in the case when Γ is a Cartesian product, μ is a product measure, and V is downward closed.

Definition 3.1 (Downward closedness). Let $\mathbb{N} := \{0, 1, \dots\}$. A set $\mathcal{I} \subset \mathbb{N}^d$ is called downward closed if $\eta \in \mathcal{I}$ implies $\eta' \in \mathcal{I}$ for any $\eta' \in \mathbb{N}^d$ with $\eta' \leq \eta$ componentwise.

A space V of polynomials on a Cartesian product domain $\Gamma = \prod_{j=1}^d I_j$ with $I_j \subset \mathbb{R}$ is called downward closed if it is the span of monomials,

$$V = \text{span}\{\gamma^\eta = \prod_{j=1}^d \gamma_j^{\eta_j} : \eta \in \mathcal{I}\},$$

for some downward closed set $\mathcal{I} \subset \mathbb{N}^d$.

Remark 2. Observe that any non-trivial downward closed polynomial space V includes the constant functions and thus satisfies the assumption of Theorem 2.1 that for all $\gamma \in \Gamma$ there exists $v \in V$ with $v(\gamma) \neq 0$.

We first discuss the case $\Gamma = [0, 1]^d$ and $\mu = \lambda$ the Lebesgue measure. For any downward closed subspace

$$V = \text{span}\{\gamma^\eta : \eta \in \mathcal{I}\} \subset L_\lambda^2([0, 1]^d)$$

with $\mathcal{I} \subset \mathbb{N}^d$ and $|\mathcal{I}| = \dim V = m$, an orthonormal basis is then given by

$$(P_\eta)_{\eta \in \mathcal{I}}$$

where

$$P_\eta(\gamma) := \prod_{j=1}^d P_{\eta_j}(\gamma_j)$$

and $(P_n)_{n \in \mathbb{N}}$ are the Legendre polynomials on $[0, 1]$, which are orthonormal with respect to the one-dimensional Lebesgue measure. By orthonormality, each P_η^2 may be interpreted as a probability density with respect to the Lebesgue measure. Thus,

$$\frac{d\nu_V^*}{d\lambda} = \rho_V^* = \frac{1}{m} \sum_{\eta \in \mathcal{I}} P_\eta^2$$

may be interpreted as mixture of m probability densities. An efficient strategy to obtain samples from ν_V^* is therefore to first choose $\eta \in \mathcal{I}$ at random and then generate a sample from the distribution with Lebesgue density P_η^2 . Since $P_\eta^2 = \prod_{j=1}^d L_{\eta_j}^2$, samples from this distribution can be generated componentwise. Finally, to obtain samples from the univariate distributions with Lebesgue densities P_n^2 , $n \in \mathbb{N}$, we use a rejection sampling method with the arcsine proposal density p_1^∞ . By [28, Theorem 1] the Legendre polynomials satisfy

$$|P_n(\gamma)|^2 \leq 4ep_1^\infty(\gamma) \quad \forall \gamma \in [0, 1] \quad \forall n \in \mathbb{N}. \quad (14)$$

Therefore, the theory of rejection sampling [11, Chapter 4.5] ensures that if we repeatedly generate $\gamma \sim p_1^\infty$ and $U \sim \text{Unif}(0, 1)$ until $U \leq |P_n(\gamma)|^2 / (4ep_1^\infty(\gamma))$ holds, then the resulting sample is exactly distributed according to P_n^2 and the required number of iterations until acceptance has a geometric distribution with mean $4e$. The total expected computational work for the generation of N samples from ν_V^* is thus $4eNd$, if we assume that the computation of $P_n^2(\gamma)$ is $O(1)$. In practice, a 3-term recurrence formula whose work is bounded by $3n$ can be used to compute $P_n(\gamma)$. This increases the upper bound for the expected work to $12eN \frac{1}{m} \sum_{\eta \in \mathcal{I}} |\eta|_1$.

Equation (14) holds more generally for probability measures on $[0, 1]$ with Lebesgue densities of the form $\frac{d\mu}{d\lambda} = C(\alpha, \beta)\gamma^\alpha(1-\gamma)^\beta$, $\alpha, \beta \geq -1/2$ [28, Theorem 1]. The bound on the associated orthogonal polynomials $(P_n^{\alpha, \beta})_{n \in \mathbb{N}}$, which are commonly called Jacobi polynomials, is

$$|P_n^{\alpha, \beta}(\gamma)|^2 \frac{d\mu}{d\lambda} \leq 2e(2 + \sqrt{\alpha^2 + \beta^2})p_1^\infty(\gamma) \quad \forall \gamma \in [0, 1] \quad \forall n \in \mathbb{N}.$$

Even more generally, the same inequality holds with a constant C_μ independent of γ and n for orthogonal polynomials with respect to a wide class of measures μ that are absolutely continuous with respect to the Lebesgue measure on $[0, 1]$ [31, Theorem 12.1.4]. When C_μ is unknown, however, rejection sampling cannot be applied. As a substitute, we could use MCMC sampling (which we also discuss below as an alternative method to sample directly from ν^* in cases when no product structure of Γ or μ can be exploited). The error resulting from the fact that the resulting samples would not be distributed exactly according to $|P_n|^2$ can be controlled using Proposition 3.2 below.

For orthonormal polynomials $(H_n)_{n \in \mathbb{N}}$ with respect to rapidly decaying measures supported on the whole real line, such as Gaussian measures, it is shown in [22] that $|H_n(\gamma)|^2 \frac{d\mu}{d\lambda}$ is exponentially concentrated in an interval $[-a_n, a_n]$ with $C^{-1}n^b \leq a_n \leq Cn^b$ for some $b > 0$ and $C > 0$ depending on μ , and that for some C_μ

$$|H_n(\gamma)|^2 \frac{d\mu}{d\lambda} \leq C_\mu \frac{a_n}{4} \left|1 - \frac{\gamma}{a_n}\right|^{-1/2} \quad \forall \gamma \in [-a_n, a_n] \quad \forall n \in \mathbb{N}.$$

Together with the stability result in Proposition 3.2 below, this shows that the previous results can be transferred to measures on the real line, if we simply ignore the mass outside $[-a_n, a_n]$ and apply rejection sampling or Markov chain methods with the proposal density $\frac{a_n}{4} \left|1 - \frac{\gamma}{a_n}\right|^{-1/2}$. Alternatively, a different result in [22] shows that on $[-a_n, a_n]$ the density $|H_n(\gamma)|^2 \frac{d\mu}{d\lambda}$ is bounded by the uniform probability density up to a factor that grows sublinearly in the polynomial degree n .

For the remainder of this subsection, we assume that a polynomial space V is fixed and use the simplified notation $\nu^* = \nu_V^*$, $\rho^* = \rho_V^*$, $w^* = w_V^*$.

We assume that we cannot use exact samples from the optimal distribution because, for example, Γ is not a hypercube or μ not a product measure. We therefore discuss the use of Markov chain methods for the generation of samples. Since the resulting samples are not exactly distributed according to the optimal distribution, we need the following stability result.

Proposition 3.2 (Stability with respect to perturbations of the sampling density). *All results in Theorem 2.1 that are valid for the optimal choice ν^* with $\frac{d\nu^*}{d\mu} = \rho^*$ of the sampling distribution hold true if we instead use samples from $\tilde{\nu}$ with $\frac{d\tilde{\nu}}{d\mu} := \tilde{\rho}$ (but keep the weight function $w^* = 1/\rho^*$), provided that*

$$\|1 - \tilde{\rho}/\rho^*\|_{L^p_{\nu^*}(\Gamma)} \leq \frac{1}{6} m^{-1-1/p} \quad \text{for some } p \geq 1,$$

where $m = \dim V$, and provided that we replace κ by $(5/36 - 5/6 \log(5/6))/(1+r)$. We note that the total variation distance $\|\cdot\|_{TV}$ satisfies

$$\|\tilde{\nu} - \nu^*\|_{TV} := \frac{1}{2} \int_{\Gamma} |\tilde{\rho}(\gamma) - \rho^*(\gamma)| \mu(d\gamma) = \frac{1}{2} \|1 - \tilde{\rho}/\rho^*\|_{L^1_{\nu^*}(\Gamma)}.$$

Proof. The proof of Theorem 2.1 in [7] is based on large deviation bounds for the matrix \mathbf{G} of Equation (4). In particular, it is based on the observation that \mathbf{G} is a Monte Carlo average,

$$\mathbf{G} = \frac{1}{N} \sum_{i=1}^N \mathbf{X}_i,$$

of independent and identically distributed matrices

$$\mathbf{X}_i := (w^*(\gamma_i) B_j(\gamma_i) B_k(\gamma_i))_{j,k \in \{1, \dots, m\}} \quad \text{with } \gamma_i \sim \nu^*$$

that satisfy

$$\mathbb{E} \mathbf{X}_i = \mathbf{I}$$

by $L^2_\mu(\Gamma)$ -orthonormality of the basis polynomials B_j , $j \in \{1, \dots, m\}$. A Chernoff inequality for matrices then provides the bound on $\mathbb{P}(\|\mathbf{G} - \mathbf{I}\| \leq 1/2)$ in part (i) of Theorem 2.1 from which everything else follows. The crucial insight is that this inequality permits small perturbations of the

expected value. Indeed, if we replace ν^* by $\tilde{\nu}$ in the definition of \mathbf{X}_i , then [32, Theorem 1.1] yields the same bound on $\mathbb{P}(\|\mathbf{G} - \mathbf{I}\| \leq 1/2)$, with the new value of κ , provided that $\|\mathbb{E}\mathbf{X}_i - \mathbf{I}\| \leq 1/6$. To show that this last estimate holds, we use

$$\begin{aligned} \|\mathbb{E}\mathbf{X}_i - \mathbf{I}\| &\leq m \|\mathbb{E}\mathbf{X}_i - \mathbf{I}\|_{\max} \\ &= m \max_{j,k \in \{1, \dots, m\}} \left| \int_{\Gamma} w^*(\gamma) B_j(\gamma) B_k(\gamma) (\tilde{\rho}(\gamma) - \rho^*(\gamma)) \mu(d\gamma) \right| \\ &\leq m \max_{j,k \in \{1, \dots, m\}} \|w^* B_j B_k\|_{L^q_{\rho^* \mu}} \|1 - \tilde{\rho}/\rho^*\|_{L^p_{\rho^* \mu}}, \end{aligned}$$

where we used Hölder's inequality with $1/q = 1 - 1/p$ in the last step.

The claim now follows if we can prove that

$$\|w^* B_j B_k\|_{L^q_{\rho^* \mu}} \leq m^{1-1/q} \quad \forall q \in [1, \infty].$$

For this, we first consider the case when $q = 1$, $p = \infty$, for which

$$\begin{aligned} \|w^* B_j B_k\|_{L^1_{\rho^* \mu}} &= \int_{\Gamma} |B_j(\gamma) B_k(\gamma)| \mu(d\gamma) \\ &\leq \|B_j\|_{L^2_{\mu}} \|B_k\|_{L^2_{\mu}} \\ &\leq 1 \end{aligned}$$

by the Cauchy-Schwarz inequality and $L^2_{\mu}(\Gamma)$ -orthonormality of the functions B_j .

Next, we consider the case when $q = \infty$, $p = 1$, for which

$$\begin{aligned} \|w^* B_j B_k\|_{L^{\infty}_{\rho^* \mu}} &= \sup_{\gamma \in \Gamma} |w^*(\gamma) B_j(\gamma) B_k(\gamma)| \\ &\leq \frac{m B_j^2(\gamma) + B_k^2(\gamma)}{2 \sum_{i=1}^m B_i^2(\gamma)} \\ &\leq m, \end{aligned}$$

where we used the elementary inequality $ab \leq \frac{1}{2}a^2 + \frac{1}{2}b^2$ for the second step.

Finally the claim for $1 < q < \infty$ follows from Littlewood's interpolation inequality for L^q norms. \square

So far, we have assumed that Γ and μ exhibit product structure, which allowed us to generate samples coordinate-wise, exploiting known bounds on univariate orthogonal polynomials. For more general cases, we now briefly discuss Metropolized independent sampling, which is a simple MCMC algorithm, for the generation of samples from the optimal distribution ν^* . For an extensive treatment of the theory of MCMC algorithms we refer to [24].

The general strategy of MCMC algorithms for the generation of samples from ν^* is to construct a Markov chain for which ν^* is an invariant distribution. Ergodic theory then shows that under some assumptions the location of this Markov chain after $n \gg 1$ steps is approximately distributed according to ν^* . Metropolis-Hastings algorithms are MCMC algorithms that construct Markov chains based on user-specified *proposal densities* $p(\gamma, \cdot)$, $\gamma \in \Gamma$ (with respect to μ) and a rejection step to ensure convergence to the desired limit distribution ν^* . More specifically, the transition kernel of a Metropolis-Hastings algorithm has the form

$$K(\gamma, d\gamma') := \begin{cases} p(\gamma, \gamma') \min\{1, \frac{\rho^*(\gamma') p(\gamma', \gamma)}{\rho^*(\gamma) p(\gamma, \gamma')}\} \mu(d\gamma') & \text{if } \gamma' \neq \gamma \\ 1 - \int_{z \neq \gamma} p(\gamma, z) \min\{1, \frac{\rho^*(z) p(z, \gamma)}{\rho^*(\gamma) p(\gamma, z)}\} \mu(dz) & \text{if } \gamma' = \gamma. \end{cases} \quad (15)$$

This kernel can be interpreted (and implemented) as proposing a transition from the current state γ to a new state γ' drawn from the density $p(\gamma, \cdot)$, and rejecting this transition with a certain probability determined by the values of ρ^* and p at the current state γ and the proposed state γ' . The rejection probability is designed to ensure the detailed balance condition $\nu^*(d\gamma) K(\gamma, d\gamma') = \nu^*(d\gamma') K(\gamma', d\gamma)$, which in turn guarantees that ν^* is invariant under K .

Metropolized independent sampling is the name of the subset of Metropolis-Hastings algorithms for which the proposal density p is independent of the current state γ . If we denote the corresponding state-independent proposal density by $p(\gamma')$ and define $g := \inf_{\gamma \in \Gamma} p(\gamma)/\rho^*(\gamma)$, then it can be shown [23, Section 3.2.2] that starting from any distribution π we have the bound

$$\|K^n \pi - \nu^*\|_{TV} \leq 2(1-g)^n \quad (16)$$

for the total variation distance between the n th step probability distribution $K^n \pi$ of the Markov chain and the target distribution ν^* . This means that if the proposal density satisfies $g := \inf_{\gamma \in \Gamma} p(\gamma)/\rho^*(\gamma) > 0$, then $n := g^{-1} \log(24m^2)$ Markov chain steps suffice to ensure that

$$\|K^n \pi - \nu^*\|_{TV} \leq 2(1-g)^{g^{-1} \log(24m^2)} \leq \frac{1}{12m^2},$$

as required by Proposition 3.2. To generate $N > 0$ independent samples from $K^n \pi$, we have to run N independent copies of the Markov chain, which differs from the more common practice to use N successive, thus dependent, steps of a single Markov chain.

3.2 Sampling from the arcsine distribution

We now determine lower and upper bounds for the optimal sampling distributions associated with downward closed polynomial spaces on $[0, 1]^d$ in terms of the arcsine distribution. Namely, we show in Proposition 3.4 below, under quite general assumptions on the measure μ , the bound

$$C^{-d} \leq \frac{d\nu_V^*}{d\lambda}(\gamma) \leq C^d p_d^\infty(\gamma),$$

where p_d^∞ is the arcsine density from Equation (13).

The lower bound can be used to make the bound in Theorem 2.1 more precise. Indeed, it implies that the weight w_V^* appearing in $e_{V, w_V^*, \infty}$ satisfies

$$w_V^*(\gamma) \leq C^d \frac{d\mu}{d\lambda}(\gamma). \quad (17)$$

The upper bound provides an alternative sampling strategy: Instead of sampling from the optimal distribution, we may simply sample from the arcsine distribution with Lebesgue density p_d^∞ without using Acceptance/Rejection or Markov chain methods. Indeed, since using the arcsine distribution for sample generation corresponds to the choice $\rho = p_d^\infty \frac{d\lambda}{d\mu}$ for the density in Section 2, the associated weight function is given by

$$w := (p_d^\infty)^{-1} \frac{d\mu}{d\lambda}. \quad (18)$$

Hence, the upper bound shows that the constant $K_{V, w}$ from Theorem 2.1 satisfies

$$\begin{aligned} K_{V, w} &= \|w(\gamma) \sum_{j=1}^m B_j(\gamma)^2\|_{L^\infty(\Gamma)} \\ &= \|w(\gamma) m \frac{d\nu^*}{d\mu}(\gamma)\|_{L^\infty(\Gamma)} \\ &\leq C^d m, \end{aligned} \quad (19)$$

which is larger than the optimal value, m , only by the factor C^d , which is independent of V . The advantages are that exact and independent samples from the univariate arcsine distribution can be generated efficiently as $(\sin(X) + 1)/2$ for a uniform random variable X on $[-\pi/2, \pi/2]$, that we can use samples from the same distribution for all polynomial subspaces, and that the weight in Equation (18), which enters the error estimate in Theorem 2.1 through $e_{V, w, \infty}$, is known explicitly, is independent of V , and vanishes at the boundary if $\frac{d\mu}{d\lambda}$ is bounded.

Definition 3.3 (Doubling measures). A measure μ on $[0, 1]$ is called doubling measure if it is absolutely continuous with respect to the Lebesgue measure and if there exists $L > 0$ such that

$$\mu(2I) \leq L\mu(I)$$

for any interval $I \subset [0, 1]$. Here, $2I$ is defined as the interval with the same midpoint as I and twice the length, intersected with $[0, 1]$. We call $L > 0$ the doubling constant of μ .

Proposition 3.4 (Bounds on the optimal distribution). If $[0, 1]^d$, $d \geq 1$ is equipped with the product $\mu = \bigotimes_{j=1}^d \mu_j$ of doubling measures μ_j on $[0, 1]$, then the optimal sampling distribution ν_V^* associated with a finite-dimensional downward closed space V of polynomials on $[0, 1]^d$ satisfies

$$C^{-d} \leq \frac{d\nu_V^*}{d\lambda}(\gamma) \leq C^d p_d^\infty(\gamma) \quad (20)$$

for some constant $C > 0$ depending only on the maximal doubling constant of the measures μ_j , $j \in \{1, \dots, d\}$.

Proof. Equation (20) was shown to hold for the univariate optimal sampling distributions $\nu_k^*(\gamma)$ associated with univariate spaces of polynomials of degree less than or equal to $k \in \mathbb{N}$ in [25, Equation 7.14]. We prove the case $d > 1$ by induction.

To reduce the complexity of notation, we assume without loss of generality that $\mu_1 = \mu_j$ for all $j \in \{1, \dots, d\}$. This, together with the assumption that V is downward closed, implies that

$$V = \text{span}_{\eta \in \mathcal{I} \subset \mathbb{N}^d} \{B_\eta(\gamma) := B_{\eta_1}(\gamma_1) \cdots B_{\eta_d}(\gamma_d)\}$$

for $(B_j)_{j \in \mathbb{N}}$ the univariate orthonormal polynomials associated with μ_1 and a multi-index set $\mathcal{I} \subset \mathbb{N}^d$. For $j \in \mathbb{N}$, we define the multi-index sets $\mathcal{I}_j := \{\eta \in \mathcal{I} : \eta_1 = j\}$ and the spaces

$$\tilde{V}_j := \text{span}_{\tilde{\eta} \in \tilde{\mathcal{I}}_j \subset \mathbb{N}^{d-1}} \{\tilde{B}_{\tilde{\eta}}(\tilde{\gamma}) := B_{\tilde{\eta}_2}(\tilde{\gamma}_2) \cdots B_{\tilde{\eta}_d}(\tilde{\gamma}_d)\}$$

of polynomials on $[0, 1]^{d-1}$ with associated optimal distributions $\tilde{\nu}_j$ on $[0, 1]^{d-1}$. This allows us to write

$$\begin{aligned} \frac{d\nu_V^*}{d\lambda}(\gamma) &= \rho_V^*(\gamma) \frac{d\mu}{d\lambda}(\gamma) \\ &= \frac{1}{|\mathcal{I}|} \sum_{\eta \in \mathcal{I}} B_\eta^2(\gamma) \frac{d\mu}{d\lambda}(\gamma) \\ &= \frac{1}{|\mathcal{I}|} \sum_{j \in \mathbb{N}} B_j^2(\gamma_1) \frac{d\mu_1}{d\lambda}(\gamma_1) \sum_{\eta \in \mathcal{I}_j} \tilde{B}_\eta^2(\tilde{\gamma}) \prod_{j=2}^d \frac{d\mu_1}{d\lambda}(\gamma_j) \\ &= \frac{1}{|\mathcal{I}|} \sum_{j \in \mathbb{N}} B_j^2(\gamma_1) \frac{d\mu_1}{d\lambda}(\gamma_1) |\mathcal{I}_j| \frac{d\tilde{\nu}_j}{d\lambda}(\tilde{\gamma}), \end{aligned}$$

which by the induction hypothesis for the case $d - 1$ entails

$$C^{-(d-1)} A(\gamma_1) \leq \frac{d\nu_V^*}{d\lambda}(\gamma) \leq A(\gamma_1) C^{d-1} p_{d-1}^\infty(\tilde{\gamma}) \quad (21)$$

with

$$A(\gamma_1) := \frac{1}{|\mathcal{I}|} \sum_{j \in \mathbb{N}} B_j^2(\gamma_1) \frac{d\mu_1}{d\lambda}(\gamma_1) |\mathcal{I}_j|.$$

We now use the fact that $A(\gamma_1)$ can be written as a weighted average of the univariate densities $\frac{d\nu_k^*}{d\lambda}(\gamma_1) = \frac{1}{k} \sum_{j=1}^k B_j^2(\gamma_1) \frac{d\mu_1}{d\lambda}(\gamma_1)$:

$$A(\gamma_1) = \frac{1}{\sum_{k=1}^\infty w_k} \sum_{k=1}^\infty w_k \frac{d\nu_{w_k}^*}{d\lambda}(\gamma_1)$$

with $w_k := |\{j : |\mathcal{I}_j| \geq k\}|$. Together with the case $d = 1$ this shows

$$C^{-1} \leq A(\gamma_1) \leq Cp_1^\infty(\gamma_1),$$

which, when inserted into Equation (21), yields

$$C^{-d} = C^{-(d-1)}C^{-1} \leq \frac{d\nu_V^*}{d\lambda}(\gamma) \leq C^{d-1}Cp_1^\infty(\gamma_1)p_{d-1}^\infty(\tilde{\gamma}) = C^d p_d^\infty(\gamma).$$

□

4 Multilevel weighted least squares approximation

In this section, we define a multilevel weighted polynomial least squares method and establish convergence rates for the approximation of a function $f_\infty : \Gamma \subset \mathbb{R}^d \rightarrow \mathbb{R}$, $d \in \mathbb{N} \cup \{\infty\}$ in a normed vector space $(F, \|\cdot\|_F) \hookrightarrow (L_\mu^2(\Gamma), \|\cdot\|_{L_\mu^2(\Gamma)})$ of continuous functions on Γ , under the following assumptions.

- **A1:** (Convergence of approximations) There exist functions $f_n \in F$, $n \geq 1$ such that

$$\|f_\infty - f_n\|_F \leq C_0 n^{-\beta_s}$$

$$\|f_\infty - f_n\|_{L_\mu^2(\Gamma)} \leq C_0 n^{-\beta_w}$$

for some $C_0 > 0$, $\beta_s > 0$ and $\beta_w \geq \beta_s$.

- **A2(p):** (Polynomial approximability) There exist downward closed spaces of polynomials V_m , $m \geq 1$ on Γ such that

$$\dim V_m \leq C_1 m^\sigma,$$

$$e_{m,p}(F) := \sup_{f \in F} \frac{e_{V_m,p}(f)}{\|f\|_F} \leq C_1 m^{-\alpha}$$

for some $\sigma > 0$, $\alpha > 0$, $C_1 > 0$ and $p = 2$ or $p = \infty$. (In the latter case, we use the shorthand $e_{V_m,\infty}(f) := e_{V_m,w_m^*,\infty}(f)$, where w_m^* is the optimal weight associated with V_m .)

- **A3:** (Sample work) The work required for a single evaluation of f_n satisfies $\text{Work}(f_n) \leq C_2 n^\gamma$ for some $\gamma > 0$, $C_2 > 0$.

Remark 3. In Assumption A2(p), we have introduced the exponent σ , which in contrast to previous sections may be different from 1, to be able to apply our results with common sequences of polynomial subspaces without the need for reparametrization.

Example 1 (Polynomial approximability). • For univariate Sobolev spaces $F = H^\alpha(\Gamma)$, $\Gamma = (0, 1)$ with $\alpha > 0$, Theorem 1 in [30] shows that

$$e_{m,2}(H^\alpha(\Gamma)) \leq Cm^{-\alpha}$$

for the space V_m of univariate polynomials with degree less than m and for μ the Lebesgue measure. Analogous results also hold in higher dimensions. Here, optimal sequences of polynomial approximation spaces depend on the available smoothness. In particular, optimal polynomial approximation spaces for functions in Sobolev spaces $H^\alpha(\Gamma)$ with $\Gamma \subset \mathbb{R}^d$ and $\alpha > 0$ are of total degree type, whereas functions in Sobolev spaces $H_{\text{mix}}^\alpha(\Gamma)$ of dominating mixed smoothness can be optimally approximated by hyperbolic cross polynomial spaces [10].

Similar results for the best approximation in the supremum norm hold for functions in Hölder spaces $F = C^{s,t}(\Gamma)$, $s \in \mathbb{N}$, $t \in [0, 1]$ [2, Theorem 2] (and their dominating mixed smoothness analogues).

- Alternatively, we may simply define the space F via polynomial approximability of its elements. Assume that we have a sequence $(V_m)_{m=1}^\infty$ of downward closed polynomial spaces on $\Gamma \subset \mathbb{R}^d$ with $d \in \mathbb{N} \cup \{\infty\}$. If for some $\alpha > 0$ we define

$$F := \left\{ f: \Gamma \rightarrow \mathbb{R} : \|f\|_F := \sup_{m \in \mathbb{N}} e_{V_m, p}(f) m^\alpha < \infty \right\}$$

with the auxiliary definition $V_0 := \{0\}$, then it is easy to show that $\|\cdot\|_F$ is a norm of F and that Assumption 2(p) holds with the given α and $C_1 = 1$. The choice of the sequence of subspaces V_m can be based on truncating a orthogonal decomposition of $L_\mu^2(\Gamma)$ such as to include only basis functions whose contribution is above a given threshold in V_m . For more information on this construction, see Section 5 and [15, 9].

We now define the multilevel least squares method for a fixed number of levels $L \in \mathbb{N}$. We introduce subsequences

$$m_k := M \exp(k/(\sigma + \alpha)), \quad k \in \{0, \dots, L\} \quad (22)$$

and

$$n_l := \exp(l/(\gamma + \beta_s)), \quad l \in \{0, \dots, L\}$$

with $M := \exp(L\delta)$, $\delta := \frac{\beta_w - \beta_s}{\alpha(\gamma + \beta_s)} \geq 0$ if $\gamma/\beta_s > \sigma/\alpha$ and $M := 1$ else. For our analysis we assume that m and n can take non-integer values; in practice, rounding up to the nearest integer increases the required work only by a constant factor. By abuse of notation, we keep the simple notation V_k , $e_{k,p}$, and f_l for the quantities V_{m_k} , $e_{m_k,p}$, and f_{n_l} , respectively.

Next, we draw independent, identically distributed, random samples

$$\Gamma_k = \{\gamma_{k,1}, \dots, \gamma_{k,|\Gamma_k|}\} \subset \Gamma, \quad k \in \{0, \dots, L\}$$

with $\gamma_{k,j} \sim \nu_k^*$, where $\nu_k^* := \nu_{V_k}^*$ is the optimal sampling distribution of V_k from Equation (11). To ensure accuracy of our approximations, we couple the numbers of samples to the dimensions of the polynomial spaces via

$$m_k^\sigma \leq \kappa \frac{|\Gamma_k|}{\log |\Gamma_k|} \leq 2m_k^\sigma \quad \forall k \in \{0, \dots, L\}, \quad \text{where } \kappa := \frac{1 - \log 2}{2 + 2L}. \quad (23)$$

By Equation (9), this guarantees that the assumption of Theorem 2.1 is satisfied with $r = L$. Alternatively, we may replace κ by $C^{-d}\kappa$ with C from Proposition 3.4 if Γ and $\mu \ll \lambda$ are products and if we use the arcsine distribution to generate samples, or we may choose κ as in Proposition 3.2 if we use samples that are only approximately distributed according to the optimal distribution.

Finally, we denote by $\Pi_k: F \rightarrow V_k$ the random weighted least squares approximation using evaluations in Γ_k , $k \in \{0, \dots, L\}$ and define the multilevel method

$$\begin{aligned} \mathcal{S}_L(f_\infty) &:= \Pi_L f_0 + \sum_{l=1}^L \Pi_{L-l}(f_l - f_{l-1}) \\ &= \sum_{l=0}^L \Pi_{L-l}(f_l - f_{l-1}) \end{aligned} \quad (24)$$

where we used the auxiliary definition $f_{-1} := 0$.

For the sake of readability, we do not keep track of constants and denote by \lesssim any inequality that holds up to a factor depending only on $C_0, C_1, C_2, \alpha, \beta_s, \beta_w, \gamma$. The exact values of these constants may be determined from the proof of the next Theorem.

Theorem 4.1 (Convergence in probability). *Denote by*

$$\text{Work}(\mathcal{S}_L(f_\infty)) := |\Gamma_L| \text{Work}(f_0) + \sum_{l=1}^L |\Gamma_{L-l}| (\text{Work}(f_l) + \text{Work}(f_{l-1})) \quad (25)$$

the work that $\mathcal{S}_L(f_\infty)$ requires for evaluations of the functions f_l , $l \in \{0, \dots, L\}$. Define

$$\lambda := \begin{cases} \sigma/\alpha & \text{if } \gamma/\beta_s \leq \sigma/\alpha \\ \theta\gamma/\beta_s + (1-\theta)\sigma/\alpha \text{ with } \theta := \beta_s/\beta_w & \text{if } \gamma/\beta_s > \sigma/\alpha \end{cases}$$

and

$$t := \begin{cases} 2 & \text{if } \gamma/\beta_s < \sigma/\alpha \\ 3 + \sigma/\alpha & \text{if } \gamma/\beta_s = \sigma/\alpha \\ 1 & \text{if } \gamma/\beta_s > \sigma/\alpha \text{ and } \beta_w = \beta_s \\ 2 & \text{if } \gamma/\beta_s > \sigma/\alpha \text{ and } \beta_w > \beta_s \end{cases}.$$

Let $0 < \epsilon \lesssim 1$. If Assumptions A1, A2(∞), and A3 hold, then we may choose $L \in \mathbb{N}$ such that

$$\text{Work}(\mathcal{S}_L(f_\infty)) \lesssim \epsilon^{-\lambda} |\log \epsilon|^t \log |\log \epsilon|,$$

and such that in an event E with $\mathbb{P}(E^c) \lesssim \epsilon^{\log |\log \epsilon|}$ the multilevel approximation satisfies

$$\|f_\infty - \mathcal{S}_L(f_\infty)\|_{L_\mu^2(\Gamma)} \leq \epsilon. \quad (26)$$

Proof. The strategy of this proof is to establish bounds on $\text{Work}(\mathcal{S}_L(f_\infty))$ and $\|f_\infty - \mathcal{S}_L(f_\infty)\|_{L_\mu^2(\Gamma)}$ for arbitrary $L \in \mathbb{N}$ first, and then to show that, for the right choice of L , the latter is smaller than ϵ and the former is bounded by $\epsilon^{-\lambda} |\log \epsilon|^t \log |\log \epsilon|$.

Work bounds. We may deduce immediately from Equation (23) the rough upper bound

$$\sqrt{|\Gamma_k|} \leq \frac{|\Gamma_k|}{\log |\Gamma_k|} \leq \frac{2}{\kappa} M^\sigma \exp(k \frac{\sigma}{\sigma + \alpha}) \lesssim (L+1) M^\sigma \exp(k \frac{\sigma}{\sigma + \alpha})$$

on the number of samples at level $k \in \{0, \dots, L\}$. Using Equation (23) again and inserting the previous estimate, we obtain the finer estimate

$$\begin{aligned} |\Gamma_k| &\leq (L+1) M^\sigma \exp(k \frac{\sigma}{\sigma + \alpha}) \log |\Gamma_k| \\ &\lesssim (L+1) M^\sigma (\log(L+1) + \log M^\sigma) \exp(k \frac{\sigma}{\sigma + \alpha}) (k+1). \end{aligned}$$

Since

$$\text{Work}(f_l) + \text{Work}(f_{l-1}) \lesssim \exp(l \frac{\gamma}{\gamma + \beta_s})$$

by Assumption A3, we may conclude that

$$\begin{aligned} \text{Work}(\mathcal{S}_L(f_\infty)) &\lesssim (L+1) M^\sigma (\log(L+1) + \log M^\sigma) \sum_{l=0}^L \exp((L-l) \frac{\sigma}{\sigma + \alpha}) (L-l+1) \exp(l \frac{\gamma}{\gamma + \beta_s}) \\ &= (L+1) M^\sigma (\log(L+1) + \log M^\sigma) \exp(L \frac{\sigma}{\sigma + \alpha}) \sum_{l=0}^L \exp\left(-l \left(\frac{\sigma}{\sigma + \alpha} - \frac{\gamma}{\gamma + \beta_s}\right)\right) (L-l+1). \end{aligned} \quad (27)$$

We now distinguish three cases.

- (a) $\gamma/\beta_s < \sigma/\alpha$: In this case $\sigma/(\sigma + \alpha) > \gamma/(\gamma + \beta_s)$. Thus, the sum on the right-hand side of Equation (27) satisfies

$$\begin{aligned} \sum_{l=0}^L \exp\left(-l \left(\frac{\sigma}{\sigma + \alpha} - \frac{\gamma}{\gamma + \beta_s}\right)\right) (L-l+1) &\lesssim (L+1) \sum_{l=0}^L \exp\left(-l \left(\frac{\sigma}{\sigma + \alpha} - \frac{\gamma}{\gamma + \beta_s}\right)\right) \\ &\lesssim L+1. \end{aligned}$$

Together with the fact that $M = 1$ in the case under consideration, this shows that

$$\text{Work}(\mathcal{S}_L(f_\infty)) \lesssim \exp(L \frac{\sigma}{\sigma + \alpha}) (L+1)^2 \log(L+1).$$

- (b) $\gamma/\beta_s = \sigma/\alpha$: In this case $\sigma/(\sigma + \alpha) = \gamma/(\gamma + \beta_s)$. Thus, the sum on the right-hand side of Equation (27) equals $\sum_{l=0}^L (L - l + 1) \lesssim (L + 1)^2$ and we obtain

$$\text{Work}(\mathcal{S}_L(f_\infty)) \lesssim \exp\left(L \frac{\sigma}{\sigma + \alpha}\right) (L + 1)^3 \log(L + 1).$$

since $M = 1$.

- (c) $\gamma/\beta_s > \sigma/\alpha$: In this case $\sigma/(\sigma + \alpha) < \gamma/(\gamma + \beta_s)$. Thus, the sum on the right-hand side of Equation (27) satisfies

$$\begin{aligned} & \sum_{l=0}^L \exp\left(-l\left(\frac{\sigma}{\sigma + \alpha} - \frac{\gamma}{\gamma + \beta_s}\right)\right) (L - l + 1) \\ &= \exp\left(L\left(\frac{\gamma}{\gamma + \beta_s} - \frac{\sigma}{\sigma + \alpha}\right)\right) \sum_{l=0}^L \exp\left(-l\left(\frac{\gamma}{\gamma + \beta_s} - \frac{\sigma}{\sigma + \alpha}\right)\right) (l + 1) \\ &\lesssim \exp\left(L\left(\frac{\gamma}{\gamma + \beta_s} - \frac{\sigma}{\sigma + \alpha}\right)\right). \end{aligned}$$

If $\beta_w = \beta_s$, then $M = 1$ and we obtain

$$\begin{aligned} \text{Work}(\mathcal{S}_L(f_\infty)) &\lesssim (L + 1) M^\sigma (\log(L + 1) + \log M^\sigma) \exp\left(L \frac{\gamma}{\gamma + \beta_s}\right) \\ &\lesssim \exp\left(L\left(\frac{\gamma}{\gamma + \beta_s}\right)\right) (L + 1) \log(L + 1). \end{aligned}$$

If instead $\beta_w > \beta_s$, then $M = \exp(\delta L)$ and we obtain

$$\begin{aligned} \text{Work}(\mathcal{S}_L(f_\infty)) &\lesssim (L + 1) M^\sigma (\log(L + 1) + \log M^\sigma) \exp\left(L \frac{\gamma}{\gamma + \beta_s}\right) \\ &\lesssim \exp\left(L\left(\frac{\gamma}{\gamma + \beta_s} + \sigma\delta\right)\right) (L + 1)^2 \log(L + 1). \end{aligned}$$

Residual bounds. First, we show that with high probability

$$\|\text{Id} - \Pi_k\|_{F \rightarrow L_\mu^2(\Gamma)} \lesssim M^{-\alpha} \exp(-k\alpha/(\sigma + \alpha)) \quad \forall k \in \{0, \dots, L\}. \quad (28)$$

By part (ii) of Theorem 2.1 together with Assumption A2(∞), it suffices to show that the event

$$E := \{\|\mathbf{G}_k - \mathbf{I}_k\| \leq 1/2 \quad \forall k \in \mathbb{N}\}$$

has a high probability, where \mathbf{G}_k is the Gramian matrix from Equation (4). But by the first part of the same theorem, the complementary probability that $\|\mathbf{G}_k - \mathbf{I}_k\| \leq 1/2$ for a fixed $k \in \mathbb{N}$ decays as the number of samples $|\Gamma_k|$ increases. Since the sets Γ_k grow exponentially in k , by Equation (23), we may conclude using a crude zeroth moment estimate and a geometric series bound:

$$\begin{aligned} \mathbb{P}(E^c) &= \mathbb{P}(\exists k \in \mathbb{N} : \|\mathbf{G}_k - \mathbf{I}_k\| > 1/2) \\ &\leq \sum_{k=0}^{\infty} \mathbb{P}(\|\mathbf{G}_k - \mathbf{I}_k\| > 1/2) \\ &\leq 2 \sum_{k=0}^{\infty} |\Gamma_k|^{-L} \\ &\leq 2\kappa^L M^{-\sigma L} \sum_{k=0}^{\infty} \exp\left(-kL \frac{\sigma}{\sigma + \alpha}\right) \\ &= \frac{2\kappa^L M^{-\sigma L}}{1 - \exp\left(-L \frac{\sigma}{\sigma + \alpha}\right)} \\ &\lesssim L^{-L}. \end{aligned} \quad (29)$$

Assuming now that the samples Γ_k , $k \in \mathbb{N}$ are such that Equation (28) holds for the associated operators Π_k , we obtain

$$\begin{aligned}
\|f_\infty - \mathcal{S}_L(f_\infty)\|_{L^2_\mu(\Gamma)} &= \left\| f_\infty - \left(\sum_{l=0}^L (f_l - f_{l-1}) - \sum_{l=0}^L (\text{Id} - \Pi_{L-l})(f_l - f_{l-1}) \right) \right\|_{L^2_\mu(\Gamma)} \\
&\leq \|f_\infty - f_L\|_{L^2_\mu(\Gamma)} + \sum_{l=0}^L \|\text{Id} - \Pi_{L-l}\|_{F \rightarrow L^2_\mu(\Gamma)} \|f_l - f_{l-1}\|_F \\
&\lesssim \exp\left(-L \frac{\beta_w}{\gamma + \beta_s}\right) + M^{-\alpha} \sum_{l=0}^L \exp\left(-\frac{\alpha}{\sigma + \alpha}(L-l)\right) \exp\left(-l \frac{\beta_s}{\gamma + \beta_s}\right) \\
&= \exp\left(-L \frac{\beta_w}{\gamma + \beta_s}\right) + M^{-\alpha} \exp\left(-L \frac{\alpha}{\sigma + \alpha}\right) \sum_{l=0}^L \exp\left(l \left(\frac{\alpha}{\sigma + \alpha} - \frac{\beta_s}{\gamma + \beta_s}\right)\right),
\end{aligned} \tag{30}$$

where we used Assumption A1. Again, we distinguish the cases (a)-(c).

- (a) $\gamma/\beta_s < \sigma/\alpha$: In this case $\alpha/(\sigma + \alpha) < \beta_s/(\gamma + \beta_s)$. Thus, the sum on the right-hand side of Equation (30) is uniformly bounded in L and we obtain

$$\begin{aligned}
\|f_\infty - \mathcal{S}_L(f_\infty)\|_{L^2(\mu)} &\lesssim \exp\left(-L \frac{\beta_w}{\gamma + \beta_s}\right) + \exp\left(-L \frac{\alpha}{\sigma + \alpha}\right) \\
&\lesssim \exp\left(-L \frac{\alpha}{\sigma + \alpha}\right),
\end{aligned}$$

where we used the fact that $\beta_w \geq \beta_s$ for the last inequality.

- (b) $\gamma/\beta_s = \sigma/\alpha$: In this case $\alpha/(\sigma + \alpha) = \beta_s/(\gamma + \beta_s)$. Thus, the sum on the right-hand side of Equation (27) equals $L + 1$ and we obtain

$$\begin{aligned}
\|f_\infty - \mathcal{S}_L(f_\infty)\|_{L^2(\mu)} &\lesssim \exp\left(-L \frac{\beta_w}{\gamma + \beta_s}\right) + \exp\left(-L \frac{\alpha}{\sigma + \alpha}\right)(L + 1) \\
&\lesssim \exp\left(-L \frac{\alpha}{\sigma + \alpha}\right)(L + 1),
\end{aligned}$$

where we used the fact that $\beta_w \geq \beta_s$ for the last inequality.

- (c) $\gamma/\beta_s > \sigma/\alpha$: In this case $\alpha/(\sigma + \alpha) > \beta_s/(\gamma + \beta_s)$. Thus, the sum on the right-hand side of Equation (27) is a divergent geometric series and we obtain

$$\begin{aligned}
\|f_\infty - \mathcal{S}_L(f_\infty)\|_{L^2(\mu)} &\lesssim \exp\left(-L \frac{\beta_w}{\gamma + \beta_s}\right) + M^{-\alpha} \exp\left(-L \frac{\beta_s}{\gamma + \beta_s}\right) \\
&\lesssim \exp\left(-L \frac{\beta_w}{\gamma + \beta_s}\right),
\end{aligned}$$

where we used the definition of $M = \exp(L\delta)$ and δ in the case $\gamma/\beta_s > \sigma/\beta_w$ in the last inequality.

Conclusion. It remains to choose L such that the residual bound equals ϵ and insert this choice of L into the work bound. For simplicity, we assume L can be any real number. In practice, rounding up to the next largest value decreases the residual and increases the work only by a constant factor. One final time, we distinguish the cases (a)-(c).

- (a) $\gamma/\beta_s < \sigma/\alpha$: Defining L as the solution of

$$\exp\left(-L \frac{\alpha}{\sigma + \alpha}\right) = \epsilon,$$

we obtain the second inequality in the following estimate:

$$\begin{aligned}
\text{Work}(\mathcal{S}_L(f_\infty)) &\lesssim \exp\left(L \frac{\sigma}{\sigma + \alpha}\right)(L + 1)^2 \log(L + 1) \\
&\lesssim \epsilon^{-\lambda} |\log \epsilon|^2 \log |\log \epsilon|.
\end{aligned}$$

(b) $\gamma/\beta_s = \sigma/\alpha$: Since we assumed that $\epsilon \lesssim 1$ there is a unique positive solution of

$$\exp(-L \frac{\alpha}{\sigma + \alpha})(L + 1) = \epsilon.$$

With this choice of L we obtain the second inequality in the following estimate:

$$\begin{aligned} \text{Work}(\mathcal{S}_L(f_\infty)) &\lesssim \exp(L \frac{\sigma}{\sigma + \alpha})(L + 1)^3 \log(L + 1) \\ &\lesssim \epsilon^{-\lambda} |\log \epsilon|^{3+\lambda} \log |\log \epsilon|. \end{aligned}$$

(c) $\gamma/\beta_s > \sigma/\alpha$: We assume $\beta_w > \beta_s$, the case $\beta_w = \beta_s$ can be treated analogously. Defining L as the solution of

$$\exp(-L \frac{\beta_w}{\gamma + \beta_s}) = \epsilon,$$

we obtain the second inequality in the following estimate:

$$\begin{aligned} \text{Work}(\mathcal{S}_L(f_\infty)) &\lesssim \exp\left(L \left(\frac{\gamma}{\gamma + \beta_s} + \sigma\delta\right)\right)(L + 1)^2 \log(L + 1) \\ &\lesssim \epsilon^{-\lambda} |\log \epsilon|^2 \log |\log \epsilon|. \end{aligned}$$

In all cases we chose L such that $L \geq |\log \epsilon|$, thus $\mathbb{P}(E^c) \lesssim L^{-L} \lesssim \epsilon^{|\log |\log \epsilon||}$ by Equation (29). \square

Remark 4. *The proof does not exploit independence of samples across different Γ_k , $k \in \{0, \dots, L\}$, but instead relies on a simple union bound (see Equation (29)). Thus, we could alternatively first create Γ_L and then define all Γ_l with $l < L$ as subsets of it.*

Remark 5. *To determine the polynomial coefficients of $\Pi_{L-l}(f_l - f_{l-1})$, $l \in \{0, \dots, L\}$, after the functions $f_l - f_{l-1}$ have been evaluated in all $\gamma \in \Gamma_{L-l}$, we need to solve linear systems of the form*

$$\mathbf{G}_k \mathbf{v}_k = \mathbf{c}_k, \quad k \in \{0, \dots, L\} \quad (31)$$

as in Equation (4). In the event E in which the residual estimate of the previous theorem holds, the condition numbers of all matrices \mathbf{G}_k are bounded by 3. Therefore, using a suitable iterative solver we can determine all coefficients \mathbf{v}_k to an accuracy of $\epsilon > 0$ with $\mathcal{O}(\log \epsilon)$ iterations. Since matrix vector products with \mathbf{G}_k require $\mathcal{O}(m_k^{2\sigma})$ operations by Remark 1, the associated computational work is, up to logarithmic factors, given by

$$\sum_{k=0}^L m_k^{2\sigma} = \sum_{k=0}^L M^{2\sigma} \exp(2k\sigma/(\sigma + \alpha)) \lesssim M^{2\sigma} \exp(2L\sigma/(\sigma + \alpha)).$$

Inspection of the proof of the previous theorem shows that, even if we include this cost in the work specification, the conclusion holds true with slightly different logarithmic factors and the exponent

$$\tilde{\lambda} := \begin{cases} 2\sigma/\alpha & \text{if } \gamma/\beta_s \leq 2\sigma/\alpha \\ \gamma/\beta_s & \text{if } \gamma/\beta_s > 2\sigma/\alpha, \end{cases}$$

instead of λ (assuming for simplicity that $\beta_s = \beta_w$), provided that we change the definition of the subsequence m_k in Equation (22) to

$$m_k := \exp(k/(2\sigma + \alpha)).$$

To obtain mean square convergence, we replace the least squares approximations Π_k by the stabilized versions Π_k^c from part (iii) of Theorem 2.1, and define

$$\mathcal{S}_L^c(f_\infty) := \Pi_L^c f_0 + \sum_{l=1}^L \Pi_{L-l}^c (f_l - f_{l-1}). \quad (32)$$

Theorem 4.2 (Mean square convergence). *Let $0 < \epsilon \lesssim 1$. If Assumptions A1, A2(2), and A3 hold, then we may choose $L \in \mathbb{N}$ such that*

$$\mathbb{E}\|f_\infty - \mathcal{S}_L^c(f_\infty)\|_{L_\mu^2(\Gamma)}^2 \leq \epsilon^2 \quad (33)$$

and

$$\text{Work}(\mathcal{S}_L^c(f_\infty)) \lesssim \epsilon^{-\lambda} |\log \epsilon|^t \log |\log \epsilon|,$$

with λ and t as in Theorem 4.1.

Proof. The work bounds from the proof of Theorem 4.1 hold unchanged.

We next establish residual bounds for arbitrary $L \in \mathbb{N}$ as before, using the error representation

$$f_\infty - \mathcal{S}_L^c(f_\infty) = f_\infty - f_L + \sum_{l=0}^L (\text{Id} - \Pi_{L-l}^c)(f_l - f_{l-1}).$$

The triangle inequality of the norm $(\mathbb{E}\|\cdot\|_{L_\mu^2(\Gamma)}^2)^{1/2}$ implies that

$$\begin{aligned} \left(\mathbb{E}\|f_\infty - \mathcal{S}_L^c(f_\infty)\|_{L_\mu^2(\Gamma)}^2\right)^{1/2} &\leq \left(\|f_\infty - f_L\|_{L_\mu^2(\Gamma)}^2\right)^{1/2} + \sum_{l=0}^L \left(\mathbb{E}\|(\text{Id} - \Pi_{L-l}^c)(f_l - f_{l-1})\|_{L_\mu^2(\Gamma)}^2\right)^{1/2} \\ &\lesssim \|f_\infty - f_L\|_{L^2(\mu)} + \sum_{l=0}^L \left(e_{V_{L-l},2}^2(f_l - f_{l-1}) + \|f_l - f_{l-1}\|_{L_\mu^2(\Gamma)}^2 |\Gamma_{L-l}|^{-2\alpha/\sigma}\right)^{1/2} =: (\star) \end{aligned}$$

where we used part (iii) of Theorem 2.1 together with the fact that $L \geq 2\alpha/\sigma$ for small enough ϵ for the second inequality. We observe that

- by Assumption A1, we have

$$\|f_\infty - f_L\|_{L_\mu^2(\Gamma)} \lesssim \exp\left(-L \frac{\beta_w}{\gamma + \beta_s}\right)$$

- by Assumptions A1 and A2(2), we have

$$e_{V_{L-l},2}^2(f_l - f_{l-1}) \lesssim \left(M^{-\alpha} \exp\left(-\frac{\alpha}{\sigma + \alpha}(L-l)\right) \exp\left(-l \frac{\beta_s}{\gamma + \beta_s}\right)\right)^2$$

- by Equation (23) and Assumption A1, we have

$$\|f_l - f_{l-1}\|_{L_\mu^2(\Gamma)}^2 |\Gamma_{L-l}|^{-2\alpha/\sigma} \lesssim \left(M^{-\alpha} \exp\left(-l \frac{\beta_w}{\gamma + \beta_s}\right) \exp\left(-\frac{\alpha}{\sigma + \alpha}(L-l)\right)\right)^2.$$

Combining these observations we arrive at

$$\begin{aligned} (\star) &\lesssim \exp\left(-L \frac{\beta_w}{\gamma + \beta_s}\right) + M^{-\alpha} \sum_{l=0}^L \exp\left(-\frac{\alpha}{\sigma + \alpha}(L-l) - l \frac{\beta_s}{\gamma + \beta_s}\right) \\ &\lesssim \exp\left(-L \frac{\beta_w}{\gamma + \beta_s}\right) + M^{-\alpha} \exp\left(-L \frac{\alpha}{\sigma + \alpha}\right) \sum_{l=0}^L \exp\left(l \left(\frac{\alpha}{\sigma + \alpha} - \frac{\beta_s}{\gamma + \beta_s}\right)\right). \end{aligned}$$

From here, the proof may be concluded exactly as that of Theorem 4.1. \square

5 An adaptive algorithm

We introduce in this section an adaptive algorithm for the case when an optimal sequence of polynomial subspaces, the rate of convergence $f_l \rightarrow f_\infty$, or the cost for evaluations of f_l are unknown.

To describe our algorithm, we restrict ourselves to the case when $\Gamma = [0, 1]^d$, $d \in \mathbb{N}$ and when $\mu = \lambda$ is the Lebesgue measure. By the results in Section 3.2, we may then use samples and weights from the arcsine distribution instead of the optimal distributions. This allows us to keep previous samples when we extend the polynomial subspaces, whereas using the optimal distribution, which depends on the polynomial subspace, would require throwing away all samples each time the polynomial subspace is extended.

We next describe the building blocks that are used by our adaptive algorithm to select polynomial approximation subspaces.

Definition 5.1 (Multivariate Legendre polynomials).

(i) We denote by $(P_i)_{i \in \mathbb{N}}$ the univariate $L_\lambda^2([0, 1])$ -orthonormal Legendre polynomials and define their tensor products

$$P_\eta := \bigotimes_{j=1}^d P_{\eta_j} : [0, 1]^d \rightarrow \mathbb{R},$$

$$P_\eta(\boldsymbol{\gamma}) := \prod_{j=1}^d P_{\eta_j}(\gamma_j)$$

for $\eta \in \mathbb{N}^d$.

(ii) For each multi-index $\mathbf{k} \in \mathbb{N}^d$, we define the polynomial subspace

$$\mathcal{P}_{\mathbf{k}} := \text{span}\{P_\eta : 2^{\mathbf{k}} - 1 \leq \eta < 2^{\mathbf{k}+1} - 1\} \subset L^2([0, 1]^d, \lambda).$$

Remark 6 (Orthonormal decomposition). Since polynomials are dense in $L_\lambda^2([0, 1]^d)$, the subspaces $(\mathcal{P}_{\mathbf{k}})_{\mathbf{k} \in \mathbb{N}^d}$ form an orthonormal decomposition of $L_\lambda^2([0, 1]^d)$. We use exponentially large subspaces instead of the simpler, one-dimensional subspaces $\mathcal{P}_{\mathbf{k}} = \mathbb{R} \cdot P_{\mathbf{k}}$ to avoid computational overhead resulting from slow construction of large polynomial subspaces.

We use the notation $f_{-1} := 0$ to avoid separate treatment of the term corresponding to $l = 0$ in the following. To describe a multilevel approximation, we need to construct a sequence $(V_k)_{k=0}^L$ of polynomial subspaces, such that the difference $f_l - f_{l-1}$ is projected onto V_{L-l} using weighted least squares approximation. The final approximation is then defined as

$$\sum_{l=0}^L \Pi_{L-l}(f_l - f_{l-1}). \quad (34)$$

where Π_k projects onto V_k for $0 \leq k \leq L$. As in Section 4, if the samples used by Π_k are distributed according to the optimal distribution of V_k , then we require that the number of samples N_k satisfy

$$\kappa \frac{N_k}{\log N_k} \geq \dim V_k \quad (35)$$

for some $\kappa > 0$. However, we then have to throw away all samples each time the space V_k is extended, as this changes the optimal distribution. As an alternative, we may use samples from the arcsine distribution, which is independent of the polynomial subspaces V_k and thus allows us to reuse samples and corresponding function evaluations. By Section 3.2, this increases the number of required samples only by a constant factor (that depends exponentially on the dimension d , however).

To construct the sequence of polynomial subspaces in an adaptive fashion, our algorithm constructs a (finite) downward closed multi-index set $\mathcal{I} \subset \mathbb{N}^{d+1}$. Given such a set, we let

$$V_k := \bigoplus_{\mathbf{k} \in \mathbb{N}^d: (\mathbf{k}, L-k) \in \mathcal{I}} \mathcal{P}_{\mathbf{k}} \quad 0 \leq k \leq L,$$

where

$$L := \max\{l \in \mathbb{N} : \exists \mathbf{k} \in \mathbb{N}^d \text{ s.t. } (\mathbf{k}, l) \in \mathcal{I}\} < \infty,$$

which means that we project the difference $f_l - f_{l-1}$ onto the subspace V_{L-l} that is determined by the slice $\mathcal{I}_l := \{\mathbf{k} \in \mathbb{N}^d : (\mathbf{k}, l) \in \mathcal{I}\}$ of the multi-index set \mathcal{I} . To construct \mathcal{I} , starting with $\mathcal{I} = \{\mathbf{0}\}$, our algorithm adds one multi-index at a time according to the following procedure, which resembles algorithms for adaptive sparse grid integration [18, 12]. We call

$$\mathcal{A} := \{(\mathbf{k}, l) \in \mathbb{N}^{d+1} \setminus \mathcal{I} : \mathcal{I} \cup \{(\mathbf{k}, l)\} \text{ is downward closed}\}$$

the set of *admissible multi-indices*.

- (i) For each admissible multi-index (\mathbf{k}, l) , we estimate the norm of the projection of $f_l - f_{l-1}$ onto $\mathcal{P}_{\mathbf{k}}$. This estimate represents the gain that is made by adding (\mathbf{k}, l) to \mathcal{I} . Furthermore, we estimate the work that adding this multi-index would incur.
- (ii) We expand \mathcal{I} by the multi-index that maximizes the ratio between the gain and work estimates.

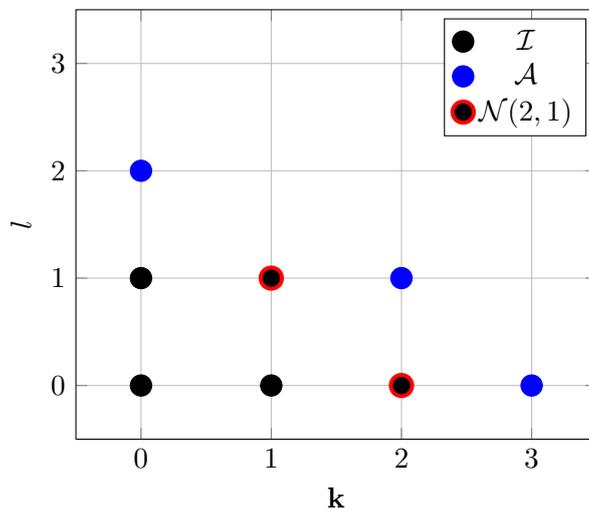


Figure 1: Example with $d = 1$ of a multi-index set \mathcal{I} and the associated set of admissible multi-indices \mathcal{A} , as well as neighbors $\mathcal{N}(2, 1)$ of $(2, 1) \in \mathcal{A}$. In this example $L = 1$, $V_1 = \text{span}\{1, \gamma, \dots, \gamma^6\} = \text{span}\{P_0(\gamma), \dots, P_6(\gamma)\}$, and $V_0 = \text{span}\{1, \gamma, \gamma^2\} = \text{span}\{P_0(\gamma), P_1(\gamma), P_2(\gamma)\}$.

We next explain how we arrive at the gain and work estimates that are required in step (i). To estimate the norm of the orthogonal projection of $f_l - f_{l-1}$ onto $\mathcal{P}_{\mathbf{k}}$, we compute the arithmetic average of corresponding estimates for the neighbors $\mathcal{N}(\mathbf{k}, l) = \{(\mathbf{k}^{(1)}, l^{(1)}), \dots\}$ of (\mathbf{k}, l) in \mathcal{I} . Here, by neighbor we mean elements of \mathcal{I} that differ from (\mathbf{k}, l) in a single entry by 1, see Figure 1. For each such neighbor, we estimate the norm of the orthogonal projection $\text{Proj}_{\mathbf{k}^{(j)}}(f_{l^{(j)}} - f_{l^{(j)}-1})$ of $f_{l^{(j)}} - f_{l^{(j)}-1}$ onto $\mathcal{P}_{\mathbf{k}^{(j)}}$ simply by computing the Euclidean norm of those basis coefficients of $\Pi_{L-l^{(j)}}(f_{l^{(j)}} - f_{l^{(j)}-1})$ that belong to $\mathcal{P}_{\mathbf{k}^{(j)}}$. (Recall that $\Pi_{L-l^{(j)}}$ is a discrete projection onto the space $V_{L-l^{(j)}}$ of which $\mathcal{P}_{\mathbf{k}^{(j)}}$ is a subspace since $\mathbf{k}^{(j)} \in \mathcal{I}_{l^{(j)}}$.) The final estimate can be expressed as

$$\frac{1}{|\mathcal{N}(\mathbf{k}, l)|} \sum_{j=1}^{|\mathcal{N}(\mathbf{k}, l)|} \|\text{Proj}_{\mathbf{k}^{(j)}} \Pi_{L-l^{(j)}}(f_{l^{(j)}} - f_{l^{(j)}-1})\|_{L_\lambda^2([0,1]^d)}.$$

To estimate the work that adding (\mathbf{k}, l) to \mathcal{I} incurs, we observe that Equation (35) tells us exactly how many new samples are needed. More specifically, if we denote by $N(\mathcal{I}_l)$ the minimal solution of Equation (35) for the polynomial subspace determined by \mathcal{I}_l , then the required number of new samples of $f_l - f_{l-1}$ is $N(\mathcal{I}_l \cup \{\mathbf{k}\}) - N(\mathcal{I}_l)$. It therefore remains to determine the work per sample, $\text{Work}(f_l - f_{l-1})$. If this work is unknown, then we store for each level l an estimate, which

we update with the observed computational work divided by the number of generated samples each time \mathcal{I}_l changes. The final estimate of the work associated with (\mathbf{k}, l) is

$$\text{Work}(f_l - f_{l-1}) \cdot (N(\mathcal{I}_l \cup \{\mathbf{k}\}) - N(\mathcal{I}_l)).$$

Algorithm 1 gives a summary of our algorithm in pseudocode.

Algorithm 1. Adaptive multilevel algorithm.

```

1: function MLA( $((f_l)_{l \in \mathbb{N}}, \text{STEPS})$ )
2:    $\mathcal{I} \leftarrow \{\mathbf{0}\}$ 
3:    $X_l \leftarrow \emptyset \forall l \in \mathbb{N}$ 
4:    $\Delta_l \leftarrow 0 \forall l \in \mathbb{N}$ 
5:   for  $0 \leq i < \text{STEPS}$  do
6:      $(\mathbf{k}, l) \leftarrow \arg \max_{(\mathbf{k}, l) \in \mathcal{A}} \frac{\text{GAIN}((\mathbf{k}, l), (\Delta_l)_{l \in \mathbb{N}}, \mathcal{I})}{\text{WORK}((\mathbf{k}, l), \mathcal{I})}$ 
7:      $N_+ \leftarrow N(\mathcal{I}_l \cup \{\mathbf{k}\}) - N(\mathcal{I}_l)$ 
8:      $\mathcal{I} \leftarrow \mathcal{I} \cup \{(\mathbf{k}, l)\}$ 
9:     for  $0 \leq j < N_+$  do
10:      Generate  $\gamma \sim p_d^\infty$ 
11:       $y \leftarrow (f_l - f_{l-1})(\gamma)$ 
12:       $X_l \leftarrow X_l \cup \{(\gamma, y)\}$ 
13:    end for
14:     $\Delta_l \leftarrow \Pi_{L-l}(f_l - f_{l-1})$ 
15:  end for
16:  return  $\sum_{0 \leq l \leq L} \Delta_l$ 
17: end function

18: function GAIN( $((\mathbf{k}, l), (\Delta_l)_{l \in \mathbb{N}}, \mathcal{I})$ )
19:    $s = 0$ 
20:   for  $(\mathbf{k}^{(j)}, l^{(j)}) \in \mathcal{N}(\mathbf{k}, l)$  do
21:      $s \leftarrow s + \|\text{Proj}_{\mathbf{k}^{(j)}} \Delta_{l^{(j)}}\|_{L_\lambda^2}$ 
22:   end for
23:   return  $s/|\mathcal{N}(\mathbf{k}, l)|$ 
24: end function

25: function WORK( $((\mathbf{k}, l), \mathcal{I})$ )
26:   return  $\text{Work}(f_l - f_{l-1}) \cdot (N(\mathcal{I}_l \cup \{\mathbf{k}\}) - N(\mathcal{I}_l))$ 
27: end function

```

6 Application to parametric PDE

We assume in this section that $u(\cdot, \gamma)$ is the solution of some partial differential equation (PDE) with parameters $\gamma \in \Gamma \subset \mathbb{R}^d$ and that we are interested in the *response surface*

$$\gamma \mapsto f_\infty(\gamma) := Q(u(\cdot, \gamma)) \in \mathbb{R},$$

where $Q(u(\cdot, \gamma))$ is a real-valued quantity of interest, such as a point evaluation, a spatial average, or a maximum. In most situations, we cannot evaluate $f_\infty(\gamma)$ exactly, as this would require an analytic solution of the PDE. Instead, we have to work with discretized solutions $u_n(\cdot, \gamma)$ for each γ , which yield approximate response surfaces

$$\begin{aligned} f_n &: \Gamma \rightarrow \mathbb{R} \\ \gamma &\mapsto Q(u_n(\cdot, \gamma)). \end{aligned}$$

For example, if we employ finite element discretizations with maximal element diameter $h := n^{-1}$, then the work required for evaluations of f_n grows like $h^{-\gamma} = n^\gamma$ for some $\gamma > 0$. To apply the multilevel method of Section 4, we need to verify the remaining Assumptions A1 and A2 from there.

As a motivating example, we consider a linear elliptic second order PDE, which has been extensively studied in recent years [17, 4, 6, 1],

$$\begin{aligned} -\nabla \cdot (a(x, \gamma) \nabla u(x, \gamma)) &= g(x) & \text{in } U \subset \mathbb{R}^D \\ u(x, \gamma) &= 0 & \text{on } \partial U, \end{aligned} \quad (36)$$

with $a : U \times \Gamma \rightarrow \mathbb{R}$ and $\Gamma := [0, 1]^d$.

Proposition 6.1. *For any $n \in \mathbb{N}$, let u_n be finite element approximations of order $r \geq 1$ and maximal element diameter $h := (n+1)^{-1}$, and let $f_n(\gamma) := Q(u_n(\cdot, \gamma))$. Assume that g and U are sufficiently smooth, that*

$$\inf_{x \in U, \gamma \in \Gamma} a(x, \gamma) > 0, \quad (37)$$

and that Q is a continuous linear functional on $L^2(U)$.

(i) *If $a \in C^r(U \times \Gamma)$ for some $r \geq 1$, then*

$$\|f_\infty - f_n\|_{L^2(\Gamma)} \lesssim h^{r+1}$$

and

$$\|f_\infty - f_n\|_{C^{r-1}(\Gamma)} \lesssim h^2.$$

(ii) *If for some $r, s \geq 1$ we have*

$$a \in C^r(U) \otimes C^s(\Gamma) := \{a : U \times \Gamma \rightarrow \mathbb{R} : \|\partial_x^r \partial_\gamma^s a\|_{C^0(U \times \Gamma)} < \infty \forall |\mathbf{r}|_1 \leq r, |\mathbf{s}|_1 \leq s\}, \quad (38)$$

then

$$\|f_\infty - f_n\|_{C^s(\Gamma)} \lesssim h^{r+1}.$$

Proof. In both cases, the standard theory of second order elliptic differential equations shows that $\gamma \mapsto u(\cdot, \gamma)$ is well defined as a map from Γ into $H^{r+1}(U)$, with

$$\|u\|_{L^\infty(\Gamma; H^{r+1}(U))} < \infty.$$

Next, we observe that the derivatives $\partial_{\gamma_j} u(\cdot, \gamma)$, $j \in \{1, \dots, d\}$ satisfy PDEs with the same operator as in Equation (36) but with new right-hand sides

$$\tilde{g}(x) := \nabla \cdot (\partial_{\gamma_j} a(x, \gamma) \nabla u(x, \gamma)).$$

The regularity of this right-hand side now depends on the assumptions on the coefficient a . In case (i) we have $\partial_{\gamma_j} a(\cdot, \gamma) \in C^{r-1}(U)$ and thus $\tilde{g} \in H^{r-2}(U)$. Therefore, $\partial_{\gamma_j} u(\cdot, \gamma) \in H^r(U)$ for each $\gamma \in \Gamma$ and, moreover, we have the uniform estimate

$$\|\partial_{\gamma_j} u\|_{L^\infty(\Gamma; H^r(U))} < \infty.$$

In case (ii) we have $\partial_{\gamma_j} a(\cdot, \gamma) \in C^r(U)$ and thus $\tilde{g} \in H^{r-1}(U)$. Therefore, $\partial_{\gamma_j} u(\cdot, \gamma) \in H^{r+1}(U)$ for each $\gamma \in \Gamma$ and, moreover, we have the uniform estimate

$$\|\partial_{\gamma_j} u\|_{L^\infty(\Gamma; H^{r+1}(U))} < \infty.$$

Repeatedly applying these arguments yields

$$\|u\|_{C^{r-1}(\Gamma; H^2(U))} < \infty,$$

and

$$\|u\|_{C^s(\Gamma; H^{r+1}(U))} < \infty,$$

in cases (i) and (ii), respectively. We may now conclude by using standard finite-element theory. In case (i), we have

$$\begin{aligned} \|f_\infty - f_n\|_{L^2(\Gamma)} &\leq \|Q\| \|u - u_n\|_{L^2(\Gamma; L^2(U))} \\ &\lesssim h^{r+1} \|u\|_{L^2(\Gamma; H^{r+1}(U))} \end{aligned}$$

and

$$\begin{aligned} \|f_\infty - f_n\|_{C^{r-1}(\Gamma)} &\lesssim \|u - u_n\|_{C^{r-1}(\Gamma; L^2(U))} \\ &\lesssim h^2 \|u\|_{C^{r-1}(\Gamma; H^2(U))}, \end{aligned}$$

whereas in case (ii), we have

$$\begin{aligned} \|f_\infty - f_n\|_{C^s(\Gamma)} &\lesssim \|u - u_n\|_{C^s(\Gamma; L^2(U))} \\ &\lesssim h^{r+1} \|u\|_{C^s(\Gamma; H^{r+1}(U))}, \end{aligned}$$

□

Remark 7. In case (i) of the previous proposition, differentiating with respect to γ reduces the number of available derivatives in x , which are required for convergence of the finite element method. Thus, the convergence in $L^2(\Gamma)$ is faster than that in $C^{r-1}(\Gamma)$. Case (ii), on the other hand, describes the so-called mixed smoothness of the coefficient in x and γ , meaning that differentiating in γ does not affect the differentiability with respect to x .

If the coefficients depend analytically on γ , then the same holds for f_∞ , which can be exploited to obtain algebraic polynomial approximability rates of f_∞ even in the case of infinite-dimensional parameters [4, 14], as shown below.

Proposition 6.2. Let $\Gamma := [-1, 1]^\infty$. Assume that Q is a linear and continuous functional on $L^2(U)$, that $0 < \inf_{x, \gamma} a(x, \gamma) \leq \sup_{x, \gamma} a(x, \gamma) < \infty$, and that

$$\begin{aligned} a(x, \gamma) &= \bar{a}(x) + \sum_{j=0}^{\infty} \gamma_j \psi_j(x), \\ a(x, \gamma) &= \bar{a}(x) + \left(\sum_{j=0}^{\infty} \gamma_j \psi_j(x) \right)^2, \end{aligned}$$

or

$$a(x, \gamma) = \exp \left(\sum_{j=0}^{\infty} \gamma_j \psi_j(x) \right).$$

If there exists $r_{\max} > 1$ such that

$$\|\psi_j\|_{C^r(U)} \lesssim (j+1)^{-(r_{\max}+1-r)} \quad \forall j \in \mathbb{N}, \quad 0 \leq r < r_{\max},$$

then, for any $r \in \mathbb{N}$ with $1 \leq r < r_{\max}$, finite element approximations with maximal element diameter $h := (n+1)^{-1}$ achieve

$$\|f_\infty - f_n\|_{L^\infty(\Gamma)} \leq Ch^{r+1}$$

with a constant C independent of n . Furthermore, for any such r , there is a sequence $(V_m)_{m \in \mathbb{N}}$ of downward closed polynomial spaces with $\dim V_m = m$ such that finite element approximations with order r and maximal diameter $h := (n+1)^{-1}$ achieve

$$e_{V_m, 1, \infty}(f_\infty - f_n) \leq C(m+1)^{-\alpha} h^{r+1} \quad \forall 0 < \alpha < r_{\max} - r$$

with a constant C independent of n and m .

Proof. It was shown in [4, Theorem 4.1 & Section 5] that for each $0 \leq r < r_{\max}$ there exists a set $\Gamma_r \subset \mathbb{C}^\infty$, $\Gamma \subset \Gamma_r$ such that $\|a\|_{L^\infty(\Gamma_r; C^r(U))} < \infty$ and such that $\gamma \mapsto u(\cdot, \gamma)$ may be extended to a complex differentiable map from Γ_r into $H^{1+r}(U)$ with

$$\|u\|_{L^\infty(\Gamma_r; H^{1+r}(U))} < \infty \quad (39)$$

For a detailed description of the sets Γ_r we refer to [4]. For our purposes it suffices to know that the better the summability of $(\|\psi_j\|_{C^r(U)})_{j \in \mathbb{N}}$, the larger Γ_r can be chosen; and the larger Γ_r the better the polynomial approximability properties of complex differentiable maps defined on Γ_r . In particular, the results of [4, Section 2], show that when restricted to the smaller set Γ such maps may be approximated at algebraic convergence rates within downward closed polynomial subspaces. More specifically, [4, Equation (2.27)] shows that if a function e is complex differentiable on Γ_r , then for any $m \in \mathbb{N}$ there exists a downward closed polynomial subspace V_m such that

$$\inf_{\tilde{v} \in V_m \otimes L^2(U)} \|e - \tilde{v}\|_{L^\infty(\Gamma; L^2(U))} \lesssim (m+1)^{-\alpha} \|e\|_{L^\infty(\Gamma_r; L^2(U))}$$

for all $\alpha < r_{\max} - r$. Applying this estimate with $e := u - u_n$ shows

$$\begin{aligned} \inf_{v \in V_m} \|(f_\infty - f_n) - v\|_{L^\infty(\Gamma)} &\leq \|Q\| \inf_{\tilde{v} \in V_m \otimes L^2(U)} \|(u - u_n) - \tilde{v}\|_{L^\infty(\Gamma; L^2(U))} \\ &\lesssim (m+1)^{-\alpha} \|u - u_n\|_{L^\infty(\Gamma_r; L^2(U))}. \end{aligned}$$

By standard finite element analysis we finally obtain

$$\|u - u_n\|_{L^\infty(\Gamma_r; L^2(U))} \leq Ch^{r+1} \|u\|_{L^\infty(\Gamma_r; H^{r+1}(U))}.$$

with $C = C(\|a\|_{L^\infty(\Gamma_r; C^r(U))}) < \infty$. Combining the previous two estimates with Equation (39) concludes the proof. \square

Remark 8. *Similar results can also be shown for PDEs of parabolic type and for some nonlinear PDEs [4].*

7 Numerical Experiments

To support our theoretical analysis, we performed numerical experiments on linear elliptic parametric PDEs of the form

$$\begin{aligned} -\nabla \cdot (a(x, \gamma) \nabla u(x, \gamma)) &= 1 \quad \text{in } U := [-1, 1]^D \\ u(x, \gamma) &= 0 \quad \text{on } \partial U, \end{aligned} \quad (40)$$

as in Section 6. We let

$$a(x, \gamma) = 1 + \|x\|_2^r + \|\gamma\|_2^s, \quad \gamma \in \Gamma := [-1, 1]^d$$

for $r := 1$, $s := 3$, $D := 2$ and $d \in \{2, 3, 4, 6\}$. Our goal was to approximate the response surface

$$\gamma \mapsto f(\gamma) := Q(u(\cdot, \gamma)) := 0.5 \int_U u(\cdot, \gamma) \, dx$$

in $L^2(\Gamma)$. The numerical scheme we used to solve Equation (40) employs centered finite difference approximations to the derivatives with a constant mesh size, h . Such a numerical scheme converges asymptotically at a rate of $\mathcal{O}(h^2)$ in the L^2 norm and requires a computational work of $\mathcal{O}(h^{-2})$, since the PDE is two-dimensional. This corresponds to the values $\beta_s = \beta_w = 2$ and $\gamma = 2$ for the parameters in Assumptions A2 and A3. To estimate the projection error of our estimate we evaluate the L^2 error norm using Monte Carlo sampling with $M = 1000$ samples,

$$\|f - S_L(f)\|_{L^2(\Gamma)}^2 \approx \frac{1}{M} \sum_{j=1}^M (f_{L+1}(\gamma_j) - S_L(f)(\gamma_j))^2. \quad (41)$$

In our tests we employ both the nonadaptive and the adaptive algorithms from Sections 4 and 5. As a basis for the nonadaptive algorithm, we use total degree polynomial spaces $V_m := \text{span}\{P_\eta : |\eta|_1 \leq m\}$, where P_η is a tensor product of Legendre polynomials as in Section 5. We also compare the multilevel algorithm to the straightforward, single-level approach, which for a given polynomial approximation space V_m uses samples from a fixed PDE discretization level that matches the accuracy of the polynomial best approximation in V_m . To find these matching PDE discretization levels, we consider the complexity curve of the single-level method as the lower envelope of complexity curves with different PDE discretization levels. Even though such a method is not practical, the choice of discretization level for a given tolerance is always optimal. The random points were sampled from the optimal distribution as explained in Section 3.1.

Before presenting the numerical results, let us derive some a-priori estimates of the complexity of the single-level and multilevel projection methods. From Proposition 6.1, if $a \in C^r(U) \otimes C^s(\Gamma)$, then using finite elements of order r and mesh size h would yield convergence in the space $F := C^s(\Gamma)$ with the values $\beta_s = \beta_w = r + 1$ of the parameters in Section 4, and optimal solvers would require the work $\mathcal{O}(h^{-\gamma})$, $\gamma := D$. Furthermore, since functions in $C^s(\Gamma)$ are approximable by polynomials of total degree less than or equal to k at the rate $\mathcal{O}(k^{-s})$ in the supremum norm [2], we expect at least $\alpha = s$. Even though our choice $a(x, \gamma) = 1 + \|x\|_2^2 + \|\gamma\|_2^2$ satisfies only $a \in C^{r-1,1}(U) \otimes C^{s-1,1}(\Gamma)$, we do not expect different rates than those derived above for $a \in C^r(U) \otimes C^s(\Gamma)$. Finally, the dimension of total degree polynomial spaces V_m equals $\binom{m+d}{d}$ and asymptotically we have $\binom{m+d}{d} \lesssim m^d$.

Thus, we expect the complexity of the single-level method to be $\mathcal{O}\left(\epsilon^{-\frac{D}{r+1} - \frac{d}{s}} \log(\epsilon^{-1})\right)$, while the complexity of the multilevel method is of $\mathcal{O}\left(\epsilon^{-\max(\frac{D}{r+1}, \frac{d}{s})} \log(\epsilon^{-1})^t\right)$, where

$$t = \begin{cases} 1 & \frac{D}{r+1} > \frac{d}{s}, \\ 3 + \frac{D}{r+1} & \frac{D}{r+1} = \frac{d}{s}, \\ 2 & \frac{D}{r+1} < \frac{d}{s}. \end{cases}$$

Hence, for $r = 1$ and $s = 3$, the complexity of the single-level method is $\mathcal{O}\left(\epsilon^{-1 - \frac{d}{3}} \log(\epsilon^{-1})\right)$ and the complexity of the multilevel method is $\mathcal{O}\left(\epsilon^{-\max(1, \frac{d}{3})} \log(\epsilon^{-1})^t\right)$ where

$$t = \begin{cases} 1, & d < 3, \\ 4, & d = 3, \\ 2, & d > 3. \end{cases}$$

Figure 2 shows the work estimate as defined in Equation (25) versus the L^2 error approximation in Equation (41). The results for the multilevel algorithm displayed there were obtained with the work parameter $\sigma := d/2$, which we found describes the pre-asymptotic behavior of $\dim V_m = \binom{m+d}{d}$ better than $\sigma := d$. The theoretical rates satisfactorily match the obtained numerical rates, which show an improvement of the multilevel methods over the single-level method. Note that the work estimate does *not* include the cost of generating points or the cost of assembling the projection matrix and computing the projection. On the other hand, these costs are included in Figure 3, which shows the total running time in seconds of the three different methods. While these plots still show the same complexity rates as Figure 2 for all three methods for sufficiently small errors, these plots also show the overhead of the multilevel methods, especially as d increases. The overhead of the adaptive algorithm for the multilevel method is especially significant and more work needs to be done to reduce it.

8 Conclusion

We have presented a novel multilevel projection method for the approximation of response surfaces using multivariate polynomials and random samples with different accuracies. For this purpose, we have discussed and analyzed various sampling methods for the underlying single-level approximation method. We have then presented theoretical and numerical results on our multilevel projection

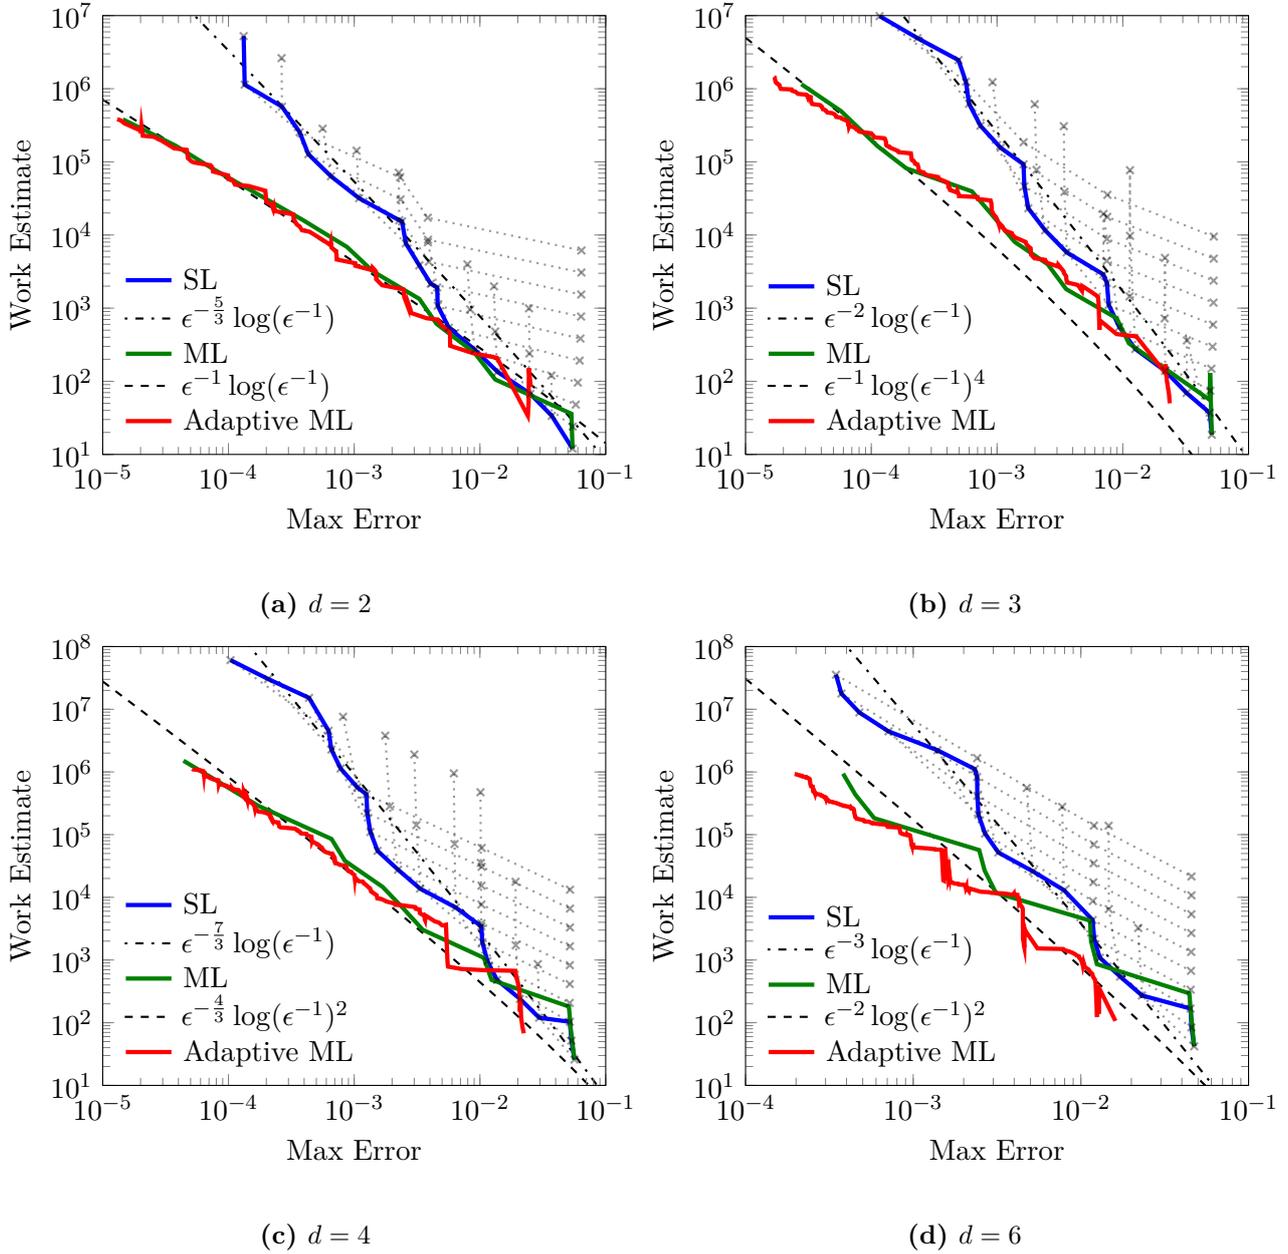


Figure 2: $L^2([-1, 1]^d)$ -error, approximated using Equation (41) vs work estimate Equation (25) of single-level (SL), multilevel (ML) and adaptive ML (ML adaptive) methods for a linear elliptic PDE with non-smooth parameter dependence. The grey dotted lines are the complexity curves of different runs of the single-level, each with a different PDE discretization level. The single-level (SL) complexity curve is then the lower envelope of all single-level complexity curves. This figure shows the agreement of the numerical results with the theoretical rates.

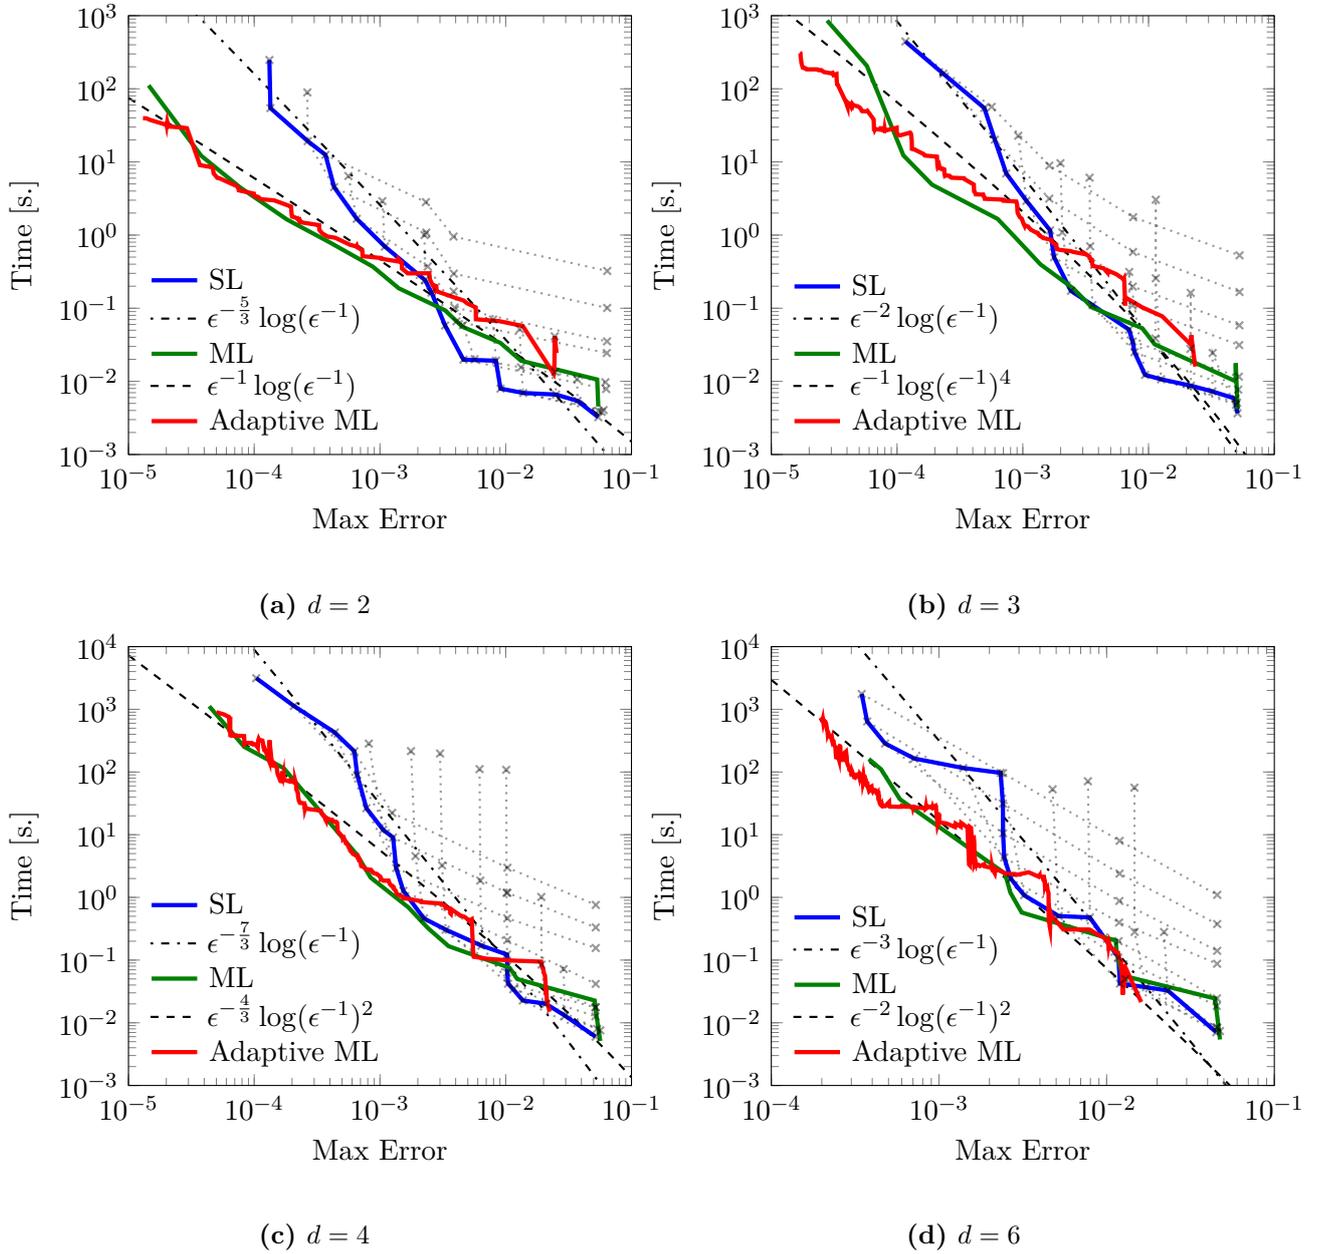


Figure 3: Similar to Figure 2, but showing the total running time of the methods instead of their work estimate. The discrepancy of the two figures is due to the overhead of sampling the points, assembling the projection matrix and computing the projection. Moreover, this plot shows the overhead of the adaptive algorithm compared to the non-adaptive one.

method for problems in which samples can be obtained at different accuracies. The numerical results show good agreement with the computational gains predicted by our theory. Future work will address the application to problems in uncertainty quantification with infinite-dimensional parameter domains and multi- or infinite-dimensional quantities of interest.

Acknowledgements F. Nobile received support from the Center for ADvanced MOdeling Science (CADMOS). R. Tempone and S. Wolfers are members of the KAUST SRI Center for Uncertainty Quantification in Computational Science and Engineering. R. Tempone received support from the KAUST CRG3 Award Ref:2281 and the KAUST CRG4 Award Ref:2584.

References

- [1] Ivo Babuska, Raúl Tempone, and Georgios E Zouraris. Galerkin finite element approximations of stochastic elliptic partial differential equations. *SIAM Journal on Numerical Analysis*, 42(2):800–825, 2004.
- [2] Thomas Bagby, Len Bos, and Norman Levenberg. Multivariate simultaneous approximation. *Constructive Approximation*, 18(4):569, December 2002.
- [3] Abdellah Chkifa, Albert Cohen, Giovanni Migliorati, Fabio Nobile, and Raul Tempone. Discrete least squares polynomial approximation with random evaluations - application to parametric and stochastic elliptic PDEs. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(3):815–837, 2015.
- [4] Abdellah Chkifa, Albert Cohen, and Christoph Schwab. Breaking the curse of dimensionality in sparse polynomial approximation of parametric PDEs. *Journal de Mathématiques Pures et Appliquées*, 103(2):400–428, 2015.
- [5] Albert Cohen, Mark A Davenport, and Dany Leviatan. On the stability and accuracy of least squares approximations. *Foundations of computational mathematics*, 13(5):819–834, 2013.
- [6] Albert Cohen, Ronald Devore, and Christoph Schwab. Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDEs. *Analysis and Applications*, 9(01):11–47, 2011.
- [7] Albert Cohen and Giovanni Migliorati. Optimal weighted least-squares methods. *arXiv*, 2016.
- [8] Manas K Deb, Ivo M Babuška, and J Tinsley Oden. Solution of stochastic partial differential equations using galerkin finite element techniques. *Computer Methods in Applied Mechanics and Engineering*, 190(48):6359–6372, 2001.
- [9] Ronald A DeVore. Nonlinear approximation. *Acta numerica*, 7:51–150, 1998.
- [10] Dinh Dũng, Vladimir N. Temlyakov, and Tino Ullrich. Hyperbolic Cross Approximation. *arXiv*, 2015.
- [11] James E. Gentle. *Random number generation and Monte Carlo methods*. Statistics and Computing. Springer, New York, second edition, 2003.
- [12] Thomas Gerstner and Michael Griebel. Dimension-adaptive tensor-product quadrature. *Computing*, 71(1):65–87, 2003.
- [13] Michael Griebel and Christian Rieger. Reproducing kernel Hilbert spaces for parametric partial differential equations. submitted to SIAM Journal on Uncertainty Quantification, also available as INS Preprint No. 1511, 2015.
- [14] Abdul-Lateef Haji-Ali, Fabio Nobile, Lorenzo Tamellini, and Raúl Tempone. Multi-index stochastic collocation convergence rates for random PDEs with parametric regularity. *Foundations of Computational Mathematics*, 16(6):1555–1605, 2016.
- [15] Abdul-Lateef Haji-Ali, Fabio Nobile, Lorenzo Tamellini, and Raúl Tempone. Multi-index stochastic collocation for random PDEs. *Computer Methods in Applied Mechanics and Engineering*, 306:95–122, 2016.

- [16] Jerrad Hampton and Alireza Doostan. Coherence motivated sampling and convergence analysis of least squares polynomial chaos regression. *Computer Methods in Applied Mechanics and Engineering*, 290:73–97, 2015.
- [17] Helmut Harbrecht, Michael Peters, and Markus Siebenmorgen. Multilevel accelerated quadrature for PDEs with log-normally distributed diffusion coefficient. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):520–551, 2016.
- [18] Markus Hegland. Adaptive sparse grids. *ANZIAM Journal*, 44((C)):C335–C353, 2002.
- [19] Stefan Heinrich. Multilevel Monte Carlo methods. In *International Conference on Large-Scale Scientific Computing*, pages 58–67. Springer, 2001.
- [20] Frances Kuo, Robert Scheichl, Christoph Schwab, Ian Sloan, and Elisabeth Ullmann. Multilevel quasi-Monte Carlo methods for lognormal diffusion problems. *Mathematics of Computation*, 2017.
- [21] Olivier Le Maître and Omar Knio. *Spectral methods for uncertainty quantification*. Springer, 2010.
- [22] Eli Levin and Doron S Lubinsky. Christoffel functions, orthogonal polynomials, and Nevai’s conjecture for Freud weights. *Constructive Approximation*, 8(4):463–535, 1992.
- [23] Jun S Liu. Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statistics and Computing*, 6(2):113–119, 1996.
- [24] Jun S Liu. *Monte Carlo strategies in scientific computing*. Springer Science & Business Media, 2008.
- [25] Giuseppe Mastroianni and Vilmos Totik. Weighted polynomial inequalities with doubling and a_∞ weights. *Constructive approximation*, 16(1):37–71, 2000.
- [26] Giovanni Migliorati, Fabio Nobile, and Raúl Tempone. Convergence estimates in probability and in expectation for discrete least squares with noisy evaluations at random points. *Journal of Multivariate Analysis*, 142:167–182, 2015.
- [27] Akil Narayan, John Jakeman, and Tao Zhou. A Christoffel function weighted least squares algorithm for collocation approximations. *Mathematics of Computation*, 86(306):1913–1947, 2017.
- [28] Paul Nevai, Tamás Erdélyi, and Alphonse P Magnus. Generalized Jacobi weights, Christoffel functions, and Jacobi polynomials. *SIAM Journal on Mathematical Analysis*, 25(2):602–614, 1994.
- [29] Fabio Nobile, Raul Tempone, and Soeren Wolfers. Sparse approximation of multilinear problems with applications to kernel-based methods in UQ. *arXiv. Submitted to Numerische Mathematik*, 2016.
- [30] Alfio Quarteroni. Some results of Bernstein and Jackson type for polynomial approximation in l^p -spaces. *Japan Journal of Applied Mathematics*, 1(1):173–181, 1984.
- [31] Gábor Szegő. *Orthogonal polynomials*. American Mathematical Society, Providence, R.I., fourth edition, 1975. American Mathematical Society, Colloquium Publications, Vol. XXIII.
- [32] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.*, 12(4):389–434, 2012.