CONSTRUCTIVE
APPROXIMATION

Optimal Sampling and Christoffel Functions on General Domains

Matthieu Dolbeault[1] · Albert Cohen[1]

Abstract

We consider the problem of reconstructing an unknown function $u \in L^2(D, \mu)$ from its evaluations at given sampling points $x^1, \ldots, x^m \in D$, where $D \subset \mathbb{R}^d$ is a general domain and μ a probability measure. The approximation is picked from a linear space V_n of interest where $n = \dim(V_n)$. Recent results (Cohen and Migliorati in SMAI J Comput Math 3:181–203, 2017, Doostan and Hampton in Comput Methods Appl Mech Eng 290:73–97, 2015, Jakeman et al. in Math Comput 86:1913–1947, 2017) have revealed that certain weighted least-squares methods achieve near best (or instance optimal) approximation with a sampling budget m that is proportional to n, up to a logarithmic factor $\ln(2n/\varepsilon)$, where $\varepsilon > 0$ is a probability of failure. The sampling points should be picked at random according to a well-chosen probability measure σ whose density is given by the inverse Christoffel function that depends both on V_n and μ. While this approach is greatly facilitated when D and μ have tensor product structure, it becomes problematic for domains D with arbitrary geometry since the optimal measure depends on an orthonormal basis of V_n in $L^2(D, \mu)$ which is not explicitly given, even for simple polynomial spaces. Therefore, sampling according to this measure is not practically feasible. One computational solution recently proposed in Adcock and Huybrechs (Approximating smooth, multivariate functions on irregular domains, forum of mathematics, sigma, Cambridge University Press, Cambridge, 2020) relies on using the restrictions of an orthonormal basis of V_n defined on a simpler bounding domain and sampling according to the original probability measure μ, in turn giving up on the optimal sampling budget $m \sim n$. In this paper, we discuss practical sampling strategies, which amounts to using a perturbed measure $\widetilde{\sigma}$ that can be computed in an offline stage, not involving the measurement of u, as recently proposed in Adcock and Cardenas (SIAM J Math Data Sci 2:607–630, 2020) and Migliorati (IMA J Numer Anal, 2020. https://doi.org/10.1093/imanum/draa023). We show that near best approximation is attained by the resulting weighted least-squares method at near-optimal sampling budget and we discuss multilevel approaches that preserve opti-

Extended author information available on the last page of the article

mality of the cumulated sampling budget when the spaces V_n are iteratively enriched. These strategies rely on the knowledge of a-priori upper bounds $B(n)$ on the inverse Christoffel function for the space V_n and the domain D. We establish bounds of the form $\mathcal{O}(n^r)$ for spaces V_n of multivariate algebraic polynomials of given total degree, and for general domains D. The exact growth rate r depends on the regularity of the domain, in particular $r = 2$ for domains with Lipschitz boundaries and $r = \frac{d+1}{d}$ for smooth domains.

Keywords Least-squares approximation · Optimal recovery · Random sampling · Christoffel functions · Multivariate polynomials

Mathematics Subject Classification 41A10 · 65D15 · 41A81 · 41A65 · 62E17

1 Introduction

1.1 Reconstruction from Point Samples

The process of reconstructing an unknown function u defined on a domain $D \subset \mathbb{R}^d$ from its sampled values

$$u^i = u(x^i)$$

at a set of point $x^1, \ldots, x^m \in D$ is ubiquitous in data science and engineering. The sampled values may be affected by noise, making critical the stability properties of the reconstruction process. Let us mention three very different settings for such reconstruction problems that correspond to different areas of applications:

(i) Statistical learning and regression: we observe m independent realizations (x^i, y^i) of a random variable $z = (x, y)$ distributed according to an unknown measure, where $x \in D$ and $y \in \mathbb{R}$, and we want to recover a function $x \mapsto v(x)$ that makes $|y - v(x)|$ as small as possible in some given sense. If we use the quadratic loss $\mathbb{E}(|y - v(x)|^2)$, the minimizer is given by the regression function

$$u(x) = \mathbb{E}(y|x).$$

and the observed y^i may be thought of as the observation of $u(x^i)$ affected by noise.

(ii) State estimation from measurements: the function u represents the distribution of a physical quantity (temperature, quantity of a contaminant, acoustic pressure) in a given spatial domain D that one is allowed to measure by sensors placed at m locations $x^1, \ldots, x^m$. These measurements can be affected by noise reflecting the lack of accuracy of the sensors.

(iii) Design of physical/computer experiments: u is a quantity of interest that depends on the solution f to a parametrized physical problem. For example, $f = f(x)$ could be the solution to a PDE that depends on a vector $x = (x_1, \ldots, x_d) \in D$ of

d physical parameters, and u could be the result of a linear form ℓ applied to f, that is, $u(x) = \ell(f(x))$. We use a numerical solver for this PDE as a black box to evaluate f, and therefore u, at m chosen parameter vectors $x^1, \ldots, x^m \in D$, and we now want to approximate u on the whole domain D from these computed values $u(x^i)$. Here, the discretization error of the solver may be considered as a noise affecting the true value.

Contrary to statistical learning, in the last two applications (ii) and (iii) the positions of the sample points x^i are not realizations of an unknown probability distribution. They can be selected by the user, which brings out the problem of choosing them in the best possible way. Indeed, measuring u at the sample points may be costly: in (ii) we need a new sensor for each new point, and in (iii) a new physical experiment or run of a numerical solver. Moreover, in certain applications, one may be interested in reconstructing many different instances of functions u. Understanding how to sample in order to achieve the best possible trade-off between the sampling budget and the reconstruction performance is one main motivation of this work. We first make our objective more precise by introducing some benchmarks for the performance of the reconstruction process and sampling budget.

1.2 Optimality Benchmarks

We are interested in controlling the distance

$$d(u, \widetilde{u}) := \|u - \widetilde{u}\|,$$

between u and its reconstruction $\widetilde{u} = \widetilde{u}(u^1, \ldots, u^m)$, measured in some given norm $\|\cdot\| = \|\cdot\|_V$, where V is a Banach function space that contains u.

For a given numerical method, the derivation of an error bound is always tied to some prior information on u. One most common way to express such a prior is in terms of membership of u to a restricted class of functions, for example a smoothness class. One alternate way is to express the prior in terms of approximability of u by particular finite dimensional spaces. It is well known that the two priors are sometimes equivalent: many classical smoothness classes can be characterized in terms of approximability in some given norm by classical approximation spaces such as algebraic or trigonometric polynomials, splines, or wavelets [10].

In this paper, we adopt the second point of view, describing u by its closeness to a given subspace $V_n \subset V$ of dimension n: defining the best approximation error

$$e_n(u) := \min_{v \in V_n} \|u - v\|,$$

our prior is that $e_n(u) \leq \varepsilon_n$ for some $\varepsilon_n > 0$. One of our motivations is the rapidly expanding field of *reduced order modeling* in which one searches for approximation spaces V_n which are optimally designed to approximate families of solutions to parametrized PDEs. Such spaces differ significantly from the above-mentioned classical examples. For example in the *reduced basis method*, they are generated by particular instances of solutions to the PDE for well-chosen parameter values. We

refer to [7] for a survey on such reduced modeling techniques and their approximation capability.

In this context, one first natural objective is to build a reconstruction map

$$(u^1, \ldots, u^m) \mapsto \widetilde{u}_n \in V_n,$$

that performs almost as good as the best approximation error. We say that a reconstruction map taking its value in V_n is *instance optimal* with constant $C_0 \geq 1$ if and only if

$$\|u - \widetilde{u}_n\| \leq C_0 \, e_n(u), \tag{1.1}$$

for any $u \in V$.

Obviously, instance optimality implies that if $u \in V_n$, the reconstruction map should return the exact solution $\widetilde{u}_n = u$. For this reason, instance optimality can only be hoped for if the sampling budget m exceeds the dimension n. This leads us to introduce a second notion of optimality: we say that the sample is *budget optimal* with constant $C_1 \geq 1$ if

$$m \leq C_1 \, n. \tag{1.2}$$

Let us stress that in many relevant settings, we do not work with a single space V_n but a sequence of nested spaces

$$V_1 \subset V_2 \subset \cdots \subset V_n \subset \ldots$$

so that $e_n(u)$ decreases as n grows. Such a hierarchy could either be fixed in advance (for example, when using polynomials of degree n), or adaptively chosen as we collect more samples (for example when using locally refined piecewise polynomials or finite element spaces). Ideally, we may wish that the constants C_0 and C_1 are independent of n. As it will be seen, a more accessible goal is that only one of the two constants is independent of n, while the other grows at most logarithmically in n.

Another way of relaxing instance optimality is to request the weaker property of *rate optimality*, which requires that for any $s > 0$ and $u \in V$,

$$\sup_{n \geq 1} n^s \, \|u - \widetilde{u}_n\|_V \leq C \sup_{n \geq 1} n^s \, e_n(u),$$

where $C \geq 1$ is a fixed constant. In other words, the approximant produced by the reconstruction method should converge at the same polynomial rate as the best approximation.

In the context where the spaces V_n are successively refined, even if the reconstruction method is instance and budget optimal for each value of n, the cumulated sampling budget until the n-th refinement step is in principle of the order

$$m(n) \sim 1 + 2 + \cdots + n \sim n^2,$$

if samples are picked independently at each step. A natural question is whether the samples used until stage k can be, at least partially, recycled for the computation of $\widetilde{u}_{k+1}$, in such a way that the cumulated sampling budget $m(n)$ remains of the optimal order $\mathcal{O}(n)$. This property will be ensured for example if for each n, the samples are picked at points $\{x^1, \ldots, x^{m(n)}\}$ that are the sections of a unique infinite sequence $\{x^m\}_{m \geq 1}$, with $m(n) \sim n$, which means that all previous samples are recycled. We refer to this property as *hierarchical* sampling. It is also referred to as *online machine learning* in the particular above-mentioned application area (i).

1.3 Objectives and Layout

The design of sampling and reconstruction strategies that combine budget and instance (or rate) optimality, together with the above progressivity prescription, turns out to be a difficult task, even for very classical approximation spaces V_n such as polynomials.

In Sect. 2, we illustrate this difficulty by first discussing the example of reconstruction by *interpolation* for which the sampling budget is optimal but instance optimality with error measured in the L^∞ norm generally fails by a large amount. We then recall recent results [4, 9, 11, 14] revealing that one can get much closer to these optimality objectives by *weighted least-squares* reconstruction methods. In this case, we estimate the approximation error in $V = L^2(D, \mu)$ where μ is an arbitrary but fixed probability measure. The sampling points are picked at random according to a different probability measure σ^* that depends on V_n and μ:

$$d\sigma^*(x) = \frac{k_n(x)}{n}\,d\mu(x).$$

Here k_n is the *inverse Christoffel function* defined by

$$k_n(x) = \sum_{j=1}^n |L_j(x)|^2,$$

where $(L_1, \ldots, L_n)$ is any $L^2(D, \mu)$-orthonormal basis of V_n. By Cauchy–Schwarz inequality, it is readily seen that this function is characterized by the extremality property

$$k_n(x) = \max_{v \in V_n} \frac{|v(x)|^2}{\|v\|^2},$$

where $\|v\| := \|v\|_V = \|v\|_{L^2(D,\mu)}$ (for notational simplicity, throughout the paper, we omit the obvious restriction $v \neq 0$ needed when optimizing quotients with numerators and denominators that are null when $v = 0$). Then, instance optimality is achieved in a probabilistic sense with a sampling budget m that is proportional to n, up to a logarithmic factor $(\ln(2n\varepsilon))$, where $\varepsilon > 0$ is a probability of failure which comes as

an additional term in the instance optimality estimate

$$\mathbb{E}(\|u - \widetilde{u}_n\|^2) \leq C_0 \, e_n(u)^2 + \mathcal{O}(\varepsilon).$$

It is important to notice that σ^* differs from μ and that the standard least-squares method using a sample drawn according to μ is generally *not* budget optimal in the sense that instance optimality requires m to be larger than the quantity

$$K_n := \|k_n\|_{L^\infty} = \sup_{x \in D} |k_n(x)| = \max_{v \in V_n} \frac{\|v\|_{L^\infty}^2}{\|v\|_{L^2}^2},$$

which may be much larger than n, for instance $\mathcal{O}(n^2)$ or worse, see [8] as well as Sect. 5.

While these results are in principle attractive since they apply to arbitrary spaces V_n, measures μ and domains D, the proposed sampling strategy is highly facilitated when D is a tensor-product domain and μ is the tensor-product of a simple univariate measure, so that an $L^2(D, \mu)$-orthonormal basis of V_n can be explicitly provided. This is the case for example when using multivariate algebraic or trigonometric polynomial spaces with μ being the uniform probability measure on $[-1, 1]^d$ or $[-\pi, \pi]^d$. For a general domain D with arbitrary—possibly irregular—geometry, the orthonormal basis cannot be explicitly computed, even for simple polynomial spaces. Therefore, sampling according to the optimal measure σ^* is not feasible.

Non-tensor product domains D come out naturally in all the above mentioned settings (i)-(ii)-(iii). For example, in design of physical/computer experiments, this reflects the fact that, while the individual parameters x_j could range in intervals I_j for $j = 1, \ldots, d$, not all values x in the rectangle $R = I_1 \times \cdots \times I_d$ are physically admissible. Therefore, the function u is only accessible and searched for in a limited domain $D \subset R$. Here we assume that the domain D is known to us, in the sense that membership in D of a point $x \in \mathbb{R}^d$ can be assessed at low numerical cost.

One practical solution proposed in [3] consists in sampling according to the measure μ and solving the least-squares problem using the restriction of an orthonormal basis of V_n defined on a simpler tensor product bounding domain, which generally gives rise to a frame. This approach is feasible for example when μ is the uniform probability measure. Due to the use of restricted bases, the resulting Gramian matrix which appears in the normal equations is ill-conditioned or even singular, which is fixed by applying a pseudo-inverse after thresholding the smallest singular values at some prescribed level. Budget optimality is generally lost in this approach since one uses μ as a sampling measure.

In this paper, we also work under the assumption that we are able to sample according to μ, but we take a different path, which is exposed in Sect. 3. In an *offline* stage, we compute an approximation $\widetilde{k}_n$ to the inverse Christoffel function, which leads to a measure $\widetilde{\sigma}$ that may be thought of as a perturbation of the optimal measure σ^*. We may then use $\widetilde{\sigma}$ to define the sampling points $\{x^1, \ldots, x^m\}$ and weights. In the *online* stage, we perform the weighted least-squares reconstruction strategy based on the measurement of u at these points. Our first result is that if $\widetilde{k}_n$ is equivalent to k_n, we

recover the stability and instance optimality results from [9] at near-optimal sampling budget $m \sim n \ln(2n/\varepsilon)$.

One approach for computing $\widetilde{k}_n$, recently proposed in [2, 18], consists in drawing a first sample $\{z^1, \ldots, z^M\}$ according to μ and defining $\widetilde{k}_n$ as the inverse Christoffel function with respect to the discrete measure associated to these points. In order to ensure an equivalence between k_n and $\widetilde{k}_n$ with high probability, the value of M needs to be chosen larger than K_n, which is unknown to us. This can be ensured by asking that M is larger than a known upper bound $B(n)$ for K_n. The derivation of such bounds for general domains is one of the objectives of this paper. We also propose an empirical strategy for choosing M that does not require the knowledge of an upper bound and appears to be effective in our numerical tests. In all cases, the size M of the offline sample could be of order substantially larger than $\mathcal{O}(n)$. However, this first set of points is only used in the offline stage to perform computations that produce the perturbed measure $\widetilde{\sigma}$, and *not* to evaluate the function u which, as previously explained, is the costly aspect in the targeted applications and could also occur for many instances of u. These more costly evaluations of u only take place in the online stage at the x^i, therefore at near-optimal sampling budget.

In the case where K_n, or its available bound $B(n)$, grows very fast with n, the complexity of the offline stage in this approach becomes itself prohibitive. In order to mitigate this defect, we introduce in Sect. 4 a multilevel approach where the approximation $\widetilde{k}_n$ of k_n is produced by successive space refinements

$$V_{n_1} \subset \cdots \subset V_{n_q}, \quad n_q = n,$$

which leads to substantial computational savings under mild assumptions. This setting also allows us to produce nested sequences of evaluation points $\{x^1, \ldots, x^{m_p}\}$ where m_p grows similar to n_p up to a logarithmic factor, therefore complying with the previously invoked prescription of hierarchical sampling. The analysis of this approach faces the difficulty that the x^i are not anymore identically distributed, and this is solved by using techniques first proposed in [17].

In Sect. 5 we turn to the study of the inverse Christoffel function k_n in the case of algebraic polynomial spaces of given total degree on general multivariate domains $D \subset \mathbb{R}^d$. We establish pointwise and global upper and lower bounds for k_n that depend on the smoothness of the boundary of D. We follow an approach adopted in [19] for a particular class of domains with piecewise smooth boundary, namely comparing D with simpler reference domains for which the inverse Christoffel function can be estimated. We obtain bounds with growth rate $\mathcal{O}(n^r)$ where the value $r = 2$ for Lipschitz domains and $r = \frac{d+1}{d}$ for smooth domains is proved to be sharp. We finally give a systematic approach that also describes the sharp growth rate for domains with cusp singularities.

We close the paper in Sect. 6 with various numerical experiments that confirm our theoretical investigations. In the particular case of multivariate algebraic polynomials, the sampling points tend to concentrate near to the exiting corner or cusp singularities of the domain, while they do not at the re-entrant singularities, as predicted by the previous analysis of the inverse Christoffel function.

2 Meeting the Optimality Benchmarks

2.1 Interpolation

One most commonly used strategy to reconstruct functions from point values is interpolation. Here we work in the space $V = \mathcal{C}(D)$ of continuous and bounded functions equipped with the L^∞ norm. For the given space V_n, and n distinct points $x^1, \ldots, x^n \in D$ picked in such way that the map $v \mapsto (v(x^1), \ldots, v(x^n))$ is an isomorphism from V_n to $\mathbb{R}^n$, we define the corresponding interpolation operator $\mathcal{I}_n : \mathcal{C}(D) \to V_n$ by the interpolation condition

$$\mathcal{I}_n u(x^i) = u(x^i), \quad i = 1, \ldots, n.$$

The interpolation operator is also expressed as

$$\mathcal{I}_n u = \sum_{i=1}^n u(x^i)\, \ell_i,$$

where $\{\ell_1, \ldots, \ell_n\}$ is the Lagrange basis of V_n defined by the conditions $\ell_i(x^j) = \delta_{i,j}$. Interpolation is obviously budget optimal since it uses $m = n$ points, that is, $C_1 = 1$ in (1.2). On the other hand, it does not guarantee instance optimality: the constant C_0 in (1.1) is governed by the *Lebesgue constant*

$$\Lambda_n = \|\mathcal{I}_n\|_{L^\infty \to L^\infty} = \max_{x \in D} \sum_{i=1}^n |\ell_i(x)|.$$

Indeed, since $\|u - \mathcal{I}_n u\|_{L^\infty} \le \|u - v\|_{L^\infty} + \|\mathcal{I}_n u - \mathcal{I}_n v\|_{L^\infty}$ for any $v \in V_n$, one has

$$\|u - \mathcal{I}_n u\|_{L^\infty} \le (1 + \Lambda_n) e_n(u).$$

The choice of the points x^i is critical to control the growth of Λ_n with n. For example in the elementary case of univariate algebraic polynomials where $D = [-1, 1]$ and $V_n = \mathbb{P}_{n-1}$, it is well known that uniformly spaced x^i results in Λ_n growing exponentially, at least like 2^n, while the slow (and optimal) growth $\Lambda_n \sim \ln(n)$ is ensured when using the Chebychev points $x^i = \cos\left(\frac{2i-1}{2n}\pi\right)$ for $i = 1, \ldots n$. Unfortunately, there is no general guideline to ensure such a slow growth for more general hierarchies of spaces $(V_n)_{n \ge 1}$ defined on multivariate domains $D \subset \mathbb{R}^d$. As an example, in the simple case of the bivariate algebraic polynomials $V_n = \mathbb{P}_p$ where $n = \frac{(p+1)(p+2)}{2}$ and a general polygonal domain D, the existence of a choice of points that would ensure a logarithmic growth of the Lebesgue constant is to our knowledge an open problem.

There exists a general point selection strategy that ensures linear behavior of the Lebesgue constant for any space V_n spanned by n functions $\{\varphi_1, \ldots, \varphi_n\}$. It consists in choosing $(x^1, \ldots, x^n)$ which maximizes over D^n the determinant of the collocation

matrix

$$M(x^1, \ldots, x^n) = (\phi_i(x^j))_{i,j=1,\ldots,n}.$$

Since the j-th element of the Lagrange basis is given by

$$\ell_j(x) = \frac{\det(M(x^1, \ldots, x^{j-1}, x, x^{j+1}, \ldots, x^n))}{\det(M(x^1, \ldots, x^n))},$$

the maximizing property gives that $\|\ell_j\|_{L^\infty} \le 1$ and therefore $\Lambda_n \le n$. In the particular case of the univariate polynomials where $D = [-1, 1]$ and $V_n = \mathbb{P}_{n-1}$, this choice corresponds to the Fekete points, which maximize the product $\prod_{i \ne j}(x^i - x^j)$.

While the above strategy guarantees the $\mathcal{O}(n)$ behavior of Λ_n, its main defect is that it is computationally unfeasible if n or d is large, since it requires solving a non-convex optimization problem in dimension $d n$. In addition to this, for a given hierarchy of spaces $(V_n)_{n \ge 1}$, the sampling points $S_n = \{x^1, \ldots, x^n\}$ generated by this strategy do not satisfy the nestedness property $S_n \subset S_{n+1}$.

A natural alternate strategy that ensures nestedness consists in selecting the points by a stepwise greedy optimization process: given S_{n-1}, define the next point x^n by maximizing over D the function $x \mapsto \det(M(x^1, \ldots, x^{n-1}, x))$. This approach was proposed in [16] in the context of reduced basis approximation and termed as *magic points*. It amounts to solving at each step a non-convex optimization problem in the more moderate dimension d, independent of n. However, there exists no general bound on Λ_n other than exponential in n. In the univariate polynomial case, this strategy yields the so-called *Leja points* for which it is only known that the Lebesgue constant grows sub-exponentially although numerical investigation indicates that it could behave linearly. In this very simple setting, the bound $\Lambda_n \le n^2$ could be established in [5], however, using a variant where the points are obtained by projections of the complex Leja points from the unit circle to the interval $[-1, 1]$.

In summary, while interpolation uses the optimal sampling budget $m = n$, it fails by a large amount in achieving instance optimality, especially when asking in addition for the nestedness of the sampling points, even for simple polynomial spaces.

2.2 Weighted Least-Squares

In order to improve the instance optimality bound, we allow ourselves to collect more data on the function u by increasing the number m of sample points, compared to the critical case $m = n$ studied before, and construct an approximation $\widetilde{u}_n$ by a least-squares fitting procedure. This relaxation of the problem gives more flexibility on the choice of the sample points: for instance, placing two of them too close will only waste one evaluation of u, whereas this situation would have caused ill-conditioning and high values of Λ_n in interpolation. It also leads to more favorable results in terms of instance optimality as we next recall.

Here, and in the rest of this paper, we assess the error in the L^2 norm

$$\|v\| = \|v\|_{L^2(D,\mu)},$$

where μ is a fixed probability measure, which can be arbitrarily chosen by the user depending on the targeted application. For example, if the error has the same significance at all points of D, one is naturally led to use the uniform probability measure

$$d\mu := |D|^{-1}dx.$$

In other applications such as uncertainty quantification, where the x variable represents random parameters that follow a more general probability law μ, the use of this specific measure is relevant since the reconstruction error may then be interpreted as the mean-square risk

$$\|u - \widetilde{u}_n\|^2_{L^2(D,\mu)} = \mathbb{E}_x(|u(x) - \widetilde{u}_n(x)|^2).$$

Once the evaluations of $u(x^i)$ are performed, the *weighted least-squares methods* defines $\widetilde{u}_n$ as the solution of the minimization problem

$$\min_{v \in V_n} \sum_{i=1}^m w(x^i)\,|u(x^i) - v(x^i)|^2, \tag{2.1}$$

where $w(x^1), \ldots, w(x^m) > 0$ are position-dependent weights. The solution to this problem is unique under the assumption that no function of $V_n \setminus \{0\}$ vanishes at all the x^i. Notice that in the limit $m = n$, the minimum in (2.1) is zero and attained by the interpolant at the points $x^1, \ldots, x^n$, which as previously discussed suffers from a severe lack of instance optimality.

The results from [9] provide with a general strategy to select the points x^i and the weight function w in order to reach instance and budget optimality, in a sense that we shall make precise. In this approach, the points x^i are drawn at random according to a probability measure σ on D, that generally differs from μ, but with respect to which μ is absolutely continuous. One then takes for w the inverse of the corresponding Radon-Nikodym derivative, so that

$$w(x)\,d\sigma(x) = d\mu(x).$$

This compatibility condition ensures that we recover a minimization in the continuous norm $\|\cdot\|$ as m tends to infinity:

$$\frac{1}{m}\sum_{i=1}^m w(x^i)\,|u(x^i) - v(x^i)|^2 \xrightarrow[m\to\infty]{a.s.} \int_D w\,|u - v|^2\,d\sigma = \int_D |u - v|^2\,d\mu = \|u - v\|^2.$$

Here we may work under the sole assumption that u belongs to the space $V = L^2(D,\mu)$, since pointwise evaluations of u and w will be almost surely well-defined. In return, since $\widetilde{u}_n$ is now stochastic, the L^2 estimation error will only be assessed in a probabilistic sense, for example by considering the mean-square error

$$\mathbb{E}(\|u - \widetilde{u}_n\|^2) = \mathbb{E}_{\otimes^m \sigma}(\|u - \widetilde{u}_n\|^2).$$

The weighted least-squares approximation may be viewed as the orthogonal projection $\widetilde{u}_n = P_n^m u$ onto V_n for the discrete ℓ^2 norm

$$\|v\|_m^2 := \frac{1}{m} \sum_{i=1}^{m} w(x^i) |v(x^i)|^2, \tag{2.2}$$

in the same way that the optimal approximation

$$u_n := \arg\min_{v \in V_n} \|u - v\| = P_n u$$

is the orthogonal projection for the continuous $L^2(D, \mu)$ norm. A helpful object for comparing these two norms on V_n is the Gramian matrix

$$G := (\langle L_j, L_k \rangle_m)_{j,k=1,\ldots,n}, \tag{2.3}$$

where $(L_1, \ldots, L_n)$ is any $L^2(D, \mu)$-orthonormal basis of V_n and $\langle \cdot, \cdot \rangle_m$ is the inner product associated with the $\|\cdot\|_m$ norm. Indeed, for all $\delta > 0$,

$$\|G - I\|_2 \leq \delta \iff (1 - \delta)\|v\|^2 \leq \|v\|_m^2 \leq (1 + \delta)\|v\|^2, \quad v \in V_n,$$

where $\|M\|_2$ denotes the spectral norm of an $n \times n$ matrix M. As noted in [8] in the case of standard least-squares, and in [9] for the weighted case, G can be seen as a mean of m independent and identically distributed matrices

$$X^i := (w(x^i) L_j(x^i) L_k(x^i))_{j,k=1,\ldots,n}$$

satisfying $\mathbb{E}(X^i) = I$, so G concentrates toward the identity as m grows to infinity. This concentration can be estimated by a matrix Chernoff bound, such as Theorem 1.1 in the survey paper [20]. As observed in [9], for the particular value $\delta = \frac{1}{2}$, this inequality can be rewritten as follows, in our case of interest.

Lemma 2.1 *For any $\varepsilon > 0$, under the sampling budget condition*

$$m \geq \gamma \|w k_n\|_{L^\infty} \ln(2n/\varepsilon),$$

where $\gamma := (3/2 \ln(3/2) - 1/2)^{-1} \approx 9.242$, one has $\Pr(\|G - I\|_2 \leq 1/2) \geq 1 - \varepsilon$.

An estimate comparing the estimator $\|u - \widetilde{u}_n\|$ with $e_n(u)$ can be obtained when imposing that $\|G - I\|_2 \leq 1/2$, as expressed in the following, which is proved in [9].

Lemma 2.2 *One has*

$$\mathbb{E}\left(\|u - \widetilde{u}_n\|^2 \chi_{\|G - I\|_2 \leq 1/2}\right) \leq \left(1 + \frac{4}{m} \|w k_n\|_{L^\infty}\right) e_n(u)^2.$$

On the other hand, the estimator $\widetilde{u}_n$ obtained by solving (2.1) is not reliable in the event where G becomes singular, which leads to modify its definition in various ways:

1. If one is able to compute $\|G - I\|_2$, one may condition the estimator to the event $\|G - I\|_2 \leq \frac{1}{2}$ by defining

$$\widetilde{u}_n^C := \widetilde{u}_n \, \chi_{\|G - I\|_2 \leq 1/2},$$

   that is, we take $\widetilde{u}_n^C = 0$ if $\|G - I\|_2 > \frac{1}{2}$.

2. If a uniform bound $\|u\|_{L^\infty(D)} \leq \tau$ is known, one may introduce a truncated estimator

$$\widetilde{u}_n^T := T_\tau \circ \widetilde{u}_n, \tag{2.4}$$

   where $T_\tau(y) := \min\{\tau, |y|\} \operatorname{sgn}(y)$.

The main results from [9], that we slightly reformulate below, show that these estimators are instance optimal in a probabilistic sense. Throughout the rest of the paper, γ denotes the same constant as in Lemma 2.1.

Theorem 2.3 *Under the sampling budget condition*

$$m \geq \gamma \, \|w \, k_n\|_{L^\infty} \, \ln(2n/\varepsilon), \tag{2.5}$$

the weighted least-squares estimator satisfies

$$\mathbb{E} \left(\|u - \widetilde{u}_n\|^2 \chi_{\|G - I\|_2 \leq 1/2} \right) \leq (1 + \eta(m)) \, e_n(u)^2. \tag{2.6}$$

The conditioned and truncated estimators satisfy the convergence bounds

$$\mathbb{E} \left(\|u - \widetilde{u}_n^C\|^2 \right) \leq (1 + \eta(m)) \, e_n(u)^2 + \|u\|^2 \, \varepsilon, \tag{2.7}$$

and

$$\mathbb{E} \left(\|u - \widetilde{u}_n^T\|^2 \right) \leq (1 + \eta(m)) \, e_n(u)^2 + 4 \, \tau^2 \, \varepsilon, \tag{2.8}$$

where $\eta(m) = \frac{4}{m} \|w \, k_n\|_{L^\infty} \leq \frac{4}{\gamma \ln(2n/\varepsilon)} \to 0$, *as* $n \to \infty$ *or* $\varepsilon \to 0$.

Proof The bound (2.6) follows directly from Lemma 2.2 and the assumption on m. In the event $\|G - I\|_2 > \frac{1}{2}$, of probability less than ε by Lemma 2.1, one can use the bounds

$$\|u - \widetilde{u}_n^C\|^2 = \|u\|^2 \quad \text{and} \quad \|u - \widetilde{u}_n^T\|^2 \leq 4\tau^2.$$

Otherwise, one has

$$\|u - \widetilde{u}_n^C\|^2 \leq \|u - \widetilde{u}_n\|^2 \quad \text{and} \quad \|u - \widetilde{u}_n^T\|^2 \leq \|u - \widetilde{u}_n\|^2.$$

This leads to (2.7) and (2.8). $\qquad\qquad\square$

Remark 2.4 The above result shows that the estimators $\widetilde{u}_n^C$ and $\widetilde{u}_n^T$ achieve instance optimality in expectation up to additional error terms of order $\mathcal{O}(\varepsilon)$, accounting for the event $\{\|G - I\|_2 > 1/2\}$. Note that ε only influences the constraint on the sampling budget logarithmically. In particular, if $e_n(u)$ decreases like n^{-r} for some $r > 0$, these estimators are rate optimal by taking ε less than n^{-2r}, which thus affects the constraint on sampling budget by a factor $\mathcal{O}(\ln(n))$. Note, however, that for exponential rates of the form $\exp(-cn^\gamma)$—that occur for example when approximating analytic functions by polynomials—imposing ε to be of this order results in a sampling budget m of sub-optimal order $n^{1+\gamma}$ up to a logarithmic factor.

Remark 2.5 One way to achieve instance optimality in expectation without an additional error term consists in redrawing the points $\{x^1, \ldots, x^m\}$ until one observes that $\|G - I\|_2 \leq \frac{1}{2}$, as proposed in [13]. We denote by u_n^* the weighted least-squares estimator corresponding to this conditioned draw. In other words u_n^* is the weighted least-squares estimator $\widetilde{u}_n$ conditioned to the event $\{\|G - I\|_2 \leq \frac{1}{2}\}$. Since, by Bayes rule,

$$
\Pr\left\{\|G - I\|_2 \leq \frac{1}{2}\right\} \mathbb{E}\left(\|u - \widetilde{u}_n\|^2 \;\middle|\; \|G - I\|_2 \leq \frac{1}{2}\right) = \mathbb{E}\left(\|u - \widetilde{u}_n\|^2 \chi_{\|G-I\|_2 \leq \frac{1}{2}}\right),
$$

we find that under the sampling budget (2.5), one has

$$
\mathbb{E}(\|u - u_n^*\|^2) = \mathbb{E}\left(\|u - \widetilde{u}_n\|^2 \;\middle|\; \|G - I\|_2 \leq \frac{1}{2}\right) \leq \frac{1}{1 - \varepsilon} \mathbb{E}\left(\|u - \widetilde{u}_n\|^2 \chi_{\|G-I\|_2 \leq \frac{1}{2}}\right),
$$

and thus

$$
\mathbb{E}(\|u - u_n^*\|^2) \leq \frac{1}{1 - \varepsilon}(1 + \eta(m))e_n(u)^2.
$$

The sampling budget condition also ensures a probabilistic control on the number of required redraws, since the probability that the event $\{\|G - I\|_2 \leq \frac{1}{2}\}$ did not occur after k redraws is less than ε^k.

Now the natural objective is to find a weight function w that makes $\|w\,k_n\|_{L^\infty}$ small in order to minimize the sampling budget. Since

$$
\|w\,k_n\|_{L^\infty} \geq \int_D w\,k_n\,d\sigma = \int_D k_n\,d\mu = n,
$$

with equality attained for the weight function

$$
w^* := \frac{n}{k_n} = \frac{n}{\sum_{j=1}^n |L_j|^2},
$$

this theorem shows that the choice of sampling measure

$$
d\sigma^* = \frac{1}{w^*}\,d\mu = \frac{k_n}{n}\,d\mu
$$

is optimal, in the sense that the above instance optimality results are achieved with a near-optimal sampling budget $m \sim n$ up to logarithmic factors.

As already explained in the introduction, when working on a general domain D, we face the difficulty that the orthonormal basis $(L_1, \ldots, L_n)$ cannot be exactly computed, and therefore, the optimal w^* and σ^* are out of reach. The next section discusses computable alternatives $\widetilde{w}$ and $\widetilde{\sigma}$ that still yield similar instance optimality results at near-optimal sampling budget.

3 Near-Optimal Sampling Strategies on General Domains

3.1 Two Steps Sampling Strategies

The sampling and reconstruction strategies that we discuss proceed in two steps:

1. In an offline stage we search for an approximation to the Christoffel function k_n. For this purpose, we sample $z^1, \ldots, z^M \in D$ according to μ, use these sampling points to compute an orthonormal basis $(\widetilde{L}_1, \ldots, \widetilde{L}_n)$ with respect to the induced discrete inner product. The approximation to the Christoffel function is then $\widetilde{k}_n = \sum_{j=1}^{n} |\widetilde{L}_j|^2$. As we explain further, one objective is to guarantee that $\widetilde{k}_n$ and k_n are pointwise equivalent. We define the sampling measure $\widetilde{\sigma}$ as proportional to $\widetilde{k}_n \mu$ and draw the points $x^1, \ldots, x^m$ according to this measure.
2. In an online stage, we evaluate u at the sampling points x^i and construct an estimate $\widetilde{u}_n$ by the weighted least-squares method.

In the offline stage M could be much larger than n; however, it should be understood that the function u is only evaluated in the online stage at the m point x^i which will be seen to have optimal cardinality $m \sim n$ up to logarithmic factors.

The two main requirements in these approaches are the access to a (non-orthogonal) basis $(\phi_1, \ldots, \phi_n)$ of V_n and the ability to sample according to measure μ. When $D \subset \mathbb{R}^d$ is a general multivariate domain, one typical setting for this second assumption to be valid is the following:

- There is a set R containing D such that μ is the restriction of a measure μ_R which can easily be sampled.
- Membership of a point x to the set D can be efficiently tested, that is, χ_D is easily computed.

This includes for instance the uniform probability measure on domains described by general systems of algebraic inequalities (such as polyhedrons and ellipsoids), by including such domains D in a rectangle $R = I_1 \times \cdots \times I_d$ on which sampling according to the uniform measure can be done componentwise. Then the z^i are produced by sampling according to μ_R and rejecting the samples that do not belong to D. The offline stage is described more precisely as follows.

Algorithm 1 Draw a certain number M of points $z^1, \ldots, z^M$ independently according to μ, and construct from $(\phi_j)_{j=1,\ldots,n}$ an orthonormal basis $(\widetilde{L}_j)_{j=1,\ldots,n}$ of V_n with

respect to the inner product

$$\langle u, v \rangle_M := \frac{1}{M} \sum_{i=1}^{M} u(z^i)\, v(z^i). \tag{3.1}$$

Then define

$$\widetilde{k}_n(x) = \sum_{j=1}^{n} |\widetilde{L}_j(x)|^2,$$

the approximate inverse Christoffel function, and the corresponding sampling measure

$$\mathrm{d}\widetilde{\sigma} := \alpha\, \frac{\widetilde{k}_n}{n}\, \mathrm{d}\mu,$$

where α is the normalization factor such that $\alpha \int_D \widetilde{k}_n\, \mathrm{d}\mu = n$.

Note that the factor α is unknown to us but its value is not needed in typical sampling strategies, such as rejection sampling or MCMC. In contrast to k_n, the function $\widetilde{k}_n$ is stochastic since it depends on the drawing of the z^i. In the online stage, we sample $x^1, \dots, x^m$ independently according to $\mathrm{d}\widetilde{\sigma}$. We then measure u at the points x^i, and define the estimator $\widetilde{u}_n \in V_n$ as the solution to the weighted least-squares problem

$$\min_{v \in V_n} \sum_{i=1}^{m} \widetilde{w}(x^i)\, |u(x^i) - v(x^i)|^2, \tag{3.2}$$

with $\widetilde{w} = \frac{n}{\alpha k_n}$. This least-squares problem can be solved explicitly by computing $\widetilde{u}_n = P_n^m u$ as the orthogonal projection of u on V_n with respect to the inner product from (2.2)

$$\langle u, v \rangle_m := \frac{1}{m} \sum_{i=1}^{m} \widetilde{w}(x^i)\, u(x^i)\, v(x^i).$$

Remark 3.1 There are now two levels of stochasticity: the draw of the z^i and the subsequent draw of the x^i. We sometimes use the symbols $\mathbb{E}_z$ and Pr_z referring to the first draw, and $\mathbb{E}_x$ and Pr_x referring to the second draw given the first one, while $\mathbb{E}$ and Pr refer to both draws.

We keep the notations G and $\widetilde{u}_n^T$ from (2.3) and (2.4). In the following section, we establish instance optimal convergence results under near optimal sample complexity m similar to (2.6) and (2.8) in Theorem 2.3. On the other hand we do not consider the conditioned estimator $\widetilde{u}_n^C$ any further since we do not have access to the matrix G which would require the knowledge of the functions L_j. The derivation of a computable estimator that satisfies a similar estimate as $\widetilde{u}_n^C$ is an open question. We also discuss the required sample complexity M of the offline stage.

3.2 Convergence Bounds and Sample Complexity

Our principle objective is to ensure the uniform framing

$$c_1 \, k_n(x) \leq \widetilde{k}_n(x) \leq c_2 \, k_n(x), \quad x \in D, \tag{3.3}$$

for some known constants $0 < c_1 \leq c_2$. Our motivation is that instance optimal convergence bounds with near-optimal sampling budget hold under this framing, as expressed by the following result.

Theorem 3.2 *Assume that (3.3) holds for some $0 < c_1 \leq c_2$ and let $c = \frac{c_2}{c_1} \geq 1$. Then, under the sampling budget condition*

$$m \geq c \, \gamma \, n \, \ln(2n/\varepsilon), \tag{3.4}$$

one has $\mathrm{Pr}_x \left(\|G - I\|_2 \geq \frac{1}{2} \right) \leq \varepsilon$. *In addition, one has the convergence bounds*

$$\mathbb{E}_x \left(\|u - \widetilde{u}_n\|^2 \chi_{\|G-I\|_2 \leq \frac{1}{2}} \right) \leq (1 + \eta(m)) \, e_n(u)^2, \tag{3.5}$$

and

$$\mathbb{E}_x(\|u - \widetilde{u}_n^T\|^2) \leq (1 + \eta(m)) \, e_n(u)^2 + 4 \, \varepsilon \, \tau^2,$$

where $\eta(m) = 4 \, c \, \frac{n}{m} \leq \frac{4}{\gamma \ln(2n/\varepsilon)}$.

Proof It is an immediate application of the results from §2.2. Indeed

$$\|\widetilde{w} \, k_n\|_{L^\infty} = \left\| \frac{n \, k_n}{\alpha \, \widetilde{k}_n} \right\|_{L^\infty} = \left\| \frac{k_n}{\widetilde{k}_n} \right\|_{L^\infty} \int_D \widetilde{k}_n \, \mathrm{d}\mu \leq \left\| \frac{k_n}{\widetilde{k}_n} \right\|_{L^\infty} \left\| \frac{\widetilde{k}_n}{k_n} \right\|_{L^\infty} \int_D k_n \, \mathrm{d}\mu \leq c \, n.$$

Therefore, the sampling condition (3.4) implies $m \geq \gamma \, \|\widetilde{w} \, k_n\|_{L^\infty} \ln(2n/\varepsilon)$, and the results follow by direct application of Lemma 2.1 and Theorem 2.3. □

We now concentrate our attention on the offline procedure which should be tuned in order to ensure that (3.3) holds with high probability. For this purpose, we introduce the Gramian matrix

$$G_M := (\langle L_j, L_k \rangle_M)_{j,k=1,\dots n},$$

where $\langle \cdot, \cdot \rangle_M$ is the inner product defined by (3.1) that uses the intermediate samples z^i which are i.i.d. according to μ. This matrix should not be confused with G defined in (2.3) that uses the inner product $\langle \cdot, \cdot \rangle_m$ based on the final samples x^i.

Lemma 3.3 *For any pair of constants $0 < c_1 \leq c_2 < \infty$, the matrix framing property*

$$c_2^{-1} \, I \leq G_M \leq c_1^{-1} \, I \tag{3.6}$$

implies the uniform framing (3.3).

Proof We use the fact that, similar to k_n, the function $\widetilde{k}_n$ is characterized by the extremality property

$$\widetilde{k}_n(x) = \max_{v \in V_n} \frac{|v(x)|^2}{\|v\|_M^2}.$$

For any $x \in D$ and $v \in V_n$, one has on the one hand

$$|v(x)|^2 \le \widetilde{k}_n(x)\,\|v\|_M^2 \le c_1^{-1}\,\widetilde{k}_n(x)\,\|v\|^2,$$

where the last inequality results from the upper one in (3.6). This shows that $c_1\,k_n(x) \le \widetilde{k}_n(x)$. On the other hand, using the lower inequality in (3.6), we find that

$$|v(x)|^2 \le k_n(x)\,\|v\|^2 \le c_2\,k_n(x)\,\|v\|_M^2,$$

which shows that $\widetilde{k}_n(x) \le c_2\,k_n(x)$. □

Remark 3.4 The matrix framing (3.6) implies the uniform framing (3.3), but the converse does not seem to hold. Finding an algebraic condition equivalent to (3.3) is an open question.

Lemma 2.1 indicates that if the amount of offline samples satisfies the condition

$$M \ge \gamma\,K_n \ln(2n/\varepsilon), \quad K_n := \|k_n\|_{L^\infty(D)}, \tag{3.7}$$

then we are ensured that

$$\mathrm{Pr}_z\left(\|G_M - I\|_2 \ge 1/2\right) \le \varepsilon,$$

and therefore the framing (3.6) holds with probability greater than $1 - \varepsilon$, for the particular values $c_1 = \frac{2}{3}$ and $c_2 = 2$. Bearing in mind that k_n is unknown to us, we assume at least that we know an upper estimate for its L^∞ norm

$$K_n \le B(n).$$

Explicit values for $B(n)$ for general domains D are established in Sect. 5 in the case where the V_n are spaces of algebraic polynomials. Therefore, given such a bound, taking M such that

$$M \ge \gamma\,B(n) \ln(2n/\varepsilon) \tag{3.8}$$

guarantees a similar framing with probability greater than $1 - \varepsilon$. We obtain the following result as a direct consequence of Theorem 3.2.

Corollary 3.5 *Assume that the amount of samples M used in the offline stage described by Algorithm 1 satisfies (3.8) for some given $\varepsilon > 0$. Then, under the sampling budget condition*

$$m \geq 3 \gamma n \ln(2n/\varepsilon),$$

for the online stage, the event $E := \{\|G - I\|_2 \leq \frac{1}{2}$ and $\|G_M - I\|_2 \leq \frac{1}{2}\}$ satisfies $\Pr(E^c) \leq 2\varepsilon$. In addition, one has the convergence bounds

$$\mathbb{E}(\|u - \widetilde{u}_n\|^2 \chi_E) \leq (1 + \eta(m)) \, e_n(u)^2 \tag{3.9}$$

and

$$\mathbb{E}(\|u - \widetilde{u}_n^T\|^2) \leq (1 + \eta(m)) \, e_n(u)^2 + 8 \, \varepsilon \, \tau^2, \tag{3.10}$$

where $\eta(m) = 12 \frac{n}{m} \leq \frac{4}{\gamma \ln(2n/\varepsilon)}$.

Proof The estimate on $\Pr(E^c)$ follows by a union bound. Since $\|G_M - I\|_2 \geq \frac{1}{2}$ ensures the framing (3.3) with $c_1 = \frac{2}{3}$ and $c_2 = 2$, the bound (3.9) follows from (3.5) in Theorem 3.2. Finally, the bound (3.10) follows from (3.9) and the probability estimate on E^c by the same argument as in the proof of Theorem 2.3. □

3.3 An Empirical Determination of the Value of *M*

In many situations, the best available bound $B(n)$ on K_n could be overestimated by a large amount. Moreover, the theoretical requirement $M \geq \gamma K_n \ln(2n/\varepsilon)$ is only a sufficient condition that guarantees that $\|G_M - I\|_2 \leq \frac{1}{2}$ with probability larger than $1 - \varepsilon$. It could happen that for smaller values of M, the matrix G_M satisfies the framing (3.6) with constants c_1 and c_2 that have moderate ratio $c = \frac{c_2}{c_1}$.

Since the computational cost of the offline stage is proportional to M, it would be desirable to use such a smaller value of M. If we could compute the matrix G_M it would suffice to raise M until the condition number

$$\kappa(G_M) = \frac{\lambda_{\max}(G_M)}{\lambda_{\min}(G_M)},$$

has value smaller than a prescribed threshold $c^* > 1$, so that (3.6) holds with $c = \kappa(G_M) \leq c^*$.

However, since the exact orthonormal basis elements L_j are generally unknown to us, we cannot compute the matrix G_M. As an alternate strategy, we propose the following method that provides an empirical determination of the value M that should be used in Algorithm 1: start from the minimal value $M = n$, and draw points $y^1, \ldots, y^M$ and $z^1, \ldots z^M$ independently according to μ. Then, defining

$$\langle u, v \rangle_y = \frac{1}{M} \sum_{i=1}^{M} u(y^i) \, v(y^i) \quad \text{and} \quad \langle u, v \rangle_z = \frac{1}{M} \sum_{i=1}^{M} u(z^i) \, v(z^i),$$

compute an orthonormal basis (L_j^y) with respect to $\langle \cdot, \cdot \rangle_y$, and define the test matrix

$$T := (\langle L_j^y, L_k^y \rangle_z)_{j,k=1,\ldots,n}.$$

If $\kappa(T) \geq c^*$, then raise the value of M by some fixed amount, and repeat this step until $\kappa(T) \leq c^*$. For this empirically found value $M = M_{\text{emp}}(n)$, use the points $\{z^1, \ldots, z^M\}$ in the offline stage described by Algorithm 1, and the constant $c = c^*$ in the sampling budget condition (3.4) used in the online stage.

The rationale for this approach is that if G_M is well conditioned with high probability, then T should also be, as shown for example by the following result.

Proposition 3.6 *If M is chosen in such a way that $\Pr(\kappa(G_M) \geq c) \leq \varepsilon$ for some $c > 1$, then*

$$\Pr(\kappa(T) \geq c^2) \leq 2\varepsilon.$$

Proof Since both matrices $G_y = (\langle L_j, L_k \rangle_y)_{j,k=1,\ldots,m}$ and $G_z = (\langle L_j, L_k \rangle_z)_{j,k=1,\ldots,m}$ are realizations of G_M, we obtain by a union bound that, with probability at least $1-2\varepsilon$, both G_y and G_z have condition numbers less than c. Under this event,

$$
\lambda_{\max}(T) = \sup_{\alpha \in \mathbb{R}^n} \frac{\|\sum_{j=1}^n \alpha_j L_j^y\|_z^2}{|\alpha|^2} \leq \sup_{\alpha \in \mathbb{R}^n} \frac{\|\sum_{j=1}^n \alpha_j L_j^y\|^2}{|\alpha|^2} \sup_{v \in V_n} \frac{\|v\|_z^2}{\|v\|^2}
$$

$$
= \sup_{v \in V_n} \frac{\|v\|^2}{\|v\|_y^2} \sup_{v \in V_n} \frac{\|v\|_z^2}{\|v\|^2} = \left(\inf_{v \in V_n} \frac{\|v\|_y^2}{\|v\|^2} \right)^{-1} \sup_{v \in V_n} \frac{\|v\|_z^2}{\|v\|^2} = \frac{\lambda_{\max}(G_z)}{\lambda_{\min}(G_y)},
$$

and

$$
\lambda_{\min}(T) = \inf_{\alpha \in \mathbb{R}^n} \frac{\|\sum_{j=1}^n \alpha_j L_j^y\|_z^2}{|\alpha|^2} \geq \inf_{\alpha \in \mathbb{R}^n} \frac{\|\sum_{j=1}^n \alpha_j L_j^y\|^2}{|\alpha|^2} \inf_{v \in V_n} \frac{\|v\|_z^2}{\|v\|^2}
$$

$$
= \inf_{v \in V_n} \frac{\|v\|^2}{\|v\|_y^2} \inf_{v \in V_n} \frac{\|v\|_z^2}{\|v\|^2} = \left(\sup_{v \in V_n} \frac{\|v\|_y^2}{\|v\|^2} \right)^{-1} \inf_{v \in V_n} \frac{\|v\|_z^2}{\|v\|^2} = \frac{\lambda_{\min}(G_z)}{\lambda_{\max}(G_y)},
$$

which implies that $\kappa(T) \leq \kappa(G_y) \kappa(G_z) \leq c^2$. $\qquad \square$

The above proposition shows that a good conditioning of G_M with high probability implies the same property for T. There is of course no theoretical guarantee that the value of M provided by the above empirical approach is sufficient to achieve good conditioning of G_M, unless the resulting M satisfies (3.8). However, in the numerical experiments of Sect. 6, we will check that the values of M for which $\kappa(T) \leq c$ do also ensure that a similar bound holds for $\kappa(G_M)$.

4 Multilevel Strategies

The sampling strategy described by Algorithm 1 provides instance optimal reconstructions of u with an optimal sampling budget up to a multiplicative factor $\ln(2n/\varepsilon)$. Thus, the execution time of the online stage, dominated by the m evaluations of u at points x^i, cannot be significantly improved. On the other hand, the complexity of the offline stage is dominated by the computation of the Gramian matrix for deriving the basis $\{\widetilde{L}_1, \ldots, \widetilde{L}_n\}$ and is therefore of order $\mathcal{O}(Mn^2)$. In particular, it depends linearly on the number of points M, which could be very large if K_n grows fast, or if its available bound $B(n)$ is over-estimated.

In this section we discuss a multilevel approach aiming at improving this offline computational cost: we produce an approximation to k_n in several iterations, by successive refinements of this function as the dimension of V_n increases. We consider a family of nested spaces $(V_{n_p})_{p \geq 1}$ of increasing dimension n_p and take an orthonormal basis $(L_j)_{j \geq 1}$ adapted to this hierarchy, in the sense that

$$V_{n_p} = \mathrm{Span}\{L_1, \ldots, L_{n_p}\}, \quad n_p \geq 1.$$

As previously, the exact functions k_{n_p} are out of reach, since we do not have access to the continuous inner product by which we would compute the basis $(L_j)_{1 \leq j \leq n_p}$. The offline stage described in Sect. 3 computes approximations $\widetilde{L}_j$ by orthogonalizing with respect to a discrete inner product with points z^i drawn according to $d\mu$. We know that a more efficient sample for performing this orthogonalization should be drawn according to $d\sigma = \frac{k_{n_p}}{n_p} d\mu$ which is, however, unknown to us. The idea for breaking this dependency loop is to replace k_{n_p} with $\widetilde{k}_{n_{p-1}}$, which was computed at the previous step. Our analysis of this strategy is based on the following assumption of proximity between $k_{n_{p-1}}$ and k_{n_p}:

There exists a known constant $\kappa > 1$ such that

$$k_{n_1}(x) \leq 3\kappa\, n_1 \text{ and } k_{n_p}(x) \leq k_{n_{p+1}}(x) \leq \kappa\, k_{n_p}(x), \quad p \geq 1, \quad x \in D. \quad (4.1)$$

The validity of this assumption can be studied through lower and upper estimates for k_n, such as those discussed in the next section. For example, Theorem 5.9 allows one to establish (4.1) for bivariate polynomial spaces of total degree p, therefore with $n_p = (p+1)(p+2)/2$, on domains with piecewise smooth boundary. Note that (4.1) allows up to exponential growth of K_n, if we simply take $n_p = p$.

Assuming that the targeted space V_n is a member of this hierarchy, that is,

$$n = n_q, \quad \text{for some} \quad q > 1,$$

we modify the offline stage as follows.

Algorithm 2 Start with $\widetilde{w}_0 = 1$ and $\widetilde{\sigma}_0 = \mu$. For $p = 1, \ldots, q$, iterate the following: draw a certain number M_p of points $z_p^1, \ldots, z_p^{M_p}$ independently according to $\widetilde{\sigma}_{p-1}$, and

construct an orthonormal basis $(L_j^p)_{1 \le j \le n_p}$ of V_{n_p} with respect to the inner product

$$\langle u, v \rangle_p := \frac{1}{M_p} \sum_{i=1}^{M_p} \widetilde{w}_{p-1}(z_p^i)\, u(z_p^i)\, v(z_p^i). \tag{4.2}$$

Then define

$$\widetilde{k}_{n_p} = \sum_{j=1}^{n_p} |L_j^p|^2, \quad \widetilde{w}_p = \frac{n_p}{\widetilde{k}_{n_p}}, \quad \text{and} \quad d\widetilde{\sigma}_p := \alpha_p \frac{\widetilde{k}_{n_p}}{n_p}\, d\mu,$$

where α_p is the normalization constant, and proceed to the next iteration. At the end of iteration q, define the perturbed Christoffel function for V_n as $\widetilde{k}_n = \widetilde{k}_{n_q}$, weight function $\widetilde{w} = \widetilde{w}_q$ and sampling measure

$$d\widetilde{\sigma} = d\widetilde{\sigma}_q = \alpha \frac{\widetilde{k}_n}{n}\, d\mu,$$

where α is a normalization factor.

The online stage remains unchanged: the samples $x^1, \ldots, x^m$ for evaluation of u are drawn i.i.d. according to $\widetilde{\sigma}$, and we solve the weighed least-squares problem (3.2). The sample size M of the offline stage is now replaced by $\overline{M} = M_1 + \cdots + M_q$. We denote by $G_p := (\langle L_j, L_k \rangle_p)_{j,k=1,\ldots,n_p}$ the Gramian matrices for the inner products (4.2), The following result shows that the conditions imposed on the M_p are less stringent than those that were imposed on M.

Theorem 4.1 *Let $\varepsilon_p > 0$ such that $\varepsilon := \sum_{p=1}^{q} \varepsilon_p < 1$, and assume that the amount of offline samples used in Algorithm 2 satisfies*

$$M_p \ge 3\kappa\,\gamma\,n_p \ln \frac{2n_p}{\varepsilon_p}, \quad p = 1, \ldots, q,$$

with κ the constant in the assumption (4.1). Then if $m \ge 3\gamma\,n \ln \frac{2n}{\varepsilon}$, the same convergence bounds (3.9) and (3.10) as in Corollary 3.5 hold, with $E := \{\|G - I\|_2 \le \frac{1}{2} \text{ and } \|G_q - I\|_2 \le \frac{1}{2}\}$ that satisfies $\Pr(E^c) \le 2\varepsilon$.

Proof We show by induction on p that the event

$$B_p := \left\{ \|G_1 - I\|_2 \le \frac{1}{2}, \ldots, \|G_p - I\|_2 \le \frac{1}{2} \right\}$$

occurs with probability at least $1 - \varepsilon_1 - \cdots - \varepsilon_p$. As

$$M_1 \ge 3\kappa\,\gamma\,n_1 \ln \frac{2n_1}{\varepsilon_1} \ge \gamma\,\|\widetilde{w}_0\,k_{n_1}\|_{L^\infty} \ln \frac{2n_1}{\varepsilon_1},$$

by Lemma 2.1,

$$\Pr(B_1) \geq 1 - \varepsilon_1.$$

For $1 \leq p < q$, under the event B_p, Lemma 3.3 gives

$$\frac{2}{3} k_{n_p}(x) \leq \widetilde{k}_{n_p}(x) \leq 2 k_{n_p}(x), \quad x \in D.$$

Therefore, using assumption (4.1), we find that

$$\left\| \frac{\widetilde{w}_p \, k_{n_{p+1}}}{\alpha_p} \right\|_{L^\infty} = \left\| \frac{n_p}{\alpha_p} \frac{k_{n_{p+1}}}{\widetilde{k}_{n_p}} \right\|_{L^\infty} \leq n_p \left\| \frac{\widetilde{k}_{n_p}}{k_{n_p}} \right\|_{L^\infty} \left\| \frac{k_{n_p}}{\widetilde{k}_{n_p}} \right\|_{L^\infty} \left\| \frac{k_{n_{p+1}}}{k_{n_p}} \right\|_{L^\infty} \leq 3 \kappa \, n_p.$$

As $M_{p+1} \geq 3 \kappa \gamma \, n_p \ln \frac{2 n_p}{\varepsilon_{p+1}}$, Lemma 2.1 applies, and combining this with the induction hypothesis:

$$\Pr(B_{p+1}) = \Pr(B_p) \Pr \left(\|G_{p+1} - I\|_2 \leq \frac{1}{2} \,\middle|\, B_p \right)$$
$$\geq (1 - \varepsilon_1 - \cdots - \varepsilon_p)(1 - \varepsilon_{p+1}) \geq 1 - \varepsilon_1 - \cdots - \varepsilon_p - \varepsilon_{p+1}.$$

Use Lemma 3.3 one last time to write, in the event B_q,

$$\frac{2}{3} k_{n_q}(x) \leq \widetilde{k}_{n_q}(x) \leq 2 k_{n_q}(x), \quad x \in D,$$

which is the framing (3.6) for the particular values $c_1 = \frac{2}{3}$ and $c_2 = 2$. Since B_q has probability larger than $1 - \varepsilon$, we conclude by the exact same arguments used in the proof of Corollary 3.5.   □

We now comment on the gain of complexity by using Algorithm 2:

1. Exponential growth of K_n: the property (4.1) might be satisfied even when K_n grows exponentially with n, by taking the choice $n_p = p$. Then, the complexity of Algorithm 1 is of order $\mathcal{O}(M n^2) \gtrsim \mathcal{O}(K_n n^2 \ln(2n/\varepsilon))$, which grows exponentially in n. In contrast, the total amount of sampling in Algorithm 2 is $\overline{M} = M_1 + \cdots + M_n \leq n M_n = \mathcal{O}(n^2 \ln(2n/\varepsilon))$, so the first stage remains of polynomial complexity $\mathcal{O}(n^4 \ln(2n/\varepsilon))$.
2. Algebraic growth of K_n: if $K_n \sim n^r$ only grows algebraically in n, one may choose $n_p = 2^p$, in which case the total number of sample points $\overline{M}$ rewrites as $M_{n_0} + \cdots + M_{n_q} \sim M_{n_q}$, giving an optimal complexity $\mathcal{O}(n^3 \ln(2n/\varepsilon))$ for the first stage. This is smaller than the complexity $\mathcal{O}(K_n n^2 \ln(2n/\varepsilon)) = \mathcal{O}(n^{2+r} \ln(2n/\varepsilon))$ encountered in Algorithm 1.

While Algorithm 2 produces a computational gain in computing a near-optimal measure $\widetilde{\sigma}$, the resulting sample $x^1, \ldots, x^m$ is specifically targeted at approximating

u in the space $V_n = V_{n_q}$. As explained in Sect. 1.2, it is sometimes desirable to obtain optimal weighted least-squares approximations $\tilde{u}_{n_p}$ for each space V_{n_p} while maintaining the cumulated number of evaluations of u until step p of the optimal order n_p up to logarithmic factors. Therefore, we would like to recycle the evaluation points $\{x^1, \ldots, x^{m_{p-1}}\}$ used until step $p - 1$ in order to create the new evaluation sample $\{x^1, \ldots, x^{m_p}\}$, for some well chosen sequence $(m_p)_{p \geq 1}$ that grows similar to $(n_p)_{p \geq 1}$ up to logarithmic factors.

Here, we assume that the family $(V_{n_p})_{p \geq 1}$ has been fixed, independently of the target function u, in contrast to being generated in an adaptive manner. Adaptive space generation brings out new difficulties: the space V_{n_p} depends on the approximation computed in the previous space $V_{n_{p-1}}$ and therefore the new evaluation samples $x^{m_{p-1}+1}, \ldots, x^{m_p}$ will not be independent from the previous ones. Maintaining optimal sample complexity in the adaptive context is, to our knowledge, an open problem.

Intuitively, since the sample should have a density proportional to k_{n_p}, most of the new points we draw at step p should be distributed according to a density proportional to $k_{n_p} - k_{n_{p-1}} = \sum_{j=n_{p-1}+1}^{n_p} |L_j|^2$. This leads us to the following algorithm.

Algorithm 3 Start with $\tilde{w}_0 = 1$ and $\tilde{\sigma}_0 = \mu$ and $m_0 = 0$. For $p = 1, 2, \ldots$, generate $z_{n_p}^i$ and compute $\tilde{w}_{n_p}$, $\tilde{\sigma}_{n_p}$ and $\tilde{k}_{n_p}$ as in Algorithm 2. When creating the orthonormal basis $(L_j^{n_p})$, ensure compatibility with the inclusion $V_{n_{p-1}} \subset V_{n_p}$, in the sense that

$$\text{Span}(L_1^{n_p}, \ldots, L_{n_{p-1}}^{n_p}) = V_{n_{p-1}}.$$

Having defined the evaluation points $\{x^1, \ldots, x^{m_{p-1}}\}$, draw the new evaluation points x^i for $i = m_{p-1} + 1, \ldots, m_p$ according to

$$d\rho_p := \frac{\alpha_p}{m_p - m_{p-1}} \left(\frac{m_p}{n_p} \sum_{j=1}^{n_p} |L_j^{n_p}|^2 - \frac{m_{p-1}}{n_{p-1}} \sum_{j=1}^{n_{p-1}} |L_j^{n_p}|^2 \right) d\mu, \quad (4.3)$$

with α_p a normalization factor.

Remark 4.2 Note that the non-negativity of ρ_p is only guaranteed when $(m_p/n_p)_{p \geq 1}$ is non-decreasing, a condition which is easily met since m_p has to grow as $n_p \ln n_p$. If we had taken m_p exactly linear with respect to the dimension n_p, the terms with $j \leq n_{p-1}$ in the expression (4.3) would cancel; hence, $d\rho_p$ would only be an approximation of $\frac{k_{n_p} - k_{n_{p-1}}}{n_p - n_{p-1}} d\mu$.

Remark 4.3 Another approach to hierarchical sampling was proposed in [17] and consists in drawing constant proportions of samples according to the measure $|L_j|^2 d\mu$. It was adapted in [2] to our setting of interest where the L_j cannot be exactly computed.

Remark 4.4 In the above algorithm, the various sections $\{x^{m_{k-1}+1}, \ldots, x^{m_k}\}$ of $\{x^1, \ldots, x^{m_p}\}$ for $k = 1, \ldots, p$ are drawn according to different probability measures. The sample $\{x^1, \ldots, x^{m_p}\}$ is thus not i.i.d. anymore, which affects the proof of the convergence theorem given below. Instead it may be thought of as a deterministic mixture of collections of i.i.d. samples, as in Theorem 2 of [17].

At any iteration q, we use the evaluations of u at all points $x^1, \ldots, x^m$ as follows to compute a least-squares approximation $\widetilde{u}_n \in V_n$, where $n := n_q$ and $m := m_q$. We denote by w the weight function defined by

$$w(x) \sum_{p=1}^{q} (m_p - m_{p-1}) \, d\rho_p = m \, d\mu,$$

and solve the weighted least-squares problem (2.1). The following result shows that instance optimality is maintained at every step q, with a cumulated sampling budget m_q that is near-optimal.

Theorem 4.5 *Take numbers $\delta_p, \varepsilon_p \in]0, 1[$ such that $\varepsilon := \sum_{p=1}^{q} \varepsilon_p < 1$ and $\delta := \sum_{p=1}^{q} \delta_p < 1/2$, and define $c_\delta = ((1 + \delta) \ln (1 + \delta) - \delta)^{-1}$. Assume that, for all $p \geq 1$,*

$$M_{n_p} \geq 2 \kappa \, c_{\delta_p} \, n_p \ln \frac{2n_p}{\varepsilon_p} \quad and \quad m_p \geq \frac{\gamma}{1 - 2\delta} \, n_p \ln \frac{2n_p}{\varepsilon},$$

with κ the constant in the assumption (4.1), and that m_p/n_p is an non-decreasing function of p. Then, with $n := n_q$ and $m := m_q$, the convergence bounds (3.9) and (3.10) simultaneously hold for all $q \geq 1$, with

$$\eta(m) = \frac{4}{(1 - 2\delta)} \frac{n}{m} \leq \frac{4}{\gamma \ln(2n/\varepsilon)}$$

and $E := \{\|G - I\|_2 \leq \frac{1}{2}$ and $\|G_p - I\|_2 \leq \delta_p, \ p \geq 1\}$, which satisfies $\Pr(E^c) \leq 2\varepsilon$.

The proof of this theorem requires a refinement of Lemma 2.1, due to the fact that the x^i are not anymore identically distributed. This uses the following tail bound, directly obtained from the matrix Chernoff bound in [20].

Proposition 4.6 *Consider a finite sequence $\{X^i\}_{i=1,\ldots,m}$ of independent, random, self-adjoint matrices with dimension n. Assume that each matrix satisfies $0 \leq X^i \leq R I$ almost surely, and that $\sum_{i=1}^{m} \mathbb{E}(X^i) = I$. Then for all $\delta \in]0, 1[$,*

$$\Pr \left(\left\| \sum_{i=1}^{m} X^i - I \right\|_2 > \delta \right) \leq 2 n \exp \left(-\frac{1}{c_\delta R} \right),$$

where $c_\delta = ((1 + \delta) \ln (1 + \delta) - \delta)^{-1}$ as in Theorem 4.5.

Proof of Theorem 4.5 By the same argument as in Theorem 4.1, we find that the event

$$B = \{\|G_p - I\|_2 \leq \delta_p, \quad p \geq 1\}$$

has probability larger than $1 - \varepsilon$, where the G_p are as in the proof of Theorem 4.1.

We then fix a value of q and for $n = n_q$ and $m = m_q$, we study the Gramian matrix G which is the sum of the independent, but not identically distributed, matrices

$$X^i := \frac{1}{m} w(x^i) \, (L_j(x^i) L_k(x^i))_{j,k=1,\ldots,n}, \quad i = 1, \ldots, m.$$

Then, with the notation $H(x) = (L_j(x) L_k(x))_{j,k=1,\ldots,n}$,

$$\sum_{i=1}^{m} \mathbb{E}(X^i) = \sum_{p=1}^{q} (m_p - m_{p-1}) \int_D \frac{1}{m} w(x) \, H(x) \, d\rho_p(x) = \int_D H(x) \, d\mu(x) = I,$$

and

$$\|X^i\|_2 = \frac{1}{m} w(x^i) \sum_{j=1}^{n} |L_j(x^i)|^2 \le \frac{1}{m} \|w \, k_n\|_{L^\infty} =: R.$$

One also has, under the event B, $\int_D |L_j^{n_p}|^2 \, d\mu \le \frac{1}{1-\delta_p}$ for $j = 1, \ldots, n_p$ so $\alpha_p \ge 1 - \delta_p$, and consequently

$$\begin{aligned}
\frac{m}{w} &= \sum_{p=1}^{q} (m_p - m_{p-1}) \frac{d\rho_p}{d\mu} \\
&= \sum_{p=1}^{q} \alpha_p \left(\frac{m_p}{n_p} \sum_{j=1}^{n_p} |L_j^{n_p}|^2 - \frac{m_{p-1}}{n_{p-1}} \sum_{j=1}^{n_{p-1}} |L_j^{n_p}|^2 \right) \\
&\ge \sum_{p=1}^{q} \left(\frac{m_p}{n_p} \frac{1 - \delta_p}{1 + \delta_p} k_{n_p} - \frac{m_{p-1}}{n_{p-1}} k_{n_{p-1}} \right) \\
&\ge \frac{m}{n} k_n - \sum_{p=1}^{q} \frac{m_p}{n_p} \frac{2\delta_p}{1 + \delta_p} k_{n_p} \\
&\ge (1 - 2\delta) \frac{m}{n} k_n,
\end{aligned}$$

so $R = \frac{1}{m} \|w \, k_n\|_{L^\infty} \le \frac{1}{(1-2\delta)} \frac{n}{m}$. Applying Proposition 4.6, we find that

$$\mathrm{Pr}_x \left(\|G - I\|_2 > \frac{1}{2} \,\Big|\, B \right) \le 2n \exp\left(-\frac{1}{\gamma \, R} \right) \le 2n \exp\left(-\frac{1 - 2\delta}{\gamma} \frac{m}{n} \right) \le \varepsilon.$$

Therefore, since $E := B \cap \{\|G - I\|_2 \le \frac{1}{2}\}$, we find that $\Pr(E) \ge 1 - 2\varepsilon$.

In order to prove the convergence bounds (3.9) and (3.10), we cannot proceed as in Corollary 3.5 by simply invoking Theorem 3.2, because the x^i are not identically distributed. This leads us to modify the statement of Lemma 2.2 and its proof given

in [9], following a strategy proposed in [17]. First, using similar arguments as in [9], we find that

$$\mathbb{E}(\|u - \tilde{u}_n\|^2 \chi_E) \leq e_n(u)^2 + 4\,\mathbb{E}\left(\sum_{k=1}^n |\langle L_k, g\rangle_m|^2 \chi_E\right), \qquad (4.4)$$

where $g = u - P_n u$ is the projection error. For each $k = 1, \ldots, n$, we define $g_k := w\,L_k\,g$ and write

$$
\begin{aligned}
&\mathbb{E}\left(|\langle L_k, g\rangle_m|^2 \chi_E\right)\\
&\quad \leq \mathbb{E}\left(|\langle L_k, g\rangle_m|^2 \chi_B\right)\\
&\quad = \frac{1}{m^2} \sum_{1 \leq i, j \leq m} \mathbb{E}\left(g_k(x^i)\,g_k(x^j)\,\chi_B\right)\\
&\quad = \frac{1}{m^2} \mathbb{E}_z\left(\chi_B\left(\sum_{1 \leq i \leq m} \mathbb{E}_x\left(|g_k(x^i)|^2\right) + \sum_{i \neq j} \mathbb{E}_x\left(g_k(x^i)\,g_k(x^j)\right)\right)\right)\\
&\quad \leq \frac{1}{m^2} \mathbb{E}_z\left(\chi_B\left(\sum_{1 \leq i \leq m} \mathbb{E}_x\left(|g_k(x^i)|^2\right) + \left(\sum_{1 \leq i \leq m} \mathbb{E}_x\left(g_k(x^i)\right)\right)^2\right)\right)\\
&\quad = \mathbb{E}_z\left(\chi_B\left(\frac{1}{m} \mathbb{E}_t\left(|g_k(t)|^2\right) + \left(\mathbb{E}_t\left(g_k(t)\right)\right)^2\right)\right),
\end{aligned}
$$

where t is a random variable distributed according to $\sum_{p=1}^q \frac{m_p - m_{p-1}}{m}\,d\rho_p = \frac{1}{w}\,d\mu$. We then note that

$$\mathbb{E}_t(g_k(t)) = \int_D g\,L_k\,d\mu = 0$$

since $g \in V_n^\perp$, and that $\sum_{k=1}^n |g_k(t)|^2 = w(t)^2\,g(t)^2\,k_n(t)$. Therefore,

$$
\begin{aligned}
\mathbb{E}\left(\sum_{k=1}^n |\langle L_k, g\rangle_m|^2 \chi_E\right) &\leq \mathbb{E}_z\left(\chi_B\,\frac{1}{m}\int_D w\,k_n\,g^2 d\mu\right)\\
&\leq \mathbb{E}_z\left(\chi_B\,R\,\|g\|^2\right)\\
&\leq \frac{1}{(1 - 2\delta)}\,\frac{n}{m}\,e_n(u)^2
\end{aligned}
$$

Combining this with (4.4), we finally obtain

$$\mathbb{E}(\|u - \tilde{u}_n\|^2 \chi_E) \leq \left(1 + \frac{4}{(1 - 2\delta)}\,\frac{n}{m}\right) e_n(u)^2.$$

$\square$

Remark 4.7 If a stopping time q is known in advance, the simplest choice is to take $\varepsilon_p = \varepsilon/q$ and $\delta_p = \delta/q$. If the stopping time q is not known in advance, we can take for instance $\varepsilon_p = \frac{6}{\pi^2} \frac{\varepsilon}{p^2}$ and $\delta_p = \frac{6}{\pi^2} \frac{\delta}{p^2}$. As $c_\delta \sim \frac{2}{\delta^2}$ when $\delta \to 0$, this choice only increases the number M_p of sample points z^i by a factor p^4, which is satisfactory in view of the previous remarks.

5 Estimates on the Inverse Christoffel Function

We have seen that the success of Algorithm 1 is based on the offline sampling condition (3.8), which means that a uniform upper bound $B(n)$ on the inverse Christoffel function k_n is needed in the first place. Likewise, the multilevel Algorithms 2 and 3 from Sect. 4 are based on the assumption (4.1), whose verification requires pointwise upper and lower estimates on $k_n(x)$. In this section we establish such bounds and pointwise estimates on general domains when the V_n are spaces of algebraic multivariate polynomials of varying total degree. Throughout this section, we assume that

$$\mu = \mu_D = |D|^{-1} \chi_D \, \mathrm{d}x$$

is the uniform measure over D, which is thus assumed to have finite Lebesgue measure $|D|$.

5.1 Comparison Strategies

Our vehicle for estimating the Christoffel function is a general strategy, first introduced in [15]: compare D with reference domains R for which the Christoffel function can be estimated. For simplicity, we use the notation

$$L^2(R) = L^2(R, \mu_R),$$

for any domain R where $\mu_R = |R|^{-1} \chi_R \, \mathrm{d}x$ is the uniform measure over R. In order to make clear the dependence on the domain, we define

$$k_{n,R}(x) = \max_{v \in V_n} \frac{|v(x)|^2}{\|v\|_{L^2(R)}^2}$$

and

$$K_{n,R} = \|k_n\|_{L^\infty(R)} = \max_{v \in V_n} \frac{\|v\|_{L^\infty(R)}^2}{\|v\|_{L^2(R)}^2}.$$

We first state a pointwise comparison result.

Lemma 5.1 *For $x \in D$, let R be such that $x \in R \subset D$ and $\beta\,|D| \leq |R|$ for some $\beta \in]0, 1]$. Then*

$$k_{n,D}(x) \leq \beta^{-1}\,k_{n,R}(x).$$

Conversely, let S be such that $D \subset S$ and $\beta\,|S| \leq |D|$ for some $\beta \in]0, 1]$. Then

$$k_{n,D}(x) \geq \beta\,k_{n,S}(x).$$

Proof For any $v \in V_n$, we have

$$|v(x)|^2 \leq k_{n,R}(x)\,\|v\|_{L^2(R)}^2 \leq k_{n,R}(x)\,\frac{|D|}{|R|}\,\|v\|_{L^2(D)}^2,$$

and

$$|v(x)|^2 \leq k_{n,D}(x)\,\|v\|_{L^2(D)}^2 \leq k_{n,D}(x)\,\frac{|S|}{|D|}\,\|v\|_{L^2(S)}^2.$$

Optimizing over v gives the upper and lower estimates of $k_{n,D}(x)$. □

Obviously, a framing on $K_{n,D}$ can be readily derived as follows, by application of the above lemma to any point in D.

Proposition 5.2 *Assume that there exist a family $\mathcal{R}$ of reference domains with the following properties:*

(i) For all $x \in D$ there exist $R_x \in \mathcal{R}$ such that $x \in R_x \subset D$.
(ii) There exists a constant $\beta \in]0, 1]$ such that $|R| \geq \beta\,|D|$ for all $R \in \mathcal{R}$.

Then, one has

$$K_{n,D} \leq \beta^{-1} \sup_{x \in D} k_{n,R_x}(x) \leq \beta^{-1} \sup_{R \in \mathcal{R}} K_{n,R}.$$

Likewise, for any $S \in \mathcal{R}$ such that $D \subset S$ and $|D| \geq \beta\,|S|$, one has

$$K_{n,D} \geq \beta \sup_{x \in D} k_{n,S}(x).$$

In what follows, we apply this strategy to spaces V_n of multivariate algebraic polynomials. Throughout this section, we consider

$$V_n = \mathbb{P}_\ell := \mathrm{span}\{x \mapsto x^\nu = x_1^{\nu_1} \dots x_d^{\nu_d} \,:\, |\nu| = \nu_1 + \dots + \nu_d \leq \ell\}, \quad (5.1)$$

the space of polynomials with total degree less or equal to ℓ, for which we have

$$n = \binom{d+\ell}{\ell}.$$

We assume D is a bounded open set of $\mathbb{R}^d$.

It is important to note that V_n is invariant by affine transformation. As a consequence, if A is any affine transformation, one has

$$R' = A(R) \implies k_{n,R'}(A(x)) = k_{n,R}(x), \quad x \in R,$$

and in particular $K_{n,R'} = K_{n,R}$.

5.2 Lipschitz Domains

In the case of the cube $Q = [-1, 1]^d$, we may express $k_{n,Q}$ by using tensorized Legendre polynomials, that is

$$k_{n,Q}(x) = \sum_{|v| \leq \ell} |L_v(x)|^2, \quad L_v(x) = \prod_{i=1}^{d} L_{v_i}(x_i),$$

where the univariate polynomials $t \mapsto L_j(t)$ are normalized in $L^2([-1, 1], \frac{dt}{2})$. Using this expression, it can be proved by induction on the dimension d that

$$K_{n,Q} \leq n^2, \quad n \geq 1,$$

see Lemma 1 in [6]. Therefore, by affine invariance,

$$K_{n,R} \leq n^2, \quad n \geq 1,$$

for any d-dimensional parallelogram R. Using this result, we may bound the growth of Christoffel functions from above for a general class of domains.

Definition 5.3 An open set $D \subset \mathbb{R}^d$ satisfies the inner cone condition if there exist $\bar{r} > 0$ and $\theta \in (0, \pi)$, such that for all $x \in \overline{D}$, there exists a unit vector u such that the cone

$$C_{\bar{r}, \theta}(x, u) := \{x + r\, v, \, 0 \leq r \leq \bar{r}, \, |v| = 1, \, u \cdot v \geq \cos(\theta)\}$$

is contained in $\overline{D}$. In particular, any Lipschitz domain $D \subset \mathbb{R}^d$ satisfies the inner cone condition (see, e.g., §4.11 in [1]).

Theorem 5.4 *Let $D \subset \mathbb{R}^d$ be a bounded domain that satisfies the inner cone condition. Then, one has*

$$K_n \leq C_D\, n^2, \quad n \geq 1, \tag{5.2}$$

where C_D depends on d, $|D|$, and on $\bar{r}$ and θ in the previous definition.

Proof The uniform cone condition ensures that there exists $\kappa = \kappa(\bar{r}, \theta, d) > 0$ such that for any $x \in D$, there exists a parallelogram R such that $x \in R \subset D$ and $|R| = \kappa$. Therefore, applying Proposition 5.2 with $\mathcal{R}$ the family of all parallelograms of area κ, one obtains (5.2) with $C_D = \frac{|D|}{\kappa}$.   □

Remark 5.5 The bound $K_{n,Q} \leq n^2$ is actually established in [6] for the more general class of polynomial spaces of the form

$$V_n = \mathbb{P}_\Lambda := \mathrm{span}\{x \mapsto x^\nu \ : \ \nu \in \Lambda\}, \quad \#(\Lambda) = n,$$

where $\Lambda \in \mathbb{N}^d$ is downward closed, i.e., such that

$$\nu \in \Lambda \quad \text{and} \quad \widetilde{\nu} \leq \nu \implies \widetilde{\nu} \in \Lambda.$$

These spaces are, however, not invariant by affine transformation and so one cannot apply the above method to treat general domains with inner cone condition. On the other hand, these spaces are invariant by affine transformation of the form $x \mapsto x_0 + Mx$ where M is a diagonal matrix, therefore transforming the cube Q into an arbitrary rectangle R aligned with the coordinate axes. As observed in [3], this leads to a bound of the form (5.2) for any domain D that satisfies the following geometrical property: for all $x \in D$ there exists a rectangle R aligned with the coordinate axes such that $x \in R \subset D$ and $|R| \geq \beta |D|$. Note that this property does not readily follows from a smoothness property of the boundary, in particular there exists smooth domains for which this property does not hold.

5.3 Smooth Domains

We next investigate smooth domains. For this purpose, we replace parallelograms by ellipsoids as reference domains. In the case of the unit ball $B := \{|x| \leq 1\}$, it is known [21] that the Christoffel function reaches its maximum on the unit sphere $S := \{|x| = 1\}$, where we have

$$k_{n,B}(x) = \binom{\ell + d + 1}{\ell} + \binom{\ell + d - 2}{\ell - 1}.$$

In order to estimate how this quantity scales with $n = \binom{\ell+d}{\ell}$ we use the fact that for any integer m, one has

$$e \left(\frac{m}{e}\right)^m \leq m! \leq m^m.$$

For the lower bound, we bound from below the first term

$$\binom{\ell + d + 1}{\ell} = \binom{\ell + d}{\ell} \frac{\ell + d + 1}{d + 1} = n \frac{\ell + d + 1}{d + 1}$$

$$\geq \frac{n}{d\,e^{1/d}}(\ell + d + 1) \geq \frac{n}{e\,(d!)^{1/d}}\left(\frac{(\ell + d)!}{\ell!}\right)^{1/d} = e^{-1}\,n^{\frac{d+1}{d}},$$

which leads to

$$K_{n,B} \geq k_{n,B}(x) \geq e^{-1}\,n^{\frac{d+1}{d}}, \quad x \in S. \tag{5.3}$$

For the upper bound, we write

$$\binom{\ell + d + 1}{\ell} + \binom{\ell + d - 2}{\ell - 1} = \binom{\ell + d}{\ell}\left(\frac{\ell + d + 1}{d + 1} + \frac{\ell d}{(\ell + d)(\ell + d - 1)}\right)$$

$$\leq n\left(\frac{\ell + d + 1}{d + 1} + 1\right) = n\left(\frac{\ell}{d + 1} + 2\right),$$

Since

$$n^{1/d} = (d!)^{-1/d}\left(\frac{(\ell + d)!}{\ell!}\right)^{1/d} \geq \frac{\ell + 1}{d} \geq \frac{\ell}{d + 1},$$

we find that

$$k_{n,B}(x) \leq \left(n^{1/d} + 2\right)n \leq 3\,n^{\frac{d+1}{d}}.$$

By affine invariance, we thus obtain

$$e^{-1}\,n^{\frac{d+1}{d}} \leq K_{n,E} \leq 3\,n^{\frac{d+1}{d}}, \tag{5.4}$$

for all ellipsoids E. This leads to the following result.

Theorem 5.6 *Assume $D \subset \mathbb{R}^d$ is a bounded domain with C^2 boundary. Then, one has*

$$K_n \leq C_D\,n^{\frac{d+1}{d}}, \quad n \geq 1, \tag{5.5}$$

where C_D depends on D.

Proof Since the boundary of D has finite curvature, we are ensured that there exists a $\beta > 0$ such that for any $x \in D$, there exist an ellipsoid E such that $x \in E \subset D$ and $|E| \geq \beta\,|D|$. Therefore, applying Proposition 5.2 with $\mathcal{R}$ the family of ellipsoids with area larger than $\beta\,|D|$, we obtain (5.5) with $C_D = 3\,\beta^{-1}$. $\square$

Remark 5.7 In the above argument, one could simply use balls instead of ellipsoids, however, at the price of diminishing the value of β and thus raising the constant C_D.

We next give a general lower bound for K_n showing that the above rate for smooth domains is sharp.

Theorem 5.8 *Let $D \subset \mathbb{R}^d$ be an arbitrary bounded domain, and let B be its Chebychev ball, that is, the smallest closed ball that contains D. Then, one has*

$$K_{n,D} \geq e^{-1} \frac{|B|}{|D|} n^{\frac{d+1}{d}}, \quad n \geq 1.$$

Proof As $\overline{D}$ is compact and B is the smallest possible ball containing $\overline{D}$, there exists a point $x \in \overline{D} \cap \partial B$, and by Lemma 5.1 one has

$$K_{n,D} \geq k_{n,D}(x) \geq \frac{|D|}{|B|} k_{n,B}(x) \geq e^{-1} \frac{|B|}{|D|} n^{\frac{d+1}{d}},$$

where the last inequality follows from (5.3) and affine invariance. □

5.4 Pointwise Bounds for Piecewise Smooth Domains

As already observed, it may be needed to get sharper bounds on $k_n(x)$ that depend on the point x, in particular when checking the validity of (4.1). In the case of algebraic polynomials in dimension $d = 2$, so $n = \frac{(\ell+1)(\ell+2)}{2}$, such bounds have been obtained for a particular class of piecewise smooth domains with exiting corners, in the following result from [19].

Theorem 5.9 *Let $D \subset \mathbb{R}^2$ be a bounded open set such that $\partial D = \cup_{i=1}^K \Gamma_i$, where the Γ_i are one-to-one C^2 curves that intersect only at their extremities, at which points the interior angles belong to $(0, \pi)$. Then, there exists a constant C_D that only depends on D such that, for all $x \in D$,*

$$C_D^{-1} k_n(x) \leq n \min_{(i,j) \in S} \rho_i(x) \, \rho_j(x) \leq C_D \, k_n(x), \quad n \geq 1,$$

where S consists of the (i, j) such that Γ_i and Γ_j intersects, and $\rho_i(x) := \min\left(\ell, d(x, \Gamma_i)^{-1/2}\right)$.

For the square domain $D = Q = [-1, 1]^2$, this implies that $k_{n,Q}(x) \sim n \ell^2 \sim n^2$ when x is close enough to a corner, and we retrieve the bound $K_{n,Q} \leq C_D n^2$ from (5.2). It is also proved that $k_n(x) \sim n \min\left(\ell, d(x, \partial D)^{-1/2}\right)$ for bidimensional domains with $\mathcal{C}^2$ boundary, which is consistent with the global bound (5.5) in the case $d = 2$.

5.5 Rate of Growth of $K_{n,D}$ and Order of Cuspitality

We end this section by a more technical but systematic approach which allows us to estimate the rate of growth of the inverse Christoffel function in a sharp way for domains D that could either be smooth, of α-Hölder boundary, or even with cusps of a given order. It is based on using the following more elaborate reference domain that describes a certain order of smoothness at the origin.

Definition 5.10 For $\alpha_1, \ldots, \alpha_{d-1} \in]0, 2]$, denote $R_{\alpha_1,\ldots,\alpha_{d-1}}$ the reference domain

$$R_{\alpha_1,\ldots,\alpha_{d-1}} := \left\{ x \in [-1, 1]^d, \max_{1 \leq i \leq d-1} |x_i|^{\alpha_i} \leq x_d \right\}.$$

We shall establish upper and lower bounds for $K_{n,D}$ based on comparisons between D and affine transformations of this reference domain, by adapting certain techniques and results from [12]. The upper bound is as follows.

Theorem 5.11 *Let D be a bounded domain. Assume there exist $\alpha_1, \ldots, \alpha_{d-1} \in]0, 2]$ and $\beta > 0$ such that, for all $x \in D$, one can find an affine map A such that $A(0) = x$, $A(R_{\alpha_1,\ldots,\alpha_{d-1}}) \subset \overline{D}$ and $|A(R_{\alpha_1,\ldots,\alpha_{d-1}})| \geq \beta |D|$. Then*

$$K_{n,D} \leq C_D \, n^{\frac{1}{d}\left(2 + \sum_{i=1}^{d-1} 2/\alpha_i\right)}, \tag{5.6}$$

where C_D is a constant depending only on D.

This result is obtained with the extension strategy proposed in [12], which consists in combining Proposition 5.13 below with a comparison of domains. Such a method was applied in the same paper to the case of smooth domains, polytopes, some two-dimensional domains, and l^α balls in $\mathbb{R}^d$, which all correspond to the situation $\alpha_1 = \cdots = \alpha_{d-1} \in [1, 2]$ in our theorem. We give below a series of intermediate results that lead to the proof of Theorem 5.11.

Lemma 5.12 *For $\alpha \in]0, 2]$ and $n \geq 1$, the function $f : x \mapsto \frac{1}{9\ell^2} + \beta x^2 - |x|^\alpha$ remains non-negative on $\mathbb{R}$ as soon as $\beta \geq \frac{\alpha}{2} \left(\frac{9}{2} (2 - \alpha) \ell^2 \right)^{\frac{2-\alpha}{\alpha}}$.*

Proof As f is symmetric, one only has to consider this function on $\mathbb{R}_+$. For $x > 0$, $f'(x) = 2\beta x - \alpha x^{\alpha-1}$ cancels only at $x_0 = \left(\frac{\alpha}{2\beta} \right)^{\frac{1}{2-\alpha}}$, so

$$\min_{x \in \mathbb{R}} f(x) = f(x_0) = \frac{1}{9\ell^2} - \frac{2-\alpha}{2} \left(\frac{\alpha}{2\beta} \right)^{\frac{\alpha}{2-\alpha}},$$

which is non-negative if and only if $\beta \geq \frac{\alpha}{2} \left(\frac{9}{2} (2 - \alpha) \ell^2 \right)^{\frac{2-\alpha}{\alpha}}$. □

The following result is Theorem 5.2 from [12].

Proposition 5.13 *Suppose $D \subset \mathbb{R}^d$ is a compact set and T is an affine transformation of $\mathbb{R}^d$ such that $T(B(0, 1)) \subset D$. Then*

$$k_{n,D} \left(T \left(0, \ldots, 0, 1 + \frac{1}{3\ell^2} \right) \right) \leq c \, |\det T|^{-1} \ell^{d+1}.$$

where c depends only on d.

Lemma 5.14 *For $\alpha_1, \ldots, \alpha_{d-1} \in \,]0, 2]$, one has*

$$k_{n, R_{\alpha_1, \ldots, \alpha_{d-1}}}(0) \leq C \ell^{2 + \sum_{i=1}^{d-1} 2/\alpha_i},$$

where C depends only on d.

Proof Define $\beta_i = \frac{\alpha_i}{2} \left(\frac{9}{2} (2 - \alpha_i) \ell^2 \right)^{\frac{2-\alpha_i}{\alpha_i}}$ for $1 \leq i \leq d - 1$, and let T be the affine transformation

$$T : x = (x_1, \ldots, x_d) \mapsto \left(\frac{x_1}{\sqrt{3\beta_1}}, \ldots, \frac{x_{d-1}}{\sqrt{3\beta_{d-1}}}, \frac{1}{3} \left(1 + \frac{1}{3\ell^2} - x_d \right) \right).$$

Then, for all $x \in B(0, 1)$, $T(x)_d \in [0, 1]$, and for $(1 \leq i \leq d-1)$, using Lemma 5.12,

$$T(x)_d = \frac{1}{3} \left(1 + \frac{1}{3\ell^2} - x_d \right) \geq \frac{1}{3} \left(\frac{1}{3\ell^2} + x_i^2 \right) = \frac{1}{9\ell^2} + \beta_i T(x)_i^2 \geq T(x)_i^{\alpha_i},$$

so $\max_{1 \leq i \leq d-1} |T(x)_i|^{\alpha_i} \leq T(x)^d$, which implies that $T(B(0, 1)) \subset R_{\alpha_1, \ldots, \alpha_{d-1}}$.

As $T \left(0, \ldots, 0, 1 + \frac{1}{3\ell^2} \right) = 0$, a direct application of Proposition (5.13) gives

$$k_{n, R_{\alpha_1, \ldots, \alpha_{d-1}}}(0) \leq c \, |\det T|^{-1} \ell^{d+1} = 3 \, c \prod_{i=1}^{d} \sqrt{3\beta_i} \, \ell^{d+1}$$

$$\leq C \ell^{d+1 + \sum_{i=1}^{d-1} \frac{2-\alpha_i}{\alpha_i}} = C \ell^{2 + \sum_{i=1}^{d-1} \frac{2}{\alpha_i}}.$$

$\square$

Proof of Theorem 5.11 One simply applies Proposition 5.2 to the family $\mathcal{R}$ of all domains of the form $A(R_{\alpha_1, \ldots, \alpha_{d-1}})$ where A is an affine map such that $|\det A| \, |R_{\alpha_1, \ldots, \alpha_{d-1}}| > \beta \, |D|$. As $\ell \leq L \, n^{1/d}$ for some $L > 0$, we obtain (5.6) with $C_D = \beta^{-1} L^{2 + \sum_{i=1}^{d-1} 2/\alpha_i} C$, the constant C coming from the lemma above. $\square$

We now prove a lower bound based on the same reference domain.

Theorem 5.15 *Let D be a bounded domain. Assume there exist $\bar{x} \in \overline{D}$, $0 < r_1 \leq r_2$, $\alpha_1, \ldots, \alpha_{d-1} \in \,]0, 2]$ and an affine transformation A with $A(0) = \bar{x}$ such that*

$$D \subset A(R_{\alpha_1, \ldots, \alpha_{d-1}}) \cup \left(\overline{B(\bar{x}, r_2)} \setminus B(\bar{x}, r_1) \right).$$

Then

$$K_{n, D} \geq c_D \, n^{\frac{1}{d} \left(2 + \sum_{i=1}^{d-1} 2/\alpha_i \right)},$$

where c_D is a constant depending only on D.

The proof follows the same path as in Theorem 8.1 and Remark 8.4 of [12], but with a radial polynomial centered at x instead of a planar polynomial, that is a univariate polynomial composed with an affine function. This small improvement shows that for a point x and a domain D satisfying the conditions of Theorems 5.11 and 5.15 with the same α_i, the asymptotic behavior of $k_{n,D}(x)$ only depends on D in a neighborhood of x.

We first recall Lemma 6.1 from the same article:

Lemma 5.16 *For any $\ell, m \geq 1$ and $y \in [-1, 1]$, there exists a univariate polynomial $P_{\ell,m,y}$ of degree at most ℓ such that $P_{\ell,m,y}(y) = 1$ and*

$$|P_{\ell,m,y}(x)| \leq c(m) \left(\frac{1 + \ell\sqrt{1 - y^2}}{1 + \ell\sqrt{1 - y^2} + \ell^2 |x - y|} \right)^m, \quad x \in [-1, 1].$$

Taking $y = -1$ and applying a change of variable $x \mapsto \frac{x+1}{2}$, we get as an immediate consequence:

Lemma 5.17 *For any $\ell, m \geq 1$, there exists a univariate polynomial $P_{\ell,m}$ of degree at most ℓ such that $P_{\ell,m}(0) = 1$ and*

$$|P_{\ell,m}(x)| \leq c(m) \min \left(1, \frac{1}{\ell^{2m}|x|^m} \right), \quad x \in [0, 1].$$

We also need a bound on the volume of $R_{\alpha_1,\ldots,\alpha_{d-1}}$.

Lemma 5.18 *For all $r > 0$, $|R_{\alpha_1,\ldots,\alpha_{d-1}} \cap B(0, r)| \leq c\, r^{1+\sum_{i=1}^{d-1} 1/\alpha_i}$.*

Proof Given $r \in [0, 1]$,

$$R_{\alpha_1,\ldots,\alpha_{d-1}} \cap \{x_d = r\} = [-r^{1/\alpha_1}, r^{1/\alpha_1}] \times \cdots \times [-r^{1/\alpha_{d-1}}, r^{1/\alpha_{d-1}}] \times \{r\}$$

has a $(d-1)$-volume equal to $\prod_{i=1}^{d-1} 2\, r^{-\alpha_i}$, so

$$|R_{\alpha_1,\ldots,\alpha_{d-1}} \cap \{0 \leq x_d \leq r\}| = \int_0^r \prod_{i=1}^{d-1} 2\, x_d^{1/\alpha_i} \, dx_d = c\, r^{1+\sum_{i=1}^{d-1} 1/\alpha_i}.$$

As for all $r > 0$, $R_{\alpha_1,\ldots,\alpha_{d-1}} \cap B(0, r) \subset R_{\alpha_1,\ldots,\alpha_{d-1}} \cap \{0 \leq x_d \leq \min(1, r)\}$, we obtain the desired result. $\square$

Proof of Theorem 5.15 Take $\ell_0 = \lfloor \frac{\ell}{2} \rfloor$, $m \geq \frac{1}{2} \sum_{i=1}^{d-1} \frac{1}{\alpha_i}$ and $r_3 \geq r_2$ such that $T(R_{\alpha_1,\ldots,\alpha_{d-1}}) \subset \overline{B}(\bar{x}, r_3)$, and define the multivariate polynomial

$$P(x) = P_{\ell_0,m} \left(\frac{|x - \bar{x}|^2}{r_3^2} \right), \quad x \in \mathbb{R}^d.$$

Then P has degree at most $2\ell_0 \leq \ell$ in each variable, $P(\bar{x}) = 1$, and Lemma 5.17 bounds P from above since $D \subset \overline{B}(\bar{x}, r_3)$. It remains to compute an upper bound of $\|P\|_{L^2(D)}$. For $0 < r < r_1$, one has:

$$
\begin{aligned}
|D \cap B(\bar{x}, r)| &= |T(R_{\alpha_1,\ldots,\alpha_{d-1}}) \cap B(\bar{x}, r)| \\
&\leq |\det T| \, |R_{\alpha_1,\ldots,\alpha_{d-1}} \cap T^{-1}(B(\bar{x}, r))| \\
&\leq |\det T| \, |R_{\alpha_1,\ldots,\alpha_{d-1}} \cap B(0, r\,\lambda_{\max}(T^{-1}))| \\
&\leq c' r^{1+\sum_{i=1}^{d-1} 1/\alpha_i},
\end{aligned}
$$

where in the last line we used Lemma 5.18, and with $c' = \dfrac{c\,|\det T|}{\lambda_{\min}(T)^{1+\sum_{i=1}^{d-1} 1/\alpha_i}}$. Therefore, one can compute

$$
\begin{aligned}
&\|P\|_{L^2(D)}^2 \\
&\leq \left\| c(m) \min\left(1, \frac{1}{\ell^{2m}|x|^m}\right) \right\|_{L^2(D)}^2 \\
&\leq c(m)^2 \left(\int_{D \cap B(\bar{x}, \ell^{-2})} dx + \int_{B(\bar{x}, r_3)} \frac{dx}{\ell^{4m} r_1^{2m}} \right. \\
&\quad \left. + \int_{D \cap B(\bar{x}, r_1) \setminus B(\bar{x}, \ell^{-2})} \left(\frac{1}{|x|^{2m}} - \frac{1}{r_1^{2m}} \right) \frac{dx}{\ell^{4m}} \right) \\
&= c(m)^2 \left(\left| D \cap B\left(\bar{x}, \ell^{-2}\right) \right| + \frac{|B(\bar{x}, r_3)|}{\ell^{4m} r_1^{2m}} + \int_{\ell^{-2}}^{r_1} \frac{2m}{\ell^{4m} r^{2m+1}} |D \cap B(\bar{x}, r)| \, dr \right) \\
&\leq c(m)^2 \left(c' \ell^{-2-\sum_{i=1}^{d-1} 2/\alpha_i} + \frac{|B(\bar{x}, r_3)|}{\ell^{4m} r_1^{2m}} + c' \frac{2m}{\ell^{4m}} \int_{\ell^{-2}}^{r_1} r^{\sum_{i=1}^{d-1} 1/\alpha_i - 2m} \, dr \right) \\
&\leq c'' \max\left(\ell^{-2-\sum_{i=1}^{d-1} 2/\alpha_i}, \ell^{-4m}, \ell^{-4m-2(\sum_{i=1}^{d-1} 1/\alpha_i - 2m + 1)} \right) \\
&= c'' \ell^{-2-\sum_{i=1}^{d-1} 2/\alpha_i},
\end{aligned}
$$

and conclude that $K_{n,D} \geq k_{n,D}(\bar{x}) \geq \dfrac{|P(\bar{x})|^2}{\|P\|_{L^2(D)}^2} \geq c_D \, \ell^{2+\sum_{i=1}^{d-1} 2/\alpha_i}$, with $c_D = 1/c''$.

$\square$

Remark 5.19 These theorems include the case of smooth domains : indeed, taking $\alpha_1 = \cdots = \alpha_{d-1} = 2$ and $e_d = (0, \ldots, 0, 1)$, one has

$$
\begin{aligned}
B\left(\frac{1}{2} e_d, \frac{1}{2}\right) \subset R_{2,\ldots,2} &\subset \left\{ x \in [-1, 1]^d, \; \frac{1}{d-1} \sum_{i=1}^{d-1} |x_i|^2 \leq x_d \right\} \\
&\subset B\left((d-1) e_d, (d-1)\right),
\end{aligned}
$$

so one can recover the results 5.6 and 5.8 , without explicit constants. Similarly, Lipschitz boundaries correspond to the particular values $\alpha_1 = \cdots = \alpha_{d-1} = 1$.

Example 5.20 It becomes useful to take distinct values for the α_i in the case of domains with edges but no corners. For instance, consider $D = \frac{\sqrt{3}}{2} e_d + B\left(\frac{1}{2} e_1, 1\right) \cap B\left(-\frac{1}{2} e_1, 1\right)$. Then $0 \in A(R_{1,2\ldots,2}) \subset \overline{D} \subset B(R_{1,2\ldots,2})$, where A and B are the linear maps defined by

$$A(x_1, \ldots, x_d) = \left(\frac{1}{4} x_1, \frac{1}{2\sqrt{d-2}} x_2, \ldots, \frac{1}{2\sqrt{d-2}} x_{d-1}, \frac{\sqrt{3}}{2} x_d\right)$$

and

$$B(x_1, \ldots, x_d) = \left(3 x_1, \frac{1}{\sqrt{3}} x_2, \ldots, \frac{1}{\sqrt{3}} x_{d-1}, \sqrt{3} x_d\right).$$

Thus, $K_{n,D} \geq k_{n,D}(0) \sim \ell^{d+2}$. Moreover, for all $x \in D$ there exists an affine transformation T such that $\det T \geq 2^{-d}$, $T(D) \subset D$ and $T(0) = x$, so $K_{n,D} \sim \ell^{d+2}$.

Remark 5.21 It is easily seen that for domains having a cusp that points outside, the value of K_n may grow as fast as any polynomial, depending on the order of cuspidality. For instance, given $\alpha \in]0, 2]$, according to Theorems 5.11 and 5.15, one has

$$k_{n, R_{\alpha,\ldots,\alpha}}(0) \sim \ell^{2 + \frac{2}{\alpha}(d-1)},$$

so that $K_{n, R_{\alpha,\ldots,\alpha}} \geq c \, \ell^{2 + \frac{2}{\alpha}(d-1)}$.

6 Numerical Illustration

In this section we give numerical illustrations of the offline and online sampling strategies in the particular case of algebraic polynomials and for different domains. As in the previous section, we consider spaces of polynomials of fixed total degree $V_n = \mathbb{P}_\ell$ as defined by (5.1).

The three considered domains are

1. $D := \{x_1^2 + x_2^2 \leq \frac{2}{\pi}\}$, the ball of area 2.
2. $D := \{-1 \leq x_1 \leq 1, \ |x_1| - 1 \leq x_2 \leq |x_1|\}$, a polygon with a re-entrant corner at $(0, 0)$.
3. $D := \{-1 \leq x_1 \leq 1, \ \sqrt{|x_1|} - 1 \leq x_2 \leq \sqrt{|x_1|}\}$, a domain with a re-entrant cusp at $(0, 0)$.

The measure μ for the error metric $L^2(D, \mu)$ is the uniform probability measure on the considered domain. In all three cases, the domain D is embedded in the unit cube $Q = [-1, 1]^2$, and described by algebraic inequalities. Thus, sampling according to μ is readily performed by uniform sampling on Q, which is done separately on the two coordinates, followed by rejection when $x \notin D$.
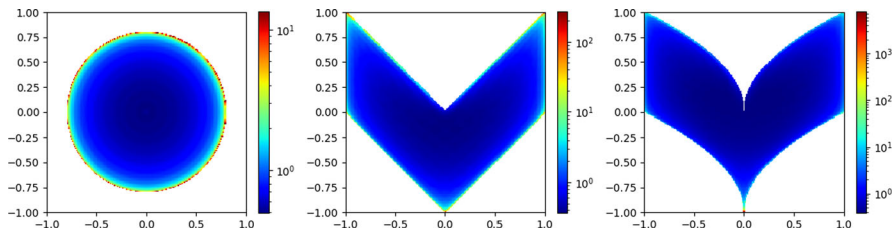
Fig. 1 The three domains (disc, polygon and cusp) and the function k_n/n for $n = 231$

The above three domains are instances of smooth, Lipschitz and cuspital domains, respectively. They are meant to illustrate how the smoothness of the boundary affects the amount of sample needed in the offline state, as rigorously analyzed in the previous section. On these particular domains, we are actually able to exactly integrate polynomials, and therefore in principle to compute the exact orthogonal polynomials L_j up to round-off error due to the orthogonalization procedure. In our numerical tests, the considered total degrees are $\ell = 0, 1, \ldots, 20$, therefore $n = n(\ell) = 1, 3, 6, \ldots, 231$. The intermediate values of n between $n(\ell)$ and $n(\ell+1)$ are treated by complementing the space V_n with the monomials $x_1^{\alpha_1} x_2^{\alpha_2}$ for $\alpha_1 + \alpha_2 = \ell + 1$ in the order $\alpha_2 = 0, \ldots, \ell + 1$. For such values, we could compute the L_j using Cholesky factorization with quadruple precision and check that $|\langle L_j, L_k \rangle - \delta_{j,k}| \leq 10^{-16}$, that is, orthonormality holds up to double precision.

We may thus compute for each value of n the exact inverse Christoffel function k_n and optimal measure $\sigma^* = \frac{k_n}{n}\mu$. Figure 1 displays the three domains and the value of k_n/n for the maximal value $n = 231$ which, as explained by the results in Sect. 5, grows near to the boundary, faster at the exiting corners (and even faster at exiting cusps), and slower in smooth regions or at re-entrant singularities.

This exact computation allows us to compare the optimal sampling strategy based on σ^* and the more realistic strategy based on $\widetilde{\sigma}$ which is computed from the approximate inverse Christoffel function $\widetilde{k}_n$ derived in the offline stage. We next show that both strategies perform similarly well in terms of instance optimality at near-optimal sampling budget. We stress, however, that for more general domains where exact integration of polynomials is not feasible, only the second strategy based on $\widetilde{k}_n$ is viable.

6.1 Sample Complexity of the Offline Stage

We first illustrate the sample complexity M in the offline stage of Algorithm 1. As discussed in §3.2, a sufficient condition to ensure the framing (3.3) between k_n and $\widetilde{k}_n$ is the matrix framing property (3.6) which expresses the fact that the condition number of G_M satisfies the bound

$$\kappa(G_M) \leq c = \frac{c_2}{c_1}.$$
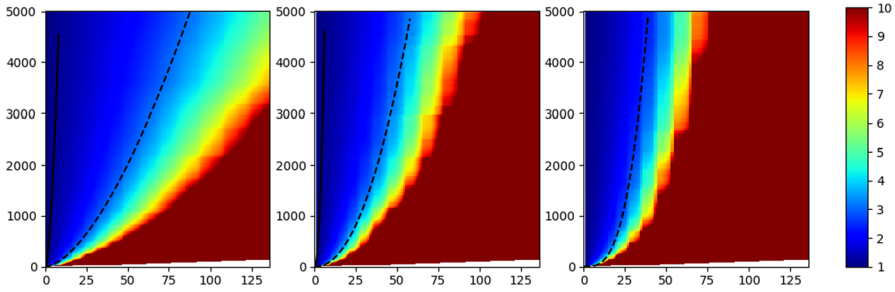
Fig. 2 Conditioning of the matrix G_M for the disc (left), polygon (center) and cusp (right) domains, averaged over 100 realizations, with theoretical value of $M_{\text{suf}}(n)$ (full curve) and adjusted value $M_{\text{adj}}(n)$ (dashed curve). The x-coordinate stands for n and the y-coordinate for M; moreover, the plotted values are saturated at 10 since we are only interested in small condition numbers

For the constants $c_1 = \frac{2}{3}$ and $c_2 = 2$, this occurs with high probability when M is larger than K_n, or a known upper bound $B(n)$, multiplied by logarithmic factors, as expressed by (3.7).

Figure 2 displays the condition number $\kappa(G_M)$, averaged over 100 realizations of the offline sample $\{z^1, \ldots, z^M\}$, as a function of n and $M \geq n$, for the three considered domains. We observe a transition region that illustrates the minimal offline sampling budget $M_{\text{min}}(n)$ that should be practically invested in order for G_M to be well conditioned. For example, if $M_{\text{min}}(n)$ is defined as the minimal value of M such that $\mathbb{E}(\kappa(G_M)) \leq 3$, this value can be estimated and visualized in Fig. 2 as the transition to the dark blue color.

We also draw in full line the value of the sufficient value

$$M_{\text{suf}}(n) := \gamma \, B(n) \ln(2n/\varepsilon),$$

for $\varepsilon = 10^{-2}$ where $B(n)$ is the upper bound for K_n derived from the theoretical analysis of Sect. 5. This upper bound is $3n^{3/2}$ for the disc in view of (5.4) and $2n^2$ for the polygonal domain by application of Proposition 5.2 with $\beta = \frac{1}{2}$, since D is the union of two parallelograms of equal size. While the sampling budget $m = M_{\text{suf}}(n)$ guarantees that $\kappa(G_M) \leq 3$ with high probability - here 0.99 - the plots reveal that this budget is by far an over-estimation of $M_{\text{min}}(n)$.

We draw in dashed line the adjusted values $M_{\text{adj}}(n) = C_{\text{adj}} M_{\text{suf}}(n)$ where the multiplicative constant is picked as small as possible with the constraint of still fitting the requirement $\mathbb{E}(\kappa(G_M)) \leq 3$, thus better fitting the minimal budget $M_{\text{min}}(n)$. We find that constant C_{adj} is approximately $\frac{1}{45}$ for the disc and $\frac{1}{120}$ for the polygon. It is even smaller for the cusp domain, for which Theorem 5.11 with $\alpha_1 = \frac{1}{2}$ yields an upper bound of the form $B(n) = Cn^3$ with a constant C that can be numerically estimated but turns out to be very pessimistic.

In summary, the offline sampling budget $M_{\text{suf}}(n)$ suggested by the theoretical analysis is always pessimistic by a large multiplicative constant. Let us remind that the value $M_{\text{min}}(n)$ is typically not accessible to us since G_M and its condition number cannot be exactly evaluated for more general domains D.
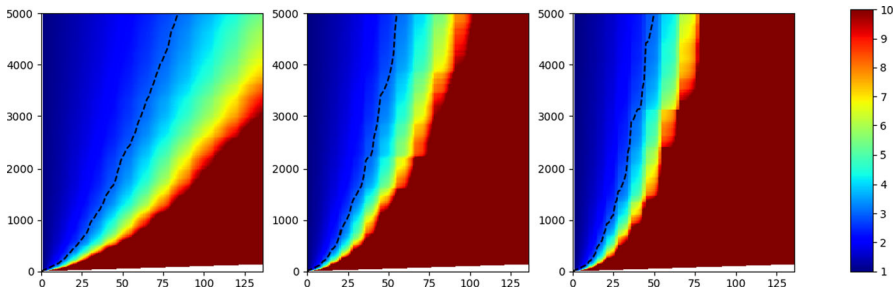
Fig. 3 Conditioning of the matrix T for the disc (left), polygon (center) and cusp (right) domains, and value of $M_{\mathrm{emp}}(n)$ (dashed curve), averaged over 100 realizations

This state of affair justifies the use of the empirical method outlined in §3.3 for selecting a good value of M. Recall that this approach consists in raising M until the conditioning of the computable matrix T becomes less than some prescribed value, for example $\kappa(T) \leq 3$. Figure 3 displays the conditioning $\kappa(T)$ again averaged over 100 realizations of the offline sample, as well as the curve showing the empirical value $M_{\mathrm{emp}}(n)$ which corresponds to the smallest value of M such that $\kappa(T) \leq 3$. It reveals the relevance of the empirical approach: due to the very good fit between $\kappa(T)$ and $\kappa(G_M)$, the value $M_{\mathrm{emp}}(n)$ appears as a much sharper estimate for $M_{\mathrm{min}}(n)$ than $M_{\mathrm{suf}}(n)$.

6.2 Sample Complexity of the Online Stage

We next study the sample complexity m of the online stage of Algorithm 1 through the conditioning of the matrix $G = (\langle L_j, L_k \rangle_m)_{j,k=1,\ldots,n}$, where $\langle \cdot, \cdot \rangle_m$ is the inner product associated to the discrete norm

$$\|v\|_m^2 := \frac{1}{m} \sum_{i=1}^m w(x^i)\,|v(x^i)|^2.$$

For the sampling measure σ and weight w, we both consider:

(i) The optimal sampling measure $d\sigma^* := \frac{k_n}{n}\,d\mu$ and weight $w^* = \frac{n}{k_n}$, which, for these particular domains, can be exactly computed from the L_j, but are not accessible for more general domains.

(ii) The empirical sampling measure $d\widetilde{\sigma} := \frac{\widetilde{k}_n}{n}\,d\mu$ and weight $\widetilde{w} = \frac{n}{\widetilde{k}_n}$ where $\widetilde{k}_n$ has been obtained from the offline stage, using the previously described empirical choice of M.

Figure 4 displays the condition number $\kappa(G)$, as a function of m and n, for both choices and the three domains. In order to illustrate the fluctuations of $\kappa(G)$, we display an averaging over 100 realizations when using k_n, and one single realization when using $\widetilde{k}_n$. While the behavior for a single realization is more chaotic, we find that in both cases, as expected, the online sampling budget $m(n)$ which ensures that G
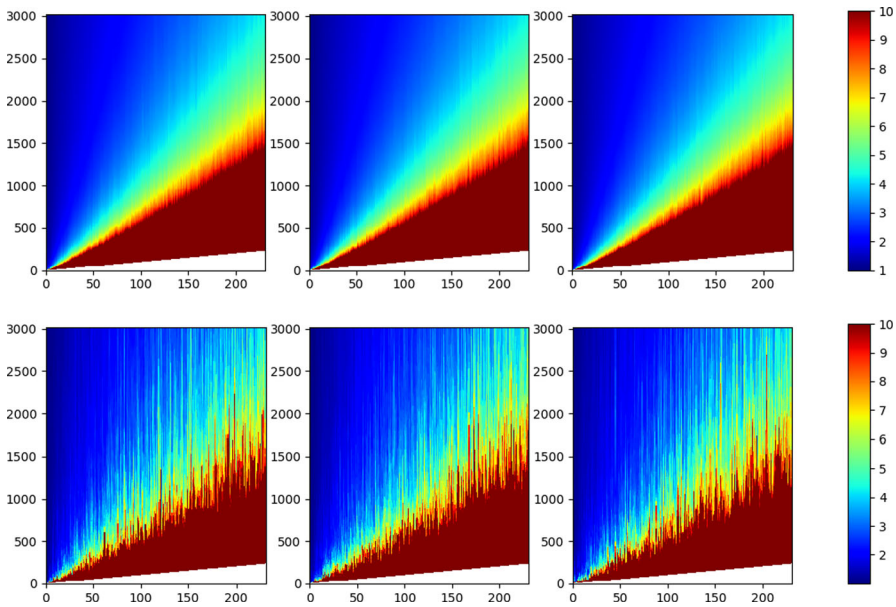
Fig. 4 Conditioning of $G = \langle L_j, L_k \rangle_m$ depending on m and n for the disc (left), polygon (center) and cusp (right) domains, using an average over 100 realizations with k_n (up), or a single realization with the estimated $\tilde{k}_n$ (down)

is well conditioned, for example $\kappa(G) \leq 3$, grows linearly with n (up to logarithmic factors), now independently of the domain shape.

6.3 Instance and Budget Optimality

In order to illustrate the achievement of our initial goal of instance and budget optimality, we consider the approximation in a polynomial space $V_n = \mathbb{P}_\ell$ of a function u that consists of a polynomial part $u_n \in V_n$ and a residual part $u_n^\perp \in V_n^\perp$ that are both explicitly given in terms of their expansions

$$u_n = \sum_{j=1}^{n} c_j L_j,$$

and

$$u_n^\perp = \sum_{j \geq n+1} c_j L_j.$$

For numerical testing, we take only finitely many non-zero c_j in this second expansion and adjust them so that $\sum_{j \geq n+1} |c_j|^2 = 10^{-4}$. Thus, the best approximation error has
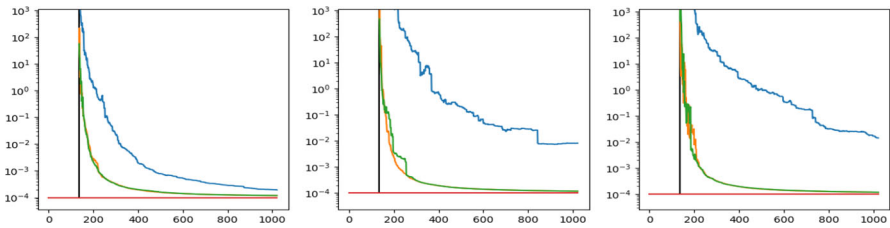
Fig. 5 Mean-square reconstruction error for the disc (left), polygon (center) and cusp (right) domains, with total polynomial degree $\ell = 15$, and sampling measures μ (blue), σ^* (orange) and $\widetilde{\sigma}$ (green). Horizontal red line: best approximation error $e_n(u)^2 = 10^{-4}$. Vertical black line: polynomial dimension $n = 136$ (Color figure online)

value

$$e_n(u) = \|u - u_n\| = \|u_n^\perp\| = 10^{-2}.$$

We study the mean-square error $\mathbb{E}(\|u - P_n^m u\|^2)$ as a function of m and compare the different sampling strategies through their ability to reach this ideal benchmark.

Figure 5 displays the error curves (obtained by averaging $\|u - P_n^m u\|^2$ over 100 realizations) for the three domains and polynomial degree $\ell = 15$ that corresponds to the dimension $n = 136$. For all domains, we observe that the best approximation error is attained up to multiplicative factor 2 with a sampling budget m that is twice larger than n, when using either the optimal sampling measure σ^* based on k_n or the measure $\widetilde{\sigma}$ based on $\bar{k}_n$ obtained in the offline stage Algorithm 1. This does not occur when sampling according to the uniform measure μ: the error remains orders of magnitude above the best approximation error and this effect is even more pronounced as the domain becomes singular. This reflects the fact that with the uniform sampling, the budget m needs to be larger than K_n which has faster growth with n for singular domains.

References

1. Adams, R.A., Fournier, J.J.F.: Sobolev Spaces. Academic Press, Cambridge (2003)
2. Adcock, B., Cardenas, J.M.: Near-optimal sampling strategies for multivariate function approximation on general domains. SIAM J. Math. Data Sci. **2**, 607–630 (2020)
3. Adcock, B., Huybrechs, D.: Approximating Smooth, Multivariate Functions on Irregular Domains, Forum of Mathematics, Sigma, vol. 8. Cambridge University Press, Cambridge (2020)
4. Arras, B., Bachmayr, M., Cohen, A.: Sequential sampling for optimal weighted least squares approximations in hierarchical spaces. SIAM J. Math. Data Sci. **1**, 189–207 (2019)
5. Chkifa, A.: On the Lebesgue constant of Leja sequences for the complex unit disk and of their real projection. J. Approx. Theory **166**, 176–200 (2013)
6. Chkifa, A., Cohen, A., Migliorati, G., Nobile, F., Tempone, R.: Discrete least squares polynomial approximation with random evaluations—application to parametric and stochastic elliptic PDEs. ESAIM Math. Modell. Numer. Anal. **49**, 815–837 (2015)
7. Cohen, A., De Vore, R.: High dimensional approximation of parametric PDEs. Acta Numerica (2015)
8. Cohen, A., Davenport, M., Leviatan, D.: On the stability and accuracy of least squares approximations. Found. Comput. Math. **13**, 819–834 (2013)

9. Cohen, A., Migliorati, G.: Optimal weighted least squares methods. SMAI J. Comput. Math. **3**, 181–203 (2017)
10. DeVore, R.A., Lorentz, G.: Constructive Approximation, vol. 303. Springer, New York (1993)
11. Doostan, A., Hampton, J.: Coherence motivated sampling and convergence analysis of least squares polynomial Chaos regression. Comput. Methods Appl. Mech. Eng. **290**, 73–97 (2015)
12. Ditzian, Z., Prymak, A.: On Nikol'skii inequalities for domains in $\mathbb{R}^d$. Constr. Approx. **44**, 23–51 (2016)
13. Haberstich, C., Nouy, A., Perrin, G.: Boosted optimal weighted least-squares. arXiv:1912.07075 (2019)
14. Jakeman, J.D., Narayan, A., Zhou, T.: A Christoffel function weighted least squares algorithm for collocation approximations. Math. Comput. **86**, 1913–1947 (2017)
15. Kroó, A.: Christoffel functions on convex and starlike domains in $\mathbb{R}^d$. J. Math. Anal. Appl. **421**, 718–729 (2015)
16. Maday, Y., Nguyen, N.C., Patera, A.T., Pau, G.S.H.: A general multipurpose interpolation procedure: the magic points. Commun. Pure Appl. Anal. **8**, 383 (2009)
17. Migliorati, G.: Adaptive approximation by optimal weighted least-squares methods. SIAM J. Numer. Anal. **7**, 2217–2245 (2019)
18. Migliorati, G.: Multivariate approximation of functions on irregular domains by weighted least-squares methods. IMA J. Numer. Anal. (2020). https://doi.org/10.1093/imanum/draa023
19. Prymak, A., Usoltseva, O.: Christoffel functions on planar domains with piecewise smooth boundary. Acta Math. Hungar. **158**, 216–234 (2019)
20. Tropp, J.: User-friendly tail bounds for sums of random matrices. Found. Comput. Math. **12**, 389–434 (2012)
21. Xu, Y.: Asymptotics for orthogonal polynomials and Christoffel functions on a ball. Methods Appl. Anal. **3**, 257–272 (1996)

Authors and Affiliations

Matthieu Dolbeault[1] · Albert Cohen[1]

✉ Albert Cohen
  albert.cohen@sorbonne-universite.fr

  Matthieu Dolbeault
  matthieu.dolbeault@sorbonne-universite.fr

[1] UPMC Univ Paris 06, CNRS, UMR 7598, Laboratoire Jacques-Louis Lions, Sorbonne Université, 4 Place Jussieu, 75005 Paris, France