WEIGHTED APPROXIMATE FEKETE POINTS: SAMPLING FOR LEAST-SQUARES POLYNOMIAL APPROXIMATION[*]

LING GUO[†], AKIL NARAYAN[‡], LIANG YAN[§], AND TAO ZHOU[¶]

Abstract. We propose and analyze a weighted greedy scheme for computing deterministic sample configurations in multidimensional space for performing least-squares polynomial approximations on L^2 spaces weighted by a probability density function. Our procedure is a particular weighted version of the approximate Fekete points method, with the weight function chosen as the (inverse) Christoffel function. Our procedure has theoretical advantages: when linear systems with optimal condition number exist, the procedure finds them. In the one-dimensional setting with any density function, our greedy procedure almost always generates optimally conditioned linear systems. Our method also has practical advantages: our procedure is impartial to the compactness of the domain of approximation and uses only pivoted linear algebraic routines. We show through numerous examples that our sampling design outperforms competing randomized and deterministic designs when the domain is both low and high dimensional.

Key words. uncertainty quantification, least-squares approximations, Fekete points, QR decomposition

AMS subject classifications. 41A10, 65A10

DOI. 10.1137/17M1140960

1. Introduction. The construction of polynomial surrogates that emulate a system response with respect to input parameters is a widely used tool in computational science. A concrete example is provided by problems in parametric uncertainty quantification (UQ), where this approach is frequently called generalized polynomial chaos [16, 36]. The standard approach is to consider a scalar function $f(y)$ depending on inputs/parameters $y \in \Gamma \subset \mathbb{R}^d$ and to approximate this function with a multivariate polynomial expansion. The parameter y is usually interpreted as a random parameter Y, and the basis chosen to perform the expansion is one whose elements ψ_j are orthonormal under the density of Y,

$$(1.1) \qquad f(y) \approx \sum_{j=1}^{N} \widehat{f}_j \psi_j(y).$$

[†]Department of Mathematics, Shanghai Normal University, Shanghai 200234, China (lguo@shnu.edu.cn).

[‡]Department of Mathematics and Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT 84112 (akil@sci.utah.edu).

[§]School of Mathematics, Southeast University, Nanjing, 210096, China (yanliang@seu.edu.cn).

[¶]Corresponding author. LSEC, Institute of Computational Mathematics, Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing 100190, China (tzhou@lsec.cc.ac.cn).

Accurately estimating the coefficients $\widehat{f}_j$ of the expansion is important since these coefficients can be easily manipulated to infer revealing properties, such as statistical moments or parametric sensitivities. Many numerical techniques on how to obtain the polynomial coefficients in UQ problems have been developed in recent years. While early development often focused on "intrusive" methods, such as stochastic Galerkin, much recent effort has concentrated on nonintrusive-type collocation methods [35, 13, 28]. In the collocation framework, one seeks to compute the expansion coefficients via point evaluations of f, and thus constructing a "good" configuration of samples in Γ has become an active area of research. Popular methods include sparse grids [15, 29, 35, 25] and polynomial interpolation [27, 9]. The particular numerical method one uses to compute expansion coefficients often influences the particular sampling strategy, as evidenced by research on sparse approximations using ℓ_1-minimization [12, 37, 38, 19].

In this paper we focus on computing coefficients via the discrete least-squares approach using point evaluations of f. While it is relatively easy to compute the coefficients via standard linear algebraic operations, least-squares approaches have known stability issues. For instance, when using Newton–Cotes quadrature abscissae (equidistant point sets) it is highly unstable even for an infinitely smooth noiseless function unless significant oversampling is performed. Several sampling strategies have been proposed in recent years [10, 22, 40, 17, 34, 26] to improve the stability for least-squares. These methods use both randomized and deterministic sampling methods.

In this paper we propose and analyze a greedily computed deterministic sample set for discrete least-squares, where the objective in the greedy process is a weighted determinant. A discrete least-squares system for computing coefficients in (1.1) from M point evaluations of f utilizes a Vandermonde-like matrix $\boldsymbol{V} \in \mathbb{R}^{M \times N}$, and our procedure greedily forms a set of points A via

$$A \leftarrow A \cup y_{n+1}, \qquad\qquad y_{n+1} = \arg\max_{y \in \Gamma} \det |\boldsymbol{W}\boldsymbol{V}\boldsymbol{V}^T\boldsymbol{W}|,$$

where $\boldsymbol{W}$ is a diagonal matrix containing the weights whose entries are the L^2 Christoffel function associated with the orthogonal polynomial family. The precise procedure is given in (3.7a), which relies on the notion of Fekete points for polynomial interpolation. The Fekete points are defined to be those points that maximize the (absolute value of the) Vandermonde determinant on the given compact set. Without the weights, the procedure is essentially the method of computing approximate Fekete points [32, 1], and so computationally this is easily implemented with pivoted linear algebraic routines. Based on these connections, we call our procedure Christoffel-weighted approximate Fekete points (CFP).

The introduction of the weights introduces mathematically nontrivial results, and our procedure results in the following theoretical and practical advantages:

- Approximate Fekete points are restricted to compact sets Γ. The CFP formulation is impartial to (non-)compactness of Γ.
- Under certain assumptions, the greedy CFP procedure generates a sequence that coincides with the result of simultaneous/global optimization (that is, maximizing the weighted determinant by varying $y_1, y_2, \ldots$ simultaneously). The required assumptions are in practice difficult to verify in multidimensional settings, but that our greedy procedure can in principle produce the same result as simultaneous optimization is a strong advantage of the method.
- For general distributions of Y in one dimension, the requisite assumptions of the previous bullet point are essentially always satisfied. Thus, in one

dimension our greedy design almost always[1] produces an optimal mesh. In this particular one-dimensional setting, the CFP algorithm produces abscissae for Gauss quadrature rules.

- Like the methodology for approximate Fekete points, the computational procedure for CFP requires essentially only pivoted linear algebra routines, in particular the QR factorization.
- Our numerical results show that CFP produces empirically superior results when compared with the deterministic sampling strategy given by approximate Fekete points and when compared against randomized Monte Carlo sampling methods. This is true for all our test cases, both in low-dimensional settings ($d = 2$) and in relatively high-dimensional settings ($d = 25$).

The paper is organized as follows. In section 2, we introduce notation and discuss least-squares problems. Our approach for grid design is introduced in section 3, along with a description of theoretical properties (with proofs provided in the appendix). Section 4 describes some details of the algorithm, and section 5 numerically investigates several examples.

1.1. Historical discussion. This section describes some previous theoretical results on sampling for polynomial least-squares. The authors in [10] provide foundational theoretical analysis for unweighted discrete least-squares using Monte Carlo sampling. For an N-term expansion in tensor-product Legendre polynomials, independent and identically distributed (iid) sampling from the uniform distribution requires $M \sim N^2$ points to guarantee the stability with high probability. With expansions in tensor-product Chebyshev polynomials, the condition can be reduced to $M \sim N^{\frac{\ln 3}{\ln 2}}$ [8]. More general results can be derived using the inequalities in [21, 20].

Weighted least-squares approaches with Monte Carlo samples have also been investigated. Analysis for general weighted procedures is given in [26], where the authors also observe that sampling from the weighted pluripotential equilibrium measure provides optimal stability and convergence estimates for approximations with asymptotically large polynomial degree. The authors in [17] propose an inexact sampling method (the Markov Chain Monte Carlo strategy) for optimal sampling in the nonasymptotic case. The results in [11] suggest an exact sampling method (sampling with the theoretical optimal measure) and show optimal convergence estimates in the nonasymptotic case. The work in [23] provides efficient computational methods for exact sampling in the nonasymptotic case.

The focus of this paper is on deterministic sampling schemes; such methods have also been investigated elsewhere [22, 37, 6]. We remark again that polynomial grids constructed via Fekete or Leja methods are closely connected to our procedure [5, 4, 1]. Another closely related approach is provided in [30], wherein the authors optimize a determinant-like objective.

2. Problem formulation. This section focuses on the introduction of notation and some background material. More details on discrete least-squares problems can be found in [31, Chaps. 10–11].

Let $y = (y_1, \ldots, y_d)^\top$ be a d-dimensional vector whose components take values in $\Gamma_i \subset \mathbb{R}$. In parametric UQ problems, each y_i corresponds to a random variable input into a system, and the goal is to understand how the system depends on these inputs. This is frequently done via a linear expansion in a basis of polynomials that is

---

[1]The initial sample for the greedy method can be any point on $\mathbb{R}$ not coinciding with $N-1$ isolated points; see Theorem 3.4 for details.

orthogonal with respect to the L^2 norm weighted by the probability density function, i.e., a polynomial chaos expansion [16, 36].

2.1. Notation. We assume that Γ is tensorial and $\rho(y)$ is a tensor-product probability density on $\Gamma \subset \mathbb{R}^d$. In parametric UQ problems, this is equivalent to assuming that the components of a random variable $Y = (Y_1, \ldots, Y_d)$ are mutually independent and that Y_i has marginal probability density $\rho_i : \Gamma_i \to [0, \infty)$. Thus $\Gamma := \prod_{i=1}^d \Gamma_i \subset \mathbb{R}^d$, and $\rho(y) = \prod_{i=1}^d \rho_i(y_i) : \Gamma \to \mathbb{R}^+$. All our results hold for variables Y_i that are discrete, or mixtures of discrete and continuous random variables, but for simplicity we assume throughout that they are continuous random variables with densities.

For each $i = 1, \ldots, d$, we can define orthogonality in terms of the ρ_i-weighted L^2 norm on Γ_i. Assuming that $0 < \mathbb{E}\,|Y_i|^k < \infty$ for all $k \in \mathbb{N}_0$, then Y_i has finite moments of all orders. This ensures existence of a family of orthogonal polynomials. We let φ_n^i denote the degree-n polynomial from the orthogonal family associated to the weight function ρ_i. Therefore,

$$\int_{\Gamma_i} \varphi_k^i(y_i)\varphi_l^i(y_i)\rho_i(y_i)dy_i = \delta_{k,l}, \qquad i = 1, \ldots, d, \ k, l = 0, 1, \ldots,$$

where $\delta_{k,l}$ is the Kronecker delta function.

Since ρ is a tensor-product weight on a tensorial domain Γ, multivariate polynomials orthogonal under ρ can be formed via tensorization: For any multi-index $\alpha \in \mathbb{N}_0^d$, the polynomials defined as

$$\psi_\alpha(y) := \prod_{j=1}^d \varphi_{\alpha_j}^j(y_j), \qquad \alpha \in \mathbb{N}_0^d,$$

satisfy

$$\langle \psi_\alpha, \psi_\beta \rangle := \int_\Gamma \psi_\alpha(y)\psi_\beta(y)\rho(y)dy = \delta_{\alpha,\beta}, \qquad \alpha, \beta \in \mathbb{N}_0^d.$$

With $\alpha = (\alpha_1, \ldots, \alpha_d) \in \mathbb{N}_0^d$ a multi-index, then $|\alpha| = \sum_{j=1}^d \alpha_j$. Various polynomial subspaces can be defined by identifying an appropriate collection of multi-indices. For example, the total degree and hyperbolic cross index sets of order k are, respectively,

$$\Lambda_k^{\mathrm{TD}} := \left\{ \alpha \ \middle| \ |\alpha| \le k \right\}, \qquad \Lambda_k^{\mathrm{HC}} := \left\{ \alpha \ \middle| \ \prod_{j=1}^d (\alpha_j + 1) \le k + 1 \right\}.$$

It will occasionally be convenient to place an ordering on indices in a finite index set. If $\Lambda \subset \mathbb{N}_0^d$ has size N, then we will assume an implicit one-to-one correspondence between the sets

$$(2.1) \qquad \{\alpha \ | \ \alpha \in \Lambda\} \longleftrightarrow \{1, \ldots, N\}.$$

We let $\alpha(j)$ denote the index corresponding to j from the above map, so that $\alpha(1), \ldots, \alpha(N)$ is an ordering of the elements of Λ. In our context, the particular choice of correspondence defined above is irrelevant.

For any index set Λ, we define the associated space of polynomials as

$$(2.2) \qquad P = P(\Lambda) = \mathrm{span}\,\{\psi_\alpha \ | \ \alpha \in \Lambda\}.$$

Via the map (2.1), the functions $\psi_j := \psi_{\alpha(j)}$, $j = 1, \ldots, N$, form an orthonormal basis for $P(\Lambda)$. We now define a set of *weighted* polynomials. With

$$K_\Lambda(y) = \sum_{\alpha \in \Lambda} \psi_\alpha^2(y),$$

$Q(\Lambda)$ is the following space of weighted polynomials:

$$(2.3) \qquad Q = Q(\Lambda) = \text{span} \left\{ \frac{\psi_\alpha}{\sqrt{K_\Lambda}} \ \Big| \ \alpha \in \Lambda \right\}.$$

Given any points $y_1, y_2, \ldots \in \Gamma$, we will use the notation

$$A_m = \{y_1, \ldots, y_m\}, \qquad\qquad m \geq 1,$$

to denote a size-m set of points. An $m \times N$ Vandermonde-like matrix for A_m on $P(\Lambda)$ using an orthonormal basis is given by

$$(2.4) \qquad \boldsymbol{V}(A_m, P) \in \mathbb{R}^{m \times N}, \qquad\qquad (V)_{j,k} := \psi_{\alpha(k)}(y_j),$$

for $1 \leq j \leq m$, and $1 \leq k \leq N$, with $\alpha(k)$ as defined in (2.1). One advantage of our using an orthonormal basis is that the Vandermonde-like matrix using *any* other orthonormal basis for $P(\Lambda)$ equals $\boldsymbol{V}(A_m, P(\Lambda)) \boldsymbol{U}$ for some orthogonal matrix $\boldsymbol{U}$.

We define $\boldsymbol{V}(A_m, Q)$ similarly, having elements $\frac{\psi_{\alpha(k)}(y_j)}{\sqrt{K_\Lambda(y_j)}}$.

2.2. Discrete least-squares problems. Given a function $f : \Gamma \to \mathbb{R}$ and multi-index set Λ, our main goal is to construct a polynomial approximation f_N from $P(\Lambda)$:

$$(2.5) \qquad f(y) \approx f_N(y) := \sum_{n=1}^{N} \widehat{f}_n \psi_n(y).$$

The $L_\rho^2(\Gamma)$-best approximation from $P(\Lambda)$ is the polynomial

$$f_N^*(y) = \sum_{n=1}^{N} \langle f, \psi_n \rangle \psi_n(y).$$

The coefficients of this polynomial clearly require significant information about the function f via the inner products $\langle f, \psi_n \rangle$. In practice such information cannot be computed directly and instead only point evaluations $f(y)$ at a discrete number y values are possible to obtain. If $A_M = \{y_1, \ldots y_M\} \subset \Gamma$ is some selection of M points, then one possible method to compute f_N is to compute the least-squares residual minimizer, which is a quadratic optimization problem whose solution is linear in the data $f(y_m)$:

$$(2.6) \qquad \left\{ \widehat{f}_n \right\}_{n=1}^{N} = \underset{f_n}{\arg\min} \sum_{m=1}^{M} \left[f(y_m) - \sum_{n=1}^{N} \widehat{f}_n \psi_n(y_m) \right]^2.$$

This problem can be written algebraically: Let $\mathbf{V}(A_M, P)$ be the $M \times N$ Vandermonde-like matrix defined in (2.4). We collect the unknown coefficients $\widehat{f}_n$ into the vector $\widehat{\mathbf{f}}$

and collect the function evaluations $f(y_m)$ into the vector $\mathbf{f} \in \mathbb{R}^M$. The least-squares approach (2.6) is equivalent to

$$(2.7) \qquad \widehat{\mathbf{f}} = \arg \min_{\mathbf{v} \in \mathbb{R}^N} \|\mathbf{V}(A_M, P)\mathbf{v} - \mathbf{f}\|_2,$$

which is a simple algebraic least-squares problem.

Recent research has shown that the unweighted least-squares formulation above is frequently inferior to a particular weighted approach [17, 26, 11]. This approach uses weights given by $1/K_\Lambda$. The algebraic formulation of the weighted approach solves

$$(2.8) \qquad \widehat{\mathbf{f}} = \arg \min_{\mathbf{v} \in \mathbb{R}^N} \|\mathbf{V}(A_M, Q)\mathbf{v} - \boldsymbol{W}\mathbf{f}\|_2,$$

where $\boldsymbol{W}$ is a diagonal matrix with entries $(W)_{m,m} = 1/\sqrt{K_\Lambda(y_m)}$, $1 \le m \le M$. This paper focuses on solving (2.8), where we use a deterministic sampling approach to compute A_M.

3. A quasi-optimal sampling strategy. Our sampling strategy relies on the notion of Fekete points for polynomial interpolation. To this end, we first review some basic definitions for the Fekete points. Throughout this section Λ is an arbitrary but fixed finite multi-index set with size N, and we use the abbreviations $P = P(\Lambda)$ and $Q = Q(\Lambda)$ as defined in (2.2) and (2.3), respectively.

The CFP method we propose in this paper is provided by the greedy optimization (3.7a), but the first three subsections below provide motivating discussion for this optimization.

3.1. Determinants and interpolation. One set of good points, in theory, for polynomial interpolation is the Fekete points. We give a brief discussion of this below, but for more details along this line we refer to [3, 4] and references therein.

Assume in this section that $\Gamma \subset \mathbb{R}^d$ is a compact set with nonempty interior. Given a set of N distinct points $\mathcal{A}_N = \{y_i\}_{i=1}^N \subset \Gamma$ and a function $f : \Gamma \to \mathbb{R}$, the polynomial interpolation problem is to find a $p \in P_\Lambda(\Gamma)$ such that

$$(3.1) \qquad p(y_i) = f(y_i) \quad \forall y_i \in \mathcal{A}_N.$$

We assume that this problem is unisolvent; this is true unless the y_i have a pathological configuration in Γ.[2] With $\{\psi_1, \psi_2, \ldots \psi_N\}$ any ordered basis for $P(\Lambda)$, then there are unique coefficients c_j satisfying a linear system that determines p:

$$(3.2) \qquad p = \sum_{j=1}^N c_j \psi_j, \qquad\qquad \mathbf{V}(A_N; P_\Lambda)\mathbf{c} = \mathbf{f}.$$

DEFINITION 3.1. *Let y_j, $j = 1, \ldots, N$, denote any set of points in Γ, so that $A_m = \{y_j\}_{j=1}^m$ for $m \le N$ is well-defined. Define the determinant modulus of the rectangular matrix $\boldsymbol{V} = \boldsymbol{V}(A_m, P)$ as*

$$|\det \boldsymbol{V}| = \sqrt{|\det(\boldsymbol{V}\boldsymbol{V}^T)|}, \qquad\qquad 1 \le m \le N.$$

---

[2]For example, if we choose y_i as realizations of a continuous random variable that is uniform on Γ, then the interpolation problem is unisolvent with probability 1.

The definition above coincides with the standard square matrix determinant (modulus) when $m = N$. Some additional notation is the operation of appending a point to a given set A_m and replacing the jth element of A_N, respectively:

$$A_{m \cup y} = \{y_1, \ldots, y_m, y\}, \qquad\qquad y \in \Gamma.$$
$$A_{N \setminus j(y)} := \{y_1, \ldots, y_{j-1}, y, y_{j+1}, \ldots, y_N\}, \qquad 1 \le j \le N, \ y \in \Gamma.$$

If the polynomial interpolation problem on P_Λ is on $\mathcal{A}_N$, then the cardinal Lagrange interpolation polynomials are given by

$$\ell_j(y) = \frac{\det \boldsymbol{V}\left(A_{N \setminus j(y)}, P_\Lambda\right)}{\det \boldsymbol{V}\left(A_N, P_\Lambda\right)}, \qquad\qquad \ell_j(y_k) = \delta_{j,k},$$

with $\delta_{j,k}$ the Kronecker delta. This allows us to explicitly construct the unique element $p \in P(\Lambda)$ that interpolates at A_N an arbitrary f continuous on A_N:

$$I_N f := \sum_{n=1}^N f(y_n) \ell_n(y).$$

The output of this operator is an element of $P(\Lambda)$, which we can view as a subspace of $C(\Gamma)$, continuous functions over Γ. This results in a popular notion of stability, the Lebesgue constant,

$$\|I_N\|_{C(\Gamma) \to C(\Gamma)} = \sup_{y \in \Gamma} \sum_{n=1}^N |\ell_n(y)|.$$

This quantity does not depend on which basis for $P(\Lambda)$ is chosen to compute the cardinal Lagrange interpolants. Finally, the conditioning of the problem of computing $\boldsymbol{c}$ from an arbitrary $\boldsymbol{f}$ is measured by

$$\kappa\left(\boldsymbol{V}\left(A_N, P\right)\right) = \frac{\sigma_1\left(\boldsymbol{V}\right)}{\sigma_N\left(\boldsymbol{V}\right)},$$

where σ_j, $j = 1, \ldots, N$, are the singular values of $\boldsymbol{V}$ in decreasing order.

3.2. Near-optimal stability and conditioning. Consider the square systems case, $M = N$, with interpolation on $P = P(\Lambda)$. A set of points that maximizes the determinant of the Vandermonde-like matrix is called a set of Fekete points:

$$(3.3) \qquad A_N^F(P(\Lambda)) := \operatorname*{arg\,max}_{A_N = \{y_1, \ldots, y_N\} \in \Gamma^N} |\det \boldsymbol{V}\left(A_N, P\right)|.$$

Classically, Fekete points on general manifolds are point configurations that minimize a Riesz energy. On a compact interval in $\mathbb{R}$, a specialization of Riesz energy coincides with the determinant of the Vandermonde matrix. Thus, in the one-dimensional setting a set of Fekete points is determined by maximizing the determinant of the Vandermonde matrix.

The utility of Fekete points for polynomial approximation is that they provide at most linear growth of the Lebesgue constant:

$$\|I_N\|_{C(\Gamma) \to C(\Gamma)} = \sup_{y \in \Gamma} \sum_{n=1}^N \left| \frac{\det \boldsymbol{V}\left(A_{N \setminus n(y)}^F, P\right)}{\det \boldsymbol{V}\left(A_N^F, P\right)} \right| \le \sup_{y \in \Gamma} \sum_{n=1}^N 1 = N.$$

In practice, logarithmic growth is observed. Thus, the computation of determinant-maximizing sample points (3.3) is of great interest.

We can similarly define an optimization problem that seeks a point configuration with minimal condition number [39]:

$$(3.4) \qquad A_N^C(P) := \underset{A_N=\{y_1,\ldots,y_N\}\in\Gamma^N}{\arg\min} \kappa\left(\boldsymbol{V}\left(A_N, P\right)\right).$$

Note that A_N^F and A_N^C are different sets in general.

3.3. Greedy designs. The optimization problems (3.3) and (3.4) are not computationally feasible in general, and so one frequently resorts to greedy algorithms. Greedy versions of these algorithms are straightforward to devise:

$$(3.5a) \qquad A_N^{F*}(P) = \left\{y_1^{F*},\ldots,y_N^{F*}\right\}, \qquad y_{n+1}^{F*} = \underset{y\in\Gamma}{\arg\max}\left|\det\boldsymbol{V}\left(A_{n\cup y}^{F*}, P\right)\right|,$$

$$(3.5b) \qquad A_N^{C*}(P) = \left\{y_1^{C*},\ldots,y_N^{C*}\right\}, \qquad y_{n+1}^{C*} = \underset{y\in\Gamma}{\arg\min}\kappa\left(\boldsymbol{V}\left(A_{n\cup y}^{C*}, P\right)\right).$$

These greedy versions are still difficult but are more feasible since they involve only repeated optimization over Γ (instead of optimization over Γ^N). The determinant-maximizing objective is easier to compute compared to the condition number objective. In practice, one often replaces exact maximization over Γ with maximization over a discrete set. Finally, there is ambiguity at each iteration if multiple locations y maximize objectives, and there is freedom in choosing the starting point y_1 in each case.

There are two major difficulties with all of our previous discussions. First, we now have four potential sets, A^F, A^C, A^{F*}, and A^{C*}, that we would like to compute. It seems unclear, for example, whether A^{F*} or A^{C*} is the better option. Our second difficulty is that none of these sets is well-defined if Γ is not compact.

3.4. Weighted greedy designs. We can partially resolve the difficulties identified at the end of the previous section by considering weighted polynomials. By doing this, we show under some assumptions that greedy designs can produce the same result as the much more burdensome simultaneous optimization designs. In addition, we show in one dimension for any ρ that this almost always happens.

We reformulate all four problems, both the optimal versions (3.3) and (3.4) as well as their greedy versions (3.5). Let $\Gamma\subset\mathbb{R}^d$, $d\geq 1$, and now assume only that Γ has an interior containing any open set with positive Lebesgue measure. Instead of working on the polynomial space $P=P(\Lambda)$, we'll use the weighted space $Q=Q(\Lambda)$ defined in (2.3). The point configurations that maximize the determinant and minimize the condition number, respectively, are defined as

$$(3.6a) \qquad A_N^F(Q) := \underset{A_N=\{y_1,\ldots,y_N\}\in\Gamma^N}{\arg\max} \left|\det\boldsymbol{V}\left(A_N, Q\right)\right|,$$

$$(3.6b) \qquad A_N^C(Q) := \underset{A_N=\{y_1,\ldots,y_N\}\in\Gamma^N}{\arg\min} \kappa\left(\boldsymbol{V}\left(A_N, Q\right)\right),$$

and the greedy versions are, for $n\geq 1$,

$$(3.7a) \qquad A_N^{F*}(Q) = \left\{y_1^{F*},\ldots,y_N^{F*}\right\}, \qquad y_{n+1}^{F*} = \underset{y\in\Gamma}{\arg\max}\left|\det\boldsymbol{V}\left(A_{n\cup y}^{F*}, Q\right)\right|,$$

$$(3.7b) \qquad A_N^{C*}(Q) = \left\{y_1^{C*},\ldots,y_N^{C*}\right\}, \qquad y_{n+1}^{C*} = \underset{y\in\Gamma}{\arg\min}\kappa\left(\boldsymbol{V}\left(A_{n\cup y}^{C*}, Q\right)\right).$$

In the above, the starting values y_1^{F*} and y_1^{C*} can take arbitrary values in $\mathbb{R}$, but this choice affects the final result of the greedy pursuit. Furthermore, at each n there

may be multiple points y_{n+1}^{F*} that extremize the objective. We assume that any one of these extremizers is chosen for the procedure, and we refer below to this potential nonuniqueness of the sequence as a *branch* of the iterative optimization.

The greedy algorithms working on Q are just as computationally feasible as those working on P. However, the advantage of this particular weighted approach is that if an optimal solution exists, then all of the four approaches (3.6) and (3.7) give the optimal solution.

THEOREM 3.2. *Let $\rho : \Gamma \to [0, \infty)$ be a probability density function on $\mathbb{R}^d$. Let Λ be an arbitrary multi-index set of size N defining Q. A point configuration A_N satisfies*

$$|\det \boldsymbol{V}\left(A_N, Q\right)| = 1 \tag{3.8a}$$

if and only if

$$\kappa\left(\boldsymbol{V}\left(A_N, Q\right)\right) = 1. \tag{3.8b}$$

Thus, solutions to (3.6) attaining optimal objective values coincide. If A_N is a set that satisfies either (hence both) of the optimal objective values above, then
- *if $y_1^{F*} \in A_N$, then the iteration (3.7a) has a branch such that $A_N^{F*} = A_N$;*
- *if $y_1^{C*} \in A_N$, then the iteration (3.7b) has a branch such that $A_N^{C*} = A_N$.*

Proof. See Appendix A. □

The strength of this result is twofold. First, it suggests that we may optimize either the determinant or the condition number and obtain equivalent answers. Second, it shows that greedy optimization recovers the global optimum. These conclusions give us great flexibility in computational procedures since we may propose a method for a greedy determinant maximization and this plausibly gives results comparable to global minimization of the condition number.

The unfortunate caveat in the result above is that we require existence of a point configuration with optimal condition number and determinant. It is initially unclear whether this is a reasonable assumption. However, nontrivial multidimensional examples for when this condition is satisfied exist [18].

A somewhat surprising positive result is that in one dimension ($d = 1$), infinitely many point configurations with optimal condition number exist, and the union of all these optimal sets covers every real number, except for $N - 1$ isolated points.

LEMMA 3.3. *Let $\rho(y)$ be any probability density function on $\Gamma = \mathbb{R}$, and let $\Lambda = \{0, \ldots, N-1\}$ for any $N \geq 1$. Recall that $\varphi_N(\cdot)$ denotes the degree-N orthonormal polynomial with respect to the ρ-weighted L^2 inner product on Γ. We use $\varphi_{N-1}^{-1}(0)$ to denote the zero set of the polynomial $\varphi_{N-1}(y)$, which is always a set of $N-1$ distinct points in $\mathbb{R}$. Then*
1. *for any $y \notin \varphi_{N-1}^{-1}(0)$, there is a set $A_N = A_N(y)$ that satisfies (3.8);*
2. *the set $A_N(y)$ defined above is unique as a function of y;*
3. *the set $A_N(y)$ is given by*

$$A_N(y) = r_N^{-1}\left(r_N(y)\right),$$

where r_N is the meromorphic function

$$r_N(y) = \frac{\varphi_N(y)}{\varphi_{N-1}(y)},$$

and r_N^{-1} denotes its set-valued functional inverse;

4. *the set $A_N(y)$ contains the abscissae for an N-point positive quadrature rule exact for polynomials up to degree $2N - 2$,*

$$\int_\Gamma p(z)\rho(z)dz = \sum_{z \in A_N(y)} \frac{1}{K_\Lambda(z)} p(z), \qquad \deg p \le 2N - 2;$$

5. *if $y \in \varphi_N^{-1}(0)$, then $A_N(y)$ are the abscissae of the N-point Gauss quadrature rule.*

Proof. See Appendix B. □

Note that in the lemma above we must consider ρ a density on $\Gamma = \mathbb{R}$, even if its support lies on a compact set. This result shows that in one dimension (for *any* ρ) there are many optimal point configurations.

THEOREM 3.4. *Let ρ be any probability density function on $\Gamma = \mathbb{R}$, and let $\Lambda = \{0, \dots, N - 1\}$ for any $N \ge 1$. Let $y \in \varphi_{N-1}^{-1}(0)$ be arbitrary but fixed, and let $A_N(y)$ denote the unique set defined in Lemma 3.3. Set $A_N^F(Q) = A_N^C(Q) = A_N(y)$, which satisfy the maximization problems* (3.6). *Set the initial values for the greedy iterations* (3.7) *as*

$$y_1^{F*} = y_1^{C*} = y.$$

Then

$$A_N^C(Q) = A_N^{C*}(Q) = A_N^{F*}(Q) = A_N^F(Q)$$

and

$$\left|\det \boldsymbol{V}\left(A_N^F, Q\right)\right| = \left|\det \boldsymbol{V}\left(A_N^{F*}, Q\right)\right| = 1 = \kappa\left(\boldsymbol{V}\left(A_N^{C*}, Q\right)\right) = \kappa\left(\boldsymbol{V}\left(A_N^C, Q\right)\right).$$

Proof. See Appendix C. □

We emphasize that the above result holds for *any* univariate density ρ, even those with noncompact support. This result completely characterizes the greedy scheme's behavior in one dimension and shows that it achieves an optimal condition number for almost any starting point y.

One final observation is that the set A_N^{F*} produced by the one-dimensional greedy iteration produces optimal quadrature rules.

COROLLARY 3.5. *Let $y \notin \varphi_{N-1}^{-1}(0)$. Then the greedy iteration* (3.7a) *with $y_1^{F*} = y$ produces the unique positive L_ρ^2 quadrature rule abscissae with optimal polynomial accuracy. In particular, if $y \in \varphi_N^{-1}(0)$, then they produce the abscissae of the ρ-weighted Gauss quadrature rule.*

The theoretical results of this section give strong motivation for using the greedy weighted determinant design $A_N^{F*}(Q)$: in one dimension we produce optimally conditioned point sets for almost any starting location in the greedy design. The one-dimensional greedy designs coincide with the more onerous simultaneous optimization designs and even coincide with designs based on condition number optimization. For multiple dimensions we retain all the previous properties but must make the assumption that a point set with unit condition number exists and that our starting location lies in this set. Under this existence assumption, greedy designs again produce optimal sets.

The remainder of this paper investigates the computational performance of the set $A_N^{F*}(Q)$.

4. CFP algorithmic details. The CFP strategy given by (3.7a) can be used to construct a sample set A having the same size as the dimension of the (weighted) polynomial space Q. However, in least-squares problems we wish the sample count $M = |A|$ to dominate the polynomial space dimension $N = \dim Q$. To achieve this, we start with a specified space Q and enrich it with weighted polynomials of a higher degree. This procedure is largely ad hoc, so we cannot claim optimality for our specific strategy, but our numerical results indicate that our procedure works very well.

Let Λ be a given multi-index set of size N. All our examples will use a downward-closed set Λ, but this is mainly for algorithmic convenience. An index set Λ is called a downward-closed set if and only if $\boldsymbol{\nu} \in \Lambda$ and $\boldsymbol{\mu} \leq \boldsymbol{\nu} \implies \boldsymbol{\mu} \in \Lambda$, where $\boldsymbol{\mu} \leq \boldsymbol{\nu}$ if and only if $\mu_k \leq \nu_k$ for all $1 \leq k \leq d$. The weighted polynomial space Q is defined in (2.3). We wish to compute a set of samples $A = \{y_1, \ldots, y_M\}$ that we use to solve the weighted least-squares problem (2.8). Our procedure to accomplish this enriches Λ to a size of $M > N$, and we subsequently compute M CFP points associated to this enriched Λ. We then compute the least-squares solution (2.8) using the original index set Λ. We describe the details of this below.

4.1. Choosing Q. Given an enrichment size ΔN, we define $\widetilde{\Lambda}$ of size $N + \Delta N$ and satisfy $\widetilde{\Lambda} \supset \Lambda$ by adding to Λ elements based on total-degree graded reverse lexicographic ordering of the set $\Lambda_n \backslash \Lambda$:

1. Compute $n = \max\{|\alpha| \mid \alpha \in \Lambda\}$.
2. Compute $S := \Lambda_n^{\mathrm{TD}} \backslash \Lambda$. Set $n \leftarrow n + 1$ if $\Lambda = \Lambda_n^{\mathrm{TD}}$.
3. Impose a total order on S: graded (partially ordered) based on total degree, and ordering within a grade defined by reverse lexicographic ordering.
4. Extract the first ΔN elements from S and append those to Λ, creating the set $\widetilde{\Lambda}$.

The input to the procedure above is simply Λ (assumed downward-closed) and an enrichment size ΔN. In all our numerical tests, we choose $\Delta N = \lfloor 0.05N \rfloor$, representing a small enrichment of Λ. We now define the $(N + \Delta N)$-dimensional space

$$\widetilde{Q} := Q_{\widetilde{\Lambda}}.$$

We will compute CFP using this enriched weighted space.

4.2. Choosing candidate sets. CFP sets are computed via (3.7a). This procedure must be discretized in practice since it is computationally difficult to optimize over the continuum Γ. We choose a large but finite-size candidate set $\widetilde{A}$ and replace the maximization over Γ by maximization over $\widetilde{A}$.

When d is "moderately small," say, $d \leq 5$, modern computational power allows us to choose a candidate mesh that "fills" this d-dimensional space. Thus, we can be reasonably sure that an intelligent choice for $\widetilde{A}$ is an effective surrogate for Γ. However, when $d \gtrsim 5$ we can no longer be reasonably confident that $\widetilde{A}$ is a fine enough mesh on Γ. In this paper we do not make any advancements with respect to this deficiency. In our numerical simulations we choose $\widetilde{A}$ as the union of two random ensembles:

$$\widetilde{A} = \left\{ R_1, \ldots, R_{\widetilde{M}/2}, S_1, \ldots, S_{\widetilde{M}/2} \right\},$$

where $\widetilde{M}$ is a reasonably large number, usually $\widetilde{M} = 10^4$. The random variables R_j, $j = 1, \ldots, \widetilde{M}/2$, are iid samples from the probability density ρ on Γ, and S_j, $j = 1, \ldots, \widetilde{M}/2$, are iid samples from a degree-asymptotic density inspired by the approach in [26].

When ρ is the uniform measure on $\Gamma = [-1,1]^d$ we choose S_j to be sampled from the tensor-product Chebyshev density, and when ρ is the Gaussian meaure on $\mathbb{R}^d$ we choose the sampling distribution of S_j to have support on the $\mathbb{R}^d$ unit ball with radius $\sqrt{2n}$, with n as chosen in section 4.1. Its density as a function of $s \in \mathbb{R}^d$ is

$$C_d \left(1 - \frac{1}{2n}\sum_{k=1}^d s_k^2\right)^{d/2}, \qquad s = (s_1,\ldots,s_d) \in \mathbb{R}^d,$$

where C_d is a normalization constant. We refer to [26] for details on this latter sampling density. Our choices above may naturally be replaced with any other candidate mesh, e.g., quasi–Monte Carlo sets, sparse grids, or tensor-product points.

We finally note that weakly admissible meshes [7] are known to be good candidate sets for *non*-weighted polynomial spaces. These meshes are good candidates for our weighted formulation as well, but we forgo their use since known constructions of these meshes exhibit very large growth for even moderate dimensions [1].

4.3. Greedy optimization. We have chosen a space $\widetilde{Q}$ of dimension $M = N + \Delta N$ and a candidate mesh $\widetilde{A}$ of size $\widetilde{M} \gg N$. Our goal is now to compute a CFP set:

$$\text{Solve (3.7a), setting } Q \leftarrow \widetilde{Q} \text{ and } \Gamma \leftarrow \widetilde{A}.$$

This optimization procedure at each iteration is equivalent to forming the matrix $\boldsymbol{V}(\widetilde{\mathcal{A}},\widetilde{Q})$ and then choosing rows that greedily maximize the spanned volume of the chosen rows. This, in turn, is easily performed by a column-pivoted QR decomposition of $\boldsymbol{V}^T$; see, e.g., [32, 2]. The ordered pivots define the choice of M points $A_M = A_M^{F*}$.

We now summarize the procedure to generate the weighted Fekete points in Algorithm 1.

Algorithm 1. A sequential algorithm to construct the weighted Fekete points.
- **Input:** A given multi-index set Λ of size N, the weighted polynomial space Q is defined in (2.3), an enrichment size ΔN, a reasonable large number $\widetilde{M}$.
- **Output:** $M = N + \Delta N$ CFP points.
1. Enrich Λ to $\widetilde{\Lambda}$ of size $N + \Delta N$ via the procedure given in section 4.1, then define the $(N+\Delta N)$-dimensional space $\widetilde{Q} := Q_{\widetilde{\Lambda}}$.
2. Generate a large candidates set $\widetilde{A}$ of size $\widetilde{M}$ following the descriptions in section 4.2.
3. Form the design matrix $\boldsymbol{V}(\widetilde{\mathcal{A}},\widetilde{Q})$, used for computing the CFP points via (3.7a).
4. Perform the QR-type decomposition with column pivoting to the matrix $\boldsymbol{V}^T$ to get M CFP points $A_M = A_M^{F*}$.

4.4. Least-squares solve. Having selected a size-M point set A_M in section 4.3, and with a size-N polynomial space $P(\Lambda)$ already defined (along with its weighted version Q), we now compute $\widehat{\boldsymbol{f}}$ as the solution to (2.8).

The coefficients in the solution vector $\widehat{\boldsymbol{f}}$ define our desired expansion shown in (2.5).

5. Numerical tests. In this section we investigate the stability and convergence properties of the CFP sampling strategy. We will generate sampling sets A_M using CFP and compare them against two popular alternatives for generating A_M: randomized Monte Carlo using iid samples from ρ, and deterministic approximate Fekete points. We are interested primarily in investigating how sampling rates of M versus the approximation space dimension N affects stable and accurate approximation. In our figures and results, we will use "MC" to denote the Monte Carlo procedure, "Fekete" to denote approximate Fekete points, and "C-Fekete" to denote the CFP algorithm of this paper.

In order to implement our proposed method, we first picked $\widetilde{M}$ random points as a candidate set, then we selected M optimal points from section 1 as CFP points. In all the following examples, we choose $\widetilde{M} = 10^4$.

5.1. Matrix stability. In this section we investigate the condition number of the matrix $\boldsymbol{V}(A_M, P)$ (MC and Fekete) and the condition number for the matrix $\boldsymbol{V}(A_M, Q)$ (C-Fekete). In all examples that follow we perform 50 trials of each procedure and report the mean condition number along with 20% and 80% quantiles.

5.1.1. Bounded domains. We first consider the Legendre polynomials for which the domain is $\Gamma = [-1, 1]^d$ and ρ is the uniform probability density. In Figure 1, the $d = 2$ condition numbers obtained by the procedures are shown for a linear over-sampling of $M = 1.05N$. The numerical results for $d = 6$ and $d = 10$ are shown in Figures 2 and 3, respectively. We notice that the CFP procedure produces point configurations that have notably more stable linear systems when compared against the approximate Fekete points (AFP) or MC designs.

5.1.2. Unbounded domains. We now let $\Gamma = \mathbb{R}^d$ with ρ a Gaussian density, $\rho(y) \propto \exp(-\|y\|_2^2)$. The associated orthonormal polynomial family is formed from tensor-product Hermite polynomials. Our tabulation of the condition numbers is shown in Figures 4–7. We note that for each case, dimension $d = 2, 6, 10$, and 25, the CFP procedure produces more stable point sets than either of the alternatives, but the improvement is modest in high dimensions.

5.2. Least-squares accuracy. In this section, we will compare the CFP, APF, and MC algorithms in terms of their ability to approximate test functions. In all examples that follow we report the numerical results over an ensemble of 50 tests.
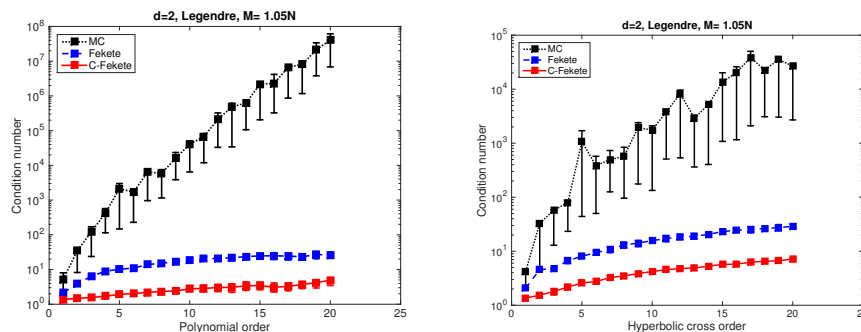

FIG. 1. *Condition number with respect to the polynomial degree in the* 2-*dimensional polynomial spaces. Left: Total degree. Right: Hyperbolic cross.*
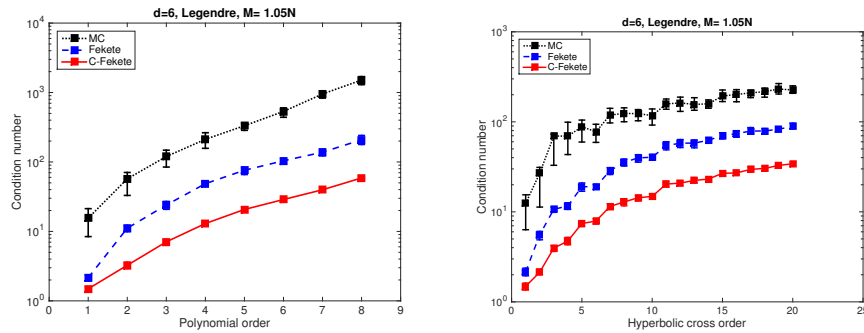
FIG. 2. *Condition number with respect to the polynomial degree in the* 6-*dimensional polynomial spaces. Left: Total degree. Right: Hyperbolic cross.*
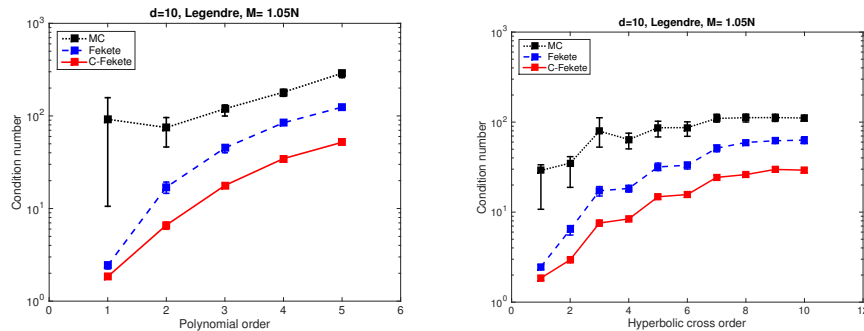

FIG. 3. *Condition number with respect to the polynomial degree in the* 10-*dimensional polynomial spaces. Left: Total degree. Right: Hyperbolic cross.*
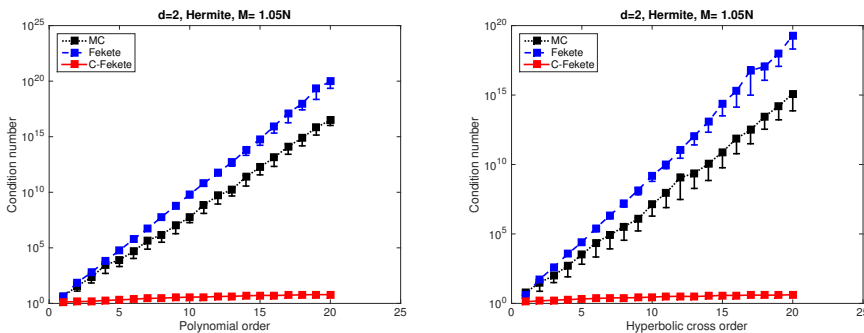

FIG. 4. *Condition number with respect to the polynomial degree in the* 2-*dimensional polynomial spaces. Left: Total degree. Right: Hyperbolic cross.*

5.2.1. Algebraic function. In Figure 8, we show the convergence rate of the least-squares approximation for Legendre approximation (ρ uniform over $[-1, 1]^d$) in the two-dimensional polynomial space for the test function $f(y) = \exp(-\sum_{j=1}^{d} y_j^2)$. We measure accuracy using the discrete ℓ_2 norm which is computed using $1{,}000$ random samples drawn from the probability measure of orthogonality. In this case the CFP and AFP procedures work comparably.
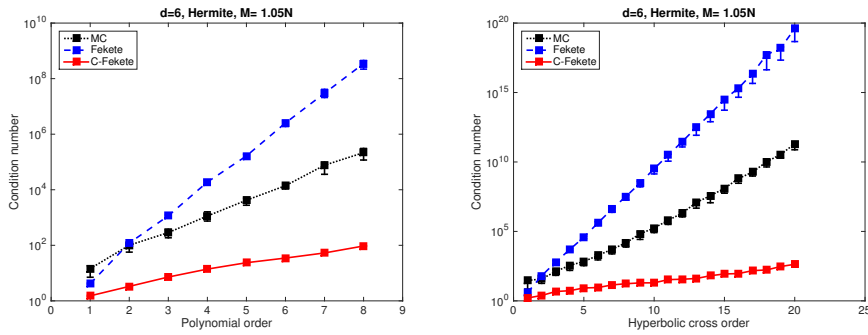
FIG. 5. *Condition number with respect to the polynomial degree in the 6-dimensional polynomial spaces. Left: Total degree. Right: Hyperbolic cross.*
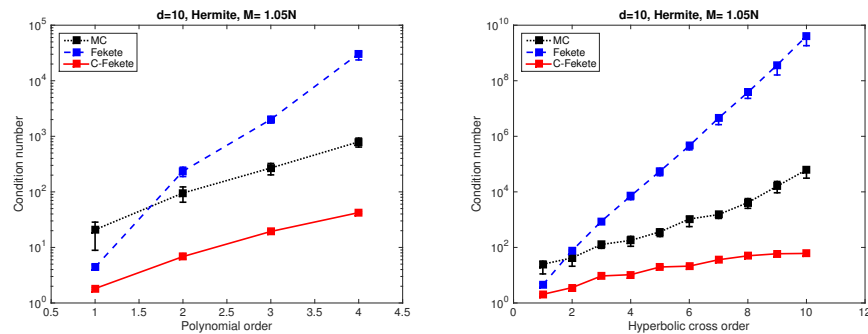

FIG. 6. *Condition number with respect to the polynomial degree in the 6-dimensional polynomial spaces. Left: Total degree. Right: Hyperbolic cross.*
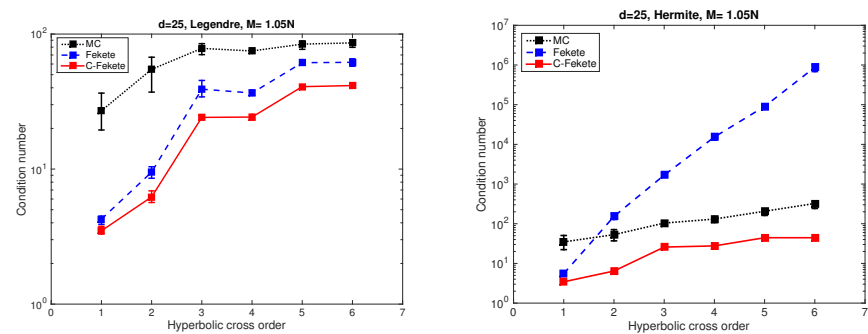

FIG. 7. *Condition number with respect to the polynomial degree in the 25-dimensional hyperbolic cross polynomial spaces. Left: Legendre. Right: Hermite.*

In Figure 9, we consider the Hermite approximation for $f(y)$ in the two-dimensional polynomial space. We observe here that CFP produces considerably better results compared with MC or AFP, especially for high-degree approximations.

5.2.2. Parameterized elliptic differential equation. We now consider the stochastic elliptic equation, one of the most used benchmark problems in UQ, in one spatial dimension,
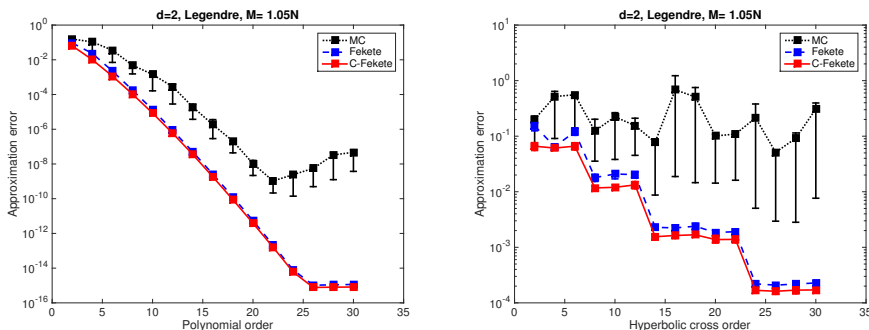
FIG. 8. *Approximation error against polynomial degree. Legendre approximation of* $f(\mathbf{Y}) = \exp(-\sum_{i=1}^{d} Y_i^2)$ *Left: Total degree. Right: Hyperbolic cross.*
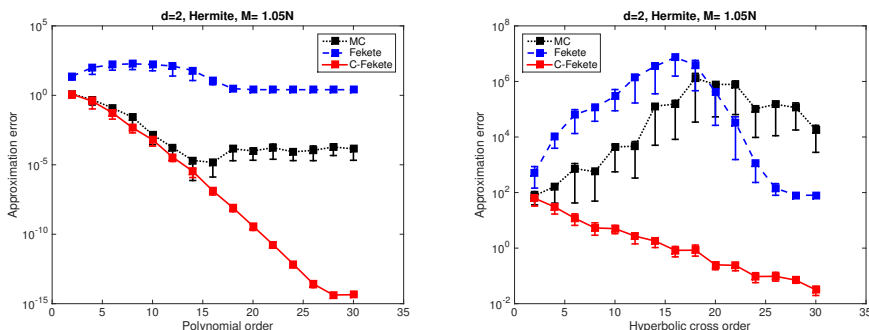

FIG. 9. *Approximation error against polynomial degree. Hermite approximation of* $f(\mathbf{Y}) = \exp(-\sum_{i=1}^{d} Y_i)$ *Left: Total degree. Right: Hyperbolic cross.*

$$(5.1) \qquad -\frac{d}{dx}[\kappa(x,y)\frac{du}{dx}(x,y)] = f, \quad (x,y) \in (0,1) \times \mathbb{R}^d,$$

with boundary conditions

$$u(0,y) = 0, \quad u(1,y) = 0,$$

and $f = 2$. The random diffusivity takes the following form:

$$(5.2) \qquad \kappa(x,y) = 1 + \sigma \sum_{k=1}^{d} \frac{1}{k^2\pi^2}\cos(2\pi k x)y_k.$$

This form resembles that of the well-known Karhunen–Loeve expansion.

We approximate the solution $u(y) = u(0.5,y)$ and let the density $\rho(y)$ be uniform over the hypercube $[-1,1]^d$ and thus use Legendre polynomials in y to approximate $u(y)$. Figures 10 and 11 compare the convergence accuracy of the least-squares approximations of the quantity of interest $u(0.5,y)$ using CFP, AFP, and MC algorithms. CFP performs comparably to, but no worse than, AFP and much better than MC.

Appendix A. Proof of Theorem 3.2. We first show the equivalence of the relations (3.8). Recall the fact that for a matrix $\boldsymbol{A} \in \mathbb{R}^{N \times N}$ having columns $\boldsymbol{a}_i \in \mathbb{R}^N$, $i = 1, \dots, N$,
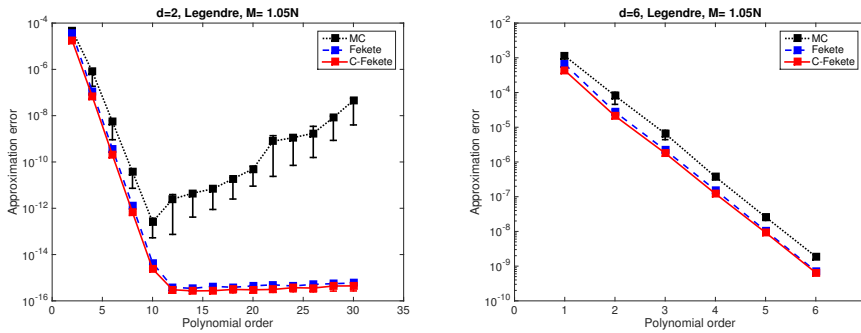
FIG. 10. *Approximation error against polynomial degree. Legendre approximation of the diffusion equation. Left: $d = 2$. Right: $d = 6$.*
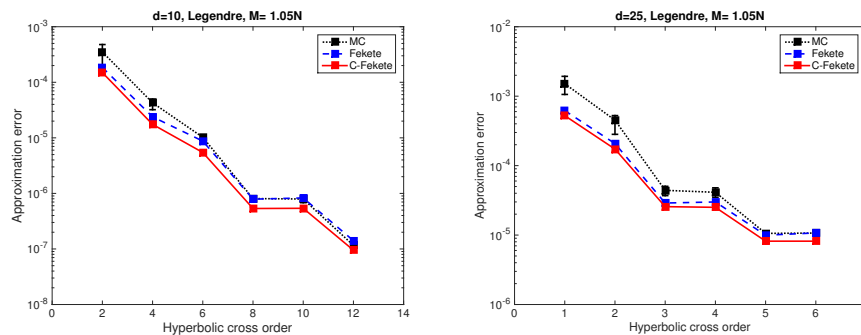

FIG. 11. *Approximation error against polynomial degree. Legendre approximation of the diffusion equation. Left: $d = 10$. Right: $d = 25$.*

$$(A.1) \qquad |\det \boldsymbol{A}| \leq \prod_{i=1}^{N} \|\boldsymbol{a}_i\|_2$$

with equality if and only if $\boldsymbol{a}_i$ are pairwise orthogonal. Now assume that (3.8a) holds with $A_N = \{y_1, \ldots, y_N\}$. The matrix $\boldsymbol{V}(A_N, Q)$ is made up of rows

$$(A.2) \qquad \boldsymbol{V} = \begin{pmatrix} \boldsymbol{\psi}^T(y_1) \\ \boldsymbol{\psi}^T(y_2) \\ \vdots \\ \boldsymbol{\psi}^T(y_N) \end{pmatrix}, \qquad \boldsymbol{\psi}^T(y_j) = \left(\frac{\psi_1(y_j)}{K_\Lambda(y_j)}, \ \frac{\psi_2(y_j)}{K_\Lambda(y_j)}, \ldots, \frac{\psi_N(y_j)}{K_\Lambda(y_j)} \right).$$

We have

$$(A.3) \qquad \|\boldsymbol{\psi}(y_j)\|_2 = \sqrt{\sum_{i=1}^{N} \frac{\psi_i^2(y_j)}{K_\Lambda}} = 1$$

for all j. Thus,

$$\left| \det \boldsymbol{V}^T \right| = \left| \det \boldsymbol{V} \right| \overset{(3.8a)}{=} 1 = \prod_{j=1}^{N} \|\boldsymbol{\psi}(y_j)\|_2 \, ,$$

thus showing equality in (A.1) for $\boldsymbol{A} = \boldsymbol{V}^T$. Therefore, the $\boldsymbol{\psi}(y_j)$ are pairwise orthogonal and are furthermore orthonormal because of (A.3). A square matrix $\boldsymbol{V}$ with pairwise orthonormal rows is an orthogonal matrix and hence has all of its singular values equal to 1. Therefore, $\kappa(\boldsymbol{V}) = 1$, showing (3.8b).

Now assume that (3.8b) holds. Then $\boldsymbol{V}$, being a square matrix, is an orthogonal matrix. Hence,

$$|\det \boldsymbol{V}|^2 = \left|\det \boldsymbol{V}^T \boldsymbol{V}\right| = |\det \boldsymbol{I}| = 1,$$

showing that $|\det \boldsymbol{V}| = 1$, and hence (3.8a). This completes the proof of the equivalence of relations (3.8).

Now let $A_N = \{y_1, \ldots, y_N\}$ be a set that satisfies either (hence, both) of the conditions (3.8), and choose $y_1^{F*} \in A_N$; we further assume without loss that $y_1^{F*} = y_1$. (If not, then relabel the elements in A_N.) We proceed by induction, showing that for each $n = 1, \ldots, N-1$, the greedy optimization (3.7a) has a branch yielding the new point $y_{n+1}^{F*} = y_{n+1}$. Note that the initialization step of induction is true by assumption. We will show the induction step is true by noting that (3.7a) varies a single $\mathbb{R}^N$ vector in an attempt to maximize an $(n+1)$-dimensional volume spanned by unit vectors. One way to maximize this volume is by choosing the vectors to correspond to the set A_N, which makes the volume an $(n+1)$-dimensional orthotope. The details are as follows.

Since conditions (3.8) are true, the vectors $\boldsymbol{\psi}(y_j)$, $j = 1, \ldots, N$, are pairwise orthonormal. For some $n \geq 1$, assume $y_j^{F*} = y_j$ for $1 \leq j \leq n$. The iteration (3.7a) at index $n \geq 1$ takes the form

$$y_{n+1}^{F*} = \arg\max_{y \in \Gamma} \left|\det \boldsymbol{V}\left(A_{n\cup y}^{F*}, Q\right) \boldsymbol{V}^T \left(A_{n\cup y}^{F*}, Q\right)\right|,$$

where

$$\boldsymbol{V}\left(A_{n\cup y}^{F*}, Q\right) = \begin{pmatrix} \boldsymbol{\psi}^T(y_1^{F*}) \\ \vdots \\ \boldsymbol{\psi}^T(y_n^{F*}) \\ \boldsymbol{\psi}^T(y) \end{pmatrix}.$$

Since $y_j^{F*} = y_j$ for $1 \leq j \leq n$, then $\boldsymbol{\psi}^T(y_j^{F*})$ for $1 \leq j \leq n$ are pairwise orthonormal. This allows us to easily compute the QR decomposition of $\boldsymbol{V}^T$:

$$\boldsymbol{V}^T = \boldsymbol{QR} = \begin{pmatrix} \boldsymbol{\psi}(y_1^{F*}) & \boldsymbol{\psi}(y_2^{F*}) & \cdots & \boldsymbol{\psi}(y_n^{F*}) & \widetilde{\boldsymbol{\psi}}(y) \end{pmatrix} \begin{pmatrix} \boldsymbol{I}_n & \boldsymbol{b} \\ \boldsymbol{0} & r \end{pmatrix},$$

where $\boldsymbol{I}_n$ is the $n \times n$ identity matrix, the vector $\boldsymbol{b} \in \mathbb{R}^n$ has entries

$$(b)_j = \boldsymbol{\psi}^T\left(y_j^{F*}\right) \boldsymbol{\psi}(y),$$

and $r = \sqrt{1 - \|\boldsymbol{b}\|_2^2}$. Note by Bessel's inequality that $\|\boldsymbol{b}\|_2 \leq 1$, and hence $r \in [0, 1]$. When $r > 0$, the last column of $\boldsymbol{Q}$ is

$$\widetilde{\boldsymbol{\psi}}(y) = \frac{1}{r}\left[\boldsymbol{\psi}(y) - \sum_{j=1}^{n} b_j \boldsymbol{\psi}(y_j^{F*})\right],$$

and when $r = 0$ we let $\widetilde{\boldsymbol{\psi}}(y)$ be any unit vector. Since the matrix $\boldsymbol{Q}$ is orthogonal, we have

$$\left| \det \boldsymbol{V} \left(A_{n \cup y}^{F*}, Q \right) \boldsymbol{V}^T \left(A_{n \cup y}^{F*}, Q \right) \right| = \left| \det(\boldsymbol{R}^T \boldsymbol{Q}^T \boldsymbol{Q} \boldsymbol{R}^T) \right| = \left| \det(\boldsymbol{R}^T \boldsymbol{R}) \right| = r^2.$$

This is maximized when $r = 1$, which happens when $\boldsymbol{\psi}(y)$ is orthonormal to $\boldsymbol{\psi}(y_j^{F*})$ for $1 \le j \le n$. This is achievable by setting $y \leftarrow y_{n+1}$ since $\boldsymbol{\psi}(y_{n+1})$ is a vector satisfying the desired orthonormality condition. Thus one maximizer of (3.7a) is $y_{n+1}^{F*} = y_{n+1}$. This completes the inductive step, showing that $y_j^{F*} = y_j$ for $1 \le j \le N$.

Nearly the same argument shows that the greedy iteration (3.7b) has a solution branch equal to A_N; we omit the proof.

Appendix B. Proof of Lemma 3.3. The results of this lemma are essentially well-known (e.g., historically [33, 14], or see [24] for a modern compilation of these results). However, these results are scattered so we provide a proof here in order to be self-contained.

There are five enumerated statements in Lemma 3.3. In this univariate $d = 1$ case, our orthonormal polynomials $\varphi_n(z)$ satisfy a three-term recurrence and the Christoffel–Darboux relation,

$$z\varphi_n(z) = \sqrt{b_n}\varphi_{n-1}(z) + a_n\varphi_n(z) + \sqrt{b_{n+1}}\varphi_{n+1}(z),$$

(B.1)
$$\sum_{j=0}^{n-1} \varphi_j(x)\varphi_j(z) = \sqrt{b_n}\frac{\varphi_n(x)\varphi_{n-1}(z) - \varphi_n(z)\varphi_{n-1}(x)}{x - z}$$

(B.2)
$$= \frac{\sqrt{b_n}}{\varphi_{n-1}(x)\varphi_{n-1}(z)}\left(\frac{r_n(x) - r_n(z)}{x - z}\right),$$

where the constants a_n and b_n depend on the polynomial moments of ρ. We define the set

$$A_N(y) = r_N^{-1}\left(r_N(y)\right),$$

where $r_N(y) \in \mathbb{R}$ since by assumption $y \notin \varphi_{N-1}^{-1}(0)$. The function r_N is meromorphic with $N - 1$ distinct poles on $\mathbb{R}$, and a straightforward computation shows that $r_N'(z)$ is continuous and positive everywhere except at its poles. This, coupled with the fact that $\lim_{z \to \pm\infty} r_N(z) = \pm\infty$, shows that the set $A_N(y)$ defined above consists of N distinct points on $\mathbb{R}$. We label these points as y_j:

$$A_N(y) = \{y_1, \ldots, y_N\}.$$

By definition, $r_N(y_j)$ is the same number for any $1 \le j \le N$. We now consider the matrix $\boldsymbol{V}\left(A_N(y); Q\right)$, whose rows are $\boldsymbol{\psi}^T(y_j)$, which is defined in (A.2). Note that $\boldsymbol{\psi}(y_j)$ is a unit vector for each j. We have, for $j \ne k$,

$$\boldsymbol{\psi}^T(y_j)\boldsymbol{\psi}(y_k) = \frac{1}{\sqrt{K_\Lambda(y_j)K_\Lambda(y_k)}}\sum_{j=0}^{N-1}\varphi_j(y_j)\varphi_j(y_k)$$

$$\overset{(B.1)}{=} \frac{\sqrt{b_N}}{\varphi_{N-1}(y_j)\varphi_{N-1}(y_k)\sqrt{K_\Lambda(y_j)K_\Lambda(y_k)}}\left(\frac{r_N(y_j) - r_N(y_k)}{y_j - y_k}\right) = 0,$$

where the last equality holds since $r_N(y_j) = r_N(y_k)$ with $y_j \ne y_k$. Thus, $\boldsymbol{V}$ is a square matrix with orthonormal rows; therefore it is an orthogonal matrix and has modulus

determinant 1 and condition number 1. Therefore, $A_N(y)$ satisfies (3.8), proving the first statement in the lemma. The second statement, uniqueness of $A_N(y)$, is straightforward given the construction above. The third statement, defining $A_N(y)$ as level sets of r_N, is our explicit construction above.

To show the fourth statement, consider the matrix $\boldsymbol{V} = \boldsymbol{V}(A_N(y), Q)$, which we have already shown is an orthogonal matrix, and hence

$$\boldsymbol{V}^T \boldsymbol{V} = \boldsymbol{I}_N.$$

For row and column indices i and j, respectively, the componentwise equality above reads

$$\text{(B.3)} \qquad \sum_{q=1}^{N} \frac{1}{K_\Lambda(y_q)} \varphi_i(y_q)\varphi_j(y_q) = \delta_{i,j} = \int_\Gamma \varphi_{i-1}(z)\varphi_{j-1}(z)\rho(z)dz$$

for $0 \le i, j \le N-1$, where the final equality is just orthonormality of the polynomials φ_i. This shows that the quadrature rule whose abscissae are collocated at $A_N(y)$ exactly integrates products $\varphi_i\varphi_j$ for $0 \le i, j \le N-1$. Let p be an arbitrary polynomial of degree $2N-2$ or less. Euclidean division of this polyomial by φ_{N-1} yields

$$p = \varphi_{N-1}q + r,$$

where q is a polynomial of degree $N-1$ or less, and r is a polynomial of degree $N-2$ or less. We now use the fact that $\{\varphi_j\}_{j=0}^{N-1}$ is a basis for polynomials of degree $N-1$ and less, and so there exist constants c_j and d_j such that

$$q(z) = \sum_{j=0}^{N-1} c_j\varphi_j(z), \qquad\qquad r(z) = \sum_{j=0}^{N-1} d_j\varphi_j(z),$$

with $d_{N-1} = 0$. Since ρ is a probability density, then $\varphi_0(z) = 1$, implying

$$p(z) = \sum_{j=0}^{N-1} c_j\varphi_j(z)\varphi_{N-1}(z) + \sum_{j=0}^{N-1} d_j\varphi_j(z)\varphi_0(z).$$

Since the quadrature rule in (B.3) is linear and can exactly integrate products $\varphi_i\varphi_j$, it can exactly integrate p, which equals a sum of such products. This shows statement 4 of the lemma.

The final statement is straightforward: if $y \in \varphi_N^{-1}(0)$, then $r_N(y) = 0$, and so $A_N = r_N^{-1}(0)$. The zero level set $r_N^{-1}(0)$ coincides with the zero level set $\varphi_N^{-1}(0)$ since the roots of φ_{N-1} and the roots of φ_N are disjoint sets. This shows that $A_N(y)$ is the zero set of φ_N; therefore, these points are the Gaussian quadrature nodes, proving statement 5 of the lemma.

Appendix C. Proof of Theorem 3.4. This proof is essentially a combination of Lemma 3.3 and Theorem 3.2. By Lemma 3.3, the set A_N has determinant modulus and condition number 1, satisfying (3.8). Thus, Theorem 3.2 guarantees that both greedy algorithms (3.7) produce the same optimal set A_N. Since the set A_N containing y is unique, this is true regardless of which solution branches are taken during the greedy iterations (3.7).

REFERENCES

[1] L. Bos, J. P. Calvi, N. Levenberg, A. Sommariva, and M. Vianello, *Geometric weakly admissible meshes, discrete least squares approximation and approximate Fekete points*, Math. Comp., 80 (2011), pp. 1623–1638.

[2] L. Bos, S. De Marchi, A. Sommariva, and M. Vianello, *Computing multivariate Fekete and Leja points by numerical linear algebra*, SIAM J. Numer. Anal., 48 (2010), pp. 1984–1999, https://doi.org/10.1137/090779024.

[3] L. Bos and N. Levenberg, *On the calculation of approximate Fekete points: The univariate case*, Electron. Trans. Numer. Anal., 30 (2008), pp. 377–397.

[4] L. Bos, S. D. Marchi, A. Sommariva, and M. Vianello, *Computing multivariate Fekete and Leja points by numerical linear algebra*, SIAM J. Numer. Anal., 48 (2010), pp. 1984–1999.

[5] L. P. Bos and N. Levenberg, *On the calculation of approximate Fekete points: The univariate case*, Electron. Trans. Numer. Anal., 30 (2008), pp. 377–397.

[6] J.-P. Calvi and N. Levenberg, *Uniform approximation by discrete least squares polynomials*, J. Approx. Theory, 152 (2008), pp. 82–100, https://doi.org/10.1016/j.jat.2007.05.005.

[7] J. P. Calvi and N. Levenberg, *Uniform approximation by discrete least squares polynomials*, J. Approx. Theory, 152 (2008), pp. 82–100.

[8] A. Chkifa, A. Cohen, G. Migliorati, F. Nobile, and R. Tempone, *Discrete least squares polynomial approximation with random evaluations application to parametric and stochastic elliptic pdes*, ESAIM Math. Model. Numer. Anal., 49 (2015), pp. 815–837.

[9] A. Chkifa, A. Cohen, and C. Schwab, *High-dimensional adaptive sparse polynomial interpolation and applications to parametric PDEs*, Found. Comput. Math., 14 (2014), pp. 601–633.

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

[11] A. Cohen and G. Migliorati, *Optimal Weighted Least-Squares Methods*, arXiv:1608.00512 [math, stat], (2016).

[12] A. Doostan and H. Owhadi, *A non-adaptive sparse approximation for pdes with stochastic inputs*, J. Comput. Phys., 230 (2011), pp. 3015–3034.

[13] M. Eldred, *Recent advances in non-intrusive polynomial chaos and stochastic collocation methods for uncertainty analysis and design*, in Proceedings of the 50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, 2009.

[14] G. Freud, *Orthogonal Polynomials*, Pergamon Press, Oxford, UK, 1971.

[15] B. Ganapathysubramanian and N. Zabaras, *Sparse grid collocation methods for stochastic natural convection problems*, J. Comput. Phys., 225 (2007), pp. 652–685.

[16] R. G. Ghanem and P. D. Spanos, *Stochastic Finite Elements: A Spectral Approach*, Springer, New York, 1991, http://portal.acm.org/citation.cfm?id=103013.

[17] J. Hampton and A. Doostan, *Coherence motivated sampling and convergence analysis of least squares polynomial chaos regression*, Comput. Methods Appl. Mech. Eng., 290 (2015), pp. 73–97.

[18] J. Jakeman and A. Narayan, *Generation and Application of Multivariate Polynomial Quadrature Rules*, preprint, arXiv:1711.00506, 2017.

[19] J. D. Jakeman, A. Narayan, and T. Zhou, *A generalized sampling and preconditioning scheme for sparse approximation of polynomial chaos expansions*, SIAM J. Sci. Comput., 39 (2017), pp. A1114–A1144.

[20] G. Migliorati, *Multivariate Markov-type and Nikolskii-type inequalities for polynomials associated with downward closed multi-index sets*, J. Approx. Theory, 189 (2015), pp. 137–159, https://doi.org/10.1016/j.jat.2014.10.010.

[21] G. Migliorati, F. Nobile, E. von Schwerin, and R. Tempone, *Approximation of quantities of interest in stochastic PDEs by the random discrete L^2 projection on polynomial spaces*, SIAM J. Sci. Comput., 35 (2013), pp. A1440–A1460.

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

[23] A. Narayan, *Computation of Induced Orthogonal Polynomial Distributions*, arXiv:1704.08465 [math], 2017.

[24] A. Narayan, *Polynomial Approximations by Sampling from the Spectral Distribution*, preprint, 2017.

[25] A. Narayan and J. Jakeman, *Adaptive Leja sparse grid constructions for stochastic collocation and high-dimensional approximation*, SIAM J. Sci. Comput., 36 (2014), pp. A2952–A2983, https://doi.org/10.1137/140966368.

[26] A. NARAYAN, J. D. JAKEMAN, AND T. ZHOU, *A Christoffel function weighted least squares algorithm for collocation approximations*, Math. Comput., 86 (2017), pp. 1913–1947.

[27] A. NARAYAN AND D. XIU, *Stochastic collocation methods on unstructured grids in high dimensions via interpolation*, SIAM J. Sci. Comput., 34 (2012), pp. A1729–A1752, https://doi.org/10.1137/110854059.

[28] A. NARAYAN AND T. ZHOU, *Stochastic collocation on unstructured multivariate meshes*, Commun. Comput. Phys., 18 (2015), pp. 1–36, https://doi.org/10.4208/cicp.020215.070515a.

[29] F. NOBILE, R. TEMPONE, AND C. G. WEBSTER, *An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data*, SIAM J. Numer. Anal., 46 (2008), pp. 2411–2442, https://doi.org/10.1137/070680540.

[30] Y. SHIN AND D. XIU, *Nonadaptive quasi-optimal points selection for least squares linear regression*, SIAM J. Sci. Comput., 38 (2015), pp. A385–A411.

[31] R. C. SMITH, *Uncertainty Quantification: Theory, Implementation, and Applications*, Comput. Sci. Eng. 12, SIAM, Philadelphia, 2013.

[32] A. SOMMARIVA AND M. VIANELLO, *Computing approximate Fekete points by QR factorizations of Vandermonde matrices*, Comput. Math. Appl., 57 (2009), pp. 1324–1336, https://doi.org/10.1016/j.camwa.2008.11.011.

[33] G. SZEGÖ, *Orthogonal Polynomials*, 4th ed., AMS, Providence, RI, 1975.

[34] T. TANG AND T. ZHOU, *On discrete least square projection in unbounded domain with random evaluations and its application to parametric uncertainty quantification*, SIAM J. Sci. Comput., 36 (2014), pp. A2272–A2295.

[35] D. XIU AND J. S. HESTHAVEN, *High-order collocation methods for differential equations with random inputs*, SIAM J. Sci. Comput., 27 (2005), pp. 1118–1139.

[36] D. XIU AND G. KARNIADAKIS, *The Wiener-Askey polynomial chaos for stochastic differential equations*, SIAM J. Sci. Comput., 24 (2002), pp. 619–644.

[37] Z. XU AND T. ZHOU, *On sparse interpolation and the design of deterministic interpolation points*, SIAM J. Sci. Comput., 36 (2014), pp. A1752–A1769.

[38] L. YAN, L. GUO, AND D. XIU, *Stochastic collocation algorithms using ℓ_1-minimization*, Int. J. Uncertain. Quantif., 2 (2012), pp. 279–293.

[39] J. J. YE AND J. ZHOU, *Minimizing the conditiona number to construct design points for polynomial regression models*, SIAM J. Optim., 23 (2013), pp. 666–686.

[40] T. ZHOU, A. NARAYAN, AND D. XIU, *Weighted discrete least-squares polynomial approximation using randomized quadratures*, J. Comput. Phys., 298 (2015), pp. 787–800.