



A Low Discrepancy Sequence on Graphs

A. Cloninger¹ · H. N. Mhaskar²

Received: 28 September 2020 / Revised: 27 May 2021 / Accepted: 28 May 2021 /

Published online: 25 August 2021

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

Many applications such as election forecasting, environmental monitoring, health policy, and graph based machine learning require taking expectation of functions defined on the vertices of a graph. We describe a construction of a sampling scheme analogous to the so called Leja points in complex potential theory that can be proved to give low discrepancy estimates for the approximation of the expected value by the empirical expected value based on these points. In contrast to classical potential theory where the kernel is fixed and the equilibrium distribution depends upon the kernel, we fix a probability distribution and construct a kernel (which represents the graph structure) for which the equilibrium distribution is the given probability distribution. Our estimates do not depend upon the size of the graph.

Keywords Equal weight quadrature on graphs · Potential theory on graphs · Density approximation on graphs · Leja points on graphs

1 Introduction

In many applications, the data is not represented by points in a high dimensional Euclidean space, but instead as vertices of networks with pairwise relations. A common problem that arises is estimating the mean or some integral of a function on those vertices, under the assumption that the function is smooth in an appropriate sense

Communicated by Isaac Pesenson.

✉ H. N. Mhaskar
hrushikesh.mhaskar@cgu.edu

A. Cloninger
acloninger@ucsd.edu

¹ Department of Mathematics and Halicioglu Data Science Institute, University of California San Diego, San Diego, CA 92093, USA

² Institute of Mathematical Sciences, Claremont Graduate University, Claremont, CA 91711, USA

with respect to the network. For example, prior to elections, polls are used to estimate the opinions of a networked population about the candidates in expectation [17]. In environmental monitoring, estimating average temperature or water quality from a network of sensors allows governments to take preventative action [22]. In health policy, monitoring the average level of health and happiness in various populations requires extensive random polling [10,40], which could be vastly downsampled with knowledge of the social networks of the population. Finally, in semi-supervised and active learning using Graph Neural Networks, the generalization error is specified as the average of the loss function associated with all vertices in a graph [20,49].

In real world networks, evaluating the function at every vertex, or even at a large number of vertices, is very impractical. Therefore, an important question is how to choose a smaller subset of the vertices at which to evaluate the function and estimate the desired integral. Approaches such as Monte-Carlo integration are independent of the network, and do not exploit the assumption that the function is smooth with respect to the network. Methods that exploit the geometry of the domain, including potential theory estimates [14,14] and quasi-Monte-Carlo sampling [12], are only defined for intervals and compact vector spaces, and only for particular measures.

The purpose of this paper is to investigate deterministic constructions of low discrepancy sequences on undirected graphs, where the integration is taken with respect to an arbitrary measure, supported on the entire vertex set of the graph. We will use potential theory ideas, in particular, study a construction analogous to the so called Leja points. In contrast to classical potential theory, we start with a given measure and a graph, and modify the weights on the graph so that the given measure is the equilibrium measure.

1.1 Related Work

Let Ω be any measure space, μ be a probability measure on Ω , f be a random variable on Ω with $|f(x)| \leq R$ for almost all $x \in \Omega$, $\mathbb{E}_\mu(|f|^2) = V$. Then if $\delta \in (0, 1)$, M is sufficiently large, and x_1, \dots, x_M are random samples from Ω , then it is well known (for example, using Bernstein concentration inequality) that with μ -probability $\geq 1 - \delta$,

$$\left| \frac{1}{M} \sum_{k=1}^M f(x_k) - \int_{\Omega} f(x) d\mu(x) \right| \leq \sqrt{\frac{2(V + R) \log(2/\delta)}{M}}. \quad (1.1)$$

Of course, deterministic variants of this inequality can be obtained using different kinds of deterministic assumptions on f . For example, an equivalent formulation of a theorem of Erdős and Turán [14] is the following. Let $\{P_M(z) = \prod_{k=1}^M (z - x_{k,M})\}$ be a sequence of monic polynomials, and $A_M = 2^M \max_{z \in [-1, 1]} |P_M(z)|$. Then for any function $f : [-1, 1] \rightarrow \mathbb{R}$ having a bounded total variation $\|f\|_{TV}$ on $[-1, 1]$, we have

$$\left| \frac{1}{M} \sum_{k=1}^M f(x_{k,M}) - \frac{1}{\pi} \int_{-1}^1 \frac{f(x)}{\sqrt{1-x^2}} dx \right| \leq \frac{8}{\log 3} \sqrt{\frac{\log A_M}{M}} \|f\|_{TV}. \quad (1.2)$$

Theorems of this kind are known as discrepancy theorems. Although this estimate cannot be improved in general, it was improved by Blatt [4] with additional assumptions as follows. Let the points $\{x_{k,M}\} \subset [-1, 1]$, and $2^M |P'_M(x_{k,M})| \geq B_M^{-1}$. Then there exists a constant $c > 0$ such that

$$\left| \frac{1}{M} \sum_{k=1}^M f(x_{k,M}) - \frac{1}{\pi} \int_{-1}^1 \frac{f(x)}{\sqrt{1-x^2}} dx \right| \leq c \log \max(A_M, B_M, M) \frac{\log M}{M} \|f\|_{TV}. \quad (1.3)$$

Both of these estimates depend heavily on potential theory estimates; in particular, the fact that $1/2$ is the logarithmic capacity of $[-1, 1]$, and the measure of integration above is the equilibrium distribution for $[-1, 1]$. There is a large amount of research devoted to generalization of this work including those involving potential theory in higher dimensions. In particular, an analogue of (1.3) in the case of arbitrary measures rather than equilibrium measures is given in [5]. A survey can be found in the book [1] of Andrievskii and Blatt.

The problem is of interest also in the theory of information based complexity where one seeks to approximate an integral over high dimensional spaces with averages of samples of the integrand. These have a different flavor, where instead of thinking in terms of zeros of polynomials and potential theory arguments, the interest is in devising quasi-Monte-Carlo systems with low discrepancy; i.e., system of points for which an estimate analogous to (1.3) holds, especially where the dependence of the constants on dimension are desired to have a polynomial growth with respect to the dimension. Most of these estimates are in the context of integration of 1-periodic functions on $[0, 1]^q$ with respect to the Lebesgue measure and the total variation is taken in the sense of the so called Hardy-Kraus variation. A survey can be found in the book [12] of Dick and Pillichshammer. Methods have also been proposed to create a low discrepancy sequence through an accept/reject model for uniform random variables [13]. Existence theorems in the context of general measures and domains are also known in the literature, based mostly on probability theory ideas [31, 32]. We note finally that the emphasis here is on approximation of an integral by an **unweighted average** of the samples of the integrand, not on quadrature formulas where a suitably weighted average of the samples can yield substantially better estimates under various smoothness assumptions on the integrand.

The question of function approximation based on samples of the target function on a graph are well studied, especially in the context of band-limited functions. For example, the papers [2, 8] discuss algorithms for obtaining points on a graph with the property that it is possible to reconstruct band-limited functions on the graph exactly using samples at these points. The question is studied from the point of view of compressive sensing in [38]. The paper [36] presents a detailed study of the space of band-limited functions on a graph and the sets of uniqueness for such functions. A key fact that characterizes such sets is that a Marcinkiewicz-Zygmund inequality holds for the space of band-limited functions involved. In turn, this leads to quadrature formulas exact to for integration of these spaces. In this paper, we deal with signals that are not necessarily band-limited; indeed, the notion of spectral decomposition of

the graph Laplacian plays no role in our theory. On the other hand, our interest is in approximating integrals of functions rather than the functions themselves.

A standard reference on potential theory is the book [25] by Landkof. We are not aware of any prior work specifically for potential theory on graphs. However, there are a number of papers dealing with potential theory on locally compact spaces, e.g., [16,30,34]. The notion of a sequence of Leja points was introduced in [26] in the context of approximation of the equilibrium measure of a compact subset of the complex plane. A detailed analysis of the rate of convergence of the sequence of measures σ_n that associates the mass $1/n$ with the first n points in this sequence to the equilibrium measure is given most recently by Pritsker [37]. This notion of discrepancy theorems can be generalized in other, more general contexts. A survey can be found in [11] by De Marche, where computational issues are discussed. Discrepancy theorems for Leja points in the context of hyper-spheres is analyzed by Götz in [18].

After the submission of our paper, we came across a paper by Brown [6], where the author has given a construction of a sequence of good discrepancy points on a graph based on the Green function of a power of the graph Laplacian, where the discrepancy is measured in terms of the Wasserstein metric. Naturally, the error in integration is estimated analogously to the Hlawka-Koksma inequality [23] in terms of the Lipschitz constant of the function. In the current paper, we propose a construction independently of the eigen-decomposition of the graph Laplacian, and estimate the discrepancy using the matrix involved in our construction.

There has been recent interest in coresets selection on general domains, including graphs. A general selection of points that are well distributed on graphs has been found through randomized QR decompositions [3], and through random walk sampling [19,33], however these results do not provide estimates on the error in approximation of an integral. There has also been significant work on sample selection for full reconstruction of the underlying signal (see [45] for a review of methods). In [44], Steinerberger provided bounds to guarantee the existence of quadrature formulas (known in the paper as graphical designs), which find exactly the averages of eigenfunctions of the graph Laplacian corresponding to large eigenvalues. In [27], Linderman and Steinerberger provide bounds for the quadrature error in computing the average of function values based on the values of the function at arbitrary points **and quadrature weights**. These bounds apply only for spectrally band limited functions. The bounds depend upon the spectral band, the L^2 norm of the target function, and certain powers of the graph Laplacian applied to the vector of the quadrature weights involved. Algorithms to choose these points and weights can be found in [28] for manifolds and [46] for a greedy algorithm of point and weight selection on graphs.

1.2 Motivating Example

As an example, in Fig. 1 we display a quasi-Monte-Carlo sampling scheme on the unit square and compare it to the various proposed sampling schemes from this paper. For the quasi-Monte-Carlo sampling, we use a Halton sequence on the unit square, skipping the first 1000 samples and with a leap of 100 [47]. For our proposed methods, we begin with 10,000 points uniformly sampled on the unit square $[0, 1]^2$, and build a 50-

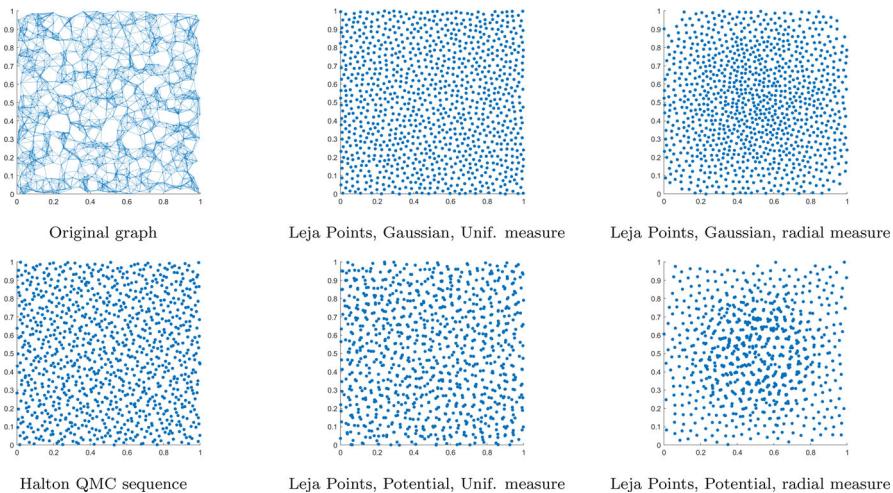


Fig. 1 Examples of various sampling schemes on the unit square. We select 1000 points for sampling. For Leja points, this is a chosen subset of 10,000 iid uniform distribution points. See Sect. 1.2 for a full description. Note, original graph image was computed on smaller number of points/edges for visualization purposes only

nearest-neighbor graph. The weighted edges are computed using: a Gaussian weight $A(x_i, x_j) = e^{-\|x_i - x_j\|^2/0.01^2}$, or a log potential weight $A(x_i, x_j) = \log(1/(\|x_i - x_j\| + \varepsilon))$ with $\varepsilon = 10^{-5}$. We select 1000 Leja points using the algorithm and construction proposed in Sects. 2.3 and 4.1. For each graph adjacency type, this is done for two different measures that we wish to integrate against: a uniform distribution on $[0, 1]$, or a nonuniform radial measure $v^*(x) \propto e^{-\|x - c\|^2/0.25^2}$ where $c = (0.5, 0.5)$.

It is clear from this example that the proposed graph Leja points with a uniform distribution recover a similar low discrepancy layout to quasi-Monte-Carlo sampled points. And beyond this, the graph Leja points are able to generalize to a non-uniform distribution with the same framework. But the true benefit of this proposed Leja point construction is that the algorithm applies in the case of an arbitrary graph, and will still result in well-spaced points.

1.3 Outline of the Paper

We describe our main theorem in Sect. 2 after reviewing some basic facts from potential theory in the abstract. As mentioned earlier, our approach is not to start with a matrix and work with whatever equilibrium measure is associated with it, but rather to start with a graph adjacency matrix and a given probability distribution on the vertices of the graph, and construct another matrix with the same properties as the graph for which this distribution is the equilibrium distribution. We describe three such constructions in Sect. 3. The theory is illustrated with various synthetic and real world examples in Sect. 4. The proof of the results in Sect. 2 are given in Sect. 5.

2 Main Theorem

We develop some basic notation in Sect. 2.1, and review a fundamental theorem in potential theory in Sect. 2.2. The notion of Leja points is defined in Sect. 2.3. Our main theorem, Theorem 2.1 shows that the sequence of Leja points is a low discrepancy sequence for integration with respect to the given measure ν^* .

2.1 Notation

Let G be a $N \times N$ symmetric matrix, \mathbb{X} be a finite set with $|\mathbb{X}| = N$. We prefer to index G with $\mathbb{X} \times \mathbb{X}$.

We consider any function $\nu : \mathbb{X} \rightarrow \mathbb{R}$ to be a measure on \mathbb{X} as well as a function on \mathbb{X} , as well as a vector. So, for any function $f : \mathbb{X} \rightarrow \mathbb{R}$,

$$\int f d\nu = \sum_{x \in \mathbb{X}} f(x) \nu(x) = \int \nu df. \quad (2.1)$$

For measures ν, μ on \mathbb{X} , we use the notation

$$G(x, \nu) = \int G(x, y) d\nu(y), \quad G(\mu, \nu) = \int G(x, y) d\mu(x) d\nu(y). \quad (2.2)$$

In particular, $G(x, \nu)$ is the x -th component of $G\nu$. The class of all probability measures on \mathbb{X} is denoted by \mathcal{P} . The class of all measures $\nu \in \mathcal{P}$ for which $\nu(x) > 0$ for all $x \in \mathbb{X}$ is denoted by \mathcal{P}_+ .

We denote the vector $(1, \dots, 1)^\top$ by $\mathbf{1}$.

A matrix G is conditionally positive semi-definite (c.p.s.d.) if $v^\top G v \geq 0$ for all v with $v^\top \mathbf{1} = 0$, and conditionally positive definite (c.p.d.) if it is c.p.s.d. and $v^\top G v = 0$, $v^\top \mathbf{1} = 0$ together imply that $v = 0$.

For a symmetric matrix A , we denote

$$\begin{aligned} M(A) &= \max_{x \in \mathbb{X}} |A(x, x)|, & \|A\| &= \max_{x \in \mathbb{X}} \sum_{y \in \mathbb{X}} |A(x, y)|, \\ \|A\|' &= \max_{x \in \mathbb{X}} \sum_{\substack{y \in \mathbb{X} \\ y \neq x}} |A(x, y)|. \end{aligned} \quad (2.3)$$

For a vector v , we denote the ℓ^p norm of v by $\|v\|_p$, and write $\kappa(v) = \|v\|_\infty / \min_{x \in \mathbb{X}} |v(x)|$.

2.2 Potential Theory

A measure $\nu^* \in \mathcal{P}$ is the equilibrium measure (vector) if

$$\Gamma(G) = G(\nu^*, \nu^*) = \min_{\nu \in \mathcal{P}} G(\nu, \nu). \quad (2.4)$$

In the context of complex potential theory, the quantity $\Gamma(G)$ is often called Robin's constant. In our context, we refer to the quantity $\Gamma(G)$ as the *capacity* of G .

An important characterization of the equilibrium measure is given by the following lemma, known generically as Frostman theorem. A proof can be found in almost any book on classical potential theory (e.g., [25, pp. 136–137] in the context of potentials on Euclidean spaces and [34, Theorem 2.1] in the context of locally compact spaces). We will reproduce a proof for the sake of completion in Sect. 5.

Lemma 2.1 (a) *If v^* is any equilibrium measure, and S^* is the support of v^* , then*

$$G(x, v^*) \geq \Gamma(G), \quad x \in \mathbb{X}, \quad (2.5)$$

and

$$G(x, v^*) = \Gamma(G), \quad x \in S^*. \quad (2.6)$$

(b) *Let G be conditionally positive semi-definite. If $c \in \mathbb{R}$, and any measure $\mu \in \mathcal{P}$ with support S satisfies both*

$$G(x, \mu) \geq c, \quad x \in \mathbb{X}, \quad (2.7)$$

and

$$G(x, \mu) = c, \quad x \in S, \quad (2.8)$$

then $c = \Gamma(G)$, and μ is an equilibrium measure. If v^* is any equilibrium measure, then $G(\mu - v^*, \mu - v^*) = 0$.

2.3 Leja Points

In this section, we will assume a symmetric, conditionally positive semi-definite matrix G , which has $v^* \in \mathcal{P}_+$ as an equilibrium measure. In particular,

$$G(x, v^*) = G(v^*, v^*) = \min_{v \in \mathcal{P}_+} G(v, v) = \min_{v \in \mathcal{P}_+} \max_{y \in \text{supp}(v)} G(y, v) = \Gamma(G). \quad x \in \mathbb{X}. \quad (2.9)$$

Definition 2.1 A sequence $\{a_k\}_{k=0}^{\infty}$ of points in \mathbb{X} will be called a *Leja sequence* (with respect to a matrix G) if for every $k \geq 1$,

$$a_k = \arg \min_{x \in \mathbb{X}} \sum_{j=0}^{k-1} G(x, a_j). \quad (2.10)$$

The points a_k will be referred to as Leja points.

Remark 2.1 The points a_k are in general not distinct. In fact, in order for Theorem 2.1 to hold, the points will need to be repeated to be commensurate with the measure v^* . Second, a Leja sequence is not uniquely determined by the initial point a_0 , since the definition does not require the $\arg \min$ to be unique.

We will denote the Dirac delta at a_k by v_k . For $k \geq 1$, let $\sigma_k = (1/k) \sum_{j=0}^{k-1} v_k$. Then

$$G(v_k, \sigma_k) = \min_{x \in \mathbb{X}} G(x, \sigma_k), \quad k \in \mathbb{Z}_+. \quad (2.11)$$

If f is the range of G , we write

$$\mathcal{D}_G(f) = \arg \min \{ \|w\|_1 : Gw = f \}. \quad (2.12)$$

Example 2.1 Let G be any positive definite, $N \times N$ matrix, \mathcal{F} be the class of all convex combinations of columns of G . If $f \in \mathcal{F}$, then $\|\mathcal{D}_G(f)\|_1 = 1$. \square

Example 2.2 Let A be the weighted adjacency matrix of a connected graph that does not have non-trivial, bipartite connected subgraph, \mathcal{L} be the graph Laplacian, $\{\lambda_k\}$ be the set of eigenvalues of \mathcal{G} , and $\{\phi_k\}$ be the corresponding orthonormalized eigenvectors. Then ([9, Lemma 1.7]) $\{\lambda_k\} \subset [0, 2)$. We let $G = I - (1/2)\mathcal{L}$. Then G is positive definite with non-negative entries, and positive diagonal, so that Theorem 2.1 is applicable, and the equilibrium measure for G is defined uniquely. In particular, G is invertible. If $f : V \rightarrow \mathbb{R}$, $f = \sum_k \hat{f}(k)\phi_k$, then

$$\mathcal{D}_G(f) = 2 \sum_k \frac{\hat{f}(k)}{2 - \lambda_k} \phi_k.$$

Let $\lambda < 2$ and Π_λ be the class of all λ -band-limited functions; i.e., the class of all $f : \mathbb{X} \rightarrow \mathbb{R}$ for which $\hat{f}(k) = 0$ if $\lambda_k \geq \lambda$. For $f \in \Pi_\lambda$, a crude estimate for $\|\mathcal{D}_G(f)\|_1$ is given by

$$\|\mathcal{D}_G(f)\|_1 \leq \frac{2}{2 - \lambda} \max_{k: \lambda_k < \lambda} \|\phi_k\|_1 \sum_{k: \lambda_k < \lambda} |\hat{f}(k)|.$$

\square

Our main theorem is the following.

Theorem 2.1 Let $G(x, y) \geq 0$ for all $x, y \in \mathbb{X}$, $\{a_k\}$ be a Leja sequence with respect to G , $v^* \in \mathcal{P}_+$ be an equilibrium measure for G , and f be in the range of G . Then

$$\left| \int_{\mathbb{X}} f d v^* - \frac{1}{n} \sum_{k=0}^{n-1} f(a_k) \right| \leq \frac{3 \|G\|}{n+1} \|\mathcal{D}_G(f)\|_1. \quad (2.13)$$

Remark 2.2 In contrast to the estimates given in [27, 46], our estimates do not require the function f to be band-limited. \square

Remark 2.3 We note that the rate of convergence of the equal weight quadrature formulas is much faster than what is expected for quasi-Monte Carlo methods, demonstrating that the sequence of Leja points is a low discrepancy sequence in the sense of [12]. It is not clear what conditions on the matrix G will ensure an equilibrium measure in \mathcal{P}_+ . In Sect. 3, we will give a variety of possible constructions to modify either an arbitrary matrix or a matrix with certain properties to ensure that any given vector $v^* \in \mathcal{P}_+$ is an equilibrium measure of the resulting matrix. Thus, rather than taking the viewpoint that G is given and v^* is its equilibrium measure, we will start with v^* and construct G with v^* as the equilibrium measure. Then the Leja sequence depends upon v^* via G . \square

The proof of Theorem 2.1 mimics the standard proof of the convergence of the measures v_n in classical potential theory. Thus, we note that the mapping $(\mu, v) \rightarrow G(\mu, v)$ is a semi-inner product on the space of all measures on \mathbb{X} . We prove first that the sequence v_n converges to v^* in the sense of the semi-norm defined by this semi-inner product. This implies in turn that the Cesáro means $(C, 2)$ of the sequence v_n converges to v^* . A tauberian argument then completes the proof. Of course, we need to keep track of the rates of convergence at each stage.

3 Constructions

Let B be any symmetric matrix, and $v \in \mathcal{P}_+$. We want to construct a matrix G such that G is conditionally positive semi-definite, and $Gv = c\mathbf{1}$, so that v is the equilibrium vector for G . We describe three constructions.

3.1 Diagonal Modification

This construction gives a modified graph Laplacian, and works with any symmetric matrix B with no further assumptions. First, we construct a symmetric matrix B_1 so that $B_1v = 0$.

We define the v -Laplacian $L_v(B)$ by

$$w(x) = \frac{(Bv)(x)}{v(x)}, \quad W = \text{diag}(w(x)), \quad L_v(B) = W - B. \quad (3.1)$$

Then with $B_1 = L_v(B)$,

$$\|B_1\| \leq (\kappa(v) + 1)\|B\|. \quad (3.2)$$

Let $V = \text{diag}(v(x)/\|v\|_\infty)$, and

$$G = 2\|B_1\|V^{-1} - B_1. \quad (3.3)$$

Then

$$G(x, x) - \sum_{y \neq x} |G(x, y)| \geq 2\|v\|_\infty \|B_1\| / v(x) - \|B_1\| \geq \|B_1\| > 0.$$

So, G is diagonal dominant, hence, positive definite, and

$$Gv = 2\|v\|_\infty \|B_1\| \mathbf{1}^\top, \quad (3.4)$$

so that v is the unique equilibrium measure of G and

$$\Gamma(G) = 2\|v\|_\infty \|B_1\|. \quad (3.5)$$

Moreover,

$$\|G\|/2 \leq \max_{x, y \in \mathbb{X}} |G(x, y)| \leq \|G\| \leq (2\kappa(v) + 1)\|B_1\|. \quad (3.6)$$

Remark 3.1 The construction of the v -Laplacian works for *all* symmetric matrices B with no further assumptions. In the case when B is the (weighted) adjacency matrix with non-negative weights, then with $V = \text{diag}(v(x))$, $V B V$ is another adjacency matrix for the same graph so that $B(x, y) = 0$ if and only if $(V B V)(x, y) = 0$. The matrix $L(V B V) = V L_v(B) V$ is the non-normalized graph Laplacian for $V B V$. Thus, the eigenvalues of L_v are the same as those of $L(V B V)$, and if ϕ is an eigenvector of $L(V B V)$, then the eigenvector for L_v for the same eigenvalue is $V\phi$. In particular, $L_v(B)$ is positive semi-definite, and the graph is connected if and only if v is the unique null vector for $W - B$.

If B is the weighted adjacency matrix of a connected graph (with positive weights) and B_1 is the v -Laplacian defined in (3.1), the matrix G , defined by (3.3), has the same graph structure as B except for self-loops, and $G(x, y) \geq 0$ for all $x, y \in \mathbb{X}$. \square

3.2 Householder Transform

In the case when B is not invertible, there is another way to construct B_1 , which essentially preserves the eigenstructure of B itself (rather than a graph Laplacian for B). This construction is also applicable for every non-invertible B which may have negative entries. If u and w are unit vectors, the Householder transform is defined by

$$H[w, u] = I - 2 \frac{(w - u)(w - u)^\top}{\|w - u\|_2^2}. \quad (3.7)$$

Clearly, $H[w, u]$ is a symmetric unitary matrix, $H[w, u](w) = u$, and $H[w, u] = H[u, w]$, so that $H[w, u](u) = w$. We have

$$\|H[w, u]\| \leq 1 + 2 \frac{\|w - u\|_\infty \|w - u\|_1}{\|w - u\|_2^2} \leq 1 + 2\kappa(w - u). \quad (3.8)$$

Since B is not invertible, there is a unit vector u such that $Bu = 0$. Writing \tilde{v} for the unit vector along v , we set

$$B_1 = H[u, \tilde{v}]BH[u, \tilde{v}]. \quad (3.9)$$

Clearly, $B_1 v = 0$,

$$\|B_1\| \leq \|H[u, \tilde{v}]\|^2 \|B\|. \quad (3.10)$$

The matrix B_1 has the same eigenvalues as B , and if $B = U\Lambda U^\top$ is the spectral decomposition of B , then the eigenvectors of B are obtained simply by replacing every eigenvector u_j of B by Hu_j . However, this construction depends upon B having a null vector u and our ability of compute $H[u, \tilde{v}]$, and the norms will depend upon the norms of this Householder transform.

With the matrix B_1 defined in (3.9), we construct G as in (3.3).

3.3 Symmetric Scaling

This construction applies only to symmetric, **non-negative** matrices B , and results in a matrix with the same eigenvalues as those of B . For a matrix B , and subsets $S_1, S_2 \subseteq \mathbb{X}$, we denote by $B[S_1, S_2]$ the sub-matrix of B obtained by extracting rows of B indexed by S_1 and columns of B indexed by S_2 .

Let B be a non-negative, symmetric matrix, $v \in \mathcal{P}_+$ satisfy the Brualdi condition
For all partitions S_1, S_2, S_3 of \mathbb{X} such that $B[S_2 \cup S_3, S_3] = 0$, we have

$$\sum_{x \in S_1} v(x) \geq \sum_{y \in S_3} v(y), \quad (3.11)$$

with equality holding if and only if $B[S_1, S_1 \cup S_2] = 0$.

Then a theorem of Brualdi [7] states that there exists a diagonal matrix D with positive entries such that

$$DBD\mathbf{1} = v. \quad (3.12)$$

Let V be the diagonal matrix $(v(x))$. Then $G = V^{-1}DBDV^{-1}$ satisfies

$$Gv = \mathbf{1}. \quad (3.13)$$

Clearly, G represents a graph with the same vertices and edges as B , except for different edge weights. We note that when B has a positive diagonal, then the Brualdi condition is satisfied vacuously for every $v \in \mathcal{P}_+$. The full algorithmic steps for finding such a bistochastic normalization of symmetric matrices can be found in [24,29].

4 Experiments

4.1 Set Up

One matrix B that is trivially row-scalable as in (3.12) is the matrix of all positive entries; i.e., the graph is complete with self-loops at each vertex. Considering a complete matrix prevents a possible blow-up of D on singleton vertices, where otherwise a singleton vertex x would have $D(x, x) = (B(x, x)v(x))^{-1/2}$. On the other hand, it is more expensive to apply the Sinkhorn algorithm [42] to a dense matrix B than applying the algorithm to a sparse matrix B , since the most expensive part of the algorithm is the matrix vector multiplications.

A sparse matrix B can be augmented to have all positive entries in a style similar to Pagerank [35] by adding a small weighted edge between any pair of vertices,

$$\tilde{B} = (1 - \alpha)B + \frac{\alpha}{N} \mathbf{1}\mathbf{1}^\top. \quad (4.1)$$

Then \tilde{B} trivially satisfies the assumptions in Sect. 3.3 since it has a positive diagonal. Moreover, because of the particular structure of \tilde{B} , it is possible to compute a matrix vector multiplication in a time that depends only on the sparsity of B . This is done using the trivial observation that

$$\tilde{B}v = (1 - \alpha)Bv + \frac{\alpha}{N} \mathbf{1}(\mathbf{1}^\top v) = (1 - \alpha)Bv + \frac{\alpha \cdot \sum_x v(x)}{N} \mathbf{1}. \quad (4.2)$$

This observation of the Pagerank modification was originally made in [21].

Similarly, the matrix G used in (2.10) can be constructed from the sparse matrix B by observing that

$$G = V^{-1}D\tilde{B}DV^{-1} = (1 - \alpha)V^{-1}DBDV^{-1} + \frac{\alpha}{N}(Dv^{*-1})(Dv^{*-1})^\top, \quad (4.3)$$

where $V = \text{diag}(v^*(x))$.

The entire algorithm for constructing the Leja points $\{a_j\}$ of the Pagerank graph in Eq. (4.1) is shown in Algorithm 1 for completeness. The computational complexity of the algorithm is mostly rooted in the Sinkhorn iterations for computing G . Let B be an $N \times N$ matrix with k edges per node (the number of edges need not be fixed, but simplifies the complexity calculation). Then each matvec operation Bv requires $O(Nk)$ flops, so computation of $\tilde{B}v$ also requires $O(Nk)$ flops. Exact bounds on the number L of Sinkhorn iterations is not known, but we can still denote the computational complexity as $O(NkL)$. Computing G similarly requires $O(Nk)$ flops and can be held in memory using sparse storage using (4.3). Finally, computing the Leja points requires summing up to n columns of G , which has computational complexity $O(nk)$. This gives a total complexity of $O(NkL) + O(nk)$.

In all the experiments below, we compare three quantities:

Algorithm 1 Graph Leja Point Selection

Input: Sparse matrix B , Pagerank factor α , Base measure v , Number Leja points n , Queryable function f

Output: Integral estimate $\widehat{\mathbb{E}}_v[f]$

- 1: $\widetilde{B} = (1 - \alpha)B + \frac{\alpha}{N} \mathbf{1}\mathbf{1}^\top$ ▷ Notational only, the full ones matrix need not be stored in memory (see (4.2))
- 2: Sinkhorn iterations to find D such that $D\widetilde{B}D\mathbf{1} = v$
- 3: $G = V^{-1}D\widetilde{B}DV^{-1}$
- 4: $a_0 = \text{random vertex } x \in \mathbb{X}$
- 5: **for** k from 1 to $n - 1$ **do**
- 6: $a_k = \arg \min_{x \in \mathbb{X}} \sum_{j=0}^{k-1} G(x, a_j)$
- 7: **end for**
- 8: $\widehat{\mathbb{E}}_v[f] = \frac{1}{n} \sum_{k=0}^{n-1} f(a_k)$

- the average $\frac{1}{n} \sum_{k=0}^{n-1} f(x_k)$, where x_k is sampled randomly (with replacement) according to a weighted probability $v^*(x_k)$. This is denoted “Random” and with blue curves in the experiments.
- the average $\frac{1}{n} \sum_{k=0}^{n-1} f(a_k)$, where a_k are the Leja points of the graph with equilibrium measure v^* . This is denoted “Leja” and with red curves in the experiments.
- the average $\frac{2}{n(n-1)} \sum_{k=0}^{n-1} (n-k) f(a_k)$ (cf. (5.7)), where the a_k are the Leja points of the graph with equilibrium measure v^* . This quantity is considered as a natural weighting of the Leja points that respects the order in which they are drawn. This is denoted “Summability” and with orange curves in the experiments.

We also note for all examples to follow in this section, that because both the Leja point selection algorithm and the Monte Carlo sampling are with replacement, it is possible to compute $n > N$ iterations without selecting all N vertices. This additional sampling can be seen as establishing weights, as a point sampled twice is given twice the weight as a point sampled once. Additionally, we will use the Pagerank kernel G as in (4.1) with $\alpha = 0.05$.

Finally, unless otherwise noted, we set v^* to be the inverse of the kernel density estimate for the graph,

$$v^*(x) \propto \frac{1}{\sum_y A(x, y)}, \quad (4.4)$$

normalized so that $\|v^*\|_1 = 1$. This is an arbitrary choice of v^* , as the only required condition is for $v^*(x) > 0$, and is meant to demonstrate that the results apply to more than $v^* = \frac{1}{N} \mathbf{1}$. This equilibrium measure is of particular interest because, in the point cloud kNN graph setting, it approximately cancels out the sampling density of the points themselves. More specifically, if the points are sampled from the density $p : \mathbb{R}^d \rightarrow \mathbb{R}_+$ that has compact support $\text{supp}(p)$, then

$$\mathbb{E}_{v^*}[f] = \int_{\text{supp}(p)} f(x) p(x) dv^*(x) \approx c_p \int_{\text{supp}(p)} f(x) dx, \quad (4.5)$$

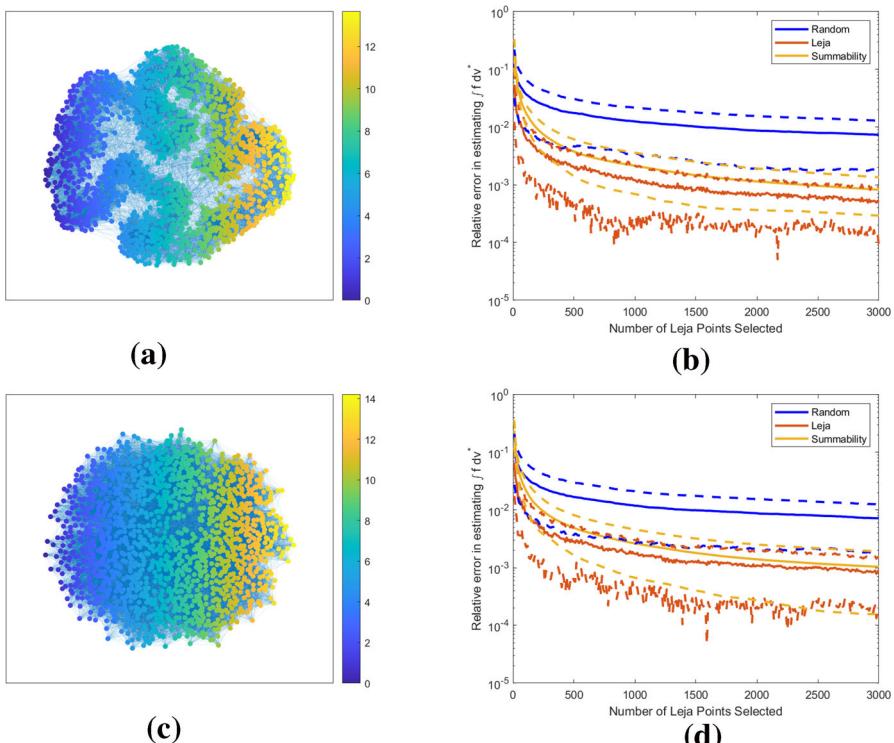


Fig. 2 Watts-Strogatz models for **a–b** $\beta = 0.05$, and **c–d** $\beta = 0.25$. We display both an example of the graph/function, and the quadrature error results over 1000 instantiations of the random graph. The dotted lines correspond to the confidence intervals around each mean curve of the same color

which would be independent of the relative heights of the sampling density. However, any positive density ν^* satisfies the necessary assumptions for the proven approximation rates.

4.2 Synthetic Graphs with Smooth Functions

A common model for social networks is a so-called *small world network* [48]. Small world networks are graphs with a small number of edges per vertex, but where any pair of vertices is likely to have a small graph distance. Mathematically, if the average graph distance between two vertices is d , a small-world graph with N vertices roughly satisfies $d \propto \log(N)$. These networks are antithetical to *nearest neighbor graphs* that are generated from a point cloud in some metric space, and have a large network diameter.

In the first set of experiments, we examine the quadrature approximation for a smooth function on a Watts-Strogatz model [48]. This is a model of random graph generation that exhibits small-world properties, including low path distances between vertices while still exhibiting a high degree of clustering. The model takes inputs of

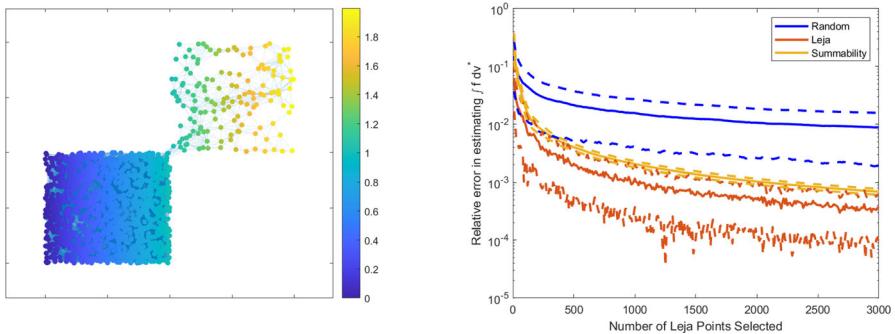


Fig. 3 Two cluster metric graph example with non-uniform density. We display both an example of the graph/function, and the quadrature error results over 1000 instantiations of the random graph. The dotted lines correspond to the confidence intervals around each mean curve of the same color

number of vertices, average number of edges per vertex, and a parameter $\beta \in [0, 1]$ that models the fraction of random connections. When $\beta = 0$ the generated graphs will exhibit lattice structure, and when $\beta = 1$ the graph will resemble an Erdős-Rényi random graph.

In this experiment, we take f to be the x -coordinate of the vertex locations after applying a force layout [15]. This results in a smooth function in the nodes, since a force layout is a heuristic algorithm that repels unconnected nodes away from one another. The functions are displayed in Fig. 2, and the graphs are constructed for various values of the parameter β , which measures the level of structure in the graph. Also, because the graph can be regenerated, we run the experiment for 1000 graphs with the same parameters. We display the mean and standard deviation across these 1000 instantiations of the graph. For each graph, we take $N = 1500$ vertices, and an average of 25 edges per node.

In a second set of experiments, we construct a nearest neighbor graph from the point cloud in Fig. 3. The graph weights are determined with a Gaussian kernel with bandwidth $\sigma = 0.1$. The interesting aspect of this data set is that the points from the two clusters are sampled in a highly non-uniform manner. Because ν^* is constructed by approximating the inverse of the local density, this results in a similar number of points sampled from both clusters, see Fig. 4. The function is taken to be the x -coordinate of the points. Again, because the graph can be regenerated, we run the experiment for 1000 graphs with the same parameters. We display the mean and standard deviation across these 1000 instantiations of the graph. For each graph, we take $N = 1500$ nodes.

We also demonstrate that the location of the Leja points for each of these data sets, using ν^* to be an inverse density estimate, in Fig. 4.

4.3 Synthetic Graphs with Non-band-Limited Functions

A novelty of the results in Theorem 2.1 is that it applies to functions that are not required to be band limited.

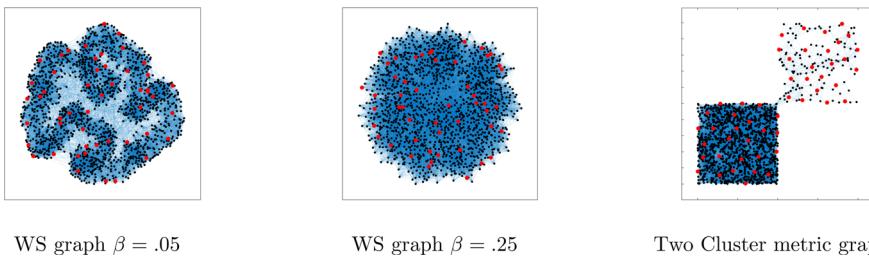


Fig. 4 Layout of 50 Leja points for various Watts-Stogatz data sets, with v^* being the inverse of the kernel density estimate

We demonstrate the strengths of the bound in Eq. (2.13), namely that the function is not required to be band limited in order for the estimation bounds to apply. To characterize this, we generate a series of experiments on the two cluster data from Fig. 3, but with a significantly more complicated function f . To construct f , we consider a spectral decomposition of the matrix $G = \Phi \Lambda \Phi^*$. Then we construct the function spectrally via

$$f = \sum_{k=1}^N \left(\xi_k \lambda_k e^{-\tau k/N} \right) \phi_k, \quad (4.6)$$

where $\xi_k \sim \text{Unif}([0, 1])$, and τ controls the rate of spectral decay with respect to the eigenvalue λ_k . In Fig. 5, we demonstrate the quadrature approximation error averaged across 1000 instantiations of the random graph and random function. We also consider these experiments across varying τ , demonstrating empirically the effect of the spectral decay rate on the overall quadrature approximation rate.

Figure 5 shows a reduction in the gap between the Monte Carlo quadrature approximation rate and the Leja point approximation rate as the spectral decay tail increases. This is to be expected, as a larger spectral tail reduces the smoothness of the function with respect to the graph. But this experiment shows that Leja point quadrature approximation rate does have a benefit over random sampling without the requirement that the target function be spectrally band-limited.

4.4 Real-World Graphs

Next, we illustrate our theory using a couple of real world graphs and functions. Since the graph is fixed in these examples and cannot be regenerated, we compute the Leja sequence for 1000 different random choices of a_0 , and average the errors. Similarly for the Monte Carlo comparison, we generate 1000 different sequences of sampled points.

In this first set of experiments, we examine the quadrature error for labels from the CORA data set [41]. The standard form of the data set comprises a digraph with 2708 publications as vertices, and edge from i to j means that paper i cited paper j . These publications are from seven areas of computer science and information theory, which

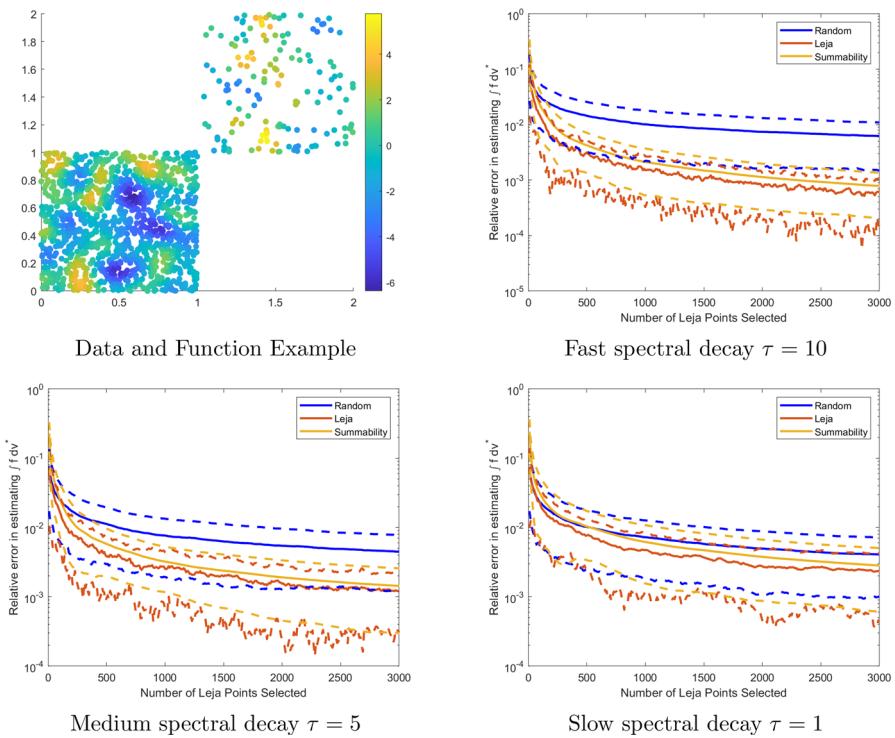


Fig. 5 Two cluster metric graph example with non-bandlimited function as in Eq. (4.6). We display the quadrature error results over 1000 instantiations of the random graph/function, and for varying τ . The dotted lines correspond to the confidence intervals around each mean curve of the same color

corresponds to the seven different classes in the data set. There are a large number of weakly connected components, many of which are singletons. The largest weakly connected component has 2485 nodes. We treat the citation links as undirected edges and construct a binary, symmetric adjacency matrix A , and take the largest connected component as the graph. This is equivalent to the largest weakly connected component, which results in an undirected graph with 2485 vertices.

Due to the sparsity of the number of edges in this graph, using the matrix G constructed directly from the graph results in a slow rate at which $\sum_k G(\cdot, a_k)$ changes as the number of Leja points increases. The process is expedited if we extend the connections between documents to two steps away by taking $B = A^2$. In this experiment in Fig. 6, we take the function f to be the indicator function of class i for $i = 1, \dots, 7$. These labels represent the field of the given document, and are related to the network due to the obvious fact that papers in the same field are more likely to cite one another.

In the second set of experiments, we examine the expected value approximation for positive association with various propositions in the Proposition dataset [43,50]. The November 2012 California ballot contained 11 initiatives, or propositions, on a variety of issues, including state taxation, corrections, and food labelling among others. The data consist of Twitter posts related to initiatives, grouped according to

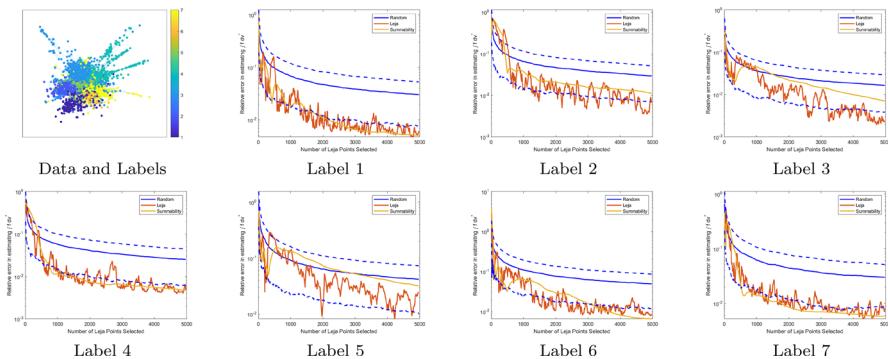


Fig. 6 Cora data set quadrature error for indicator function of various labels. The Leja and summability curves are averaged across 1000 different initializations. The dotted lines around the Monte Carlo curve are the confidence intervals

different propositions. For each proposition, the data is a directed graph with edge from i to j if the tweet originated from user i to user j . The authors of [43] have assigned an evaluation of the emotion of the sender with each tweet. We treat the edges as undirected and construct a binary, symmetric adjacency matrix A . We take the largest connected component of the graph for these experiments. Each proposition has its own graph, with a positive or negative label of sentiment on each node. The mean binary sentiment estimate could be used as a proxy for the number of people that would vote for/against the proposition.

Due to the sparsity of the number of edges in this graph, we extend the connections between documents to two steps away by taking $B = A^2$.

In this experiment, we take the function to be a binary indicator function (0 or 1) of whether the sentiment is positive. We examine two different graphs in Fig. 7, those for Proposition 30 and those for Proposition 37, as those are the graphs with the largest number of vertices, with 4436 vertices for Prop 30 and 8039 vertices for Prop 37.

As we can see in both the Cora experiments and the Proposition experiments, the Leja quadrature error performs better than the average Monte Carlo performance even for a small number of sampled points. We also wish to address the relative lack of smoothness of the quadrature approximations relative to the synthetic experiments and the Monte Carlo trials. First, recall that the labels here are incredibly non-smooth as they are binary. Second, we note that the Leja point sampling can be thought of as a distance sampling scheme, choosing consecutive points far apart from one another and points likely to have different labels. Because of these points, the approximation error will have a larger local variability than random sampling, even when averaged across multiple initializations. This is also a motivation for using the summability approximation, which applies a decaying weight to added points and does not demonstrate the same local fluctuations as the unweighted Leja point averages.

We also demonstrate that the location of the Leja points for each of these data sets, using ν^* to be an inverse density estimate, in Fig. 8.

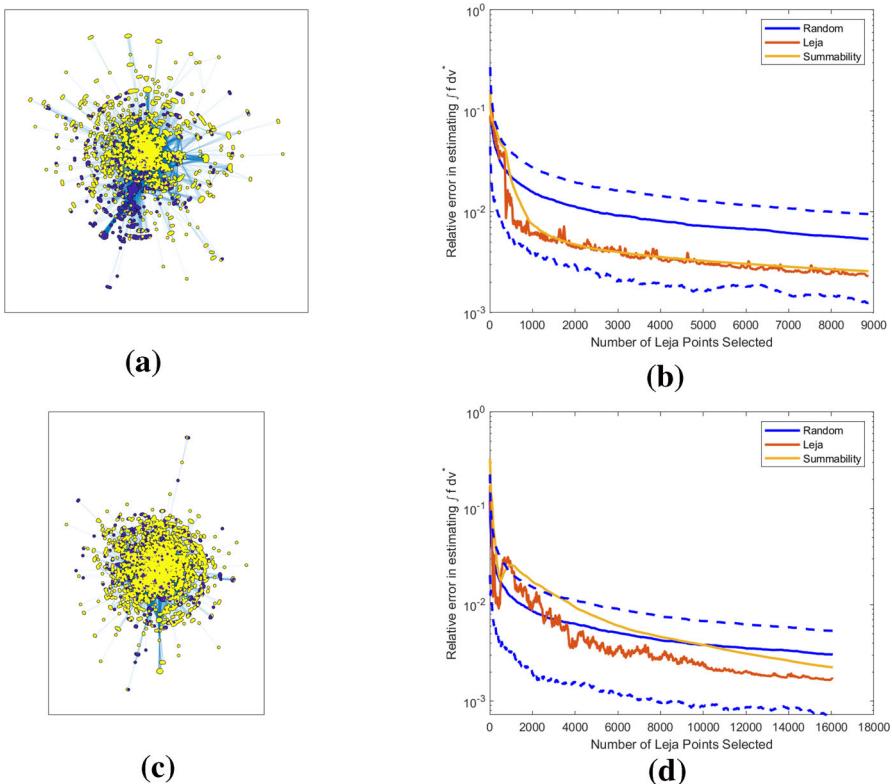


Fig. 7 Proposition data set quadrature error. **a–b** Proposition 30, **c–d** Proposition 37. The Leja and summability curves are averaged across 1000 different initializations. The dotted lines around the Monte Carlo curve are the confidence intervals

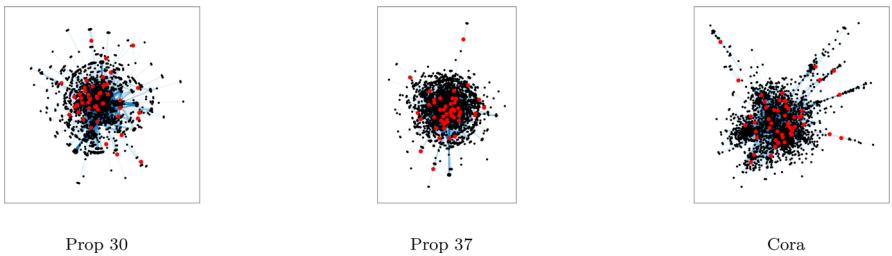


Fig. 8 Layout of 50 Leja points for various real-world data sets, with ν^* being the inverse of the kernel density estimate

4.5 Additional Comparisons

In this section, we compare our method to a natural alternative: computing a function interpolation from the sampled subset $W \subset V$ to the entire set of vertices V , and then integrating that estimate against ν^* . A recent method for this function interpolation is graph sampling set selection (GSSS) [39], which is built on taking a weighted linear

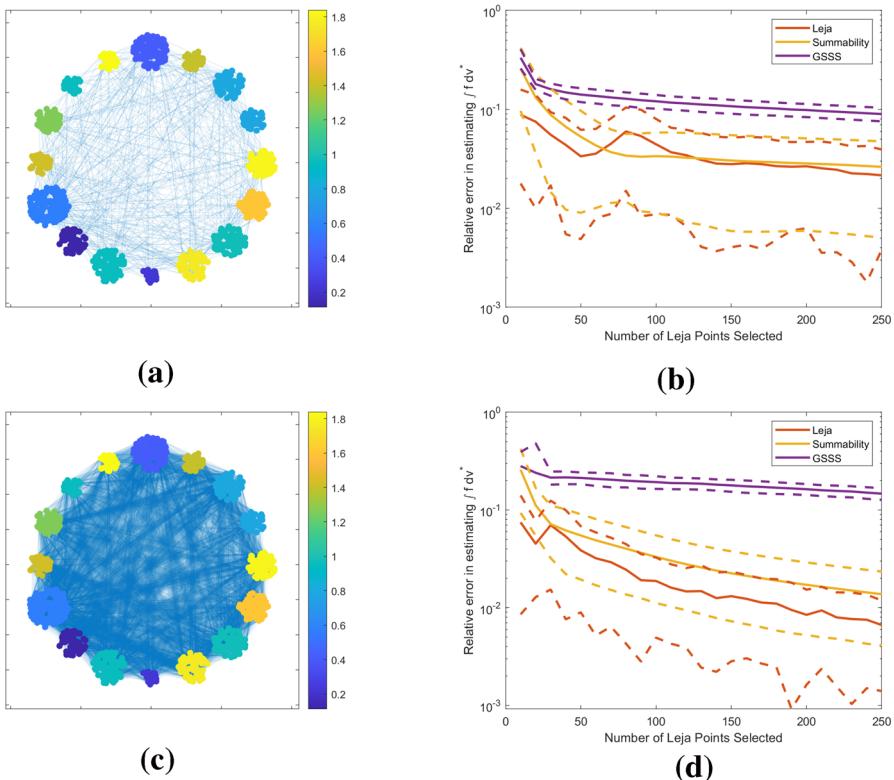


Fig. 9 a–b Adjacency of communities, c–d Adjacency matrix squared. We display both an example of the graph/function, and the quadrature error results over 1000 instantiations of the random graph. The dotted lines correspond to the confidence intervals around each mean curve of the same color

combination of columns of the spectrally filtered graph Laplacian. We compare to this method in Fig. 9 on the community graph data set proposed in [39]. We compare using both the original adjacency matrix A , and using the two-hop adjacency matrix A^2 , as the starting set of edges. We take v^* to be the inverse density of each community, normalized. Again, we compare across 1000 instantiations of the graph with $N = 1000$ points, and report the mean and standard deviation of the error at each fraction of points kept. The function being regressed in each instantiation is chosen to be constant on each community, with the value randomly chosen from a uniform distribution.

Finally, we demonstrate that the bound (2.13) on the quadrature error is independent of the number of data points for problems of fixed complexity. In this experiment, we return to the two cluster example from Fig. 3 while varying the number of data points. The graph created from this point cloud remains fixed, constructed with a Gaussian kernel of bandwidth $\sigma = 0.1$. In Fig. 10, we vary the number of data points and plot the quadrature error for a fixed number of Leja iterations. After an initial level-off as the number of points grows, the error is more or less constant. This shows that the

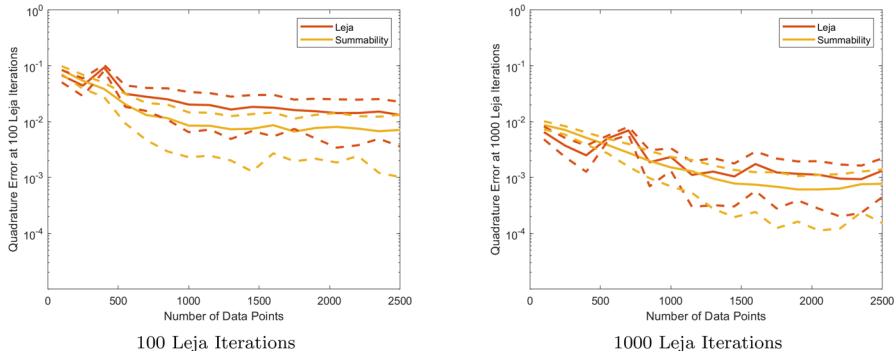


Fig. 10 Two cluster dataset with a fixed number of Leja points and a fixed Gaussian kernel bandwidth $\sigma = 0.1$. The quadrature error results are over 1000 instantiations of the random graph. The dotted lines correspond to the confidence intervals around each mean curve of the same color

error does not have a strong dependence on N for neighbor graphs built from a fixed sampling distribution.

5 Proofs

Proof of Lemma 2.1. Let v^* be an equilibrium measure. If $v \in \mathcal{P}$, then for all $t \in [0, 1]$, $(1-t)v^* + tv \in \mathcal{P}$, and therefore, $t = 0$ is a minimum for

$$\begin{aligned} f(t) &= G((1-t)v^* + tv, (1-t)v^* + tv) = (1-t)^2 G(v^*, v^*) \\ &\quad + t^2 G(v, v) + 2t(1-t)G(v, v^*). \end{aligned}$$

So, $f'(0) = -2G(v^*, v^*) + 2G(v, v^*) \geq 0$. This implies (2.5). Let $\epsilon > 0$ and, in this proof only, $L = \{x \in \mathbb{X} : G(x, v^*) > \Gamma(G) + \epsilon\}$. Then

$$\begin{aligned} \Gamma(G) &= G(v^*, v^*) = \int_L G(x, v^*) d\nu^* + \int_{\mathbb{X} \setminus L} G(x, v^*) d\nu^*(x) \\ &\geq v^*(L)(\Gamma(G) + \epsilon) + \Gamma(G)v^*(\mathbb{X} \setminus L) = \Gamma(G) + v^*(L)\epsilon. \end{aligned}$$

Therefore, $v^*(L) = 0$. Since \mathbb{X} is a finite set, this implies (2.6).

Next, let G be conditionally positive semi-definite, and μ satisfy the conditions of part (b). Then (2.8) shows that $G(\mu, \mu) = c$, and (2.7) shows that $G(v^*, \mu) \geq c$. So,

$$0 \leq G(v^* - \mu, v^* - \mu) = G(v^*, v^*) + G(\mu, \mu) - 2G(v^*, \mu) \leq G(v^*, v^*) - c.$$

Thus, $c = G(\mu, \mu) \leq G(v^*, v^*)$. Since v^* is an equilibrium measure, it follows that $c = G(\mu, \mu) = G(v^*, v^*) = \Gamma(G)$, so that μ is also an equilibrium measure. Also, $G(v^* - \mu, v^* - \mu) = 0$, so that $\mu - v^*$ is in the null space of G . \square

The next theorem shows in a standard manner than the sequence σ_n associated with the Leja points converge to the equilibrium measure. In the sequel, we write $M = \max_{x \in \mathbb{X}} G(x, x)$.

Theorem 5.1 *We have*

$$G(\sigma_n - v^*, \sigma_n - v^*) \leq \frac{1}{n} (M - G(v^*, v^*)), \quad n \in \mathbb{Z}_+. \quad (5.1)$$

Any weak-star limit of a subsequence of $\{\sigma_n\}$ is also an equilibrium measure.

Proof We have for $n \geq 2$,

$$\begin{aligned} G(\sigma_n, \sigma_n) &= \frac{1}{n^2} \sum_{j,k=0}^{n-1} G(v_j, v_k) = \frac{1}{n^2} \sum_{j=0}^{n-1} G(v_j, v_j) + \frac{2}{n^2} \sum_{0 \leq j < k \leq n-1} G(v_j, v_k) \\ &= \frac{1}{n^2} \sum_{j=0}^{n-1} G(v_j, v_j) + \frac{2}{n^2} \sum_{k=1}^{n-1} G\left(\sum_{j=0}^{k-1} v_j, v_k\right) \\ &= \frac{1}{n^2} \sum_{j=0}^{n-1} G(v_j, v_j) + \frac{2}{n^2} \sum_{k=1}^{n-1} k G(\sigma_k, v_k). \end{aligned} \quad (5.2)$$

In view of the definition (2.11) of Leja points,

$$G(\sigma_k, v_k) = G(v_k, \sigma_k) \leq G(x, \sigma_k), \quad x \in \mathbb{X}. \quad (5.3)$$

Integrating both sides with respect to v^* and using (2.9), we get for all $k \in \mathbb{Z}_+$,

$$G(\sigma_k, v_k) \leq G(\sigma_k, v^*) = G(v^*, v^*).$$

Therefore, (5.2) leads to

$$G(\sigma_n, \sigma_n) \leq M/n + \frac{2}{n^2} G(v^*, v^*) \sum_{k=1}^{n-1} k = G(v^*, v^*) + (M - G(v^*, v^*))/n. \quad (5.4)$$

In view of (2.9),

$$\begin{aligned} G(\sigma_n - v^*, \sigma_n - v^*) &= G(\sigma_n, \sigma_n) + G(v^*, v^*) - 2G(\sigma_n, v^*) \\ &= G(\sigma_n, \sigma_n) - G(v^*, v^*). \end{aligned}$$

Therefore, (5.4) implies (5.1). In turn, if $\{\sigma_n\}_{n \in \Lambda}$ be any subsequence of $\{\sigma_n\}$ converging weak-star to a probability measure σ , then (5.1) shows that $G(\sigma - v^*, \sigma - v^*) = 0$. Therefore, σ is also an equilibrium measure. \square

We note a couple of corollaries of the proof of Theorem 5.1.

Corollary 5.1 For $n \geq 2$,

$$\frac{n}{n-1}G(v^*, v^*) - M/(n-1) \leq \frac{2}{n(n-1)} \sum_{k=1}^{n-1} kG(v_k, \sigma_k) \leq G(v^*, v^*). \quad (5.5)$$

We formulate the next corollary as a theorem in its own right.

Theorem 5.2 Let $G(v^*, v^*) \geq 0$. Then for all $x \in \mathbb{X}$, $n \geq 2$,

$$G(v^*, v^*) - M/(n-1) \leq \frac{2}{n(n-1)} \sum_{k=1}^{n-1} kG(x, \sigma_k) \leq \frac{2\|G\|}{n} \leq G(v^*, v^*) + \frac{2\|G\|}{n}. \quad (5.6)$$

In particular,

$$\max_{x \in \mathbb{X}} \left| \frac{2}{n(n-1)} \sum_{j=0}^{n-1} (n-j)G(x, v_j) - G(x, v^*) \right| \leq \frac{2\|G\|}{n}. \quad (5.7)$$

Proof The first estimate in (5.6) follows from the first estimate in (5.5) and (5.3). The second estimate in (5.6) follows from the fact that

$$G(x, \sigma_k) = \frac{1}{k} \sum_{j=0}^{k-1} G(x, v_k) = \frac{1}{k} \sum_{j=0}^{k-1} G(x, a_k) \leq \frac{1}{k} \sum_{j=0}^{k-1} |G(x, a_k)| \leq \frac{2\|G\|}{k}. \quad (5.8)$$

The final inequality in (5.6) following from the fact that $G(v^*, v^*) \geq 0$. We observe now that

$$\sum_{k=1}^{n-1} kG(x, \sigma_k) = \sum_{k=0}^{n-1} kG(x, \sigma_k) = \sum_{k=0}^{n-1} \sum_{j=0}^{n-1} G(x, v_j) = \sum_{j=0}^{n-1} (n-j)G(x, v_j). \quad (5.9)$$

Therefore, (5.7) follows from (5.6) and the fact that $G(x, v^*) = G(v^*, v^*)$ for all $x \in \mathbb{X}$. \square

Theorem 5.3 If $G(x, y) \geq 0$ for all $x, y \in \mathbb{X}$, we have for $n \geq 2$,

$$\max_{x \in \mathbb{X}} |G(x, v^*) - G(x, \sigma_n)| \leq \frac{3\|G\|}{n+1}. \quad (5.10)$$

Proof For $\ell \geq 2$, we have

$$\begin{aligned}
 \ell G(v_\ell, \sigma_\ell) &= \sum_{j=0}^{\ell-1} G(a_\ell, a_j) = G(a_\ell, a_{\ell-1}) + \sum_{j=0}^{\ell-2} G(a_\ell, a_j) \\
 &= G(a_\ell, a_{\ell-1}) + (\ell - 1)G(v_\ell, \sigma_{\ell-1}) \\
 &\geq (\ell - 1)G(v_{\ell-1}, \sigma_{\ell-1}) \geq \ell G(v_{\ell-1}, \sigma_{\ell-1}) - \frac{1}{\ell - 1} \sum_{j=0}^{\ell-2} G(a_{\ell-1}, a_j) \\
 &\geq \ell G(v_{\ell-1}, \sigma_{\ell-1}) - \frac{\|G\|}{\ell - 1};
 \end{aligned}$$

i.e.,

$$G(v_{\ell-1}, \sigma_{\ell-1}) - G(v_\ell, \sigma_\ell) \leq \|G\| \left(\frac{1}{\ell - 1} - \frac{1}{\ell} \right).$$

Therefore, for any k with $2 \leq k \leq n$, a summation in the above inequality leads to

$$G(v_k, \sigma_k) - G(v_n, \sigma_n) \leq \|G\| \frac{n - k}{nk}.$$

Since $G(v^*, v^*) \geq 0$, Corollary 5.1 now shows that

$$\begin{aligned}
 \frac{n(n+1)}{2} \left(G(v^*, v^*) - \frac{M}{n} \right) &\leq \sum_{k=1}^n k G(v_k, \sigma_k) \leq \frac{n(n+1)}{2} G(v_n, \sigma_n) \\
 &\quad + G(v_1, \sigma_1) + \frac{\|G\|}{n} \sum_{k=2}^n (n - k) \\
 &= \frac{n(n+1)}{2} G(v_n, \sigma_n) \\
 &\quad + G(v_1, \sigma_1) + \frac{\|G\| (n-1)(n-2)}{n}.
 \end{aligned}$$

Rearranging,

$$G(v^*, v^*) - G(v_n, \sigma_n) \leq \frac{M}{n} + \frac{2G(v_1, \sigma_1)}{n(n+1)} + \frac{\|G\| (n-1)(n-2)}{n(n+1)}.$$

For $n \geq 2$, $2n + (n-1)(n-2) \leq n^2$, and we deduce that

$$G(v^*, v^*) - G(v_n, \sigma_n) \leq \frac{2M + \|G\|}{n+1} \leq \frac{3\|G\|}{n+1}.$$

For any $x \in \mathbb{X}$, we have $G(x, v^*) = G(v^*, v^*)$ and $G(v_n, \sigma_n) \leq G(x, \sigma_n)$. Therefore,

$$G(x, v^*) - G(x, \sigma_n) \leq \frac{3\|G\|}{n+1}.$$

In the reverse direction, we recall (5.8). Together with (5.6), this leads to (5.10). \square

Proof of Theorem 2.1. The estimate (2.13) follows easily from (5.10) and the observation that

$$\begin{aligned} \int_{\mathbb{X}} f d\nu^* - \frac{1}{n} \sum_{k=0}^{n-1} f(a_k) &= \int_{\mathbb{X}} f d\nu^* - \int_{\mathbb{X}} f d\sigma_n \\ &= \int_{\mathbb{X}} \left\{ \int_{\mathbb{X}} G(x, y) d\nu^*(y) - \int_{\mathbb{X}} G(x, y) d\sigma_n(y) \right\} d\mathcal{D}_G(f)(x). \end{aligned}$$

\square

Acknowledgements The work of AC was supported in part by NSF DMS Grants 2012266, 1819222, and Sage Foundation Grant 2196. The work of HNM is supported in part NSF DMS Grant 2012355 and ARO Grant W911NF2110218. We thank Professors Percus at Claremont Graduate University for his help in securing the Proposition data set, which was sent to us by Dr. Linhong Zhu at USC Information Sciences Institute in Marina Del Ray, California.

References

1. Andrievskii, V.V., Blatt, H.-P.: Discrepancy of Signed Measures and Polynomial Approximation. Springer Science & Business Media, New York (2013)
2. Anis, A., Gadde, A., Ortega, A.: Efficient sampling set selection for bandlimited graph signals using graph spectral proxies. *IEEE Trans. Signal Process.* **64**(14), 3775–3789 (2016)
3. Bermanis, A., Averbuch, A., Coifman, R.R.: Multiscale data sampling and function extension. *Appl. Comput. Harmon. Anal.* **34**(1), 15–29 (2013)
4. Blatt, H.-P.: On the distribution of simple zeros of polynomials. *J. Approx. Theory* **69**(3), 250–268 (1992)
5. Blatt, H.-P., Mhaskar, H.N.: A general discrepancy theorem. *Arkiv för Matematik* **31**(2), 219–246 (1993)
6. Brown, L.: Sequences of well-distributed vertices on graphs and spectral bounds on optimal transport. *J. Fourier Anal. Appl.* **27**(2), 1–27 (2021)
7. Brualdi, R.A.: The DAD theorem for arbitrary row sums. *Proc. Am. Math. Soc.* **45**(2), 189–194 (1974)
8. Chen, S., Varma, R., Sandryhaila, A., Kovačević, J.: Discrete signal processing on graphs: sampling theory. *IEEE Trans. Signal Process.* **63**(24), 6510–6523 (2015)
9. Chung, F.R.K.: Spectral Graph Theory, vol. 92. American Mathematical Soc., Providence (1997)
10. Cloninger, A., Roy, B., Riley, C., Krumholz, H.M.: People mover’s distance: class level geometry using fast pairwise data adaptive transportation costs. *Appl. Comput. Harmon. Anal.* **47**(1), 248–257 (2019)
11. De Marchi, S.: On Leja sequences: some results and applications. *Appl. Math. Comput.* **152**(3), 621–647 (2004)
12. Dick, J., Pillichshammer, F.: Digital Nets and Sequences: Discrepancy Theory and Quasi-Monte Carlo Integration. Cambridge University Press, Cambridge (2010)
13. Dwivedi, R., Feldheim, O.N., Gurel-Gurevich, O., Ramdas, A.: The power of online thinning in reducing discrepancy. *Probab. Theory Relat. Fields* **174**(1–2), 103–131 (2019)
14. Erdős, P., Turán, P.: On the uniformly-dense distribution of certain sequences of points. *Ann. Math.* **41**, 162–173 (1940)

15. Fruchterman, T., Reingold, E.: Graph drawing by force-directed placement. *Software* **21**(11), 1129–1164 (1991)
16. Fuglede, B.: On the theory of potentials in locally compact spaces. *Acta Mathematica* **103**(3–4), 139–215 (1960)
17. Gayo-Avello, D.: A meta-analysis of state-of-the-art electoral prediction from twitter data. *Soc. Sci. Comput. Rev.* **31**(6), 649–679 (2013)
18. Götz, M.: On the distribution of Leja-Górski points. *J. Comput. Anal. Appl.* **3**(3), 223–241 (2001)
19. Jin, L., Chen, Y., Hui, P., Ding, C., Wang, T., Vasilakos, A. V., Deng, B., Li, X.: Albatross sampling: robust and effective hybrid vertex sampling for social graphs. In: Proceedings of the 3rd ACM international workshop on MobiArch, pp. 11–16 (2011)
20. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: Proceedings of ICLR (2017)
21. Knight, P.A.: The sinkhorn-knopp algorithm: convergence and applications. *SIAM J. Matrix Anal. Appl.* **30**(1), 261–275 (2008)
22. Krause, A., Leskovec, J., Guestrin, C., VanBriesen, J., Faloutsos, C.: Efficient sensor placement optimization for securing large water distribution networks. *J. Water Resour. Plan. Manag.* **134**(6), 516–526 (2008)
23. Kuipers, L., Niederreiter, H.: Uniform Distribution of Sequences. Courier Corporation. Wiley, New York (2012)
24. Landa, B., Coifman, R.R., Kluger, Y.: Doubly stochastic normalization of the gaussian kernel is robust to heteroskedastic noise. *SIAM J. Math. Data Sci.* **3**(1), 388–413 (2021)
25. Landkof, N.S.: Foundations of Modern Potential Theory, vol. 180. Springer, New York (1972)
26. Leja, F.: Sur certaines suites liées aux ensembles plans et leur application à la représentation conforme. *Annales Polonici Mathematici* **1**(4), 8–13 (1957)
27. Linderman, G., Steinerberger, S.: Numerical integration on graphs: where to sample and how to weigh. *Math. Comput.* **89**(324), 1933–1952 (2020)
28. Lu, J., Sachs, M., Steinerberger, S.: Quadrature points via heat kernel repulsion. *Construct. Approx.* **51**(1), 27–48 (2020)
29. Marshall, N.F., Coifman, R.R.: Manifold learning with bi-stochastic kernels. *IMA J. Appl. Math.* **84**(3), 455–482 (2019)
30. Mhaskar, H.N.: Weighted polynomials, radial basis functions and potentials on locally compact spaces. *Numer. Funct. Anal. Optim.* **11**(9–10), 987–1017 (1990)
31. Mhaskar, H.N.: On the tractability of multivariate integration and approximation by neural networks. *J. Complex.* **20**(4), 561–590 (2004)
32. Mhaskar, H.N.: Dimension independent bounds for general shallow networks. *Neural Netw.* **123**, 142–152 (2020)
33. Nazi, A., Zhou, Z., Thirumuruganathan, S., Zhang, N., Das, G.: Walk, not wait: faster sampling over online social networks. *Proc. VLDB Endow.* **8**(6), 678–689 (2015)
34. Ohtsuka, M., et al.: On potentials in locally compact spaces. *J. Sci. Hiroshima Univ. Ser. AI (Math.)* **25**(2), 135–352 (1961)
35. Page, L., Brin, S., Motwani, R., Winograd, T.: The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab (1999)
36. Pesenson, I.: Sampling in paley-wiener spaces on combinatorial graphs. *Trans. Am. Math. Soc.* **360**(10), 5603–5627 (2008)
37. Pritsker, I.E.: Equidistribution of points via energy. *Arkiv för Matematik* **49**(1), 149–173 (2011)
38. Puy, G., Tremblay, N., Gribonval, R., Vandergheynst, P.: Random sampling of bandlimited signals on graphs. *Appl. Comput. Harmon. Anal.* **44**(2), 446–475 (2018)
39. Sakiyama, A., Tanaka, Y., Tanaka, T., Ortega, A.: Eigendecomposition-free sampling set selection for graph signals. *IEEE Trans. Signal Process.* **67**(10), 2679–2692 (2019)
40. Sears, L.E., Agrawal, S., Sidney, J.A., Castle, P.H., Rula, E.Y., Coberley, C.R., Witters, D., Pope, J.E., Harter, J.K.: The well-being 5: development and validation of a diagnostic instrument to improve population well-being. *Popul. Health Manag.* **17**(6), 357–365 (2014)
41. Sen, P., Namata, G., Bilgic, M., Getoor, L., Galligher, B., Eliassi-Rad, T.: Collective classification in network data. *AI Mag.* **29**(3), 93–93 (2008)
42. Sinkhorn, R., Knopp, P.: Concerning nonnegative matrices and doubly stochastic matrices. *Pac. J. Math.* **21**(2), 343–348 (1967)

43. Smith, L.M., Zhu, L., Lerman, K., Kozareva, Z.: The role of social media in the discussion of controversial topics. In: Proceedings of the 2013 International Conference on Social Computing, pp. 236–243. IEEE (2013)
44. Steinerberger, S.: Generalized designs on graphs: sampling, spectra, symmetries. *J. Graph Theory* **93**(2), 253–267 (2020)
45. Tanaka, Y., Eldar, Y.C., Ortega, A., Cheung, G.: Sampling signals on graphs: from theory to applications. *IEEE Signal Process. Mag.* **37**(6), 14–30 (2020)
46. Vahidian, S., Mirzasoleiman, B., Cloninger, A.: Coresets for estimating means and mean square error with limited greedy samples. In: Proceedings of the Conference on Uncertainty in Artificial Intelligence, pp. 350–359. PMLR (2020)
47. Wang, X., Hickernell, F.J.: Randomized halton sequences. *Math. Comput. Modell.* **32**(7–8), 887–899 (2000)
48. Watts, D.J., Strogatz, S.H.: Collective dynamics of “small-world” networks. *Nature* **393**(6684), 440–442 (1998)
49. Wu, Y., Xu, Y., Singh, A., Yang, Y., Dubrawski, A.: Active learning for graph neural networks via node feature propagation. [arXiv:1910.07567](https://arxiv.org/abs/1910.07567) (2019)
50. Zhu, L., Galstyan, A., Cheng, J., Lerman, K.: Tripartite graph clustering for dynamic sentiment analysis on social media. In: Proceedings of the 2014 ACM SIGMOD international conference on Management of data, pp. 1531–1542 (2014)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.