

Multiscale estimates for the condition number of non-harmonic Fourier matrices

Weilin Li*

October 1, 2024

Abstract

This paper studies the extreme singular values of non-harmonic Fourier matrices. Such a matrix of size $m \times s$ can be written as $\Phi = [e^{-2\pi i j x_k}]_{j=0,1,\dots,m-1, k=1,2,\dots,s}$ for some set $\mathcal{X} = \{x_k\}_{k=1}^s$. Its condition number controls the stability of inversion, which is of great importance to super-resolution and nonuniform Fourier transforms. Under the assumption $m \geq 6s$ and without any restrictions on \mathcal{X} , the main theorems provide explicit lower bounds for the smallest singular value $\sigma_s(\Phi)$ in terms of distances between elements in \mathcal{X} . More specifically, distances exceeding an appropriate scale τ have modest influence on $\sigma_s(\Phi)$, while the product of distances that are less than τ dominates the behavior of $\sigma_s(\Phi)$. These estimates reveal how the multiscale structure of \mathcal{X} affects the condition number of Fourier matrices. Theoretical and numerical comparisons indicate that the main theorems significantly improve upon classical bounds and recover the same rate for special cases but with relaxed assumptions.

2020 Math Subject Classification: 15A12, 15A60, 42A05, 42A15, 65F22

Keywords: Fourier matrix, singular values, trigonometric interpolation, density, local sparsity

1 Introduction

1.1 Motivation

For any set $\mathcal{X} = \{x_k\}_{k=1}^s \subseteq \mathbb{T} := \mathbb{R}/\mathbb{Z}$ and natural number $m \geq s$, a (*non-harmonic*) *Fourier matrix* of size $m \times s$ is defined as

$$\Phi := \Phi(m, \mathcal{X}) := \left[e^{-2\pi i j x_k} \right]_{j=0,1,\dots,m-1, k=1,2,\dots,s}.$$

This definition generalizes the discrete Fourier transform matrix, whereby $m = s$ and \mathcal{X} consist of s equally spaced points in \mathbb{T} . Throughout the expository portions of this paper, we will implicitly assume that $|\mathcal{X}| = s$, and we impose $m \geq s > 1$ to avoid trivialities.

Fourier matrices are classical objects that appear in numerous areas of mathematics. They provide a fundamental connection between linear algebra and trigonometric interpolation, which can be traced back to the work of Newton and Lagrange. They are matrix representations of the Fourier transform, so they naturally appear in the analysis of Fourier series [43], exponential sums [41], and nonuniform Fourier transforms [19]. Since Φ is also a Vandermonde matrix, it has full rank whenever $m \geq s$. When the rows of Φ are viewed as elements of \mathbb{C}^s , then the squared extreme singular values are the upper and lower frame constants [12, 18].

*CUNY City College. Email: wli6@ccny.cuny.edu

For numerical applications, we require quantitative estimates for the extreme singular values of Φ to ensure that it can be inverted in numerical schemes without incurring significant error. We provide two motivational examples. The first is a simplified model for one-dimensional super-resolution, whereby \mathcal{X} represents the support of a discrete measure and Φ is the matrix representation of the Fourier transform of the measure sampled at m consecutive integers. The problem is to recover \mathcal{X} through noisy Fourier data, and even though this is a nonlinear inverse problem, the condition number of Φ is closely related to performance of algorithms [27, 28, 29] and fundamental limits of recovery [17, 16, 27, 6]. The second is a connection to the nonuniform discrete Fourier transform of Type I. Here, Φ is the NUDFT Type I matrix that evaluates the first m Fourier series coefficients of the measure on \mathcal{X} for the periodic domain \mathbb{T} , see [19, 35] for an overview. This is an over-determined linear system, so the condition number of Φ is pertinent to the least squares solution of fitting data by exponential sums.

While classical papers such as [21, 15, 9] concentrated on square matrices, tall ones, where m may be significantly larger than s , tend to be better conditioned. Tall Fourier matrices have received considerable interest recently [30, 31, 30, 1, 27, 3, 5, 22, 23]. This is partly due to modern applications in signal and image processing, where rectangular matrices appear more frequently, since m represents the number of measurements or parameters, while s corresponds to the number of constraints, see [17, 20, 8, 28, 29, 14] and references therein. It is worth noting that parallel to this line of research, the approximation properties of trigonometric interpolation in the $m \geq s$ regime has received interest [25, 40, 36] due to connections with over-parameterization in machine learning.

The condition number $\kappa(\Phi(m, \mathcal{X}))$ greatly depends on the “Rayleigh length” $\frac{1}{m}$ versus the “geometry” of \mathcal{X} . The latter can be partially described by the *minimum separation* of \mathcal{X} , defined as

$$\Delta(\mathcal{X}) := \min_{j \neq k} |x_j - x_k|_{\mathbb{T}}, \quad \text{where} \quad |t|_{\mathbb{T}} := \min_{n \in \mathbb{Z}} |t - n|.$$

Letting $\sigma_k(\Phi)$ denote the k -th largest singular value of Φ , it was shown in [1] that

$$\text{if } \Delta(\mathcal{X}) > \frac{1}{m}, \quad \text{then} \quad \sqrt{m - \frac{1}{\Delta(\mathcal{X})}} \leq \sigma_s(\Phi(m, \mathcal{X})) \leq \sigma_1(\Phi(m, \mathcal{X})) \leq \sqrt{m + \frac{1}{\Delta(\mathcal{X})}}. \quad (1.1)$$

The intuition behind this inequality is that the columns of Φ are almost orthogonal. This result and a closely related one in [31], are proved using analytic number theory methods.

On the other hand, when $\Delta(\mathcal{X}) \leq \frac{1}{m}$, simple numerical experiments, see [27, 3], show that $\sigma_s(\Phi(m, \mathcal{X}))$ does not follow the behavior in (1.1). This makes intuitive sense since if $|x_k - x_{\ell}|_{\mathbb{T}}$ is small, then the k -th and ℓ -th columns of Φ are highly correlated, which results in a large condition number. If it is significantly larger than \sqrt{s} , then the smallest singular value is the culprit because we have the trivial estimate $\sigma_1(\Phi) \leq \|\Phi\|_F = \sqrt{ms}$, where $\|\cdot\|_F$ denotes the Frobenius norm.

Accurate bounds for the smallest singular value have been obtained under specific scenarios, namely when \mathcal{X} can be partitioned in subsets called “clumps”, where each clump is contained in an interval whose length is on the order of $\frac{1}{m}$. In contrast to the separation condition required in (1.1), without any conditions on $\Delta(\mathcal{X})$ relative to m , the results in [27, 3, 22, 5, 4] roughly state that if each clump has cardinality at most λ and the clumps are sufficiently far away from each other, then

$$c(\lambda, s)\sqrt{m} (m\Delta(\mathcal{X}))^{\lambda-1} \leq \sigma_s(\Phi(m, \mathcal{X})) \leq C(\lambda, s)\sqrt{m} (m\Delta(\mathcal{X}))^{\lambda-1}. \quad (1.2)$$

Since the exponent $\lambda - 1$ may be significantly smaller than $s - 1$, this bound shows that $\sigma_s(\Phi)$ depends on the local geometry of \mathcal{X} . It captures the intuition that columns of Φ which correspond

to different clumps are almost orthogonal with respect to each other, so we expect the conditioning of Φ to only depend on each clump separately.

1.2 A motivational multiscale example

To better illustrate the limitations of prior work, let us consider a typical set with multiscale structure such as

$$\begin{aligned} \mathcal{X} &:= \mathcal{X}_1 \cup \mathcal{X}_2 \cup \mathcal{X}_3, \quad \text{where} \\ \mathcal{X}_1 &:= \{0, \frac{1}{90}, \frac{2}{90}, \frac{3}{90}\}, \quad \mathcal{X}_2 = \frac{1}{3} + \{0, \frac{1}{200}, \frac{2}{200}\}, \quad \text{and} \quad \mathcal{X}_3 = \frac{2}{3} + \{0, \frac{1}{500}\}. \end{aligned} \quad (1.3)$$

We have defined in \mathcal{X} this way to emphasize that its three disjoint subsets \mathcal{X}_1 , \mathcal{X}_2 , and \mathcal{X}_3 each have significantly different minimum separations and should be treated as sets with completely different scales. The set \mathcal{X} is shown in Figure 1. Our problem is to determine behavior of $\sigma_s(\Phi(m, \mathcal{X}))$ as a function of m . Clearly $\Phi(m, \mathcal{X})$ only has full rank when $m \geq s = 9$ and (1.1) kicks in when $m > \frac{1}{\Delta(\mathcal{X})} = 500$. What about the missing range $m \in [9, 500] \cap \mathbb{N}$?

For the range $m \in [9, 600] \cap \mathbb{N}$, none of the bounds in [27, 3, 22, 5] are applicable. The reason is that while \mathcal{X} consists of three “clumps” \mathcal{X}_1 , \mathcal{X}_2 , and \mathcal{X}_3 , they are too close to each other and do not satisfy these theorems’ assumptions, or such theorems have implicit constants in their separation criterion that cannot be explicitly determined. It is important to remark that the aforementioned papers concentrated on the super-resolution limit, whereby either m is sufficiently large and there is a sequence of \mathcal{X}_m for which $\Delta(\mathcal{X}_m) \rightarrow 0$, or alternatively, $m \rightarrow \infty$ and $\Delta(\mathcal{X}_m) \rightarrow 0$ with some relationship between m and $\Delta(\mathcal{X}_m)$. Hence, it is not that surprising they cannot be directly used for fixed \mathcal{X} and m .

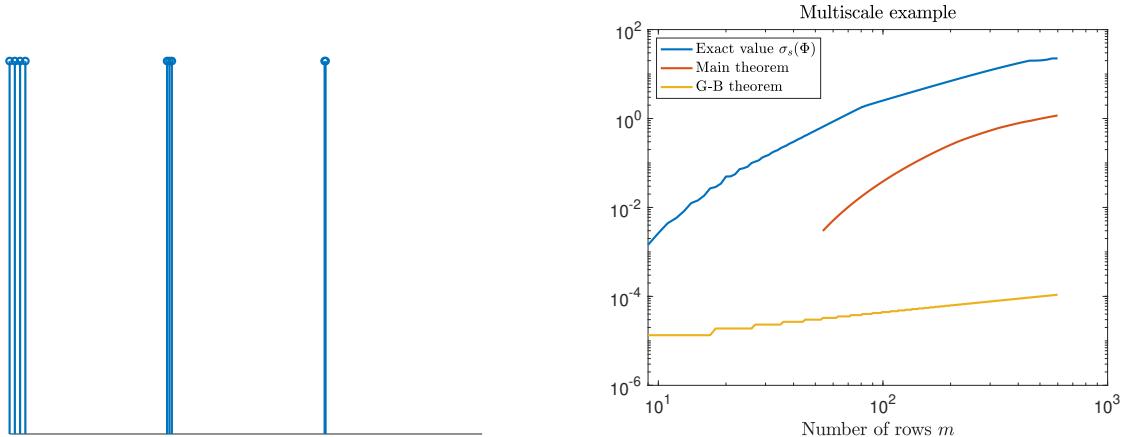


Figure 1: Left: Plot of \mathcal{X} defined in (1.3). Right: Plot of $\sigma_s(\Phi(m, \mathcal{X}))$ and two different lower bounds as functions of m for \mathcal{X} defined in (1.3).

In contrast, this experiment is for fixed \mathcal{X} and variable m within a finite range. Determining $\sigma_s(\Phi(m, \mathcal{X}))$ is naturally a discrete problem – there are no large or small parameters to exploit. We are only aware of one prior work that applies to this example. It results from combining Gautschi [21] and Bazán [7] to obtain

$$\sigma_s(\Phi(m, \mathcal{X})) \geq \sqrt{\frac{\lfloor \frac{m}{s} \rfloor}{s}} \min_{1 \leq k \leq s} \left\{ \prod_{j \neq k} \frac{|e^{2\pi i x} - e^{2\pi i x_k}|}{2} \right\}. \quad (1.4)$$

We refer to this as the Gautschi-Bazán theorem, and will provide more details on its derivation in the comparisons section. Aside from this inequality, a main result of this paper, Theorem 1, is applicable to this discrete problem, but with a mild restriction that $m \geq 6s$. The true $\sigma_s(\Phi(m, \mathcal{X}))$, our main theorem, and the Gautschi-Bazán bound are displayed in Figure 1. We see that our theorem offers a substantial improvement and better captures the true behavior. Just to highlight a dramatic improvement offered by Theorem 1, when $m = 400$, our main theorem underestimates the true value by a multiplicative factor of 21.0038, whereas the Gautschi-Bazán bound is off by a factor of 1.9687e+05.

2 Main results

The goal of this paper is to provide explicit, interpretable, and accurate bounds for $\sigma_s(\Phi(m, \mathcal{X}))$ for arbitrary \mathcal{X} when $\Delta(\mathcal{X}) < \frac{1}{m}$. Doing so is a tricky balancing act. We require conditions on \mathcal{X} that are not too restrictive, yet are sufficiently informative enough that a resulting lower bound is not too loose. We avoid restrictive assumptions by working with general geometric notions.

Definition 2.1. Let $\tau \in (0, \frac{1}{2}]$ and $\mathcal{X} \subseteq \mathbb{T}$ be a finite set. The τ local sparsity of \mathcal{X} is

$$\nu(\tau, \mathcal{X}) := \max_{x \in \mathcal{X}} |\{x' \in \mathcal{X}: |x - x'|_{\mathbb{T}} \leq \tau\}| \quad \text{and} \quad \nu(\tau, \emptyset) := 0.$$

The τ local sparsity is the maximum number of elements in \mathcal{X} contained with a τ neighborhood of any $x \in \mathcal{X}$. By definition, $\nu(\tau, \mathcal{X}) = 1$ if and only if $\Delta(\mathcal{X}) > \tau$. Importantly, we have $\nu(\tau, \mathcal{U}) \leq \nu(\tau, \mathcal{X})$ whenever $\mathcal{U} \subseteq \mathcal{X}$. If $\tau = \frac{1}{2}$, then $\nu(\tau, \mathcal{X}) = |\mathcal{X}| = s$, but it is possible that the local sparsity is significantly smaller than s .

Definition 2.2. For any $m \in \mathbb{N}_+$ and $\tau \in (0, \frac{1}{2}]$, we say a finite set $\mathcal{X} \subseteq \mathbb{T}$ satisfies the (m, τ) density criterion if

$$\frac{3\nu(\tau, \mathcal{X})}{\tau} \leq m.$$

We call it the density criterion because $\frac{\nu(\tau, \mathcal{X})}{2\tau}$ can be interpreted as the density of \mathcal{X} at scale 2τ , so the assumption asserts that it cannot be bigger than $\frac{m}{6}$. This criterion is not difficult to fulfill for some τ . Indeed, if we assume $m \geq 6s$ and select $\tau = \frac{1}{2}$, then $\frac{3\nu(\tau, \mathcal{X})}{\tau} = 6s \leq m$, so the density criterion is satisfied. However, if \mathcal{X} satisfies the (m, τ) density criterion, then there may be infinitely many other τ that are also valid, and the choice of τ will influence the below estimates.

We are almost ready to present our first main result. When interpreting the expressions in this paper, we use the standard convention that the product over an empty set is defined as 1. To simplify some of the notation that will appear in this paper, we define the subsets,

$$\mathcal{B}(x, \tau, \mathcal{X}) := \{x' \in \mathcal{X}: |x - x'|_{\mathbb{T}} \leq \tau\}, \quad \text{and} \quad \mathcal{G}(x, \tau, \mathcal{X}) := \{x' \in \mathcal{X}: |x - x'|_{\mathbb{T}} > \tau\}.$$

We will refer to these as the “bad” and “good” sets respectively, and this terminology will make sense later. We define a special function $\phi: [1, \infty) \rightarrow [1, \infty)$ by

$$\phi(t) := \frac{t}{\lfloor t \rfloor}. \tag{2.1}$$

This function appears in several bounds since our methods depend on number theoretic properties of several quantities. Note that $\phi(t) \leq \min\{2, 1 + \frac{1}{t-1}\}$, so in particular, $\phi(t) \rightarrow 1$ as $t \rightarrow \infty$. Terms that involve $\phi(t)$ are inconsequential when t is large.

Theorem 1. Let $m, s \in \mathbb{N}_+$ such that $s \geq 2$ and $m \geq 6s$. Suppose $\mathcal{X} = \{x_k\}_{k=1}^s \subseteq \mathbb{T}$ and pick any τ such that \mathcal{X} satisfies the (m, τ) density criterion. For each $k = 1, 2, \dots, s$, define

$$\mathcal{G}_k := \mathcal{G}(x_k, \tau, \mathcal{X}), \quad \alpha_k := \frac{|\mathcal{B}(x_k, \tau, \mathcal{X})|}{2m - 4\nu(\tau, \mathcal{G}_k)\tau^{-1}},$$

and the subsets

$$\mathcal{I}_k := \{x \in \mathcal{X} : 0 < |x - x_k| \leq \alpha_k\} \quad \text{and} \quad \mathcal{J}_k := \{x \in \mathcal{X} : \alpha_k < |x - x_k| \leq \tau\}.$$

Then we have

$$\frac{1}{\sigma_s^2(\Phi(m, \mathcal{X}))} \leq \sum_{k=1}^s 2^{\nu(\tau, \mathcal{G}_k)} \frac{2\alpha_k}{1 - 2\alpha_k} \prod_{x \in \mathcal{J}_k} \phi\left(\frac{1}{2|x - x_k|_{\mathbb{T}}}\right)^2 \prod_{x \in \mathcal{I}_k} \frac{\alpha_k^2}{(1 - 2\alpha_k)^2 |x - x_k|_{\mathbb{T}}^2}, \quad (2.2)$$

and in particular,

$$\sigma_s(\Phi(m, \mathcal{X})) \geq \min_{1 \leq k \leq s} \left\{ \frac{1}{\sqrt{4s\alpha_k}} \frac{1}{\sqrt{2^{\nu(\tau, \mathcal{G}_k)}}} \frac{1}{2^{|\mathcal{J}_k|}} \prod_{x \in \mathcal{I}_k} \frac{|x - x_k|_{\mathbb{T}}}{2\alpha_k} \right\}. \quad (2.3)$$

Both inequalities in this theorem provide “multiscale” lower bounds for $\sigma_s(\Phi)$. Let us explain what this terminology means in the context of (2.3), which is a simplified version of (2.2). Here, we fix a reference point $x_k \in \mathcal{X}$. The ranges $(\tau, \frac{1}{2}]$, $(\alpha_k, \tau]$, and $(0, \alpha_k]$ consist of the coarse, intermediate, and small scales respectively, and the elements in \mathcal{X} whose distance to x_k lie in these three ranges are \mathcal{G}_k , \mathcal{J}_k , and \mathcal{I}_k , respectively. An example is shown in Figure 2.

- (a) The set \mathcal{G}_k contributes a factor of $2^{-\nu(\tau, \mathcal{G}_k)/2}$ to the smallest singular value. Note that $\nu(\tau, \mathcal{G}_k)$ may be significantly smaller than $s-1$ for many types of \mathcal{X} , including the motivational example. Hence, points in \mathcal{X} that are τ away from x_k typically have little influence on the lower bound.
- (b) The set \mathcal{J}_k contributes a factor of $2^{-|\mathcal{J}_k|}$, so each element in \mathcal{J}_k contributes a multiplicative factor of $\frac{1}{2}$ to the lower bound. Note that $\alpha_k \leq \frac{\tau}{2}$ due to the (m, τ) density criterion, so α_k is naturally the next smaller scale following τ .
- (c) The set \mathcal{I}_k has the greatest amount of influence on the lower bound. Notice that each term inside the product in (2.3) is at most $\frac{1}{2}$, but may be significantly smaller depending on the structure of \mathcal{X} near x_k . For instance, if we let $s_{k,\ell} \in \mathbb{N}$ denote the number of elements in \mathcal{X} whose distance to x_k lies in $(2^{-\ell-1}\alpha_k, 2^{-\ell}\alpha_k]$, then

$$\log_2 \left(\prod_{x \in \mathcal{I}_k} \frac{|x - x_k|_{\mathbb{T}}}{2\alpha_k} \right) \asymp - \sum_{\ell=0}^{\infty} s_{k,\ell} \ell.$$

This illustrates that the product term is not equally influenced by all distances less than α_k and that it actually depends on the structure of \mathcal{X} near x_k at infinitely many finer scales.

If there is a $\tau \ll \frac{1}{2}$ for which the density criterion holds, then this theorem effectively communicates a localization phenomenon. Even though the Fourier transform is non-local, in the sense that all elements of \mathcal{X} participate, only those whose distances are closer than τ substantially contribute. On the other hand, if $\tau = \frac{1}{2}$ is selected, then there is no localization.

Motivated by inverse problems where only weak information about \mathcal{X} is known or can be reasonably assumed, we provide a different lower bound for $\sigma_s(\Phi)$ in terms of any lower bound for $\Delta(\mathcal{X})$. The following is our second main result.

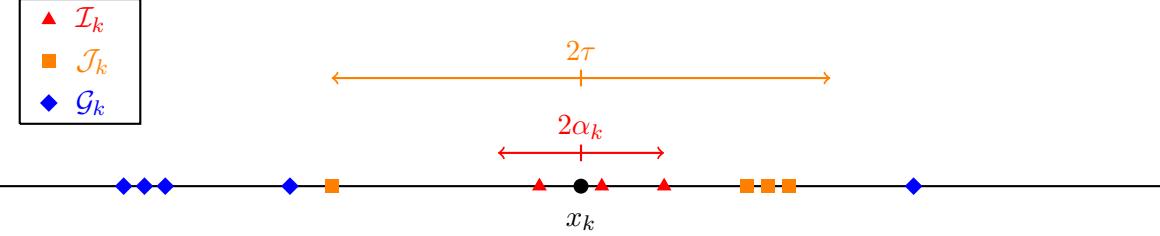


Figure 2: An example of the three sets \mathcal{I}_k , \mathcal{J}_k and \mathcal{G}_k defined in Theorem 1.

Theorem 2. *Let $m, s \in \mathbb{N}_+$ such that $s \geq 2$ and $m \geq 6s$, and let $\delta \in (0, \frac{1}{m}]$. Suppose $\mathcal{X} \subseteq \mathbb{T}$ is a set of cardinality s with $\Delta(\mathcal{X}) \geq \delta$, and pick any τ such that \mathcal{X} satisfies the (m, τ) density criterion. For each $k = 1, 2, \dots, s$, define*

$$\mathcal{G}_k := \mathcal{G}(x_k, \tau, \mathcal{X}), \quad r_k := |\mathcal{B}(x_k, \tau, \mathcal{X})|, \quad \text{and} \quad n_k := \left\lfloor m - \frac{2\nu(\tau, \mathcal{G}_k)}{\tau} \right\rfloor.$$

Then we have

$$\frac{1}{\sigma_s^2(\Phi(m, \mathcal{X}))} \leq \frac{4e^2}{\pi^2} \sum_{k=1}^s 2^{\nu(\tau, \mathcal{G}_k)} \frac{\phi(\frac{n_k}{r_k})}{r_k n_k} \left(\frac{2e\phi(\frac{n_k}{r_k})}{\sin(\frac{\pi}{2}n_k\delta)} \right)^{2r_k-2}, \quad (2.4)$$

and in particular,

$$\sigma_s(\Phi(m, \mathcal{X})) \geq \frac{\pi}{2e} \sqrt{\frac{m}{6s}} \min_{1 \leq k \leq s} \left\{ \sqrt{\frac{r_k}{2^{\nu(\tau, \mathcal{G}_k)}}} \left(\frac{m\delta}{12e} \right)^{r_k-1} \right\}. \quad (2.5)$$

We emphasize that δ is an independent parameter, so the theorem is applicable to sets for which $\Delta(\mathcal{X})$ is arbitrarily small. This theorem is written from the perspective of δ . In (2.5), the exponent on δ is $r_k - 1$, which shows that interactions between x and x_k at scales smaller than τ are most significant. This theorem assumes that $\delta \leq \frac{1}{m}$, which can be relaxed by adapting this theorem's proof, but with some additional technical complications. This is not a prohibitive assumption since the estimate (1.1) can be used whenever $\delta > \frac{1}{m}$.

Compared to Theorem 1, Theorem 2 is easier to employ since it requires less information about \mathcal{X} , but it generally yields a looser bound. This is expected since Theorem 1 contains product terms that depend on the pairwise distances between elements, whereas Theorem 2 has effectively replaced all of these distances by δ . Both theorems give similar predictions if all small scales are approximately δ . For sets with many scales between δ and τ , it is generally advisable to use Theorem 1 instead.

To use Theorems 1 and 2, one first needs to select an appropriate τ for which the (m, τ) density criterion holds, and as mentioned earlier, there may be infinitely many choices. For certain sets, one could select τ heuristically based on m and properties of \mathcal{X} . Another option is to sweep through a collection \mathcal{T} of τ and pick a τ that maximizes whichever bound one would like to use. We will discuss the computational cost at the end of this article.

Clumps models for \mathcal{X} were independently introduced in [27, 3] and were used to control the condition number of tall Fourier matrices. There are some subtle differences between the definitions in these papers, so to facilitate the presentation and to avoid giving two separate definitions, we work with the following boarder definition that encapsulates both frameworks. For sets $\mathcal{U}, \mathcal{V} \subseteq \mathbb{T}$, we define the diameter and distance,

$$\text{diam}(\mathcal{U}) := \sup_{u, u' \in \mathcal{U}} |u - u'|_{\mathbb{T}} \quad \text{and} \quad \text{dist}(\mathcal{U}, \mathcal{V}) := \inf_{u \in \mathcal{U}, v \in \mathcal{V}} |u - v|_{\mathbb{T}}.$$

Definition 2.3. A set $\mathcal{X} \subseteq \mathbb{T}$ consists of *separated clumps* with parameters $(s, \delta, r, \lambda, \alpha, \beta)$ if the following hold. We have $|\mathcal{X}| = s$, $\Delta(\mathcal{X}) \geq \delta$, and there is a disjoint union

$$\mathcal{X} = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \cdots \cup \mathcal{C}_r,$$

where each \mathcal{C}_k is called a *clump* such that

$$\max_{1 \leq k \leq r} |\mathcal{C}_k| = \lambda, \quad \max_{1 \leq k \leq r} \text{diam}(\mathcal{C}_k) \leq \alpha, \quad \text{and} \quad \min_{j \neq k} \text{dist}(\mathcal{C}_j, \mathcal{C}_k) > \beta > \alpha \quad \text{if } r \geq 2.$$

A few remarks are in order. For a fixed \mathcal{X} , the choice of parameters is not unique and it is usually advisable to select valid parameters that minimize λ . If $r = 1$, then $s = \lambda$ and β is not a meaningful parameter since there is only a single clump. This is why the clump separation requirement is necessary only when $r \geq 2$. Notice $\beta > \alpha$ is included in the assumption so that distances between clumps exceeds within a clump. We will see that λ plays the role of $\max_{1 \leq k \leq s} r_k$, where r_k was defined in Theorem 2.

There are natural situations where a set consisting of separated clumps also satisfies the requirements of our main results, as shown in the next proposition.

Corollary 1. *Let $m, s \in \mathbb{N}_+$ such that $s \geq 2$ and $m \geq 6s$. Suppose \mathcal{X} consists of separated clumps with parameters $(s, r, \delta, \lambda, \alpha, \beta)$ such that $\delta \leq \frac{1}{m}$. If $r > 1$, also assume that $\beta \geq \frac{3\lambda}{m}$. Set $\tau = \frac{1}{2}$ if $r = 1$, otherwise let $\tau = \beta$. Then \mathcal{X} satisfies the (m, τ) density criterion, and the conclusions of Theorems 1 and 2 hold. In particular, we have*

$$\sigma_s(\Phi(m, \mathcal{X})) \geq \frac{\pi}{2e} \sqrt{\frac{m}{12s}} \left(\frac{m\delta}{12\sqrt{2}e} \right)^{\lambda-1}.$$

The condition that β scales linearly in $\frac{\lambda}{m}$ is the best one can expect without imposing further restrictions on \mathcal{X} . Indeed, if we allow $\beta < \frac{\lambda}{m}$, then it may occur that $m < s$. Although Corollary 1 provides the same or worse estimate compared to Theorem 2, we included the corollary in order to compare with prior results for clumps.

While this paper focuses on the smallest singular value, the techniques developed in this paper provide a straightforward and nontrivial upper bound for the largest singular value.

Theorem 3. *Let $m, s \in \mathbb{N}_+$ such that $m \geq s$. For any $\mathcal{X} \subseteq \mathbb{T}$ of cardinality s and $\tau \in (0, \frac{1}{2}]$ such that $m > \frac{1}{\tau}$, we have*

$$\sigma_1(\Phi(m, \mathcal{X})) \leq \sqrt{\nu(\tau, \mathcal{X}) \left(m + \frac{1}{\tau} \right)}.$$

For comparison purposes, recall the trivial bound $\sigma_1(\Phi) \leq \|\Phi\|_F = \sqrt{ms}$. Observe that Theorem 3 provides a significantly better upper bound if the τ local sparsity of \mathcal{X} is much smaller than s . For example, if we were to apply the above theorem for $\tau = \frac{2}{m}$, then $\sigma_1(\Phi) = \sqrt{3\nu(\tau, \mathcal{X})m/2}$, which is an improvement over the trivial bound if $\nu(\tau, \mathcal{X}) < \frac{2s}{3}$.

Organization

Section 3 provides detailed comparisons with prior work on the condition number of Fourier matrices, and serves as an expanded version of Section 1.1. There, we will see that our main theorems capture the scaling and localization phenomena that are missing from the classical Gautschi-Bazán

theorem. In the case of clumps, we will see that Corollary 1 is equivalent to the lower bound in (1.2) modulo universal constants, while holding under tremendously weaker separation assumptions.

Section 4 is dedicated to examples and numerical simulations, with comparisons to the predictions provided by this paper. It also provides more details regarding the motivational example in Section 1.2. There, we provide some extreme examples that illustrate when localization does (not) occur. One on hand, there are examples where $\tau = \frac{C}{m}$ like in the clumps model, while in other examples, $\tau = \frac{1}{2}$ such as for sparse spike trains. They illustrate the effectiveness and flexibility of the main results.

The remaining portions deal with proofs. Section 5 develops the main tool called the polynomial method and introduces two specific trigonometric interpolation problems that are connected to the main theorems. This section also outlines the main strategy for proving the main theorems without technical details. Section 6 addresses the “good” and “bad” interpolation problems, and how the resulting polynomials are related to other interpolation strategies. Section 7 contains proofs of the main results stated in Section 2.

3 Comparison with prior art

3.1 Comparison to classical estimates

Classical versus modern papers on Fourier matrices centers on the differences between square versus rectangular. A $s \times s$ Fourier matrix is perfectly conditioned if and only if \mathcal{X} is some shift of the uniform lattice $\{\frac{k}{s}\}_{k=1}^s$, see [9]. It is natural to wonder whether it is possible to relax both sides of this characterization. It would be a delicate task, since [15] established that if \mathcal{X} consists of the first s terms of the Van Der Corput sequence, then $\kappa(\Phi(s, \mathcal{X})) = 1$ only if $\log_2(s)$ is an integer, but grows like \sqrt{s} otherwise. This example illustrates that it is possible for $\kappa(\Phi(s, \mathcal{X}))$ to be unbounded in s even if \mathcal{X} is “spread out” in \mathbb{T} . Stability of the discrete Fourier transform matrix to perturbations of its nodes via the Kadec- $\frac{1}{4}$ theorem were derived in [42].

The results listed in the previous paragraph illustrate the brittleness of square matrices, while rectangular ones are much more robust. Any $m \times s$ sub-matrix of the $m \times m$ discrete Fourier transform matrix is perfectly conditioned even though the nodes are not uniformly spaced on the circle. More generally, notice from inequality (1.1) that $\Delta(\mathcal{X}) > \frac{1}{m}$ implies the conditioning of $\Phi(m, \mathcal{X})$ can be bounded uniformly in both m and s . It is important to mention that this inequality only applies to rectangular matrices since $\Delta(\mathcal{X}) > \frac{1}{m}$ and $|\mathcal{X}| = s$ imply that $m > s$. These observations should be compared with the ones listed in the previous paragraph for square matrices.

Results for square matrices can be used to deduce bounds for rectangular ones, beyond the trivial relationship $\sigma_s(\Phi(m, \mathcal{X})) \geq \sigma_s(\Phi(s, \mathcal{X}))$. We first start with Gautschi [21, Theorem 1] for square matrices,

$$\|\Phi(s, \mathcal{X})^{-1}\|_\infty \leq \max_{1 \leq k \leq s} \left\{ \prod_{j=1, j \neq k}^s \frac{2}{|e^{2\pi i x_j} - e^{2\pi i x_k}|} \right\}.$$

Here, $\|\cdot\|_p$ denotes the $\ell^p \rightarrow \ell^p$ operator norm. Next, Bazán [7, Theorem 1] showed that whenever $\frac{m}{s} \in \mathbb{N}$, then

$$\|\Phi(m, \mathcal{X})^\dagger\|_2 \leq \sqrt{\frac{s}{m}} \|\Phi(s, \mathcal{X})^{-1}\|_2.$$

Combining the above two inequalities, that $\|A\|_2 \leq \sqrt{s}\|A\|_\infty$ if $A \in \mathbb{C}^{s \times s}$, and $\sigma_s(\Phi(m, \mathcal{X})) \geq \sigma_s(\Phi(\lfloor \frac{m}{s} \rfloor s, \mathcal{X}))$, we obtain the Gautschi-Bazán theorem, which was stated in inequality (1.4).

Comparing the Gautschi-Bazán theorem with Theorem 1, we see that there are two main differences. First, the former does not exhibit *localization* since the product in (1.4) is taken over all $x \in \mathcal{X} \setminus \{x_k\}$, whereas in the latter, the product is over all $x \in \mathcal{I}_k$, while the further away elements are less significant. Note that if $|x - x_k|_{\mathbb{T}}$ is small, then $|e^{2\pi i x} - e^{2\pi i x_k}|$ is comparable to $|x - x_k|_{\mathbb{T}}$. Second, the former does not exhibit the correct *scaling* in front of $|e^{2\pi i x} - e^{2\pi i x_k}|$. In the latter, notice that each term has a helpful α_k factor, which could be on the order of $\frac{1}{m}$ depending on \mathcal{X} . The localization and scaling phenomenon manifest when we consider rectangular Fourier matrices, which are absent for square ones. Examining the proof of [7, Theorem 1], we see that it treats tall Fourier matrices as $\lfloor \frac{m}{s} \rfloor$ independent $s \times s$ blocks, and does not fully exploit the algebraic structure of tall Fourier matrices.

Continuing the remarks made in the previous paragraph, there is an elementary explanation for why tall Fourier matrices should behave differently from square ones. Notice that $(\Phi^* \Phi)_{j,k} = D_m(x_j - x_k)$, where $D_m(t) := \sum_{k=0}^{m-1} e^{2\pi i k t}$ is the Dirichlet kernel. We easily see that $|D_m(t)|$ is on the order of m on the interval $[-\frac{1}{m}, \frac{1}{m}]$ and decays at a rate of $|t|^{-1}$ away from 0. This means that the Gram matrix $\Phi^* \Phi$, for fixed \mathcal{X} and increasing m , becomes increasingly diagonally dominant. Basic and generic tools such as the Gershgorin circle theorem fail to provide any meaningful results when $\Delta(\mathcal{X}) < \frac{1}{m}$ and $s \geq 3$ because the diagonal entries of $\Phi^* \Phi$ are m while the ℓ^1 norm of off-diagonal rows and columns of $\Phi^* \Phi$ exceed m . Instead, the proof methods used in this paper specifically take advantage of the algebraic structure of Fourier matrices and are able to obtain finer results.

3.2 Comparison to clumps

In this part, we compare Corollary 1 with the results in [27, 3, 22, 5]. As usual, we let m and s denote the number of rows and columns of Φ . We will only compare the general scaling of the model parameters and do not compare universal constants, since the latter can be improved by optimizing their proofs or by providing more accurate but complicated expressions. When comparing our main theorems with other papers, we will generally ignore distinctions between $m - 1$, m , and $2m$, since the extraneous factors can be absorbed into other constants.

The result [27, Theorem 2.7] shows that if \mathcal{X} consists of separated clumps with parameters $(s, \delta, r, \lambda, \frac{1}{m}, \beta)$ such that

$$m \geq s^2 \quad \text{and} \quad \beta \geq \frac{20}{m-1} \sqrt{\frac{s\lambda^5}{(m-1)\delta}} \quad \text{if} \quad r > 1, \quad (3.1)$$

then there exist explicit universal constants $C > 0$ and $c \in (0, 1)$ such that

$$\sigma_s(\Phi(m, \mathcal{X})) \geq C \sqrt{\frac{m-1}{\lambda}} (c(m-1)\delta)^{\lambda-1}.$$

One main drawback of condition (3.1) is that $\beta \rightarrow \infty$ as $\delta \rightarrow 0$, so for sufficiently small δ , the theorem only applies when there is only a single clump. Some improvements to the explicit constants and variations of this inequality can be found in [22]. All results in this paper also require separation conditions for which $\beta \rightarrow \infty$ as $\delta \rightarrow 0$.

Corollary 1 shows that under the same hypotheses (3.1), this paper's main results are applicable and they yield the same $C\sqrt{m}(cm\delta)^{\lambda-1}$ estimate with different constants. However, Corollary 1 requires significantly weaker clump separation assumptions and relationship between m versus s . Importantly, it removes the artificial behavior that β explodes in the limit that δ goes to zero.

Moving on, [5, Theorem 2.2] shows that if \mathcal{X} consists of separated clumps with parameters $(s, \delta, r, \lambda, \alpha, \beta)$ such that

$$\frac{c_1(\lambda)s}{\pi\beta} \leq m \leq \frac{c_2(\lambda)}{\pi s\alpha}, \quad (3.2)$$

for some $c_1(\lambda), c_2(\lambda) > 0$ depending only on λ , then there is an explicit universal constant $C' > 0$ such that

$$\sigma_s(\Phi(2m+1, \mathcal{X})) \geq C' \sqrt{m} \left(\frac{m\delta}{16e} \right)^{\lambda-1}.$$

Some of the constants in these expression are different than those in [5] since that paper identifies \mathbb{T} with $[-\pi, \pi)$ as opposed to $[0, 1)$ in this paper.

Although $c_1(\lambda)$ and $c_2(\lambda)$ are not given explicitly, [5, Section 6.3] shows that $c_1(\lambda) > C(1 + \lambda^9)$ for some universal $C > 0$ and $c_2(\lambda) \leq 2\pi$. Hence, condition (3.2) requires, for all sufficiently large λ ,

$$\alpha \leq \frac{2}{sm} \quad \text{and} \quad \beta \geq \frac{C(1 + \lambda^9)s}{m} \geq \frac{3\lambda}{m}.$$

This establishes that Corollary 1 provides a similar lower bound, but again, under significantly weaker assumptions.

Clumps models were also introduced in [3] to bound the smallest singular value of a “continuous” analogue, whereby Φ is replaced with an integral operator. In particular, [3, Corollary 3.6] assumes that \mathcal{X} consists of separated clumps and with the additional requirement that $\text{diam}(\mathcal{X}) \leq \frac{1}{\pi s^2}$. It is not possible rescale this result to avoid this requirement, so we cannot provide a reasonable comparison. Nevertheless, the restriction that \mathcal{X} is contained in an interval of length $\frac{1}{\pi s^2}$ is removed in a follow-up result [5], which we already compared to.

3.3 Other related work

The “colliding nodes” model, where \mathcal{X} can be decomposed into clumps where each one has exactly two elements, was studied in [23]. This is much more restrictive than the clumps model and can be treated with specialized tools that cannot be extend to more complicated and general sets.

There is a plethora of papers that examine sub-matrices of the discrete Fourier transform matrix, see [2] and references therein. This would correspond to the situation where $\mathcal{X} \subseteq \{\frac{k}{n}\}_{k=1}^n$ and n is a large parameter that can be selected independent of m, s . This setting is more specialized since there are cancellation properties and explicit formulas that are not available in the general case.

4 Numerical simulations and examples

4.1 Setup and definitions

When comparing the true value of $\sigma_s(\Phi)$ and our estimated one, we use our more accurate estimate (2.2) from Theorem 1. The software that reproduces the figures in this paper are publicly available on the author’s Github repository ¹, which is also linked to the author’s personal website ².

The behavior of Fourier matrices was numerically evaluated in [27, 5, 3, 22] under the *super-resolution limit*, whereby m is sufficiently large and there is a family of $\{\mathcal{X}_m\}_{m=1}^\infty$ for which $\Delta(\mathcal{X}_m) \rightarrow 0$, or alternatively, $m \rightarrow \infty$ and $\Delta(\mathcal{X}_m) \rightarrow 0$ with some relationship between m and $\Delta(\mathcal{X}_m)$. This is an important scaling in the theory of super-resolution and the behavior of $\sigma_s(\Phi)$ greatly simplifies

¹<https://github.com/weilinlimath>

²<https://weilinli.ccny.cuny.edu>

in this scenario. The main results of this paper can be used for the super-resolution limit as well and would give equivalent predictions up to implicit constants, see Corollary 1. One can consider a complementary scaling, called the *well-separated case*, whereby \mathcal{X} is fixed and $m \rightarrow \infty$. In this case, (1.1) is applicable.

Rather than look at either scaling again, we look at more challenging examples. In the absence of a large or small parameter, $\sigma_s(\Phi)$ is naturally a discrete quantity. Nonetheless, even though our main theorem is proved using analytic tools, it only requires a weak assumption that $m \geq 6s$, so it is applicable to a greater variety of examples. We are only aware of one other result with this generality, which is the Gautschi-Bazán theorem in (1.4).

Since we provide lower bounds for the smallest singular value, it makes sense to quantify the quality of approximation by a multiplicative factor. That is, we define the

$$\text{inaccuracy factor} := \frac{\text{true value}}{\text{estimated value}}.$$

Of course, this quantity is lower bounded by 1.

4.2 The motivational example revisited

Here, we provide additional details for the motivational example in Section 1.2. First, notice that for a fixed \mathcal{X} , the set of τ for which the (m, τ) density criterion is satisfied are nested increasing sets as m increases. More precisely, if we define

$$\mathcal{S}(m, \mathcal{X}) := \{\tau : \mathcal{X} \text{ satisfies the } (m, \tau) \text{ density criterion}\},$$

then $\mathcal{S}(m, \mathcal{X}) \subseteq \mathcal{S}(m+1, \mathcal{X})$. So as m increases, we have the option of choosing τ smaller in order to reduce the number $x \in \mathcal{X}$ that are close to each x_k . However, we do not simply define τ as the infimum of $\mathcal{S}(m, \mathcal{X})$ because α_k may increase when τ decreases. Choosing an optimal τ is beyond the scope of this paper. It is not difficult to select reasonable a candidate based on intuition, or trial and error.

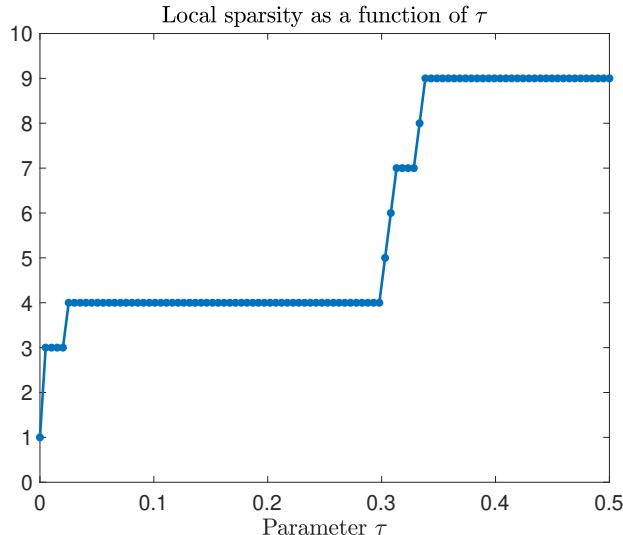


Figure 3: Plot of $\nu(\tau, \mathcal{X})$ as a function of $\tau \in (0, \frac{1}{2}]$.

Returning back to the motivational example, after some calculations, we define

$$\tau := \begin{cases} \frac{2}{30} & \text{if } m \in (450, 600], \\ \frac{3}{10} & \text{if } m \in [54, 450]. \end{cases}$$

For these corresponding values of τ , it can be easily checked that $\nu(\tau, \mathcal{X})$ is 3 or 4 and that \mathcal{X} satisfies the (m, τ) density criterion. To visualize the former, we have plotted $\nu(\tau, \mathcal{X})$ as a function of τ in Figure 3. Note that our choices for τ are not optimized, but were chosen according to reasonable heuristics.

As shown in Figure 1, our theorem yields a significantly more accurate prediction, which becomes more apparent as m increases. This occurs because the effective scale τ should be chosen to decrease in m and the distance between nearby elements is scaled according to α_k , neither of which are captured in the Gautschi-Bazán theorem. Additionally, the results in [27, 22] are not applicable for any $m \in [9, 600] \cap \mathbb{N}$, because the separation condition (3.1) is not fulfilled, while [3] cannot be used since \mathcal{X} is not contained in an interval of length $\frac{1}{\pi s^2}$, and it is unclear whether [5, 3] can be used since they contain implicit constants in their separation criterion.

4.3 Another multiscale example

Unlike the motivational example in Section 1.2 where \mathcal{X} was fixed and m varies, we consider the reverse situation where m is fixed and we have a family of sets \mathcal{X}_ε parameterized by a $\varepsilon \in (0, 1]$. Consider the set

$$\begin{aligned} \mathcal{X}_\varepsilon &:= \mathcal{X}_{1,\varepsilon} \cup \mathcal{X}_{2,\varepsilon} \cup \mathcal{X}_{3,\varepsilon}, \quad \text{where} \\ \mathcal{X}_{1,\varepsilon} &= \varepsilon \{0, \frac{1}{90}, \frac{2}{90}, \frac{3}{90}\}, \quad \mathcal{X}_{2,\varepsilon} = \frac{1}{3} + \varepsilon \{0, \frac{1}{200}, \frac{2}{200}\}, \quad \text{and} \quad \mathcal{X}_{3,\varepsilon} = \frac{2}{3} + \varepsilon \{0, \frac{1}{500}\}. \end{aligned} \quad (4.1)$$

We have defined \mathcal{X}_ε in this way to emphasize that while ε controls the minimum separation since $\Delta(\mathcal{X}_\varepsilon) = \frac{\varepsilon}{500}$, the three sets \mathcal{X}_1 , \mathcal{X}_2 , and \mathcal{X}_3 are still of different scales for each ε .

Since $\Delta(\mathcal{X}_\varepsilon) \leq \frac{1}{500}$ for any ε , we consider only $m \leq 500$. If we pick $\tau = \frac{3}{10}$, then $\nu(\tau, \mathcal{X}_\varepsilon) = 4$ for all ε . For two separate experiments, we select $m = 400$ and $m = 100$. Note that \mathcal{X}_ε satisfies the (m, τ) density criterion for all values of ε .

The results are shown in Figure 4. We see from the simulations that our lower bound matches the true behavior of the smallest singular value. Notice that for both experiments, $\sigma_s(\Phi)$ is piecewise linear, which is expected. Indeed, our theory states that the only significant interactions between $x, x' \in \mathcal{X}$ are those for which $|x - x'| \leq \tau$. For $\tau = \frac{3}{10}$, the sets \mathcal{X}_1 , \mathcal{X}_2 , and \mathcal{X}_3 do not have significant interactions due to our choice of τ . They also have cardinality 4, 3, and 2 respectively for all ε , and the interactions between elements in each \mathcal{X}_k scales linearly with ε . Hence, according to Theorem 1, we expect

$$\sigma_s(\Phi(m, \mathcal{X}_\varepsilon)) \gtrsim c_1 \varepsilon^3 + c_2 \varepsilon^2 + c_3 \varepsilon^1,$$

for some universal constants that can be explicitly computed. Hence, as ε varies, depending on the regime of ε and the size of c_1, c_2, c_3 , the dominant term in this inequality changes. In fact, Figure 4 shows that $\sigma_s(\Phi)$ appears to consist of three power-law pieces; on a log-log plot, they have slopes approximately 3.0055, 2.0321, and 0.9485, which is consistent with our prediction.

4.4 Sparse spike train

In this example, we consider an extreme situation where τ cannot be chosen on the order of $\frac{1}{m}$ even though the number of elements in an interval of length $\frac{2}{m}$ is at most 3. For any $\varepsilon \in (0, 1]$ and

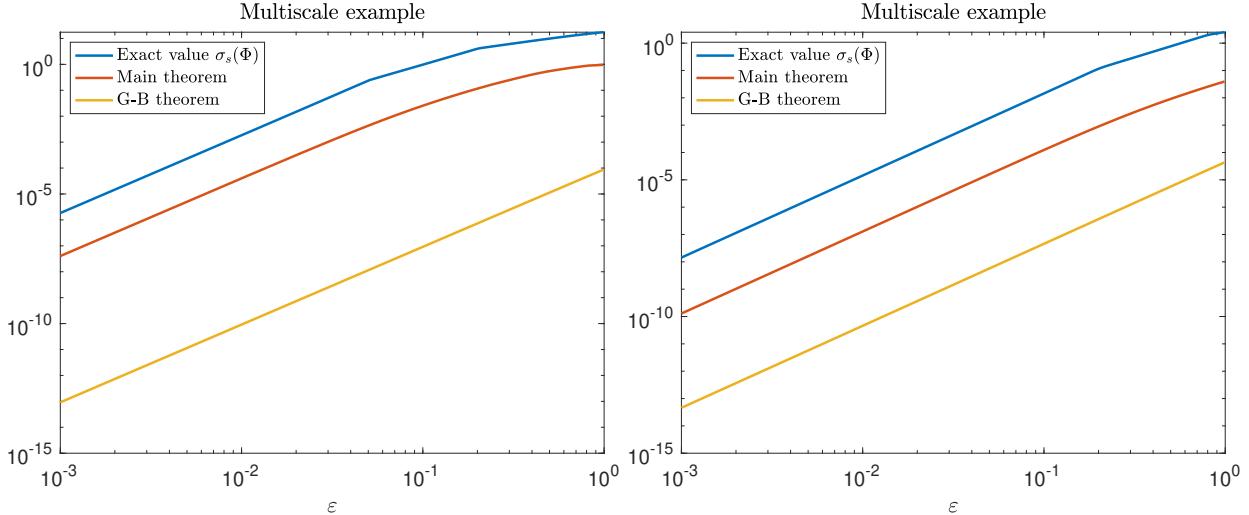


Figure 4: For \mathcal{X} defined in (4.1), plot of $\sigma_s(\Phi(m, \mathcal{X}_\varepsilon))$, main theorem, and Gautschi-Bazán bound, as a function of ε , with $m = 400$ on the left and $m = 100$ on the right.

$s \in [5, 30] \cap \mathbb{N}_+$, we set $m = 200$ and consider the following set,

$$\mathcal{X}_{s,\varepsilon} = \varepsilon \left\{ 0, \frac{1}{m}, \dots, \frac{s-1}{m} \right\}. \quad (4.2)$$

Our choices of m and s here are arbitrary and we could have considered larger or smaller m provided that s is sufficiently small compared to m .

It was shown in [31] that for fixed ε , provided that $s \geq C \log m$, then $\sigma_s(\Phi(m, \mathcal{X}_{s,\varepsilon})) \leq C' e^{-c\varepsilon s}$ for some unspecified universal $C, C', c > 0$. An improved estimate and with explicit C, C', c were derived in [2]. Provided that $\frac{m}{\varepsilon}$ is an integer, [2, Theorem 10] provides an explicit upper bound, which for sufficiently large $\frac{m}{\varepsilon}$, simplifies to

$$\sigma_s(\Phi(m, \mathcal{X}_{s,\varepsilon})) \leq C e^{-\pi(1-\varepsilon)s/2}. \quad (4.3)$$

There is another bound [2, Lemma 19], which is better for $\varepsilon \leq \varepsilon_* \approx 0.117$ and takes the form

$$\sigma_s(\Phi(m, \mathcal{X}_{s,\varepsilon})) \leq \frac{2\sqrt{ms}}{1 - (e\pi(m-1)\varepsilon/4m)} \left(\frac{e\pi(m-1)\varepsilon}{4m} \right)^{s-1}.$$

These upper bounds have important implications. First, they show that even if \mathcal{X} consists of clumps, they need to be sufficiently far apart for the lower bound in (1.2) to be valid, otherwise there is a contradiction. However, it does not provide a quantitative bound on the clump separation. Second, they imply that if $\Delta(\mathcal{X}) < \frac{1}{m}$, then we need to put some restrictions on s otherwise $\kappa(\Phi)$ may grow exponentially. Indeed, without an upper bound on s , we can let \mathcal{X} be the set in (4.2) with s on the order of m . This also explains why we cannot substantially relax the (m, τ) density criterion. We will provide more details related to the second point below.

Notice that $\nu(\frac{1}{m}, \mathcal{X}_{s,\varepsilon}) \leq 3$ for all $\varepsilon \in (0, \frac{1}{2}]$, so at first glance, it may be tempting to set τ on the order of $\frac{1}{m}$. However, it is not hard to see that there is no $\tau < \frac{3s}{m}$ for which $\mathcal{X}_{s,\varepsilon}$ satisfies the (m, τ) density criterion. On the other hand, if $\tau \geq \frac{3s}{m}$, then $\nu(\tau, \mathcal{X}_{s,\varepsilon}) = s$ and consequently, \mathcal{X} satisfies the (m, τ) density criterion. The graph of $\nu(\tau, \mathcal{X}_{s,\varepsilon})$ as a function of τ is shown in Figure 5. Thus, we are in the extreme case where we should just pick $\tau = \frac{1}{2}$. Intuitively, we think of $\mathcal{X}_{s,\varepsilon}$ as a high density set.

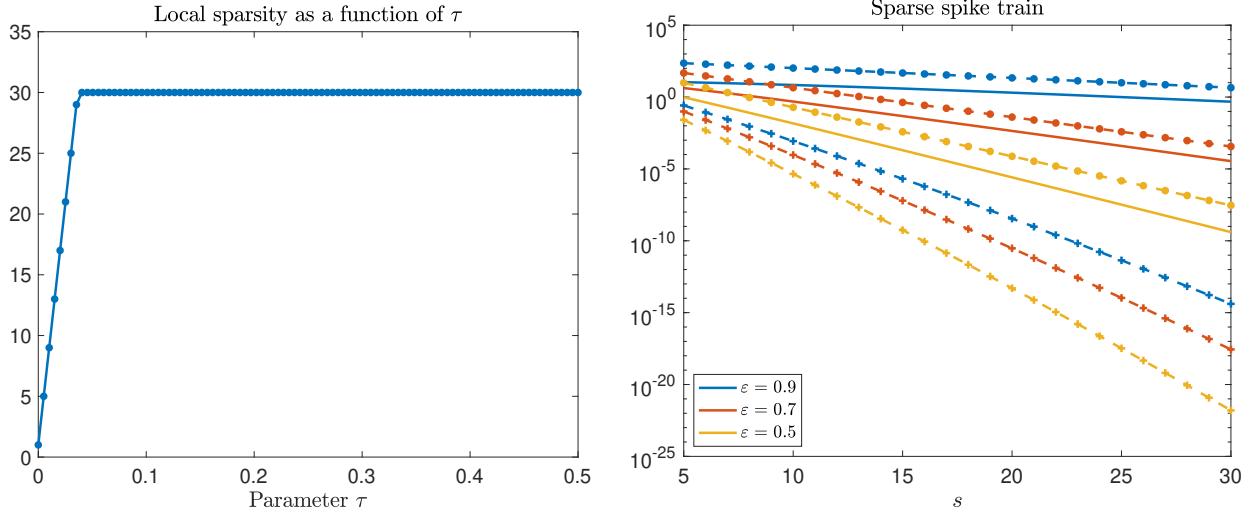


Figure 5: Left: Plot of the local sparsity $\nu(\tau, \mathcal{X}_{s,\varepsilon})$ as a function of τ for $\varepsilon = \frac{1}{2}$. Right: For $\mathcal{X}_{s,\varepsilon}$ defined in (4.2), $m = 200$, and $\varepsilon = 0.9, 0.7, 0.5$, the graphs of $\sigma_s(\Phi(m, \mathcal{X}_{s,\varepsilon}))$, upper bound (4.3) with $C = 500$, and lower bound (2.2) are shown in solid, dashed circular markers, and dashed plus sign markers, respectively. The slopes of the upper bounds are $-\frac{\pi(1-\varepsilon)s}{2}$. The slopes of the lower bounds are -1.2729 , -1.5259 , and -1.8625 , found by best linear fit.

Figure 5 plots the numerically computed $\sigma_s(\Phi(m, \mathcal{X}_{s,\varepsilon}))$ and our main theorem as functions of s and for $\varepsilon = 0.9, 0.7, 0.5$. The numerical simulations indicate that $\log(\sigma_s(\Phi(m, \mathcal{X}_{s,\varepsilon})))$ is well approximated by an affine function of s , namely,

$$\log(\sigma_s(\Phi(m, \mathcal{X}_{s,\varepsilon}))) = b_{\varepsilon,m} - c_{\varepsilon,m}s + \text{less significant terms depending on } s.$$

This is consistent with Theorem 1, which predicts that the dominant term in $\log(\sigma_s(\Phi))$ is affine in s . To see why, using the notation defined in the theorem, there is a k for which $\mathcal{I}_k = \mathcal{X} \setminus \{x_k\}$, so there is at least one term on the right side of (2.2) that contains a product of $s-1$ terms. They are the dominant terms since \mathcal{I}_k exerts the greatest influence on the lower bound.

If one wants a provable (but worse) result for this example, we recommend Theorem 2. Following the notation of that theorem, whenever $s \geq 2$, $m \geq 6s$, and $\varepsilon \in (0, 1]$, we set $\tau = \frac{1}{2}$ and $\delta = \frac{\varepsilon}{m}$. Then for each k , we have $\mathcal{G}_k = \emptyset$, $\nu(\tau, \mathcal{G}_k) = 0$, $n_k = m$, $\mathcal{B}(x_k, \tau, \mathcal{X}) = \mathcal{X}$, and $r_k = s$. Since $m \geq 6s$, we have $\phi(\frac{m}{s}) \leq 1 + \frac{1}{5}$. Using (2.4), we have

$$\sigma_s(\Phi(m, \mathcal{X}_{s,\varepsilon})) \geq \frac{\pi}{2e} \sqrt{\frac{m}{\phi(\frac{m}{s})}} \left(\frac{\sin(\frac{\pi\varepsilon}{2})}{2e\phi(\frac{m}{s})} \right)^{s-1} \geq \frac{\pi}{2e} \sqrt{\frac{5m}{6}} \left(\frac{5\sin(\frac{\pi\varepsilon}{2})}{12e} \right)^{s-1}. \quad (4.4)$$

The lower bound for $\log(\sigma_s(\Phi))$ provided by (4.4) has slope -1.8879 , -1.9909 and -2.2220 for $\varepsilon = 0.9, 0.7, 0.5$ respectively.

4.5 Colliding clumps

Here we introduce an example where there are two localized sets that are progressive being pushed towards each other. To make this notion more precise, we fix $m = 100$ and for sufficiently small $\beta > 0$, define

$$\mathcal{C}_1 := \{0, \frac{1}{2m}, \frac{2}{2m}\}, \quad \mathcal{C}_2(\beta) := \beta + \frac{1}{m} + \mathcal{C}_1, \quad \text{and} \quad \mathcal{X}(\beta) := \mathcal{C}_1 \cup \mathcal{C}_2(\beta). \quad (4.5)$$

As $\beta \rightarrow 0$, the two sets \mathcal{C}_1 and $\mathcal{C}_2(\beta)$ become closer and $\sigma_s(\Phi(m, \mathcal{X}(\beta))) \rightarrow 0$ as $\beta \rightarrow 0$. Note we can think of \mathcal{C}_1 and $\mathcal{C}_2(\beta)$ as clumps with separation β . Eventually for sufficiently small β , we should think of $\mathcal{X}(\beta)$ as just a single clump as opposed to two separate ones.

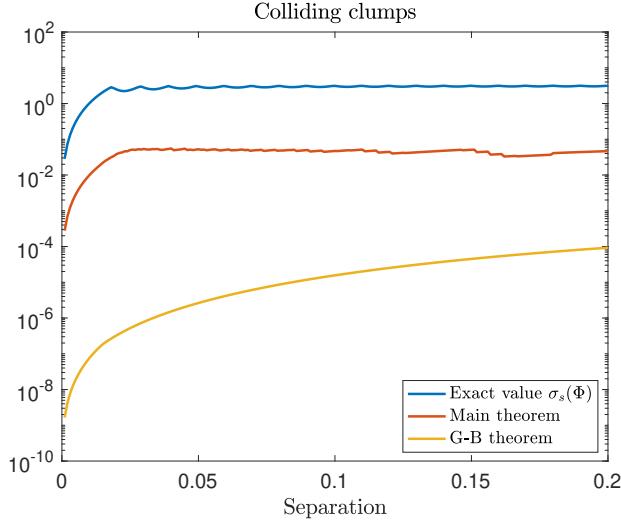


Figure 6: For $m = 100$ and $\mathcal{X}(\beta)$ defined in (4.5), plot of $\sigma_s(\Phi(m, \mathcal{X}(\beta)))$, our main theorem, and the Gautschi-Bazán bound as a function of $\beta \in [\frac{0.1}{m}, \frac{20}{m}]$.

To employ Theorem 1, we pick $\tau = \beta$ if $\beta \geq \frac{18}{m}$, otherwise we set $\tau = \frac{1}{2}$. Our choice of τ is consistent with Corollary 1, but we picked a bigger constant (18 instead of 9) in front of $\frac{1}{m}$ to temper the growth of several implicit constants in our estimates. For this example, the clumps bounds [27, 22] do not apply since \mathcal{C}_1 and $\mathcal{C}_2(\beta)$ are too close together. The results of this experiment are shown in Figure 6. This behavior of $\sigma_s(\Phi(m, \mathcal{X}(\beta)))$ undergoes a phase transition at $\beta \approx \frac{1}{m}$ since for $\beta \gg \frac{1}{m}$, it is intuitive that $\mathcal{X}(\beta)$ should be treated as two clumps instead of just one. There are some fluctuations in the graph of $\sigma_s(\Phi)$ due to number theoretic reasons since $\mathcal{X}(\beta)$ is a partial unitary matrix if it is a subset of the lattice with spacing $\frac{1}{2m}$. Our estimate is significantly better than the Gautschi-Bazán theorem. For example, at $\beta = 0.1 = \frac{10}{m}$, the former has an inaccuracy factor of 66.1225, while the latter is 1.9916e+05.

5 Proof strategy

5.1 The polynomial method

The torus is defined as $\mathbb{T} := \mathbb{R}/\mathbb{Z}$, which we normally identify with $[0, 1)$ via the map $x \mapsto \text{mod}(x, 1)$. The canonical basis vectors for \mathbb{R}^d is denoted $\{e_k\}_{k=1}^d$. We let $\|\cdot\|_p$ and $\|\cdot\|_{L^p}$ denote the ℓ^p and L^p norms respectively, for $1 \leq p \leq \infty$. The Fourier transform of a $f \in L^2(\mathbb{T})$ is denoted $\widehat{f}: \mathbb{Z} \rightarrow \mathbb{C}$, where $\widehat{f}(k) := \int_{\mathbb{T}} f(x) e^{-2\pi i k x} dx$ for each $k \in \mathbb{Z}$. We say f is a trigonometric polynomial of degree $m-1$ if its Fourier transform is supported in $\{0, 1, \dots, m-1\}$, and we let \mathcal{P}_m be the set of all trigonometric polynomials of degree at most $m-1$.

The primary method that we will use to lower bound $\sigma_s(\Phi)$, or more precisely, upper bound $1/\sigma_s(\Phi)$, is through a “dual” relationship with minimum norm trigonometric interpolation. This duality was introduced in [27, Proposition 2.12]: For any integers $m \geq s \geq 1$, finite set $\mathcal{X} \subseteq \mathbb{T}$ of cardinality s , and unit vector $v \in \mathbb{C}^s$ such that $\|\Phi(m, \mathcal{X})v\| = \sigma_s(\Phi(m, \mathcal{X}))$ (i.e., v is any right

singular vector corresponding to the smallest singular value of $\Phi(m, \mathcal{X})$), we have

$$\sigma_s(\Phi(m, \mathcal{X})) = \max \left\{ \frac{1}{\|f\|_{L^2}} : f \in \mathcal{P}_m \text{ and } f(x_k) = v_k \text{ for each } k \right\}. \quad (5.1)$$

We refer to this equation as the *duality principle*, since it provides a connection between the smallest singular value to minimum norm trigonometric interpolation. There are related concepts [17, 11, 10] for \mathbb{R} instead of \mathbb{T} , but one main difference is that our \mathcal{X} is arbitrary and can be completely nonuniform.

It may be helpful to explain the intuition behind this duality principle. Suppose v is a unit norm right singular vector corresponding to the smallest singular value of $\Phi(m, \mathcal{X})$ and we examine all solutions u to $\Phi^*u = v$. Due to the singular value decomposition, any minimum norm vector u that is consistent with this system will have norm $1/\sigma_s(\Phi)$. Note that Φ^* is the matrix representation of the linear transform that maps Fourier coefficients of functions in \mathcal{P}_m to their restriction on \mathcal{X} . Using the Plancherel's theorem allows us to pass from the Fourier coefficients to polynomials. It follows from this discussion that a $f \in \mathcal{P}_m$ which achieves equality in (5.1) is necessarily a f whose Fourier coefficients are $u/\sigma_s(\Phi)$ where u is any unit norm left singular vector of Φ which corresponds to $\sigma_s(\Phi)$.

The duality principle provides a natural and constructive avenue for lower bounding $\sigma_s(\Phi)$. However, since we have no exploitable information on the right singular vectors of Φ , we construct interpolants for arbitrary v , and then estimate them in L^2 uniformly in v . This leads us to the subsequent definition and lemma.

Definition 5.1. For any set $\mathcal{X} = \{x_k\}_{k=1}^s \subseteq \mathbb{T}$, we say $\{f_k\}_{k=1}^s$ is a family of Lagrange interpolants for \mathcal{X} if $f_k(x_\ell) = \delta_{k,\ell}$ for each $1 \leq k, \ell \leq s$.

Lemma 5.2. For any $m, s \in \mathbb{N}_+$ with $m \geq s$ and $\mathcal{X} \subseteq \mathbb{T}$ of cardinality s , if $\{f_k\}_{k=1}^s \subseteq \mathcal{P}_m$ is a family of Lagrange interpolants for \mathcal{X} , then

$$\frac{1}{\sigma_s^2(\Phi(m, \mathcal{X}))} \leq \sum_{k=1}^s \|f_k\|_{L^2}^2.$$

Proof. Let $v \in \mathbb{C}^s$ be any unit norm vector such that $\|\Phi v\| = \sigma_s(\Phi)$. Since $f = \sum_{k=1}^s v_k f_k$ interpolates v on \mathcal{X} , by equation (5.1) and Cauchy-Schwarz, we have

$$\frac{1}{\sigma_s(\Phi)} \leq \|f\|_{L^2} \leq \sum_{k=1}^s |v_k| \|f_k\|_{L^2} \leq \left(\sum_{k=1}^s \|f_k\|_{L^2}^2 \right)^{1/2} \left(\sum_{k=1}^s |v_k|^2 \right)^{1/2} = \left(\sum_{k=1}^s \|f_k\|_{L^2}^2 \right)^{1/2}.$$

□

This lemma was implicitly used in [27], and allows us to reduce the problem of lower bounding $\sigma_s(\Phi)$ into constructing Lagrange interpolants. One strength of this method is that it does not require any information about the singular vectors of Φ , which is usually more difficult to analyze than the singular values. However, if we had additional information about them, such as localization properties, then the L^2 estimate provided here can be improved.

The next proposition serves as a converse to Lemma 5.2. It shows that any lower bound on the smallest singular value provides the existence of polynomials with prescribed interpolation properties.

Proposition 5.3. *If $\mathcal{X} \subseteq \mathbb{T}$ is a non-empty finite set with cardinality s , then for any $w \in \mathbb{C}^s$ and integer $m \geq s$, there exists $f \in \mathcal{P}_m$ such that $f|_{\mathcal{X}} = w$,*

$$\|f\|_{L^2} \leq \frac{\|w\|_2}{\sigma_s(\Phi(m, \mathcal{X}))} \quad \text{and} \quad \|f\|_{L^\infty} \leq \frac{\sqrt{m} \|w\|_2}{\sigma_s(\Phi(m, \mathcal{X}))}.$$

Proof. Note that $m \geq s$ implies $\Phi := \Phi(m, \mathcal{X})$ is injective due to the Vandermonde determinant theorem. We have the singular value decomposition $\Phi = \sum_{k=1}^s \sigma_k u_k v_k^*$, where the v_k 's and u_k 's are orthonormal and the σ_k 's are the nonzero singular values of Φ .

For any $w \in \mathbb{C}^s$, we have $w = \sum_{k=1}^s b_k v_k$ for some $b \in \mathbb{C}^s$ such that $\|b\|_2 = \|w\|_2$. For each k , we define $f_k \in \mathcal{P}_m$ such that $\hat{f}_k = u_k / \sigma_k$. Using again that Φ^* is the matrix representation of the operator that maps the Fourier coefficients of a function in \mathcal{P}_m to its values on \mathcal{X} , a direct calculation then yields that $f_k|_{\mathcal{X}} = \Phi^* \hat{f}_k = v_k$.

From here we see that $f := \sum_{k=1}^s b_k f_k$ satisfies $f|_{\mathcal{X}} = w$. Moreover, since the u_k 's are orthonormal, an application of Parseval's shows that the f_k 's are L^2 orthogonal, and so

$$\|f\|_{L^2}^2 = \sum_{k=1}^s |b_k|^2 \|f_k\|_{L^2}^2 = \sum_{k=1}^s \frac{|b_k|^2}{\sigma_k^2} \leq \frac{\|w\|_2^2}{\sigma_s^2}.$$

For the L^∞ bound, we use that $f \in \mathcal{P}_m$, and Cauchy-Schwarz, to get

$$\|f\|_{L^\infty} \leq \|\hat{f}\|_{\ell^1} \leq \sqrt{m} \|\hat{f}\|_{\ell^2} = \sqrt{m} \|f\|_{L^2}.$$

□

This proposition can be rephrased as a result for an interpolation operator. Consider the operator $T(\mathcal{X})$ that takes $f^* \in C(\mathbb{T})$ and produces the $f \in \mathcal{P}_m$ guaranteed by this proposition such that $f^* = f$ on \mathcal{X} . Equipping both $C(\mathbb{T})$ and \mathcal{P}_m with the L^∞ norm yields the estimate

$$\|T(\mathcal{X})\|_{L^\infty \rightarrow L^\infty} \leq \frac{\sqrt{ms}}{\sigma_s(\Phi(m, \mathcal{X}))}.$$

We loosely refer to the strategy provided by the results of this subsection as the *polynomial method*. A primary usefulness of this connection between $\sigma_s(\Phi)$ and trigonometric interpolation is that it can be used employ tools from Fourier analysis and polynomial approximation, instead of solely working with matrices. While this connection is helpful, it is only useful if one can construct Lagrange interpolants with small norm, otherwise the resulting lower bounds for $\sigma_s(\Phi)$ would be quite loose.

5.2 Outline of the main proofs from an abstract perspective

The proofs of Theorems 1 and 2 are based on the following general recipe. Due to the polynomial method, we only need to provide the existence of Lagrange interpolants with suitably small norms and degree at most $m - 1$. Note that \mathcal{P}_m enjoys numerous algebraic properties. In addition to being vector space, if $f \in \mathcal{P}_m$ and $g \in \mathcal{P}_n$, then $fg \in \mathcal{P}_{m+n-1}$. It is also a shift invariant space, namely, $f \in \mathcal{P}_m$ if and only if $f(\cdot - t) \in \mathcal{P}_m$ for any $t \in \mathbb{T}$.

We start the proof by fixing any $x_k \in \mathcal{X}$ and concentrate on establishing a polynomial $f_k \in \mathcal{P}_m$ such that $f_k(x_k) = 1$ and vanishes on $\mathcal{X} \setminus \{x_k\}$. Constructing a Lagrange interpolant of this data

is straightforward, but doing so in a naive manner leads to loose estimates. We use the standard Lagrange interpolant $\ell_k \in \mathcal{P}_s$ as a benchmark. Note that it has a pointwise upper bound,

$$\|\ell_k\|_{L^\infty} \leq \prod_{j \neq k} \frac{2}{|e^{2\pi i x_k} - e^{2\pi i x_j}|} = 2^{s-1} \prod_{j \neq k} \frac{1}{|e^{2\pi i x_k} - e^{2\pi i x_j}|}. \quad (5.2)$$

The right hand side grows exponentially in s and it contains a product of $s-1$ terms. It is significantly larger compared to the norms of polynomials that we will construct later. The main deficiency of ℓ_k is that $\deg(\ell_k) = s-1$, so it does not take advantage of the possibility that interpolants can be selected from \mathcal{P}_m where m can be significantly larger than s . From this point of view, we interpret m as the number of parameters or degrees of freedom, and s as the number of constraints.

Addition of two polynomials results in polynomial whose degree is the max, while multiplication adds their degrees. It is intuitive that points in \mathcal{X} near x_k require larger norm polynomials to interpolate, since we need $f_k(x_k) = 1$, yet f_k can potentially have many nearby zeros, whereas further away points require smaller norms. Hence, it is natural to decompose

$$\mathcal{X} = \mathcal{B}_k \cup \mathcal{G}_k,$$

where the “bad” \mathcal{B}_k and “good” \mathcal{G}_k sets contain the points near and far away from x_k , respectively. The scale τ at which these sets are selected is important and determined by the density criterion. Hence the original interpolation problem can be solved by finding and multiplying Lagrange interpolants b_k and g_k where $b_k(x_k) = g_k(x_k)$, b_k vanishes on $\mathcal{B}_k \setminus \{x_k\}$, and g_k vanishes on \mathcal{G}_k .

The interpolation problem for the good set will be handled in Section 6.1. Although each element of \mathcal{G}_k is sufficiently far away from x_k , points in \mathcal{G}_k can still be close together. Hence, it is not clear that there is even any advantage of splitting \mathcal{X} into the good and bad sets. To deal with this, we will employ Proposition 6.1 to further decompose \mathcal{G}_k as

$$\mathcal{G}_k = \mathcal{G}_{k,1} \cup \dots \cup \mathcal{G}_{k,\nu_k}, \quad \text{where } \nu_k := \nu(\tau, \mathcal{G}_k),$$

such that $\Delta(\mathcal{G}_{k,j})$ is suitably controlled from below. By using Proposition 5.3, we can recast inequality (1.1) as an interpolation statement. Doing so, we obtain the existence of ν_k many interpolants, which are multiplied together to obtain a desired g_k . Interpolation for the good set will require a budget of roughly $2\nu(\tau, \mathcal{G}_k)/\tau$, which is guaranteed to be at most $2m/3$ in view of the density criterion.

The interpolation problem for the bad set will be handled in Section 6.2. The starting point is a basic observation that the standard Lagrange interpolant for the bad set can be pointwise bounded by the distances between elements of \mathcal{B}_k and x_k as seen in (5.2). Note that if $q|t|_{\mathbb{T}} \leq \frac{1}{2}$ for some $q \in \mathbb{N}_+$, then $|qt|_{\mathbb{T}} = q|t|_{\mathbb{T}}$. Hence, if we shift and dilate the elements of \mathcal{B}_k , and use a Lagrange interpolant for the dilated points, such as

$$b_k(x) = \prod_{x_j \in \mathcal{B}_k \setminus \{x_k\}} \frac{e^{2\pi i q_j x} - e^{2\pi i q_j x_k}}{e^{2\pi i q_j x_k} - e^{2\pi i q_j x_j}},$$

then this polynomial will have significantly smaller norm and larger degree compared to the standard Lagrange interpolant ℓ_k . Here, each q_j will need to be chosen so that the degree of b_k is not too large. Interpolation for the bad set will use the remaining portion of our budget consisting of roughly $m-1-2\nu(\tau, \mathcal{G}_k)/\tau$.

Finally, the desired Lagrange interpolant is $f_k := b_k g_k$. Doing this for each $x_k \in \mathcal{X}$ yields a family of Lagrange interpolants for \mathcal{X} in \mathcal{P}_m , allowing us to employ Lemma 5.2, which completes the proof.

Carrying out these steps requires exploiting the advantages of several seemingly disparate approaches. The polynomial method for estimating the smallest singular value of Fourier matrices was introduced in [27] and was inspired by interpolation techniques [17, 11]. It is further refined in this paper to handle more abstract sets, beyond clumps and subsets of lattices, by incorporating density ideas. Although we were unable to find a prior reference that uses exactly the same density criterion, there are strong connections between sampling and density, such as [24].

The initial decomposition of \mathcal{X} into $\mathcal{B}_k \cup \mathcal{G}_k$ and further decompositions of \mathcal{G}_k into $\mathcal{G}_{k,1}, \dots, \mathcal{G}_{k,\nu_k}$, are inspired by the classical Calderón-Zygmund decomposition. Our method for dealing with the good set requires the lower bound in (1.1), which was proved in [1] by using powerful machinery developed for analytic number theory [39, 38, 33, 32]. Finally, the method for dealing with the bad set using local dilation methods was originally employed in [27], for which we make significant improvements to.

6 Two trigonometric interpolation problems

6.1 Small norm Lagrange interpolants

In this section, we study the interpolation problem for the “good” set $\mathcal{G} \subseteq \mathbb{T}$ where all elements of \mathcal{G} are away from zero, and we would like to find a trigonometric polynomial that vanishes on \mathcal{G} and equals 1 at 0. Since we do not want to place any assumptions on $\Delta(\mathcal{G})$, which we allow to be arbitrarily small, this is a delicate problem. We will construct a polynomial that is significantly better behaved than the standard Lagrange interpolant.

A key observation is the following *sparsity decomposition* which essentially states that a set can be decomposed into disjoint sets, each with unit local sparsity and minimum separation that is well-controlled. The key is that the number of sets equals the local sparsity of the original set, and not the cardinality. While this decomposition is intuitive, some care is taken with the proof due to the periodic boundary conditions that are imposed on us due to working with the torus.

Proposition 6.1. *For any $\tau \in (0, \frac{1}{2}]$ and non-empty $\mathcal{W} \subseteq \mathbb{T}$, letting $\nu := \nu(\tau, \mathcal{W})$, there exist non-empty disjoint subsets $\mathcal{W}_1, \mathcal{W}_2, \dots, \mathcal{W}_\nu \subseteq \mathcal{W}$ such that their union is \mathcal{W} and $\Delta(\mathcal{W}_k) > \tau$ for each k .*

Proof. Since the statement we are proving is invariant under periodic shifts, we can assume that $w_1 := 0 \in \mathcal{W}$. We sort the elements of \mathcal{W} by w_1, w_2, \dots, w_n sorted counterclockwise and provide a greedy method for generating the desired sets $\mathcal{W}_1, \mathcal{W}_2, \dots, \mathcal{W}_\nu$. We initialize these sets to be empty and we add $w_\ell \in \mathcal{W}$ to one of these sets until all elements of \mathcal{W} have been exhausted. We say w_ℓ has been *assigned* if it has been placed in a \mathcal{W}_k , and *unassigned* otherwise. We start by placing $w_1 \in \mathcal{W}_1$. For each unassigned $w_\ell \in \mathcal{W}$, we consider the set of $\mathcal{U}_\ell := \mathcal{W} \cap [w_\ell - \tau, w_\ell + \tau]$ and place w_ℓ in an arbitrary \mathcal{W}_k that does not contain any assigned elements in \mathcal{U}_ℓ . This is always possible since $|\mathcal{U}_\ell| \leq \nu$ for all ℓ . By construction, $\mathcal{W}_1, \dots, \mathcal{W}_\nu$ are disjoint and $\Delta(\mathcal{W}_k) > \tau$. To see why $\mathcal{W}_1, \mathcal{W}_2, \dots, \mathcal{W}_\nu$ are each nonempty, by definition of the τ density, there is a ℓ such that $[w_\ell - \tau, w_\ell + \tau]$ contains exactly ν elements of \mathcal{W} and they are necessarily placed in different \mathcal{W}_k ’s. \square

There is a stark conceptual distinction between the clumps decomposition in Definition 2.3, which groups the elements of \mathcal{X} by their spatial locations, versus Proposition 6.1, which decomposes \mathcal{X} into disjoint subsets that each have unit local sparsity. An example is shown in Figure 7.

The usefulness of this decomposition for controlling the condition number of Fourier matrices is not obvious, but it will be made more clear in the following proof.

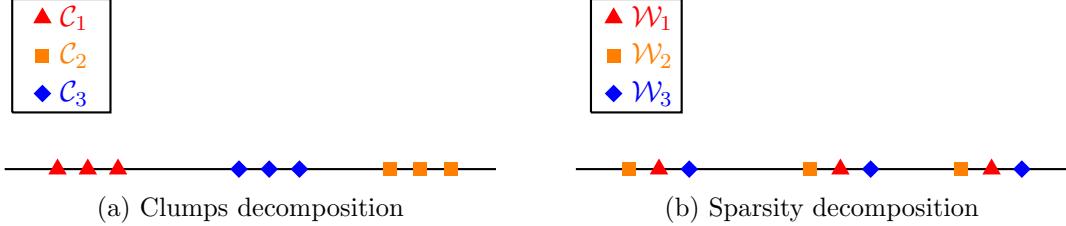


Figure 7: Clumps versus sparsity decomposition of the same set.

Proposition 6.2. *Let $\mathcal{G} \subseteq \mathbb{T}$ be a non-empty finite set such that for some $\tau \in (0, \frac{1}{2}]$, we have $|w|_{\mathbb{T}} > \tau$ for all $w \in \mathcal{G}$. Suppose $m, r \in \mathbb{N}_+$ such that $\nu(\tau, \mathcal{G}) \leq r$ and $m > \frac{1}{\tau}$. Then there is $f \in \mathcal{P}_{r(m-1)+1}$ such that $f(0) = 1$, f vanishes on \mathcal{G} , and*

$$\|f\|_{L^2} \leq \frac{1}{\sqrt{m}} \left(1 + \frac{1}{m\tau - 1}\right)^{r/2} \quad \text{and} \quad \|f\|_{L^\infty} \leq \left(1 + \frac{1}{m\tau - 1}\right)^{r/2}.$$

Proof. Let $\nu := \nu(\tau, \mathcal{G})$. By Proposition 6.1, there exists a disjoint decomposition,

$$\mathcal{G} = \mathcal{G}_1 \cup \mathcal{G}_2 \cup \dots \cup \mathcal{G}_\nu, \quad \text{where} \quad \mathcal{G}_k \neq \emptyset \quad \text{and} \quad \Delta(\mathcal{G}_k) > \tau.$$

The assumption that $|w|_{\mathbb{T}} > \tau$ for all $w \in \mathcal{G}$ implies $\Delta(\mathcal{G}_k \cup \{0\}) > \tau$. Using the assumption $m > \frac{1}{\tau}$, we invoke the lower bound in (1.1), which implies

$$\sigma_{\min}(\Phi(m, \mathcal{G}_k \cup \{0\})) \geq \sqrt{m - \frac{1}{\tau}}.$$

By Proposition 5.3, applied to the data points $(\{0\} \cup \mathcal{G}_k, e_1)$, there exists a $f_k \in \mathcal{P}_m$ such that

$$f_k(0) = 1, \quad f_k|_{\mathcal{G}_k} = 0, \quad \|f_k\|_{L^2} \leq \frac{1}{\sqrt{m}} \sqrt{\frac{m\tau}{m\tau - 1}} \quad \text{and} \quad \|f_k\|_{L^\infty} \leq \sqrt{\frac{m\tau}{m\tau - 1}}. \quad (6.1)$$

Let f be the product of f_1, f_2, \dots, f_ν . It follows immediately from (6.1) that $f(0) = 1$ and $f|_{\mathcal{G}} = 0$. The claimed bounds for $\|f\|_{L^2}$ and $\|f\|_{L^\infty}$ follow from Hölder's inequality. Moreover, we readily see that

$$\deg(f) = \sum_{k=1}^{\nu} \deg(f_k) \leq \nu(m-1) \leq r(m-1).$$

□

These polynomials can be numerically computed. First, we compute the decomposition of \mathcal{G} outlined in Proposition 6.1, which can be done constructively using the greedy method described in its proof. Second, for each \mathcal{G}_k in this decomposition, we find an interpolant f_k of the data $(\{0\} \cup \mathcal{G}_k, e_1)$ via Proposition 5.3. This can also be done numerically since \hat{f}_k is precisely a scaled left singular vector of $\Phi(m, \{0\} \cup \mathcal{G}_k)$, see the discussion following (5.1). Finally, these interpolants are then multiplied together to yield the desired f .

We refer to a f generated by this proposition as a *small norm Lagrange interpolant*. While each f_k is found by minimizing a L^2 norm with interpolation constraints, it is not necessarily true that f is also a minimum L^2 norm interpolant. Nonetheless, it is the pointwise bound that is important for this paper, and it is not clear if any of the f_k 's or f are extremal in the L^∞ norm.

The interpolant in Proposition 6.2 enjoys many favorable and surprising properties. First, in the absence of additional assumptions, it is degree-optimal. Notice that $\Delta(\{0\} \cup \mathcal{G}_k) > \tau$ and $0 \notin \mathcal{G}$ imply that $(|\mathcal{G}_k| + 1)\tau < 1$. This in turn establishes

$$|\mathcal{G}| \leq \nu \max_{1 \leq k \leq \nu} |\mathcal{G}_k| \leq r \left\lfloor \frac{1}{\tau} - 1 \right\rfloor.$$

This inequality is sharp since it is possible to provide an example of a \mathcal{G} such that these inequalities are achieved. On the other hand, with the only stipulation that $m > \frac{1}{\tau}$, the theorem provides a polynomial of degree $r(m-1)$ that has up to $r\lfloor \frac{1}{\tau} - 1 \rfloor$ zeros. Hence it is not possible to reduce the degree of this interpolant in general.

Aside from degree optimality, there is a second significance of small norm Lagrange interpolants, and to explain this, let us look at an example. For any $\varepsilon > 0$ that will be made sufficiently small, consider the set,

$$\mathcal{G}_\varepsilon = \left(\frac{1}{4} + \{0, \varepsilon, 2\varepsilon\} \right) \cup \left(\frac{1}{2} + \{0, \varepsilon\} \right) \cup \left(\frac{3}{4} + \{0, \varepsilon, \dots, 3\varepsilon\} \right). \quad (6.2)$$

Whenever ε is small enough, this set consists of 9 elements. In order to find a Lagrange interpolant that interpolates $(0, 1)$ and vanishes on \mathcal{G}_ε , we could use the usual Lagrange interpolant ℓ of degree $|\mathcal{G}_\varepsilon| = 9$, which satisfies the inequality

$$\|\ell\|_{L^\infty} \leq 2^{|\mathcal{G}_\varepsilon|} \prod_{w \in \mathcal{G}_\varepsilon} \frac{1}{|1 - e^{2\pi i w}|}.$$

One advantage of the Lagrange interpolant is that its degree does not depend on ε . However, $\|\ell\|_{L^\infty}$ grows exponentially in $|\mathcal{G}_\varepsilon|$ and the product term may be large, which makes it rather unappealing for our purposes; this could be an artifact of estimating its sup-norm by upper bounding each term individually, but it is difficult to circumvent.

On the other hand, there are interpolation methods that yield interpolants with larger degrees but smaller norms. For instance, [13, Theorem 2.1] proves that there is a trigonometric interpolant with good control over its norm, but requires the interpolant to have degree that scales inversely proportional to the minimum separation of the nodes. Hence, this result gives us an interpolant of any data defined on $\{0\} \cup \mathcal{G}_\varepsilon$ with degree that proportional to $\frac{1}{\varepsilon}$. Related interpolation results in [34] that are proved using functional analysis also exhibit similar behavior.

In contrast to the above types of interpolants, the small norm Lagrange interpolants enjoy both advantages. Set $\tau = \frac{1}{5}$ and for all sufficiently small ε , say $\varepsilon \leq \frac{1}{300}$, we have $\nu(\tau, \mathcal{G}_\varepsilon) = 4$. Pick $m = \frac{2}{\tau} = 10$ so that Proposition 6.2 is applicable. Notice that $\deg(f_\varepsilon) = 4 \cdot 9 = 36$ does not depend on ε , hence does not explode as $\varepsilon \rightarrow 0$. We also have $\|f_\varepsilon\|_{L^\infty} \leq 2^2 = 4$ for all ε . This is because the norm of f does not grow exponentially in $|\mathcal{G}_\varepsilon|$, but only grows exponentially in $\nu(\tau, \mathcal{G}_\varepsilon)$. This is crucial for the purposes of this paper, since we do not want high degree interpolants or large norms, while allowing the minimum separation to be arbitrary. Graphs of the real part of the four interpolants f_1, f_2, f_3, f_4 generated by Proposition 6.2 are shown in Figure 8.

We suspect that the advantages offered by Proposition 6.2 is highly dependent on the signs of the data that are being interpolated, that is, $1, 0, \dots, 0$ on $\{0\} \cup \mathcal{G}_\varepsilon$. The number ‘0’ is not necessarily crucial though since we can interpolate $1, a, \dots, a$ as well by multiplying the small norm Lagrange interpolants by $1 - a$ and then adding a . The important part here is that $0, \dots, 0$ are all equal. In contrast, the other interpolation methods work for arbitrary data, which is perhaps why they do not enjoy the same advantages.

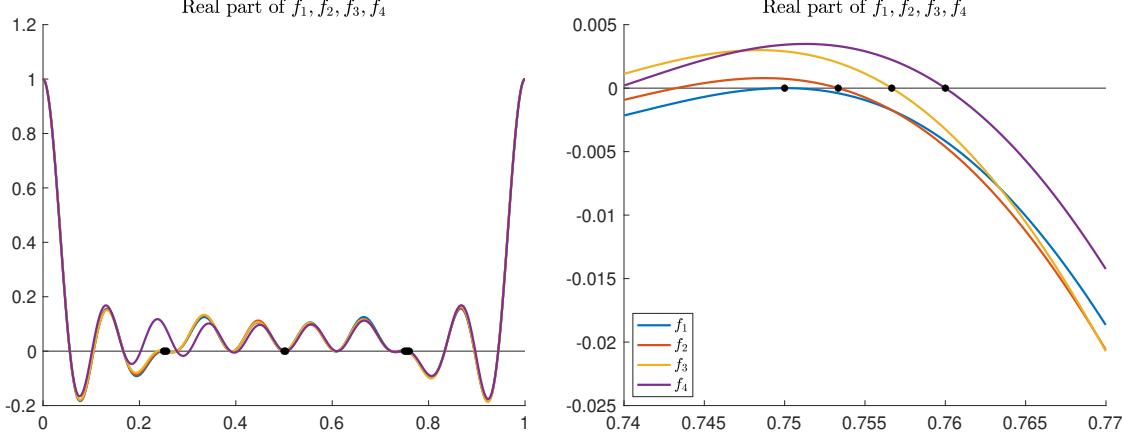


Figure 8: For the zero set \mathcal{G}_ε defined in (6.2) with $\varepsilon = \frac{1}{300}$, the four small norm Lagrange interpolants each of degree 9 are displayed.

6.2 Construction of local Lagrange polynomials

In this section, we study the interpolation problem for the “bad” set \mathcal{B} . It contains 0, all other elements in \mathcal{B} are close to zero, and we place no assumptions on $\Delta(\mathcal{B})$. We seek a trigonometric polynomial f such that $f(0) = 1$ and f vanishes on $\mathcal{B} \setminus \{0\}$.

Throughout this paper, let $\psi: [-\frac{1}{2}, \frac{1}{2}] \rightarrow \mathbb{R}$ be the sinc kernel restricted to $[-\frac{1}{2}, \frac{1}{2}]$,

$$\psi(t) := \begin{cases} \frac{\sin(\pi t)}{\pi t} & \text{if } t \neq 0, \\ 1 & \text{if } t = 0. \end{cases} \quad (6.3)$$

One can provide lower bounds for ψ via Taylor expansions, but the point here is that multiplication by $\psi(t)$ is negligible whenever $t \approx 0$. It naturally appears from the following calculation. For all $|t| \leq \frac{1}{2}$, we have

$$|1 - e^{2\pi i t}|^2 = 2 - 2 \cos(2\pi t) = 4\pi^2 t^2 \left(\frac{\sin(\pi t)}{\pi t} \right)^2 = 4\pi^2 t^2 \psi(t)^2. \quad (6.4)$$

We have the basic bound that $\psi(t) \geq \psi(\frac{1}{2}) = \frac{2}{\pi}$ since it is decreasing away from zero in its domain.

We have our first result for the bad set, which will be used in the proof of Theorem 1.

Lemma 6.3. *Suppose \mathcal{B} is a set of at most r points in \mathbb{T} with $0 \in \mathcal{B}$ and $\mathcal{B} \subseteq [-\tau, \tau]$ for some $\tau \in (0, \frac{1}{2}]$. For any $n \in \mathbb{N}_+$ such that $n \geq r$, define the subsets*

$$\mathcal{I} := \left\{ w \in \mathcal{B}: 0 < |w|_{\mathbb{T}} \leq \frac{r}{2n} \right\} \quad \text{and} \quad \mathcal{J} := \mathcal{B} \setminus (\mathcal{I} \cup \{0\}).$$

Then there exists a $f \in \mathcal{P}_n$ such that f vanishes on $\mathcal{B} \setminus \{0\}$, $f(0) = 1$, and

$$\|f\|_{L^2} \leq \frac{1}{\sqrt{\lfloor \frac{n}{r} \rfloor}} \prod_{w \in \mathcal{J}} \frac{1}{2 \lfloor \frac{1}{2|w|_{\mathbb{T}}} \rfloor} \prod_{w \in \mathcal{I}} \frac{1}{2 \lfloor \frac{n}{r} \rfloor} |w|_{\mathbb{T}}.$$

Proof. We first deal with the $\mathcal{B} = \{0\}$ case, in which case $\mathcal{I} = \mathcal{J} = \emptyset$. Then we define f to be a normalized Dirichlet kernel,

$$f(x) := \frac{1}{n} \sum_{\ell=0}^{n-1} e^{2\pi i \ell x}.$$

Notice that $f(0) = 1$ and $f \in \mathcal{P}_n$. Moreover, an application of Parseval establishes that for any $r \in \mathbb{N}_+$, we have

$$\|f\|_{L^2} \leq \frac{1}{\sqrt{n}} \leq \frac{1}{\sqrt{\lfloor \frac{n}{r} \rfloor}}.$$

This takes care of the $\mathcal{B} = \{0\}$ case.

From here onward, assume that $|\mathcal{B}| \geq 2$. We first deal with the case that $\mathcal{I} \neq \emptyset$. For each $w \in \mathcal{B}$, we define the natural number $q(w)$ where

$$q(w) := q := \left\lfloor \frac{n}{r} \right\rfloor \quad \text{if } w \in \mathcal{I}, \quad \text{and} \quad q(w) := \left\lfloor \frac{1}{2|w|_{\mathbb{T}}} \right\rfloor \quad \text{if } w \in \mathcal{J}.$$

We readily verify that $q(w)|w|_{\mathbb{T}} \leq \frac{1}{2}$ for all $w \in \mathcal{I} \cup \mathcal{J}$, which implies

$$|q(w)w|_{\mathbb{T}} = q(w)|w|_{\mathbb{T}} \quad \text{for each } w \in \mathcal{I} \cup \mathcal{J}. \quad (6.5)$$

In particular, this implies $|q(w)w|_{\mathbb{T}} \neq 0$ whenever $w \in \mathcal{I} \cup \mathcal{J}$. This enables us to define the polynomials,

$$h_0(x) := \prod_{w \in \mathcal{I}} \frac{e^{2\pi i q(w)x} - e^{2\pi i q(w)w}}{1 - e^{2\pi i q(w)w}}, \quad \text{and} \quad g(x) := \prod_{w \in \mathcal{J}} \frac{e^{2\pi i q(w)x} - e^{2\pi i q(w)w}}{1 - e^{2\pi i q(w)w}}.$$

By construction, h_0 vanishes on \mathcal{I} and $h_0(0) = 1$. For each $\ell = 1, 2, \dots, q-1$, we define the functions

$$h_{\ell}(x) := e^{2\pi i \ell x} h_0(x) \quad \text{and} \quad h := \frac{1}{q} \sum_{\ell=0}^{q-1} h_{\ell}.$$

Using that $|h_{\ell}| = |h_0|$, we see that $h(0) = 1$ and h also vanishes on \mathcal{I} . Note that h_{ℓ} is a trigonometric polynomial whose frequencies are in $\ell + \{0, q, \dots, |\mathcal{I}|q\}$. This implies that $\{h_{\ell}\}_{\ell=0}^{q-1}$ are L^2 orthogonal, $\deg(h) \leq |\mathcal{I}|q + q - 1$, and by orthogonality,

$$\|h\|_{L^2}^2 = \frac{1}{q^2} \sum_{\ell=1}^{q-1} \|h_{\ell}\|_{L^2}^2 = \frac{1}{q^2} \sum_{\ell=1}^{q-1} \|h_0\|_{L^2}^2 \leq \frac{1}{q} \|h_0\|_{L^{\infty}}^2. \quad (6.6)$$

We define $f := gh$. By construction, $f(0) = 1$ and f vanishes on $\mathcal{B} \setminus \{0\}$. We bound the degree of f . Note for each $w \in \mathcal{J}$, we have $|w|_{\mathbb{T}} > \frac{r}{n}$ and so $q(w) \leq \frac{n}{2r}$. This implies, together with the assumption $|\mathcal{B}| \leq r$, that

$$\deg(f) = |\mathcal{I}|q + q - 1 + \sum_{w \in \mathcal{J}} q(w) \leq \frac{n(|\mathcal{I}| + 1)}{r} - 1 + \frac{n|\mathcal{J}|}{r} \leq \frac{n|\mathcal{B}|}{r} - 1 \leq n - 1.$$

It remains to obtain the desired bound for $\|f\|_{L^2}$. Using (6.6), we get

$$\|f\|_{L^2} \leq \|h\|_{L^2} \|g\|_{L^{\infty}} \leq \frac{1}{\sqrt{q}} \|h_0\|_{L^{\infty}} \|g\|_{L^{\infty}} \leq \frac{1}{\sqrt{q}} \prod_{w \in \mathcal{I} \cup \mathcal{J}} \frac{2}{|1 - e^{2\pi i q(w)w}|}.$$

Combining this with (6.4), (6.5), and that $\psi(t) \geq \frac{2}{\pi}$ for $|t| \leq \frac{1}{2}$, we obtain

$$\|f\|_{L^2} \leq \frac{1}{\sqrt{q}} \prod_{w \in \mathcal{B}} \frac{1}{\pi|q(w)w|_{\mathbb{T}} \psi(|q(w)w|_{\mathbb{T}})} \leq \frac{1}{\sqrt{q}} \prod_{w \in \mathcal{B}} \frac{1}{2q(w)|w|_{\mathbb{T}}}.$$

Using the definition of $q(w)$ yields the claimed upper bound for $\|f\|_{L^2}$ when $\mathcal{I} \neq \emptyset$.

Finally, for the remaining case where $|\mathcal{B}| \geq 2$ and $\mathcal{I} = \emptyset$, we use the same g as above but with a different h . We instead use a normalized Dirichlet kernel,

$$h(x) := \frac{1}{\lfloor n/r \rfloor} \sum_{\ell=0}^{\lfloor n/r \rfloor - 1} e^{2\pi i \ell x}.$$

We have $h \in \mathcal{P}_{\lfloor n/r \rfloor}$ and that $\|h\|_{L^2} = 1/\sqrt{\lfloor n/r \rfloor}$. Setting $f := gh$, we see that

$$\deg(f) = \deg(g) + \deg(h) \leq |\mathcal{J}| \frac{n}{r} + \frac{n}{r} - 1 \leq n - 1,$$

where the final inequality used that since $\mathcal{I} = \emptyset$ and $0 \in \mathcal{B}$, we have $|\mathcal{J}| = |\mathcal{B}| - 1 \leq r - 1$. Hence, $f \in \mathcal{P}_n$ and it satisfies the desired interpolation properties. Combining the previous upper bounds for $\|g\|_{L^\infty}$ and $\|h\|_{L^2}$ completes the proof. \square

The following is our second result for the bad set, which will be used in the proof of Theorem 2.

Lemma 6.4. *Let $n, r \in \mathbb{N}_+$ such that $n \geq r$ and $\delta \in (0, \frac{1}{n}]$. Suppose $\mathcal{B} \subseteq \mathbb{T}$ is a finite set such that $0 \in \mathcal{B}$, $|\mathcal{B}| = r$ and $\delta \leq \Delta(\mathcal{B}) \leq \frac{1}{n}$. Then there exists a $f \in \mathcal{P}_n$ such that f vanishes on $\mathcal{B} \setminus \{0\}$, $f(0) = 1$, and*

$$\|f\|_{L^2} \leq \frac{2e}{\pi r} \frac{1}{\sqrt{\lfloor \frac{n}{r} \rfloor}} \left(\frac{4e}{\pi \psi(\frac{n\delta}{2}) r \lfloor \frac{n}{r} \rfloor \delta} \right)^{r-1}.$$

Proof. We define the subsets,

$$\mathcal{I} := \left\{ w \in \mathcal{B} : 0 < |w|_{\mathbb{T}} \leq \frac{r}{2n} \right\} \quad \text{and} \quad \mathcal{J} := \mathcal{B} \setminus (\mathcal{I} \cup \{0\}).$$

We first deal with the $\mathcal{B} = \{0\}$ case, in which case $r = 1$ and $\mathcal{I} = \mathcal{J} = \emptyset$. Similar to the proof of Lemma 6.3, we define $f(x) := \frac{1}{n} \sum_{\ell=0}^{n-1} e^{2\pi i \ell x}$. Then $f(0) = 1$, $f \in \mathcal{P}_n$, and $\|f\|_{L^2} \leq \frac{1}{\sqrt{n}}$. Since $\frac{2e}{\pi} \geq 1$, this proves the $\mathcal{B} = \{0\}$ case.

From here onward assume that $r \geq 2$. We enumerate the elements of \mathcal{B} as $0 = w_0, w_1, \dots, w_{r-1}$ where $|w_k|_{\mathbb{T}} \leq |w_{k+1}|_{\mathbb{T}}$ for each k . For reasons that will become apparent later, for any $d \in \mathbb{N}_+$, we define the following sequence, $0, -1, 1, -2, 2, -3, 3, \dots$, which we enumerate by a_0, a_1, \dots . We define the natural numbers q_1, \dots, q_{r-1} as

$$q_k := q := \left\lfloor \frac{n}{r} \right\rfloor \quad \text{if } w_k \in \mathcal{I}, \quad \text{and} \quad q_k := \left\lfloor \frac{q|a_k|\delta}{|w_k|_{\mathbb{T}}} \right\rfloor \quad \text{if } w_k \in \mathcal{J}.$$

We need to set the stage before we explicitly construct f . For each $w_k \in \mathcal{I}$, we immediately get $q_k|w_k|_{\mathbb{T}} \leq \frac{1}{2}$ by definition of q and \mathcal{I} . For each $w_k \in \mathcal{J}$, we use that $|a_k| \leq \frac{r}{2}$ regardless of the parity of r and the assumption $\delta \leq \frac{1}{n}$ to see that $q_k|w_k|_{\mathbb{T}} \leq q|a_k|\delta \leq \frac{1}{2}$. This implies

$$|q_k w_k|_{\mathbb{T}} = q_k |w_k|_{\mathbb{T}} \quad \text{for each } w_k \in \mathcal{B} \setminus \{0\}. \tag{6.7}$$

This enables us to define the polynomials,

$$h_0(x) := \prod_{w \in \mathcal{I}} \frac{e^{2\pi i q w x} - e^{2\pi i q w}}{1 - e^{2\pi i q w}}, \quad \text{and} \quad g(x) := \prod_{w_k \in \mathcal{J}} \frac{e^{2\pi i q_k w_k x} - e^{2\pi i q_k w_k}}{1 - e^{2\pi i q_k w_k}}.$$

We repeat the same sub-argument that appeared in the proof of Lemma 6.3. We define the function $h(x) := \frac{1}{q} \sum_{\ell=0}^{q-1} e^{2\pi i \ell x} h_0(x)$, and we see that $\|h\|_{L^2} \leq \frac{1}{\sqrt{q}} \|h_0\|_{L^\infty}$. Thus, we define $f = gh$ and so

$$\|f\|_{L^2} \leq \frac{1}{\sqrt{q}} \|h_0\|_{L^\infty} \|g\|_{L^\infty} \leq \frac{1}{\sqrt{q}} \prod_{w \in \mathcal{I}} \frac{2}{|1 - e^{2\pi i q w}|} \prod_{w_k \in \mathcal{J}} \frac{2}{|1 - e^{2\pi i q_k w_k}|}. \quad (6.8)$$

By construction, f satisfies the desired interpolation properties. We argue that $f \in \mathcal{P}_n$. Notice that $\deg(h_0) = q|\mathcal{I}|$ and $\deg(h) = q|\mathcal{I}| + q - 1$. On the other hand, for each $w_k \in \mathcal{J}$, using that $|a_k| \leq \frac{r}{2}$, $\delta \leq \frac{1}{n}$, and $|w_k|_{\mathbb{T}} > \frac{r}{2n}$, we see that $q_k \leq q$. Thus,

$$\deg(f) = q(|\mathcal{I}| + 1) - 1 + \sum_{w_k \in \mathcal{J}} q_k \leq q(|\mathcal{I}| + |\mathcal{J}| + 1) - 1 \leq \frac{n|\mathcal{B}|}{r} - 1 \leq n - 1.$$

It remains to upper bound $\|f\|_{L^2}$. By (6.7), we see that $|q_k w_k|_{\mathbb{T}} = q_k |w_k|_{\mathbb{T}} \geq q |a_k| \delta / 2$. Using this inequality on the right side of (6.8), we get

$$\|f\|_{L^2} \leq \frac{1}{\sqrt{q}} \prod_{w \in \mathcal{I}} \frac{2}{\sqrt{2 - 2 \cos(2\pi q|w|_{\mathbb{T}})}} \prod_{w_k \in \mathcal{J}} \frac{2}{\sqrt{2 - 2 \cos(2\pi q|a_k|\delta/2)}}. \quad (6.9)$$

We let $\ell := |\mathcal{I}|$ and claim that

$$\prod_{w \in \mathcal{I}} \frac{2}{\sqrt{2 - 2 \cos(2\pi q|w|_{\mathbb{T}})}} \leq \prod_{k=1}^{\ell} \frac{2}{\sqrt{2 - 2 \cos(2\pi q|a_k|\delta)}} \quad (6.10)$$

Recall that $\Delta(\mathcal{I} \cup \{0\}) \geq \delta$ and that $0 \notin \mathcal{I}$. We define the auxiliary function,

$$\gamma(t_1, t_2, \dots, t_\ell) := \sum_{k=1}^{\ell} \frac{1}{1 - \cos(2\pi q|t_k|)},$$

where $\delta \leq |t_k| \leq \frac{1}{2q}$ for each k and $|t_j - t_k| \geq \delta$ for each $j \neq k$. Clearly this function increases if any t_j is made smaller while the remaining t_k 's are fixed. We claim that γ is maximized precisely when t_1, t_2, \dots, t_ℓ is $a_1\delta, a_2\delta, \dots, a_\ell\delta$. To see this, we list t_1, t_2, \dots, t_ℓ as $\{u_1, \dots, u_a, v_1, \dots, v_b\}$ where $a + b = \ell$ and

$$u_a < \dots < u_1 \leq -\delta < 0 < \delta \leq v_1 < \dots < v_b.$$

If $b \geq 1$, we can assume that $v_1 = \delta$ since a shift of all v_1, \dots, v_b by the same amount towards 0 increases the value of γ . If $a \geq 1$, we can likewise assume that $u_1 = -\delta$. Finally, γ is further increased if all the u 's and v 's are fixed except v_2 is replaced with 2δ , then v_3 is moved to 3δ , etc. Likewise, γ is increased if u_2 is replaced to -2δ , etc. Hence, we see that $\gamma(t_1, t_2, \dots, t_\ell)$ is dominated by $\gamma(-a\delta, \dots, -\delta, \delta, \dots, b\delta)$. If $|a - b| > 1$, then by reflecting elements across the origin and shifting again, we see that $\gamma(-a\delta, \dots, -\delta, \delta, \dots, b\delta)$ is further dominated by $\gamma(a_1\delta, a_2\delta, \dots, a_\ell\delta)$. This establishes inequality (6.10).

We continue with the upper bound for $\|f\|_{L^2}$. Note that $q|a_k|\delta \leq \frac{n\delta}{2}$ for each $k = 1, \dots, r-1$ and that ψ is decreasing on $[0, \frac{1}{2}]$. Using (6.4), (6.9), and (6.10), we see that

$$\|f\|_{L^2} \leq \frac{1}{\sqrt{q}} \prod_{k=1}^{\ell} \frac{1}{\pi\psi(q|a_k|\delta)q|a_k|\delta} \prod_{w_k \in \mathcal{J}} \frac{2}{\pi\psi(q|a_k|\delta/2)q|a_k|\delta} \leq \frac{1}{\sqrt{q}} \prod_{k=1}^{r-1} \frac{2}{\pi\psi(\frac{n\delta}{2})q|a_k|\delta}. \quad (6.11)$$

To control the product over k , first note that

$$C_d := \prod_{k=1}^d |a_k| = \begin{cases} \left(\frac{d}{2}!\right)\left(\frac{d}{2}!\right) & \text{if } d \text{ is even,} \\ \left(\frac{d+1}{2}!\right)\left(\frac{d-1}{2}!\right) & \text{if } d \text{ is odd.} \end{cases}$$

Recall the well known inequalities $k! \geq \sqrt{2\pi k} \left(\frac{k}{e}\right)^k$ and $1+t \leq e^t$ for all $t \geq 0$. We have $C_1 = 1$, and for $d \geq 2$, we have

$$\begin{aligned} \frac{(d+1)^d}{C_d} &\leq \frac{(2e)^d}{\pi d} \left(\frac{d+1}{d}\right)^d \leq \frac{(2e)^d 2e}{\pi(d+1)} \quad \text{if } d \text{ is even,} \\ \frac{(d+1)^d}{C_d} &= \frac{2(d+1)^{d-1}}{\left(\frac{d-1}{2}!\right)^2} \leq \frac{2(2e)^{d-1}}{\pi(d-1)} \left(\frac{d+1}{d-1}\right)^{d-1} \leq \frac{2(2e)^{d-1} e^2}{\pi(d-1)} \leq \frac{(2e)^d 2e}{\pi(d+1)} \quad \text{if } d \text{ is odd.} \end{aligned}$$

Using these and the definition of q in (6.11) completes the proof. \square

It may be worthwhile to mention a subtle technical part of this proof. We do not prove, or claim, that the “worst case” \mathcal{B} (up to trivial invariances) satisfying the hypotheses of Lemma 6.4 is

$$\delta \left\{ -\left\lfloor \frac{r-1}{2} \right\rfloor, \dots, -1, 0, 1, \dots, \left\lceil \frac{r-1}{2} \right\rceil \right\}. \quad (6.12)$$

In this proof, we showed that for any \mathcal{B} satisfying the hypotheses, there is an explicit Lagrange interpolant f with some upper bound for $\|f\|_{L^2}$. Then we showed that this upper bound for $\|f\|_{L^2}$ is maximized precisely when \mathcal{B} is equal to (6.12). This is the quantity reported in the conclusion of Lemma 6.4.

7 Proofs of the main results

7.1 Proof of Theorem 1

Proof. Let us first set the stage and discuss several immediate implications of the assumptions. Fix any $1 \leq k \leq s$ and for convenience, we define

$$n_k := \left\lfloor m - \frac{2\nu(\tau, \mathcal{G}_k)}{\tau} \right\rfloor.$$

Note that $\nu(\tau, \mathcal{G}_k) \leq \nu(\tau, \mathcal{X})$ since $\mathcal{G}_k \subseteq \mathcal{X}$. Also using the assumptions $m \geq 6s$ and $3\nu(\tau, \mathcal{X}) \leq \tau m$, we have

$$m - \frac{2\nu(\tau, \mathcal{G}_k)}{\tau} \geq m - \frac{2\nu(\tau, \mathcal{X})}{\tau} \geq \frac{m}{3} \geq 2s. \quad (7.1)$$

As immediate consequences of this inequality, we have $n_k \in \mathbb{N}_+$ and that $\alpha_k > 0$. Since $\nu(\tau, \mathcal{B}_k) \leq \nu(\tau, \mathcal{X})$ due to $\mathcal{B}_k \subseteq \mathcal{X}$, we use the assumption $3\nu(\tau, \mathcal{X}) \leq \tau m$ to see that

$$\alpha_k \leq \frac{\nu(\tau, \mathcal{X})}{2m - 4\nu(\tau, \mathcal{X})\tau^{-1}} \leq \frac{\nu(\tau, \mathcal{X})}{6\nu(\tau, \mathcal{X})\tau^{-1} - 4\nu(\tau, \mathcal{X})\tau^{-1}} \leq \frac{\tau}{2}. \quad (7.2)$$

We first deal with the “good” set \mathcal{G}_k . If $\mathcal{G}_k = \emptyset$, then $\nu(\tau, \mathcal{G}_k) = 0$ and we set $g_k := 1$. Otherwise, we assume $\mathcal{G}_k \neq \emptyset$. We apply Proposition 6.2, where $\mathcal{G}_k - x_k$, $\lceil \frac{2}{\tau} \rceil$, and $\nu(\tau, \mathcal{G}_k)$ play the roles of \mathcal{B} , m and r respectively, in the referenced proposition’s notation. This provides us with a polynomial,

which after shifting by x_k , we call it $g_k \in \mathcal{P}_{\nu(\tau, \mathcal{G}_k)(\lceil \frac{2}{\tau} \rceil - 1) + 1}$ such that $g_k(x_k) = 1$, g_k vanishes on \mathcal{G}_k , and

$$\|g_k\|_{L^\infty} \leq \left(1 + \frac{1}{\lceil \frac{2}{\tau} \rceil \tau - 1}\right)^{\nu(\tau, \mathcal{G}_k)/2} \leq 2^{\nu(\tau, \mathcal{G}_k)/2}. \quad (7.3)$$

Note that this statement is still valid in the corner case that $\mathcal{G}_k = \emptyset$ since in this case, we have $\nu(\tau, \mathcal{G}_k) = 0$, which is consistent with $\|g_k\|_{L^\infty} = 1$ and $g_k \in \mathcal{P}_1$.

Now we deal with the “bad” set $\mathcal{B}_k := \mathcal{B}(x_k, \tau, \mathcal{X})$. We use the shorthand notation $r_k := |\mathcal{B}_k|$. We are ready to employ Lemma 6.3, where $\mathcal{B}_k - x_k$, n_k , and r_k play the roles of \mathcal{B} , n , and r respectively, in the referenced lemma’s notation. Note that $n_k \geq s$ from (7.1). The lemma provides us with a polynomial, and after shifting by x_k , we call it $b_k \in \mathcal{P}_{n_k}$ such that b_k vanishes on $\mathcal{B}_k \setminus \{x_k\}$, $b_k(x_k) = 1$, and b_k enjoys the estimates,

$$\|b_k\|_{L^2} \leq \sqrt{\frac{1}{\lceil \frac{n_k}{r_k} \rceil}} \prod_{x \in \mathcal{J}_k} \frac{1}{2^{\lfloor \frac{1}{2|x-x_k|_{\mathbb{T}}} \rfloor} |x - x_k|_{\mathbb{T}}} \prod_{x \in \mathcal{I}_k} \frac{1}{2^{\lfloor \frac{n_k}{r_k} \rfloor} |x - x_k|_{\mathbb{T}}}.$$

Now we perform some algebraic manipulations and simplifications. First note that $\frac{1}{\alpha_k} \leq \frac{2n_k+2}{r_k}$, and $\lfloor \frac{n_k}{r_k} \rfloor \geq \frac{n_k}{r_k} - \frac{r_k-1}{r_k}$ since $r_k \in \mathbb{N}_+$. Together, they imply that

$$\frac{1}{2^{\lfloor \frac{n_k}{r_k} \rfloor}} \leq \frac{1}{\frac{1}{\alpha_k} - 2} = \frac{\alpha_k}{1 - 2\alpha_k}.$$

Using this observation and the definition of ϕ in the previous upper bound for $\|b_k\|_{L^2}$, we have

$$\|b_k\|_{L^2} \leq \sqrt{\frac{2\alpha_k}{1 - 2\alpha_k}} \prod_{x \in \mathcal{J}_k} \phi\left(\frac{1}{2|x - x_k|_{\mathbb{T}}}\right) \prod_{x \in \mathcal{I}_k} \frac{\alpha_k}{(1 - 2\alpha_k)|x - x_k|_{\mathbb{T}}}. \quad (7.4)$$

We next define $f_k := b_k g_k$. We have $f_k \in \mathcal{P}_m$ because

$$\deg(f_k) = \deg(b_k) + \deg(g_k) \leq \nu(\tau, \mathcal{G}_k) \left(\left\lceil \frac{2}{\tau} \right\rceil - 1\right) + n_k - 1 \leq \frac{2\nu(\tau, \mathcal{G}_k)}{\tau} + n_k - 1 \leq m - 1.$$

Together with the interpolation properties of g_k and b_k , we see that $\{f_k\}_{k=1}^s \subseteq \mathcal{P}_m$ is a family of Lagrange interpolants for \mathcal{X} . We use Hölder’s inequality and the upper bounds (7.3) and (7.4) to get

$$\|f_k\|_{L^2} \leq \sqrt{\frac{2\alpha_k 2^{\nu(\tau, \mathcal{G}_k)}}{1 - 2\alpha_k}} \prod_{x \in \mathcal{J}_k} \phi\left(\frac{1}{2|x - x_k|_{\mathbb{T}}}\right) \prod_{x \in \mathcal{I}_k} \frac{\alpha_k}{(1 - 2\alpha_k)|x - x_k|_{\mathbb{T}}}.$$

We apply Lemma 5.2 to complete the proof of (2.2).

Now we proceed to further upper bound the right side of (2.2). Using that $\phi(t) \leq 2$ for all $t \geq 1$, that $\alpha_k \leq \frac{\tau}{2} \leq \frac{1}{4}$ due to (7.2), we obtain

$$\begin{aligned} \frac{1}{\sigma_s^2(\Phi(m, \mathcal{X}))} &\leq \sum_{k=1}^s 2^{\nu(\tau, \mathcal{G}_k)} 4\alpha_k 4^{|\mathcal{J}_k|} \prod_{x \in \mathcal{I}_k} \frac{4\alpha_k^2}{|x - x_k|_{\mathbb{T}}^2} \\ &\leq s \max_{1 \leq k \leq s} \left\{ 2^{\nu(\tau, \mathcal{G}_k)} 4\alpha_k 4^{|\mathcal{J}_k|} \prod_{x \in \mathcal{I}_k} \frac{4\alpha_k^2}{|x - x_k|_{\mathbb{T}}^2} \right\}. \end{aligned}$$

Rearranging this inequality and taking the square root completes the proof of (2.3). \square

7.2 Proof of Theorem 2

Proof. Fix any $1 \leq k \leq s$. Note (7.1) showed that $n_k \geq \lfloor \frac{m}{3} \rfloor$, while $n_k \leq m$ immediately by definition. The proof is analogous to the proof of Theorem 1, but with a different function for the bad set. We carry over the same definitions of n_k and r_k . There we constructed the function $g_k \in \mathcal{P}_{\nu(\tau, \mathcal{G}_k)(\lceil \frac{2}{\tau} \rceil - 1) + 1}$ for the “good” set \mathcal{G}_k such that

$$\|g_k\|_{L^\infty} \leq 2^{\nu(\tau, \mathcal{G}_k)/2}. \quad (7.5)$$

For the “bad” set $\mathcal{B}_k := \mathcal{B}(x_k, \tau, \mathcal{X})$, we note that $\Delta(\mathcal{B}_k) \geq \Delta(\mathcal{X}) \geq \delta$. We use Lemma 6.4, where $\mathcal{B}_k - x_k$, n_k and r_k play the roles of \mathcal{B} , n , and r respectively, in the referenced lemma’s notation. Also note that $n_k \geq s \geq r_k$ due to (7.1), and that $\delta \leq \frac{1}{m} \leq \frac{1}{n_k}$. The lemma provides us with a polynomial, and after shifting by x_k , we call it $b_k \in \mathcal{P}_{n_k}$ such that b_k vanishes on $\mathcal{B}_k \setminus \{x_k\}$, $b_k(x_k) = 1$, and b_k enjoys the estimate,

$$\|b_k\|_{L^2} \leq \frac{2e}{\pi r_k} \frac{1}{\sqrt{\lfloor \frac{n_k}{r_k} \rfloor}} \left(\frac{4e}{\pi \psi(\frac{n_k \delta}{2}) r_k \lfloor \frac{n_k}{r_k} \rfloor \delta} \right)^{r_k-1}. \quad (7.6)$$

We define $f_k = g_k b_k$. Repeating the same argument as in the proof of Theorem 1 shows that $f_k \in \mathcal{P}_m$. By construction, $\{f_k\}_{k=1}^s$ is a family of Lagrange polynomials for \mathcal{X} . Using Hölder’s inequality, (7.5), (7.6), and the definition of ϕ , we get

$$\|f_k\|_{L^2} \leq \frac{2e}{\pi} \sqrt{\frac{2^{\nu(\tau, \mathcal{G}_k)} \phi(\frac{n_k}{r_k})}{r_k n_k}} \left(\frac{4e \phi(\frac{n_k}{r_k})}{\pi \psi(\frac{n_k \delta}{2}) n_k \delta} \right)^{r_k-1}.$$

Finally, using Lemma 5.2 completes the proof of (2.4).

We proceed to make numerous simplifications of the right hand side of (2.4). Since $\delta \leq \frac{1}{m}$ and $n_k \leq m$, we have $\psi(\frac{n_k \delta}{2}) \geq \psi(\frac{1}{2}) = \frac{2}{\pi}$. Note that (7.1) and the assumption $m \geq 6s$ imply

$$r_k \left\lfloor \frac{n_k}{r_k} \right\rfloor \geq r_k \left(\frac{n_k}{r_k} - \frac{r_k - 1}{r_k} \right) = n_k - r_k + 1 \geq m - \frac{2\nu(\tau, \mathcal{X})}{\tau} - r_k \geq \frac{m}{3} - s \geq \frac{m}{6}.$$

Using these observations in (2.4) now establishes

$$\frac{1}{\sigma_s^2(\Phi(m, \mathcal{X}))} \leq \frac{24e^2}{\pi^2 m} \sum_{k=1}^s \frac{2^{\nu(\tau, \mathcal{G}_k)}}{r_k} \left(\frac{12e}{m \delta} \right)^{2r_k-2} \leq \frac{24e^2 s}{\pi^2 m} \max_{1 \leq k \leq s} \left\{ \frac{2^{\nu(\tau, \mathcal{G}_k)}}{r_k} \left(\frac{12e}{m \delta} \right)^{2r_k-2} \right\}.$$

Rearranging this inequality and taking the square root completes the proof of (2.5). \square

7.3 Proof of Corollary 1

Proof. We first claim that \mathcal{X} satisfies the (m, τ) density criterion. This trivially holds when $r = 1$ because then $\tau = \frac{1}{2}$ and so $\nu(\tau, \mathcal{X}) = |\mathcal{X}| = s = \lambda$. From here onward, assume that $r > 1$. For any $x \in \mathcal{X}$, let \mathcal{C}_x be the clump that x belongs to. Since $\beta > \alpha$ by definition, we see that

$$\mathcal{C}_x = \{x' \in \mathcal{X} : |x - x'|_{\mathbb{T}} \leq \beta\},$$

otherwise it would contradict the assumption that any two clumps are separated by distances strictly larger than β and that $\text{diam}(\mathcal{C}_x) \leq \alpha < \beta$. This shows that $\nu(\tau, \mathcal{X}) \leq \lambda$, and since there

is a clump that has cardinality exactly equal to λ , we see that $\nu(\tau, \mathcal{X}) = \lambda$. We have $\beta \geq \frac{3\lambda}{m}$ by assumption, and so

$$\frac{3\nu(\tau, \mathcal{X})}{\tau} = \frac{3\lambda}{\beta} \leq m.$$

We have shown that \mathcal{X} satisfies the (m, τ) density criterion. This shows that the assumptions of Theorem 1 and Theorem 2 are satisfied. For each k in the right side of (2.5), we use that $r_k \leq \nu(\tau, \mathcal{X}) = \lambda$ and $\nu(\tau, \mathcal{G}_k) \leq \nu(\tau, \mathcal{X}) = \lambda$ to complete the proof. \square

7.4 Proof of Theorem 3

Proof. Letting $\nu := \nu(\tau, \mathcal{X})$, by the decomposition given in Proposition 6.1, we have a disjoint union

$$\mathcal{X} = \mathcal{X}_1 \cup \mathcal{X}_2 \cup \dots \cup \mathcal{X}_\nu, \quad \text{where } \Delta(\mathcal{X}_k) \geq \tau \quad \text{for each } k = 1, 2, \dots, \nu.$$

Since the singular values of $\Phi(m, \mathcal{X})$ are invariant under permutations of its columns, after reshuffling,

$$\Phi(m, \mathcal{X}) = [\Phi(m, \mathcal{X}_1) \ \Phi(m, \mathcal{X}_2) \ \dots \ \Phi(m, \mathcal{X}_\nu)].$$

Let $u \in \mathbb{C}^{|\mathcal{X}|}$ be any unit norm vector, and likewise, we partition u into sub-vectors u_1, u_2, \dots, u_ν such that $u_k \in \mathbb{C}^{|\mathcal{X}_k|}$. Since $\Delta(\mathcal{X}_k) \geq \tau$ for each k and $m > \frac{1}{\tau}$, we use the upper bound in (1.1) to get

$$\|\Phi(m, \mathcal{X})u\|_2 = \left\| \sum_{k=1}^{\nu} \Phi(m, \mathcal{X}_k)u_k \right\|_2 \leq \sqrt{m + \frac{1}{\tau}} \sum_{k=1}^{\nu} \|u_k\|_2.$$

Using Cauchy-Schwarz and that u has unit norm, we obtain

$$\sum_{k=1}^{\nu} \|u_k\|_2 \leq \sqrt{\nu} \left(\sum_{k=1}^{\nu} \|u_k\|_2^2 \right)^{1/2} = \sqrt{\nu} \|u\|_2 = \sqrt{\nu}.$$

Combining the above inequalities completes the proof. \square

Computational costs and benefits

Selecting reasonable τ to use in the main theorems can be done computationally. Let $\mathcal{T} \subseteq (0, \frac{1}{2}]$ be a collection of τ for which we would like to evaluate the main estimates. We proceed to analyze the naive time and storage complexity of evaluating the main estimates for each $\tau \in \mathcal{T}$. All of the main inequalities are simple expressions that can be evaluated once the various parameters are determined, which importantly, only depend on τ and distances between elements in \mathcal{X} .

As overhead, we first compute the distances between all $\binom{s}{2}$ distinct pairs of elements in \mathcal{X} , which requires $O(s^2)$ operations and storage. As we will see, computation of the local sparsity is the main bottleneck. For an arbitrary and finite set $\mathcal{U} \subseteq \mathbb{T}$, computation of $\nu(\tau, \mathcal{U})$ requires $O(|\mathcal{U}|^2)$ operations since we need to enumerate through each $u \in \mathcal{U}$ and find the number of $u' \in \mathcal{U}$ such that $|u - u'| \leq \tau$.

Coming back to the task at hand, we perform a first loop over each $x_k \in \mathcal{X}$ and $\tau \in \mathcal{T}$, in order to compute $|\mathcal{B}(x_k, \tau, \mathcal{X})|$ and $\nu(\tau, \mathcal{G}(x_k, \tau, \mathcal{X}))$. This loop requires $O(s^3|\mathcal{T}|)$ time and $O(s|\mathcal{T}|)$ storage. After the first loop executes, we perform a second loop through each $x_k \in \mathcal{X}$ and $\tau \in \mathcal{T}$. In this loop, we calculate the necessary parameters (e.g., α_k , n_k) and each of the s terms in the

main inequalities. This second loop requires $O(s^2|\mathcal{T}|)$ time for Theorem 1 and $O(s|\mathcal{T}|)$ time for Theorem 2. Then we perform a final minimization or summation (depending on which inequality is evaluated), and then find the the optimal $\tau \in \mathcal{T}$. Ultimately, computation of the optimal $\tau \in \mathcal{T}$ for any of these theorems requires $O(s^3|\mathcal{T}|)$ operations and $O(s^2 + |\mathcal{T}|)$ storage.

Let us quickly remark that these complexities are for arbitrary \mathcal{X} and \mathcal{T} , and the computations can be sped up under special cases. Suppose \mathcal{X} consists of clumps and $\mathcal{T} \subseteq [\alpha, \beta]$, where α and β are clump parameters as in Definition 2.3. Notice that $|\mathcal{B}(x_k, \tau, \mathcal{X})|$ is the cardinality of the clump that contains x_k , while $\mathcal{G}(x_k, \tau, \mathcal{X})$ consists of all the other clumps that do not contain x_k , and $\nu(\tau, \mathcal{G}(x_k, \tau, \mathcal{X}))$ is the cardinality of the largest clump that does not contain x_k . All of them do not vary with $\tau \in \mathcal{T}$. So these quantities can be computed with $O(s^2)$ operations instead of $O(s^3|\mathcal{T}|)$.

Another possibility of improvement is to take advantage of redundancies in these calculations when looping through $\tau \in \mathcal{T}$. Notice that $|\mathcal{B}(x_k, \tau, \mathcal{X})|$ and $\nu(\tau, \mathcal{G}(x_k, \tau, \mathcal{X}))$ are piece-wise constant functions with at most $s - 1$ break points, so there is no need to enumerate through \mathcal{T} . So the $|\mathcal{T}|$ terms in the above complexities can be replaced with s^2 whenever $|\mathcal{T}| \geq s^2$. However, in our experience, we usually have a good sense for what range of τ is appropriate, so only a few different τ need to be evaluated. Moreover, the loop over \mathcal{T} can be parallelized to further reduce the scaling in $|\mathcal{T}|$.

In comparison, the singular value decomposition of $\Phi(m, \mathcal{X})$ requires $O(m^2s)$ operations and $O(ms)$ storage; recall that $m \geq s$ and we require $m \geq 6s$ to use the main theorems. Depending on the relationship between m , s , and $|\mathcal{T}|$, it may be significantly faster to compute the expressions of the main theorems to obtain approximations rather than numerically computing the actual smallest singular value. This gain is particularly noticeable for many signal and image processing problems where $m \gg s$.

Our theorems offer significant computational benefits when there is uncertainty in \mathcal{X} or freedom to choose \mathcal{X} . For instance, consider a situation where each x_k lies in some interval $[x_k^* - \varepsilon_k, x_k^* + \varepsilon_k]$. This may occur if we only have some rough idea of what x_k is or if we are forced to select x_k in this interval due to some constraint. Hence, we are dealing with infinitely many possible sets \mathcal{X} . In principle, we could N many possible \mathcal{X} , and use the SVD to calculate the singular values of $\Phi(m, \mathcal{X})$, which requires $O(Nm^2s)$ operations.

On the other hand, our theorems give interpretable bounds for the smallest singular value in terms of the geometry of \mathcal{X} , so we have some understanding of which sets lead to poor condition numbers. This greatly reduces the number of sets \mathcal{X} that would be considered, from large N to much smaller n . Hence, it suffices to compute n many SVDs, which requires $O(nm^2s)$ operations, or evaluate our main theorems for n many \mathcal{X} and $|\mathcal{T}|$ many τ , which requires $O(ns^3|\mathcal{T}|)$ operations.

Conclusion and future work

This paper presented multiscale estimates for the condition number of Fourier matrices for general \mathcal{X} provided that there is a modicum of redundancy, $m \geq 6s$. The main results are completely new whenever $\Delta(\mathcal{X}) < \frac{1}{m}$ and \mathcal{X} does not consist of separated clumps. Even in the clump framework, the main results significantly reduce sufficient conditions of prior works and achieve similar estimates. The main results also greatly improve upon classical estimates and provide a unified framework for dealing with a disparate collection of sets, which were previously treated on a case-by-case basis.

We state one immediate consequence of the main results. It was shown in [28] that the stability of a foundational algorithm called ESPRIT [37] used for signal processing enjoys (under suitable conditions) the error estimate

$$\text{error} \leq \frac{C_s \cdot \text{noise}}{\sigma_s^2(\Phi(m, \mathcal{X}))}.$$

A significance of this inequality is that it establishes ESPRIT is near min-max optimal. Since this paper greatly enlarges the collection of \mathcal{X} for which we have accurate estimates for $\sigma_s(\Phi(m, \mathcal{X}))$, it yields significant practical implications for ESPRIT and related signal processing algorithms and applications such as [29]. These improvements and their implications will be discussed in a separate article.

Returning to discussion of main results, a natural question is the selection of an optimal scale parameter τ for which to invoke the main inequalities in Theorems 1 and 2. This is not a simple task (for theoretical purposes) and greatly depends on \mathcal{X} . We saw examples where the best effective scale τ is on the order of $\frac{1}{m}$ such as for clumps, whereas $\tau = \frac{1}{2}$ for sparse spike trains. These polarizing examples illustrate that the optimal effective scale does not only depend on m , s , and/or $\Delta(\mathcal{X})$, but on more complicated relationships depending on \mathcal{X} .

Regarding the main theorems' assumptions, they can be weakened to $m \geq 3s$ and $\frac{2\nu(\tau, \mathcal{X})}{\tau} + s \leq m - 1$ without significant modifications to the main proofs. However, doing so would change the numerical constants in a rather undesirable way. For this reason, we decided to state the main results with a stronger than necessary conditions. The techniques introduced in this paper are unable to deal with the extreme case where $m \geq s$ and $\frac{\nu(\tau, \mathcal{X})}{2\tau} \leq m$. This is due to splitting the good and bad sets into separate problems, which comes at a cost of making the interpolants' degrees larger than necessary. To circumvent this, one can handle the good and bad sets in a unified manner and construct interpolants in a completely different way. Our construction of these alternative polynomials have horribly large norms, which in turn, yields a lower bound for $\sigma_s(\Phi)$ that appears to have limited use outside of special contexts.

Many of the techniques and ideas in this paper, including the polynomial method, are flexible. They can be altered to deal with more restricted classes of \mathcal{X} if desired and can be extended to multivariate Fourier matrices. Such a matrix has the form $\Phi = [e^{2\pi i j \cdot x_k}]_{j \in \Omega, x_k \in \mathcal{X}}$ for some $\mathcal{X} = \{x_k\}_{k=1}^s \subseteq \mathbb{T}^d$ and $\Omega \subseteq \mathbb{R}^d$. There are many open questions about the condition number of multivariate Fourier matrices and their behavior greatly depends on the structure of both \mathcal{X} and Ω . From the dual perspective, interpolation by multivariate polynomials is also much more involved. Due to these added technical difficulties and important differences, we deal with multidimensional matrices in a separate article [26].

Acknowledgments

WL is supported by NSF-DMS Award #2309602, a PSC-CUNY grant, and a start-up fund from the Foundation for City College. The author thanks John J. Benedetto, Albert Fannjiang, Wenjing Liao, and Kui Ren for helpful feedback and suggestions. The author also thanks an anonymous reviewer for his/her careful reading and insightful suggestions that led to many improvements of the initial draft.

References

- [1] Céline Aubel and Helmut Bölcskei. Vandermonde matrices with nodes in the unit disk and the large sieve. *Applied and Computational Harmonic Analysis*, 47(1):53–86, 2019.
- [2] Alex H Barnett. How exponentially ill-conditioned are contiguous submatrices of the Fourier matrix? *SIAM Review*, 64(1):105–131, 2022.

- [3] Dmitry Batenkov, Laurent Demanet, Gil Goldman, and Yosef Yomdin. Conditioning of partial nonuniform Fourier matrices with clustered nodes. *SIAM Journal on Matrix Analysis and Applications*, 41(1):199–220, 2020.
- [4] Dmitry Batenkov, Benedikt Diederichs, Gil Goldman, and Yosef Yomdin. The spectral properties of Vandermonde matrices with clustered nodes. *Linear Algebra and its Applications*, 609:37–72, 2021.
- [5] Dmitry Batenkov and Gil Goldman. Single-exponential bounds for the smallest singular value of Vandermonde matrices in the sub-Rayleigh regime. *Applied and Computational Harmonic Analysis*, 55:426–439, 2021.
- [6] Dmitry Batenkov, Gil Goldman, and Yosef Yomdin. Super-resolution of near-colliding point sources. *Information and Inference: A Journal of the IMA*, 10(2):515–572, 2021.
- [7] Fermín SV Bazán. Conditioning of rectangular Vandermonde matrices with nodes in the unit disk. *SIAM Journal on Matrix Analysis and Applications*, 21(2):679–693, 2000.
- [8] John J Benedetto and Weilin Li. Super-resolution by means of Beurling minimal extrapolation. *Applied and Computational Harmonic Analysis*, 48(1):218–241, 2020.
- [9] Lihu Berman and Arie Feuer. On perfect conditioning of Vandermonde matrices on the unit circle. *The Electronic Journal of Linear Algebra*, 16:157–161, 2007.
- [10] Arne Beurling. Balayage of Fourier-Stieltjes transforms. *The Collected Works of Arne Beurling*, 2:341–350, 1989.
- [11] Arne Beurling. Interpolation for an interval in \mathbb{R}^1 . *The Collected Works of Arne Beurling*, 2:351–365, 1989.
- [12] Peter G Casazza and Gitta Kutyniok. *Finite frames: Theory and applications*. Springer Science & Business Media, 2012.
- [13] Shivkumar Chandrasekaran, Karthik R Jayaraman, and Hrushikesh Narhar Mhaskar. Minimum Sobolev norm interpolation with trigonometric polynomials on the torus. *Journal of Computational Physics*, 249:96–112, 2013.
- [14] Charles K Chui. Super-resolution wavelets for recovery of arbitrarily close point-masses with arbitrarily small coefficients. *Applied and Computational Harmonic Analysis*, 61:202–253, 2022.
- [15] Antonio Córdova, Walter Gautschi, and Stephan Ruscheweyh. Vandermonde matrices on the circle: spectral properties and conditioning. *Numerische Mathematik*, 57(1):577–591, 1990.
- [16] Laurent Demanet and Nam Nguyen. The recoverability limit for superresolution via sparsity. *arXiv preprint arXiv:1502.01385*, 2015.
- [17] David L. Donoho. Superresolution via sparsity constraints. *SIAM Journal on Mathematical Analysis*, 23(5):1309–1331, 1992.
- [18] Richard J Duffin and Albert C Schaeffer. A class of nonharmonic Fourier series. *Transactions of the American Mathematical Society*, 72(2):341–366, 1952.

- [19] Alok Dutt and Vladimir Rokhlin. Fast Fourier transforms for nonequispaced data. *SIAM Journal on Scientific computing*, 14(6):1368–1393, 1993.
- [20] Albert C Fannjiang, Thomas Strohmer, and Pengchong Yan. Compressed remote sensing of sparse objects. *SIAM Journal on Imaging Sciences*, 3(3):595–618, 2010.
- [21] Walter Gautschi. On inverses of vandermonde and confluent vandermonde matrices. *Numerische Mathematik*, 5:425–430, 1963.
- [22] Stefan Kunis and Dominik Nagel. On the smallest singular value of multivariate vandermonde matrices with clustered nodes. *Linear Algebra and its Applications*, 604:1–20, 2020.
- [23] Stefan Kunis and Dominik Nagel. On the condition number of vandermonde matrices with pairs of nearly-colliding nodes. *Numerical Algorithms*, 87:473–496, 2021.
- [24] Henry J. Landau. Necessary density conditions for sampling and interpolation of certain entire functions. *Acta Mathematica*, 117:37–52, 1967.
- [25] Weilin Li. Generalization error of minimum weighted norm and kernel interpolation. *SIAM Journal on Mathematics of Data Science*, 3(1):414–438, 2021.
- [26] Weilin Li. Nonharmonic multivariate fourier transforms and matrices: condition numbers and hyperplane geometry. *arXiv preprint arXiv:2407.10313*, 2024.
- [27] Weilin Li and Wenjing Liao. Stable super-resolution limit and smallest singular value of restricted fourier matrices. *Applied and Computational Harmonic Analysis*, 51:118–156, 2021.
- [28] Weilin Li, Wenjing Liao, and Albert Fannjiang. Super-resolution limit of the ESPRIT algorithm. *IEEE Transactions on Information Theory*, 66(7):4593–4608, 2020.
- [29] Weilin Li, Zengying Zhu, Weiguo Gao, and Wenjing Liao. Stability and super-resolution of MUSIC and ESPRIT for multi-snapshot spectral estimation. *IEEE Transactions on Signal Processing*, 70:4555–4570, 2022.
- [30] Wenjing Liao and Albert Fannjiang. Music for single-snapshot spectral estimation: Stability and super-resolution. *Applied and Computational Harmonic Analysis*, 40(1):33–67, 2016.
- [31] Ankur Moitra. Super-resolution, extremal functions and the condition number of Vandermonde matrices. *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, 2015.
- [32] Hugh L Montgomery. The analytic principle of the large sieve. *Bulletin of the American Mathematical Society*, 84(4):547–567, 1978.
- [33] Hugh Lowell Montgomery and Robert Charles Vaughan. The large sieve. *Mathematika*, 20(2):119–134, 1973.
- [34] Francis J. Narcowich and Joseph D. Ward. Scattered-data interpolation on \mathbb{R}^n : error estimates for radial basis and band-limited functions. *SIAM Journal on Mathematical Analysis*, 36(1):284–300, 2004.
- [35] Gerlind Plonka, Daniel Potts, Gabriele Steidl, and Manfred Tasche. *Numerical Fourier Analysis*. Springer, 2018.

- [36] Kui Ren, Yunan Yang, and Björn Engquist. A generalized weighted optimization method for computational learning and inversion. In *International Conference on Learning Representations*, 2022.
- [37] Richard Roy and Thomas Kailath. ESPRIT-estimation of signal parameters via rotational invariance techniques. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(7):984–995, 1989.
- [38] Atle Selberg. *Collected papers*, volume 1. Springer, 1989.
- [39] Jeffrey D. Vaaler. Some extremal functions in Fourier analysis. *Bulletin of the American Mathematical Society*, 12(2):183–216, 1985.
- [40] Yuege Xie, Hung-Hsu Chou, Holger Rauhut, and Rachel Ward. Overparameterization and generalization error: weighted trigonometric interpolation. *SIAM Journal on Mathematics of Data Science*, 4(2):885–908, 2022.
- [41] Robert M. Young. *An introduction to nonharmonic Fourier series*. Academic press, 1981.
- [42] Annan Yu and Alex Townsend. On the stability of unevenly spaced samples for interpolation and quadrature. *BIT Numerical Mathematics*, 63(2):23, 2023.
- [43] Antoni Zygmund. *Trigonometric Series*, volume 1. Cambridge University Press, 1959.