

Sparse Frequency Estimation: Stability and Algorithms

Dissertation with the aim of achieving a doctoral degree
at the Faculty of Mathematics, Informatics and Natural Sciences
Department of Mathematics
of Universität Hamburg

submitted by Benedikt Diederichs

Hamburg, 2018

Als Dissertation angenommen vom Fachbereich Mathematik der Universität Hamburg auf Grund der Gutachten von:

Prof. Dr. Armin Iske
Prof. Dr. Tomas Sauer
Prof. Dr. Daniel Potts

Datum der Disputation: 17.08.2018

Abstract

The thesis at hand is concerned with the problem of sparse frequency estimation. It can be described as follows: Presented with a finite number of samples of an exponential sum, one wishes to calculate its spectrum, which is a discrete set. The focus is on the higher dimensional case, which attracted considerable attention in the last few years.

In the first part of this thesis, we prove that in the one and two dimensional case, the sparse frequency problem is conditionally well-posed. More precisely, we give rather sharp estimates, which guarantee that if two exponential sums have well separated frequencies and their samples are close, so are their frequencies. Further, we give a posteriori error estimates. To prove that, we rely on special band-limited functions, satisfying certain sign patterns. And while such functions are known for quite some time, none of them satisfies an additional property we need. Therefore, we give a construction of such a function and start by reviewing the necessary results from sampling theory.

In the second part, we turn to algorithms to actually solve the estimation problem. We quickly review classical univariate methods and then turn to so-called projection based methods. They cleverly reduce the higher dimensional problem to multiple one dimensional ones, by sampling the exponential sum along several lines. We give recovery guarantees for scattered as well as for parallel lines. For the latter case, we propose a new ESPRIT-like algorithm, combining the estimates along several lines into a single step.

Finally, we turn to other multivariate methods. By explicitly considering the signal space, we can quite naturally deduce higher dimensional analogs of Prony's method, ESPRIT and MUSIC. That allows us to extend Sauer's sampling set, originally proposed for Prony's method, to ESPRIT and MUSIC, which reduces the number of necessary samples as well as the computational complexity significantly.

Acknowledgments

First and foremost I want to thank my advisor, Armin Iske, for his constant support and encouragement. He gently pushed me into the world of mathematical research and without his guidance this thesis would not exist.

Then I want to cordially thank my colleagues, in particular my working group, for all the interesting discussions, and the friendly and supportive atmosphere. A special thanks goes to Matthias and Sara for proofreading parts of my thesis.

I like to thank Tomas Sauer and Daniel Potts for writing reports on my thesis (sacrificing something priceless, their time).

Finally, I want to thank my friends and my family for not only supporting, but also for distracting me from time to time.

Contents

Abstract	i
Acknowledgments	iii
1 Introduction	1
2 Stability and Well-Posedness	5
2.1 Sampling Theorems	9
2.2 Localizing Functions	17
2.3 Conditional Well-Posedness	25
3 Parameter Estimation Techniques	39
3.1 Review of Univariate Methods	40
3.2 Projection-based Methods	50
3.2.1 Projection on Parallel Lines	58
3.3 Other Multivariate Methods	65
3.3.1 Multivariate Prony-type Methods	66
3.3.2 Multivariate MUSIC and ESPRIT-type Algorithms	70
3.3.3 Numerical Examples	73
Bibliography	76
Zusammenfassung	82

Chapter 1

Introduction

Parameter estimation problems are prevalent in all natural sciences. Whenever one has a mathematical model, depending on a set of parameters, and observations, hopefully fitting the model, estimating these parameters is necessary. For this reason, it is not surprising that the parameter estimation of exponential sum has a long and fruitful history, dating back as long as 1795, when Gaspard de Prony published his famous method in [23]. Indeed, exponential sums may be used as a model for the superposition of waves with different frequencies, which is one reason for their common occurrence.

Exponential sums, in their simplest form, are given by

$$f(x) = \sum_{j=1}^M c_j e^{2\pi i y_j x},$$

where $y_j \in [0, 1)$ are the frequencies of f and $c_j \in \mathbb{C}$ is the coefficient to y_j . To obtain a unique representation of f , we always assume that the frequencies are mutually distinct and that the coefficients are non-zero. Often, we denote the frequencies of f by Y^f and the coefficients by $c = c^f \in \mathbb{C}^{Y^f}$, i.e. we use Y^f as an indexing set for the coefficients, which results in

$$f(x) = \sum_{y \in Y^f} c_y e^{2\pi i y x}.$$

The parameter estimation problem for such an exponential sum f is asking to calculate the frequencies and their coefficients from a finite set of samples $f(k)$, $k = 0, \dots, N$, where usually the samples are corrupted by noise.

It is important to keep in mind that we are mostly interested in the frequencies Y^f and not in approximating the given data $f(k)$, $k = 0, \dots, N$ by an exponential sum. To calculate such an approximation is closely related and interesting in its own right, see for example [68, 74].

In the second half of the last century, this problem was studied intensely, as it occurs in direction-of-arrival estimation, where one uses a sensor array to determine the direction of a superposition of wavefronts. Many prominent algorithms, like MUSIC and ESPRIT date back to that time. For an introduction see [61]. In recent years, this estimation problem again attracted a lot of attention. On the one hand, estimating the frequencies of an exponential sum was proposed as a mathematical model for super-resolution [16]. Super-resolution is a somewhat vague term, describing multiple methods, on the technical as well as on the theoretical side, to overcome the natural diffraction limit of imaging technologies. For seminal work in this area, the Nobel Prize in chemistry was awarded to Eric Betzig, Stefan Hell and William Moerner in 2014.

On the other hand, the phenomenon of sparsity has become one of the central aspects in many areas of applied mathematics. Sparsity can be described informally like this: Any element of a vector space is given by a linear combination of a basis. In particular, to describe an element of a d dimensional space, d numbers, called coefficients, are needed. But for some elements, most of these coefficients vanish. If only s coefficients are non-zero, only s numbers are necessary (or one could argue $2s$: s for the coefficients themselves, s to describe which are non-zero). Such elements are called (s -)sparse. It turns out, that quite often, the sparse (or almost sparse) elements are of interest. The challenge then is to exploit the fact, that less information are needed to describe these sparse elements. For example, compression problems can be phrased as finding a basis where all signals we

wish to compress, are approximately sparse. And in sampling, i.e. identifying an element f of a vector space from measurements, having the a-priori information that f is sparse, might enable us to use significantly less measurements. This latter idea is the main theme of a young research area, called compressed sensing (see [32] for an introduction). On a technical level, sparsity turns out to be quite difficult to exploit. One of the main challenge is the missing structure: The sum of two sparse signals is in general less sparse. Hence, the set of sparse signals are not a linear space (or even a convex set).

Now exponential sums can also be seen as sparse signals: They are sparse in the frequency domain. Indeed, they are given as the Fourier transform of a sum of Dirac measures. One reoccurring theme of this thesis is to flesh out the similarities and differences between exponential sums and the standard finite dimensional setting of compressed sensing.

There are a lot possible extensions of the sparse frequency estimation problem. Here, we are mostly interested in the multivariate case, where one wants to calculate the frequency vectors $y_j \in [0, 1)^d$ of a multivariate exponential sum

$$f(x) = \sum_{j=1}^M c_j e^{2\pi i y_j \cdot x} \quad (1.1)$$

from a finite number of samples $f(g)$, $g \in G$, where typically $G \subset \mathbb{Z}^d$. But there are a lot of other variants. We give a few examples.

- Sometimes, instead of coefficients c_j one is confronted with unknown polynomials p_j . In the one dimensional case, this means that f takes the form

$$f(x) = \sum_{j=1}^M p_j(x) e^{2\pi i y_j x}.$$

This problem is sometimes known as the confluent Prony problem, see for example [6].

- Another perspective is the following: f is the Fourier transform of measure of the form

$$\mu = \sum_{j=1}^M c_j \delta_{y_j},$$

where δ_{y_j} are of course the Dirac measure, located at y_j . In the multivariate Prony problem, y_j are on a d dimensional torus. This can be generalized to the case that y_j is on a different manifold, like a sphere or the rotation group $SO(3)$. That is the topic of the PhD thesis of Kristof Schröder [88].

- A more algebraic formulation is possible as well. To this end, let $z_j = e^{2\pi i y_j}$. Then f becomes

$$f(x) = \sum_{j=1}^M c_j z_j^x.$$

Now if we sample only on \mathbb{N}^n , the right-hand side is well defined for c_j, z_j in an arbitrary field. Such extensions are discussed in the PhD thesis of Ulrich von der Ohe [96].

Exponential sums are themselves fascinating objects, which are prevalent in many mathematical fields. They are sometimes called non-harmonic Fourier series, see the book by Young [101]. For certain choices of frequencies y , the span of $e^{2\pi i y \cdot}$ form a frame. Such frames are sometimes called exponential or Fourier frames and indeed the concept of frames was introduced to describe them. Since then, the abstract idea of a frame has become an important generalization of bases, with applications in many fields of mathematics, see the book by Christensen [21] for a thorough introduction.

Another field, where exponential sums are heavily used, is analytic number theory. Efficient estimates of special exponential sums give rise to number theoretical results, like Vinogradov's theorem, to give one of the first important instances. Later, such estimates became important in a subfield called sieve theory. We are particularly interested in sharp estimates for the so-called large sieve, see [64]. One possible proof is due to Atle Selberg, relying on extremal functions in Fourier analysis. We discuss this approach in Chapter 2, though without mentioning the sieve theoretical background.

Contributions

This thesis is divided in two parts. In the first part, covered in Chapter 2, we prove well-posedness of and a posteriori error estimates for the one and two dimensional frequency estimation problem. In Chapter 3 we discuss various existing and new methods to solve it.

The question of well-posedness can be quickly motivated as follows. Assume we are given noisy measurements

$$\tilde{f}(k) = f(k) + \varepsilon_k, \quad k \in G,$$

where f is a multivariate exponential sum as in (1.1), $G \subset \mathbb{Z}^d$ is a finite set and ε_k is some unknown noise. Then we run an algorithm, which uses $\tilde{f}(k)$ to give us an exponential sum g . As we are interested in the frequencies of f , we hope that the frequencies of g and of f are reasonably close. And because we only have the given samples of f , this leads to the question whether

$$\sum_{k \in G} |f(k) - g(k)|^2 \ll 1$$

implies that f and g have similar frequencies. That is the main question, we will answer in Chapter 2.

In general, this does not have to be the case. For example

$$f_1(x) = \delta_1, \quad f_2(x) = 1 - e^{2\pi i \delta_2 x}$$

are both very close to zero, if δ_1 and δ_2 are small and are therefore almost indistinguishable from $g(x) = 0$ on any fixed sampling set G . To circumvent this problem, we assume that the modulus of all coefficients of f and g are larger than c_{\min} and that all frequencies of f and g are well-separated. We measure the distance of two frequency vectors $y, w \in [0, 1)^d$ in the wrap-around distance

$$\|y - w\|_{\mathbb{T}^d} = \min_{n \in \mathbb{Z}^d} \|y - w - n\|_{\infty}.$$

Then we can phrase our first main result as follows: If f and g have well-separated frequencies and

$$\sum_{k=-N}^N |f(k) - g(k)|^2 \leq c_{\min}^2 (N+1),$$

for any frequency y of f there is exactly one frequency $n(y)$ of g . Furthermore, we have that

$$\sum_{y \in Y^f} |y - n(y)|_{\mathbb{T}^d}^2 \lesssim c_{\min}^{-2} N^{-3} \sum_{k=-N}^N |f(k) - g(k)|^2 \lesssim N^{-2}. \quad (1.2)$$

We can interpret this result as a conditional well-posedness property of the frequency estimation problem: If f and g satisfy our model assumptions and their samples are close, so are their frequencies. The actual result, stated in Theorem 2.26, is a little bit stronger. An extension to the two dimensional case is given in Theorem 2.31.

As we only know $\tilde{f}(k)$, we cannot directly use (1.2) to get an a posteriori error estimate. However, it is a routine exercise to estimate $|f(k) - g(k)|$ by $|\tilde{f}(k) - g(k)|$. To this end, we need more assumptions on the noise ε_k . We later state a posteriori error estimates for bounded noise $\|(\varepsilon_k)_k\|_{\leq \eta}$ and for Gaussian noise.

In Chapter 3 we start by reviewing classical univariate methods, namely Prony's method, MUSIC, ESPRIT and the matrix pencil method. Our presentation here differs slightly from the literature. We start by considering the signal space

$$\text{Sig}(f, N) := \text{span} \{(f(k), \dots, f(k+N-1)) : k \in \mathbb{Z}\} \subset \mathbb{C}^N.$$

After proving a few basic facts about $\text{Sig}(f, N)$, we are able to quickly deduce the aforementioned methods.

We then consider methods cleverly using multiple instances of the univariate problem to solve the multivariate problem. The basic idea is that restricting a bivariate exponential sum f to a line

$\ell = \{tv + bv^\perp, t \in \mathbb{R}\}$, where $v, v^\perp \in \mathbb{R}^2$ are orthonormal vectors and $b \in \mathbb{R}$, gives a univariate exponential sum:

$$f(tv + bv^\perp) = f|_\ell(t) = \sum_{y \in Y^f} c_y e^{2\pi i b v^\perp \cdot y} e^{2\pi i v \cdot y t} = \sum_{y^\ell \in v \cdot Y^f} \left(\sum_{\substack{y \in Y^f \\ y \cdot v = y^\ell}} c_y e^{2\pi i b v^\perp \cdot y} \right) e^{2\pi i y^\ell t}.$$

This approach is due to Potts and Tasche [75]. The question is, for which choices of lines ℓ_1, \dots, ℓ_K we can reconstruct f from $f|_{\ell_j}$. In the case of pairwise non-parallel lines, previous work relied on the additional assumption

$$\sum_{\substack{y \in Y^f \\ y \cdot v = y^\ell}} c_y e^{2\pi i b v^\perp \cdot y} \neq 0$$

for all possible choice of y^ℓ and all lines ℓ . We are able to remove this assumption, though at heavy computational costs, a result already published in [25] by Iske and the author.

Next, we turn our attention to the case of multiple parallel lines. We give an algorithm able to recover Y^f from $\mathcal{O}_d(M)$ number of samples (up to logarithmic factors), where $M = \text{ord } f$. This is optimal, but unfortunately the algorithm has exponential complexity. For the more classical case of sampling on $G_N = [-N, N]^d \cap \mathbb{Z}^d$, we give efficient algorithms including a variation of ESPRIT, which combines the estimation along several lines into a single step. Parts of these results are published in [26], again by Iske and the author.

Finally, we turn to multivariate extensions of Prony's method, MUSIC and ESPRIT. Building on our results for the one dimensional case, we again start by defining the multivariate signal space of f . All methods then follow quite naturally. In particular, we rederive Sauer's sampling set (as introduced in [85]), which is of order $\mathcal{O}_d(M^2)$ (up to logarithmic factors), for Prony's method and show how similar sampling sets for MUSIC and ESPRIT can be constructed. With that, we reduce the size of sampling sets for MUSIC and ESPRIT significantly (from $\mathcal{O}_d(M^d)$ to $\mathcal{O}_d(M^2)$, up to logarithmic factors). Furthermore, the computational complexity of ESPRIT reduces from $\mathcal{O}_d(M^{3d})$ to $\mathcal{O}_d(M^3)$ (again, up to logarithmic factors). We conclude with a few numerical examples.

Notation

We use the notations $a \lesssim b$, if $|a| \leq Cb$ for a positive constant C and $a \sim b$ if $a \lesssim b \lesssim a$. If C depends on some parameters of interest, say on d , we write $a \lesssim_d b$. Matrices are denoted by $A = (a_{i,j})_{i,j=1}^{M,N} \in \mathbb{C}^{M \times N}$. A matrix A with columns v_1, \dots, v_N is denoted by $A = [v_1 \dots v_N]$. A^\dagger is the Moore-Penrose pseudo-inverse, A^T the transpose and A^H the conjugate transpose.

The real part of a complex number z is denoted by $\Re z$, the imaginary part by $\Im z$.

All function spaces are complex-valued, e.g. $L^2(\mathbb{R}) = \{f : \mathbb{R} \rightarrow \mathbb{C} : \int |f|^2(x) dx < \infty\}$ (where f is of course measurable).

For a measurable set $A \subset \mathbb{R}^d$, we use χ_A to denote the characteristic function of A .

We use the following normalization of the Fourier transform of a function $f \in L^1(\mathbb{R}^d)$

$$\mathcal{F}[f](w) = \mathcal{F}f(w) = \hat{f}(w) = \int_{\mathbb{R}^d} f(x) e^{-2\pi i x \cdot w} dx.$$

Of course, \mathcal{F} can be extended to $L^1(\mathbb{R}^d) \cup L^2(\mathbb{R}^d)$ by density. In our normalization, we have that \mathcal{F} is an isometry on L^2 , i.e.

$$\|\mathcal{F}f\|_2 = \|f\|_2$$

and that $\mathcal{F}^{-1}[\cdot](w) = \mathcal{F}[\cdot](-w)$. We assume that the reader is familiar with basic results, like Plancherel's theorem, the Fourier inversion formula and Poisson's summation formula. These topics can be found in many textbooks, two examples the author is particularly fond of are [36] and [50].

Chapter 2

Stability and Well-Posedness

In this chapter we discuss stability and (conditional) well-posedness for the frequency estimation problem. To motivate this, we consider frequency estimation as an inverse problem. Let

$$\mathcal{S}^d = \left\{ f(x) = \sum_{y \in Y} c_y e^{2\pi i y \cdot x} : Y \subset [0, 1)^d \text{ finite and } c_y \in \mathbb{C} \setminus \{0\} \right\}$$

be the set of all exponential sums $f : \mathbb{R}^d \rightarrow \mathbb{C}$. We denote the set of frequencies of such an f by Y^f , and the coefficients by c^f . Now we collect samples from an unknown ground truth $f^* \in \mathcal{S}^d$ we wish to identify, by applying the sampling operator

$$\mathcal{P}_G : \mathcal{S}^d \rightarrow \mathbb{C}^G, \quad \mathcal{P}_G f = (f(g))_{g \in G},$$

where $G \subset \mathbb{R}^d$ is the finite sampling set. Typically, $G \subset \mathbb{Z}^d$ and the most popular choice is $G_N = [-N, N]^d \cap \mathbb{Z}^d$. We abbreviate $\mathcal{P}_{G_N} = \mathcal{P}_N$. Now we can try to solve the problem

$$\text{Find } f \in \mathcal{S}^d \text{ with } \mathcal{P}_G f = \mathcal{P}_G f^*,$$

but this turns out to be impossible. As \mathcal{P}_G is a linear map from an infinite dimensional to a finite dimensional space, we are always able to find infinitely many f with $\mathcal{P}_G f = \mathcal{P}_G f^*$. To have any chance to identify f^* , we need additional information. To this end, we exploit the sparsity of f^* in the frequency domain. For an $f \in \mathcal{S}^d$, we call the number of frequencies of f the order of f , i.e., $\text{ord } f = |Y^f|$. Then finding the solution of

$$\min_{f \in \mathcal{S}^d} \text{ord } f \quad \text{subject to } \mathcal{P}_G f = \mathcal{P}_G f^*$$

will recover f^* exactly, at least if G is sufficiently large. Precise statements on the minimal size of G are covered in the next chapter, one simple result is that G_N with $N \geq \text{ord } f^*$ suffices.

The question we are interested in is what happens when we only have noisy data available, i.e., $\mathcal{P}_G f^* + \varepsilon$, where $\varepsilon \in \mathbb{C}^G$ is a vector containing the unknown noise. At this point, we do not want to place any other assumption on the noise, except that $\|\varepsilon\|_2$ is small. To get a feeling what kind of stability results are expected, we quickly recall the finite dimensional case. Everything we discuss here is well-known, see for example [32].

In finite dimensions the sampling operator is a linear map

$$A : \mathbb{C}^N \rightarrow \mathbb{C}^m \quad A(x) = Ax, \quad A \in \mathbb{C}^{m \times N}$$

and we want to find the ground truth $x^* \in \mathbb{C}^N$, given measurements $Ax^* + \varepsilon$, where $\varepsilon \in \mathbb{C}^m$ is a noise vector. We are interested in the case $N > m$ and again, even in the noise-free case

$$\text{Find } x \in \mathbb{C}^N \text{ with } Ax = Ax^*,$$

we have infinitely many solutions. Again, we try to exploit sparsity of x^* and therefore assume that $\|x^*\|_0 = s < N$, where $\|x\|_0$ denotes the number of non-zero entries of a vector $x \in \mathbb{C}^N$. We call

vectors with $\|x\|_0 \leq s$ s -sparse. This leads to the problem

$$\min_{x \in \mathbb{C}^N} \|x\|_0 \quad \text{subject to } Ax = Ax^*. \quad (2.1)$$

But when does this recover x^* exactly? We denote one solution of this problem by \tilde{x} and clearly we have that $\|\tilde{x}\|_0 \leq \|x^*\|_0$ and $A(\tilde{x} - x^*) = 0$. Furthermore, if a vector $x \in \mathbb{C}^N$ with $x \neq x^*$ and $\|x\|_0 \leq \|x^*\|_0$ and $A(x - x^*) = 0$ exists, we cannot recover x^* uniquely. Therefore, x^* is the unique solution to (2.1) if and only if any $2s$ columns of A are linearly independent.

This leads to the restriction $2s \geq m$, which is sharp. Indeed, Vandermonde matrices of the form

$$\begin{pmatrix} 1 & 1 & \dots & 1 \\ z_1 & z_2 & \dots & z_N \\ z_1^2 & z_2^2 & \dots & z_N^2 \\ \vdots & \vdots & \ddots & \vdots \\ z_1^{m-1} & z_2^{m-1} & \dots & z_N^{m-1} \end{pmatrix}$$

where $z_1, \dots, z_N \in \mathbb{C}$ are arbitrary, mutually distinct numbers, satisfy this condition – restricting these matrices to a subspace spanned by m unit vectors gives an injective map.

Next we are interested to recover x^* from noisy measurements. It turns out that we have to have a quantitative property instead of a qualitative one (the injectivity) to deal with the noisy case. To this end, we say that a matrix A satisfies the restricted isometric property (RIP) of order s if there is a $\delta < 1$ such that for all s -sparse x we have that

$$(1 - \delta)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta)\|x\|_2^2. \quad (2.2)$$

Note that if any s columns of A are linearly independent, we can find such a δ (at least if we are only interested in the lower bound), simply because there is only a finite number of possible choices.

The lower bound gives rise to stability. We consider

$$\min_{x \in \mathbb{C}^N} \|x\|_0 \quad \text{subject to } \|A(x^* - x) + \varepsilon\|_2 \leq \eta, \quad (2.3)$$

where $\eta \in \mathbb{R}_{\geq 0}$. Assume that $\|\varepsilon\|_2 \leq \eta$ and that A satisfies the RIP of order $2s$ with constant $\delta < 1$. Further, denote one of the solutions to this problem by \tilde{x} . As x^* is admissible, we have that $\|\tilde{x}\|_0 \leq \|x^*\|_0 = s$. Now, and this is important, we clearly have that $x^* - \tilde{x}$ is $2s$ -sparse. Therefore, we can use the RIP and obtain

$$\|x^* - \tilde{x}\|_2^2 \leq \frac{\|A(x^* - \tilde{x})\|_2^2}{1 - \delta} \leq \frac{(\eta + \|\varepsilon\|_2)^2}{1 - \delta} \leq \frac{4\eta^2}{1 - \delta}. \quad (2.4)$$

This implies that every solution to (2.3) is reasonably close to the ground truth. If additional assumptions on the modulus of the smallest non-zero entry of x^* are available, one can even conclude that the support of x^* is recovered exactly. Indeed, if all non-zero entries of x^* and \tilde{x} have a modulus larger than $2\eta/\sqrt{1 - \delta}$, we obtain $\text{supp } x^* = \text{supp } \tilde{x}$ (where the support of a vector $x \in \mathbb{C}^N$ is the set of indices j with $|x_j| \neq 0$).

We conclude: If we have an a-priori bound s on the sparsity of x^* , we see that every vector $x \in \mathbb{C}^N$, satisfying our model assumption (being s -sparse) with samples close to the measurements, is close to the ground truth. A property, which can be interpreted as a conditional well-posedness of the problem.

However, solving (2.3) is more challenging. In general, such problems are NP-hard. One frequently used approach is to relax $\|x\|_0$ to $\|x\|_1$. Of course this can result in a completely different solution. The result that under suitable conditions on the isometric constant δ the relaxation actually does not change the solution at all, came as a surprise. This was first observed by Candès and Tao in [17], one of the papers which started an active field of research in applied mathematics, called compressed sensing. A thorough introduction in this field is given in [32].

Now we return to the frequency estimation problem. One can ask whether there is a property of \mathcal{P}_G assembling the lower bound in (2.2). It turns out that such an estimate is not possible, at least if one wishes to obtain a bound like

$$(1 - \delta)\|c^f\|_2^2 \leq \|\mathcal{P}_G f\|_2^2 \quad (2.5)$$

uniformly over all f with $\text{ord } f \leq s$. This is actually easy to see, as no such bound can hold for the family of functions

$$f_\lambda(x) = e^{2\pi i x \cdot y} - e^{2\pi i x \cdot (y+\lambda)}$$

for any $y \in [0, 1)^d$ and $\lambda \in \mathbb{R}$ small. We state explicit results in this direction later on.

However, it is possible to prove a bound assembling (2.5) uniformly over all exponential sums having well-separated frequencies. To measure closeness of frequencies, we use the wrap-around distance

$$\|y - w\|_{\mathbb{T}^d} = \min_{n \in \mathbb{Z}^d} \|y - w - n\|_\infty,$$

which accounts for the periodicity of the frequency parameter. Further, we use the notation $|y - w|_{\mathbb{T}} = \|y - w\|_{\mathbb{T}^1}$. Such estimates are actually known for a long time. In the univariate case, we prove a particular strong result in Theorem 2.18 (improving on a result due to Moitra, see [62]), namely

$$\left(2N + 2 - \frac{1}{q}\right) \|c^f\|_2^2 \leq \|\mathcal{P}_N f\|_2^2 \quad (2.6)$$

for all $f \in \mathcal{S}^1$ with $\text{sep } f \geq q$, where

$$\text{sep } f = \max_{y \in Y^f} \min_{\substack{w \in Y^f \\ w \neq y}} \|y - w\|_{\mathbb{T}^d} \quad \text{for } f \in \mathcal{S}^d.$$

In the case of $|Y^f| = 1$, we set $\text{sep } f = 1$. Multivariate extensions are available as well, though less sharp, see Proposition 2.29.

Unfortunately, in contrast to the case of sparse vectors, this does not give any kind of estimate resembling (2.4). The reason is simple. Even if we use well-separated frequency vectors as a prior, we cannot use (2.6) to bound the error, as no information about the separation of the frequencies of $f^* - f$ is available or even possible, as the whole point of our estimate is to have frequencies close to f^* .

In this chapter we overcome this difficulty and prove that if $f, g \in \mathcal{S}^1$ have reasonably well-separated frequencies compared to N , and if

$$\|\mathcal{P}_N(f) - \mathcal{P}_N(g)\|_2^2 \leq c_{\min}^2(N+1),$$

where c_{\min} is a lower bound on the modulus of the coefficients of f and g , the frequencies of f and g have to be very close. Namely, for each frequency y of f there is exactly one close frequency $n(y) \in Y^g$, such that

$$\sum_{y \in Y^f} |y - n(y)|_{\mathbb{T}}^2 \lesssim \frac{\|\mathcal{P}_N(f) - \mathcal{P}_N(g)\|_2^2}{c_{\min}^2 N^3} \lesssim N^{-2}.$$

The precise statement, including all constants, is actually a little bit stronger and given in Theorem 2.26. We show that these estimates are reasonably sharp and in particular that the given orders in N and $|y - n(y)|_{\mathbb{T}}$ are optimal. For $f, g \in \mathcal{S}^2$, a very similar statement is given in Theorem 2.31.

As all constants are reasonably small, these results can be used to obtain a posteriori error estimates. Assume that we have a candidate f , close to the (unknown) ground truth f^* . Further, assume that we are given noisy samples $v = \mathcal{P}_N(f^*) + \varepsilon$. Then we know $\|\mathcal{P}_N(f) - v\|_2^2$, which gives us an estimate for $\|\mathcal{P}_N(f) - \mathcal{P}_N(f^*)\|_2^2$, at least if we have some information about ε . If only $\|\varepsilon\|_2 \leq \eta$ is known, we have to use the triangle inequality. However, if some stochastic information on ε are available, we can do better. We give the details for Gaussian noise as an example in Corollary 2.27 for the univariate and in Corollary 2.32 for the bivariate case.

To prove these results, we rely on specific functions ψ , satisfying

$$\psi(x) \leq \chi_{[-N, N]^d}(x), \quad \text{supp } \hat{\psi} \subset [-1, 1]^d, \quad d = 1, 2.$$

At least in the case $d = 1$, such functions are known for a long time. Their construction is due to Atle Selberg, who built on work of Arne Beurling. Recently, progress for $d > 1$ was made by Carruth, Gonçalves and Kelly [18]. We, however, need an additional property, namely that $\hat{\psi}$ has its global

maximum in zero (and an exact analysis of the growth of $\hat{\psi}(0) - \hat{\psi}(w)$). While in the case $d = 1$ such a function is known (though difficult to find in the literature), we give a new construction for $d = 2$.

This construction relies on a derivative sampling expansion of bandlimited functions (as in [18]). Therefore, we start by giving an overview over the necessary results from sampling theory. One might argue that it would suffice to quote these results. The reason we give them in full detail is threefold. First, the author is very fond of this theory. Secondly, we wish to point out that these expansions were known in the sampling community before they were rediscovered by Vaaler [94] in the univariate case and in [18] in the multivariate case. And finally, reproving and reevaluating known results is an important part of the mathematical science. Often a new approach to an old problem gives some new insight. For example, we give a rather general criterion on when a generalized sampling expansion (as introduced by Papoulis) converges locally uniformly. This result relies mostly on Lemma 2.6, which we prove using frame theoretical results.

Related Literature

One of the first works in a similar direction is due to Donoho [27]. However, his assumptions are different to our setting in multiple ways. First, he assumes that the spectrum of the underlying signal is known on an entire interval, and not only at a finite number of sampling points, i.e., $f|_I$, $I = [-N, N] \subset \mathbb{R}$ instead of $(f(k))_{k \in G}$ is given. Secondly, he makes the a priori assumption that the frequencies are on a grid, i.e., $Y^f \subset \Delta\mathbb{Z}$, $\Delta \ll 1$. That influences his closeness measure. Namely, for two signals f, \tilde{f} the quantity

$$\|c^f - c^{\tilde{f}}\|_2,$$

is bounded from above, where $c^f \in \mathbb{C}^{\Delta\mathbb{Z}}$ is the vector of coefficients of f . In particular, two frequencies are considered close if and only if they are equal. We, however, aim for arbitrary frequencies. Finally, Donoho considers only the univariate case.

Another interesting direction is due to Batenkov [6] (extending previous work by Batenkov and Yomdin [7]). Therein Batenkov considers the so-called confluent setting, which is more general. We give a very short sketch of his results for the special case of exponential sums. The main idea is to consider the mapping

$$\tilde{\mathcal{P}}_N : \mathbb{C}^{2M} \ni (y_1, c_1, y_2, c_2, \dots, y_M, c_M) \mapsto \left(\sum_{j=1}^M c_j e^{2\pi i y_j k} \right)_{k=0, \dots, N-1} \in \mathbb{C}^N.$$

Note that in contrast to our analysis, the space of exponential functions is parametrized explicitly. Denote by $v = \tilde{\mathcal{P}}_N(f^*) + \varepsilon$ the noisy samples of our ground truth $f^* \in \mathbb{C}^{2M}$. Now one could try to recover f^* by solving

$$\text{Find } g \in \mathbb{C}^{2M} \text{ minimizing } \|v - \tilde{\mathcal{P}}_N(g)\|_2.$$

For small $\|\varepsilon\|_2$, a reasonable proxy is the linearized least-squares problem

$$\text{Find } g \in \mathbb{C}^{2M} \text{ minimizing } \|v - \tilde{\mathcal{P}}_N(f^*) - d\tilde{\mathcal{P}}_N(f^*)(g - f^*)\|_2. \quad (2.7)$$

Denote by $g^* \in \mathbb{C}^{2M}$ the solution to this problem. Then one has that $g^* - f^* = d\tilde{\mathcal{P}}_N^\dagger(f^*)\varepsilon$. Therefore, as usual, the conditioning of this problem is described by $\tilde{\mathcal{P}}_N^\dagger(f^*)$. Batenkov then proceeds by estimating

$$\kappa_{k,N}(f^*) = \sum_{j=1}^N \left| \tilde{\mathcal{P}}_N^\dagger(f^*)_{k,j} \right| |\tilde{\mathcal{P}}_N(f^*)_j| \lesssim N^{-1}$$

whenever $N \text{ sep } f^* > K$ for a constant K . This component-wise conditioning describes the local behavior of $\tilde{\mathcal{P}}_N$ with respect to perturbations. In particular, if $|\varepsilon_k|/|\tilde{\mathcal{P}}_N(f^*)_k| \leq \varepsilon$, the estimate

$$|f^* - g^*|_k \leq \kappa_{k,N}\varepsilon$$

holds true. A similar result can be given in case of absolute perturbations. Note that (2.7) cannot be

solved in practice, as f^* is unknown.

This analysis is inherently local, which is an important difference to our results, as we strive for global estimates. Furthermore, it is currently unknown how to lift these explicit estimates to higher dimensions.

In the univariate case, more is known if one assumes a specific noise model or wants to analyze a specific algorithm (or both). For example the Cram er-Rao bound in case of Gaussian noise is derived in [66] (i.e., a lower bound on the variance of any unbiased estimator, satisfying some weak additional assumptions). Asymptotic results for ESPRIT and MUSIC are, again in the case of Gaussian noise, available, see [79, 66] and the references therein. Additional results on MUSIC are given in [55].

Results beyond such a restrictive noise model are known for methods relying on total variation minimization. That is in fact one reason for their increasing popularity over the last few years. In [15], the error of the reconstruction and the ground truth on a larger frequency band than the observed one is estimated. Closer to our analysis is [30], where it is shown that the solution to an infinite dimensional convex optimization problem (which can be solved using semidefinite programming, by a duality argument), recovers the frequencies reasonably well even in the noisy case. A thorough analysis extending this result is carried out in [28]. All results are in one dimension, though (in theory) an extension to the higher dimensional case is possible.

2.1 Sampling Theorems

In this section, we summarize results from sampling theory we rely on in the next section. We start with one-dimensional results and then extend them to higher dimensions. Almost everything in this section is well known and can be found in the literature, though our exposition here may vary slightly, as we develop the results with our applications in mind. An excellent starting point for the univariate theory is [42] or [89].

We are mostly interested in Hermite interpolation, first introduced by Jagerman and Fogel in 1956 [47]. A slightly more general result is then used by Vaaler in [94] to derive extremal Fourier functions, an approach we will follow in the next section. Note that if one wants to construct a function, which interpolates data and satisfies a sign condition, Hermitian interpolation is more useful than normal interpolation (as in the Shannon sampling theorem). Interestingly, this strategy is also used to construct a dual certificate in [16], though they only have to interpolate a finite number of data points.

However, instead of proving the derivative sampling expansion directly, we will establish it as a simple corollary of the generalized sampling theorem by Papoulis [67]. For completeness, we give a proof of Papoulis' result. We generalize it to higher dimensions (a result, which might be interesting in its own right) and again obtain a derivative sampling extension – a theorem first stated in [18].

The natural setting for the theory are the so called Paley-Wiener spaces.

Definition 2.1 (Paley-Wiener Space). *For $\delta > 0$ and $p \in [1, 2]$, the Paley-Wiener space $\text{PW}_\delta^p(\mathbb{R}^d)$ is defined by*

$$\text{PW}_\delta^p(\mathbb{R}^d) = \left\{ f : \mathbb{R}^d \rightarrow \mathbb{C} : f \in L^p(\mathbb{R}^d) \text{ and } \text{supp } \hat{f} \subset [-\delta, \delta]^d \right\}.$$

Often, we let $\delta = 1$ or $p = 2$ and drop them from the notation. Further, we use the abbreviation $\text{PW} = \text{PW}(\mathbb{R}) = \text{PW}_1^2(\mathbb{R})$.

By the following theorem of Paley and Wiener, $\text{PW}_\delta(\mathbb{R})$ can be characterized as a Hilbert space of analytic functions, that satisfy a certain growth condition. In particular, point evaluations and derivatives of $f \in \text{PW}$ are well-defined.

Theorem 2.2 (Paley-Wiener Theorem). *Let $g \in L^2(\mathbb{R})$ with $\text{supp } g \subset [-a, a]$, where $a \in \mathbb{R}_{>0}$. Then $f = \mathcal{F}g$ is an entire analytic function obeying*

$$|f(z)| \lesssim e^{2\pi a |\Im z|}$$

and $f|_{\mathbb{R}} \in L^2(\mathbb{R})$. Conversely, every entire f satisfying these conditions has a Fourier transform with support in $[-a, a]$.

Proof. See [92], Theorem 11.1.2. □

Remarks. 1. No choice of normalization of the Fourier transform is perfect for everything. And while our choice here gives slightly cleaner results later on, it is a little bit unusual in the sampling community. One inconvenience here is that if $\text{supp } \hat{g} \subset [-a, a]$, the function g is not of exponential type a , but of exponential type $2\pi a$.

2. Clearly, changing δ (by dilation) results in isometrically isomorph spaces.

3. By Hausdorff-Young, we know that if $f \in \text{PW}_\delta^p(\mathbb{R}^d)$, the Fourier transform of f lives in $L^{p'}([-\delta, \delta]^d)$, where p' is the Hölder conjugate of p , i.e., $1/p + 1/p' = 1$. Furthermore, for $p_1 < p_2$, we have that $L^{p_1}([-\delta, \delta]^d) \subset L^{p_2}([-\delta, \delta]^d)$, which results in the embedding

$$\text{PW}_\delta^{p_1}(\mathbb{R}^d) \subset \text{PW}_\delta^{p_2}(\mathbb{R}^d), \quad p_1 \leq p_2.$$

We start with a proof of the Shannon sampling theorem. Note that this theorem was actually known before Shannon's groundbreaking work [90], more information on its interesting history can be found in [42].

First, we calculate the Fourier series of $e^{2\pi i x \cdot} \in L^2[-1, 1]$ for an arbitrary, fixed $x \in \mathbb{R}$

$$e^{2\pi i x w} = \sum_{k \in \mathbb{Z}} \frac{\sin(\pi(2x - k))}{\pi(2x - k)} e^{\pi i w k}. \quad (2.8)$$

Now we only have to apply the Fourier inversion formula, which holds pointwise for any $f \in \text{PW}$ (indeed, it holds almost everywhere as $f \in L^2$, which extends to everywhere by continuity of f):

$$f(x) = \int_{-1}^1 \hat{f}(w) e^{2\pi i x w} dw = \sum_{k \in \mathbb{Z}} \frac{\sin(\pi(2x - k))}{\pi(2x - k)} \int_{-1}^1 \hat{f}(w) e^{\pi i w k} dw = \sum_{k \in \mathbb{Z}} \text{sinc}(2x - k) f(k/2),$$

where sinc is defined by

$$\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x} \quad \forall x \in \mathbb{R} \setminus \{0\}$$

and $\text{sinc}(0) = 1$. On the first glance it is not clear, why we are allowed to exchange integral and sum. But actually, this is just a manifestation of Parseval's theorem. If $g, h \in L^2[-1, 1]$, then

$$\frac{1}{2} \int_{-1}^1 g(w) \overline{h(w)} dw = \sum_{k \in \mathbb{Z}} \hat{g}(k) \overline{\hat{h}(k)},$$

where $\hat{g}(k), \hat{h}(k)$ are the Fourier coefficients of g and h . We summarize:

Theorem 2.3 (Shannon Sampling Theorem). *For any $f \in \text{PW}$, we have the pointwise representation*

$$f(x) = \sum_{k \in \mathbb{Z}} \text{sinc}(2x - k) f(k/2).$$

Furthermore, the series on the right-hand side converges absolutely and locally uniformly.

Proof. We already proved that the formula holds pointwise. To prove locally uniform convergence, we use the Weierstrass M-test. First, we note that $(f(k/2))_k \in \ell^2(\mathbb{Z})$, as they are the Fourier coefficients (up to a factor of two) of $\hat{f} \in L^2[-1, 1]$. For any compact set $C \subset \mathbb{R}$, we find a $K \in \mathbb{N}_{>0}$ such that for all $k \in \mathbb{Z}$ with $|k| \geq K$ we have

$$\frac{|\sin(\pi(2x - k))|}{\pi|2x - k|} \lesssim \frac{1}{|k|} \quad \text{uniformly over } x \in C.$$

This immediately proves the claim, as $\left(\frac{f(k/2)}{|k|+1}\right)_k \in \ell^1(\mathbb{Z})$ by Hölder's inequality. \square

Next, we generalize this result beyond point evaluations. A fairly general idea is to consider

samples of the form

$$b[f](t) = \int_{-1}^1 B(w) \hat{f}(w) e^{\pi i w t} dw,$$

where B is a function in $L^\infty[-1, 1]$. Clearly, $b[f](-k)$ is the k -th Fourier coefficient of $2\hat{f}(w)B(w) \in L^2[-1, 1]$. In particular, $(b[f](k))_k \in \ell^2(\mathbb{Z})$.

Example 2.4. *Choosing $B(w) = 1$ results in $b[f](t) = f(t/2)$. A second important example is $B(w) = 2\pi i w$. A short calculation shows that*

$$b[f](t) = \int_{\mathbb{R}} 2\pi i w \hat{f}(w) e^{\pi i w t} dw = f'(t/2). \quad (2.9)$$

The key insight of Papoulis' work [67] is to replace the expansion (2.8) by a more general one. As an example, assume for the moment we wish to reconstruct an $f \in \text{PW}$ from samples $b_1[f](2k)$ and $b_2[f](2k)$, $k \in \mathbb{Z}$, where b_1, b_2 are of the form (2.9). Note that we pick every second sample, as we have actually two types of samples at each point $2k$.

If we had an expansion of the form, where $x \in \mathbb{R}$ and $w \in [-1, 1]$,

$$e^{2\pi i x w} = B_1(w) \sum_{k \in \mathbb{Z}} y_k^{(1)}(x) e^{2\pi i w k} + B_2(w) \sum_{k \in \mathbb{Z}} y_k^{(2)}(x) e^{2\pi i w k}, \quad (2.10)$$

we were done – indeed, using the same idea as for the Shannon sampling theorem above, we obtain

$$\begin{aligned} f(x) &= \int_{-1}^1 \hat{f}(w) e^{2\pi i x w} dw = \sum_{k \in \mathbb{Z}} \left(y_k^{(1)}(x) \int_{-1}^1 B_1(w) \hat{f}(w) e^{2\pi i w k} dw + y_k^{(2)}(x) \int_{-1}^1 B_2(w) \hat{f}(w) e^{2\pi i w k} dw \right) \\ &= \sum_{k \in \mathbb{Z}} y_k^{(1)}(x) b_1[f](2k) + y_k^{(2)}(x) b_2[f](2k). \end{aligned}$$

To understand how to obtain an expansion like (2.10), we let

$$Y_r(w, x) = \sum_{k \in \mathbb{Z}} y_k^{(r)}(x) e^{2\pi i w k}, \quad r = 1, 2.$$

Clearly, Y_r is 1-periodic in w and $Y_r(\cdot, x)$ is given as a Fourier series for every fixed $x \in \mathbb{R}$ with Fourier coefficients $y_k^{(r)}(x)$. Therefore, (2.10) can be written as a system of equations

$$\begin{aligned} e^{2\pi i x w} &= B_1(w) Y_1(w, x) + B_2(w) Y_2(w, x) \\ e^{2\pi i x} e^{2\pi i x w} &= B_1(w+1) Y_1(w, x) + B_2(w+1) Y_2(w, x) \end{aligned} \quad (2.11)$$

for all $x \in \mathbb{R}$ and $w \in [-1, 0]$. This allows an interesting observation. Multiplying the system with $e^{-2\pi i x w}$, we see that the left-hand side is 1-periodic in x . If we assume that the system is regular, i.e., that

$$\begin{pmatrix} B_1(w) & B_2(w) \\ B_1(w+1) & B_2(w+1) \end{pmatrix}$$

is non-singular for all $w \in [-1, 0]$, the solution $Y_r(w, x) e^{-2\pi i x w}$ has to be 1-periodic in x as well. But then

$$y_k^{(r)}(x) = \int_0^1 Y_r(w, x) e^{-2\pi i w k} dw = \int_0^1 Y_r(w, x-k) dw = y_0^{(r)}(x-k) =: y_r(x-k)$$

and we get

$$f(x) = \sum_{k \in \mathbb{Z}} y_1(x-k) b_1[f](2k) + y_2(x-k) b_2[f](2k).$$

The general statement is now quite natural: If we have sampling operators b_1, \dots, b_m and samples $b_j[f](mk)$, $k \in \mathbb{Z}$, the functions Y_r will be $2/m$ -periodic and (2.11) becomes a system of m equations,

which has to be satisfied on $[-1, -1 + 2/m]$.

Theorem 2.5 (Papoulis' Generalized Sampling Theorem). *Given $B_1, \dots, B_m \in L^\infty[-1, 1]$ such that the matrix*

$$T(w) = \begin{pmatrix} B_1(w) & B_2(w) & \dots & B_m(w) \\ B_1(w + 2/m) & B_2(w + 2/m) & \dots & B_m(w + 2/m) \\ \vdots & \vdots & \ddots & \vdots \\ B_1(w + 2(m-1)/m) & B_2(w + 2(m-1)/m) & \dots & B_m(w + 2(m-1)/m) \end{pmatrix}$$

is invertible for all $w \in [-1, -1 + 2/m]$ and furthermore that $(T^{-1}(w))_{j,k} \in L^2[-1, -1 + 2/m]$ for all $j, k = 1, \dots, m$. Then for all $f \in \text{PW}$ we have the expansion

$$f(x) = \sum_{k \in \mathbb{Z}} y_1(x - km/2) b_1[f](km) + \dots + y_m(x - km/2) b_m[f](km), \quad (2.12)$$

which converges absolutely and locally uniformly. Here, $(b_j[f](n))_n \in \ell^2(\mathbb{Z})$ is given by

$$b_j[f](n) = \int_{-1}^1 B_j(w) \hat{f}(w) e^{\pi i w n} dw$$

and $y_j \in \text{PW}$ by

$$y_j(x) = \frac{m}{2} \int_{-1}^{-1+2/m} Y_j(w, x) dw,$$

where $Y_j(w, x) : [-1, -1 + 2/m] \times \mathbb{R} \rightarrow \mathbb{C}$ is the solution of

$$T(w)(Y_j(w, x))_{j=1, \dots, m} = (e^{2\pi i x(w+2(j-1)/m)})_{j=1, \dots, m}.$$

Proof. By assumption, $Y_j(w, x)$ is given by

$$Y_j(w, x) = t_{j,1}(w) e^{2\pi i x w} + \dots + t_{j,m}(w) e^{2\pi i x(w+2(m-1)/m)}, \quad (2.13)$$

where $t_{j,s} \in L^2[-1, -1 + 2/m]$. Therefore, y_j is of the form

$$\begin{aligned} y_j(x) &= \frac{m}{2} \int_{-1}^{-1+2/m} t_{j,1}(w) e^{2\pi i x w} + \dots + t_{j,m}(w) e^{2\pi i x w} e^{2\pi i x(w+2(m-1)/m)} dw \\ &= \frac{m}{2} \left(\int_{-1}^{-1+2/m} t_{j,1}(w) e^{2\pi i x w} dw + \dots + \int_{-1+2/m}^1 t_{j,m}(w - 2(m-1)/m) e^{2\pi i x w} dw \right) \end{aligned}$$

and in particular is an element of the Paley-Wiener space PW.

Next we observe that $Y_j(w, x) e^{-2\pi i x w}$ has period $m/2$ in x (at least for almost all w), due to (2.13). Thus, for all x we obtain (extending $Y_j(w, x)$ periodically in w)

$$\begin{aligned} y_j(x - km/2) &= \frac{m}{2} \int_{-1}^{-1+2/m} Y_j(w, x - km/2) e^{-2\pi i(x - km/2)w} e^{2\pi i(x - km/2)w} dw \\ &= \frac{m}{2} \int_{-1}^{-1+2/m} Y_j(w, x) e^{-\pi i k m w} dw = \frac{1}{2} \int_{-1}^1 Y_j(w, x) e^{-\pi i k m w} dw. \end{aligned}$$

We see that for any fixed $x \in \mathbb{R}$, the Fourier coefficients of $Y(\cdot, x) \in L^2[-1, -1 + 2/m]$ are given by $(y_j(x - km/2))_k \in \ell^2(\mathbb{Z})$. This gives for all fixed $x \in \mathbb{R}$

$$e^{2\pi i x w} = \sum_{j=1}^m B_j(w) Y_j(w, x) = \sum_{j=1}^m B_j(w) \sum_{k \in \mathbb{Z}} y_j(x - km/2) e^{\pi i k m w}$$

in the $L^2[-1, 1]$ -sense as functions in w . This already gives the pointwise equality, just as before by

applying Fourier inversion and Parseval's theorem:

$$\begin{aligned} f(x) &= \int_{-1}^1 \hat{f}(w) e^{2\pi i w x} dw = \sum_{j=1}^m \int_{-1}^1 \hat{f}(w) B_j(w) Y_j(w, x) dw \\ &= \sum_{j=1}^m 2 \sum_{k \in \mathbb{Z}} \left(\frac{1}{2} \int_{-1}^1 \hat{f}(w) B_j(w) e^{-\pi i k w} dw \right) \overline{\left(\frac{1}{2} \int_{-1}^1 Y_j(w, x) e^{-\pi i k w} dw \right)} \\ &= \sum_{j=1}^m \sum_{k \in \mathbb{Z}} y_j(x + km/2) b_j[f](-km). \end{aligned}$$

Here we used that only terms of the form $e^{\pi i k m w}$ do appear in the Fourier expansion of $Y_j(w, x)$.

Now we prove the locally uniform convergence. Again, relying on the Weierstrass M-test, we need to show that for any $x \in \mathbb{R}$ we find $\varepsilon > 0$ with

$$\max_{z \in [x-\varepsilon, x+\varepsilon]} |y_j(z - km/2)| \leq M_k,$$

where $(M_k) \in \ell^2(\mathbb{Z})$ may depend on y_j and x . Then we obtain that on $[x - \varepsilon, x + \varepsilon]$, we have $|y_j(z - km/2) b[f](km)| \leq M_k |b[f](km)|$ and $(M_k |b[f](km)|)_k \in \ell^1(\mathbb{Z})$. Then, Weierstrass gives us that

$$\sum_{k \in \mathbb{Z}} y_j(x - km/2) b_1[f](km)$$

converges absolutely and locally uniformly. We prove the existence of such $(M_k)_k$ for general functions in PW in the following lemma. \square

Lemma 2.6. *Let $g \in \text{PW}$ be given. Then for any closed interval $I \subset \mathbb{R}$ we have that*

$$\sum_{n \in \mathbb{Z}} \max_{x \in I} |g(x + n/2)|^2 \lesssim_{|I|} \int_{\mathbb{R}} |g(y)|^2 dy.$$

Proof. First, it is clear that we may assume that $I = [-\frac{1}{9}, \frac{1}{9}]$. Indeed, any longer interval can be split into several subintervals I_1, \dots, I_r covering I . Then we clearly have that

$$\sum_{n \in \mathbb{Z}} \max_{x \in I} |g(x + n/2)|^2 \leq \sum_{j=1}^r \sum_{n \in \mathbb{Z}} \max_{x \in I_j} |g(x + n/2)|^2.$$

Furthermore, it is obvious that if the bound holds true for one interval, it holds true for all of its translates (which is equivalent to translating g).

For a fixed g and $I = [-\frac{1}{9}, \frac{1}{9}]$, let $x_n/2 \in n/2 + I$ be chosen such that

$$|g(x_n/2)| = \max_{x \in I} |g(x + n/2)|.$$

We obtain

$$\sum_{n \in \mathbb{Z}} \max_{x \in I} |g(x + n/2)|^2 = \sum_{n \in \mathbb{Z}} |g(x_n/2)|^2 = \sum_{n \in \mathbb{Z}} \left| \int_{-1}^1 \hat{g}(w) e^{\pi i w x_n} dw \right|^2 = 4 \sum_{n \in \mathbb{Z}} |\langle \hat{g}, e^{-\pi i x_n \cdot} \rangle_{L^2[-1,1]}|^2.$$

If we knew that $(e^{-\pi i x_n \cdot})_n$ were a frame in $L^2[-1, 1]$ with an upper frame bound $B \in \mathbb{R}$ independent of $(x_n)_n$, we were done:

$$\sum_{n \in \mathbb{Z}} |\langle \hat{g}, e^{-2\pi i x_k \cdot} \rangle_{L^2[-1,1]}|^2 \leq B \|\hat{g}\|_{L^2[-1,1]}^2 = B \int_{\mathbb{R}} |g(y)|^2 dy.$$

As $\sup_n |x_n - n| \leq \frac{2}{9} < \frac{1}{4}$, this is indeed the case, a result of the extension of Kadec's $\frac{1}{4}$ -Theorem to frames, which is proved in [5] and [20], see also [21] Theorem 9.8.6. \square

Remarks. 1. Our proof of Theorem 2.5 follows the original proof by Papoulis, see [67]. Papoulis'

generalized sampling theorem is well-known and found its way into textbooks, see for example [29].

2. In [11], stability of the sampling expansion is discussed and $y_j \in \text{PW}$ proved. The author did not find a source, which discusses the locally uniform convergence.
3. To check whether $(T^{-1}(w))_{j,k} \in L^2[-1, -1 + 2/m]$, a convenient method is to check if one can find a $c > 0$ such that

$$|\det T(w)| \geq c > 0 \quad \text{for almost all } w \in [-1, -1 + 2/m],$$

which even gives $(T^{-1}(w))_{j,k} \in L^\infty([-1, -1 + 2/m])$ (and in particular $Y_j(w, x) \in L^\infty([-1, -1 + 2/m] \times \mathbb{R})$), by Cramer's rule. This is discussed in [11].

4. Later, we need a result like Lemma 2.6, but for the L^1 -norm. See Lemma 2.10, where we also list some related literature.
5. Note that

$$\int_{\mathbb{R}} |g(y)|^2 dy \lesssim \sum_{n \in \mathbb{Z}} \max_{x \in I} |g(x + n/2)|^2$$

holds trivially, as shrinking I decreases the right-hand side and the inequality holds with equality (up to a factor of two) for $I = \{0\}$. This is due to the fact that $g(n/2)$ is the n -th Fourier coefficient of \hat{g} (again up to a factor of two), a result we already used. However, by using the lower frame bound it is possible to prove a much stronger result, namely

$$\int_{\mathbb{R}} |g(y)|^2 dy \lesssim_I \sum_{n \in \mathbb{Z}} \min_{x \in I} |g(x + n/2)|^2,$$

at least if $I \subset (-\frac{1}{8}, \frac{1}{8})$ is satisfied.

Example 2.7. We now give an example, which will be of great importance later on. We pick

$$B_1(w) = 1 \quad \text{and} \quad B_2(w) = 2\pi iw.$$

As already discussed in Example 2.4, this results in $b_1[f](t) = f(t/2)$ and $b_2[f](t) = f'(t/2)$. To determine y_j , we solve (2.11), which results in

$$Y_2(w, x) = \frac{1}{2\pi i} e^{2\pi i w x} (e^{2\pi i x} - 1) \quad \text{and} \quad Y_1(w, x) = e^{2\pi i w x} (1 + w - w e^{2\pi i x}).$$

Calculating y_1 and y_2 is straightforward:

$$y_2(x) = \frac{1}{2\pi i} \int_{-1}^0 e^{2\pi i x w} (e^{2\pi i x} - 1) dw = \frac{1}{4\pi^2 i^2 x} (e^{2\pi i x} - 1) (1 - e^{-2\pi i x}) = \frac{\sin^2(\pi x)}{\pi^2 x}$$

$$y_1(x) = \int_{-1}^0 e^{2\pi i x w} (1 + w - w e^{2\pi i x}) dw = \frac{\sin^2(\pi x)}{\pi^2 x^2}.$$

Thus, we proved the expansion

$$f(x) = \sum_{k \in \mathbb{Z}} \frac{\sin^2(\pi(x-k))}{\pi^2(x-k)^2} f(k) + \frac{\sin^2(\pi(x-k))}{\pi^2(x-k)} f'(k) = \frac{\sin^2(\pi x)}{\pi^2} \sum_{k \in \mathbb{Z}} \frac{f(k)}{(x-k)^2} + \frac{f'(k)}{x-k},$$

which is (up to a different scaling) exactly the formula introduced by Jagerman and Fogel in [47]. The series converges absolutely and locally uniformly.

Remarks. 1. There are more examples of sampling expansions, which can be derived using Papoulis' generalized sampling theorem. In particular, instead of using only the first derivative, it is straightforward (though not easy to compute) to obtain expansions including higher derivatives. These expansions predate Papoulis' theorem and were first introduced by Linden and Abramson in [56]. Another example is bunched sampling, i.e., using samples of the form $f(km/2 + \alpha_j)$,

where α_j , $j = 1, \dots, m$ are pairwise distinct numbers in $] - 1, 1[$. A sampling theorem covering this case was first proved by Yen in [100] and again can be recovered using Papoulis' result. These examples are already given in Papoulis' work [67].

2. One might wish to obtain similar expansions for functions, which are entire but not in $L^2(\mathbb{R})$. At least for bounded, entire function of exponential type 2π , one can easily use the presented results, by noting that if $f : \mathbb{R} \rightarrow \mathbb{C}$ is bounded and of exponential type π , we may consider

$$g(z) = \begin{cases} f'(0) & \text{if } z = 0 \\ \frac{f(z) - f(0)}{z} & \text{otherwise.} \end{cases}$$

Then $g \in \text{PW}$, by the well-known fact that an entire function f of exponential type τ , which is bounded on the real line, satisfies $|f(z)| \lesssim e^{\tau|\Im z|}$ – see [10], Theorem 6.2.4, and the Paley-Wiener Theorem 2.2. Hence, we can apply the aforementioned results. Just to give an example, using Shannon's sampling theorem results in

$$f(z) = \frac{\sin(2\pi z)}{\pi} \left(\frac{f(0)}{2z} + f'(0) + \sum_{\substack{k \in \mathbb{Z} \\ k \neq 0}} (-1)^k f(k/2) \left(\frac{1}{k} + \frac{1}{2z - k} \right) \right).$$

This formula is well-known, see for example [102], Theorem 7.19. Using g in the derivative expansion from Example 2.7 results in another useful representation, see [94], Theorem 10.

Now we give a multivariate extension of Papoulis' sampling theorem. We only need an expansion using tensor-type sampling operators of the form

$$b[f](k) = \int_{[-1,1]^d} B^{(1)}(w_1) \cdots B^{(d)}(w_d) \hat{f}(w_1, \dots, w_d) e^{2\pi i(w_1, \dots, w_d) \cdot k} \mathfrak{d}(w_1, \dots, w_d)$$

for $k \in \mathbb{Z}^d$ and $f \in \text{PW}(\mathbb{R}^d)$. Therefore, the multivariate theorem is an easy corollary of the univariate result.

Theorem 2.8 (Papoulis' Sampling Theorem - Multivariate Case). *Given $B_j^{(k)} \in L^\infty[-1, 1]$, $j = 1, \dots, m$ and $n = 1, \dots, d$ such that for all k the functions $B_j^{(k)}$, $j = 1, \dots, m$ satisfy the conditions given in Theorem 2.5. Then for all $f \in \text{PW}(\mathbb{R}^d)$ we obtain*

$$f(x) = \sum_{k \in \mathbb{Z}^d} \sum_{j_1, \dots, j_d=1}^m y_{j_1}^{(1)}(x_1 - mk_1/2) \cdots y_{j_d}^{(d)}(x_d - mk_d/2) b_{j_1, \dots, j_d}[f](km), \quad (2.14)$$

which converges absolutely and locally uniformly.

Here, for each $n = 1, \dots, d$ the functions $y_j^{(n)} \in \text{PW}$ are related to $B_j^{(n)}$ as in the univariate case and $b_{j_1, \dots, j_d}[f]$ is given by

$$b_{j_1, \dots, j_d}[f](k) = \int_{[-1,1]^d} B_{j_1}^{(1)}(w_1) \cdots B_{j_d}^{(d)}(w_d) \hat{f}(w_1, \dots, w_d) e^{2\pi i(w_1, \dots, w_d) \cdot k} \mathfrak{d}(w_1, \dots, w_d).$$

Proof. First, note that $b_{j_1, \dots, j_d}[f](k/2)$ is the k -th Fourier coefficient of $B_{j_1}^{(1)}(w_1) \cdots B_{j_d}^{(d)}(w_d) \hat{f}(w) \in L^2[-1, 1]^d$ and therefore in $\ell^2(\mathbb{Z}^d)$. Further, we already know that $(y_j^{(n)}(x_n - k/2))_k \in \ell^2(\mathbb{Z})$. Thus, the right-hand side of (2.14) converges absolutely in each point $x \in \mathbb{R}^d$. Now we have the decomposition

$$\begin{aligned} e^{2\pi i x \cdot w} &= e^{2\pi i x_1 w_1} \cdots e^{2\pi i x_d w_d} = \prod_{n=1}^d \left(\sum_{j=1}^m B_j^{(n)}(w_n) Y_j^{(n)}(w_n, x_n) \right) \\ &= \sum_{j_1, \dots, j_d=1}^m B_{j_1}^{(1)}(w_1) Y_{j_1}^{(1)}(w_1, x_1) \cdots B_{j_d}^{(d)}(w_d) Y_{j_d}^{(d)}(w_d, x_d). \end{aligned}$$

Using Fourier inversion and Parseval's theorem gives

$$\begin{aligned}
f(x) &= \int_{[-1,1]^d} \hat{f}(w) e^{2\pi i x \cdot w} \, dw = \sum_{j_1, \dots, j_d=1}^m \int_{[-1,1]^d} \hat{f}(w) B_{j_1}^{(1)}(w_1) Y_{j_1}^{(1)}(w_1, x_1) \cdots B_{j_d}^{(d)}(w_d) Y_{j_d}^{(d)}(w_d, x_d) \, dw \\
&= \sum_{j_1, \dots, j_d=1}^m \sum_{k \in \mathbb{Z}^d} b_{j_1, \dots, j_d}[f](-km) \left(\frac{1}{2^d} \prod_{n=1}^d \int_{-1}^1 Y_{j_n}^{(n)}(w_n, x_n) e^{\pi i w_n k_n} \, dw_n \right) \\
&= \sum_{j_1, \dots, j_d=1}^m \sum_{k \in \mathbb{Z}^d} b_{j_1, \dots, j_d}[f](-km) \prod_{n=1}^d y_{j_n}^{(n)}(x_n + k_n m/2).
\end{aligned}$$

The proof of the locally uniform convergence of (2.14) works exactly as in the univariate case, by relying on the Weierstrass M-test. However, no multivariate version of Lemma 2.6 is necessary, as for $Q = I_1 \times \cdots \times I_d$

$$\left(\max_{(x_1, \dots, x_d) \in Q} |y_{j_1}^{(1)}(x_1 - k_1/2) \cdots y_{j_d}^{(d)}(x_d - k_d/2)| \right)_{k_1, \dots, k_d} \in \ell^2(\mathbb{Z}^d)$$

follows directly from the univariate result. \square

Remarks. 1. It is easy (though a notational burden) to use different numbers m_1, \dots, m_d . Indeed, we used the expansions in each dimension separately. It is also possible to cover the case that $\text{supp } \hat{f} \subset [-a_1, a_1] \times \cdots \times [-a_d, a_d]$, where $a_j \in \mathbb{R}_{>0}$ for $j = 1, \dots, d$.

2. Sometimes it is of interest to obtain sampling expansions for f with $\text{supp } \hat{f} \subset K$, where K is some set in \mathbb{R}^d . Note that Theorem 2.8 can be used to cover the case of K being a parallelepiped by applying it to $f(Q \cdot)$ for a suitable matrix Q . A more detailed discussion can be found in [19].
3. This result is closely related to a more general result of Brown and Sa-Ngsari in [12]. They state an extension of the linear system (2.11) to the higher dimensional case. However, they do not give any example and do not mention the tensor-product case. Our result on the mode of convergence seems to be new.

Examples 2.9. We start by stating the multivariate analog of Shannon's sampling expansion. This is known since at least 1962, see [69]. Every $f \in \text{PW}(\mathbb{R}^d)$ can be written as

$$f(x) = \sum_{k \in \mathbb{Z}^d} \text{sinc}(2x_1 - k_1) \cdots \text{sinc}(2x_d - k_d) f(k/2),$$

where the right-hand side converges absolutely and locally uniformly.

Next, we state the derivative sampling expansion. We let $B_1^{(k)}(w_k) = 1$ and $B_2^{(k)}(w_k) = 2\pi i w_k$ for all $k = 1, \dots, d$. We use the notation $\mathbf{1} = (1, \dots, 1) \in \mathbb{Z}^d$ and $\mathbf{2} = 2\mathbf{1}$. This results in

$$b_j[f](k) = \int_{[-1,1]^d} (2\pi i)^{|j-1|} w^{j-1} \hat{f}(w) e^{2\pi i w \cdot k} \, dw = \partial_{j-1} f(k), \quad j \in \{1, 2\}^d.$$

Here, we use the usual notation: For $j \in \mathbb{N}_0^d$ and $w \in \mathbb{R}^d$, $w^j = w_1^{j_1} \cdots w_d^{j_d}$, $|j| = j_1 + \cdots + j_d$ and $\partial_j = \partial_1^{j_1} \cdots \partial_d^{j_d}$. We obtain for all $f \in \text{PW}(\mathbb{R}^d)$

$$f(x) = \frac{\sin^2(\pi x_1) \cdots \sin^2(\pi x_d)}{\pi^{2d}} \sum_{k \in \mathbb{Z}^d} \sum_{j \in \{0,1\}^d} \frac{\partial_j f(k)}{(x-k)^{2^{-j}}}, \quad (2.15)$$

where the right-hand side converges absolutely and locally uniformly. This expansion was first proved in [19] (explicitly stated only for $d = 2$ and only pointwise) and reproved in [18] (Theorem 11), where the mode of convergence was explicitly given.

2.2 Localizing Functions

We now turn to localizing functions. To state what we understand under a localizing function, let

$$\square_N(x) = \begin{cases} 1 & x \in (-N, N) \\ 0 & x \notin [-N, N] \\ \frac{1}{2} & x = \pm N \end{cases},$$

where $N \in \mathbb{N}_{>0}$. Any function $f \in \text{PW}$ with $f(x) \leq \square_N(x)$ for all $x \in \mathbb{R}$ is called localizing function (for \square_N). Mainly, we are interested in finding functions $f \in \text{PW}$ minimizing

$$\int_{\mathbb{R}} \square_N(x) - f(x) dx. \quad (2.16)$$

These functions are sometimes called extremal functions. Of course, this question is interesting not only for \square_N . In fact, the first results in this direction are given by Arne Beurling in [9]. He gave an entire function of exponential type 2π with $B(x) \geq \text{sgn}(x)$, minimizing $\int_{\mathbb{R}} B(x) - \text{sgn}(x) dx$. Clearly, $B \notin \text{PW}^p$, as $\text{sgn} \notin L^p(\mathbb{R}^d)$ for all $p \in [1, 2]$.

In the 1970s Atle Selberg noted that Beurling's construction can be used to obtain extremal functions for \square_N . Indeed,

$$c_N(x) = -\frac{1}{2}(B(x-N) + B(-N-x)) \leq \square_N(x)$$

gives

$$\int_{\mathbb{R}} \square_N(x) - c_N(x) dx = 1,$$

which turns out to be optimal. The case where one wishes to find extremal functions in PW smaller (or greater) than the characteristic function of an arbitrary interval is more tricky and was first announced by Logan [59] in 1977. Much later, in 2013, a complete proof was given by Littmann [58], though we do not need his result here.

We prove which $f \in \text{PW}$ are extremal for \square_N without taking the detour of finding extremal functions to sgn , which has the technical advantage that we can work completely in PW. The extremal function is not unique, but it will turn out that only one of them has an additional property we need later on - namely, a global maximum of \hat{f} in zero. Beurling's function can then be recovered by a limiting process.

In the higher dimensional case extremal functions are not known. However, we construct localizing functions, which suffice for our purpose. This builds on work of Carruth, Gonçalves and Kelly [18]. We give some extensions of their work (which is limited to $N = 1$) and in particular construct a localizing function F such that \hat{F} has its global maximum in zero. The function enables us to prove the main results of this chapter in the next section.

But what makes localizing functions so interesting? Their usefulness comes from the fact that they allow to overcome a basic uncertainty principle: No function can be too localized in the spatial and in the frequency domain simultaneously. There are many results making this abstract principle precise like Heisenberg's uncertainty principle or Hardy's uncertainty principle. The Paley-Wiener theorems are also connected to this observation. The Fourier transform of a perfectly localized (i.e., compactly supported) function is an entire function (and therefore in particular not compactly supported).

Now \square_N is perfectly localized in the spatial domain and therefore not localized in the frequency domain at all ($\mathcal{F}\square_N(w)$ decays only like w^{-1}). However, by estimating $\square_N \geq f$, where f is any localizing function, we replaced \square_N by a function perfectly localized in the frequency domain.

As a preparation, we give an instance of Poisson's summation formula, slightly different from most results in literature. To this end, we need the following lemma.

Lemma 2.10. *Let $f \in \text{PW}_{\delta}^1(\mathbb{R}^d)$, then*

$$\sum_{k \in \mathbb{Z}^d} |f(k)| \lesssim_{\delta, d} \int_{\mathbb{R}^d} |f(x)| dx.$$

Proof. We extend a technique used by Wiener in [98], Lemma 67. Pick a smooth function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ with $\text{supp } \phi \subset [-2\delta, 2\delta]^d$ and $\phi(x) = 1$ for all $x \in [-\delta, \delta]$. Then

$$\begin{aligned} f(x) &= \int_{\mathbb{R}^d} \hat{f}(w) e^{2\pi i x \cdot w} dw = \int_{\mathbb{R}^d} \hat{f}(w) \phi(w) e^{2\pi i x \cdot w} dw \\ &= \int_{\mathbb{R}^d} \hat{f}(w) \int_{\mathbb{R}^d} \hat{\phi}(u - x) e^{2\pi i u w} du dw \\ &= \int_{\mathbb{R}^d} \hat{\phi}(u - x) f(u) du. \end{aligned}$$

Now, due to the smoothness of ϕ , we know that $\sum_{k \in \mathbb{Z}^d} |\hat{\phi}(u - k)| \in L^\infty(\mathbb{R}^d)$ and we obtain

$$\sum_{k \in \mathbb{Z}^d} |f(k)| \leq \int_{\mathbb{R}^d} |f(u)| \sum_{k \in \mathbb{Z}^d} |\hat{\phi}(u - k)| du \lesssim_{\delta, d} \int_{\mathbb{R}^d} |f(u)| du.$$

□

Remarks. 1. Note that if we additionally assume that the ϕ in the proof is even, we have that

$$f(x) = f * \hat{\phi}(x).$$

This is not surprising, as ϕ does not filter the frequencies in $[-\delta, \delta]^d$.

2. It is easy to see that one can use the same idea to prove a stronger result. Namely, let $D \subset \mathbb{R}^d$ be compact, $f \in \text{PW}_\delta^1(\mathbb{R}^d)$. Then

$$\sum_{k \in \mathbb{Z}^d} \max_{x \in D} |f(x + k)| \lesssim_{D, \delta, d} \int_{\mathbb{R}^d} |f(x)| dx.$$

In the case $d = 1$, this is exactly the result by Wiener in [98], Lemma 67. Note the relation to Lemma 2.6, where a similar result for $f \in \text{PW}_\delta^2(\mathbb{R})$ is given.

3. There is a result covering arbitrary, positive $p > 0$, by Plancherel and Polya [71] (Theorem 3, formula 127). Namely, they prove that for any F of exponential type and any positive p ,

$$\sum_{k \in \mathbb{Z}^d} |F(k)|^p \lesssim \int_{\mathbb{R}^d} |F(x)|^p dx.$$

However, the author feels that using the elementary argument given above is interesting in its own right and makes this work more self contained.

Theorem 2.11 (Poisson's Summation Formula on PW_δ^1). *Let $f \in \text{PW}_\delta^1(\mathbb{R}^d)$ be arbitrary. Then*

$$\sum_{k \in \mathbb{Z}^d} f(k) e^{-2\pi i k \cdot x} = \sum_{k \in \mathbb{Z}^d} \hat{f}(x - k)$$

holds for all $x \in \mathbb{R}^d$, where both sums converge uniformly on compact sets.

Proof. By assumption, \hat{f} is continuous with support in $[-\delta, \delta]^d$. We denote the periodification by

$$G(x) = \sum_{k \in \mathbb{Z}^d} \hat{f}(x - k). \tag{2.17}$$

The Fourier coefficients of G are given by

$$\begin{aligned} c_n &= \int_{[0, 1]^d} G(x) e^{-2\pi i x \cdot n} dx = \int_{[0, 1]^d} \sum_{k \in \mathbb{Z}^d} \hat{f}(x - k) e^{-2\pi i (x - k) \cdot n} dx \\ &= \int_{\mathbb{R}^d} \hat{f}(x) e^{-2\pi i x \cdot n} dx = f(-n). \end{aligned}$$

By Lemma 2.10, the Fourier coefficients of G are summable. Hence, its Fourier series is absolutely convergent and a continuous functions, almost everywhere equal to G . But due to the bounded support of \hat{f} , the sum in (2.17) is actually finite and G is continuous as well. Therefore,

$$G(x) = \sum_{k \in \mathbb{Z}^d} \hat{f}(x - k) = \sum_{k \in \mathbb{Z}^d} f(k) e^{-2\pi i k \cdot x}.$$

□

Remarks. 1. The special case $x = 0$ and $\delta = 1$ is stated in [18].

2. Usually, Poisson's summation formula (for functions in $L^1(\mathbb{R}^d)$) relies on additional decay properties of f and \hat{f} , like $|f|(x), |\hat{f}|(x) \lesssim (1+|x|)^{-d-\varepsilon}$, where $\varepsilon > 0$ (see for example [36], Proposition 1.4.2). This is necessary, as there are functions $f \in L^1(\mathbb{R}) \cap C(\mathbb{R})$ with $\hat{f} \in L^1(\mathbb{R}) \cap C(\mathbb{R})$ such that both series converge absolutely and the formula still fails to hold, see [50], p. 130.

We are now ready to prove lower bounds for (2.16).

Proposition 2.12. *For any $f \in \text{PW}$ and any $N \in \mathbb{N}_{>0}$ with $f(x) \leq \square_N(x)$ (making the tacit assumption that f is real valued), we have that*

$$\int_{\mathbb{R}} \square_N(x) - f(x) dx \geq 1.$$

Equality holds if and only if f satisfies

$$f(k) = \begin{cases} 1 & k \in \mathbb{Z}, |k| < N \\ 0 & k \in \mathbb{Z}, |k| \geq N \end{cases} \quad \text{and} \quad f'(k) = 0 \quad \text{for all } k \in \mathbb{Z} \setminus \{\pm N\}.$$

Proof. First, we notice that

$$\int_{\mathbb{R}} \square_N(x) - f(x) dx = 2N - \int_{\mathbb{R}} f(x) dx = 2N - \hat{f}(0).$$

Now, if $f \notin L^1(\mathbb{R})$, the inequality holds trivially. We therefore assume that $f \in \text{PW}^1$, which implies that \hat{f} is continuous. Thus, $\hat{f}(k) = 0$, $k \in \mathbb{Z} \setminus \{0\}$. Poisson's summation formula yields

$$\int_{\mathbb{R}} \square_N(x) - f(x) dx = 2N - \sum_{k \in \mathbb{Z}} \hat{f}(k) = 2N - \sum_{k \in \mathbb{Z}} f(k) \geq 1,$$

where the last inequality follows from $f(k) \leq 1$ if $|k| < N$ and $f(k) \leq 0$ if $|k| \geq N$.

This already explains the first set of conditions. The second set follows easily from the fact that f has to have local maxima at $k \in \mathbb{Z} \setminus \{\pm N\}$, again due to $f(x) \leq \square_N(x)$. □

Now we have to construct such an $f \in \text{PW}$. To this end, we invoke the expansion of Jagerman and Fogel, as given in Example 2.7. Plugging in the necessary conditions we just proved, we obtain

$$f(x) = \frac{\sin^2(\pi x)}{\pi^2} \left(\frac{f'(-N)}{x+N} + \frac{f'(N)}{x-N} + \sum_{|k| < N} \frac{1}{(x-k)^2} \right).$$

As we wish to have $f \in L^1(\mathbb{R})$, we have to have that $-f'(N) = f'(-N)$. Therefore, we have to find for which $\alpha \in \mathbb{R}$ the function

$$\frac{\sin^2(\pi x)}{\pi^2} \left(\frac{\alpha}{x+N} - \frac{\alpha}{x-N} + \sum_{|k| < N} \frac{1}{(x-k)^2} \right) = \frac{\sin^2(\pi x)}{\pi^2} \left(\frac{2N\alpha}{N^2 - x^2} + \sum_{|k| < N} \frac{1}{(x-k)^2} \right)$$

is pointwise smaller than \square_N . This is a little bit tricky and requires some non-trivial inequalities, as we will now see.

Proposition 2.13. For $\alpha \in [\frac{2N-1}{2N}, \frac{2N+1}{2N}]$, the function

$$F_{\alpha,N}(x) = \frac{\sin^2(\pi x)}{\pi^2} \left(\frac{2N\alpha}{N^2 - x^2} + \sum_{|k| < N} \frac{1}{(x-k)^2} \right)$$

satisfies $F_{\alpha,N}(x) \leq \square_N(x)$ and minimizes (2.16) among all $f \in \text{PW}$ minorizing \square_N . If $\alpha < \frac{2N-1}{2N}$, this is no longer true.

Proof. We have to prove that $F_{\alpha,N}(x) \leq \square_N(x)$ for $\alpha \in [\frac{2N-1}{2N}, \frac{2N+1}{2N}]$. We start by proving $F_{\alpha,N}(x) \leq 0$ if $|x| \geq N$. As $F_{\alpha,N}$ is even, we can assume that $x \geq N$. Further, for $x = N$ this is obvious and for $x > N$ the claim is equivalent to

$$\sum_{|k| < N} \frac{1}{(x-k)^2} \leq \frac{2N\alpha}{x^2 - N^2}. \quad (2.18)$$

Multiplying with $x^2 - N^2$ and letting $x \rightarrow \infty$, we see that $\alpha \geq \frac{2N-1}{2N}$ is necessary. We show that it is sufficient as well. To this end, we can restrict ourselves to the case $\alpha = \frac{2N-1}{2N}$. We note that

$$(x-k)^2 > \left(x-k-\frac{1}{2}\right) \left(x-k+\frac{1}{2}\right)$$

and therefore

$$\begin{aligned} \sum_{|k| < N} \frac{1}{(x-k)^2} &< \sum_{|k| < N} \frac{1}{\left(x-k-\frac{1}{2}\right) \left(x-k+\frac{1}{2}\right)} = \sum_{|k| < N} \frac{1}{x-k-\frac{1}{2}} - \frac{1}{x-k+\frac{1}{2}} \\ &= \frac{1}{x-N+\frac{1}{2}} - \frac{1}{x+N-\frac{1}{2}}. \end{aligned}$$

Elementary rearrangements give

$$\begin{aligned} \frac{2N-1}{x^2 - N^2} &\geq \frac{1}{x-N+\frac{1}{2}} - \frac{1}{x+N-\frac{1}{2}} && \Leftrightarrow \\ (2N-1) \left(x-N+\frac{1}{2}\right) \left(x+N-\frac{1}{2}\right) &\geq (x^2 - N^2) \left(x+N-\frac{1}{2} - \left(x-N+\frac{1}{2}\right)\right) && \Leftrightarrow \\ x^2 - \left(N-\frac{1}{2}\right)^2 &\geq x^2 - N^2 \end{aligned}$$

and (2.18) is proved.

Now we like to prove that $F_{\alpha,N}(x) \leq 1$ for $|x| < N$. This is clearly true when $x \in \mathbb{Z}$ and by symmetry we can assume that $x > 0$. Using the well-known formula

$$1 = \frac{\sin^2(\pi x)}{\pi^2} \sum_{k \in \mathbb{Z}} \frac{1}{(x-k)^2}$$

we actually only have to prove that

$$\frac{2N\alpha}{N^2 - x^2} \leq \sum_{|k| \geq N} \frac{1}{(x-k)^2}.$$

Clearly, it suffices to check the case $\alpha = (2N+1)/(2N)$. We use the fact that the trapezoidal rule overestimates convex functions to obtain

$$\begin{aligned} \sum_{|k| \geq N} \frac{1}{(x-k)^2} &\geq \frac{1}{2} \frac{1}{(x-N)^2} + \frac{1}{2} \frac{1}{(x+N)^2} + \int_N^\infty \frac{1}{(x-y)^2} dy + \int_{-\infty}^{-N} \frac{1}{(x-y)^2} dy \\ &= \frac{1}{2} \frac{1}{(x-N)^2} + \frac{1}{2} \frac{1}{(x+N)^2} + \frac{1}{N-x} + \frac{1}{N+x}. \end{aligned}$$

Again we are left with a few elementary rearrangements, we state them for completeness:

$$\begin{aligned} \frac{1}{2} \frac{1}{(x-N)^2} + \frac{1}{2} \frac{1}{(x+N)^2} + \frac{1}{N-x} + \frac{1}{N+x} &\geq \frac{2N+1}{N^2-x^2} && \Leftrightarrow \\ \frac{1}{2} \left(\frac{N+x}{N-x} + \frac{N-x}{N+x} \right) + N+x+N-x &\geq 2N+1, \end{aligned}$$

which is true as $a + a^{-1} \geq 2$ for all $a \in \mathbb{R}_{>0}$. □

Remarks. 1. Selberg's minorant, as stated in [94], corresponds to the case $\alpha = 1$. The author was not able to find the result of Proposition 2.13 in the literature, though the non-uniqueness of extremal minorants is certainly known (it is for example mentioned in [18]). In [34], p. 289, the corresponding result for majorizing functions is stated and credited to unpublished work of A. Selberg, but without a proof.

2. While our approach can be easily extended to find minorants of an interval $[A, B]$ with $A, B \in \mathbb{Z}$, the general case is more difficult. However, one can use this approach to get Beuerling's majorant of the sgn-function. Indeed, it is not difficult to see that

$$F_{1,N}(x-N) \rightarrow F(x)$$

locally uniformly. Note that $\alpha = 1$ is the only possible choice working for all $N \in \mathbb{N}_{>0}$. Then

$$\begin{aligned} B(x) = 2F(x) - 1 &= \frac{\sin^2(\pi x)}{\pi^2} \left(\frac{2}{x} + \sum_{k=1}^{\infty} \frac{2}{(x-k)^2} \right) - 1 \\ &= \frac{\sin^2(\pi x)}{\pi^2} \left(\frac{2}{x} + \sum_{k=1}^{\infty} \frac{1}{(x-k)^2} - \sum_{k=0}^{\infty} \frac{1}{(x+k)^2} \right) \end{aligned}$$

is a minorant of sgn. The corresponding majorant $-B(-x)$ is exactly Beuerling's majorizing function and satisfies

$$\int_{\mathbb{R}} -B(-x) - \operatorname{sgn}(x) \, dx = 1.$$

For any interval $[\alpha, \beta] \subset \mathbb{R}$, one then obtains a minorant by considering

$$\psi_{\alpha,\beta}(x) = \frac{1}{2} (B(x-\alpha) + B(-x+\beta)) \leq \chi_{[\alpha,\beta]}(x). \quad (2.19)$$

This is Selberg's original approach. He then proved that

$$\int_{\mathbb{R}} \chi_{[\alpha,\beta]}(x) - \psi_{\alpha,\beta}(x) \, dx = 1.$$

Note that we are only able to obtain $F_{1,N}$ this way and not the functions we need later on. See [94] for more information and proofs. However, if $\beta - \alpha \notin \mathbb{Z}$, this minorant is not extremal. Extremal functions in this case were announced by Logan in [59]. A proof appeared much later, by Littmann [58].

This concludes the univariate theory of extremal functions we need. In higher dimensions, the results are more complicated. We formulate the problem as follows:

$$\text{Find } f \in \text{PW}^d \text{ minimizing } \int_{\mathbb{R}^d} \square_N^d(x) - f(x) \, dx \text{ subject to } f \leq \square_N^d, \quad (2.20)$$

where $\square_N^d = \chi_{[-N,N]^d}$. Up to today, no extremal solution are known. However, it is known that there are functions attaining the minimum, a result due to Carruth, Gonçalves and Kelly [18].

Theorem 2.14. *The minimum in (2.20) is attained by a function $F \in \text{PW}^d$. Furthermore, there exists a solution with the symmetries of a square, i.e., for all permutations σ of d elements it holds*

that

$$F(x_1, \dots, x_d) = F(|x_{\sigma(1)}|, \dots, |x_{\sigma(d)}|), \quad \text{for all } x_1, x_2, \dots, x_d \in \mathbb{R}.$$

The existence part follows directly from Lemma 12 in [18]. That symmetry can be assumed is stated implicitly and follows easily from the observation that the symmetrization of an $F \in \text{PW}$

$$F_{\text{sym}}(x) = \frac{1}{2^d d!} \sum_{\varepsilon \in \{\pm 1\}^d} \sum_{\sigma} F(\varepsilon_1 x_{\sigma(1)}, \dots, \varepsilon_d x_{\sigma(d)})$$

is still admissible and attains the same value in (2.20).

In the next propositions, we highlight two of the main differences between the one and the higher dimensional case. Already the higher dimensional analog of Proposition 2.12 shows an important difference, as the lower bound growth with N .

Proposition 2.15. *For any $f \in \text{PW}(\mathbb{R}^d)$ and any $N \in \mathbb{N}_{>0}$ with $f(x) \leq \square_N^d(x)$, we have that*

$$\int_{\mathbb{R}^d} \square_N^d(x) - f(x) \, dx \geq (2N)^d - (2N - 1)^d \geq d(2N - 1)^{d-1}.$$

Proof. The proof is a straightforward generalization of the one dimensional result. Clearly, we may assume $f \in \text{PW}^1(\mathbb{R}^d)$. As $f(k) \leq 1$ if $\|k\|_{\infty} \leq N - 1$ and $f(k) \leq 0$ otherwise, using Poisson's summation formula, we obtain

$$\int_{\mathbb{R}^d} \square_N^d(x) - f(x) \, dx = (2N)^d - \hat{f}(0) = (2N)^d - \sum_{k \in \mathbb{Z}^d} f(k) \geq (2N)^d - (2N - 1)^d.$$

The second inequality is a direct application of the mean value theorem. \square

But the main problem is that this lower bound cannot be achieved. Again, this result is due to Carruth, Gonçalves and Kelly [18] in the case $N = 1$. We give a modified proof for arbitrary N .

Theorem 2.16. *Let $F \in \text{PW}^d$ be given, minorizing \square_N^d , where $d \geq 2$. Then*

$$\hat{F}(0) = \int_{\mathbb{R}^d} F(x) \, dx < (2N - 1)^d.$$

Proof. As already explained, we may assume that F obeys the symmetry of a square. Furthermore, the previous proposition already established $\hat{F}(0) \leq (2N - 1)^d$. We start with the case $d = 2$ and assume that $\hat{F}(0) = (2N - 1)^2$. This however is only possible if

$$F(m, n) = \begin{cases} 1 & \text{if } |m|, |n| < N \\ 0 & \text{otherwise.} \end{cases}$$

As $F \leq \square_N^d$, all points (m, n) not of the form $(\pm N, n)$ or $(m, \pm N)$ with $|n|, |m| \leq N$ are local maxima of F . In particular, the partial derivatives vanish. They vanish at (m, n) with $|m| = |n| = N$ as well, due to the surrounding sign pattern and $F(m, n) = 0$. For the same reason, we have $\partial_2 F(\pm N, n) = 0$ and $\partial_1 F(m, \pm N) = 0$ whenever $|m|, |n| < N$.

Now we use the representation of F by the derivative sampling expansion, as stated in (2.15). We see that the functions $x \mapsto F(x, n)$ and $y \mapsto F(m, y)$ are zero for all $|n|, |m| \geq N$. Therefore, $\partial_1^2 F(n, m) = \partial_2^2 F(n, m) = 0$ for all $|n|, |m| \geq N$. But these (m, n) are local maxima and the second-order necessary condition is given by

$$\partial_1^2 F(m, n) \partial_2^2 F(m, n) - (\partial_1 \partial_2 F(m, n))^2 \geq 0.$$

Hence, $\partial_1 \partial_2 F(m, n) = 0$ at such points. Up to now, we followed the proof given in [18] closely. Now

we plug all gathered information into the derivative sampling expansion (2.15):

$$F(x, y) = \frac{\sin^2(\pi x) \sin^2(\pi y)}{\pi^4} \left(\sum_{|m|, |n| < N} \frac{1}{(x-m)^2 (y-n)^2} + \sum_{|n| < N} \frac{2N \partial_1 F(N, n)}{(x^2 - N^2)(y-n)^2} \right. \\ \left. + \sum_{|m| < N} \frac{2N \partial_2 F(m, N)}{(y^2 - N^2)(x-m)^2} + \sum_{0 \leq m, n \leq N} \frac{4mn \partial_1 \partial_2 F(m, n)}{(x^2 - m^2)(y^2 - n^2)} \right)$$

Next, for $\varepsilon > 0$, we consider

$$G(k, \varepsilon) = F(\sqrt{N^2 + \varepsilon}, \sqrt{k^2 + \varepsilon}) \frac{\pi^4 \varepsilon^2}{\sin^2(\pi \sqrt{N^2 + \varepsilon}) \sin^2(\pi \sqrt{k^2 + \varepsilon})} \leq 0.$$

Letting ε go to zero, we obtain for all $|k| \leq N$

$$\partial_1 \partial_2 F(N, k) \leq 0. \quad (2.21)$$

Now we form the limes

$$0 \geq \lim_{k \rightarrow \infty} G(k, \varepsilon) k^2 = \varepsilon \sum_{|n| < N} 2N \partial_1 F(N, n) + \varepsilon 4N \sum_{n=0}^N n \partial_1 \partial_2 F(N, n) + \mathcal{O}(\varepsilon^2).$$

Note that this inequality also holds for $\varepsilon < 0$. Therefore, we have

$$\sum_{|n| < N} 2N \partial_1 F(N, n) + 4N \sum_{n=0}^N n \partial_1 \partial_2 F(N, n) = 0. \quad (2.22)$$

But $x \mapsto F(x, n)$ is in PW^1 , minorizes \square_N and is equal to one at $x = -N + 1, \dots, N - 1$. Therefore, this function is actually of the type $F_{\alpha, N}$ as presented in Proposition 2.13 with

$$\alpha = -\partial F_1(N, n) \geq \frac{2N - 1}{2N}$$

This leads to a contradiction: $\partial_1 F(N, n) < 0$ and (2.21) contradict (2.22).

The case $d > 2$ can be easily deduced from the case $d = 2$. Indeed, if $F \in \text{PW}^d$ minorized \square_N^d and $\hat{F}(0) = (2N - 1)^d$, the function $F(x, y, 0, \dots, 0) \in \text{PW}^2$ would contradict the case $d = 2$. \square

Remark. One of the main results in [18] is, that for all dimensions d larger than a critical threshold, the function $F \in \text{PW}^d$ minorizing \square_1^d and maximizing $\hat{F}(0)$ is the zero function. However, no bound on the critical dimension is currently known.

Though no extremal functions are known in the higher dimensional case, there are some explicit constructions of $F \in \text{PW}^d$ minorizing \square_N^d . In the following discussion, we focus on the two dimensional case. It is intriguing to start with a tensor product ansatz, for example

$$F(x, y) = F_{\alpha, N}(x) F_{\alpha, N}(y).$$

However, this is not a minorant, as $F(x, y) \geq 0$ if $|x|, |y| \geq N$. This unfortunately shows that the function proposed in [54] is actually not a minorant.

To fix this, one has to add a correction term. One suggestion is to consider

$$F(x, y) = F_{1, N}(x) F_{1, N}(y) - \frac{\sin^2(\pi x) \sin^2(\pi y)}{\pi^4} \sum_{m, n = \pm N} \frac{1}{(x-m)^2 (y-n)^2},$$

i.e., subtract a function, which is minus one at the vertices of $[-N, N]^2$ and zero at all other points in \mathbb{Z}^2 (also, all partial derivatives occurring in (2.15) are zero at all points in \mathbb{Z}^2). This is Selberg and Montgomery's construction (which agree in this instance), see [18] for a more detailed discussion and additional references.

To improve on this construction, we consider functions of the form

$$F_N(x, y) = F_{\frac{2N-1}{2N}, N}(x)F_{\frac{2N-1}{2N}, N}(y) - C^2 G_N(x)G_N(y), \quad (2.23)$$

where $C > 0$ is a constant and

$$G_N(x) = \frac{\sin^2(\pi x)}{\pi^2} \left(\frac{1}{(x-N)^2} + \frac{1}{(x+N)^2} - \frac{2}{x^2 - N^2} \right).$$

Note that G_N and its first derivative are zero at all $k \in \mathbb{Z}$ except $k = \pm N$.

Proposition 2.17. *Let $F_N \in \text{PW}^2$ be given by (2.23). Then $F_N \leq \square_{\frac{1}{N}}^d$ if $C \geq \frac{1}{2}$ and*

$$\hat{F}_N(0) = (2N-1)^2 - 1.$$

In the special case $N = 1$, we can choose $C \geq \frac{1}{4}$, which results in $\hat{F}_1(0) = 3/4$. We subsequently use F_N to denote the function with $C = \frac{1}{4}$ for $N = 1$ and $C = \frac{1}{2}$ otherwise.

Proof. Note that $G_N(x) \geq 0$ for all $x \in \mathbb{R}$. The only thing left to prove is $F_N(x, y) \leq 0$ if $|x|, |y| \geq 0$ and $C = \frac{1}{2}$. This follows from

$$\left| F_{\frac{2N-1}{2N}, N}(x) \right| \leq \frac{G_N(x)}{2}, \quad \text{for all } |x| \geq N.$$

Equivalently, we have to show that

$$\frac{2N-1}{x^2 - N^2} - \sum_{|k| < N} \frac{1}{(x-k)^2} \leq \frac{1}{2(x-N)^2} + \frac{1}{2(x+N)^2} - \frac{1}{x^2 - N^2}, \quad \forall x > N.$$

As already in the proof of Proposition 2.13, we use that the trapezoidal rule overestimates convex functions, which implies

$$\frac{2N}{x^2 - N^2} = \int_{-N}^N \frac{1}{(x-y)^2} dy \leq \frac{1}{2(x+N)^2} + \sum_{|k| < N} \frac{1}{(x-k)^2} + \frac{1}{2(x-N)^2}$$

and the claim follows. To calculate $\hat{F}_N(0)$, we again apply the Poisson summation formula and obtain

$$\hat{F}_N(0) = \sum_{k \in \mathbb{Z}^d} F_N(k) = (2N-1)^d - 4C^2.$$

In the special case $N = 1$, we apply a more direct approach. Note that

$$|F_{\frac{1}{2}, 1}(x)| = \frac{1}{x^2(x^2-1)} \leq C \left(\frac{1}{(x-1)^2} + \frac{1}{(x+1)^2} + \frac{2}{1-x^2} \right) = \frac{4C}{(x^2-1)^2},$$

which is equivalent to

$$\frac{x^2-1}{4x^2} \leq C.$$

This, however, holds true for all $|x| \geq 1$ whenever $C \geq \frac{1}{4}$. □

Remark. The case of majorants is a lot easier. Indeed, if H_N is a majorant for \square_N , the tensor product

$$H(x_1, \dots, x_d) = H_N(x_1) \cdots H_N(x_d)$$

majorizes \square_N^d and is even extremal, as virtually the same argument used in Proposition 2.15 shows. This tensor product construction was first given by Liao [54].

Finally, we give a short list of other approaches. Closest to the here presented minorant is the function constructed in [18] (covering only the case $N = 1$), which is given by

$$\tilde{F}(x, y) = F_{\frac{1}{2}, 1}(x)F_{\frac{1}{2}, 1}(y) - \frac{1}{16}\tilde{G}(x)\tilde{G}(y),$$

where

$$\tilde{G}(x) = \frac{\sin^2(\pi x)}{\pi^2} \left(\frac{1}{(x-1)^2} + \frac{1}{(x+1)^2} + \frac{2}{x^2-1} \right).$$

This results in a better (that is larger) $\hat{F}(0) = \frac{63}{64}$, we gave F_1 just for reference. However, the property we need later on, that \hat{F} has a global maximum at zero, is not satisfied (even when the paper actually states that they numerically observed this). F_2 , on the other hand, has this property. Whether such a function exists for the case $N = 1$ remains open.

For the already mentioned functions constructed by Selberg and Montgomery, we refer the interested reader again to [18].

If one changes the problem a bit, a lot more is known. Most notably, if one wishes to find functions satisfying

$$\text{supp } \hat{F} \subset B_1^2, \quad F \leq \chi_{B_1^2}, \quad (2.24)$$

where B_r^p is the ball $\{x \in \mathbb{R}^d : \|x\|_p \leq r\}$, extremal results are given by Holt and Vaaler in [43].

A very nice and explicit construction in the case of

$$\text{supp } \hat{F} \subset B_{2r}^p, \quad F \leq \chi_{B_R^2} \quad (2.25)$$

for $p \in [1, \infty]$ is due to Komornik and Loreti, see [51]. The basic idea is as follows: Pick the function $G \in H_0^1(B_r^p)$ (where $H_0^1(B_1^p)$ denotes the usual Sobolev space), which is an eigenfunction of $-\Delta$ (where Δ is of course the Laplace operator) corresponding to the smallest eigenvalue of $-\Delta$. This function is known to be positive. Then

$$\begin{aligned} \hat{F} &= (R^2 + \Delta)G * G, \\ F(x) &= (R^2 - |x|^2)|\hat{G}|^2(x) \end{aligned}$$

satisfies (2.25). For more details we refer to [51]. One interesting extension of this idea is to replace the Laplace operator with $\sum_j \frac{\partial^k}{\partial x_j^k}$. Then a similar construction yields functions satisfying

$$\text{supp } \hat{F} \subset B_{2r}^\infty, \quad F \leq \chi_{B_R^q}.$$

This idea has been investigated in [52].

2.3 Conditional Well-Posedness

We are now ready to prove the main results of this chapter. In the first part of this section, we restrict ourselves to the univariate setting. Recall that we are interested to give a good lower bound to

$$\|\mathcal{P}_N(f - g)\|_2^2 = \sum_{k=-N}^N |(f - g)(k)|^2,$$

where $f, g \in \mathcal{S}^1$ have well-separated frequencies. We collect such exponential sums in

$$\mathcal{S}^d(q) = \{f \in \mathcal{S}^d : \text{sep } f \geq q\}.$$

We start by the now classical results for $\mathcal{P}_N(f)$. We use the dilation operator, defined by

$$\text{Dil}_\alpha f(x) = f(\alpha x).$$

Recall that

$$\mathcal{F} \text{Dil}_\alpha = \frac{1}{\alpha^d} \text{Dil}_{\frac{1}{\alpha}} \mathcal{F},$$

a relation we will frequently use.

Theorem 2.18. *Let $f \in \mathcal{S}^1(q)$ for $q \in (0, 1]$ and two integers $A < B$ be given. Then*

$$\sum_{k=A}^B |f(k)|^2 \geq \left(B - A + 2 - \frac{1}{q} \right) \|c^f\|_2^2$$

Proof. Let $\psi_{q(A-1), q(B+1)}$ be defined as in (2.19). Then

$$\Psi(x) := \text{Dil}_q \psi_{q(A-1), q(B+1)}(x) \leq \begin{cases} 1 & x \in (A-1, B+1) \\ 0 & x \notin (A-1, B+1) \end{cases}$$

and we can calculate, using Poisson's summation formula, as given in Theorem 2.11

$$\begin{aligned} \sum_{k=A}^B |f(k)|^2 \Psi(k) &\geq \sum_{k \in \mathbb{Z}} |f(k)|^2 \Psi(k) \\ &= \sum_{y, w \in Y^f} c_y \overline{c_w} \sum_{k \in \mathbb{Z}} e^{2\pi i(y-w)k} \Psi(k) \\ &= \sum_{y, w \in Y^f} c_y \overline{c_w} \sum_{k \in \mathbb{Z}} \hat{\Psi}(w - y - k). \end{aligned}$$

But $\hat{\Psi}(w) = \frac{1}{q} \hat{\psi}_{q(A-1), q(B+1)}(w/q)$, in particular it is a continuous function with $\text{supp } \hat{\Psi} \subset [-q, q]$, which implies that for all k and all $w \neq y \in Y^f$ we have that $\hat{\Psi}(w - y - k) = 0$. Further,

$$\hat{\Psi}(0) = \frac{1}{q} \hat{\psi}_{q(A-1), q(B+1)}(0) = \frac{1}{q} ((B - A + 2)q - 1).$$

This finishes the proof. \square

Remarks. 1. The case $q = 1$ is only possible if $f(x) = ce^{2\pi iyx}$. Note that in this instance, the lower bound is exact. Furthermore, the lower bound is only non-trivial if $q > 1/(B - A + 2)$. In particular, $\text{ord } f \leq B - A + 1$. This is sharp, as for every $Y \subset [0, 1)$ with $|Y| = B - A + 2$ we can find a $f \in \mathcal{S}^1$ with $Y^f = Y$ and $f(k) = 0$ for all $k = A, A + 1, \dots, B$.

2. This clearly gives for all $f \in \mathcal{S}^1(q)$ the lower bound

$$\|\mathcal{P}_N(f)\|_2^2 \geq \left(2N + 2 - \frac{1}{q} \right) \|c^f\|_2^2.$$

3. Note that this bound is slightly better than the original bound, obtain by Moitra in [62] by more or less the same method. This improvement is due to the fact, that we used $\psi_{q(A-1), q(B+1)}$ instead of $\psi_{qA, qB}$. It also improves on a result given in [4], which is itself an improvement on Moitra's result, relying on a different method.

4. Virtually the same argument, but using the majorant instead, gives an upper bound, namely

$$\sum_{k=A}^B |f(k)|^2 \leq \left(B - A + \frac{1}{q} \right) \|c^f\|_2^2, \quad \forall f \in \mathcal{S}^1(q).$$

Details can be found again in Moitra's work [62] as well as in [94]. Note that in this case we have to use the majorant of $[qA, qB]$. Also, the case $q = 1$ again results in equality.

This result can be reinterpreted as given bounds on the singular value of Vandermonde matrices. This was exactly the main motivation in Moitra's work. We repeat the argument, again obtaining slightly better results for the lower bound. We remark that the more general case of nodes in the unit disc is covered in [4].

Definition 2.19. *For $N \in \mathbb{N}^*$ and $y \in [0, 1)$, let*

$$v_N(y) = \left(1, e^{2\pi iy}, \dots, e^{2\pi iy(N-1)} \right)^T \in \mathbb{C}^N.$$

For $y_1, \dots, y_M \in [0, 1)$ we define the Vandermonde matrix

$$V_N(y_1, \dots, y_M) = [v_N(y_1) \cdots v_N(y_M)] \in \mathbb{C}^{N \times M}.$$

If y_1, \dots, y_M are clear, we use the notation V_N . Recall that if y_j are mutually distinct, $V_N(y_1, \dots, y_M)$ has full rank.

Corollary 2.20. *Let $y_1, \dots, y_M \in [0, 1)$ have separation distance q . Let σ_{\min} and σ_{\max} be the smallest and largest singular value of $V_N(y_1, \dots, y_M)$ respectively. Then the following bounds hold true:*

$$\begin{aligned} \sigma_{\min}^2 &\geq \left(N + 1 - \frac{1}{q} \right) \\ \sigma_{\max}^2 &\leq \left(N - 1 + \frac{1}{q} \right). \end{aligned}$$

Proof. For any $c \in \mathbb{C}^N$ we let

$$f(x) = \sum_{j=1}^M c_j e^{2\pi i x \cdot y_j}$$

which gives

$$\|V_N(y_1, \dots, y_M)c\|_2^2 = \sum_{k=0}^{N-1} |f(k)|^2.$$

Using the aforementioned results together with the fact that $\sigma_{\min}(A) = \min_{x \neq 0} \|Ax\|_2 / \|x\|_2$ and the corresponding result for the largest singular value gives the state bounds. \square

Remark. Again, it is interesting to compare with Moitra [62], who gave the lower bound

$$\sigma_{\min}^2 \geq \left(N - 1 - \frac{1}{q} \right)$$

and Aubel and Bölcskei [4], who proved

$$\sigma_{\min}^2 \geq \left(N + \frac{1}{2} - \frac{1}{q} \right).$$

While our gain seems to be small, it covers all possible q and in particular more cases of square Vandermonde matrices, as only if $q \leq \frac{1}{N}$ one can find N frequencies, which are q -separated.

Now we extend this technique to estimate the difference $f - g$ of two well-separated exponential sums. Let Ψ as in the proof of Theorem 2.18 and

$$f(x) = \sum_{y \in Y^f} c_y^f e^{2\pi i y x}, \quad g(x) = \sum_{y' \in Y^g} c_{y'}^g e^{2\pi i y' x}.$$

Following the calculation in Theorem 2.18, we obtain

$$\begin{aligned} \sum_{k=A}^B |f(k) - g(k)|^2 &\geq \sum_{k=A}^B |f(k) - g(k)|^2 \Psi(k) \\ &= \sum_{y, w \in Y^f} c_y^f \overline{c_w^f} \sum_{k \in \mathbb{Z}} e^{2\pi i (y-w)k} \Psi(k) + \sum_{y', w' \in Y^g} c_{y'}^g \overline{c_{w'}^g} \sum_{k \in \mathbb{Z}} e^{2\pi i (y'-w')k} \Psi(k) \\ &\quad + 2\Re \left(\sum_{y \in Y^f} \sum_{y' \in Y^g} c_y^f \overline{c_{y'}^g} \sum_{k \in \mathbb{Z}} e^{2\pi i (y-y')k} \Psi(k) \right). \end{aligned}$$

If $f, g \in \mathcal{S}^1(2q)$, for any $y \in Y^f$ there is at most one $y' \in Y^g$ with $|y - y'|_{\mathbb{T}} < q$. We call $y' = n(y)$ the

neighbor of y and collect all $y \in Y^f$ with a neighbor $N(Y^f)$. The sum simplifies to

$$\sum_{k=A}^B |f(k) - g(k)|^2 \geq \hat{\Psi}(0) \left(\sum_{y \in Y^f} |c_y^f|^2 + \sum_{y' \in Y^g} |c_{y'}^g|^2 \right) + 2\Re \left(\sum_{y \in N(Y^f)} c_y^f \overline{c_{n(y)}^g} \hat{\Psi}(|y - n(y)|_{\mathbb{T}}) \right) \quad (2.26)$$

Now we need the Fourier transform of Ψ . We calculate it in the following lemma.

Lemma 2.21. *The Fourier transform of*

$$F_{\alpha, N}(x) = \frac{\sin^2(\pi x)}{\pi^2} \left(\frac{2N\alpha}{N^2 - x^2} + \sum_{|k| < N} \frac{1}{(x - k)^2} \right)$$

is given by

$$\hat{F}_{\alpha, N}(w) = \begin{cases} \frac{\alpha}{\pi} \sin(2\pi N|w|) + (1 - |w|)D_{N-1}(w) & \text{for } |w| \leq 1 \\ 0 & \text{otherwise} \end{cases},$$

where D_{N-1} is the Dirichlet kernel of order $N - 1$, which is given by

$$D_{N-1}(w) = 1 + 2 \sum_{k=1}^{N-1} \cos(2\pi kw).$$

Proof. We only sketch the proof, which is an easy calculation (at least if one uses the theory of distributions). Using the Fourier convolution theorem gives

$$\begin{aligned} \mathcal{F} \left[\sum_{|k| < N} \frac{\sin^2(\pi x)}{\pi^2(x - k)^2} \right] (w) &= \mathcal{F} \left[\text{sinc}^2 * \sum_{|k| < N} \delta_k \right] (w) = \mathcal{F}[\text{sinc}^2](w) \cdot \mathcal{F} \left[\sum_{|k| < N} \delta_k \right] (w) \\ &= (1 - |w|)D_{N-1}(w). \end{aligned}$$

The other part follows by

$$\begin{aligned} \mathcal{F} \left[\frac{\sin^2(\pi(x + N))}{\pi(x + N)} \right] (w) &= e^{2\pi i N w} (\mathcal{F}[\text{sinc}] * \mathcal{F}[\sin(\pi x)])(w) = \frac{e^{2\pi i N w}}{2i} \left(\chi_{[-\frac{1}{2}, \frac{1}{2}]} * (\delta_{\frac{1}{2}} - \delta_{-\frac{1}{2}}) \right) (w) \\ &= \begin{cases} \frac{e^{2\pi i N w}}{2i} & \text{if } w \in [0, 1], \\ -\frac{e^{2\pi i N w}}{2i} & \text{if } w \in [-1, 0], \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

□

Furthermore, we need reasonably sharp estimates of sin and cos.

Lemma 2.22. *We have that*

$$\begin{aligned} \sin(x) &\geq x - \frac{x^3}{6} && \forall x \geq 0 \\ \cos(x) &\geq 1 - \frac{x^2}{2} && \forall x \in \mathbb{R}. \end{aligned}$$

Furthermore, the following upper bounds hold true:

$$\begin{aligned} \sin(x) &\leq x - x^3 \left(\frac{1}{6} - \frac{\pi^2}{120} \right) && \forall x \in [0, \pi] \\ \cos(x) &\leq 1 - x^2 \left(\frac{1}{2} - \frac{\pi^2}{96} \right) && \forall x \in \left[-\frac{\pi}{2}, \frac{\pi}{2} \right]. \end{aligned}$$

Proof. We sketch the basic proofs for the reader's convenience. The first two inequalities follow from integrating $\sin(x) \leq x$ once (respectively twice).

The upper bound for cosine can be derived by using the Taylor formula, which gives the existence of a $\xi \in [-\frac{\pi}{2}, \frac{\pi}{2}]$ such that

$$\cos(x) - 1 + \frac{x^2}{2} = \frac{\cos(\xi)}{24}x^4 \leq \frac{\pi^2}{96}x^2, \quad \forall x \in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right].$$

The corresponding result for sine follows similarly. \square

Now consider $F_{\alpha,1}(x)$. Clearly, $\hat{F}_{\alpha,1}$ is even and

$$\hat{F}'_{\alpha,1}(w) = 2\alpha \cos(2\pi w) - 1 \quad \text{for } w > 0.$$

Therefore, $\hat{F}_{\alpha,1}$ has a local minimum in zero except when $\alpha = \frac{1}{2}$, the smallest possible choice. $\hat{F}_{\frac{1}{2},1}$ is monotonically decreasing in $[0, 1]$. We use this function to obtain our first stability result.

Proposition 2.23. *Let $f, g \in \mathcal{S}^1(2q)$ be given. Assume that $q \geq \frac{1}{N+1}$ and that*

$$\|\mathcal{P}_N(f - g)\|_2^2 = \sum_{k=-N}^N |f(k) - g(k)|^2 < c_{\min}^2(N+1), \quad (2.27)$$

where c_{\min} is a lower bound on the modulus of the coefficients of f and g . Then for every $y \in Y^f$ we find exactly one $y' = n(y) \in Y^g$ with $|y - y'| < \frac{1}{2(N+1)}$ and vice versa.

Furthermore, we get

$$3(N+1)^4 \sum_{y \in Y^f} \left((|c_y^f|^2 + |c_{n(y)}^g|^2) |y - n(y)|_{\mathbb{T}}^3 \right) + \frac{N+1}{2} \sum_{y \in Y^f} |c_y^f - c_{n(y)}^g|^2 \leq \|\mathcal{P}_N(f - g)\|_2^2. \quad (2.28)$$

Proof. We use

$$\Psi(x) = \text{Dil}_{\frac{1}{N+1}} F_{1, \frac{1}{2}}(x),$$

which satisfies $\Psi(x) \leq 1$ and $\Psi(k) \leq 0$ for all $|k| \geq N+1$. Now we use (2.26) and the algebraic fact

$$(|a|^2 + |b|^2)C - 2c\Re(ab) = (C - c)(|a|^2 + |b|^2) + |a - b|^2c \quad \text{for all } a, b, c, C \in \mathbb{C}$$

to obtain

$$\begin{aligned} \sum_{k=-N}^N |f(k) - g(k)|^2 &\geq \sum_{y \in N(Y^f)} (|c_y^f|^2 + |c_{n(y)}^g|^2) (\hat{\Psi}(0) - \hat{\Psi}(|y - n(y)|_{\mathbb{T}})) + |c_y^f - c_{n(y)}^g|^2 \hat{\Psi}(|y - n(y)|_{\mathbb{T}}) \\ &\quad + \hat{\Psi}(0) \left(\sum_{y \in Y^f \setminus N(Y^f)} |c_y^f|^2 + \sum_{\substack{y' \in Y^g \\ \#n^{-1}(y')}} |c_{y'}^g|^2 \right). \end{aligned}$$

As $\hat{\Psi}(0) = N+1$, assumption (2.27) gives that $Y^f = N(Y^f)$ and $Y^g = \{n(y) : y \in Y^f\}$, i.e., for every $y \in Y^f$ we find exactly one $y' \in Y^g$ with $|y - y'| < q$ and vice versa.

Now we estimate

$$\hat{\Psi}(0) - \hat{\Psi}(x) = (N+1) \left((N+1)|x| - \frac{1}{2\pi} \sin(2\pi(N+1)|x|) \right) \quad \text{for } |x| \leq \frac{1}{N+1}.$$

Using Lemma 2.22 gives

$$\hat{\Psi}(0) - \hat{\Psi}(x) \geq 4\pi^2 \left(\frac{1}{6} - \frac{\pi^2}{120} \right) (N+1)^4 |x|^3 \geq 3(N+1)^4 |x|^3 \quad \text{for } |x| \leq \frac{1}{2(N+1)} \quad (2.29)$$

while $\sin(x) \leq 0$ on $[\pi, 2\pi]$ gives

$$\hat{\Psi}(0) - \hat{\Psi}(x) \geq (N+1)^2|x| \quad \text{for } |x| \in \left[\frac{1}{2(N+1)}, \frac{1}{N+1} \right].$$

Now if there were a $y^* \in Y^f$ with $|y^* - n(y^*)|_{\mathbb{T}} \in \left[\frac{1}{2(N+1)}, \frac{1}{N+1} \right]$ we would get

$$\sum_{k=-N}^N |f(k) - g(k)|^2 \geq \left(|c_{y^*}^f|^2 + |c_{n(y^*)}^g|^2 \right) (\hat{\Psi}(0) - \hat{\Psi}(|y^* - n(y^*)|_{\mathbb{T}})) \geq 2c_{\min}^2 (N+1)^2 |y^* - n(y^*)|_{\mathbb{T}}.$$

Again, (2.27) gives rise to a contradiction. Therefore, we can use (2.29), which yields

$$\sum_{k=-N}^N |f(k) - g(k)|^2 \geq 3(N+1)^4 \sum_{y \in Y^f} (|c_y^f|^2 + |c_{n(y)}^g|^2) |y - n(y)|_{\mathbb{T}}^3 + |c_y^f - c_{n(y)}^g|^2 \hat{\Psi}(|y - n(y)|_{\mathbb{T}}).$$

Finally, the monotonicity of $\hat{\Psi}$ together with $|y - n(y)|_{\mathbb{T}} < \frac{1}{2(N+1)}$ and $\hat{\Psi}\left(\frac{1}{2(N+1)}\right) = \frac{N+1}{2}$ gives the result. \square

Remarks. There are some possible variations of this result. For example, if one observes a very pronounced frequency $y \in Y^f$, i.e., $|c_y| \gg c_{\min}$, and one is only interested to see if a close frequency in Y^g exists, it is possible to relax (2.27) to $\|\mathcal{P}_N(f - g)\|_2^2 \lesssim N|c_y|^2$, to get an error estimate for $d(y, Y^g)$. We refrain from giving the details, trusting that any reader interested in such results can carry out the necessary analysis him/herself, following the proofs presented here.

The condition (2.27) ensures that f and g are close enough so that the frequencies of f and g can be related. As

$$\|\mathcal{P}_N(f - (f + c_{\min}e^{2\pi iy}))\|_2^2 = (2N+1)c_{\min}^2$$

shows, it is (up to a factor of two) necessary.

A similar observation shows that the rate in $c_y^f - c_{n(y)}^g$ given in (2.28) is optimal. Indeed,

$$\|\mathcal{P}_N(c\epsilon e^{2\pi iy} - (c + c_1)e^{2\pi iy})\|_2^2 = (2N+1)|c_1|^2.$$

Of course, we are mostly interested in the frequencies. We consider a similar test case in the following lemma.

Lemma 2.24. *The following estimate holds true:*

$$\|\mathcal{P}_N(e^{2\pi i(y+\epsilon)\cdot} - e^{2\pi iy\cdot})\|_2^2 = \sum_{k=-N}^N |e^{2\pi i(y+\epsilon)k} - e^{2\pi iyk}|^2 \leq \frac{4\pi^2}{3}(N+1)^3\epsilon^2.$$

More generally, for two integers $A < B$ and $\epsilon > 0$ sufficiently small we obtain

$$\sum_{k=A}^B |e^{2\pi i(y+\epsilon)k} - e^{2\pi iyk}|^2 \sim_{A,B} \epsilon^2.$$

Proof. This is the result of a direct calculation, using the mean value theorem:

$$\sum_{k=-N}^N |e^{2\pi i(y+\epsilon)k} - e^{2\pi iyk}|^2 = \sum_{k=-N}^N |e^{2\pi i\epsilon k} - 1|^2 \leq \sum_{k=-N}^N 4\pi^2 k^2 \epsilon^2 = 4\pi^2 \epsilon^2 \frac{N(N+1)(2N+1)}{6}.$$

For the second claim, $\lesssim_{A,B}$ follows by the same calculation. The lower bound on the other hand follows with

$$|e^{2\pi i(y+\epsilon)k} - e^{2\pi iyk}|^2 = 2 - 2\cos(2\pi\epsilon k) \geq 8\pi k^2 \epsilon^2$$

for ϵ sufficiently small, due to $1 - \cos(\pi t) \geq \pi t^2$ for all $|t| \leq \frac{1}{2}$. \square

Remarks. 1. The Lemma proves that any family of exponentials $(e^{2\pi iy})_{y \in Y}$ satisfying a lower bound like

$$\|c\|_2^2 \lesssim \left\| \mathcal{P} \left(\sum_{y \in Y} c_y e^{2\pi iy} \right) \right\|_2^2, \quad \text{for all } c \in \mathbb{C}^Y,$$

has to be reasonably well-separated. For a precise statement, see [57]. Extensions to multivariate exponential sums are available as well, see [75].

2. While not our concern here, sometimes one is interested in the case of distinguishing l near-colliding nodes (the lemma discusses the case $l = 1$). A result covering this case is given in [1].

3. The necessity for separation is even more important than the lemma indicates. In fact, there are two $f_1, f_2 \in \mathcal{S}^1(q)$ such that exponentially small noise makes them indistinguishable, if only $\mathcal{P}_N(f_j)$ are known, where $qN < (1 - \varepsilon)$. For a precise statement, see [62], Corollary 3.2.

Unfortunately, (2.28) only gives

$$6(N+1)^4 \varepsilon^2 \leq \|\mathcal{P}_N(e^{2\pi i(y+\varepsilon)\cdot} - e^{2\pi iy\cdot})\|_2^2$$

which is strictly worse (as $\varepsilon < \frac{1}{2(N+1)}$ whenever (2.27) holds). Before we give an estimate realizing this order, we state an interesting consequence for functions in PW minorizing an interval.

Proposition 2.25. *Let $\psi \in \text{PW}$ be given, satisfying*

$$\psi(x) \leq \chi_{[A,B]}(x).$$

for any real numbers $A < B$. Then there is a neighborhood I of zero and a constant $C > 0$ such that

$$\hat{\psi}(0) - \hat{\psi}(w) \leq C|w|^2, \quad \text{for all } w \in I.$$

Proof. Following the lines of the proof of Theorem 2.18 with $f(x) = 1$ and $g(x) = e^{2\pi i \varepsilon x}$ gives, using Lemma 2.24,

$$\varepsilon^2 \gtrsim_{A,B} \sum_{k=A}^B |f(k) - g(k)|^2 \geq 2(\hat{\psi}(0) - \hat{\psi}(\varepsilon)).$$

□

Now we improve the order in (2.28) to the optimal one. The proof is quite similar, only that we use $F_{2, \frac{3}{4}}$ instead of $F_{1, \frac{1}{2}}$. However, we need Proposition 2.23 to find for each $y \in Y^f$ a close neighbor in Y^g .

Theorem 2.26. *Let $f, g \in \mathcal{S}^1(2q)$. Assume that $q \geq \frac{2}{N+1}$ and*

$$\|\mathcal{P}_N(f - g)\|_2^2 = \sum_{k=-N}^N |f(k) - g(k)|^2 < c_{\min}^2(N+1), \quad (2.30)$$

where c_{\min} is a lower bound on the modulus of the coefficients of f and g . Then for every $y \in Y^f$ we find exactly one $y' = n(y) \in Y^g$ with $|y - y'|_{\mathbb{T}} < \frac{1}{2(N+1)}$ and vice versa.

Furthermore, the following estimate holds true:

$$\frac{\pi^2}{3}(N+1)^3 \sum_{y \in Y^f} (|c_y^f|^2 + |c_{n(y)}^g|^2) |y - n(y)|_{\mathbb{T}}^2 + \frac{3}{8}(N+1) \sum_{y \in Y^f} |c_y^f - c_{n(y)}^g|^2 \leq \|\mathcal{P}_N(f - g)\|_2^2$$

Proof. Let

$$\Psi(x) = \text{Dil}_{\frac{2}{N+1}} F_{2, \frac{3}{4}}(x) \leq \begin{cases} 1, & \text{if } |x| \leq N+1 \\ 0, & \text{if } |x| \geq N+1 \end{cases}.$$

By again using $\sin(x) \leq x - \frac{x^3}{\pi^2}$ for $x \in [0, \pi]$ and

$$\cos(x) \leq 1 - \frac{x^2}{\pi} \quad \text{for } x \in [0, \pi]$$

and Lemma 2.21 we obtain for $w \in [0, \frac{1}{4}]$

$$\begin{aligned} \hat{F}_{2, \frac{3}{4}}(0) - \hat{F}_{2, \frac{3}{4}}(w) &= 3 - \frac{3}{4\pi} \sin(4\pi w) - (1-w)(1+2\cos(2\pi w)) \\ &\geq 3 - \frac{3}{4\pi} \left(4\pi w - 4^3 \pi^3 \left(\frac{1}{6} - \frac{\pi^2}{96} \right) w^3 \right) - (1-w) \left(3 - 8\pi^2 \left(\frac{1}{2} - \frac{\pi^2}{96} \right) w^2 \right) \\ &= 8\pi^2 \left(\frac{1}{2} - \frac{\pi^2}{96} \right) w^2 + 8\pi^2 \left(6 \left(\frac{1}{6} - \frac{\pi^2}{96} \right) - \frac{1}{2} + \frac{\pi^2}{96} \right) w^3 \\ &\geq \frac{8\pi^2}{3} w^2. \end{aligned}$$

Therefore,

$$\hat{\Psi}(0) - \hat{\Psi}(w) \geq \frac{\pi^2}{3} (N+1)^3 w^2 \quad \text{for } w \in \left[0, \frac{1}{2(N+1)} \right].$$

It is a routine exercise to check that $\hat{F}_{2, \frac{3}{4}}$ is monotonically decreasing on $[0, \frac{1}{4}]$. Indeed,

$$\begin{aligned} \hat{F}'_{2, \frac{3}{4}}(w) &= 3 \cos(4\pi w) - 1 - 2 \cos(2\pi w) - 4\pi \sin(2\pi w)(1-w) \\ &\leq 3(\cos^2(2\pi w) - \sin^2(2\pi w)) - 1 - 2 \cos(2\pi w) \\ &\leq \cos(2\pi w)(3 \cos(2\pi w) - 2) - 1 \leq 0. \end{aligned}$$

By Proposition 2.23 we find for each $y \in Y^f$ exactly one $n(y) \in Y^g$ with $|y - n(y)|_{\mathbb{T}} < \frac{1}{2(N+1)}$ and hence

$$\hat{\Psi}(|y - n(y)|_{\mathbb{T}}) = \frac{N+1}{2} \hat{F}_{2, \frac{3}{4}} \left(\frac{|y - n(y)|_{\mathbb{T}}(N+1)}{2} \right) \geq \frac{N+1}{2} \hat{F}_{2, \frac{3}{4}} \left(\frac{1}{4} \right) = \frac{3}{8} (N+1).$$

Repeating verbatim the argument given in the proof of Proposition 2.23 yields the claim. \square

Remarks. 1. As already remarked, the exponents in the error term are optimal in N , $|y - n(y)|_{\mathbb{T}}^2$ and $|c_y^f - c_{n(y)}^g|^2$, while all constants are explicitly given and reasonably large.

2. This result can be interpreted as a conditional well-posedness property of the frequency estimation problem: If we use the model of well-separated exponential sums and have sufficiently many samples available, closeness of the measurements guarantees closeness of the parameters of interest.

One actual application of this result are a-posteriori error estimates. Assume we have given noisy samples

$$\tilde{s}_N(f^*) = \mathcal{P}_N(f^*) + \varepsilon,$$

where ε is some noise vector satisfying $\|\varepsilon\|_2 \leq \eta$. Further, assume that we used any recovery algorithm, resulting in f . We cannot apply Theorem 2.26 directly, as we do not know $\mathcal{P}_N(f^*)$. But we can estimate

$$(\|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2 + \eta)^2 \geq \|\mathcal{P}_N(f) - \mathcal{P}_N(f^*)\|_2^2$$

and (at least if (2.27) is satisfied) give a rather tight error estimate. Note that we usually expect η to be of order $\mathcal{O}(N)$. If, for example, all $|\varepsilon_j|$ have values in $[0, c_{\min}/4]$, we can choose $\eta = \frac{\sqrt{2N+1}}{4} c_{\min}$.

Now we estimate the error of the frequencies. We have that

$$\begin{aligned} \frac{\pi^2}{3}(N+1)^3 \sum_{y \in Y^f} (|c_y^f|^2 + |c_{n(y)}^g|^2) |y - n(y)|_{\mathbb{T}}^2 + \frac{3}{8}(N+1) \sum_{y \in Y^f} |c_y^f - c_{n(y)}^g|^2 \\ \geq \frac{2\pi^2}{3}(N+1)^3 c_{\min}^2 \sum_{y \in Y^f} |y - n(y)|_{\mathbb{T}}^2. \end{aligned}$$

This gives for the frequency error an estimate of the form

$$\sum_{y \in Y^f} |y - n(y)|_{\mathbb{T}}^2 \leq \frac{3}{4\pi^2(N+1)} \left(\frac{\|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2}{c_{\min}(N+1)} + \frac{\sqrt{2N+1}}{2(N+1)} \right)^2.$$

The right-hand side is of order N^{-2} . Note that we did not assume that ε_j are independent, this result holds true even if the noise is of the form $\varepsilon_j = \frac{c_{\min}}{4} e^{2\pi i y j}$ for a $y \in [0, 1)$.

To give at least one result with a more specific noise model, we consider the case of complex white noise, i.e., $\varepsilon_j \sim X_{j,1} + iX_{j,2}$ where $X_{j,l} \sim \mathcal{N}(0, \sigma^2)$ for $j = -N, \dots, N$ and $l = 1, 2$ are independent, identically distributed normal random variables with mean zero and variance σ^2 . This results in

$$\begin{aligned} \|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2^2 &= \|\mathcal{P}_N(f) - \mathcal{P}_N(f^*)\|_2^2 + \sum_{j=-N}^N (X_{j,1}^2 + X_{j,2}^2) - 2\Re\langle \mathcal{P}_N(f - f^*), \varepsilon \rangle \\ &= \|\mathcal{P}_N(f) - \mathcal{P}_N(f^*)\|_2^2 + \sigma^2 Y - 2\sigma \|\mathcal{P}_N(f - f^*)\|_2 Z, \end{aligned}$$

where $Z \sim \mathcal{N}(0, 1)$ (due to the rotational invariance of a Gaussian random vector) and $Y \sim \chi_{4N+2}^2$, i.e., Y is χ^2 -distributed with $4N + 2$ degrees of freedom. Rearranging gives

$$(\|\mathcal{P}_N(f) - \mathcal{P}_N(f^*)\|_2 - \sigma Z)^2 = \|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2^2 - \sigma^2 Y + \sigma^2 Z^2. \quad (2.31)$$

Now we use standard tools from probability theory to bound Y and Z from above. The tail of a Gaussian random variable can be simply estimated by:

$$\Pr(|Z| \geq t_1) \leq e^{-\frac{t_1^2}{2}} \quad \text{for all } t_1 > 0, \quad (2.32)$$

see for example [32], Proposition 7.5.

Further, the expected value of Y is given by $\mathbb{E}Y = 4N + 2$ and

$$\Pr(|Y - (4N + 2)| \geq (4N + 2)t_2) \leq 2e^{-(4N+2)t_2^2/8} \quad \text{for all } t_2 \in (0, 1) \quad (2.33)$$

by the Bernstein inequality for sums of Gamma random variables, see [8], Theorem 2.57.

Corollary 2.27. *Let $f, f^* \in \mathcal{S}^1(2q)$ be fixed with $q \geq \frac{2}{N+1}$. Further, assume that we know f and*

$$\tilde{s}_N(f^*) = \mathcal{P}_N(f^*) + \varepsilon \in \mathbb{C}^{2N+1},$$

where $\varepsilon_j = X_{j,1} + iX_{j,2}$ and $X_{j,l}$ are pairwise independent, normally distributed random variables with mean zero and variance σ^2 . Then for any $\delta \in (0, 1)$, it already

$$\left| \|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2^2 - \sigma^2(4N + 2) \right|^{\frac{1}{2}} + (2 + \sqrt{2})\sigma(2N + 1)^{(1+\delta)/4} \leq c_{\min}(N + 1)^{1/2} \quad (2.34)$$

the following error estimate holds true

$$\begin{aligned} \frac{\pi^2}{3}(N+1)^3 \sum_{y \in Y^f} (|c_y^f|^2 + |c_{n(y)}^g|^2) |y - n(y)|_{\mathbb{T}}^2 + \frac{3}{8}(N+1) \sum_{y \in Y^f} |c_y^f - c_{n(y)}^g|^2 \\ \leq \left(\left| \|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2^2 - \sigma^2(4N + 2) \right|^{\frac{1}{2}} + (2 + \sqrt{2})\sigma(2N + 1)^{(1+\delta)/4} \right)^2, \end{aligned}$$

with probability of at least

$$1 - e^{-(2N+1)^{(1+\delta)/2}/2} - 2e^{-(2N+1)^\delta/8}.$$

Proof. Rearranging (2.31) and the subadditivity of the square root give

$$\begin{aligned} \|\mathcal{P}_N(f) - \mathcal{P}_N(f^*)\|_2 &\leq \left(\|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2^2 - \sigma^2 Y + \sigma^2 Z^2 \right)^{\frac{1}{2}} + \sigma|Z| \\ &\leq \left| \|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2^2 - \sigma^2(4N+2) \right|^{\frac{1}{2}} + 2\sigma|Z| + \sigma|Y - (4N+2)|^{\frac{1}{2}}. \end{aligned}$$

Next we choose in (2.32) and in (2.33)

$$t_1 = (2N+1)^{(1+\delta)/4}, \quad t_2 = \frac{(2N+1)^{\delta/2}}{(4N+2)^{1/2}}.$$

Then, with probability at least

$$1 - e^{-t_1^2/2} - 2e^{-(4N+2)t_2^2/8} = 1 - e^{-(2N+1)^{(1+\varepsilon)/2}/2} - 2e^{-(2N+1)^\varepsilon/8}$$

the estimate

$$\|\mathcal{P}_N(f) - \mathcal{P}_N(f^*)\|_2 \leq \left| \|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2^2 - \sigma^2(4N+2) \right|^{\frac{1}{2}} + (2 + \sqrt{2})\sigma(2N+1)^{(1+\delta)/4}$$

holds true. The claim follows directly from Theorem 2.26. \square

Remarks. 1. If one is only interested in the frequency error, we see that with high probability, we obtain

$$\sum_{y \in Y^f} |y - n(y)|_{\mathbb{T}}^2 = \mathcal{O}(N^{-2}),$$

at least if the condition (2.34) is satisfied. However, understating the result as an asymptotic estimate is a little bit misleading, as only $f = f^*$ satisfies (2.34) for all N .

2. It is possible to carry out similar calculations for different noise models. This is meant to be a prototypical example, how knowledge of the noise gives rise to an a-posteriori error estimate.

Bivariate Results

Now we extend the univariate stability results to the bivariate case. Conceptually, we have little work to do. Indeed, virtually the same proof strategy works in this case as well. We always consider a sampling set of the form

$$G_N = [-N, N]^2 \cap \mathbb{Z}^2.$$

We start with a lower bound for singular values of multivariate Vandermonde matrices.

Definition 2.28. For a finite set $G \subset \mathbb{Z}^d$ and $y \in [0, 1)^d$ we define a Vandermonde vector by

$$v_G(y) = (e^{2\pi i y \cdot n})_{n \in G} \in \mathbb{C}^G.$$

Further, for mutually distinct $y_1, \dots, y_M \subset [0, 1)^d$, we define the Vandermonde matrix

$$V_G(Y) = (e^{2\pi i n \cdot y})_{n \in G, y \in Y} = [v_G(y) : y \in Y] \in \mathbb{C}^{G \times Y}.$$

Proposition 2.29. Let $f \in \mathcal{S}^2(q)$ for $q = K/(N+1)$, $K, N \in \mathbb{N}_{>0}$ be given. Then

$$\sum_{k \in G_N} |f(k)|^2 \geq \begin{cases} (N+1)^2 \left(\left(2 - \frac{1}{K}\right)^2 - \frac{1}{K^2} \right) \|c^f\|_2^2 & \text{if } K \neq 1, \\ \frac{63}{64} (N+1)^2 \|c^f\|_2^2 & \text{if } K = 1. \end{cases}$$

In particular, we get the following estimate for the smallest singular value of $V_{G_N}(Y^f)$:

$$\sigma_{\min}^2 \geq \begin{cases} (N+1)^2 \left(\left(2 - \frac{1}{K}\right)^2 - \frac{1}{K^2} \right) & \text{if } K \neq 1, \\ \frac{63}{64}(N+1)^2 & \text{if } K = 1. \end{cases}$$

Proof. We start with the case $K \neq 1$. Let F_K be defined as in Proposition 2.17 and

$$\Psi(x) = \text{Dil}_q F_K(x) \leq \chi_{[-N-1, N+1]^2}(x).$$

The same calculation as given in Theorem 2.18 results in

$$\sum_{k \in G_N} |f(k)|^2 \geq \hat{\Psi}(0) \|c^f\|_2^2 = (N+1)^2 \left(\left(2 - \frac{1}{K}\right)^2 - \frac{1}{K^2} \right).$$

In the special case $K = 1$, we use \tilde{F} , as given in (2.24), which gives $\hat{\Psi}(0) = (N+1)^2 \frac{63}{64}$. \square

Remarks. 1. Upper bounds can be easily provided by using a tensor product majorant, which results in

$$\sum_{k \in G_N} |f(k)|^2 \leq \left(2N + \frac{1}{q}\right)^2 \|c^f\|_2^2.$$

This was done in [54].

2. The condition $q = K/(N+1)$ is due to the fact that we only constructed minorants of $\chi_{[-N, N]^2}$ with $N \in \mathbb{N}_{>0}$. Using minorants of a more general type removes this restriction. For possible choices, see the discussion after Proposition 2.17.

Now we turn to the lower bounds for $\mathcal{P}_N(f-g)$. All we have to do is extend the local analysis of $\hat{F}_{\frac{3}{4}, 2}$ to the function \hat{F}_2 , as given in Proposition 2.17.

We start by calculating the Fourier transform of F_2 , using Lemma 2.21:

$$\begin{aligned} \hat{F}_2(w_1, w_2) &= \hat{F}_{\frac{3}{4}, 2}(w_1) \hat{F}_{\frac{3}{4}, 2}(w_2) - \frac{1}{4} \hat{G}_2(w_1) \hat{G}_2(w_2) \\ &= \prod_{j=1}^2 \left(\frac{3}{4\pi} \sin(4\pi|w_j|) + (1 - |w_j|)(1 + 2\cos(2\pi w_j)) \right) \\ &\quad - \frac{1}{4} \prod_{j=1}^2 \left(\frac{1}{2\pi} \sin(4\pi|w_j|) + 2(1 - |w_j|) \cos(4\pi w_j) \right) \end{aligned}$$

A direct calculation gives

Lemma 2.30. *If $|w_j| \geq \frac{1}{4}$ for $j = 1$ or $j = 2$, we have that*

$$\hat{F}_2(w_1, w_2) \leq 3.$$

On the other hand, for $0 \leq w_1, w_2 \leq \frac{1}{4}$, we have that

$$\partial_j \hat{F}_2 \leq 0, \quad j = 1, 2,$$

with equality if and only if $w_1 = w_2 = 0$.

Proof. First, note that due to symmetry, we can always assume that $w_1, w_2 \geq 0$. We calculate

$$\begin{aligned} \hat{F}'_{\frac{3}{4}, 2}(w) &= 3\cos(4\pi w) - (1 + 2\cos(2\pi w)) - 4\pi(1-w)\sin(2\pi w) \\ \hat{G}'_2(w) &= 2\cos(4\pi w) - 2\cos(4\pi w) - 8\pi(1-w)\sin(4\pi w) = -8\pi(1-w)\sin(4\pi w) \end{aligned}$$

for $w \geq 0$. We start with the first term. Clearly, \hat{G}_2 has critical points exactly at $n/4$, $n = 0, \dots, 4$. We can easily check that $[\hat{G}_2]$ has its global maximum on $[0, 1]$ at zero and on $[\frac{1}{4}, 1]$ at $\frac{1}{4}$. Furthermore,

we already established that $\hat{F}_{\frac{3}{4},2}$ is decreasing on $[0, \frac{1}{4}]$, see the proof of Theorem 2.26. It is routine (though slightly annoying) to check that the maximum of $|\hat{F}_{\frac{3}{4},2}|$ on $[\frac{1}{4}, 1]$ is at $\frac{1}{4}$. This results in

$$\hat{F}_2(w_1, w_2) \leq \hat{F}_{\frac{3}{4},2}(0)\hat{F}_{\frac{3}{4},2}\left(\frac{1}{4}\right) + \frac{1}{4}\hat{G}_2(0)\hat{G}_2\left(\frac{1}{4}\right) = 3 \quad \text{for all } (w_1, w_2) \in \left[0, \frac{1}{4}\right] \times [0, 1].$$

The other cases follow by symmetry.

To prove the second claim, it suffices to check that

$$\partial_1 \hat{F}_2(w_1, w_2) = \hat{F}'_{\frac{3}{4},2}(w_1)\hat{F}_{\frac{3}{4},2}(w_2) - \frac{1}{4}\hat{G}'_2(w_1)\hat{G}_2(w_2) \leq 0,$$

where equality holds if and only if $w_1 = 0$. The case $w_1 = 0$ is clear, we therefore assume that $w_1 > 0$ and the claim is equivalent to

$$\frac{\hat{F}'_{\frac{3}{4},2}(w_1)}{\hat{G}'_2(w_1)} \geq \frac{1}{4} \frac{\hat{G}_2(w_2)}{\hat{F}_{\frac{3}{4},2}(w_2)}.$$

First, we prove that the right-hand side is smaller than $\frac{1}{4}$ for all $w_2 \in (0, \frac{1}{4})$, which is equivalent to $\hat{F}_{\frac{3}{4},2} > \hat{G}_2$, which again is clearly true, as

$$\frac{1}{4\pi} \sin(4\pi w) + (1-w)(1+2\cos(2\pi w) - 2\cos(4\pi w)) > 0.$$

Next, we prove that the left-hand side is at least $\frac{1}{4}$. This is equivalent to

$$4\hat{F}'_{\frac{3}{4},2}(w) \leq \hat{G}'_2(w) \quad \Leftrightarrow$$

$$0 \leq 8\pi(1-w)(2\sin(2\pi w) - \sin(4\pi w)) + 4(1+2\cos(2\pi w) - 3\cos(4\pi w)),$$

which is again true (as $\sin(4\pi w) = 2\sin(2\pi w)\cos(2\pi w)$). \square

Now we are ready to prove the two dimensional analog to Theorem 2.26.

Theorem 2.31. *Let $f, g \in \mathcal{S}^2(2q)$. Further, assume that $q \geq \frac{2}{N+1}$ and*

$$\|\mathcal{P}_N(f - g)\|_2^2 < \frac{5}{4}(N+1)^2 c_{\min}^2, \quad (2.35)$$

where c_{\min} is a lower bound on the modulus of the coefficients of f and g . Then for every $y \in Y^f$ we find exactly one $y' = n(y) \in Y^g$ with $\|y - y'\|_{\mathbb{T}^2} < \frac{1}{2(N+1)}$ and vice versa.

Furthermore, the following estimate holds true:

$$\frac{15}{16}(N+1)^4 \sum_{y \in Y^f} (|c_y^f|^2 + |c_{n(y)}^g|^2) \|y - n(y)\|_{\mathbb{T}^2}^2 + \frac{3(N+1)^2}{4} \sum_{y \in Y^f} |c_y^f - c_{n(y)}^g|^2 \leq \|\mathcal{P}_N(f - g)\|_2^2$$

Proof. We use

$$\Psi(x) = \text{Dil}_{\frac{2}{N+1}} F_2(x) \leq \chi_{[-N, N]^2}(x).$$

The Fourier transform of Ψ is given by $\frac{(N+1)^2}{4} \hat{F}_2((N+1)w/2)$. Exactly as we derived (2.26), we get,

using that by Lemma 2.30 the function $|\hat{\Psi}|$ has its global maximum in zero,

$$\begin{aligned}
\sum_{k \in G_N} |f(k) - g(k)|^2 &\geq \sum_{k \in \mathbb{Z}^2} |f(k) - g(k)|^2 \Psi(k) \\
&\geq \hat{\Psi}(0) \left(\sum_{y \in Y^f} |c_y^f|^2 + \sum_{y' \in Y^g} |c_{y'}^g|^2 \right) + 2\Re \left(\sum_{y \in N(Y^f)} c_y^f \overline{c_{n(y)}^g} \hat{\Psi}(y - n(y)) \right) \\
&\geq \hat{\Psi}(0) \left(\sum_{y \in Y^f} |c_y^f|^2 + \sum_{y' \in Y^g} |c_{y'}^g|^2 \right) - 2 \sum_{y \in N(Y^f)} |c_y^f| |c_{n(y)}^g| |\hat{\Psi}(y - n(y))| \\
&\geq \hat{\Psi}(0) \left(\sum_{y \in Y^f \setminus N(Y^f)} |c_y^f|^2 + \sum_{\substack{y' \in Y^g \\ \nexists n^{-1}(y')}} |c_{y'}^g|^2 \right).
\end{aligned}$$

Now, as $\hat{\Psi}(0) = 2(N+1)^2$, by (2.35), we see that $Y^f \setminus N(Y^f) = \emptyset$, i.e., we find for each $y \in Y^f$ exactly one $n(y) \in Y^g$ with $\|y - n(y)\|_{\mathbb{T}^2} < 2/(N+1)$.

Now assume we find a pair $y \in Y^f$ and $n(y) \in Y^g$ with $\|y - n(y)\|_{\mathbb{T}^2} \geq \frac{1}{2(N+1)}$. By Lemma 2.30, this implies that

$$\hat{\Psi}(0) - \hat{\Psi}(y - n(y)) = \frac{(N+1)^2}{4} \left(\hat{F}_2(0) - \hat{F}_2((N+1)(y - n(y))/2) \right) \geq \frac{(N+1)^2}{4} (8-3).$$

Again, (2.35) yields a contradiction, which proves the first claim.

Next, we analyze the behavior of $\hat{F}_2(0) - \hat{F}_2(w)$ for $w \in [0, \frac{1}{4}]^2$ (by symmetry, we can always assume that $y - n(y) \in [0, \frac{1}{4}]^2$). By the sign pattern of $\nabla \hat{F}_2$, proved in Lemma 2.30, we see that

$$\hat{F}_2(w) \leq \hat{F}_2(\|w\|_\infty, 0).$$

This enables us to give the following estimate, using Lemma 2.22:

$$\begin{aligned}
\hat{F}_2(0) - \hat{F}_2(w) &\geq 8 - \hat{F}_2(\|w\|_\infty, 0) \geq 8 - \hat{F}_{\frac{3}{4}, 2}(\|w\|_\infty) \hat{F}_{\frac{3}{4}, 2}(0) + \frac{1}{4} \hat{G}_2(\|w\|_\infty) \hat{G}_2(0) \\
&\geq 8 - 3\hat{F}_{\frac{3}{4}, 2}(\|w\|_\infty) + \frac{1}{2} \hat{G}_2(\|w\|_\infty) \\
&= 8 - \frac{2}{\pi} \sin(4\pi\|w\|_\infty) - 3(1 - \|w\|_\infty)(1 + 2\cos(2\pi\|w\|_\infty)) + (1 - \|w\|_\infty) \cos(4\pi\|w\|_\infty) \\
&\geq 8 - \frac{2}{\pi} \left(4\pi\|w\|_\infty - 4^3\pi^3\|w\|_\infty^3 \left(\frac{1}{6} - \frac{\pi^2}{120} \right) \right) - 3(1 - \|w\|_\infty)(1 + 2 - 8\pi^2\|w\|_\infty^2 \left(\frac{1}{2} - \frac{\pi^2}{96} \right)) \\
&\quad + (1 - \|w\|_\infty)(1 - 8\pi^2\|w\|_\infty^2) \\
&= \left(24\pi^2 \left(\frac{1}{2} - \frac{\pi^2}{96} \right) - 8\pi^2 \right) \|w\|_\infty^2 + \left(2^7\pi^2 \left(\frac{1}{6} - \frac{\pi^2}{120} \right) + 8\pi^2 - 24\pi^2 \left(\frac{1}{2} - \frac{\pi^2}{96} \right) \right) \|w\|_\infty^3 \\
&> 15\|w\|_\infty^2.
\end{aligned}$$

Hence,

$$\hat{\Psi}(0) - \hat{\Psi}(w) \geq \frac{(N+1)^2}{4} 15(N+1)^2 \|w\|_\infty^2 / 4.$$

The same technique we already applied in the univariate case then give

$$\begin{aligned}
\sum_{k \in G_N} |f(k) - g(k)|^2 &\geq \sum_{y \in Y^f} (|c_y^f|^2 + |c_{n(y)}^g|^2) (\hat{\Psi}(0) - \hat{\Psi}(y - n(y))) + |c_y^f - c_{n(y)}^g|^2 \hat{\Psi}(y - n(y)) \\
&\geq \frac{15}{16} (N+1)^4 \sum_{y \in Y^f} (|c_y^f|^2 + |c_{n(y)}^g|^2) \|y - n(y)\|_{\mathbb{T}^2}^2 + |c_y^f - c_{n(y)}^g|^2 \hat{\Psi}(y - n(y))
\end{aligned}$$

We apply Lemma 2.30 one last time to get

$$\begin{aligned}\hat{\Psi}(y - n(y)) &\geq \hat{\Psi}(\|y - n(y)\|_{\mathbb{T}^2}, \|y - n(y)\|_{\mathbb{T}^2}) \geq \hat{\Psi}\left(\frac{1}{2(N+1)}, \frac{1}{2(N+1)}\right) \\ &= \frac{(N+1)^2}{4} \hat{F}_2(1/4, 1/4) = \frac{3(N+1)^2}{4}.\end{aligned}$$

□

Finally, we give an a-posteriori error estimate for complex Gaussian noise. This is the bivariate extension of Corollary 2.27.

Corollary 2.32. *Let $f, f^* \in \mathcal{S}^2(2q)$ be fixed with $q \geq \frac{2}{N+1}$. Further, assume that we know f and*

$$\tilde{s}_N(f^*) = \mathcal{P}_N(f^*) + \varepsilon \in \mathbb{C}^{(2N+1)^2},$$

where $\varepsilon_j = X_{j,1} + iX_{j,2}$ and $X_{j,l}$ are pairwise independent, normally distributed random variables with mean zero and variance σ^2 . Then for any $\delta \in (0, 1)$, it already

$$\|\|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2^2 - 2\sigma^2(2N+1)^2\|^{1/2} + (2 + \sqrt{2})\sigma(2N+1)^{(2+\delta)/4} \leq \frac{\sqrt{5}}{2} c_{\min}(N+1) \quad (2.36)$$

the following error estimate holds true

$$\begin{aligned}\frac{15}{16}(N+1)^4 \sum_{y \in Y^f} (|c_y^f|^2 + |c_{n(y)}^g|^2) \|y - n(y)\|_{\mathbb{T}^2}^2 + \frac{3(N+1)^2}{4} \sum_{y \in Y^f} |c_y^f - c_{n(y)}^g|^2 \\ \leq \left(\|\|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2^2 - 2\sigma^2(2N+1)^2\|^{1/2} + (2 + \sqrt{2})\sigma(2N+1)^{(2+\delta)/4} \right)^2,\end{aligned}$$

with probability at least

$$1 - e^{-(2N+1)^{(2+\delta)/2}/2} - 2e^{-(2N+1)^\delta/8}.$$

Proof. Again,

$$\|\|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2^2 = \|\mathcal{P}_N(f) - \mathcal{P}_N(f^*)\|_2^2 + \sigma^2 Y - 2\sigma \|\mathcal{P}_N(f - f^*)\|_2 Z,$$

with $Z \sim \mathcal{N}(0, 1)$ and $Y \sim \chi_{2(2N+1)^2}^2$. The proof is completely analogous to Corollary 2.27. □

With that we conclude this chapter. Explicit examples are given at the end of the next chapter, after we introduced several strategies to actually estimate frequencies of exponential sums. It would have been interesting to get well-posedness results beyond the two dimensional case. And while the strategy presented here works independent of the dimension, suitable localizing functions are currently unknown.

But there is another issue, which makes these results less interesting. Namely, with increasing dimension, well-separated exponential sums become worse and worse proxies to sparse exponential sums. The class $\mathcal{S}^d(q)$ contains exponential sums f up to an order of q^{-d} and (given suitable localizing functions), we could hope for stability if we have samples on $G_{q^{-1}}^d$, i.e., $\mathcal{O}(q^{-d})$ samples.

On the other hand, to recover an exponential sum of order at most M , efficient algorithms for sampling sets of order $\mathcal{O}_d(M^2)$ (up to logarithmic factors) exist. If $d > 2$, these sets do not contain subsets G_N^d of any reasonable size N . It would be more interesting, to establish stability for such sets. However, there are currently no techniques available to prove such estimates.

Chapter 3

Parameter Estimation Techniques

In this chapter we discuss various techniques to solve the parameter estimation problem for exponential sums. We reiterate some of the one dimensional methods and identify their essential ideas, which are then extended to the multivariate case. This inspires a multivariate Prony method. We explain the connection to Sauer's version [84, 85] and prove that d dimensional MUSIC and ESPRIT-type methods only need $\mathcal{O}_d(M^2)$ instead of $\mathcal{O}_d(M^d)$ samples (up to logarithmic factors) and are therefore on a par with Prony's method. Furthermore, we explain and expand a class of so-called projection-based methods, which cleverly combine multiple one dimensional problems to obtain a result for the multivariate case.

Again, we denote by

$$f(x) = \sum_{j=1}^M c_j e^{2\pi i x \cdot y_j}$$

our unknown exponential function we wish to identify from given samples. The coefficients c_j are in $\mathbb{C} \setminus \{0\}$, the frequency vectors $y_j \in [0, 1)^d$ are assumed to be mutually distinct. Due to the linear independence of the exponential functions, M is uniquely determined and called the order of f . We collect the frequencies of f in Y^f and use the notation

$$f(x) = \sum_{y \in Y^f} c_y e^{2\pi i x \cdot y}$$

which does not fix any enumeration of Y^f . The set of all such f is denoted by \mathcal{S}^d , all exponential sums of order at most M are collected in \mathcal{S}_M^d .

As we wish to obtain spectral information of f from the given samples, this is an instance of a spectral estimation problem. Furthermore, we have an explicit model of f , the harmonic model. We only consider methods, which use this model explicitly, i.e., we choose a so-called model based or parametric approach. Of course, general spectral estimation methods are of great importance and can be applied here as well, but as they do not exploit the explicit model, their performance is inferior to model based methods. In fact, by using an explicit model, we can overcome resolution limits of general methods. This is the reason why these methods are able to achieve super-resolution. For an introduction to general spectral estimation techniques, we refer to the books [61, 93].

Note that we do not rely on additional assumptions on f . In many applications, for example in array processing, the coefficients are assumed to be of the form

$$c_j = |c_j| e^{2\pi i \phi_j},$$

where ϕ_j are uncorrelated, uniformly in $[0, 1)$ distributed random variables. This models the assumption that the sources emitting the different wavefronts are uncorrelated. Again, we refer to [61] for an introduction in this application. However, when applied to certain imaging techniques, a reasonable assumption is that $c_j \in \mathbb{R}_{>0}$, as suggested in [65]. Then ϕ_j are not uncorrelated but perfectly correlated. To obtain results covering all cases, we make no further assumptions on the coefficients.

Another assumption sometimes considered is genericity. A precise definition of the notion of genericity we use here is the following.

Definition 3.1. We say that a property P of an $f \in \mathcal{S}^d$ holds generically, if for all $M \in \mathbb{N}$ choosing M points (y_1, \dots, y_M) randomly in $[0, 1]^d$ with uniform probability, the probability that there is a $f \in \mathcal{S}^d$ with frequencies $Y^f = \{y_1, \dots, y_M\}$ such that P does not hold for f is zero.

Of course, different notations of genericity are possible, e.g. one could choose the coefficients randomly as well or even only them. But in the following discussion, it becomes clear that the definition used here is reasonable. This notion is used in [48] as well. Though less focus of this work, we comment sporadically on the generic case.

There are a few general problems with properties that hold generically. The first is that it restricts our model which might cause problems in some applications. For example if the frequencies of the unknown f are known to have some geometric structure, e.g. are on a grid, the generic model is meaningless.

More problematic is that while the probability that a generic f does not satisfy P is zero, the probability that a \tilde{f} very close to f does not satisfy P is (generally) not. In particular, no numerical stability can be expected, when the reconstruction scheme relies on a generically true property.

Finally, we are of course interested in the effect of noise. However, there are a lot of different noise models (like Gaussian or Poisson noise) which are used, depending on the situation at hand. The aim of this work is not to discuss all these situations. Therefore, we (mostly) stick with a very simple noise model, namely, that we have a bound on the ℓ^2 norm of the noise vector. This has the advantage that it is very general and serves as a prototype of such estimates. However, it is important to keep in mind that all methods presented here might need a tweak to fully take advantage of additional knowledge about the noise.

This chapter is organized as follows. We start with a review of one dimensional techniques, with a focus on subspace methods (like Prony, MUSIC and ESPRIT). While this section is classical, we change the perspective a little bit by starting with the signal space and then presenting the different techniques.

Following this, we consider the projection-based results, first on scattered lines, then on parallel lines. We derive several small sampling sets, cumulating in a sampling set of size $\mathcal{O}_d(M)$ (up to logarithmic factors). However, using them results in an algorithm which has exponential complexity. Furthermore, we propose an ESPRIT-type algorithm to stabilize computations in case of parallel lines.

Finally, we discuss multivariate extension of classical algorithms. As already announced, we use our framework, developed in the one dimensional case, to give very natural extensions of Prony's method, MUSIC and ESPRIT to the higher dimensional case. While these are (in one form or another) already known in the literature, we hope that our unified approach clarifies similarities and differences.

Furthermore, we are able to prove that multivariate MUSIC and ESPRIT-type methods only need $\mathcal{O}_d(M^2)$ samples (up to logarithmic factors) and are still able to recover all $f \in \mathcal{S}_M^d$. Previous methods of that type use $\mathcal{O}_d(M^d)$ samples. Such sampling sets were first introduced by Sauer in [85], who proposed a Prony-type method relying on them. Our ESPRIT method has the additional advantage of having a computational complexity of $\mathcal{O}_d(M^3)$ (again up to logarithmic factors), a clearly improvement over previous ESPRIT methods, which have a complexity of $\mathcal{O}_d(M^{3d})$.

3.1 Review of Univariate Methods

For the reader's convenience, we give a quick recap of one dimensional methods to estimate the frequency of an unknown exponential sum

$$f(x) = \sum_{j=1}^M c_j e^{2\pi i y_j x},$$

where, as usual, $y_j \in [0, 1)$ are mutually distinct and $c_j \in \mathbb{C} \setminus \{0\}$. This section is completely classical and a reader experienced with Prony and ESPRIT-type methods is invited to skim through or skip it.

We remark that the restriction $y_j \in [0, 1)$ is necessary, if we sample at $G \subset \mathbb{Z}$, due to the periodicity of the exponential function. Of course, if $y_j \in [0, \alpha)$, we can rescale the sampling set $G \subset \frac{1}{\alpha}\mathbb{Z}$.

Reflecting the common appearance of exponential sums, there is a large number of methods available. We give a non exhausted list: Prony's method [23] dating back to 1795, Pisarenko's method [70] and its generalization, MUSIC [87], ESPRIT [83, 82] and the related matrix pencil method [45],

OPUC, which is based on orthogonal polynomials on the unit circle [31] and methods using TV minimization (which we will not discuss here), as in [16, 28, 88].

More methods, including non parametric approaches, can be found in [93].

Further, estimating frequencies of exponential sums is closely related to Padé approximation and the annihilation filter method, see [72]. The connection between Prony's approach and Sylvester's method for solving the Waring problem for binary forms is explained in [96]. Also, the connection to tensor decomposition is discussed there.

This section is inspired not only by the original references listed above, but also by the book [61], which gives a more detailed introduction to MUSIC and ESPRIT and by [76], where the connection between Prony's original method, ESPRIT and matrix pencil method is clarified. But our approach here differs from these references in multiple ways. Most importantly, we start by defining the signal space and try to establish the methods coming from there.

Definition 3.2. Let $f \in \mathcal{S}_M^1$ and a window size $N \in \mathbb{N}^*$ be given. We call

$$\text{Sig}(f, N) := \text{span} \left\{ (f(k), \dots, f(k+N-1))^T : k \in \mathbb{Z} \right\} \subset \mathbb{C}^N$$

the signal space of f and define $s_N(j)$ to be one batch of samples

$$s_N(j) = (f(j), \dots, f(j+N-1))^T.$$

Further, when $N \geq M$ let $T : \mathbb{C}^N \rightarrow \mathbb{C}^N$ be the linear map, which shifts the window by one. More precisely, T is defined by

$$Ts_N(j) = s_N(j+1) \quad \text{for all } j \in \mathbb{Z} \quad \text{on } \text{Sig}(f, N)$$

and extended by zero on $\text{Sig}(f, N)^\perp$, the orthogonal complement of the signal space.

The idea is that $\text{Sig}(f, N)$ is in fact a lower dimensional subspace of \mathbb{C}^N , which we can identify from the given samples and that T carries information on f . Next, we show that T is well-defined and establish the precise relation between T and f . Recall that we defined for $y \in [0, 1)$ and $N \in \mathbb{N}_{>0}$

$$v_N(y) = \left(1, e^{2\pi iy}, \dots, e^{2\pi iy(N-1)} \right)^T \in \mathbb{C}^N,$$

and for $y_1, \dots, y_M \in [0, 1)$

$$V_N(y_1, \dots, y_M) = [v_N(y_1) \cdots v_N(y_M)] \in \mathbb{C}^{N \times M}.$$

Lemma 3.3. Let $f \in \mathcal{S}_M^1$ be given with $Y^f = \{y_1, \dots, y_M\}$, coefficients $c = (c_1, \dots, c_M)^T$ and order M . Further, let $j \in \mathbb{Z}$ and $N \in \mathbb{N}$ be arbitrarily chosen. Then the following statements hold true:

(1) Let D be the diagonal matrix $\text{diag}(\exp(2\pi iy_1), \dots, \exp(2\pi iy_M))$. Then

$$s_N(j) = V_N D^j c.$$

(2) If $N \geq M$, a basis of $\text{Sig}(f, N)$ is given by $(s_N(j), \dots, s_N(j+M-1))$. In particular,

$$\dim \text{Sig}(f, N) = M = \text{ord } f$$

and $Ts_N(j) := s_N(j+1)$, $j = 0, \dots, M-1$ is well-defined.

(3) If $N \geq M$, the vectors $(v_N(y_j))_{j=1, \dots, M}$ form a basis \mathcal{V} of $\text{Sig}(f, N)$.

(4) If $N \geq M$, T is diagonalized by \mathcal{V} and $v_N(y_j)$ is an eigenvector of T with eigenvalue $e^{2\pi iy_j}$. Further, $Ts_N(j) = s_N(j+1)$ holds for all $j \in \mathbb{Z}$.

Proof. (1) This is a simple calculation, as

$$s_N(j) = \sum_{k=1}^M c_k e^{2\pi iy_k j} v_N(y_k).$$

(2) By (1), $\text{Sig}(f, N)$ is contained in the span of $v_N(y_j)$, $j = 1, \dots, M$ and $\dim \text{Sig}(f, N) \leq M$. It remains to show that $s_N(j), \dots, s_N(j + M - 1)$ are linearly independent. Without loss of generality, let $N = M$. Assume $\alpha_1, \dots, \alpha_M \in \mathbb{C}$ are given such that

$$0 = \sum_{k=1}^M \alpha_k s_N(k + j - 1) = V_N D^j \sum_{k=1}^M \alpha_k D^{k-1} c.$$

Now $V_N = V_M$ is invertible and $c_j \neq 0$, thus

$$\sum_{k=1}^M \alpha_k e^{2\pi i y_j (k-1)} = 0 \quad \text{for } j = 1, \dots, M.$$

But this is equivalent to $(\alpha_1, \dots, \alpha_M) \perp v_N(y_j)$, which then implies that $\alpha_j = 0$ and we are done.

(3) As $\dim \text{Sig}(f, N) = M$, they form a basis.

(4) Note that

$$\sum_{k=1}^M c_k e^{2\pi i y_k j} e^{2\pi i y_k} v_N(y_k) = s_N(j + 1) = T s_N(j) = \sum_{k=1}^M c_k e^{2\pi i y_k j} T v_N(y_k).$$

As \mathcal{V} and $(s(j))_{j=1, \dots, M}$ are bases of $\text{Sig}(f, N)$, the claim $T v_N(y_k) = e^{2\pi i y_k} v_N(y_k)$ follows. Then we have that

$$T V_N = V_N D$$

which establishes the last claim. \square

There is an interesting perspective on f and T , connecting them to linear difference equations. In fact, the matrix representing T (for $N = M$) is given by

$$\begin{pmatrix} 0 & 0 & \dots & 0 & -p_0 \\ 1 & 0 & \dots & 0 & -p_1 \\ 0 & 1 & \dots & 0 & -p_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & -p_{M-1} \end{pmatrix}^T,$$

where $p_j \in \mathbb{C}$, $j = 0, \dots, M - 1$. This is of course the transposed of the companion matrix to the polynomial $P(z) = z^M + p_{M-1}z^{M-1} + \dots + p_0$. Further, f satisfies the linear difference equation with constant coefficients

$$\sum_{k=0}^{M-1} -p_k f(j+k) = f(j+M) \quad \text{for all } j \in \mathbb{Z}. \quad (3.1)$$

As we just calculated the eigenvalues of T , we established

$$P(z) = z^M + p_{M-1}z^{M-1} + \dots + p_0 = \prod_{k=1}^M (z - e^{2\pi i y_k}). \quad (3.2)$$

This polynomial is called the Prony polynomial. For $N \geq M$, the matrix becomes the transpose of the companion matrix of $z^{N-M}P(z)$, as we extended T by zero on $\text{Sig}(f, N)^\perp$.

From this point of view, we observe a solution to a linear difference equation and want to learn which linear difference equation this is. Furthermore, the Prony polynomial only has simple roots, a restriction which makes sense in our model, but from the perspective of linear finite difference equation is rather restrictive. Without this restriction, one has to consider the more general model, where coefficients c_j are allowed to be polynomials. This model has also found multiple applications. For further information on this, we refer to the literature, for example [6, 40].

These insights already give rise to Prony's method. To state (3.1) concisely in matrix form, we define Hankel matrices.

Definition 3.4. Let $f \in \mathcal{S}_M^1$ be given. For $N, L \in \mathbb{N}_{\geq 0}$ and $j \in \mathbb{Z}$ we define the Hankel matrix

$$\begin{aligned} H_{N,L}^f(j) &= H_{N,L}(j) = [s_N(j) \quad s_N(j+1) \quad \dots \quad s_N(j+L-1)] \\ &= \begin{pmatrix} f(j) & f(j+1) & \dots & f(j+L-1) \\ f(j+1) & f(j+2) & \dots & f(j+L) \\ \vdots & \vdots & \ddots & \vdots \\ f(j+N-1) & f(j+N) & \dots & f(j+L+N-2) \end{pmatrix} \in \mathbb{C}^{N \times L}. \end{aligned}$$

Note that by Lemma 3.3 (2), if $N, L \geq M$ this matrix has rank M . Further, we use the notation $H_{N,L}^f = H_{N,L}^f(0)$.

Theorem 3.5 (Prony's Method (Prony, 1795)). Let $f \in \mathcal{S}_M^1$ be an exponential sum and $j \in \mathbb{Z}$. Assume that we know $f(j), \dots, f(j+2M-1)$ (and hence M). The following method recovers frequencies and coefficients of f :

1. Calculate the solution $p \in \mathbb{C}^M$ of

$$H_{M,M}(j)x = -s_M(j+M). \quad (3.3)$$

2. Form the Prony polynomial (3.2) and solve for its roots. Call them z_1, \dots, z_M . Determine $y_k \in [0, 1)$ with $z_k = e^{2\pi i y_k}$.

3. Solve the (overdetermined) linear system

$$V_{2M}D^j c = s_{2M}(j), \quad (3.4)$$

where $D = \text{diag}(\exp(2\pi i y_1), \dots, \exp(2\pi i y_M))$.

Proof. Due to Lemma 3.3 (2), the matrix $H_{M,M}$ has full rank and hence the solution to (3.3) is unique. By the preceding discussion, it is clear that the roots of P are exactly $e^{2\pi i y_j}$. Further, the linear system (3.4) has a unique solution, which is the vector of coefficients - this follows from Lemma 3.3 (1). \square

Usually, the order of f is not known, only an upper bound $N \geq M$ and samples $(f(j+k))_{k=0, \dots, 2N-1}$ are given. Still, Prony's method can be applied. To obtain M , one has to estimate the rank of $H_{N,N}$. Unfortunately, this is rather unstable, we discuss this issue in more detail below. Then one can proceed as sketched above. Of course, it is advisable to replace (3.3) and (3.4) by

$$\begin{aligned} H_{2N-M,M}x &= -s_{2N-M}(j+M) \\ V_{2N}D^j c &= s_{2N}(j) \end{aligned}$$

to use the additional samples to stabilize the scheme.

One of the most important challenges in signal processing is how to address noise. Parameter estimation of exponential sum is no exception. We use the notation

$$\tilde{f}(j) = f(j) + \varepsilon_j,$$

where ε_j is some small perturbation. Analogously, we let $\tilde{s}_N(j) = (\tilde{f}(j), \dots, \tilde{f}(j+N-1))$ and $\tilde{H}_{N,L}(j) = [\tilde{s}_N(j) \dots \tilde{s}_N(j+L-1)]$.

Now estimating $\text{ord } f$ by estimating the rank of $\tilde{H}_{N,L}(j)$ has to be done with care, as $\tilde{H}_{N,L}(j)$ usually has full rank. The most used method is performing a singular value decomposition of $\tilde{H}_{N,L}(j)$ (where we drop for ease of notation the dependence on j):

$$H_{N,L}(j) = U\Sigma W^H \quad \tilde{H}_{N,L}(j) = \tilde{U}\tilde{\Sigma}\tilde{W}^H \quad (3.5)$$

Here $W, \tilde{W} \in \mathbb{C}^{L \times L}$ and $U, \tilde{U} \in \mathbb{C}^{N \times N}$ are unitary matrices and $\Sigma, \tilde{\Sigma} \in \mathbb{C}^{L \times N}$ are diagonal matrices with diagonal entries $\sigma_1 \geq \dots \geq \sigma_{\min\{L,N\}}$ resp. $\tilde{\sigma}_1 \geq \dots \geq \tilde{\sigma}_{\min\{L,N\}}$. We assume that $L, N \geq M$. Then

$$\sigma_1 \geq \dots \geq \sigma_M > \sigma_{M+1} = 0 = \dots = \sigma_{\min\{L,N\}}$$

and using standard perturbation results for singular values (details are given below), one expects a gap in the singular values of $\tilde{H}_{N,L}(j)$:

$$\tilde{\sigma}_1 \geq \cdots \geq \tilde{\sigma}_M \gg \tilde{\sigma}_{M+1} \geq \tilde{\sigma}_{\min\{L,N\}}.$$

To give a quantitative estimate for this spectral gap, we need quantitative estimates on σ_M (and assumptions on ε_j , of course). This can be achieved by a factorization of $H_{N,L}(j)$ and estimates on the singular values of Vandermonde matrices.

Lemma 3.6. *For $f \in \mathcal{S}_M^1$ with frequencies y_1, \dots, y_M and coefficients c_1, \dots, c_M , we have that*

$$H_{N,L}^f(j) = V_L(y_1, \dots, y_M) \text{diag} (c_1 e^{2\pi i j y_1}, \dots, c_M e^{2\pi i j y_M}) V_N(y_1, \dots, y_M)^T.$$

Proof. This well-known factorization (see for example [45]) is derived by a direct calculation. Using the notation from Lemma 3.3, we obtain

$$\begin{aligned} H_{N,L}^f(j) &= [s_N(j) \quad \cdots \quad s_N(j+L-1)] = [V_N D^j c \quad \cdots \quad V_N D^{j+L-1} c] \\ &= V_N [D^j c \quad \cdots \quad D^{j+L-1} c] = V_N \text{diag} (c_1 e^{2\pi i j y_1}, \dots, c_M e^{2\pi i j y_M}) V_L^T. \end{aligned}$$

□

We now present estimates on the singular values of $H_{N,L}^f$. A version of this result is given in [78], but instead of relying on estimates for singular values of Vandermonde matrices presented in [55], we rely on the improved estimates given in Corollary 2.20.

Theorem 3.7. *Let $f \in \mathcal{S}_M^1(q)$ be given, where $q \in (0, 1)$. Then we obtain the following bounds for the smallest non-zero singular value σ_{\min} and the largest singular value σ_{\max} of $H_{N,L}^f$:*

$$\begin{aligned} \sigma_{\min}^2 &\geq c_{\min}^2 \left(N + 1 - \frac{1}{q} \right) \left(L + 1 - \frac{1}{q} \right) \\ \sigma_{\max}^2 &\leq c_{\max}^2 \left(N - 1 + \frac{1}{q} \right) \left(L - 1 + \frac{1}{q} \right), \end{aligned}$$

where c_{\min} and c_{\max} are the modulus of the smallest and largest coefficient of f , respectively.

Proof. This is a direct consequence of the factorization of $H_{N,L}^f$, as we just proved in Lemma 3.6 and the estimates of singular values of Vandermonde matrices, stated in Corollary 2.20. □

This can be used to give some coarse estimates under which the correct order of f can be estimated. Indeed, by perturbation results for singular values [33], Corollary 8.6.2, we have that

$$|\sigma_k(\tilde{H}_{N,L}^f) - \sigma_k(H_{N,L}^f)| \leq \sigma_1(H_{N,L}^f - \tilde{H}_{N,L}^f) = \|H_{N,L}^f - \tilde{H}_{N,L}^f\|_2,$$

if $N \geq L$. Let $E = H_{N,L}^f - \tilde{H}_{N,L}^f$ be the matrix containing the noise. Thus, if

$$\sigma_M(H_{N,L}^f) \geq 2\|E\|_2, \tag{3.6}$$

we can choose $\text{tol} = \|E\|_2$ and are guaranteed to recover $\text{ord } f = M$.

If we only know $|\varepsilon_j| \leq \eta$, we have the bound

$$\|E\|_2 \leq \sqrt{NL}\eta.$$

(3.6) holds true if

$$c_{\min}^2 \left(1 + \frac{1}{N} - \frac{1}{Nq} \right) \left(1 + \frac{1}{L} - \frac{1}{Lq} \right) \geq 4\eta^2.$$

Of course, we usually assume some statistical knowledge on ε_j . Note that E is a structured random matrix. Understanding spectral properties of structured random matrices has become an active research topic in the last years. However, non-asymptotic bounds seem to be not known, even in the

simplest case (e.g. if ε_j are i.i.d. normally distributed random variables). For asymptotic results, see [13].

Next, we give a quick recap of (a variation of) the Multiple Signal Classification (short: MUSIC) algorithm by Schmidt [87]. It is based on the fact that for $N \geq M + 1$ we have that

$$\text{Sig}(f, N) = \text{span}\{v_N(y_j) : j = 1, \dots, M\}$$

and $v_N(y) \notin \text{Sig}(f, N)$ for all $y \in [0, 1) \setminus Y^f$. Thus, finding the frequencies of f is equivalent to finding all $y \in [0, 1)$ with $v_N(y) \in \text{Sig}(f, N)$. Now assume we are given noisy samples, enough to build $\tilde{s}_N(j)$, $j = 0, \dots, L - 1$, where $N \geq M + 1, L \geq M$. Further assume that we know M (or have determined it by finding the gap in the singular values of $\tilde{H}_{N,L}^f = \tilde{H}_{N,L}^f(0)$). We could now try to formalize and solve the following problem:

$$\begin{aligned} & \text{Find } \tilde{y}_1, \dots, \tilde{y}_M \text{ such that } \text{span}\{v_N(\tilde{y}_j) : j = 1, \dots, M\} \\ & \text{is as close as possible to } \text{span}\{\tilde{s}_N(j) : j = 0, \dots, L - 1\}. \end{aligned}$$

And while it is possible to pursue this idea (see for example the discussion on the maximum likelihood estimator in [82]), it leads to a difficult, nonlinear optimization problem. The idea of MUSIC is the following simplification:

1. Calculate an estimation $\tilde{\text{Sig}}(f, N)$ of the signal space.
2. Find the M frequencies $\tilde{y}_j \in [0, 1)$ for which $v_N(\tilde{y}_j)$ are closest to $\tilde{\text{Sig}}(f, N)$.

For the first step, usually a SVD is used. Indeed, $\text{span}\{\tilde{s}_N(j) : j = 0, \dots, L - 1\}$ is equal to the range of $\tilde{H}_{L,N}^f$ and performing a SVD like in (3.5), we let

$$\tilde{\text{Sig}}(f, N) = \text{span}\{\tilde{u}_1, \dots, \tilde{u}_M\} \quad (3.7)$$

where \tilde{u}_j is the j th column vector of \tilde{U} . Note that in the noise free case, $\tilde{\text{Sig}}(f, N) = \text{Sig}(f, N)$.

Remark. As $(H_{L,N}^f)^T = H_{N,L}^f$, we could use the first M right singular vectors as well. There is no difference (as long as one adjusts L and N accordingly), though one additional conjugation is then necessary.

Now we consider the second step. MUSIC uses the following closeness indicator:

$$R(y) = \frac{1}{\sum_{k=M+1}^N |\tilde{u}_k^H v_N(y)|^2}. \quad (3.8)$$

$\sum_{k=M+1}^N |\tilde{u}_k^H v_N(y)|^2$ is of course the norm of the projection of $v_N(y)$ on the space spanned by $\tilde{u}_{M+1}, \dots, \tilde{u}_N$. In the noiseless case, $R(y)$ has singularities at y_j , while in the noisy case (hopefully) pronounced peaks at locations \tilde{y}_j close to y_j are visible. R is sometimes called MUSIC pseudospectrum. We summarize:

Algorithm 3.8 (MUSIC). *Input:* $\tilde{f}(j)$, $j = 0, \dots, N + L - 2$, $N, L \in \mathbb{N}$ and $M \in \mathbb{N}$

- (Incomplete) SVD of $\tilde{H}_{N,L}$: Determine M left singular vectors $\tilde{u}_1, \dots, \tilde{u}_M$ corresponding to the largest singular values.
- Find M largest local maxima of $R(y)$, as given in (3.8), $\tilde{y}_1, \dots, \tilde{y}_M$.

Output: $\tilde{y}_1, \dots, \tilde{y}_M$

An explicit evaluation of the computational complexity of MUSIC is not possible, due to the search for local maxima of $R(y)$.

Remark. In the original version in [87], the signal space is estimated from the (estimated) covariance matrix. However, the approach using SVD is long known as well (see for example the discussion in [82]). We quickly point to a few other variations. A modified version, taking the singular values $\tilde{\sigma}_k$ into account is frequently used as well (see [49, 61]). Also, a reformulation, where roots instead of singularities are considered, is known (the so called root-MUSIC, see the discussion in [61]). Pisarenko's method [70], predating MUSIC, corresponds to the special case $N = M + 1$.

We do not comment on stability properties of MUSIC, instead we point to the literature. A good starting point is [55].

Now we turn to ESPRIT, an acronym for Estimation of Signal Parameter via Rotational Invariance Techniques. This very successful method was published by Roy et al. in 1986 [83], later extended by Roy and Kailath in [82]. We follow [61] closely.

As the name indicates, ESPRIT exploits a rotational invariance of the signal space. Namely, let

$$V_N = \begin{bmatrix} V_N^{(1)} \\ * \\ V_N^{(2)} \\ * \end{bmatrix},$$

i.e., $V_N^{(1)}$ is formed from V_N by discarding the last, $V_N^{(2)}$ by discarding the first row. Again, let $D = \text{diag}(e^{2\pi i y_1}, \dots, e^{2\pi i y_M})$. Then

$$V_N^{(1)} D = V_N^{(2)}. \quad (3.9)$$

This is the rotation ESPRIT uses. The key observation is that it transfers to different bases of $\text{Sig}(f, N)$. Indeed, given a basis $B = [b_1, \dots, b_M]$ of $\text{Sig}(f, N)$, there is a basis transformation matrix $S \in \mathbb{C}^{M \times M}$, S invertible, such that

$$V_N = BS.$$

Now forming $B^{(1)}$ and $B^{(2)}$ by discarding the last resp. first row, (3.9) transforms to

$$B^{(1)} SDS^{-1} = B^{(2)}. \quad (3.10)$$

Also, as $B^{(1)} = V_N^{(1)} DS^{-1}$, the matrix has full rank and hence, if $N \geq M + 1$, the matrix SDS^{-1} is uniquely determined by (3.10). But SDS^{-1} is similar to D and therefore shares the eigenvalues with D , which are precisely what we need to determine Y^f . Thus, any basis of $\text{Sig}(f, N)$ suffices to determine the frequencies of f , at least if $N \geq M + 1$.

In the noisy case, this idea can be used as well. As an estimate for a basis of $\text{Sig}(f, N)$, we again use a SVD and use the left singular vectors corresponding to the largest singular values, see (3.7). We call the matrix $\tilde{U}_M = [\tilde{u}_1 \dots \tilde{u}_M]$, thus (3.10) becomes

$$\tilde{U}_M^{(1)} F = \tilde{U}_M^{(2)}, \quad (3.11)$$

where $F = SDS^{-1}$ for an unknown, invertible matrix S . The idea is now to solve for F and then determine its eigenvalues. There are two approaches to solve this system. Either we use a least square approach, which tries to minimize $\|E_2\|_F$ (the Frobenius norm) subject to

$$\tilde{U}_M^{(1)} F - \tilde{U}_M^{(2)} = E_2.$$

This results in

$$F = \left(\tilde{U}_M^{(1)} \right)^\dagger \tilde{U}_M^{(2)},$$

where A^\dagger denotes the Moore-Penrose inverse (or pseudoinverse) of a matrix A .

But this assumes that $\tilde{U}_M^{(1)}$ is not perturbed. Hence, a total least squares (TLS) approach, which minimizes $\|[E_1 \ E_2]\|_F$ subject to

$$(\tilde{U}_M^{(1)} + E_1)F = \tilde{U}_M^{(2)} + E_2$$

is usually preferable, as it treats both $\tilde{U}_M^{(1)}$ and $\tilde{U}_M^{(2)}$ equally. An extensive discussion of the TLS problem, including many references, is given in [33].

We summarize the algorithm.

Algorithm 3.9 (ESPRIT). *Input:* $\tilde{f}(j)$, $j = 0, \dots, N + L - 2$, $N, L \in \mathbb{N}$ with $N \geq M + 1$, $L \geq M$ and $\text{tol} \geq 0$

- (Incomplete) SVD of $\tilde{H}_{N,L}$: Determine M left singular vectors $\tilde{u}_1, \dots, \tilde{u}_M$ corresponding to the singular values larger than tol .
- Form $\tilde{U}_M^{(1)}$ and $\tilde{U}_M^{(2)}$. Solve (3.11) in the least (or total least) square sense.

- Determine the eigenvalues $e^{2\pi i \tilde{y}_k}$ of F .

Output: $\tilde{y}_1, \dots, \tilde{y}_M$

The computational complexity of ESPRIT is largely due to the complexity of the singular value decomposition, which is $\mathcal{O}(\min(LN^2, L^2N))$. Indeed, the other operations are of lower order: The least square system needs $\mathcal{O}(M^2N)$ flops and determining the eigenvalues requires $\mathcal{O}(M^3)$ flops (see for example [33]). Note that we do not need to calculate the left singular vectors. A more detailed discussion, including an approach using Lanczos bidiagonalization, is given in [77]. Results concerning asymptotic convergence (assuming Gaussian noise) are given in [79, 66].

We mention one final improvement of ESPRIT. While the aforementioned algorithm works even when the frequencies have an real part (e.g. when the signal decays), the following exploits that we only have imaginary frequencies. Namely, for $f \in \mathcal{S}_M^1$ we have that

$$\overline{f(n)} = \overline{\sum_{y \in Y^f} c_y e^{2\pi i y n}} = \sum_{y \in Y^f} \overline{c_y} e^{2\pi i y (-n)} =: g(-n).$$

Then $g \in \mathcal{S}_M^1$ has the same frequencies and we can easily calculate $g(-N), \dots, g(0)$ if we know $f(0), \dots, f(N)$. Therefore, the Hankel matrix $H_{N,L}^f$ can be replaced with

$$\begin{bmatrix} s_N^f(0) & \dots & s_N^f(L-1) & s_N^g(-N+1) & \dots & s_N^g(-L-N+2) \end{bmatrix} = \begin{bmatrix} H_{N,L}^f & J_N \overline{H_{N,L}^f} \end{bmatrix},$$

where J_N is the matrix with ones on the antidiagonal and zeros elsewhere. The matrix $\begin{bmatrix} H_{N,L}^f & J_N \overline{H_{N,L}^f} \end{bmatrix}$ has an extra structure, it is centro-Hermitian. Such matrices can be easily transformed to real matrices of the same size, which has computational advantages. The improved ESPRIT taken advantage of these ideas is known as Unitary ESPRIT and was introduced in [37].

ESPRIT is actually closely related to another technique, the matrix pencil method. This method is introduced in [45], the close connection to ESPRIT (and several variations of ESPRIT) is explained in [46, 76]. We give a quick introduction to the main idea of the matrix pencil method.

As the name indicates, the method recovers the frequencies by considering a matrix pencil. A matrix pencil of two matrices $A_1, A_2 \in \mathbb{C}^{L \times N}$ is a linear combination of them,

$$A_1 - zA_2,$$

where $z \in \mathbb{C}$ is a free variable. As we are interested in possibly singular A_1, A_2 we are considering a singular matrix pencil. Now for any $v \in \mathbb{C}^N$ with $v \notin \ker A_1 \cap \ker A_2$, we have that $(A_1 - zA_2)v = 0$ for at most one $z \in \mathbb{C}$. All $z \in \mathbb{C}$ for which such a v exists are called rank reducing numbers or generalized eigenvalues (GE's) of $A_1 - zA_2$, the corresponding v is called right generalized eigenvector. The connection to the usual eigenvalues and eigenvectors is that if $A_1 \in \mathbb{C}^{N \times N}$ and A_2 is the identity matrix, we obtain $A_1 v - z v = 0$, i.e., v is an eigenvector to the eigenvalue z .

It is clear that if $B \in \mathbb{C}^{N \times P}$ is surjective and $C \in \mathbb{C}^{Q \times L}$ is injective, the GE's of $A_1 - zA_2$ are equal to the GE's of $CA_1B - zCA_2B$ (the only nontrivial observation is that elements in the kernel of B are not allowed to be chosen as generalized eigenvectors).

Now we can use the factorization of H^f , derived in Lemma 3.6, to establish that

$$H_{N,L}^f(1) - zH_{N,L}^f(0) \tag{3.12}$$

has exactly $e^{2\pi i y}$, $y \in Y^f$, as generalized eigenvalues, at least if $N, L \geq M$. Indeed,

$$\begin{aligned} & H_{N,L}^f(1) - zH_{N,L}^f(0) \\ &= V_N(y_1, \dots, y_M) \left(\text{diag}(c_1 e^{2\pi i y_1}, \dots, c_M e^{2\pi i y_M}) - z \text{diag}(c_1, \dots, c_M) \right) V_L(y_1, \dots, y_M)^T \end{aligned}$$

and assuming $N, L \geq M$, V_L and V_N have rank M and hence this matrix pencil has the same GE's as

$$\text{diag}(c_1 e^{2\pi i y_1}, \dots, c_M e^{2\pi i y_M}) - z \text{diag}(c_1, \dots, c_M).$$

But in this case the GE's are obviously equal to $e^{2\pi i y_j}$, $j = 1, \dots, M$.

However, to actually compute the GE's of $H_{N,L}^f(1) - zH_{N,L}^f(0)$, we have to work a little bit more. Due to the fact that the involved matrices are not of full rank, we cannot simply apply a standard

algorithm (e.g. the QZ algorithm, see [33]). One possibility to overcome this difficulty is to compute the reduced singular value decompositions of $H_{N,L}^f(j)$, $j = 0, 1$. Let

$$H_{N,L}^f(j) = U_j \Sigma_j W_j^H, \quad j = 0, 1,$$

where $U_j \in \mathbb{C}^{N \times M}$, $W_j \in \mathbb{C}^{N \times M}$ have orthogonal columns and $\Sigma_j \in \mathbb{C}^{M \times M}$ are diagonal matrices, containing the non-zero singular values of $H_{N,L}^f(j)$. Thus, the matrix pencil (3.12) has the same GE's as

$$\Sigma_0^{-1} U_0^H U_1 \Sigma_1 W_1^H W_0 - z I_M$$

and we have reduced the problem to an eigenvalue problem of a (non-symmetric) matrix in $\mathbb{C}^{M \times M}$.

In the noisy case, we replace $\tilde{H}_{N,L}^f(j)$ by its M -truncated singular value decomposition.

The connection to ESPRIT is not clear from these considerations. Furthermore, performing two SVDs is a bit wasteful. A different idea is to start with a reduced SVD

$$H_{N,L+1}^f = U \Sigma W^H,$$

where $U \in \mathbb{C}^{N \times M}$, $W \in \mathbb{C}^{(L+1) \times M}$ and $\Sigma \in \mathbb{R}^{M \times M}$. We denote by $W(0)$ and $W(1)$ the matrix formed by deleting the last and first row respectively. This results in two factorizations

$$H_{N,L}^f(j) = U \Sigma W(j)^H, \quad j = 1, 2.$$

Note that while $W(j)$ have no longer orthogonal columns, they still have rank M , if $N, L \geq M$, as $H_{N,L}^f(j)$ has rank M . Thus, we can transform (3.12) to

$$W(1)^H - z W(0)^H.$$

If we find a matrix $F \in \mathbb{C}^{M \times M}$ solving

$$W(0)^H F = W(1)^H$$

we see that the matrix pencil has generalized eigenvalues equal to the eigenvalues of F . This, however, is just like equation (3.11) only using the right singular vectors instead of the left ones of $H_{N,L}^f$. A similar discussion about the connection of matrix pencil methods and ESPRIT is already given in [46].

A lot of further results, including perturbation theory, can be found in the literature, see [45, 46, 76].

This concludes our review of univariate methods, relying on an estimation of the signal space.

Numerical Examples

Finally, we give numerical examples. The aim is not to compare these well-known methods, but instead to demonstrate the a posteriori error estimates developed in the last chapter.

We start with a deterministic example, to show the power of Theorem 2.26. We consider the exponential sum f with frequencies

$$Y^f = \{0.1k : k = 0, \dots, 9\}$$


and coefficients $c_{0.1k} = (-1)^k$. As a second exponential sum, we consider g_λ , which is constructed by $Y^{g_\lambda} = \lambda + Y^f$. The coefficients are then determined in a least square sense, such that

$$\|\mathcal{P}_N(f - g_\lambda)\|_2^2$$

is minimized. We pick $N = 39$ and as $f, g_\lambda \in \mathcal{S}^1(0.1)$, we have that $\text{sep } f, \text{sep } g_\lambda \geq \frac{4}{N+1}$ as required in Theorem 2.26.

We consider the frequency error, given by

$$\frac{\pi^2}{3} (N+1)^3 \sum_{y \in Y^f} (|c_y^f|^2 + |c_{n(y)}^{g_\lambda}|^2) |y - n(y)|_{\mathbb{T}}^2,$$

Figure 3.1: Comparison of the errors for different g_λ .

where again $n(y)$ is the frequency of g_λ closest to $y \in Y^f$, the coefficient error

$$\frac{3}{8}(N+1) \sum_{y \in Y^f} |c_y^f - c_{n(y)}^{g_\lambda}|^2$$

and their sum, which we call total error. $\|\mathcal{P}_N(f - g_\lambda)\|_2^2$ is called l^2 error. For various choices of λ , these values are reported in Figure 3.1.

Clearly, the estimate is very sharp and the l^2 error is only approximately 20% larger than the total error. For the last two data points, condition (2.27) is not satisfied. And indeed, for the last data point, the assertion of the theorem does not hold.

Next we consider the a posteriori estimates. We consider randomly chosen exponential sums f with $\text{sep } f \geq 0.1$. We do that by picking the first frequency randomly between 0 and 0.05, the next one is picked with a distance to the first randomly chosen between 0.1 and 0.15 etcetera. The coefficients are then chosen as

$$c_j = r_j e^{2\pi i \phi_j},$$

where r_j and ϕ_j are uniformly distributed in $[0.2, 1.2]$ and $[0, 1]$ respectively.

We then sample f at $0, 1, \dots, 2N$ and add noise of the form

$$\varepsilon_j = X_{1,j} + iX_{2,j},$$

where $X_{k,j} \sim \mathcal{N}(0, \sigma^2)$ are independent and normal distributed random variables. Again we choose $N = 39$. Then we apply ESPRIT, as given in Algorithm 3.9, to the noisy samples (choosing $\text{tol} = 0.1$ and $L = N$), which gives us an estimate \tilde{f} for f .

The frequency error, coefficient error and total error are defined as before. Finally, the error estimate of Corollary 2.27 is given by

$$\left(\left| \|\mathcal{P}_N(f) - \tilde{s}_N(\tilde{f})\|_2^2 - \sigma^2(4N+2) \right|^{\frac{1}{2}} + (2 + \sqrt{2})\sigma(2N+1)^{(1+\delta)/4} \right)^2.$$




Figure 3.2: Comparison of the errors at different noise levels. x -axis: Standard deviation of the noise. All results are averaged over 25 runs.

We choose $\delta = 0.8$, resulting in a probability of at least 95% that our estimate holds true.

As Figure 3.2 indicates, the qualitative behavior of the errors is captured quite well by our estimate.

3.2 Projection-based Methods

In this section we introduce one family of reconstruction techniques for the frequency estimation problem. The main idea is to use univariate methods, which are significantly better understood than their higher dimensional analogs, to obtain an estimate for the multivariate problem. We will focus mostly on the two dimensional problem and comment on the situation in higher dimensions along the lines.

Now we pick a univariate method of our choice. The choice is plenty and always extending. MUSIC, ESPRIT and the closely related matrix-pencil method we already discussed and anyone of them is fine. In our numerical examples later on, we use ESPRIT.

Next we choose a line, parametrized by

$$\ell = \{tv + bv^\perp, t \in \mathbb{R}\} \subset \mathbb{R}^2.$$

Whenever we parametrize a line, we assume that v, v^\perp are orthonormal, $b \in \mathbb{R}$. Of course, v, v^\perp are unique only up to sign. Also, v^\perp is not uniquely determined by v . To avoid this problem, we simply define $v^\perp = (-v_2, v_1)$, where $v = (v_1, v_2)$.

Restricting f to ℓ , we obtain

$$f(tv + bv^\perp) = f|_\ell(t) = \sum_{y \in Y^f} c_y e^{2\pi i b v^\perp \cdot y} e^{2\pi i v \cdot y t} = \sum_{y^\ell \in v \cdot Y^f} \left(\sum_{\substack{y \in Y^f \\ y \cdot v = y^\ell}} c_y e^{2\pi i b v^\perp \cdot y} \right) e^{2\pi i y^\ell t}. \quad (3.13)$$

Hence, $f|_\ell$ is a univariate exponential sum with frequencies $Y^{f|_\ell} \subset v \cdot Y^f$, which are the projections

of the frequencies of f onto the line ℓ . This is the reason why methods exploiting this fact are called projection-based. One of the mayor challenges is, that this inclusion can be strict. In fact, it is easily possible that the occurring coefficient is zero. Nonetheless, sampling $f|_\ell$ at $0, \alpha, \dots, (2M-1)\alpha$ with $\alpha \in \mathbb{R}_{>0}$ allows us to apply our univariate method of choice to obtain (a possibly perturbed estimate of) $Y^{f|_\ell}$ and the corresponding coefficients.

Here we have to choose the step size α with care. We may assume without loss of generality that $v_1 \geq 0$. Then, if $v_2 \geq 0$, we have that $[0, 1)^2 \cdot v \subset [0, v_1 + v_2)$ and therefore choosing $\alpha \leq \frac{1}{\sqrt{2}}$ is fine. On the other hand, if $v_2 < 0$, we have $[0, 1)^2 \cdot v \subset (v_2, v_1)$. Again, choosing $\alpha \leq \frac{1}{\sqrt{2}}$ is possible, but we have to translate the frequencies such that they are in (v_2, v_1) after applying our univariate method.

The question now is for which choices of lines $\ell_j = \{tv_j + b_j v_j^\perp, t \in \mathbb{R}\}$, $j = 1, \dots, K$ a reconstruction of f from $f|_{\ell_1}, \dots, f|_{\ell_K}$ is possible. In the work of Potts and Tasche [75], where projection-based methods were first introduced, it is shown that under the assumption that ℓ_j are pairwise non-parallel and that $Y^{f|_{\ell_j}} = v_j \cdot Y^f$, only $M+1$ lines are sufficient. This can be deduced from an older result by Renyi [80]. In the author's master thesis [24], the assumption $Y^{f|_{\ell_j}} = v_j \cdot Y^f$ has been removed, though at serious computational costs. This result was then published in [25].

We give a complete overview over all projection-based methods using scattered (i.e., pairwise non-parallel) lines in the first part of this chapter, including polished proofs of the aforementioned results. Unfortunately, stability results are neither known nor expected. In fact, we already discussed in the first part of this thesis that the separation distance of the frequencies is vital to obtain stability. Now even if Y^f is perfectly separated, we cannot hope for any separation of $Y^{f|_\ell}$ or $v \cdot Y^f$. Thus, it remains unclear, why the univariate method applied along such a line should give reasonable results, at least if we expect a certain noise level.

To remedy this problem, we consider in the second part of this section what happens if we project on multiple parallel lines. First, we discuss why this still gives us the possibility to reconstruct f . On the first glance, no better stability properties of this method are expected, for the same reasons as above. However, by introducing a method which does not apply a univariate method along each line individually, but rather on all parallel lines simultaneously, these problems are overcome. The proposed method relies on a new variation of the ESPRIT algorithm. At least numerically, we show that this method behaves well. A part of these results is published in [26].

Projection on Scattered Lines

We start by considering the special case of lines passing through the origin. In this case, the offset b in our parametrization of a line ℓ is zero, i.e., $\ell = \{tv, t \in \mathbb{R}\}$. In particular, (3.13) becomes

$$f(tv) = f|_\ell(t) = \sum_{y^\ell \in v \cdot Y^f} \left(\sum_{\substack{y \in Y^f \\ y \cdot v = y^\ell}} c_y \right) e^{2\pi i y^\ell t},$$

i.e., we just have to sum the coefficients of all frequencies which are projected on a single point. As the univariate method gives us all y^ℓ with non-zero coefficients, this problem is not really connected to exponential sums, but simply the question when a discrete, weighted point cloud can be recovered from a finite number of projections. Not surprisingly, this problem did already occur in other areas of mathematics. Most notably, it was considered in [80] and [41], where the aim was to recover a two dimensional discrete probability distribution from projections. This is an easier situation, as then all coefficients are positive and they cannot cancel out. In particular $v \cdot Y^f = Y^{f|_\ell}$ holds true.

To stress that in the case $b = 0$ we actually have projections of weighted point clouds, we restate the problem as follows: Given an unknown, discrete set $X \subset \mathbb{R}^2$, $|X| = M$ and a weight function $w : \mathbb{R}^2 \rightarrow \mathbb{C}$, satisfying

$$w(x) \neq 0 \Leftrightarrow x \in X,$$

how many projections are necessary to recover X uniquely? Obviously, such a w can be constructed from a given f by letting $X = Y^f$ and $w(y) = c_y$.

Note that if w is known, so is X : $\text{supp } w = X$. A projection of w on $v \in \mathbb{R}^2$ defines the weight

function on \mathbb{R}

$$[P_v(w)](x) = w_v(x) = \sum_{\substack{y \in X \\ v \cdot y = x}} w(y). \quad (3.14)$$

This is of course the weight function corresponding to the univariate exponential sum $f|_\ell$.

Then the support of w_v is a subset of the projections of X onto the line $\{tv, t \in \mathbb{R}\}$:

$$\text{supp } w_v \subset \{v \cdot x, x \in X\} = v \cdot X.$$

We obtain the following inequality:

$$|\text{supp } w_v| \leq |v \cdot X| \leq |X| = M.$$

When the inequality on the right-hand side holds strictly, some points in X feature the same projection on $\{tv, t \in \mathbb{R}\}$. Then, their weights are added up. If their weights add up to zero, the inequality on the left-hand side is strict. On the other hand, if equality holds on the right-hand side, so it does on the left-hand side.

Also note that P_v is linear, i.e., for all $t \in \mathbb{R}$, $v \in \mathbb{R}^2$ we have

$$(w + t\tilde{w})_v = w_v + t\tilde{w}_v.$$

Scaling a v results in a scaling of w_v :

$$w_{tv}(x) = \sum_{\substack{y \in X \\ tv \cdot y = x}} w(y) = \sum_{\substack{y \in X \\ v \cdot y = x/t}} w(y) = w_v(x/t).$$

Remark. Another interpretation of a weighted point cloud would be identifying it with a weighted sum of Dirac delta distributions. Instead of $w : \mathbb{R}^2 \rightarrow \mathbb{C}$ we obtain

$$d_w = \sum_{x:w(x) \neq 0} w(x)\delta_x.$$

If the weights are positive and sum up to one, d_w is a probability distribution. The projections are then marginal distributions and the corresponding exponential sum is of course exactly the Fourier transform of this measure. But as no properties of measures are used in this section, we stick with the simpler weight functions.

Projections of Discrete Point Clouds

In this section we summarize all results we are aware of. Note that some of them are stated in a quite different context, we reformulate them to fit our setting.

The first result is given in [80]:

Theorem 3.10. (*Renyi, 1952*) *Assume projections on $M + 1$ pairwise linearly independent $v_k \in \mathbb{R}^2$ are given and that $\text{supp } w_{v_k} = v_k \cdot X$. Then X is uniquely determined and a direct reconstruction of X and w is possible.*

Proof. We follow an idea of Heppes [41], which can be generalized easily to higher dimension. Consider the set

$$\tilde{X} = \{x \in \mathbb{R}^2 : v_j \cdot x \in v_j \cdot X \text{ for all } j = 1, \dots, M + 1\}.$$

Obviously, $X \subset \tilde{X}$. We prove the other inclusion. Let $x \in \tilde{X}$ be arbitrary. By definition, for each $j = 1, \dots, M + 1$ we find a $x_j \in X$ with $v_j \cdot x = v_j \cdot x_j$. As $M + 1$ points x_j are chosen, by the pigeonhole principle, we may find $j \neq k$ with $x_j = x_k$. But then $x \cdot v_j = x_j \cdot v_j$ and $x \cdot v_k = x_j \cdot v_k$. As v_j, v_k are linearly independent, $x = x_j \in X$.

Note that we are able to construct \tilde{X} explicitly. We are left with calculating the weights. For each x we have $M + 1$ projections but only $M - 1$ points in $X \setminus \{x\}$. In particular, we may find a v such that $w_v(x \cdot v) = w(x)$. \square

Remark. We give a few remarks.

1. $M + 1$ cannot be improved, it is optimal in certain examples. Consider for example a regular $2M$ polygon and let v_j be the normals to the sides. Then choosing X as every other vertex will give the same projection as the remaining vertices. More examples are given in [80].
2. This is exactly the result mentioned in [75], where the connection between projection of point clouds and parameter estimation of exponential sums was first established.
3. If we choose a fixed line ℓ , the property $|Y^{f|\ell}| = |Y^f|$ holds generically. Hence, the weight function corresponding to f will satisfy $\text{supp } w_{v_k} = v_k \cdot X$ generically. Even better, only three lines suffice to reconstruct a generic f , by the same argument we will use to prove Theorem 3.12.

If one drops the assumption that $\text{supp } w_{v_k} = v_k \cdot X$, X is still uniquely determined but a direct reconstruction is no longer known, see [25]:

Theorem 3.11. (*Diederichs, Iske, 2015*) *Assume projections on $M + 1$ pairwise linearly independent $v_k \in \mathbb{R}^2$ are given. Then X is uniquely determined.*

Proof. We prove that if $w \neq 0$ and $w_{v_j} = 0$ for $j = 1, \dots, M$ pairwise linearly independent v_j , then $|\text{supp } w| \geq 2M$.

Let such a w be given. Note that for each $x \in X$ and each v_j we have to have a $y \in X$, $y \neq x$ with $x \cdot v_j = y \cdot v_j$. Now consider for each v_j the strip

$$S_j = \{x \in \mathbb{R}^2 : \min v_j \cdot X \leq x \cdot v_j \leq \max v_j \cdot X\}.$$

We have $X \subset S_j$ and each connected component of the boundary of S_j is a line containing at least two elements of X . Thus $\bigcap_{j=1}^M S_j$ contains X and is a convex polygon with at least $2M$ sides. Each side contains at least two elements of X , hence X has to have at least $2M$ elements, the extremal case being when X is exactly the set of vertices of $\bigcap S_j$.

The theorem follows then easily, because if w, \tilde{w} feature equal projections on v_k , $k = 1, \dots, M + 1$, $w - \tilde{w}$ has to have a support of at least $2M + 2$ or $w = \tilde{w}$ as claimed. \square

Remark. There is an easier proof: Choose $x \in X$ and assume that $\text{Re}(w(x)) > 0$. For $w_{v_k}(x \cdot v_k) = 0$ we have to find $x_k \in X$ with $\text{Re}(w(x_k)) < 0$ and $(x - x_k) \cdot v_k = 0$. This gives $x_1, \dots, x_M \in X$, which are pairwise distinct by construction. Repeating this procedure for x_1 gives $\tilde{x}_1, \dots, \tilde{x}_M \in X$ with $\text{Re}(w(\tilde{x}_k)) > 0$. Thus we have found $2M$ distinct points in X .

Unfortunately, this proof does not generalize to the case of arbitrary lines. Interestingly, it does not work if $w : \mathbb{R}^2 \rightarrow \mathbb{F}_2$ either (where \mathbb{F}_2 is the field with two elements).

An adaptive choice of the v_k can reduce the number of necessary lines significantly (see [73]):

Theorem 3.12. (*Plonka, Wischerhoff, 2013*) *Let v_1, v_2 be linearly independent. If $\text{supp } w_{v_k} = v_k \cdot X$, $k = 1, 2$, there exists a vector v_3 such that X and w can be reconstructed from w_{v_j} , $j = 1, 2, 3$.*

Proof. Consider

$$\tilde{X} = \{x \in \mathbb{R}^2 : v_j \cdot x \in v_j \cdot X, j = 1, 2\}.$$

Then $X \subset \tilde{X}$ and $|\tilde{X}| \leq M^2$. Choose a v_3 such that $|v_3 \cdot \tilde{X}| = |\tilde{X}|$. Then

$$X = \{x \in \mathbb{R}^2 : v_j \cdot x \in v_j \cdot X, j = 1, 2, 3\}$$

and $w(x) = w_{v_3}(x \cdot v_3)$. \square

Given enough lines, it is even possible to prove that on a few of them all $x \in X$ have distinct projections. This observation is used in a proof in [14]. We present their techniques isolated from its original purpose.

Theorem 3.13. (*Buhmann, Pinkus, 1995*) *Let v_1, \dots, v_s be in general position, where $s = \binom{M}{2} + 2$. Then there are at least two v_{r_1}, v_{r_2} on which all $x \in X$ have distinct projections, i.e., $|v_{r_j} \cdot X| = M$, $j = 1, 2$.*

Proof. We count. Given $x, y \in X$, $x \neq y$ with $v_j \cdot x = v_j \cdot y$, we see that $v_j \in \text{span}(x - y)^\perp$. As v_1, \dots, v_s are in general position, no other v_k is in this one dimensional subspace. Hence for each pair in X at most one v_j cannot distinguish them. We have $\binom{M}{2}$ such pairs, thus two v_j can distinguish all $x \in X$. \square

Corollary 3.14. (*Buhmann, Pinkus, 1995*) *Let v_1, \dots, v_s be in general position, where $s = \binom{M}{2} + 2$. Assume further that $w(x) \neq w(y)$ for all $x \neq y$ in X . Then X is uniquely determined and can be reconstructed.*

Proof. Using Theorem 3.13, we find v_{r_1}, v_{r_2} such that $|v_{r_j} \cdot X| = M$, $j = 1, 2$. Then for each $x \in \text{supp } w_{v_{r_1}} = v_{r_1} \cdot X$ we find one $y \in v_{r_2} \cdot X$ such that $w_{v_{r_1}}(x) = w_{v_{r_2}}(y)$. Let \tilde{x} be the unique solution to $\tilde{x} \cdot v_{r_1} = x$ and $\tilde{x} \cdot v_{r_2} = y$. Then $\tilde{x} \in X$. \square

In \mathbb{R}^d a reconstruction from projection onto lines is possible as well. Again, many formulations exit. Note that the definition of a projection of a weight function (3.14) can be used verbatim in \mathbb{R}^d as well. We give one particularly useful example, first proved in [35]:

Theorem 3.15. (*Griesmaier, Schmiedecke, 2017*) *Let $v_j \in \mathbb{R}^d$, $j = 1, \dots, K$ be in general position, i.e., any d of them form a basis. Assume $w : \mathbb{R}^d \rightarrow \mathbb{C}$ is a weight function with $\text{supp } w = X$ and $|X| = M$. If $K > (d - 1)(2M - 1)$, we can reconstruct w from w_{v_j} . More specifically,*

$$X = \{x \in \mathbb{R}^d : x \cdot v_j \in \text{supp } w_{v_j} \text{ for at least } (d - 1)M + 1 \text{ different } j\},$$

and for each $x \in X$ there is at least one line where x is projected on an individual point, which gives the weight.

Proof. We call the set on the right-hand side \tilde{X} . Then $X \subset \tilde{X}$. Indeed, if for an $x \in X$ we have $x \cdot v_j \notin \text{supp } w_{v_j}$, there has to be a distinct $x' \in X$ with $x \cdot v_j = x' \cdot v_j$. Suppose this happens for $(d - 1)(M - 1) + 1$ different v_j . By the pigeonhole principle, there is one $x' \in X \setminus \{x\}$ and j_1, \dots, j_d with $x \cdot v_{j_k} = x' \cdot v_{j_k}$. But v_j being in general position yields a contradiction. Note that this also proves the second claim (in fact, this proves that we find d different v_j where x is projected to an individual point).

Now we show $\tilde{X} \subset X$. If $\tilde{x} \in \tilde{X}$, we find for $(d - 1)M + 1$ different j an $x_j \in X$ with $\tilde{x} \cdot v_j = x_j \cdot v_j$. Using the pigeonhole principle once more, we see that there is at least one $x \in X$ with $\tilde{x} \cdot v_j = x \cdot v_j$ for d different v_j . Again, $x = \tilde{x}$ follows. \square

This theorem is also interesting in the two dimensional case. Note that in contrast to Theorem 3.11, a specific (and simple) reconstruction strategy is given, but the number of necessary lines increases from $M + 1$ to $2M$.

General Scattered Lines

Now we again consider the general case. The restriction of f to a line $\ell = \{tv + bv^\perp, t \in \mathbb{R}\}$ gives an exponential sum, which corresponds to the one dimensional weight function

$$w_{v,b}(x) = \sum_{\substack{y \in X \\ v \cdot y = x}} w(y) e^{2\pi i b y \cdot v^\perp}.$$

We sum up a modulation of the weights, corresponding to all points with the same projection. Again, note that it is equivalent whether we are looking for w and know $w_{v,b}$ or whether we are looking for f knowing $f|_\ell$.

Now we are able to translate results from the previous section to this application.

Corollary 3.16. *Let w be a weight function with $\text{supp } w = X$ and $|X| = M$. If w_{v_j, b_j} , $j = 1, \dots, M + 1$ are given where v_j are pairwise non-parallel and $\text{supp } w_{v_j, b_j} = v_j \cdot X$, then w is uniquely determined and a reconstruction is possible.*

Proof. Just note that the proof of Theorem 3.10 is still applicable. This observation was first made in [75]. \square

Remark. $M + 1$ is optimal in this case as well. The $2M$ polygon gives again an example. A detailed discussion can be found in [99].

Corollary 3.17. *Let w be a weight function with $\text{supp } w = X$ and $|X| = M$. If w_{v_j, b_j} , $j = 1, \dots, M+1$ are given where v_j are pairwise non-parallel, then w is uniquely determined.*

Proof. The proof of Theorem 3.11 applies here as well (but not the shorter proof!). \square

Unfortunately, it does not give a direct reconstruction. However, a reformulation as an optimization problem can be given, where we switch back to the setting of parameter estimation of exponential sums.

Theorem 3.18. *(Diederichs, Iske, 2015) Let $f : \mathbb{R}^2 \rightarrow \mathbb{C}$ be an exponential sum of order M and $\ell_j = \{tv_j + b_jv_j^\perp, t \in \mathbb{R}\}$, $j = 1, \dots, M+1$ a collection of pairwise non-parallel lines. If G is a set of samples on the lines ℓ_j sufficient to determine each of the restrictions $f|_{\ell_j}$, f is the unique solution to*

$$\min \|\mathbf{c}\|_0 \quad (3.15)$$

$$\text{subject to: } \sum_{y \in \tilde{Y}} c_y e^{2\pi i x \cdot y} = f(x), \quad \forall x \in G. \quad (3.16)$$

Here, \tilde{Y} is the following finite set containing all frequency vectors:

$$\tilde{Y} = \{y \in \mathbb{R}^2 : \exists j_1, j_2 \text{ such that } y \cdot v_{j_k} \in \text{supp } w_{v_{j_k}, b_{j_k}} \text{ for } k = 1, 2\}.$$

Proof. Pick any $y \in Y^f$. As there are $M-1$ other frequencies in Y^f , but $M+1$ non-parallel lines, we have at least two lines where y has a different projection than all other frequency vectors. In particular, $y \in \tilde{Y}$. Hence, f is admissible and the optimization problem has a M -sparse solution.

Now assume that there is a second M -sparse solution g . Due to the uniqueness of the univariate problem, we have that $f|_{\ell_j} = g|_{\ell_j}$. By Corollary 3.17, $f = g$. \square

Remark. There is another formulation of this theorem, which captures more the essence of its proof. Assume we have lines ℓ_1, \dots, ℓ_K as above. Let $H = \cup \ell_j$ the union of these lines. Then the restriction operator

$$\cdot|_H : \mathcal{S}^2 \rightarrow \mathcal{C}(H, \mathbb{C})$$

is an algebra homomorphism and hence its kernel is an ideal. We then proved

$$f \in \ker \cdot|_H \setminus \{0\} \quad \Rightarrow \quad \text{ord } f \geq 2K.$$

As restricting f to a line gives an exponential sum of order at most M , instead of H a suitable discrete set can be chose, namely G .

Returning to Theorem 3.18, we observe that the stated optimization problem (3.15) is quite costly to solve. One difficulty is that \tilde{Y} is rather large, an upper bound is given by $|\tilde{Y}| \leq \binom{M}{2} M^2$. Further, such non-convex problems are in general hard to solve. In fact, minimizing $\|\cdot\|_0$ under a linear constraint $Ax = b$ is known to be NP-hard, see [32] Section 2.3. Hence, usually a convex relaxation is considered, namely replacing $\|\cdot\|_0$ by $\|\cdot\|_1$. In some cases, it can be proved that the relaxation does not change the solution. Here however, elements in \tilde{Y} can be arbitrarily close which results in strongly correlated columns in the matrix A , which is problematic.

By taking samples on more lines, this problem can be remedied by relying on Theorem 3.15. Indeed, Theorem 3.15 directly applies to arbitrary lines, without any modification of the stated proof. This works in the higher dimensional case as well, however as a line is then parametrized by

$$\ell = \{tv + \tilde{v}, t \in \mathbb{R}\},$$

where $\tilde{v} \perp v$, the projection on ℓ of a weight function is defined by

$$w_{v, \tilde{v}}(x) = \sum_{\substack{y \in X \\ v \cdot y = x}} w(y) e^{2\pi i y \cdot \tilde{v}}.$$

For completeness, we state the result (which follows directly from Theorem 3.15, proved by Griesmaier and Schmiedecke in [35]).

Corollary 3.19. *Let $v_j \in \mathbb{R}^d$, $j = 1, \dots, K$ be in general position. Assume $w : \mathbb{R}^d \rightarrow \mathbb{C}$ is a weight function with $\text{supp } w = X$ and $|X| = M$. If $K > (d-1)(2M-1)$, it is possible to reconstruct w from w_{v_j, \tilde{v}_j} , where $\tilde{v}_j \in \mathbb{R}^d$ has to be orthogonal to v_j but can otherwise be arbitrarily chosen. Again,*

$$X = \{x \in \mathbb{R}^d : x \cdot v_j \in \text{supp } w_{v_j, \tilde{v}_j} \text{ for at least } (d-1)M + 1 \text{ different } j\}.$$

Finally, we consider the generic situation in d dimensions. As expected, everything works out smoothly.

Theorem 3.20. *Let $\ell_j = \{tv_j + \tilde{v} : t \in \mathbb{R}\}$, $j = 1, \dots, d+1$ be lines in \mathbb{R}^d , where v_j are in general position. Then for a generic f , it is possible to recover it from $f|_{\ell_j}$. Explicitly,*

$$Y^f = \left\{ y \in \mathbb{R}^d : v_j \cdot y \in Y^{f|_{\ell_j}} \text{ for } j = 1, \dots, d+1 \right\}.$$

Proof. We call the set on the right-hand side Y . Obviously, $Y^f \subset Y$ and we are left to prove $Y \subset Y^f$. Consider

$$\tilde{Y} = \left\{ y \in \mathbb{R}^d : v_j \cdot y \in Y^{f|_{\ell_j}} \text{ for } j = 1, \dots, d \right\} \supset Y.$$

This is a finite set (as v_1, \dots, v_d are linearly independent). By symmetry it suffices to prove for one $y \in Y^f$ that the probability of $x \in \tilde{Y} \setminus \{y\}$ satisfying $y \cdot v_{d+1} = x \cdot v_{d+1}$ is zero. By assumption, there is an index $j^* \in \{1, \dots, d\}$ such that $x \cdot v_{j^*} = z \cdot v_{j^*}$ with $z \in Y^f \setminus \{y\}$. Further, if all other frequencies are already chosen, $y \cdot v_{d+1} = x \cdot v_{d+1}$ is only satisfied if z lies in a proper affine linear space of \mathbb{R}^d , an event with probability zero. \square

Numerical Verification

We close this section by giving at least a small numerical example. Due to the very special structure of the sampling points, the only available methods are the projection-based ones. Neither methods based on TV minimization and convex optimization (like [16]) nor multivariate Prony-type methods (as discussed in the next section) can handle these sampling sets, as they cannot be rescaled to be a subset of \mathbb{Z}^d . Hence, the connection to multivariate (trigonometric) polynomials is not available. Other methods, like the multivariate matrix pencil method presented in [44, 81] rely on sampling sets contained in \mathbb{Z}^d as well.

Now assume that we sample an unknown exponential sum f on $\ell_1 = \{tv_1 + b_1 v_1^\perp : t \in \mathbb{R}\}, \dots, \ell_K$. The frequency sets $Y^{f|_{\ell_j}}$ can be estimated by applying a univariate method to the samples on ℓ_j . In our numerical experiments we use least squares ESPRIT, Algorithm 3.9.

However, in general the estimates $Y_{\text{est}}^{f|_{\ell_j}}$ may differ substantially from $Y^{f|_{\ell_j}}$, in particular in the presence of noise. To account for this fact, we have to introduce a tolerance $\varepsilon \geq 0$ in such a way that all frequencies in $Y^{f|_{\ell_j}}$ are at most perturbed by ε in our estimation.

Depending on a priori assumptions on f (in particular assuming $\text{ord } f \leq M$), we choose one of the following alternatives:

1. If we know (or assume) that $v_j \cdot Y^f = Y^{f|_{\ell_j}}$ and $K \geq M+1$, we can reconstruct Y^f by letting

$$Y_{\text{est}} = \left\{ x \in \mathbb{R}^2 : \text{dist}(x \cdot v_j, Y_{\text{est}}^{f|_{\ell_j}}) < \varepsilon \text{ for all } j = 1, \dots, K \right\}.$$

This approach was proposed in [75] and can be used when K is smaller than $M+1$ as well, though no guarantee to recover Y^f is possible.

2. If $K \geq 2M$, we estimate Y^f by

$$Y_{\text{est}} = \left\{ x \in \mathbb{R}^2 : \text{dist}(x \cdot v_j, Y_{\text{est}}^{f|_{\ell_j}}) < \varepsilon \text{ for at least } M+1 \text{ different } j = 1, \dots, K \right\}.$$

Note that $Y = Y^f$ by Corollary 3.19. This strategy is in the spirit of [35].

3. Finally, if $K \geq M+1$, we form the candidate set

$$\tilde{Y} = \left\{ x \in \mathbb{R}^2 : \exists j_1, j_2 \text{ such that } \text{dist}(x \cdot v_{j_k}, Y^{f|_{\ell_{j_k}}}) < \varepsilon \text{ for } k = 1, 2 \right\}.$$




Figure 3.3: ℓ_j , $j = 1, \dots, 5$. On the left-hand side: Y^f and projections of Y^f onto the lines. On the right-hand side: Candidate set \tilde{Y} and the projections we actually observe.

We denote the solution to the optimization problem (3.15) by $c^* \in \mathbb{R}^Y$. Then

$$Y_{\text{est}} = \left\{ y \in \tilde{Y} : |c_y^*| \geq \varepsilon_1 \right\}$$

gives Y^f by Theorem 3.18. This strategy was proposed by Iske and the author in [25]. Note that we can choose any $\varepsilon_1 > 0$, but if some lower bound of $|c_y|$ is available, we should use it here. Assuming such a lower bound is necessary for stability guarantees anyway, as discussed in the previous chapter.

Of course, using stronger assumptions on f gives better results. This means that when applicable, the first method is preferable over the second, which again is preferable over the third.

Note that if $\varepsilon = 0$ and $Y_{\text{est}}^{f|\ell_j} = Y^{f|\ell_j}$, all three possibilities will recover f .

We give an example, where the third strategy is necessary. To this end, we choose the five lines

$$\begin{aligned} \ell_1 &= \{te_1 : t \in \mathbb{R}\}, & \ell_2 &= \{te_2 : t \in \mathbb{R}\}, \\ \ell_3 &= \{t(1,1)/\sqrt{2} + (-1,1)/\sqrt{2} : t \in \mathbb{R}\}, & \ell_4 &= \{t(1,-1)/\sqrt{2} + (1,1)/\sqrt{2} : t \in \mathbb{R}\}, \\ \ell_5 &= \{t(\sin(\pi/8), \cos(\pi/8)) : t \in \mathbb{R}\}. \end{aligned}$$

As frequencies we choose

$$\begin{aligned} y_1 &= (0.1, 0.1) & y_2 &= (0.1, 0.6) \\ y_3 &= (0.6, 0.1) & y_4 &= (0.6, 0.6) \end{aligned}$$

with coefficients $(c_1, c_2, c_3, c_4) = (1, -1, -1, 1)$. This example is build in such a way that

$$Y^{f|\ell_1} = \emptyset, \quad Y^{f|\ell_2} = \emptyset, \quad Y^{f|\ell_3} = \{0.14, 0.85\}, \quad Y^{f|\ell_4} = \{-0.35, 0.35\},$$

rounded up to two decimal places. In particular $Y^{f|\ell_j} \not\subseteq v_j \cdot Y^f$ for $j = 1, \dots, 4$ and a reconstruction of f using an arbitrary number of samples on ℓ_j , $j = 1, \dots, 4$ is not possible. See figure 3.3 for a depiction of the situation.

In this instance, the non-convex optimization problem (3.15) can be solved quite easily, as the number of candidates is smaller than the number of samples and we therefore expect that only one solution to the linear constraint exists. We test how stable the procedure is. To this end, we add noise

$$\tilde{f}_j(k/\sqrt{2}) = f|_{\ell_j}(k/\sqrt{2}) + n_{j,k,1} + in_{j,k,2},$$

where $n_{j,k,l}$ are independent and uniformly distributed in $\delta[-0.5, 0.5]$ for various choices of $\delta \in \mathbb{R}_{>0}$.

Table 3.1: Results of the example

δ	N	L	tol	c_{\min}	$d_H(Y^f, Y_{est})$	fails/100
0	5	4	0.5	0.01	1e-30	-
0	10	9	0.5	0.01	1e-30	-
1e-6	5	4	0.5	0.01	1e-14	-
1e-6	10	9	0.5	0.01	2e-15	-
1e-2	5	4	0.5	0.01	1e-06	-
1e-2	10	9	0.5	0.01	2e-07	-
0.5	5	4	2	0.5	3e-03	10
0.5	10	9	2	0.5	5e-04	11

As an error measure we give

$$d_H(Y^f, Y_{est}) = \max \left\{ \max_{y \in Y^f} \min_{\tilde{y} \in Y_{est}} \|y - \tilde{y}\|_2, \max_{\tilde{y} \in Y_{est}} \min_{y \in Y^f} \|y - \tilde{y}\|_2 \right\},$$

the Hausdorff distance, whenever the number of frequencies matches, otherwise we consider the recovery failed. Our findings are collected in Table 3.1. All results are averaged over 100 runs. The parameter N, L, tol are as described in Algorithm 3.9. In this example, the proposed method works well when the noise is reasonable small. In the high noise regime, one has to carefully choose tol and have a reasonable guess for c_{\min} to still get useful results.

To conclude this numerical test, when Y^f is very small, the proposed method can be used. However, even for moderately large Y^f , the candidate set is too large and it is advisable to collect enough samples to follow the strategy presented by Griesmaier and Schmiedecke [35], which was introduced after the author's method and is for all practical purposes preferable. Finally, we remark that Griesmaier and Schmiedecke also gave a MUSIC-like pairing scheme, which looks promising. However, this is not the place to further discuss their results.

3.2.1 Projection on Parallel Lines

Now we turn to sampling sets on multiple, parallel lines. In this section, we always assume without loss of generality that we sample along lines

$$\ell_j = \{te_1 + b_j e_2, t \in \mathbb{R}\}, \quad j = 1, \dots, K,$$

where $e_1 = (1, 0)$, $e_2 = (0, 1)$ and $b_j \in \mathbb{R}$ pairwise distinct. This covers the most prominent sampling set, both in applications and theory, namely

$$G_N = \{n \in \mathbb{Z}^2 : \|n\|_\infty \leq N\},$$

where $N \in \mathbb{N}$, or subsets thereof. Indeed, by choosing $K = 2N + 1$ and $b_j = -N - 1 + j$, the set G_N provides us with $2N + 1$ equispaced samples on each line ℓ_j , $j = 1, \dots, 2N + 1$. If $N \geq M = \text{ord } f$, this gives us sufficient samples to determine $f|_{\ell_j}$ by any univariant method.

But why does sampling along parallel lines help us? We have that

$$f|_{\ell_j}(t) = \sum_{y \in Y^f} c_y e^{2\pi i b_j y_2} e^{2\pi i y_1 t} = \sum_{y_1 \in e_1 \cdot Y^f} \left(\sum_{y_2: (y_1, y_2) \in Y^f} c_{(y_1, y_2)} e^{2\pi i b_j y_2} \right) e^{2\pi i y_1 t}. \quad (3.17)$$

Here we denote the j th component of $y \in \mathbb{R}^2$ by y_j . The idea when using scattered lines is that we obtain multiple different projections. Here, we only see one projection. But note that the corresponding coefficients vary for different b_j . Even if for one particular b_j a critical cancellation occurs and one frequency vanishes, we might observe it for different values of b_j .

In the first part of this section, we make this observation precise and prove under which conditions we are able to recover f from $f|_{\ell_j}$. In the second part we turn to an efficient implementation of this idea. We introduce a variation of ESPRIT, which can be applied to all parallel lines at once, which greatly stabilizes the scheme.

Partly, these results are published in [26].

Recovery Guarantees for Parallel Lines

A close look on (3.17) reveals, that the coefficients of the one dimensional frequencies of $f|_{\ell_j}$ form an exponential sum (in b_j) themselves. Indeed, we have that

$$c_{y_1}(b_j) = \sum_{y_2: (y_1, y_2) \in Y^f} c_{(y_1, y_2)} e^{2\pi i b_j y_2}. \quad (3.18)$$

Now as we know $f|_{\ell_j}$, we know $c_{y_1}(b_j)$ for $j = 1, \dots, K$. If b_j form an arithmetic progression with step size one, we can apply a univariate method to $c_{y_1}(\cdot)$, which gives us the frequencies and their coefficients. In particular, the sampling set G_N can be used here.

In the following, we use the projection on the last $d - 1$ coordinates

$$P_{d-1} : \mathbb{R}^d \rightarrow \mathbb{R}^{d-1}, \quad P_{d-1}(x_1, \dots, x_d) = (x_2, \dots, x_d).$$

Theorem 3.21. (Diederichs, Iske, 2017) *Let f be an exponential sum of order at most M . Then f can be uniquely determined (under all exponential sums of order at most M) from samples taken on*

$$G_N = \{(n_1, \dots, n_d) \in \mathbb{Z}^d : |n_j| \leq N \forall j = 1, \dots, d\}$$

for any $N \geq M$ by only solving a finite number of one dimensional parameter estimation problems.

Proof. We use induction in d . The case $d = 1$ is clear. We rewrite f as

$$f(n_1, x) = \sum_{y \in Y^f} c_y e^{2\pi i y_1 n_1} e^{2\pi i P_{d-1} y \cdot x} = \sum_{\tilde{y} \in P_{d-1} Y^f} \left(\sum_{y_1: (y_1, \tilde{y}) \in Y^f} c_{(y_1, \tilde{y})} e^{2\pi i y_1 n_1} \right) e^{2\pi i \tilde{y} \cdot x}.$$

Using the induction hypothesis, we are able to determine for each $|n_1| \leq N$ the frequencies and coefficients of $f(n_1, \cdot) \in \mathcal{S}_M^{d-1}$. We therefore know for each $\tilde{y} \in P_{d-1} Y^f$ and each $|n_1| \leq N$

$$c_{\tilde{y}}(n_1) = \sum_{y \in P_{d-1}^{-1}\{\tilde{y}\} \cap Y^f} c_{(y_1, \tilde{y})} e^{2\pi i y_1 n_1}.$$

Note that $c_{\tilde{y}}(\cdot)$ cannot vanish for all $|n_1| \leq N$, as by assumption $c_{\tilde{y}} \in \mathcal{S}_M^1$. Further, we can recover $c_{(y_1, \tilde{y})}$ and y_1 for each \tilde{y} by applying any univariate method. Now clearly

$$Y^f = \{(y_1, \tilde{y}) : \tilde{y} \in P_{d-1} Y^f \text{ and } y_1 \in Y^{c_{\tilde{y}}}\},$$

while the coefficients are given by $c_{(y_1, \tilde{y})}$. □

Remark. We give a few comments on this theorem.

1. Cuyt and Lee pursue a similar idea in [22], developed independently from the author's result. They also use the property that the coefficients again form an exponential sum (though without having projections in mind). However, the focus of this work is different: Instead of looking for a fixed sampling set, an adaptive strategy is presented, which in most cases uses significantly less samples.
2. One of the important advantages of this method is that no matching of different projections is required. They are matched automatically, as the second component of the frequencies occur in the coefficient of the first component. Nonetheless, it is necessary to match projections on parallel lines. Indeed when we use the induction hypothesis, we implicitly make the assumption that we can match correctly the frequencies of $f(n, \cdot)$ and $f(m, \cdot)$. While we just have to match equal numbers in the noise-free case, if we only have perturbed frequencies this might cause problems. In the next section, we are able to overcome these problems. The matching is one of the main disadvantages of taking scattered lines, but also of direct multivariate version of the matrix pencil method [44]. In other multivariate methods, this matching comes with the cost of having to diagonalize certain matrices simultaneously. We discuss this approach later in more detail.
3. We do not state computational cost, as we give an efficient algorithm later on.

4. In [26], Iske and the author considered a slightly smaller sampling set. We refrained here from this small gain, as we give a significantly (in fact optimal) smaller sampling set next.

We turn now to minimal sampling sets. To this end, a little bit of preparation is necessary.

It is well known that whenever $N \geq M$ and $y_1, \dots, y_M \in [0, 1)$ are mutually distinct, the vectors

$$v_N(y_j) = (e^{2\pi i y_j k})_{k=0, \dots, N-1}$$

are linearly independent. Indeed, the Vandermonde matrix with columns $v_N(y_j)$ has full rank. We prove an analog property for multivariate Vandermonde vectors.

Definition 3.22. *We define*

$$\Gamma_N^d = \left\{ (n_1, \dots, n_d) \in \mathbb{N}_0^d : \prod_{j=1}^d (n_j + 1) \leq N \right\},$$

which is the non-negative part of the hyperbolic cross. Further, for $y \in [0, 1)^d$, we let

$$v_{\Gamma_N^d}(y) = v_N(y) = (e^{2\pi i y \cdot k})_{k \in \Gamma_N^d}.$$

Later, we use $v_G(y) = (e^{2\pi i y \cdot n})_{n \in G} \in \mathbb{C}^G$ for arbitrary, finite $G \subset \mathbb{Z}^d$.

Now we prove linear independence of M Vandermonde vectors $v_{\Gamma_M^d}$. This has been proved by Sauer in [85], though our proof here differs significantly.

Lemma 3.23. *Let $Y \subset [0, 1)^d$ be a finite set, $|Y| = M$. Then for $N \geq M$, the vectors $v_N(y)$, $y \in Y$ are linearly independent.*

Proof. We use induction over d . The case $d = 1$ is clear. Assume we have $c_y \in \mathbb{C}$ such that

$$\sum_{y \in Y} c_y v_N(y) = 0.$$

Then we have that

$$\sum_{y \in Y} c_y e^{2\pi i y \cdot k} = 0, \quad \forall k \in \Gamma_N^d. \quad (3.19)$$

Let $P_{d-1} : \mathbb{C}^d \rightarrow \mathbb{C}^{d-1}$ be the projection on the last $d-1$. We can rewrite (3.19) as

$$\sum_{y \in P_{d-1}Y} \left(\sum_{y_1 : (y_1, y) \in Y} c_{(y_1, y)} e^{2\pi i y_1 k_1} \right) e^{2\pi i y \cdot k} = 0, \quad \forall (k_1, k) \in \Gamma_N^d.$$

As $(0, k) \in \Gamma_N^d$ for all $k \in \Gamma_N^{d-1}$, we can apply the induction hypothesis and obtain that

$$\sum_{y_1 : (y_1, y) \in Y} c_{(y_1, y)} = 0, \quad \forall y \in P_{d-1}Y. \quad (3.20)$$

Thus, for all $y \in P_{d-1}Y$ with $P_{d-1}^{-1}\{y\} \cap Y = \{y_1\}$, we know that $c_{(y_1, y)} = 0$ and hence we can assume that no such $y \in P_{d-1}Y$ exist (by possibly reducing M).

But this implies that $|P_{d-1}^{-1}\{y\} \cap Y| \geq 2$ for all $y \in Y$ and hence that $|P_{d-1}Y| \leq \lfloor \frac{M}{2} \rfloor$. Now we can use that $(1, k) \in \Gamma_N^d$ for all $k \in \Gamma_{\lfloor \frac{N}{2} \rfloor}^{d-1}$ and hence, by induction hypothesis once more, that

$$\sum_{y_1 : (y_1, y) \in Y} c_{(y_1, y)} e^{2\pi i y_1} = 0, \quad \forall y \in P_{d-1}Y. \quad (3.21)$$

However, this implies (by (3.20), (3.21) and the case $d = 1$) that for all $y \in P_{d-1}Y$ with

$$|P_{d-1}^{-1}\{y\} \cap Y| \leq 2$$

the coefficients vanish. Therefore, we may assume that $|P_{d-1}^{-1}\{y\} \cap Y| \geq 3$. Repeating this argument gives that all c_y have to be zero, (using $(r, k) \in \Gamma_N^d$ for all $k \in \Gamma_{\lfloor \frac{N}{r} \rfloor}^{d-1}$ and $r \in \mathbb{N}_{>0}$). \square

This proof works for more general frequency vectors in $[0, 1)^d + i\mathbb{R}^d$ as well. With this in mind, we comment quickly on the point of view in [85]. Consider the space of polynomials

$$\Pi_{\Gamma_M^d} = \text{span}\{z^n = z_1^{n_1} \dots z_d^{n_d} : n \in \Gamma_M^d\}.$$

Clearly, Lemma 3.23 implies that for any $X \subset \mathbb{R}^d + i[0, 1)^d$ with $|X| \leq M$ and any $c \in \mathbb{C}^X$ there is a polynomial $p \in \Pi_{\Gamma_M^d}$ with $p(x) = c_x$ for all $x \in X$. Indeed, the interpolation matrix has full rank and therefore is surjective. Of course it cannot be injective as $|\Gamma_M^d| > M$ for $d > 1$. This does not come as a surprise, as having a unique interpolant as well would contradict the Mairhuber-Curtis theorem, see for example [97]. The set Γ_M^d is in certain ways an optimal choice among all $G \subset \mathbb{N}_0^d$ with this universal interpolation property, see [85] for details.

Furthermore, we remark that this improves a lemma in [54], where full rank of Vandermonde matrices on a full grid $[0, N]^d \cap \mathbb{Z}^d$ were proved. In [54], this result is then used to give recovery guarantees for a multivariate MUSIC algorithm, which can be improved using Lemma 3.23. We return to this later.

Theorem 3.24. *Let $f \in \mathcal{S}_M^d$ be an exponential sum. Further, let*

$$\tilde{\Gamma}_M^d := \left\{ n \in \mathbb{N}_0^d : \prod_{j=1}^d \left\lceil \frac{n_j + 1}{2} \right\rceil \leq M \right\}.$$

Then f is uniquely determined by $(f(k))_{k \in \tilde{\Gamma}_M^d}$ and Y^f , as well as the coefficients, can be recovered using univariate methods alone.

Proof. First, we establish uniqueness. Clearly, $\Gamma_{2M}^d \subset \tilde{\Gamma}_M^d$. Assume that $f, g \in \mathcal{S}_M^d$ are equal on Γ_{2M}^d . We have that

$$(f(k) - g(k))_{k \in \Gamma_{2M}^d} = \sum_{y \in Y^f} c_y^f v_{2M}(y) - \sum_{y \in Y^g} c_y^g v_{2M}(y)$$

and as $|Y^f \cup Y^g| \leq 2M$ the preceding lemma gives $f = g$.

To recover f , we give a strategy which at the end proposes a finite number of candidates. We then evaluate all of them on Γ_{2M}^d . f will always be one of the candidates and can hence, by the first part of this proof, be identified uniquely. Furthermore, we use an induction on d , where the case $d = 1$ follows trivially.

To this end, we use the following property, relating $\tilde{\Gamma}_M^d$ to $\tilde{\Gamma}_M^{d-1}$:

$$(2k_1 - 1, k) \in \tilde{\Gamma}_M^d \Leftrightarrow k \in \tilde{\Gamma}_{\lfloor \frac{M}{k_1} \rfloor}^{d-1} \quad \forall k_1 \in \mathbb{N}_{>0}, k \in \mathbb{N}_0^{d-1}.$$

The general idea is similar to the preceding lemma. We write f as

$$f(k_1, k) = \sum_{P_{d-1} Y^f} \left(\sum_{y_1: (y_1, y) \in Y^f} c_{(y_1, y)} e^{2\pi i y_1 k_1} \right) e^{2\pi i k \cdot y}.$$

Now we note that if $k_1 = 1$ we have $(2(k_1 - 1), k), (2k_1 - 1, k) \in \tilde{\Gamma}_M^d$ for all $k \in \tilde{\Gamma}_M^{d-1}$ and hence, by induction hypothesis, we are able to recover all $y \in P_{d-1} Y^f$ for which the two coefficients

$$\left(\sum_{y_1: (y_1, y) \in Y^f} c_{(y_1, y)} e^{2\pi i y_1 n_1} \right), \quad n_1 = 0, 1$$

do not vanish simultaneously. Assume for the moment that we knew for which $y \in P_{d-1} Y^f$ only one y_1 with $(y_1, y) \in Y^f$ exists (and collect them in Y^1), we could calculate $c_{(y_1, y)}$ for all such y and then

consider

$$f^{(R)}(x_1, x) = f(x_1, x) - \sum_{y \in Y^1} c_{(y_1, y)} e^{2\pi i(x_1 y_1 + x \cdot y)}.$$

Note that we know $f^{(R)}|_{\tilde{\Gamma}_M^d}$. Further, $f^{(R)}$ has the property that for all $y \in P_{d-1}Y^{f^{(R)}}$, there are at least two y_1, \tilde{y}_1 with $(y_1, y), (\tilde{y}_1, y) \in Y^{f^{(R)}}$. This on the other hand implies that $|P_{d-1}Y^{f^{(R)}}| \leq \lfloor \frac{M}{2} \rfloor$ and hence using the induction hypothesis, we are able to determine the coefficients of all $y \in P_{d-1}Y^f$ for which

$$\left(\sum_{y_1: (y_1, y) \in Y^f} c_{(y_1, y)} e^{2\pi i y_1 n_1} \right), \quad n_1 = 0, 1, 2, 3$$

do not vanish simultaneously; indeed $(2(k_1 - 1), k), (2k_1 - 1, k) \in \tilde{\Gamma}_M^d$ for all $k \in \tilde{\Gamma}_{\lfloor \frac{M}{2} \rfloor}^{d-1}$ and $k_1 = 1, 2$. But as these are again exponential sums, we are able to determine y_1 and $c_{(y_1, y)}$ for all y where only at most two y_1, \tilde{y}_1 with $(y_1, y) \in Y^{f^{(R)}}$ exist. Again assuming we somehow knew which $y \in P_{d-1}Y^{f^{(R)}}$ where of this kind, we could again subtract them, obtain a new exponential sum and iterating this will give (after at most M steps) f .

But we do not know the set Y^1 beforehand (nor the similarly formed sets in the subsequent steps). On the other hand, we can of course take any guess for Y^1 and proceed. As for each guess only a finite number of possibilities exist, and only M steps are necessary, we are able to create a finite number of candidates. As described earlier, we are done. \square

Remarks. 1. This procedure is so outrageously inefficient, that the author does not try to give the precise computational complexity. Note however, that if $|Y^1| = M/2$, we have to build (at least) one candidate function for each subset of Y^1 . This alone gives $\mathcal{O}(2^{M/2})$. Therefore, Theorem 3.24 should be considered as being of theoretical interest only.

2. As $\Gamma_{2M}^d \subset \tilde{\Gamma}_M^d \subset \Gamma_{2^d M}^d$ and $|\Gamma_M^d| \leq M \log^{d-1}(M)$ (see Lemma 1.4, p. 71 in [60]), only $\mathcal{O}_d(M \log^{d-1}(M))$ sampling points are needed. This improves a result by Sauer [85], which gives a set G with $|G| \leq (d+1)M^2 \log^{2d-2}(M)$. However, Sauer's method, as well as multivariate ESPRIT-type methods are of polynomial runtime, as we will soon see.

3. In the uniqueness part, we actually proved that

$$\cdot|_{\Gamma_{2M}^d} : \mathcal{S}_M^d \rightarrow \mathbb{C}^{\Gamma_{2M}^d}, \quad f \mapsto f|_{\Gamma_{2M}^d}$$

is one-to-one. Thus, reconstructing f from $f|_{\Gamma_{2M}^d}$ is possible. But the only reconstruction scheme is searching for f in \mathcal{S}_M^d . An algorithm that only needs finite time is unknown.

4. While Theorem 3.24 includes Theorem 3.21 as a special case, the procedure given in the proof of Theorem 3.21 is more practical and the foundation of an algorithm presented later on.

Now we consider once again the generic situation, where only $2dM$ samples are required - this is closer to the setting of Cuyt and Lee in [22]. In fact, this theorem can be easily and directly deduced from their considerations. We give a proof more in line with the rest of this chapter.

Theorem 3.25. *If we sample a $f \in \mathcal{S}_M^d$ on*

$$G = \bigcup_{k=1}^d \{(n, \delta_{2,k}, \dots, \delta_{d,k}) : -M < n \leq M\},$$

we are generically able to reconstruct it using univariate methods alone.

Proof. In the generic situation, all frequencies of $f \in \mathcal{S}_M^d$ have a different first coordinate. We apply a univariate method to $f(k, 0, \dots, 0)$, $k = -M + 1, \dots, M$. As

$$f(x, 0, \dots, 0) = \sum_{y \in Y^f} c_y e^{2\pi i x y_1}$$

this gives us the first coordinate of all frequencies as well as their coefficients. Applying a univariate method to $f(k, 1, 0, \dots, 0)$. $k = -M + 1, \dots, M$, we see that

$$f(x, 1, 0, \dots, 0) = \sum_{y \in Y^f} c_y e^{2\pi i y_2} e^{2\pi i x y_1}$$

and hence we can easily calculate the second component. Analogously, we identify all other components of the frequencies of f . \square

Remark. Related results can be found in [91], extended in [48]. There, sampling sets of the form $G = [0, I_1] \times \dots \times [0, I_d] \cap \mathbb{Z}^d$ are considered and various criteria under which a generic $f \in \mathcal{S}_M^d$ is uniquely determined by $f|_G$ are given.

Simultaneous Frequency Estimation

Now we give an efficient algorithm to estimate the frequencies of a multivariate sum $f \in \mathcal{S}_M^d$ using samples taken on parallel lines. Recall that in the strategy suggested by Theorem 3.21 we consider lines

$$\ell_m = \{(x, m) : x \in \mathbb{R}\},$$

where $m \in \mathbb{Z}^{d-1}$. Then we wish to estimate the set

$$e_1 \cdot Y^f = \bigcup_{m \in G_N^{d-1}} Y^{f|_{\ell_m}}.$$

We let $f_m = f|_{\ell_m} \in \mathcal{S}_M^1$. Note that we expect to find the same frequencies in Y^{f_m} and Y^{f_n} for most $m, n \in \mathbb{Z}^{d-1}$.

Definition 3.26. For $f_1, \dots, f_L \in \mathcal{S}_M^1$ and $N \in \mathbb{N}$ we define their joint signal space of window size N by

$$\text{Sig}(f_1, \dots, f_L, N) := \text{Sig}(f_1, N) + \dots + \text{Sig}(f_L, N).$$

We can immediately transfer the following results from Lemma 3.3.

Lemma 3.27. For $f_1, \dots, f_L \in \mathcal{S}_M^1$ and $N \in \mathbb{N}$ with $N \geq M$ let

$$Y = \bigcup_{j=1}^L Y^{f_j}.$$

Then the following holds true:

- (1) The dimension of $\text{Sig}(f_1, \dots, f_L, N)$ is equal to $|Y|$ and a basis is given by $(v_N(y))_{y \in Y}$.
- (2) A spanning set is given by

$$s_N^{f_k}(j) \quad k = 1, \dots, L, \quad j = j_0, \dots, j_0 + M - 1,$$

where $j_0 \in \mathbb{Z}$ can be arbitrarily chosen.

Proof. Both claims are direct consequences of Lemma 3.3. \square

Now it is quite clear, that the idea of ESPRIT still applies. Indeed, ESPRIT can be interpreted as estimating the parameters y_1, \dots, y_M of a space spanned by $v_N(y_1), \dots, v_N(y_M)$ when $N \geq M + 1$ from any spanning set of this space.

We summarize the proposed algorithm, which uses samples of an unknown $f \in \mathcal{S}_M^d$ on the set $G_N^d = \{n \in \mathbb{Z}^d : \|n\|_\infty \leq N\}$ to determine Y^f .

Algorithm 3.28. Input: $N \in \mathbb{N}$ with $N \geq M$, a tolerance $\text{tol} \geq 0$ and $f(k)$, $k \in G_N$, where $f \in \mathcal{S}_M^d$ is unknown. Let $K = (2N + 1)^{d-1}$.

- If $d = 1$, apply ESPRIT. Otherwise, fix any enumeration j_1, \dots, j_K of G_N^{d-1} , let $f_{j_k}(x) = f(x, j_k)$ and build the matrix

$$H = \begin{bmatrix} H_{N+1, N+1}^{f_{j_1}}(-N) & \dots & H_{N+1, N+1}^{f_{j_K}}(-N) \end{bmatrix}.$$

- Perform an incomplete SVD of H to determine its numerical rank, i.e., the number M_1 of singular values larger than tol . Store the left singular vectors u_1, \dots, u_{M_1} associated with the M_1 largest singular values of H in a matrix $U = [u_1 \dots u_{M_1}]$.
- Form $U^{(1)}$ and $U^{(2)}$ by discarding the last resp. first row of U . Solve

$$U^{(1)}W = U^{(2)}$$

in the (total) least square sense.

- Determine the eigenvalues $e^{2\pi i y_j}$, $j = 1, \dots, M_1$ of W . $Y_1 := \{y_j : j = 1, \dots, M_1\}$ is the estimate of $e_1 \cdot Y^f$.
- For each $y \in Y_1$, determine the coefficients $c_y(j_k)$ of y in the exponential sums f_{j_k} , $k = 1, \dots, K$. Apply this algorithm to $c_y(j_k)$, $j_k \in G_N^{d-1}$ (which are samples of an exponential sum in \mathcal{S}_M^{d-1}) and save the result in Y_y .

Output: $\{(y, z) \in [0, 1) \times [0, 1)^{d-1} : y \in Y_1, z \in Y_y\}$.

Remarks. 1. In the case of noisy data, the algorithm is applied to \tilde{H} , storing the noisy samples. It is then crucial that tol is chosen such that the rank, and hence M is correctly determined. Otherwise, the algorithm will not give any reasonable results. Unfortunately, this holds true for the subsequent applications to $c_y(\cdot)$. But the coefficients are more prone to perturbation, as they arise from solving a perturbed Vandermonde system.

2. It is possible to use a different block size in H , i.e., using $H_{L_1, L_2}^{f_{j_k}}$ blocks, as long as $L_1 \geq M + 1$ and $L_2 \geq M$, exactly as in ESPRIT. To use all samples, one should choose L_1, L_2 such that $L_1 + L_2 = 2N + 2$. Different choices result in slight performance differences.
3. The largest computational cost is due to the (incomplete) SVD of H , which costs $\mathcal{O}_d(N^{d-1}NN^2) = \mathcal{O}_d(N^{d+2})$. Solving for W and determining its eigenvalues is of lower order (as $W \in \mathbb{C}^{M_1 \times M_1}$). When determining the coefficients, it is crucial to note that they arise all from the same system matrix. Indeed, fixing an enumeration y_1, \dots, y_{M_1} of Y_1 , we see that $c(j_k) = (c_{y_1}(j_k), \dots, c_{y_{M_1}}(j_k))^T$ is given by

$$\begin{pmatrix} e^{-N2\pi i y_1} & \dots & e^{-N2\pi i y_{M_1}} \\ e^{-(N-1)2\pi i y_1} & \dots & e^{-(N-1)2\pi i y_{M_1}} \\ \vdots & \ddots & \vdots \\ e^{N2\pi i y_1} & \dots & e^{N2\pi i y_{M_1}} \end{pmatrix} c(j_k) = \begin{pmatrix} f_{j_k}(-N) \\ \vdots \\ f_{j_k}(N) \end{pmatrix}.$$

Thus, one can determine a QR factorization of the matrix (which is just a shifted Vandermonde matrix), resulting in $\mathcal{O}(N^3 + N^{d-1}N^2)$ operations. Furthermore, it is clear that when we apply the algorithm in dimension $\tilde{d} < d$ at most M times (as each application has to give at least one frequency). Thus, overall, we have a complexity of at most $\mathcal{O}_d(N^{d+2})$.

4. When not enough samples to form f_j , $j \in G_N^{d-1}$ are available, one can still use this algorithm, though even in absence of noise a recovery of Y^f cannot be guaranteed. One can for example use a sampling set as suggested in Theorem 3.25 (where f_j , $j = 0, e_1, \dots, e_d$ are available).
5. Numerical examples are given in [26].

Corollary 3.29. For all $f \in \mathcal{S}_M^d$ Algorithm 3.28 recovers Y^f , if one chooses $\text{tol} = 0$.

Proof. We use induction over d . The case $d = 1$ is clear, as we just apply ESPRIT. Let $P_1 : \mathbb{R}^d \rightarrow \mathbb{R}$, $P_1 x = x_1$. Then for $(x_1, \tilde{x}) \in \mathbb{R} \times \mathbb{R}^{d-1}$ we have that

$$f(x_1, \tilde{x}) = \sum_{(y_1, \tilde{y}) \in Y^f} c_{(y_1, \tilde{y})} e^{2\pi i y_1 x_1} e^{2\pi i \tilde{x} \cdot \tilde{y}} = \sum_{y_1 \in P_1 Y^f} \left(\sum_{\tilde{y} \in P_1^{-1}\{y_1\}} c_{(y_1, \tilde{y})} e^{2\pi i \tilde{y} \cdot \tilde{x}} \right) e^{2\pi i x_1 y_1},$$

which is a univariate exponential sum in x_1 . As

$$c_{y_1}(\tilde{x}) = \sum_{\tilde{y} \in P_1^{-1}\{y_1\}} c_{(y_1, \tilde{y})} e^{2\pi i \tilde{y} \cdot \tilde{x}}$$

is in \mathcal{S}_M^{d-1} , it cannot vanish on G_N^{d-1} due to Theorem 3.21 and hence each $y_1 \in P_1 Y^f$ has to appear with a non-zero coefficient in at least one $f(\cdot, j_k)$. Lemma 3.27 implies that $\text{Sig}(f_{j_1}, \dots, f_{j_K}, N)$ is the range of H and that $(v_N(y) : y \in e_1 \cdot Y^f = P_1 Y^f)$ form a basis. Further, $M_1 = |e_1 \cdot Y^f|$ and u_1, \dots, u_{M_1} is a basis of $\text{Sig}(f_{j_1}, \dots, f_{j_K}, N)$ as well. $U^{(1)}W = U^{(2)}$ gives, just as explained around (3.10), a matrix W with eigenvalues $e^{2\pi i y}$, $y \in e_1 \cdot Y^f$.

Using that this algorithm works in dimension $d - 1$ by assumption, that each frequency \tilde{y} of c_{y_1} gives an $(y_1, \tilde{y}) \in Y^f$ and that all $y \in Y^f$ are of this form, the claim is clear. \square

3.3 Other Multivariate Methods

Instead of reducing the problem to a set of one dimensional ones, it is possible to directly transfer Prony's method, MUSIC, ESPRIT and the matrix pencil method to the multivariate setting. In this section we discuss these ideas. We start by defining a multivariate, windowed signal and the signal space.

Definition 3.30. Let $f \in \mathcal{S}_M^d$ and $G \subset \mathbb{Z}^d$ be given. We define the G -windowed signal of f at $j \in \mathbb{Z}^d$ by

$$s_G^f(j) = (f(n + j))_{n \in G}.$$

Further, we call

$$\text{Sig}(f, G) = \text{span}_{j \in \mathbb{Z}^d} \left\{ s_G^f(j) \right\}$$

the signal space of f (with window G).

In Lemma 3.3, we used the linear independence of the Vandermonde vectors to derive several useful properties of the signal space. Similarly, we can use Lemma 3.23 to derive properties of the multivariate signal space, at least if $\Gamma_M^d \subset G$. To give concise statements, it is useful to define a notion of matrices, which do not fix an enumeration of a finite set.

Definition 3.31. For finite sets X, Y , we define a $X \times Y$ matrix $A \in \mathbb{C}^{X \times Y}$ as

$$A = (a_{x,y} : x \in X, y \in Y),$$

where $a_{x,y} \in \mathbb{C}$ for all $(x, y) \in X \times Y$. Of course, such a matrix can be multiplied with a vector $v \in \mathbb{C}^Y$ just like a normal matrix:

$$Av = \left(\sum_{y \in Y} a_{x,y} v_y \right)_{x \in X} \in \mathbb{C}^X.$$

Further, we define $A^T = (a_{x,y})_{y \in Y, x \in X} \in \mathbb{C}^{Y \times X}$ and $A^H = (\overline{a_{x,y}})_{y \in Y, x \in X} \in \mathbb{C}^{Y \times X}$.

Recall the definition of multivariate Vandermonde matrices, which are given by

$$V_G(Y) = (e^{2\pi i n \cdot y})_{n \in G, y \in Y} = [v_G(y) : y \in Y] \in \mathbb{C}^{G \times Y},$$

where $G \subset \mathbb{Z}^d$ and $Y \subset [0, 1)^d$ are finite subsets. Further, we define for any $j \in \mathbb{Z}^d$ the diagonal matrix

$$D_Y(j) = \text{diag}(e^{2\pi i j \cdot y} : y \in Y) \in \mathbb{C}^{Y \times Y}.$$

Lemma 3.32. Let $f \in \mathcal{S}_M^d$ with frequencies Y^f , coefficients $c \in \mathbb{C}^{Y^f}$ and order M be given. Let $j \in \mathbb{Z}^d$. Then the following statements hold true:

(1) We have the identity

$$s_G^f(j) = V_G(Y^f)D_{Y^f}(j)c.$$

(2) If $\Gamma_M^d \subset G$, a spanning set of $\text{Sig}(f, G)$ is given by

$$(s_G^f(j+k))_{k \in \Gamma_M^d}$$

and

$$\dim \text{Sig}(f, G) = M.$$

(3) If $\Gamma_M^d \subset G$, a basis of $\text{Sig}(f, N)$ is given by

$$(v_G(y))_{y \in Y^f}.$$

Proof. The first identity follows easily (and just like in the one dimensional case) from a direct calculation:

$$s_G^f(j) = \sum_{y \in Y^f} c_y (e^{2\pi i(j+n) \cdot y})_{n \in G} = \sum_{y \in Y^f} c_y e^{2\pi i j \cdot y} v_G(y) = V_G(Y^f)D_{Y^f}(j)c.$$

To prove that $\text{Sig}(f, G)$ is spanned by $s_G^f(j+k)$, $k \in \Gamma_M^d$ and has dimension M , it clearly suffices to prove that $D_{Y^f}(k)c$, $k \in \Gamma_M^d$ span a space of dimension M , due to the injectivity of $V_G(Y^f)$ (a consequence of Lemma 3.23 and $\Gamma_M^d \subset G$). But

$$[D_{Y^f}(k)c : k \in \Gamma_M^d] = \text{diag}(c) [v_{\Gamma_M^d}(y)^T : y \in Y^f]^T \quad (3.22)$$

and Lemma 3.23 gives that the transposed Vandermonde matrix on the right-hand side has rank M , which proves the claim.

Now (1) directly implies that $\text{Sig}(f, G) \subset \text{span}(v_G(y) : y \in Y^f)$ and due to Lemma 3.23 the span has dimension M , hence (2) gives the claim. \square

3.3.1 Multivariate Prony-type Methods

Now, proceeding similarly to the section on one dimensional methods, we can use these properties to describe the different approaches to the multivariate frequency estimation problem. We start with Prony's method. Instead of one shift T , we now have one shift for each dimension.

Definition 3.33. For $k = 1, \dots, d$, an $f \in \mathcal{S}_M^d$ and a finite set $G \subset \mathbb{Z}^d$ with $\Gamma_M^d \subset G$, we define the linear map T_k as the unique extension of

$$s_G^f(j) \mapsto s_G^f(j + e_k)$$

to $\text{Sig}(f, G)$. T_k is called k -shift operator. The extension of T_k to \mathbb{C}^G by zero on the orthogonal complement of the signal space is again denoted by T_k .

Again, it is straight forward to check that the eigenvalues of T_k give the frequencies of f . The additional difficulty is that the eigenvalues of T_k are the projection of the frequencies of f onto the subspace $e_k \cdot \mathbb{C}^d$. To match them, we use that T_1, \dots, T_d commute and hence have a common basis of eigenvectors. These eigenvectors induce a matching - one eigenvector corresponds to d eigenvalues $z_j = e^{2\pi i y_j}$ of T_j and gives rise to the frequency vector (y_1, \dots, y_d) .

Lemma 3.34. For $f \in \mathcal{S}_M^d$, a set $G \subset \mathbb{Z}^d$ with $\Gamma_M^d \subset G$ and the k -shift operator T_k defined above, the following statements hold true:

(1) For any left inverse L of $V_G(Y^f)$ we have that

$$T_k = V_G(Y^f)D_{Y^f}(e_k)L,$$

in particular T_k is well-defined.

(2) T_1, \dots, T_d commute.

(3) T_k has eigenvalues $e^{2\pi iy}$ with $y \in e_k \cdot Y^f$. For one such y let $Y = \{\tilde{y} \in Y^f : y = e_k \cdot \tilde{y}\}$. Then the geometric (and algebraic) multiplicity of $e^{2\pi iy}$ is given by $|Y|$ and the eigenvectors are given by $v_G(\tilde{y})$, $\tilde{y} \in Y$.

Proof. The first claim is clear due to Lemma 3.32 (1) and the fact that $D_{Y^f}(e_k)D_{Y^f}(j) = D_{Y^f}(e_k + j)$, the second claim is obvious. Finally, for any $\tilde{y} \in Y$ we have that

$$T_k v_G(\tilde{y}) = V_G(Y^f) D_{Y^f}(e_k) L v_G(\tilde{y}) = V_G(Y^f) D_{Y^f}(e_k) e_{\tilde{y}} = e^{2\pi i \tilde{y} \cdot e_k} V_G(Y^f) e_{\tilde{y}} = e^{2\pi iy} v_G(\tilde{y}).$$

This gives M linearly independent eigenvectors, which is the dimension of $\text{Sig}(f, G)$. \square

Now all we have to do is to find a basis of $\text{Sig}(f, G)$, to represent T_j , $j = 1, \dots, d$ in this basis and to calculate a joint eigenbasis of the eigenspaces of T_j . Then, as described above, we have estimated Y^f . We are thus left with a little bit of linear algebra.

But before we describe the necessary linear algebra, we count the minimal number of sampling points we need. First of all we choose $G = \Gamma_M^d$ (which is the minimal choice). By Lemma 3.32, $s_G^f(k)$, $k \in \Gamma_M^d$ form a spanning set of $\text{Sig}(f, G)$. Therefore, T_j is uniquely determined by $T_j s_{\Gamma_M^d}^f(k)$, $k \in \Gamma_M^d$. But we also need to know $T_j s_{\Gamma_M^d}^f(k) = s_{\Gamma_M^d}^f(k + e_j)$. All combined, we need to know f on $\Gamma_M^d + (\Gamma_M^d + e_j)$ to determine T_j and forming the union over all $j = 1, \dots, d$ gives precisely the sampling set described by Sauer [85].

Definition 3.35. We define the corona of a set $G \subset \mathbb{Z}^d$ by

$$\lceil G \rceil = G \cup \bigcup_{j=1}^d G + e_j.$$

We immediately see that we need samples of f on $\Gamma_M^d + \lceil \Gamma_M^d \rceil$ to estimate T_j , $j = 1, \dots, d$.

We start by estimating $\text{Sig}(f, G)$. Assume that we know $N \in \mathbb{N}_{>0}$, an upper bound of M , the unknown order of f . We then collect a spanning set of the signal space and perform a singular value decomposition. Unfortunately, we have to leave the convenient notation of fixing no enumeration of G and $\lceil \Gamma_N^d \rceil$, as the SVD always fixes an enumeration. Let

$$[s_G^f(n) : n \in \Gamma_N^d] =: H_{G, \Gamma_N^d}^f = U \Sigma W^H.$$

We ignore the slight notational inaccuracy that the left-hand side is a matrix in $\mathbb{C}^{G \times \Gamma_N^d}$, the right-hand side in $\mathbb{C}^{|G| \times |\Gamma_N^d|}$. Estimating the rank of $H_{G, \Gamma_N^d}^f$ by thresholding the singular values at a $\text{tol} > 0$, we obtain M . The first M columns of U , denoted by $u_1, \dots, u_M \in \mathbb{C}^{|G|}$ then form an orthogonal basis of $\text{Sig}(f, G)$. Note that this estimate can also be applied when we only have noisy measurements.

But how to choose tol ? As in the univariate case, H_{G_1, G_2}^f can be factorized in Vandermonde matrices and a diagonal matrix, which then gives rise to estimates of σ_M .

Proposition 3.36. Let $f \in \mathcal{S}_M^d$ with frequencies Y^f , coefficients $c \in \mathbb{C}^{Y^f}$ and order M be given. Further, let $G_1, G_2 \subset \mathbb{Z}^d$ be two finite sets. Then

$$H_{G_1, G_2}^f = V_{G_1}(Y^f) \text{diag}(c_y : y \in Y^f) V_{G_2}(Y^f)^T.$$

Further, if $d = 2$ and $G_j = [-N_j, N_j]^2 \cap \mathbb{Z}^2$, $j = 1, 2$ and $f \in \mathcal{S}^2(q)$ with $q \geq K_j/(N_j + 1)$, where $K_1, K_2, N_1, N_2 \in \mathbb{N}_{>0}$, the smallest non-zero singular value of H_{G_1, G_2}^f can be estimated by

$$\sigma_{\min}^2 \geq c_{\min}^2 \sigma_{\min}^2(V_{G_1}(Y^f)) \sigma_{\min}^2(V_{G_2}(Y^f)) \gtrsim (K_1 K_2 N_1 N_2)^{-2},$$

where c_{\min} is a lower bound to the modulus of the coefficients of f . The precise constants are given in Proposition 2.29.

Proof. The factorization can be derived exactly as in the univariate case, using Lemma 3.32 and (3.22) (which is true for arbitrary finite sets):

$$[s_G^f(n) : n \in G_2] = [V_{G_1}(Y^f) D_{Y^f}(n) c : n \in G_2] = V_{G_1}(Y^f) \text{diag}(c_y : y \in Y^f) V_{G_2}(Y^f)^T.$$

The lower bound for σ^2 follows directly from Proposition 2.29. \square

Remarks. 1. The factorization is well-known in the literature, see for example [85].

2. It is possible to obtain lower bounds for higher dimensions as well, if one relies on Montgomery's construction, see [18] Corollary 22. For fixed q , this results in the following estimate:

$$\sigma_{\min}^2 \gtrsim_q c_{\min}^2 ((2N_1)^d - (2N_1)^{d-1} + \mathcal{O}(N_1^{d-2})) ((2N_2)^d - (2N_2)^{d-1} + \mathcal{O}(N_2^{d-2})).$$

3. Unfortunately, such estimates are unknown for sampling sets of the form $G_1 = \Gamma_N^d$, $G_2 = [\Gamma_N^d]$.
4. The discussion after Theorem 3.7 carries over to the multivariate case with only slight adjustments. Indeed, if we are given $\tilde{f}(n) = f(n) + \varepsilon_n$ and we only know that $|\varepsilon_n| \leq \eta$, we have to choose tol larger than

$$\|E\|_2 \leq \eta \sqrt{|G_1||G_2|},$$

where $E = \{\varepsilon_{n+k} : n \in G_1, k \in G_2\}$ is the matrix containing the noise, to recover $f = M$ from H_{G_1, G_2}^f . However, that is only guaranteed to work if $\sigma_M(H_{G_1, G_2}^f) \geq 2\|E\|_2$. As in the univariate case, more sophisticated estimates, using specific noise models and random matrix theory are currently unknown.

Next we consider the reduced singular value decomposition of $H_{G, \Gamma_N^d}^f$, which for readability is again denoted by $U\Sigma W^H$, but now $\Sigma \in \mathbb{C}^{M \times M}$ is a diagonal matrix with positive, decreasing diagonal entries, $U \in \mathbb{C}^{|G| \times M}$ and $W \in \mathbb{C}^{M \times |\Gamma_N^d|}$ with orthogonal columns.

Now we wish to obtain a matrix representation of T_j in the basis u_1, \dots, u_M . This results in

$$M_j := U^H T_j U = U^H T_j H_{G, \Gamma_N^d}^f W \Sigma^{-1} = U^H H_{G, \Gamma_N^d + e_j}^f W \Sigma^{-1} \in \mathbb{C}^{M \times M}. \quad (3.23)$$

We summarize the algorithm.

Algorithm 3.37 (Multivariate Prony's Method). **Input:** $f(k)$, $k \in G + [\Gamma_N^d]$ of an unknown $f \in \mathcal{S}_M^d$, $N \geq M$ (i.e., an upper bound of the order of f), $G \supset \Gamma_M^d$ and $\text{tol} \geq 0$.

- Calculate an SVD of $H_{G, \Gamma_N^d}^f$, let M be the number of singular values larger than tol . Save the reduced SVD $H_{G, \Gamma_N^d}^f = U\Sigma W^H$.
- Form the matrices M_1, \dots, M_d as in (3.23).
- Calculate a basis of joint eigenvectors v_1, \dots, v_M of M_j and denote the eigenvalue of M_j and v_k by $e^{2\pi i y_j^k}$.

Output: The frequency vectors (y_1^k, \dots, y_d^k) , $k = 1, \dots, M$.

How to calculate a basis of eigenvectors is well-known, see for example [33]. To obtain a joint eigenbasis, we usually do not need to calculate an eigenvalue decomposition of all M_j . Indeed, as the eigenspaces of one M_j are invariant under all other M_k , we simply start with an eigenspace decomposition of M_1 . All one dimensional eigenspaces are then also eigenspaces of all other matrices. For all higher dimensional eigenspaces, we proceed as follows: Let E be such an eigenspace of M_1 . We then perform an eigenspace decomposition of $M_2|_E : E \rightarrow E$. For all eigenspaces of M_2 in E with a dimension larger than one we continue by decomposing M_3 on each of them et cetera. See [63] and [86], where this idea is made precise.

An alternative approach is to simply form a linear combination M of all matrices M_j . Clearly, such a linear combination has still the same eigenvector basis. The eigenvalues of M are then the corresponding linear combination of the eigenvalues of M_j . For all but a finite number of linear combinations, the eigenvalues of M will only have one dimensional eigenspaces (an easy consequence of Lemma 3.34) and we are done. This strategy has been pursued in [81, 95]. Later, in our numerical experiments, we will use it as well.

Remarks. 1. It is easily possible to use a larger sampling set, for example $G + [G_2]$. As long as $\Gamma_M^d \subset G \cap G_2$ the method will work.

2. If one uses noiseless samples and chooses $\text{tol} = 0$, the preceding discussion shows that the algorithm will recover Y^f .
3. Concerning the computational complexity, calculating the SVD is the most costly step, namely $\mathcal{O}(|G||\Gamma_N^d|^2)$. Using the minimal set $G = \Gamma_M^d$, this results in $\mathcal{O}_d(N^3)$, up to logarithmic terms.
4. In the case $G = [-N_1, N_1]^d \cap \mathbb{Z}^d$ and $G_2 = [-N_2, N_2]^d \cap \mathbb{Z}^d$ Algorithm 3.37 is actually the same as the algorithm proposed in [40]. However, not only is the number of samples increased to $\mathcal{O}_d(N^d)$, the computational complexity increases drastically to $\mathcal{O}_d(N^{3d})$.

This variation of Prony's method is closely related to the method introduced by Sauer in [84]. Indeed, we claim that the matrices M_j , defined in (3.23) are similar to the transposed multiplication tables, as given in [84], Theorem 5. We now give a short reasoning for this claim.

To this end, we note that each T_j can be seen as a difference equation, as it gives rise to an equation of the form

$$\sum_{n \in G} t_{m,n}^{(j)} f(n+k) = f(m+k+e_j) \quad \text{for all } m \in G \text{ and all } k \in \mathbb{Z}^d$$

with $t_{m,n}^{(j)} \in \mathbb{C}$. In the one dimensional case, we were able to identify these coefficients with a polynomial, with roots equal to the frequencies. To see whether this is still possible, we consider

$$P_{m,j}(z) = z^{m+e_j} - \sum_{n \in G} t_{m,n}^{(j)} z^n \in \Pi_{G \cup \{m+e_j\}}.$$

Now we see that

$$\begin{aligned} 0 &= f(m+k+e_j) - \sum_{n \in G} t_{m,n}^{(j)} f(n+k) = \sum_{y \in Y^f} c_y e^{2\pi i y \cdot k} \left(e^{2\pi i y \cdot (m+e_j)} - \sum_{n \in G} t_{m,n}^{(j)} e^{2\pi i y \cdot n} \right) \\ &= \sum_{y \in Y^f} c_y e^{2\pi i y \cdot k} P_{m,j}(e^{2\pi i y}). \end{aligned}$$

As this equation holds for all $k \in \mathbb{Z}^d$, we conclude that $P_{m,j}(e^{2\pi i y}) = 0$ for all $y \in Y^f$. Further, note that for all $m+e_j \notin G$, clearly $P_{m,j} \neq 0$ and all these $P_{m,j}$ are linearly independent.

Denote the vanishing ideal of Y^f in the polynomial ring Π in d variables by

$$I_{Y^f} = \{p \in \Pi : p(y) = 0 \text{ for all } y \in Y^f\}.$$

Further, we denote by $[p]$ the equivalence class of $p \in \Pi$ modulo I_{Y^f} . We just proved that $P_{m,j} \in I_{Y^f}$, i.e., $[P_{m,j}] = 0$. If we identify \mathbb{C}^G with Π_G by

$$c \in \mathbb{C}^G \leftrightarrow \sum_{k \in G} c_k z^k \in \Pi_G,$$

we can define

$$\tilde{T}_j : \Pi_G \rightarrow \Pi/I_{Y^f}, \quad p \mapsto [T_j^T p].$$

Next we claim that $I_{Y^f} \cap \Pi_G$ is in the kernel of \tilde{T}_j . To prove this, we choose an arbitrary polynomial $p = (p_n)_{n \in G} \in I_{Y^f}$ and calculate (using $[P_{m,j}] = 0$)

$$\begin{aligned} [T_j^T p] &= \left[\sum_{n \in G} \left(\sum_{m \in G} t_{m,n}^{(j)} p_m \right) z^n \right] = \left[\sum_{m \in G} p_m \left(P_{m,j}(z) + \sum_{n \in G} t_{m,n}^{(j)} z^n \right) \right] \\ &= \left[\sum_{m \in G} z_j p_m z^m \right] = [z_j p(z)] = 0. \end{aligned}$$

For notational convenience, we continue to use \tilde{T}_j for the mapping $\Pi_G / (I_{Y^f} \cap \Pi_G) \rightarrow \Pi/I_{Y^f}$ induced by \tilde{T}_j .

Furthermore, by interpolation on Y^f , one can construct a mapping

$$\pi : \Pi \rightarrow \Pi_G / (I_{Y^f} \cap \Pi_G).$$

Note that while interpolating on Y^f is possible, the interpolant is not uniquely defined (see discussion after Lemma 3.23). However, it is unique modulo I_{Y^f} . π contains I_{Y^f} in its kernel, as it is constructed by interpolation.

Hence, we obtain a mapping (which is a linear mapping between vector spaces)

$$\tilde{T}_j \circ \pi : \Pi / I_{Y^f} \rightarrow \Pi / I_{Y^f}.$$

We claim that this mapping is actually equal to $[p] \mapsto [z_j p]$, the multiplication with the j -th variable. It suffices to check this on $[z^m]$, $m \in \Gamma_M^d$, which is a spanning set. This, however, can be verified by a quick calculation:

$$\tilde{T}_j \circ \pi([z^m]) = \tilde{T}_j([z^m]) = \left[\sum_{n \in G} t_{m,n}^{(j)} z^n \right] = \left[\sum_{n \in G} t_{m,n}^{(j)} z^n + P_{m,j} \right] = [z^{m+e_j}] = [z_j z^m],$$

where we used that $[P_{m,j}] = 0$. It is interesting to note that in the one dimensional case, the shift operator T can be represented by the companion matrix (which represents $[p] \mapsto [z p]$ in Π / I_{Y^f}). Analogously, we just showed that T_j^T represents multiplication by the j th variable. A matrix representation of $\tilde{T}_j \circ \pi$ can therefore be interpreted as a higher dimensional analog of the companion matrix. Finally, M_j can be seen as a matrix representing T_j in a suitable orthonormal basis. We summarize:

Proposition 3.38. *The matrices M_j^T , $j = 1, \dots, d$ as given in (3.23) represent the linear mappings*

$$\begin{aligned} \Pi / I_{Y^f} &\rightarrow \Pi / I_{Y^f} \\ [p] &\mapsto [z_j p]. \end{aligned}$$

While this theorem shows that Algorithm 3.37 and the method proposed by Sauer are closely related, there are some important differences. Most importantly, Algorithm 3.37 starts with an estimate of the signal space, using as many samples as possible. Then projecting onto the estimated signal space (hopefully) clears most of the noise from the data. This gives a rather stable algorithm. On the other hand, the method introduced by Sauer uses a nested set of sampling points $G + A_0 \subset G + A_1 \subset \dots \subset G + [\Gamma_M^d]$ and terminates at step k if $G + A_k$ suffices to recover f . It is reasonable to assume that this leads to a method more prone to noise, as not always all samples are used. However, as many $f \in \mathcal{S}_M^d$ can be estimated with fewer samples (in particular if the bound N of the order is crude), only using as many samples as necessary has computational advantages.

Another slightly different perspective is given in [40] by Harmouch, Khalil and Mourrain. Proposition 3.38 is closely related to Proposition 4.1 in [40], which covers only the case $G = [-N_1, N_1]^d \cap \mathbb{Z}^d$ and $G_2 = [-N_2, N_2]^d \cap \mathbb{Z}^d$, though a more general problem. They translate the problem to a polynomial setting (along the lines of the sketch given above) and then deduce Algorithm 3.37 for this case.

Finally, we remark that the first result rephrasing the multivariate Prony problem as finding the joint zeros of a finite number of polynomials is given by Kunis et al. in [53]. However, no method to actually compute these zeros is given. In our version, we circumvent this formulation entirely.

3.3.2 Multivariate MUSIC and ESPRIT-type Algorithms

Now we give a variation of MUSIC in arbitrary dimensions. Our method is more or less equivalent to the method presented in [54]. However, we are able to reduce the number of necessary samples significantly, from $\mathcal{O}_d(M^d)$ to $\mathcal{O}_d(M^2)$ up to logarithmic factors. This brings it on a par with the preceding multivariate Prony.

It is actually quite clear how the univariate MUSIC can be extended to the multivariate case. For any $G \supset \Gamma_{M+1}^d$, we know that

$$\forall y \in [0, 1]^d : v_G(y) \in \text{Sig}(f, G) \Leftrightarrow y \in Y^f$$

by Lemma 3.23. Indeed, we start again by calculating an orthonormal basis of $\text{Sig}(f, G)$ by employing once again the tool of our choice, the SVD. Let

$$H_{G, \Gamma_N^d}^f = U \Sigma W^H$$

and $U = [U_1 U_2]$, where U_1 are the left singular vectors associated with the non-zero singular values and U_2 the other ones. We modify the MUSIC pseudospectrum:

$$R(y) = \frac{1}{\|U_2 v_G(y)\|_2^2}.$$

Clearly, $R(y)$ has peaks exactly at $y \in Y^f$. We summarize the algorithm.

Algorithm 3.39 (Multivariate MUSIC). *Input:* Samples $f(n+k)$, $n \in G$, $k \in \Gamma_N^d$ of an unknown $f \in \mathcal{S}_M^d$ with $N \geq M$ and $G \supset \Gamma_{N+1}^d$, $\text{tol} \geq 0$.

- Form the matrix $H_{G, \Gamma_N^d}^f$ and perform a SVD. Estimate the rank of $H_{G, \Gamma_N^d}^f$ by counting the number of singular values larger than tol . Let \tilde{M} be this estimate.
- Form U_2 from all left singular vectors associated with singular values smaller or equal to tol .
- Compute $R(y)$ on $[0, 1]^d$.

Output: The \tilde{M} largest local maxima of $R(y)$.

This is almost the exact algorithm presented by Liao in [54]. However, due to the theory developed in this chapter the number of necessary samples is significantly reduced for all dimensions larger than two. Indeed, Liao uses samples of f taken on $\{n \in \mathbb{N}_0^d : \|n\|_\infty \leq 2N\}$, i.e., at least $\mathcal{O}_d(M^d)$ samples, while we only need $\mathcal{O}_d(M^2)$ samples (again, up to logarithmic factors) when we choose the minimal $G = \Gamma_{M+1}^d$.

Now we discuss a multivariate extension of the ESPRIT algorithm. Multivariate matrix pencil and ESPRIT algorithms have been discussed since the early nineties, one of the first method introduced is the Matrix Enhancement and Matrix Pencil method (MEMP) by Hua [44], where the pairing procedure of the components is performed by a MUSIC-like criterion. Using a joint diagonalization of two matrices, as in the Prony scheme given above, was then proposed in [95] and extended in [81]. Another possibility, an extension of the Unitary ESPRIT, is available if one wishes to find a joint diagonalization of two real matrices M_1, M_2 by diagonalizing $M_1 + iM_2$. This has been proposed in [39].

Most of these methods are given in the two dimensional case but can be extended to the multivariate case quite easily. However, this usually results in sampling sets of the form $[0, N]^d \cap \mathbb{Z}^d$. One example is [38], which includes additional references and applications. Recently, Andersson and Carlsson stated ESPRIT in arbitrary dimensions on general grids in [2], building on their previous work [3]. In their work, they exclude degenerate cases (basically assuming that the corresponding Vandermonde matrices have full rank), which also makes the joint eigenspace decomposition trivial (as an eigenbasis of one matrix is an eigenbasis for all matrices). However, combined with Lemma 3.23, their results are equivalent to the ESPRIT algorithm we present next.

We give the basic idea of multivariate ESPRIT. Again, we can profit from our general results for the multivariate signal space and reduce the necessary sampling set to be of order $\mathcal{O}_d(M^2)$ up to logarithmic terms.

We start by echoing the rotation property of the Vandermonde basis in the higher dimensional setting. While in the one dimensional case we simply removed the last and first row of the Vandermonde matrix, we now select submatrices by considering (for $j = 1, \dots, d$)

$$E_1 = (\delta_{x,y} : x \in G, y \in [G]) \in \mathbb{C}^{G \times [G]}, \quad E_2^{(j)} = (\delta_{x+e_j, y} : x \in G, y \in [G]) \in \mathbb{C}^{G \times [G]}.$$

With the diagonal matrices $D_j = \text{diag}(e^{2\pi i y \cdot e_j}, y \in Y^f)$ we obtain

$$E_1 V_{[G]}(Y^f) D_j = E_2^{(j)} V_{[G]}(Y^f). \quad (3.24)$$

This again transfers to other bases of the signal space $\text{Sig}(f, [G])$. Indeed, if $B = [b_1, \dots, b_M]$ is a matrix with columns $b_j \in \mathbb{C}^{[G]}$ forming a basis of the signal space, there is an invertible matrix

$S \in \mathbb{C}^{M \times M}$ (now fixing an ordering of Y^f) such that

$$V_{[G]}(Y^f) = BS.$$

For this basis, equation (3.24) implies that

$$E_1 B S D_j S^{-1} = E_2^{(j)} B.$$

As $E_1 V_{[G]}(Y^f) = V_G(Y^f)$ has rank M , provided that $\Gamma_M^d \subset G$, so has $E_1 B$ and therefore we can solve this equation for $S D_j S^{-1}$. Clearly, for all j the matrices $S D_j S^{-1}$ commute and are diagonalizable. A joint eigenbasis is given by the columns s_1, \dots, s_M of S . As in the multivariate Prony algorithm, such an eigenvector s_j corresponds to the eigenvalues $e^{2\pi i y \cdot e_j}$ of the matrices $S D_j S^{-1}$ and therefore gives rise to a $y \in Y^f$.

Algorithm 3.40 (Multivariate ESPRIT). **Input:** Samples $f(n+k)$, $n \in [G]$, $k \in \Gamma_N^d$ of an unknown $f \in \mathcal{S}_M^d$ with $N \geq M$ and $G \supset \Gamma_N^d$, $\text{tol} \geq 0$.

- Form the matrix $H_{[G], \Gamma_N^d}^f$ and perform a SVD. Estimate the rank of $H_{[G], \Gamma_N^d}^f$ by counting the number of singular values larger than tol . Let \tilde{M} be this estimate.
- Form U from all left singular vectors associated with singular values larger than tol .
- For $j = 1, \dots, d$, solve

$$E_1 U N_j = E_2^{(j)} U$$

in the least (or total least) square sense.

- Calculate a basis of joint eigenvectors $v_1, \dots, v_{\tilde{M}}$ of N_j and denote the eigenvalue of N_j and v_k by $e^{2\pi i y_j^k}$.

Output: The frequency vectors (y_1^k, \dots, y_d^k) , $k = 1, \dots, \tilde{M}$.

Remarks. 1. Similar to Prony and MUSIC, we only need $\mathcal{O}_d(M^2)$ samples (up to logarithmic factors), a significant improvement to former ESPRIT-type methods, which used $\mathcal{O}_d(M^d)$ samples. Furthermore, the same algorithm works if one uses $H_{[G], G_2}^f$ as long as $\Gamma_M^d \subset G_2$.

2. The computational costs are again dominated by the singular value decomposition and of order $\mathcal{O}(|G| |\Gamma_N^d|^2)$, i.e., at least $\mathcal{O}_d(N^3)$ (up to logarithmic factors).
3. Again, we use joint eigenvectors of commuting matrices to match the frequency components.
4. As in the one dimensional case, it is possible to derive a similar algorithm based on matrix pencils. Due to the repetitive nature, we refrain from giving the details here.
5. Generically, $|G|$ Vandermonde vectors of length $|G|$ are linearly independent on \mathbb{C}^G . Therefore, using ESPRIT on $H_{[G], G_2}^f$ with $|G|, |G_2| \geq M$ is expected to recover f . Even more, one actually only has to have full rank of $E_1 U$, which is expected to hold generically when

$$M \leq |(G + e_2) \cup \dots \cup (G + e_d)|.$$

Details can be found in [2].

There is a close relation to the multivariate Prony method and the multivariate ESPRIT algorithm. Indeed, M_j^T is similar to N_j (in the noise free case) if we choose $G = \Gamma_N^d$ in both cases. Note that Proposition 3.38 implies that N_j is actually similar to a multiplication table.

To give a few more details, recall that $H_{\Gamma_N^d, \Gamma_N^d}^f = U \Sigma W^H$ is the reduced SVD and therefore

$$M_j = U^H H_{\Gamma_N^d, \Gamma_N^d + e_j}^f W \Sigma^{-1} U^H U = U^H H_{\Gamma_N^d, \Gamma_N^d + e_j}^f \left(H_{\Gamma_N^d, \Gamma_N^d}^f \right)^\dagger U.$$

On the other hand, we have

$$\begin{aligned} E_1 H_{[\Gamma_N^d], \Gamma_N^d}^f &= H_{\Gamma_N^d, \Gamma_N^d}^f, \\ E_2^{(j)} H_{[\Gamma_N^d], \Gamma_N^d}^f &= H_{\Gamma_N^d + e_j, \Gamma_N^d}^f. \end{aligned}$$

Now given a reduced SVD $H_{[\Gamma_N^d], \Gamma_N^d}^f = \hat{U} \hat{\Sigma} \hat{W}^H$, we see that

$$E_1 \hat{U} \hat{\Sigma} \hat{W}^H \hat{W} \hat{\Sigma}^{-1} N_j = E_2^{(j)} \hat{U} \hat{\Sigma} \hat{W}^H \hat{W} \hat{\Sigma}^{-1}$$

and therefore

$$H_{\Gamma_N^d, \Gamma_N^d}^f \hat{W} \hat{\Sigma}^{-1} N_j = H_{\Gamma_N^d + e_j, \Gamma_N^d}^f \hat{W} \hat{\Sigma}^{-1}.$$

This gives (using that the columns of \hat{W} are an orthonormal basis of the orthogonal complement of the kernel of $H_{\Gamma_N^d, \Gamma_N^d}^f$ and that the range of $H_{\Gamma_N^d, \Gamma_N^d}^f$ is equal to the range of $H_{\Gamma_N^d + e_j, \Gamma_N^d}^f$, as both are the signal space $\text{Sig}(f, \Gamma_N^d)$)

$$N_j = \hat{\Sigma} \hat{W}^H \left(H_{\Gamma_N^d, \Gamma_N^d}^f \right)^\dagger H_{\Gamma_N^d + e_j, \Gamma_N^d}^f \hat{W} \hat{\Sigma}^{-1}.$$

Now we see clearly that N_j and M_j^T are similar, as $(H_{\Gamma_N^d, \Gamma_N^d + e_j}^f)^T = H_{\Gamma_N^d + e_j, \Gamma_N^d}^f$.

Does this give any indication to what method might be preferable? \hat{W} is estimated using $H_{[\Gamma_N^d], \Gamma_N^d}^f$ while U is estimated using $H_{\Gamma_N^d, \Gamma_N^d}^f$, a matrix with significantly less samples. Therefore, one might hope that in presents of noise using more samples stabilizes the scheme and that the multivariate ESPRIT is preferable. We give numerical evidence that this is indeed the case in the next section.

3.3.3 Numerical Examples

We now give numerical examples to illustrate the performance of the proposed algorithms, in particular of the multivariate Prony method, Algorithm 3.37 and of the multivariate ESPRIT Algorithm 3.40. Again, we use the Hausdorff distance between the frequencies of f and the estimated frequencies Y_{est} as an error measure, given by

$$d_H(Y^f, Y_{est}) = \max \left\{ \max_{y \in Y^f} \min_{\tilde{y} \in Y_{est}} \|y - \tilde{y}\|_2, \max_{\tilde{y} \in Y_{est}} \min_{y \in Y^f} \|y - \tilde{y}\|_2 \right\}$$

as well as

$$d_2(Y^f, Y_{est})^2 = \max \left\{ \sum_{y \in Y^f} \text{dist}(y, Y_{est})^2, \sum_{y \in Y_{est}} \text{dist}(y, Y^f)^2 \right\}.$$

However, if $|Y^f| \neq |Y_{est}|$ we consider the recovery failed.

We consider noisy samples

$$\tilde{f}(k) = f(k) + n_{k,1} + in_{k,2},$$

where $n_{k,j}$ are independent and uniformly distributed in $\delta[-0.5, 0.5]$ for multiple choices of $\delta \in \mathbb{R}_{>0}$.

Whenever M_1, \dots, M_d are a family of pairwise commuting and diagonalizable matrices and we have perturbed $\tilde{M}_1, \dots, \tilde{M}_j$, finding a joint eigenbasis of $\tilde{M}_1, \dots, \tilde{M}_d$ is in general impossible, as \tilde{M}_j are not expected to be commuting. This seems to be a problem, as finding matrices N_1, \dots, N_d which are close to $\tilde{M}_1, \dots, \tilde{M}_d$ and commute is a difficult problem.

We use the following easy strategy. We choose $\alpha_1, \dots, \alpha_d \in \mathbb{R}_{>0}$ and form the matrix

$$\sum_{j=1}^d \alpha_j \tilde{M}_j.$$

Table 3.2: Results of the first example

δ	N	K	tol	d_2^P	d_H^P	Prony fails/100	d_2^E	d_H^E	ESPRIT fails/100
1e-4	5	5	0.1	1.6e-4	1.4e-4	1	7.7e-5	6.6e-5	0
1e-1	5	5	1	3.5e-2	3.1e-2	30	2.6e-2	2.1e-2	14
1e-1	15	5	1	8.6e-3	7.2e-3	0	7.5e-3	6.4e-3	0
1e-0	25	5	10	5.3e-2	4.4e-2	1	5.1e-2	4.4e-2	0
1e-2	10	10	1	4.1e-3	3.5e-3	7	2.1e-3	1.6e-3	6
1e-2	20	20	1	6.8e-3	5.5e-3	0	1.7e-3	1.3e-3	0
5e-1	20	20	5	1.0e-1	8.4e-2	24	5.6e-2	3.9e-2	9
5e-1	50	50	10	1.1e-1	6.7e-2	9	3.9e-2	2.6e-2	6

Table 3.3: ESPRIT, random frequencies in three dimensions

δ	N	K	tol	d_2^E	d_H^E	fails/100
1e-4	5	5	1	2.5e-5	2.1e-5	0
1e-1	5	5	1	2.4e-2	2.1e-2	1
1e-0	15	5	10	7.0e-2	5.8e-2	2
1e-1	15	25	10	1.6e-2	1.1e-2	3
1e-1	50	50	10	7.6e-3	5.3e-3	0
1e-1	50	100	10	8.2e-3	5.5e-3	0

Then we determine an eigenbasis v_1, \dots, v_M of this matrix. Each v_j gives rise to a frequency y_j : y_{jk} is given by the average of $(\tilde{M}_k v_j)_n / v_{jn}$ over all $n = 1, \dots, M$ with $v_{jn} \neq 0$.

While one might argue that this is a naive method, we will see that it works reasonably well. We choose α_j randomly, namely as independent, uniformly distributed in $[1, 2]$.

As a first example, we consider K frequencies chosen randomly in $[0, 1]^2$ and coefficients given by $r_j e^{2\pi i \phi_j}$, where r_j is chosen randomly in $[0.2, 1.2]$ and ϕ_j in $[0, 1]$, all with respect to the uniform distribution. We use the sampling set $[\Gamma_N^d] + \Gamma_N^d$. For the rank estimation we use tol for both algorithms. All examples were run 100 times. The results are collected in Table 3.2, where the average errors of the Prony methods are denoted by d_2^P and d_H^P while the average errors of ESPRIT are denoted by d_2^E and d_H^E .

The algorithms were implemented in Matlab 2016b.

Clearly, the multivariate ESPRIT algorithm is a little bit more stable than the multivariate Prony method. This confirms the reasoning we gave earlier. Furthermore, we see that additional samples stabilize the problem significantly.

Next, we consider randomly generated frequencies in three dimensions. As ESPRIT already proved superior to Prony, we stick with it. However, we switch to an implementation in Julia 0.6.2, which is much quicker (at least in the author's implementation).

We give the results in Table 3.3. As we use randomly chosen frequencies, we can usually recover more than N frequencies, as the table shows. For $N = 50$ we already use 29053 samples and for $K = 100$, the computational time of the 100 repetitions on a standard computer (with a AMD Ryzen 1700 processor) was roughly three minutes.

We report results for higher dimensions in Table 3.4. ESPRIT can easily and efficiently be used even in 12 dimensions, where the choice $N = 8$ results in 204418 samples. $K = 20$ results in a computational time of approximately ten minutes, while $K = 250$ took approximately 25 minutes. If we had to use samples on $[-8, 8]^{12} \cap \mathbb{Z}^{12}$, we had more than 10^{15} samples and any algorithm would be completely unfeasible.

In all examples the value of tol is chosen by considering a few instances of $\sigma_{\text{ord } f}(H_{[\Gamma_N^d], \Gamma_N^d}^f)$ and then guessing a reasonable lower bound. That works quite well, as the examples show.

Error Estimates

Next we give examples of the error estimates presented in the previous chapter. We start with a randomized example. As we have to generate well-separated frequencies, we pick sixteen frequency vectors in $[0, 0.1]^2$ and translate them by $0.25(k_1, k_2)$, $k_1, k_2 = 0, \dots, 3$. Again, the coefficients are given by $r_j e^{2\pi i \phi_j}$, where r_j is chosen randomly in $[0.2, 1.2]$ and ϕ_j in $[0, 1]$, all with respect to the

Table 3.4: ESPRIT, random frequencies in higher dimensions

d	δ	N	K	tol	d_2^E	d_H^E	fails/100
4	1e-0	10	20	10	0.19	0.12	1
4	1e-0	10	50	10	0.44	0.21	3
5	1e-0	20	20	100	9.0e-2	6.4e-2	0
5	1e-1	20	200	10	2.1e-2	2.1e-2	0
8	1e-0	10	20	100	0.16	0.12	0
8	1e-1	10	200	10	6.8e-2	4.3e-2	0
12	1e-1	8	20	100	2.0e-2	2.6e-2	0
12	1e-1	8	250	100	4.4e-2	2.6e-2	0

uniform distribution.

To agree with our analysis in Corollary 2.32, we perturb the samples of a randomly generated f^*

$$\tilde{f}(n) = f^*(n) + X_{n,1} + iX_{n,2},$$

where all $X_{n,j} \sim \mathcal{N}(0, \sigma)$ are independent random variables. We then use ESPRIT on the sampling set $[\Gamma_{25}^2] + \Gamma_{25}^2$ to get an estimate for the frequencies. To test the error estimate, we calculate new samples on $[-30, 30]^2 \cap \mathbb{Z}^2$. We call

$$\left(\|\mathcal{P}_N(f) - \tilde{s}_N(f^*)\|_2^2 - 2\sigma^2(2N+1)^2 \right)^{\frac{1}{2}} + (2 + \sqrt{2})\sigma(2N+1)^{(2+\delta)/4}$$

the error estimator, which is an upper bound to

$$\frac{15}{16}(N+1)^4 \sum_{y \in Y^f} (|c_y^f|^2 + |c_{n(y)}^g|^2) \|y - n(y)\|_{\mathbb{T}^2}^2 + \frac{3(N+1)^2}{4} \sum_{y \in Y^f} |c_y^f - c_{n(y)}^g|^2.$$

Here, the first summand is the frequency part, the second the coefficient part of the error. We choose $\delta = 0.9$, which results in a probability of at least 99% that our bound is applicable.

Carrying out this procedure for different values of the standard deviation σ 25 times and averaging gives Figure 3.4.

As is evident from this picture, the estimator is within two orders of magnitudes of the error. Furthermore, the asymptotic behavior in σ seems to be correctly captured. Note that we cannot hope for a perfect result, as the estimation captures the worst case.

To test how sharp Theorem 2.31 is, we consider two deterministic exponential sums. Let f be the exponential sum with frequencies

$$Y^f = \{(n, m)/7 : n, m = 0, 1, \dots, 6\}$$

and coefficients $c_{(n,m)} = (-1)^{n+m}$. As a second exponential sum, we use g_λ , which has frequencies

$$Y^{g_\lambda} = \{y + (\lambda, \lambda) : y \in Y^f\}.$$

The coefficients of g_λ are determined to minimize the least squares error $\mathcal{P}_N(f - g_\lambda)$ for $N = 30$. Now we let λ vary and compare the least squares error with the total error, the latter again split into its frequency and its coefficient part.

The example shows that while the qualitative behavior in λ is well captured, we are again off by a constant of order 10^{-2} . That might be an indication that an improvement of the involved constants is possible.




Figure 3.4: Comparison of the errors at different noise levels. x -axis: Standard deviation of the noise. All results are averaged over 25 runs.




Figure 3.5: Comparison of the errors for different g_λ .

Bibliography

- [1] AKINSHIN, A., BATENKOV, D., AND YOMDIN, Y. Accuracy of spike-train Fourier reconstruction for colliding nodes. In *Sampling Theory and Applications (SampTA), 2015 International Conference on* (2015), IEEE, pp. 617–621.
- [2] ANDERSSON, F., AND CARLSSON, M. ESPRIT for multidimensional general grids. *arXiv preprint arXiv:1705.07892* (2017).
- [3] ANDERSSON, F., AND CARLSSON, M. On the structure of positive semi-definite finite rank general domain Hankel and Toeplitz operators in several variables. *Complex Analysis and Operator Theory* 11, 4 (2017), 755–784.
- [4] AUBEL, C., AND BÖLCSKEI, H. Vandermonde matrices with nodes in the unit disk and the large sieve. *Applied and Computational Harmonic Analysis* (2017).
- [5] BALAN, R. Stability theorems for Fourier frames and wavelet Riesz bases. *Journal of Fourier Analysis and Applications* 3, 5 (1997), 499–504.
- [6] BATENKOV, D. Stability and super-resolution of generalized spike recovery. *Applied and Computational Harmonic Analysis* (2016).
- [7] BATENKOV, D., AND YOMDIN, Y. On the accuracy of solving confluent Prony systems. *SIAM Journal on Applied Mathematics* 73, 1 (2013), 134–154.
- [8] BERCU, B., DELYON, B., AND RIO, E. *Concentration inequalities for sums and martingales*. Springer, 2015.
- [9] BEURLING, A. Sur les intégrales de fourier absolument convergentes et leur application à une transformation fonctionnelle. In *Ninth Scandinavian Mathematical Congress* (1938), pp. 345–366.
- [10] BOAS, R. P. *Entire functions*. Academic Press, 1954.
- [11] BROWN, J., AND CABRERA, S. On well-posedness of the Papoulis generalized sampling expansion. *IEEE Transactions on Circuits and Systems* 38, 5 (1991), 554–556.
- [12] BROWN, J., AND SA-NGSARI, K. Sampling reconstruction of n-dimensional band-limited images after multilinear filtering. *IEEE Transactions on Circuits and Systems* 36, 7 (1989), 1035–1038.
- [13] BRYC, W., DEMBO, A., AND JIANG, T. Spectral measure of large random Hankel, Markov and Toeplitz matrices. *The Annals of Probability* (2006), 1–38.
- [14] BUHMANN, M. D., AND PINKUS, A. On a recovery problem. *Annals of Numerical Mathematics* 4 (1997), 129–142.
- [15] CANDÈS, E. J., AND FERNANDEZ-GRANDA, C. Super-resolution from noisy data. *Journal of Fourier Analysis and Applications* 19, 6 (2013), 1229–1254.
- [16] CANDÈS, E. J., AND FERNANDEZ-GRANDA, C. Towards a mathematical theory of super-resolution. *Communications on Pure and Applied Mathematics* 67, 6 (2014), 906–956.
- [17] CANDÈS, E. J., AND TAO, T. Decoding by linear programming. *IEEE transactions on information theory* 51, 12 (2005), 4203–4215.
- [18] CARRUTH, J., GONÇALVES, F., AND KELLY, M. The Beurling-Selberg box minorant problem. *arXiv preprint arXiv:1702.04579* (2017).

- [19] CHEUNG, K. F. A multidimensional extension of Papoulis' generalized sampling expansion with the application in minimum density sampling. In *Advanced Topics in Shannon Sampling and Interpolation Theory*. Springer, 1993, pp. 85–119.
- [20] CHRISTENSEN, O. Perturbation of frames and applications to Gabor frames. In *Gabor analysis and algorithms*. Springer, 1998, pp. 193–209.
- [21] CHRISTENSEN, O. *An introduction to frames and Riesz bases*. Springer, 2016.
- [22] CUYT, A., AND LEE, W.-s. Multivariate exponential analysis from the minimal number of samples. *arXiv preprint arXiv:1610.06329* (2016).
- [23] DE PRONY, B. G. R. Essai expérimental et analytique: sur les lois de la dilatabilité de fluides élastique et sur celles de la force expansive de la vapeur de l'alkool, a différentes températures. *Journal de l'école polytechnique 1*, 22 (1795), 24–76.
- [24] DIEDERICHS, B. Parameter estimation for bivariate exponential sums by Prony-like methods. Master thesis, University of Hamburg, 2014.
- [25] DIEDERICHS, B., AND ISKE, A. Parameter estimation for bivariate exponential sums. In *Sampling Theory and Applications (SampTA), 2015 International Conference on* (2015), IEEE, pp. 493–497.
- [26] DIEDERICHS, B., AND ISKE, A. Projection-based multivariate frequency estimation. In *Sampling Theory and Applications (SampTA), 2017 International Conference on* (2017), IEEE, pp. 360–363.
- [27] DONOHO, D. L. Superresolution via sparsity constraints. *SIAM journal on mathematical analysis 23*, 5 (1992), 1309–1331.
- [28] DUVAL, V., AND PEYRÉ, G. Exact support recovery for sparse spikes deconvolution. *Foundations of Computational Mathematics 15*, 5 (2015), 1315–1355.
- [29] ELДАР, Y. C. *Sampling theory: Beyond bandlimited systems*. Cambridge University Press, 2015.
- [30] FERNANDEZ-GRANDA, C. Support detection in super-resolution. In *Sampling Theory and Applications (SampTA), 2013 International Conference on* (2013), pp. 145–148.
- [31] FILBIR, F., MHASKAR, H., AND PRESTIN, J. On the problem of parameter estimation in exponential sums. *Constructive Approximation 35*, 3 (2012), 323–343.
- [32] FOU CART, S., AND RAUHUT, H. *A mathematical introduction to compressive sensing*. Springer, 2013.
- [33] GOLUB, G. H., AND VAN LOAN, C. F. *Matrix computations*, 4. edition ed. JHU Press, 2013.
- [34] GRAHAM, S. W., AND VAALER, J. D. A class of extremal functions for the Fourier transform. *Transactions of the American Mathematical Society 265*, 1 (1981), 283–302.
- [35] GRIESMAIER, R., AND SCHMIEDECKE, C. A multifrequency MUSIC algorithm for locating small inhomogeneities in inverse scattering. *Inverse Problems 33*, 3 (2017), 035015.
- [36] GRÖCHENIG, K. *Foundations of time-frequency analysis*. Springer Science & Business Media, 2013.
- [37] HAARDT, M., AND NOSSEK, J. A. Unitary ESPRIT: How to obtain increased estimation accuracy with a reduced computational burden. *IEEE transactions on signal processing 43*, 5 (1995), 1232–1242.
- [38] HAARDT, M., ROEMER, F., AND DEL GALDO, G. Higher-order SVD-based subspace estimation to improve the parameter estimation accuracy in multidimensional harmonic retrieval problems. *IEEE Transactions on Signal Processing 56*, 7 (2008), 3198–3213.

- [39] HAARDT, M., ZOLTOWSKI, M. D., MATHEWS, C. P., AND NOSSEK, J. 2d unitary ESPRIT for efficient 2d parameter estimation. In *Acoustics, Speech, and Signal Processing, 1995. ICASSP-95., 1995 International Conference on* (1995), vol. 3, IEEE, pp. 2096–2099.
- [40] HARMOUCH, J., KHALIL, H., AND MOURRAIN, B. Structured low rank decomposition of multivariate Hankel matrices. *Linear Algebra and its Applications* (2017).
- [41] HEPPEES, A. On the determination of probability distributions of more dimensions by their projections. *Acta Mathematica Hungarica* 7, 3-4 (1956), 403–410.
- [42] HIGGINS, J. R. Five short stories about the cardinal series. *Bulletin of the American Mathematical Society* 12, 1 (1985), 45–89.
- [43] HOLT, J. J., AND VAALER, J. D. The Beurling-Selberg extremal functions for a ball in Euclidean space.
- [44] HUA, Y. Estimating two-dimensional frequencies by matrix enhancement and matrix pencil. *IEEE Transactions on Signal Processing* 40, 9 (1992), 2267–2280.
- [45] HUA, Y., AND SARKAR, T. K. Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise. *Acoustics, Speech and Signal Processing, IEEE Transactions on* 38, 5 (1990), 814–824.
- [46] HUA, Y., AND SARKAR, T. K. On SVD for estimating generalized eigenvalues of singular matrix pencil in noise. In *Circuits and Systems, 1991., IEEE International Symposium on* (1991), IEEE, pp. 2780–2783.
- [47] JAGERMAN, D., AND FOGEL, L. Some general aspects of the sampling theorem. *IRE transactions on information theory* 2, 4 (1956), 139–146.
- [48] JIANG, T., SIDIROPOULOS, N. D., AND TEN BERGE, J. M. Almost-sure identifiability of multidimensional harmonic retrieval. *IEEE Transactions on Signal Processing* 49, 9 (2001), 1849–1859.
- [49] JOHNSON, D., AND DEGRAAF, S. Improving the resolution of bearing in passive sonar arrays by eigenvalue analysis. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 30, 4 (1982), 638–647.
- [50] KATZNELSON, Y. *An introduction to harmonic analysis*. Cambridge University Press, 2004.
- [51] KOMORNIK, V., AND LORETI, P. *Fourier series in control theory*. Springer Science & Business Media, 2005.
- [52] KUNIS, S., MÖLLER, H. M., PETER, T., AND VON DER OHE, U. Prony’s method under an almost sharp multivariate Ingham inequality. *Journal of Fourier Analysis and Applications* (2017), 1–13.
- [53] KUNIS, S., PETER, T., RÖMER, T., AND VON DER OHE, U. A multivariate generalization of Prony’s method. *Linear Algebra and its Applications* 490 (2016), 31–47.
- [54] LIAO, W. MUSIC for multidimensional spectral estimation: stability and super-resolution. *IEEE Transactions on Signal Processing* 63, 23 (2015), 6395–6406.
- [55] LIAO, W., AND FANNJIANG, A. MUSIC for single-snapshot spectral estimation: Stability and super-resolution. *Applied and Computational Harmonic Analysis* 40, 1 (2016), 33–67.
- [56] LINDEN, D., AND ABRAMSON, N. M. A generalization of the sampling theorem. *Information and Control* 3, 1 (1960), 26–31.
- [57] LINDNER, A. A universal constant for exponential Riesz sequences. *Zeitschrift für Analysis und ihre Anwendungen* 19, 2 (2000), 553–559.
- [58] LITTMANN, F. Quadrature and extremal bandlimited functions. *SIAM Journal on Mathematical Analysis* 45, 2 (2013), 732–747.

- [59] LOGAN, B. Bandlimited functions bounded below over an interval. *Notices Amer. Math. Soc* 24 (1977), A331.
- [60] LUBICH, C. *From quantum to classical molecular dynamics: reduced models and numerical analysis*. European Mathematical Society, 2008.
- [61] MANOLAKIS, D. G., INGLE, V. K., AND KOGON, S. M. *Statistical and adaptive signal processing: spectral estimation, signal modeling, adaptive filtering, and array processing*. McGraw-Hill Boston, 2000.
- [62] MOITRA, A. Super-resolution, extremal functions and the condition number of Vandermonde matrices. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing* (2015), ACM, pp. 821–830.
- [63] MÖLLER, H. M., AND TENBERG, R. Multivariate polynomial system solving using intersections of eigenspaces. *Journal of symbolic computation* 32, 5 (2001), 513–531.
- [64] MONTGOMERY, H. L. The analytic principle of the large sieve. *Bulletin of the American Mathematical Society* 84, 4 (1978), 547–567.
- [65] MORGENSHTERN, V. I., AND CANDÉS, E. J. Super-resolution of positive sources: The discrete setup. *SIAM Journal on Imaging Sciences* 9, 1 (2016), 412–444.
- [66] OTTERSTEN, B., VIBERG, M., AND KAILATH, T. Performance analysis of the total least squares ESPRIT algorithm. *IEEE Transactions on Signal Processing* 39, 5 (1991), 1122–1135.
- [67] PAPOULIS, A. Generalized sampling expansion. *IEEE transactions on circuits and systems* 24, 11 (1977), 652–654.
- [68] PEREYRA, V., AND SCHERER, G. *Exponential data fitting and its applications*. Bentham Science Publishers, 2010.
- [69] PETERSEN, D. P., AND MIDDLETON, D. Sampling and reconstruction of wave-number-limited functions in n-dimensional Euclidean spaces. *Information and control* 5, 4 (1962), 279–323.
- [70] PISARENKO, V. F. The retrieval of harmonics from a covariance function. *Geophysical Journal International* 33, 3 (1973), 347–366.
- [71] PLANCHEREL, M., AND POLYA, G. Fonctions entières et intégrales de fourier multiples. *Commentarii mathematici Helvetici* 10, 1 (1937), 110–163.
- [72] PLONKA, G., AND TASCHE, M. Prony methods for recovery of structured functions. *GAMM-Mitteilungen* (2013).
- [73] PLONKA, G., AND WISCHERHOFF, M. How many Fourier samples are needed for real function reconstruction? *Journal of Applied Mathematics and Computing* 42, 1-2 (2013), 117–137.
- [74] POTTS, D., AND TASCHE, M. Parameter estimation for exponential sums by approximate Prony method. *Signal Processing* 90, 5 (2010), 1631–1642.
- [75] POTTS, D., AND TASCHE, M. Parameter estimation for multivariate exponential sums. *Electronic Transactions on Numerical Analysis* 40 (2013), 204–224.
- [76] POTTS, D., AND TASCHE, M. Parameter estimation for nonincreasing exponential sums by Prony-like methods. *Linear Algebra and its Applications* 439, 4 (2013), 1024–1039.
- [77] POTTS, D., AND TASCHE, M. Fast ESPRIT algorithms based on partial singular value decompositions. *Applied Numerical Mathematics* 88 (2015), 31–45.
- [78] POTTS, D., TASCHE, M., AND VOLKMER, T. Efficient spectral estimation by MUSIC and ESPRIT with application to sparse FFT. *Frontiers in Applied Mathematics and Statistics* 2 (2016), 1.
- [79] RAO, B. D., AND HARI, K. Performance analysis of ESPRIT and TAM in determining the direction of arrival of plane waves in noise. *IEEE Transactions on acoustics, speech, and signal processing* 37, 12 (1989), 1990–1995.

- [80] RENYI, A. On projections of probability distributions. *Acta Mathematica Hungarica* 3, 3 (1952), 131–142.
- [81] ROUQUETTE, S., AND NAJIM, M. Estimation of frequencies and damping factors by two-dimensional ESPRIT type methods. *IEEE Transactions on signal processing* 49, 1 (2001), 237–245.
- [82] ROY, R., AND KAILATH, T. ESPRIT-estimation of signal parameters via rotational invariance techniques. *Acoustics, Speech and Signal Processing, IEEE Transactions on* 37, 7 (1989), 984–995.
- [83] ROY, R., PAULRAJ, A., AND KAILATH, T. ESPRIT—a subspace rotation approach to estimation of parameters of cisoids in noise. *IEEE transactions on acoustics, speech, and signal processing* 34, 5 (1986), 1340–1342.
- [84] SAUER, T. Prony’s method in several variables. *Numerische Mathematik* 136, 2 (2017), 411–438.
- [85] SAUER, T. Prony’s method in several variables: symbolic solutions by universal interpolation. *Journal of Symbolic Computation* (2017).
- [86] SAUER, T. Companion matrices and joint eigenvectors of commuting families of matrices for polynomial zero finding. In *Fourteenth International Conference Zaragoza–Pau on Mathematics and its Applications: Jaca (Spain), September 12 (2018)*, vol. 41, Prensas de la Universidad de Zaragoza, pp. 171–185.
- [87] SCHMIDT, R. Multiple emitter location and signal parameter estimation. *IEEE transactions on antennas and propagation* 34, 3 (1986), 276–280.
- [88] SCHRÖDER, K. *Localized Kernels and Super-Resolution on Special Manifolds*. PhD thesis, Technische Universität München, 2018.
- [89] SEIP, K. *Interpolation and sampling in spaces of analytic functions*. No. 33. American Mathematical Soc., 2004.
- [90] SHANNON, C. E. A mathematical theory of communication. *Bell System Technical Journal* 27 (1948), 379–423.
- [91] SIDIROPOULOS, N. D. Generalizing Caratheodory’s uniqueness of harmonic parameterization to n dimensions. *IEEE Transactions on Information Theory* 47, 4 (2001), 1687–1690.
- [92] SIMON, B. *Basic Complex Analysis: A Comprehensive Course in Analysis, Part 2A*. American Mathematical Soc., 2015.
- [93] STOICA, P., MOSES, R. L., ET AL. *Spectral analysis of signals*, vol. 1. Pearson Prentice Hall Upper Saddle River, NJ, 2005.
- [94] VAALER, J. D. Some extremal functions in Fourier analysis. *Bulletin of the American Mathematical Society* 12, 2 (1985), 183–216.
- [95] VANPOUCKE, F., MOONEN, M., AND BERTHOUMIEU, Y. An efficient subspace algorithm for 2-d harmonic retrieval. In *Acoustics, Speech, and Signal Processing, 1994. ICASSP-94., 1994 IEEE International Conference on (1994)*, vol. 4, IEEE, pp. IV–461.
- [96] VON DER OHE, U. *On the reconstruction of multivariate exponential sums*. PhD thesis, Universität Osnabrück, 2017.
- [97] WENDLAND, H. *Scattered Data Approximation*. Cambridge University Press, 2005.
- [98] WIENER, N. *The Fourier integral and certain of its applications*. Dover, 1951.
- [99] WISCHERHOFF, M. *Reconstruction of Structured Functions From Sparse Fourier Data*. PhD thesis, Niedersächsische Staats-und Universitätsbibliothek Göttingen, 2015.
- [100] YEN, J. On nonuniform sampling of bandwidth-limited signals. *IRE Transactions on circuit theory* 3, 4 (1956), 251–257.

- [101] YOUNG, R. M. *An Introduction to Non-Harmonic Fourier Series, Revised Edition*. Elsevier, 2001.
- [102] ZYGMUND, A. *Trigonometric series: Vol. II. 2nd edn.(reprinted with corrections and some additions)*. Cambridge University Press, London/New York, 1968.

Zusammenfassung

Die vorliegende Arbeit befasst sich mit Frequenzschätzung von Exponentialsummen. Kurz gesagt ist die Aufgabe, aus einer endlichen Anzahl abgetasteter Funktionswerte die unbekanntes Frequenzen, also das diskrete Spektrum, einer Exponentialsumme zu berechnen. Gerade der höher dimensionale Fall hat in den letzten Jahren viel Aufmerksamkeit auf sich gezogen.

Der erste Teil dieser Arbeit behandelt die Wohlgestelltheit des Frequenzschätzungsproblem. Die Leitfrage lässt sich wie folgt formulieren: Wenn man zwei Exponentialsummen hat, deren abgetastete Funktionswerte eng beieinander liegen, was kann über ihre Frequenzen ausgesagt werden? Unter der (notwendigen) Voraussetzung, dass beide Exponentialsummen wohlseparierte Frequenzen haben, werden scharfe Abschätzungen gezeigt. Diese führen dann zu a posteriori Abschätzungen.

Für den Beweis benötigt man spezielle, bandlimitierte Funktion, die einer Vorzeichenbedingung genügen. Da die bisher bekannten Funktionen dieser Klasse nicht über eine notwendige zusätzliche Eigenschaft verfügen, wird eine geeignete Konstruktion angegeben. Dazu werden Ergebnisse aus der Sampling Theorie verwendet, weshalb das Kapitel mit einer kurzen Einführung in diese beginnt.

Der zweiten Teil wendet sich dem algorithmischen Aspekt des Problems zu. Nach einer kurzen Wiederholung einiger gängiger Methoden, werden zunächst projektionsbasierte Verfahren diskutiert. Diese reduzieren das höherdimensionale Problem auf mehrere eindimensionale Probleme, indem die multivariate Exponentialsumme entlang einiger Linien abgetastet wird. Sowohl für den Fall von parallelen, wie auch von paarweise nicht parallelen Linien werden Kriterien, die eine Wiederherstellung garantieren, bewiesen. Im Fall von parallelen Linien wird ein ESPRIT ähnliches Verfahren vorgeschlagen, dass die entstehenden eindimensionalen Probleme gleichzeitig löst.

Anschließend werden andere Zugänge zum mehrdimensionalen Frequenzschätzungsproblem besprochen. Durch Einführen des Signalraums lassen sich leicht Varianten von Pronys Verfahren, ESPRIT und MUSIC für diesen Fall entwickeln. Insbesondere erlaubt dies die Verwendung von sehr kleinen Abtastmengen, was die bisher bekannte Theorie für ESPRIT und MUSIC erweitert. Weiterhin wird dadurch die Komplexität erheblich reduziert. Solche Abtastmengen wurden vorher von Sauer für das Pronyverfahren eingeführt.

Eidesstattliche Versicherung

Hiermit erkläre ich an Eides statt, dass ich die vorliegende Dissertationsschrift selbst verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Hamburg, den

Unterschrift: