Forum of Mathematics, Sigma (2026), **Vol. 14**:e95 1–71
doi:10.1017/fms.2026.10238

CAMBRIDGE UNIVERSITY PRESS

RESEARCH ARTICLE

Spectral algorithms in higher-order Fourier analysis

Pablo Candela[1], Diego González-Sánchez[2] and Balázs Szegedy[3]

[1]ICMAT, Madrid, Spain; E-mail: pablo.candela@icmat.es (Corresponding author).
[2]Université Paris Cité, Sorbonne Université, CNRS, IMJ-PRG, Paris, France; E-mail: gonzalezsanchez@imj-prg.fr.
[3]Alfréd Rényi Institute of Mathematics, Budapest, Hungary; E-mail: szegedyb@gmail.com.

Received: 29 September 2025 UTC; **Revised:** 13 April 2026 UTC; **Accepted:** 8 May 2026 UTC

2020 Mathematical Subject Classification: *Primary* – 11B30; *Secondary* – 47S99, 47J99, 65D15

Abstract

Our goal is to provide simple and practical algorithms in higher-order Fourier analysis which are based on spectral decompositions of operators. We propose a general framework for such algorithms and provide a detailed analysis of the quadratic case. Our results reveal new spectral aspects of the theory underlying higher-order Fourier analysis. Along these lines, we prove new inverse and regularity theorems for the Gowers norms based on higher-order character decompositions. Using these results, we prove a spectral inverse theorem and a spectral regularity theorem in quadratic Fourier analysis.

Contents

1 Introduction **2**
1.1 Outline of the spectral approach . . . 4
1.2 Main results with algorithmic consequences . . . 6
1.3 Nilspace theoretic foundations for the spectral method . . . 9
1.4 Connection with the ultralimit setting . . . 11
2 Basic notions **12**
2.1 Invariant operators . . . 14
2.2 A simple instance of the spectral approach: recovering classical Fourier analysis . . . 15
2.3 Structured functions of order k . . . 15
2.4 Characters of order k . . . 18
2.5 The Fourier denoising operator K_ε . . . 21
3 Elementary motivation of the spectral approach in quadratic Fourier analysis **26**
3.1 Spectrally isolated eigenvectors of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ are (weak) quadratic characters . . . 26
3.2 Further properties of the operator K_ε . . . 28
4 Nilspace characters **32**
4.1 Uniform approximation of Lipschitz functions by vertical Fourier sums . . . 33
4.2 2-step nilspace characters are quadratic characters . . . 36
4.3 Approximate orthogonality of balanced nilspace characters . . . 39
4.4 k-step nilspace polynomials are structured functions of order k . . . 41
5 Refined regularity in terms of nilspace characters **42**
5.1 The refined regularity lemma . . . 42
5.2 Correlation, approximate diagonalization of the U^k norm, and regularization . . . 43

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

6 A structure theorem for $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ using 2-step nilspace characters **47**
7 Algorithmic consequences of the structure theorem for $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ **48**
7.1 Obtaining a quadratically structured part of f from the spectrum of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ 49
7.2 Obtaining dominant nilspace characters of f from the spectrum of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ 54
A Metrics on CFR nilspaces derived from the manifold structure **63**
References **69**

1. Introduction

The theory of higher-order Fourier analysis, initiated by the work of Gowers in arithmetic combinatorics [22], is part of a revolutionary development in mathematics from the early 2000s, which has led to a much deeper understanding of Szemerédi's famous theorem on arithmetic progressions [62], and of additive structures in general. Other parts of this development occurred in hypergraph theory (via refinements of Szemerédi's regularity lemma and its extensions for hypergraphs [23, 56, 65]), and in ergodic theory (in directions stemming from Furstenberg's approach to Szemerédi's theorem [19], especially the structure theory of characteristic factors [43, 73]). A common goal of these topics is to understand higher-order correlations and interactions in complex structures. This development has had a profound impact in several areas in addition to the aforementioned ones, particularly in number theory (a landmark being the Green-Tao theorem [33]) and in theoretical computer science [40, 70].

Progress in higher-order Fourier analysis in the last two decades includes the discovery and advancement of general foundations for this theory, which involve nilpotent structures [5, 11, 34, 35, 44]. However, on one hand the proofs are often long and intricate, which can make the field less accessible to nonexperts, and on the other hand, the development of these foundations has so far focused on the theory rather than on potential practical applications. Yet, classical Fourier analysis is one of the most important tools in modern science, and there are good reasons to believe that its higher-order extensions can also become powerful practical tools. Therefore, there is a strong motivation for finding new approaches that can advance this field, including its foundations, in more elementary and applicable directions. In this paper we introduce such an approach, with the goal of bridging the gap between theory and practice.

The theoretical part of this approach consists in connecting higher-order components of functions on abelian groups (generalizations of Fourier characters, detailed below) with eigenvectors of certain operators constructed from such functions. The construction of these operators is elementary and is outlined in Subsection 1.1 of this introduction. In the first part of the paper, we provide initial results establishing the aforementioned connection at a relatively elementary level. The second part of the paper (starting in Section 4) delves into deeper reasons for this connection, which involve what can be viewed as a constituent of the foundations of higher-order Fourier analysis, namely *nilspace theory* (discussed further below). Our results in this part include a new structure (or regularity) theorem for Gowers norms on finite abelian groups, which decomposes functions into sums of certain generalizations of Fourier characters that we call *nilspace characters*, which are nearly orthogonal (Theorem 5.1). Consequences of this result include an inverse theorem involving such nilspace characters (Theorem 5.2), and new approximate diagonalizations and Parseval identities for Gowers norms (Theorems 5.7 and 5.4). These theoretical results are discussed in more detail in Subsection 1.3 of this introduction.

To demonstrate the applicability of these results, we provide simple algorithms which use the above-mentioned spectral data to recover the higher-order decompositions, and which do not require detailed knowledge of the deeper aspects of the theory. While the general pattern for these algorithms is easily formulated for any higher order (see Remark 1.4), in this paper we focus on fully demonstrating such algorithms in the case of *quadratic Fourier analysis*. This is outlined in Subsection 1.2 in this introduction, where we present our main theorems in this case (Theorems 1.1 and 1.6) and formulate the resulting algorithms (Algorithms 1 and 2). These algorithms can be implemented straightforwardly using the spectral decomposition of self-adjoint matrices and the (fast) Fourier transform. Algorithm 1 could be used in the prediction of missing data (or to carry out what could be called *quadratic denoising*)

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

in ordered data sets and thus in time series prediction, analogously to how the usual Fourier transform is currently used for such tasks (a basic illustration is given in Figure 2). The decomposition into *quadratic characters* given by Algorithm 2 could be used to identify important components of a function, similarly to principal component analysis (see also Remark 2.14). These aspects will be further explored in a more applied follow-up work, which will focus on the practical aspects of our methods, including experiments and refinements.

To begin explaining and motivating our spectral approach in more detail, let us first recall here some basic ideas of higher-order Fourier analysis. The Gowers norms (or uniformity norms) are among the most important notions in this theory. For each integer $k \geq 2$, the k-th Gowers norm (or U^k-norm) is defined on the space of complex-valued functions f on a finite abelian group,[1] the norm $\|f\|_{U^k}$ consisting of an average over configurations known as k-dimensional *cubes* on the group (the formula is recalled in equation (1.3) below). The U^2-norm of f is equal to the ℓ^4-norm of the Fourier transform of f, thereby connecting the U^2-norm with classical (first-order) Fourier analysis. The role of the U^{k+1}-norm in k-th order Fourier analysis centers on providing a useful concept of a function being *noise of order k* (or *quasirandom of degree k*), namely, a function having small U^{k+1}-norm. This in turn yields, in a dual way, a notion of a function being *structured of order k*, namely, a function being nearly orthogonal to any noise of order k. The Gowers norms form an increasing sequence (meaning that $\|f\|_{U^k} \leq \|f\|_{U^{k+1}}$ for every k and f), and this implies that, as k increases, fewer functions are classified as noise, and the notion of structured function becomes more inclusive and subtle. Major efforts in this field have gone into describing, as precisely as possible, these higher-order structured functions, in particular by seeking *fundamental* k-th order structured functions that could act adequately as generalizations of Fourier characters. As we shall see in this paper, the above-mentioned nilspace characters are functions of this type. A fascinating aspect of this subject is that, whereas the geometric object underlying classical Fourier characters is the circle group, the higher-order structured functions can involve nonabelian nilpotent structures such as the Heisenberg nilmanifold. In these directions, two central and interrelated themes have developed: on one hand, the so-called *inverse theorems for Gowers norms*, which establish that functions with nonnegligible U^{k+1}-norm correlate nontrivially with some k-th order generalization of a Fourier character; on the other hand, *decomposition theorems* (or *regularity lemmas*), which express any bounded function essentially as a sum of a structured part and a noise part of order k.

There is by now extensive literature in higher-order Fourier analysis proving inverse theorems for Gowers norms in various families of abelian groups, often with the remarkable additional feature of giving effective bounds for the result (see for instance [12, 22, 26, 27, 32, 35, 46, 47, 51, 52, 53, 54, 67, 68]). There are also many works proving higher-order decompositions of functions and regularity lemmas for Gowers norms (e.g., [24, 28, 29, 31, 33, 42, 64]). However, works providing algorithmic implementations of the above results are fewer, with notable examples being the paper by Tulsiani and Wolf introducing quadratic analogues of the Goldreich–Levin algorithm [71], and the more recent paper of Kim, Li, and Tidor on the cubic case [48]. The algorithms in these works are probabilistic, involve iterative processes [71, See §3], and are specialized to the finite-field setting (of central interest in theoretical computer science). The new algorithms that we begin to explore in this paper are based on spectral decompositions of operators associated with functions, are more direct (in particular, essentially noniterative), and applicable on any finite abelian group.

This new spectral aspect of higher-order Fourier analysis has quite natural motivations. One such motivation is the well-known connection between classical Fourier analysis and spectral decompositions. Fourier characters on abelian groups are eigenvectors of shift operators. More generally, if a linear operator is shift-invariant (i.e., it commutes with every shift), then every Fourier character is an eigenvector of that operator. This property enables Fourier analysis to be recovered from spectral theory (see Subsection 2.2). The results in this paper describe a similar connection between higher-order Fourier analysis and spectral theory through adequate generalizations of shift-invariant operators and Fourier

[1] These norms can be defined more generally for bounded measurable functions on compact abelian groups.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

characters. Another motivation for using spectral methods comes from Szemerédi's regularity lemma for graphs [63], which is a precursor of the intricate regularity lemmas in higher-order Fourier analysis. The spectral nature of graph regularization was first highlighted by Frieze and Kannan in [18]. Subsequently, the third author showed in [61] that Szemerédi-type regular partitions (even in their stronger forms [3]) can be obtained from the dominant eigenvectors of adjacency matrices. The deep connection between eigenvectors and regularization in graph theory is part of a rich algorithmic framework that includes principal component analysis, low-rank approximations, spectral clustering, and dimensionality reduction more broadly. An important aim of this paper is to position higher-order Fourier analysis within this framework by describing its spectral aspects, with emphasis on regularization.

1.1. Outline of the spectral approach

Given a finite abelian group Z and a function $f : \mathrm{Z} \to \mathbb{C}$, a first key step in this approach consists in turning f into a matrix in $\mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ (or a Z-*matrix*, as we shall call it; see Definition 2.1). We consider various choices for how to do this, all of which proceed by applying a (typically nonlinear) transformation $K : \mathbb{C}^{\mathrm{Z}} \to \mathbb{C}^{\mathrm{Z}}$ to every so-called Z-*diagonal function* of the rank-1 matrix

$$f \otimes \overline{f} : \mathrm{Z}\times\mathrm{Z} \to \mathbb{C},\ (x, y) \mapsto f(x)\overline{f(y)},$$

a Z-diagonal being a set of entries whose indices have fixed difference in Z (see Definition 2.1).

The matrix $f \otimes \overline{f}$ is a natural and convenient object to consider in higher-order Fourier analysis, since its Z-diagonal function with entry-difference t is the multiplicative derivative

$$\Delta_t f(x) := f(x + t)\overline{f(x)},$$

which plays a key role in this theory (notably in proofs of inverse theorems). The matrix $f \otimes \overline{f}$ thus encapsulates in a useful way all the multiplicative derivatives of f.

In general, given a map (or operator) $K : \mathbb{C}^{\mathrm{Z}} \to \mathbb{C}^{\mathrm{Z}}$ and a matrix $M \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$, we denote by $\mathcal{K}(M)$ the matrix obtained by applying K to every Z-diagonal function of M. To ensure that $\mathcal{K}$ is well-behaved, we require the operator K to commute with complex conjugation and the shift action of Z on functions in $\mathbb{C}^{\mathrm{Z}}$ (in which case we call K an *invariant* operator; see Definition 2.10). If these conditions are satisfied, then $\mathcal{K}$ preserves the property of being a *self-adjoint* (or Hermitian) matrix. In particular $\mathcal{K}(f \otimes \overline{f})$ is self-adjoint, so it has real eigenvalues and orthogonal eigenvectors. (In this paper, the linear-algebraic notions pertaining to Z-matrices, including their eigenvalues, are normalized in coherence with the probability Haar measure on Z and the corresponding inner-product on the space $\mathbb{C}^{\mathrm{Z}}$, a space on which Z-matrices act as kernels of linear integral operators; see Definition 2.3 and the discussion preceding it.)

It turns out that the spectral data associated with $\mathcal{K}(f \otimes \overline{f})$ carries important information to perform k-th order Fourier analysis on f. Let us formulate this as the following principle, which captures in general terms a phenomenon that is central to this approach.

Order increment principle: *If $K : \mathbb{C}^{\mathrm{Z}} \to \mathbb{C}^{\mathrm{Z}}$ maps functions in $\mathbb{C}^{\mathrm{Z}}$ to their k-th order structured parts, then for $f \in \mathbb{C}^{\mathrm{Z}}$ the spectral decomposition of $\mathcal{K}(f \otimes \overline{f})$ can be used to obtain the $(k + 1)$-th order structured part (and corresponding useful decomposition) of f.*

This principle can be used to turn k-th order Fourier analysis into $(k + 1)$-th order Fourier analysis. A precise form of this principle was observed in the ultralimit setting in [59].

A simple choice for K is the averaging operator (mapping f to the constant function[2] $\mathbb{E}_{x\in\mathrm{Z}} f(x)$), which pertains to 0-th order Fourier analysis. In this case the eigenvectors of $\mathcal{K}(f \otimes \overline{f})$ are the Fourier characters of f, and the eigenvalues are the squares of the absolute values of the corresponding Fourier

[2]Recall the standard notation for averaging, whereby given a finite set X and a function f from X to a real or complex vector space, we denote by $\mathbb{E}_{x\in X} f(x)$ the average $\frac{1}{|X|}\sum_{x\in X} f(x)$.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

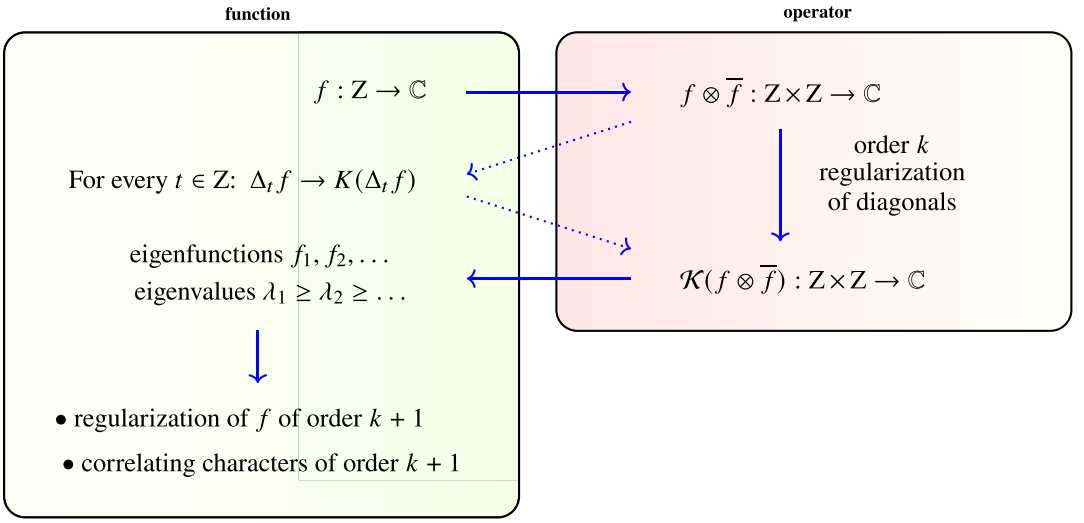


Figure 1. *Sketch of the spectral approach to higher-order Fourier analysis.*

coefficients. This illustrates how 0-th order Fourier analysis is turned into first order, classical Fourier analysis (we discuss this in more detail in Subsection 2.2).

The next logical step is to progress from first-order to quadratic Fourier analysis using the above principle. This step requires an operator K that extracts the Fourier-structured component of a function f in an adequate way. Finding such an operator is challenging because of a lack of uniqueness. Indeed, while in an appropriate limit setting (such as the setting of [59]), there exists a unique, and even linear, operator that isolates the structured part (consisting of conditional expectation relative to an appropriate σ-algebra), the counterparts in finite settings exist only in an approximate and nonlinear way. There are various possible choices, many of which depend on additional parameters. In this work, we introduce such an operator with especially useful analytic properties, which we call the *Fourier denoising operator*, and we develop several tools for its application.

To define this operator, recall first that for a character χ in the dual group $\widehat{\mathrm{Z}}$ (where Z is a finite abelian group), we define the corresponding Fourier coefficient of f by

$$\widehat{f}(\chi) = \mathbb{E}_{x\in\mathrm{Z}} f(x)\overline{\chi(x)}. \tag{1.1}$$

Then, for a fixed (typically small) constant $\varepsilon > 0$, our denoising operator is defined by

$$K_\varepsilon(f) = \sum_{\chi\in\widehat{\mathrm{Z}} :\, |\widehat{f}(\chi)|\geq\varepsilon} \frac{|\widehat{f}(\chi)|-\varepsilon}{|\widehat{f}(\chi)|}\, \widehat{f}(\chi)\, \chi. \tag{1.2}$$

Note that there is a simpler operator (that we can call the *Fourier cut-off operator*) which sets Fourier coefficients with absolute value smaller than ε to 0 and keeps the remaining terms in the Fourier decomposition unchanged. The main problem with the Fourier cut-off operator is that it is not continuous and has inconvenient properties for calculations. Instead, the operator K_ε applies the continuous function $x \mapsto \mathrm{ReLU}(x-\varepsilon)$ to the magnitudes of the Fourier coefficients while keeping their phases. Here, ReLU is the function $x \mapsto (x+|x|)/2 = \max(x, 0)$ known from machine learning. It turns out that K_ε is a contraction in L^2 and has other pleasant properties while being sufficiently similar to the cut-off operator. Related to the operator K_ε we introduce the operator $\mathcal{K}_\varepsilon : \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}} \to \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ which applies K_ε to all Z-diagonals of a matrix in $\mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$. A more detailed discussion of the Fourier denoising operator begins in Subsection 2.5.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

1.2. Main results with algorithmic consequences

Let us begin by recalling the formula for the Gowers norms from [22, Lemma 3.9] (see also [66, Definition 11.2]). For each integer $k \geq 2$, the U^k-norm of a function $f \in \mathbb{C}^{\mathrm{Z}}$ is defined by

$$\|f\|_{U^k} := \left(\mathbb{E}_{x,t_1,\dots,t_k \in \mathrm{Z}} \prod_{\nu \in \{0,1\}^k} \mathcal{C}^{|\nu|} f(x + \nu(1)\, t_1 + \nu(2)\, t_2 + \cdots + \nu(k)\, t_k)\right)^{1/2^k}, \tag{1.3}$$

where $\mathcal{C}$ denotes the complex conjugation operator and $|\nu| := \sum_{i=1}^{k} \nu(i)$.

The dual of the U^k-norm, defined by $\|f\|_{U^k}^* := \sup_{g \in \mathbb{C}^{\mathrm{Z}}, \|g\|_{U^k} \leq 1} \mathbb{E}_{x \in \mathrm{Z}} f(x)\overline{g(x)}$, is a useful tool to measure the extent to which a function is structured in the sense of $(k-1)$-th order Fourier analysis, as we will explain in more detail later (see Subsection 2.3). These dual norms have been analyzed previously, notably in the work of Host and Kra [42]. Their earlier work [45] already studied analogous dual objects for uniformity seminorms, and the work of Tao [64] introduced closely related *uniform almost-periodicity norms*, also mentioned later below. In this paper, we shall use the dual norms especially to define the quantitative and relatively elementary concept of *structured function of order* $k-1$ (see Definition 2.23).

The next theorem is one of the main results in this paper. It provides a regularity (or decomposition) result for functions in quadratic Fourier analysis, expressing the quadratically structured part of the function in terms of the dominant eigenvalues and corresponding eigenvectors of the appropriate Z-matrix. Thus it connects regularity with spectral methods in the quadratic setting. This provides a theoretical basis for Algorithm 1 below.

Theorem 1.1 (Spectral U^3-regularization). *For every $\rho_0 \in [0,1)$ there exists $\varepsilon_0 > 0$ such that the following holds. Let Z be a finite abelian group and let $f : \mathrm{Z} \to \mathbb{C}$ be a 1-bounded function. Then there exists $\rho \in [\rho_0/2, \rho_0]$ and $\varepsilon \in [\varepsilon_0, 1]$ with the following property. Let f_{reg} be the projection of f to the linear span of the eigenspaces of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ with corresponding eigenvalues at least ρ. Then $\|f - f_{reg}\|_{U^3} \leq 2\rho^{3/8}$ and there exists $h : \mathrm{Z} \to \mathbb{C}$ such that $\|f_{reg} - h\|_2 \leq \rho$ and $\|h\|_{U^3}^* = O_\rho(1)$.*

The conclusion involving the approximating function h, of bounded U^3-dual-norm, can be summarized as saying that f_{reg} is a *structured function of order 2* (see Definition 2.23). Note also that, as mentioned earlier, eigenvalues here are taken in a normalized form (see Definition 2.3).

Remark 1.2. Analogous to Szemerédi's regularity lemma for graphs [63], arithmetic regularity lemmas (such as Theorem 1.1) exist in multiple versions [24, 31], each balancing the trade-off between the strength of the statement and the quality of the bounds. For instance, the weak regularity lemma by Frieze and Kannan [17] provides a more limited approximation for graphs but achieves significantly better bounds, making it suitable for practical applications. Szemerédi's original formulation [62] occupies a middle ground in this trade-off: it does not provide effective bounds, but it has stronger theoretical implications. The primary goal of Theorem 1.1 is to present a concise, representative member of a broader family of regularity lemmas that connect higher-order Fourier analysis with spectral theory. In fact, the version of Theorem 1.1 that we actually prove, namely Theorem 7.12, is already finer and more flexible (though also more technical). In future papers, we will explore other versions of the result, including stronger formulations, to extend its applicability and impact.

The calculation of f_{reg} in Theorem 1.1 is efficient in practice. It involves $|\mathrm{Z}|$ independent applications of the operator K_ε (which can be executed in parallel) and the computation of dominant eigenvectors and eigenvalues of a single $\mathrm{Z} \times \mathrm{Z}$ matrix. Each application of K_ε can be done using a fast Fourier transform, a simple coefficient truncation, and an inverse fast Fourier transform. The whole process depends on the original function f defined on Z, as well as on two positive parameters ε and ρ. The details of this calculation are outlined in the following pseudocode.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Input: $f : \mathrm{Z} \to \mathbb{C}$, $\rho, \varepsilon \in \mathbb{R}_{>0}$
$M \leftarrow f \otimes \overline{f} \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$
for $t \in \mathrm{Z}$ **do**
  $M(\cdot + t, \cdot) \leftarrow K_\varepsilon(M(\cdot + t, \cdot))$
end
$(\mu_1, v_1), \ldots, (\mu_{|\mathrm{Z}|}, v_{|\mathrm{Z}|}) \leftarrow$ **Eigendecomposition**(M)
Output: $f_{\mathrm{reg}} \leftarrow \sum_{\mu_i \geq \rho} \langle f, v_i \rangle v_i$ and $\{(\mu_1, v_1), \ldots, (\mu_{|\mathrm{Z}|}, v_{|\mathrm{Z}|})\}$[3]

Algorithm 1: U^3-regularization algorithm.

Remark 1.3 (Choice of parameters). In practice, there is no definitive choice for the parameters ρ and ε in Algorithm 1. The optimal values depend heavily on the nature of the dataset the algorithm is applied to. Adjusting these parameters allows for analyzing the quadratic structure of a function at different levels of resolution. Practical aspects of how to choose these parameters will be investigated in a paper focusing on applications.

Remark 1.4 (Higher-order versions). Let us denote the outcome of Algorithm 1 by $H_{\varepsilon,\rho}(f)$. It is easy to see that for fixed positive ε and ρ, the operator $f \mapsto H_{\varepsilon,\rho}(f)$ commutes with shifts and conjugation. Thus, by applying it to the Z-diagonals of $f \otimes \overline{f}$, we obtain a self-adjoint matrix. By choosing a new value $\kappa > 0$ and projecting f onto the space spanned by the eigenvectors of this matrix with eigenvalue at least κ, we obtain a new operator $f \mapsto H_{\varepsilon,\rho,\kappa}(f)$ which approximates the structured part of f in cubic Fourier analysis. By further iterating this process, we obtain a spectral approach to k-th order Fourier analysis based on a regularization operator $f \mapsto H_{\varepsilon_1,\varepsilon_2,\ldots,\varepsilon_k}(f)$ with k parameters. However, we do not pursue the analysis of this general algorithm here, as its rigorous validation requires work beyond the scope of this paper. Indeed, while several central tools for such a general validation are already established below (see especially Theorem 5.1, valid for all orders), some key results and finitary notions are only established in the quadratic setting in this paper (a central example being the Fourier denoising operator and the related Proposition 3.17, which are important for the proof of the main result validating Algorithm 1, namely Theorem 6.1). Nevertheless, we strongly believe that such higher-order algorithms can be validated rigorously; this belief is supported notably by related perspectives in nonstandard analysis (see Subsection 1.4).

Remark 1.5 (Continuous versions). Note that the output $H_{\varepsilon,\rho}(f)$ of Algorithm 1 is not continuous in its input $f \in \mathbb{C}^{\mathrm{Z}}$ for fixed positive ρ, ε. To see this, fix some ρ, ε and let $\mathrm{Z} = \mathbb{Z}_p$ for some prime p, $g(x) := e^{2\pi i x^2/p}$, and $c \in \mathbb{R}_{\geq 0}$. It is easy to see that $\mathcal{K}_\varepsilon((cg) \otimes \overline{(cg)}) = [\sqrt{c^2 - \varepsilon})g] \otimes \overline{[\sqrt{c^2 - \varepsilon})g]}$ for $c^2 > \varepsilon$, and 0 otherwise, so $H_{\varepsilon,\rho}(cg(x)) = (\sqrt{c^2 - \varepsilon})g$ if $\rho < c^2 - \varepsilon$ and 0 otherwise. There is therefore a discontinuity in the output $H_{\varepsilon,\rho}(cg)$ as the parameter c in the input cg increases past $\sqrt{\rho + \varepsilon}$.

There exist continuous variants of $H_{\varepsilon,\rho}$, which are worth mentioning since these variants can be convenient for instance in computer experiments. An example of a continuous version is $H'_{\varepsilon,\rho}(f) := \mathbb{E}_{\rho' \in [\rho/2,\rho]} H_{\varepsilon,\rho'}(f)$. It is easy to see that $H'_{\varepsilon,\rho}$ can also be computed efficiently, with complexity similar to $H_{\varepsilon,\rho}$. To achieve this, the formula for the projection in $H_{\varepsilon,\rho}$ must be augmented with certain weight terms that depend on the eigenvalues. Recall that the Fourier cut-off operator mentioned earlier suffered from a similar continuity problem, and was replaced by the *continuous* denoising operator K_ε to obtain our proofs in the quadratic setting. Besides producing more stable and precise outcomes, continuous versions also appear to be more suitable when moving to higher orders via the order increment principle. These and other refinements of the algorithm lie outside the scope of this paper and will be explored in future work.

A deeper and more sophisticated version of our spectral algorithm focuses on the meaning of the individual eigenvectors of the operator $\mathcal{K}_\varepsilon(f \otimes \overline{f})$. It turns out that these eigenvectors can give further

[3]We assume that the eigenvalues are ordered decreasingly, that is, $\mu_i \geq \mu_{i+1}$ for all i.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

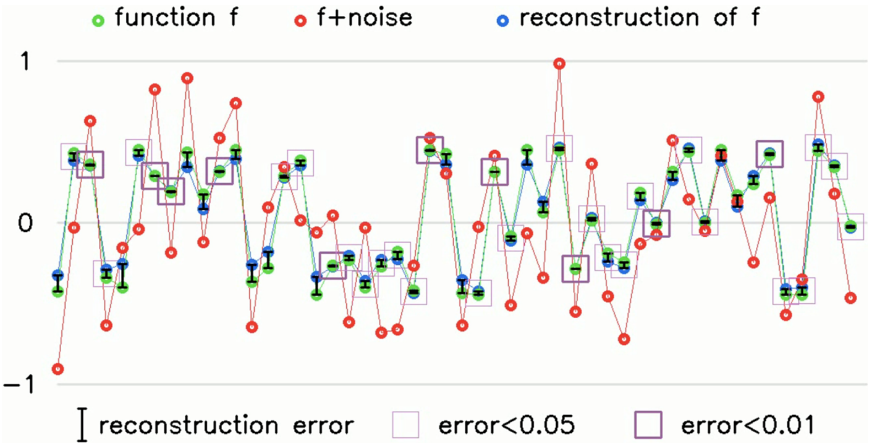


Figure 2. *Removing added random noise from a quadratic structured function by a version of Algorithm 1 on the cyclic group* $\mathbb{Z}_{500}$. *A window of length* 50 *is plotted for illustration. In this example, the function* $f(i) := \sin(8i^2 + 3i + 1)$, $i \in [500]$ *(green graph) is perturbed by random noise, resulting in the function* $g = f + r$ *(red graph). The spectral algorithm is applied to* g *and the reconstruction* f_2 *of* f *(blue graph) is obtained by the projection of* g *to the space spanned by the* 6 *leading eigenvectors of the operator constructed from* g. *The plot highlights the reconstruction error* $|f(i) - f_2(i)|$.

information on the quadratic structured part of f. In particular, if the corresponding eigenvalue is large enough and well-separated from other eigenvalues, such an eigenvector behaves as a quadratic generalization of a Fourier character, and has large inner product with f. This generalization is an interesting notion in itself and can be defined in a basic way. The initial observation for this is that a classical Fourier character χ (or a constant multiple of it) is defined by the property that for every t the multiplicative derivatives $\Delta_t \chi$ is a constant function, and that constant functions are the structured functions in 0-th order Fourier analysis (this is proved, in greater generality, in Lemma 2.26). Going one degree higher, we obtain the following basic notion: a function $f \in \mathbb{C}^{\mathbb{Z}}$ is a *quadratic character* if for every $t \in \mathbb{Z}$ the multiplicative derivative $\Delta_t f$ is *Fourier structured* (of first order) in the sense of Definition 2.23. This notion of quadratic character depends on two parameters that measure the extent to which the multiplicative derivatives are Fourier structured; more precisely f is a quadratic character of *complexity* R and *precision* δ if for every t there is $g_t \in \mathbb{C}^{\mathbb{Z}}$ such that $\|\Delta_t f - g_t\|_{L^2} \leq \delta$ and $\|g_t\|^*_{U^2} \leq R$. Let us mention straightaway that a central result in this paper connects this basic notion of a quadratic character with the 2-step case of the deeper notion of *nilspace characters* mentioned earlier (we discuss this central result in the next subsection). We also introduce basic notions of *characters of order* k for each $k \geq 2$, studied in Subsection 2.4 (see Definition 2.29).

It turns out that the quadratically structured part of f given by Theorem 1.1 can be decomposed in terms of quadratic characters of bounded complexity, and that our second algorithm is able to find (good approximations of) those characters. To state our main result validating this second algorithm, we use the following terms. For a self-adjoint $\mathbb{Z}$-matrix A, we denote by $\mathrm{Spec}_\rho(A)$ the set of eigenvalues of A larger than ρ, and we denote by $\mathrm{Eigen}_\rho(A) \subset \mathbb{C}^{\mathbb{Z}}$ the subspace spanned by eigenvectors of A corresponding to eigenvalues in $\mathrm{Spec}_\rho(A)$. We also say that a set $X \subset \mathbb{C}$ is *δ-separated* if for every $x, y \in X$ with $x \neq y$ we have $|x - y| \geq \delta$. We can now state the main result in question.

Theorem 1.6. *For every* $\rho_0 \in [0, 1/10]$, *there exists* $\varepsilon_0 > 0$ *such that for any finite abelian group* $\mathbb{Z}$ *and any 1-bounded function* $f : \mathbb{Z} \to \mathbb{C}$, *there exists* $\rho \in [\rho_0/2, \rho_0]$ *and* $\varepsilon \in [\varepsilon_0, 1]$ *satisfying the following property. Let* h *be a function in* $\mathrm{Eigen}_\rho\big(\mathcal{K}_\varepsilon(f \otimes \overline{f})\big) \subset \mathbb{C}^{\mathbb{Z}}$ *with* $\|h\|_2 \leq 1$ *such that* $\mathrm{Spec}_{\rho^7}\big(\mathcal{K}_\varepsilon(h \otimes \overline{h})\big)$ *is* ρ^7*-separated and has cardinality equal to* $|\mathrm{Spec}_\rho\big(\mathcal{K}_\varepsilon(f \otimes \overline{f})\big)|$. *Then for every unit eigenvector* v *of* $\mathcal{K}_\varepsilon(h \otimes \overline{h})$ *corresponding to an eigenvalue in* $\mathrm{Spec}_{\rho^7}\big(\mathcal{K}_\varepsilon(h \otimes \overline{h})\big)$, *there is a quadratic character* g *of complexity* $O_\rho(1)$ *and precision* ρ *satisfying* $\|v - g\|_{L^2} \leq 56\rho^{7/2}$ *and* $|\langle f, g\rangle| \geq \sqrt{\rho/4}$.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Remark 1.7. Theorem 1.6 is a simpler version of the result that we actually prove, namely Theorem 7.15. In this introductory setting, let us mention informally just one of the additional features of the latter theorem: there is in fact a *bijection* between the eigenvectors of $\mathcal{K}_\varepsilon(h \otimes \overline{h})$ with large eigenvalues and the dominant 2-step nilspace characters of f given by the new decomposition result proved in this paper (Theorem 5.1, discussed in the next subsection). In particular, every such nilspace character can be recovered from these eigenvectors (up to a small error).

A typical situation enabling a simple application of Theorem 1.6 is when the 1-bounded function $f : \mathrm{Z} \to \mathbb{C}$ has the largest eigenvalues of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ sufficiently separated. Then, taking h to be the structured part f_{reg} given by[4] Theorem 1.1, the spectral assumptions in Theorem 1.6 are satisfied, and we can then directly obtain the dominant quadratic characters of f from the eigenvectors associated with those large eigenvalues (see Lemma 7.18). This situation is actually generic, in a sense that we can use in order to handle the cases where those spectral assumptions might fail. Indeed, in such cases, we show that one can use a randomized procedure, consisting essentially in calculating the dominant eigenvectors and eigenvalues of $\mathcal{K}_\varepsilon(h \otimes \overline{h})$ where h is a random combination of the dominant eigenvectors of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$. We show that, with high probability, this procedure yields a function h to which Theorem 1.6 applies, thus resulting in a decomposition of the structured part of f into quadratic characters (see Subsection 7.2.1). The formalization of this probabilistic procedure is given in Theorem 7.24 and Remark 7.25. Theorems 1.6 and 7.24 together validate Algorithm 2 below.

Input: $f : \mathrm{Z} \to \mathbb{C}, \rho, \varepsilon, \delta \in \mathbb{R}_{>0}$
$f_{\mathrm{reg}}, (\mu_1, v_1), \dots, (\mu_{|\mathrm{Z}|}, v_{|\mathrm{Z}|}) \leftarrow$ **Algorithm** 1(f, ρ, ε)
$S \leftarrow \max(i \ : \ \mu_i \geq \rho)$
if IsSeparated$(\delta, \{\mu_1 \geq \cdots \geq \mu_S\})$ **then**
  $h \leftarrow f_{\mathrm{reg}}$
else
  $h \leftarrow$ **RandomUnitVector**$(v_1, \dots, v_S)$
end
$f'_{\mathrm{reg}}, (\mu'_1, v'_1), \dots, (\mu'_{|\mathrm{Z}|}, v'_{|\mathrm{Z}|}) \leftarrow$ **Algorithm** 1(h, δ, ε)
$S' \leftarrow \max(i \ : \ \mu'_i \geq \delta)$
Output: $\{v'_i \ : \ i \in [S]\}$, (**IsSeparated**$(\delta, \{\mu'_1 \geq \cdots \geq \mu'_S\})$ **and** $|S| = |S'|$)

Algorithm 2: Quadratic character decomposition.

For a 1-bounded function $f : \mathrm{Z} \to \mathbb{C}$, Algorithm 2 finds a decomposition of f into quadratic characters (the v'_i in the first output) if the second output equals **True** (which happens with high probability under appropriate choices of parameters, e.g., $\delta = \rho^7$ as in Theorem 1.6). The following auxiliary routines are used: **IsSeparated**$(\delta, \{\mu_1 \geq \cdots \geq \mu_S\})$, which returns true only if $|\mu_i - \mu_{i+1}| \geq \delta$ for all i, and **RandomUnitVector**$(v_1, \dots, v_S)$, which returns a uniformly chosen random unit vector in the subspace $\langle v_1, \dots, v_S \rangle$ with respect to the L^2 norm.

Question 1.8. Instead of the *random* choice made in Theorem 7.24, is there a *deterministic* procedure to find a function $h \in \mathrm{Eigen}_\rho\big(\mathcal{K}_\varepsilon(f \otimes \overline{f})\big)$ satisfying the properties required for Theorem 1.6 to be applicable?

1.3. Nilspace theoretic foundations for the spectral method

From Section 4 onwards, our main proofs rely strongly on the nilspace approach to higher-order Fourier analysis, which is summarized in Figure 3. Nilspaces are algebraic structures that emerge naturally from the large scale behavior of functions on abelian groups relative to the Gowers norms. Despite having

[4]More precisely, we use the finer version of Theorem 1.1 that we actually prove, namely Theorem 7.12.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

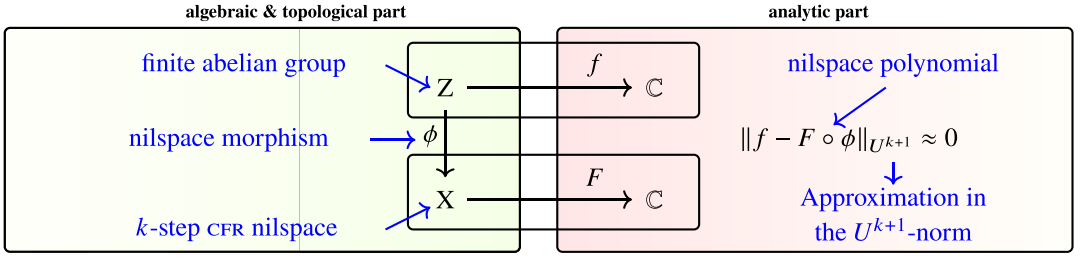


Figure 3. *The nilspace approach to k-th order Fourier analysis.*

been introduced relatively recently [5], nilspaces are treated in various works (see in particular [6, 7, 12, 37, 38, 39]). We shall recall some basic background on nilspace theory in Section 4.

The nilspace approach yields a regularity theorem for the Gowers norms which holds in particular on any finite abelian group, and which is a key ingredient in this paper [12, Theorem 1.5]. According to this theorem, a 1-bounded function f on a finite abelian group Z has a structured part of order k that is a so-called k-step *nilspace polynomial*, that is, a function of the form $F \circ \phi$, where $\phi : \mathrm{Z} \to \mathrm{X}$ is a nilspace morphism from Z (viewed as a 1-step nilspace) into a k-step nilspace X which is *compact and of finite rank* (in particular X is a finite-dimensional manifold), and F is a 1-bounded Lipschitz complex-valued function on X (see Definition 4.1). The pair (X, F) is in some sense a compressed version of (Z, f) which contains information on the k-th order structure of f at a given "resolution".

In this paper we refine this regularization technique by combining it with a generalized Fourier decomposition of F. To this end, we use the fact that the compact and finite-rank (CFR) nilspace X has the structure of an iterated *principal abelian bundle*, with *structure groups* $\mathrm{Z}_1(\mathrm{X}), \mathrm{Z}_2(\mathrm{X}), \dots, \mathrm{Z}_k(\mathrm{X})$ which are compact abelian Lie groups (see [7, Proposition 2.1.9]). The last structure group $\mathrm{Z}_k(\mathrm{X})$ acts on X freely by nilspace automorphisms. This can be used to decompose a Lipschitz function $F : \mathrm{X} \to \mathbb{C}$ uniquely as a series $F = \sum_{\chi \in \widehat{\mathrm{Z}_k(\mathrm{X})}} F_\chi$ converging uniformly. This yields the decomposition of the structured part of f as

$$F \circ \phi = \sum_{\chi \in \widehat{\mathrm{Z}_k(\mathrm{X})}} F_\chi \circ \phi. \tag{1.4}$$

Our main theoretical results describe how this additional decomposition of the structured part further enhances the framework provided by the nilspace regularity theorem [12, Theorem 1.5], especially when this decomposition is combined with the strong equidistribution property of ϕ (known as *balance*) given by this theorem. Below, we summarize the main results that we obtain in this direction, and their connection with our spectral approach.

Nilspace characters. In Section 4 we introduce the generalizations of Fourier characters mentioned previously in this introduction, namely *nilspace characters*, and we describe their fundamental properties. Nilspace characters are basically the functions of the form $F_\chi \circ \phi$ obtained in the decomposition (1.4) (for their formal description see Definition 4.5). These functions provide a nilspace theoretic generalization of Fourier characters, valid on any finite abelian group. In the specific setting of cyclic groups, they are closely related (though not identical) to the notion of *nilcharacter* from [35]. A central task in our study of nilspace characters is to connect this notion with the more elementary notion of a *character of order k* mentioned in the previous subsection (and formalized in Definition 2.29). Note that both notions are *parametric* (i.e., they depend on additional parameters, such as a *complexity* parameter). This is analogous to other concepts, such as being a "structured" or "quasirandom" function, which are also parametric, and become exact (i.e., parameter-free) only in a limit setting (see for instance [58, 59, 60]). A central result in Section 4 is Theorem 4.19, which focuses on the quadratic case and establishes, informally speaking, that 2-step nilspace characters are quadratic characters under an appropriate transformation of their parameters.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Decompositions with nilspace characters. In Section 5 we prove a refined version of the nilspace regularity theorem which implements the additional decomposition (1.4) of the structured part as a sum of nilspace characters; see Theorem 5.1. Moreover, using the balance property of the underlying morphism ϕ, this theorem guarantees that the nilspace characters satisfy approximate orthogonality properties which are crucial for our purposes. This new formulation of the nilspace regularity theorem also brings higher-order Fourier analysis more in line with classical Fourier analysis, in particular by yielding a new result constituting a generalization of Parseval's identity (see Theorem 5.7). A result of this kind has been called for in several works as a desirable tool for higher-order Fourier analysis (see in particular [25, Section 16], and [42, end of Section 1]). Another new result of this type is an approximate diagonalization formula for the Gowers norms in terms of nilspace characters (see Theorem 5.4).

The bridge between spectral decompositions and higher-order characters. In Section 6, we prove Theorem 6.1, which links the spectral properties of the operators $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ with the quadratic nilspace character components of f given by Theorem 5.1. An informal version of this theorem asserts that

$$\mathcal{K}_\varepsilon(f \otimes \overline{f}) = \sum_{\chi \in S} g_\chi \otimes \overline{g_\chi} + E$$

where ε is a small, appropriately chosen number, the functions $g_\chi := F_\chi \circ \phi$ are the dominant 2-step nilspace character components of f for some finite set S, and the matrix E is a small error. In particular, by the orthogonality properties of the functions g_χ, this theorem implies that each g_χ is a so-called *pseudoeigenvector* of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ (see Definition 7.1). As a consequence, one obtains that the corresponding pseudoeigenvalue is close to some proper eigenvalue λ. Furthermore, if λ is well separated from other eigenvalues, then the corresponding eigenvector v is close to g_χ. This phenomenon yields a direct connection between the eigenvectors of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ and the nilspace characters g_χ.

1.4. Connection with the ultralimit setting

To close this introduction, let us mention that the spectral approach to higher-order Fourier analysis is deeply motivated by an analogous construction in the ultralimit setting [58, 59, 60], which, on one hand, is more exact (not depending on various parameters), and on the other hand is intrinsically qualitative (in particular not yielding effective bounds). This construction involves measurable functions on ultraproducts of abelian groups. In particular, the Gowers norms become *semi*norms on such groups. In general, a major benefit of this ultralimit approach is its simpler and more algebraic language. This is illustrated by the following facts that hold in this setting: the notion of noise of order $k-1$ becomes exact, namely, it becomes the property of having U^k-seminorm equal to 0; then, every bounded measurable function f can be *uniquely* decomposed as $f = f_s + f_r$, where f_r is noise of order $k-1$ and f_s is a structured function of order $k-1$, that is, is orthogonal to any noise function of order $k-1$. Moreover, it can be proved that there is a sub-σ-algebra $\mathcal{F}_{k-1}$ on any such group such that a function g is $\mathcal{F}_{k-1}$-measurable if and only if g is structured of order $k-1$. Hence, the structured part f_s can be obtained as the conditional expectation $\mathbb{E}(f|\mathcal{F}_{k-1})$ relative to the σ-algebra $\mathcal{F}_{k-1}$. In this setting, the spectral approach simplifies to applying the map $g \mapsto \mathbb{E}(g|\mathcal{F}_{k-1})$ to each Z-diagonal function of $f \otimes \overline{f}$, yielding the unique self-adjoint operator $M(f,k)$.

Quite surprisingly, and nontrivially, the eigenvectors of $M(f,k)$ are structured functions of order k and their multiplicative derivatives are structured of order $k-1$ (see [58, Theorem 2] and [59]). This constitutes a limit-setting analogue of the notion of k-th order character introduced in this paper (Definition 2.29). Another useful property of $M(f,k)$ is that the k-th order structured part of f is equal to the projection of f onto the space spanned by the eigenvectors of $M(f,k)$ with nonzero eigenvalues. The resulting decomposition of the structured part of f into these eigenvectors is a higher-order generalization of the ordinary Fourier transform.

The ultralimit setting serves as a valuable framework for envisioning further results in higher-order Fourier analysis and deriving them at a qualitative level. This will be explored in future work in

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

connection with the spectral approach, especially in the general k-th order case. On the other hand, for the aim of making higher-order Fourier analysis practical, it becomes vital to introduce useful notions of a finite (usually parametric) type, and in this effort one is also aided by the deeper insights provided by the structures underlying higher-order decompositions, such as nilmanifolds and, more broadly, nilspaces. Hence this paper's focus on developing new connections between these structures, spectral properties, and higher-order decompositions.

2. Basic notions

Throughout this paper, by Z we will denote a finite abelian group, unless we explicitly say otherwise. Much of our work will involve the following special type of square matrices.

Definition 2.1 (Z-matrices). Let Z be a finite abelian group. A Z-*matrix* is a matrix in $\mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$. Given such a matrix M and an element $t \in \mathrm{Z}$, we call the set of pairs $\{(z+t, z) : z \in \mathrm{Z}\}$, viewed as a set of indices of entries of M, the Z-*diagonal* of M at level t (or the t-th Z-diagonal of M), and we denote it by $\mathrm{diag}(t)$. The Z-*diagonal function* of M at level t is the function

$$\mathfrak{D}_{M,t} : \mathrm{Z} \to \mathbb{C},\ z \mapsto M(z+t, z). \tag{2.1}$$

We then define the map $\mathfrak{D}_M : \mathrm{Z} \to L^2(\mathrm{Z})$ as $t \mapsto \mathfrak{D}_{M,t}$.

An important type of Z-matrix for us will be the rank-1 examples of the form $M = f \otimes \overline{f}$ for $f \in \mathbb{C}^{\mathrm{Z}}$, that is, $M(x, y) = f(x)\overline{f(y)}$ for $x, y \in \mathrm{Z}$. In this case, note that $\mathfrak{D}_{f\otimes\overline{f},t}(z) = \Delta_t f(z)$.

Remark 2.2. Thus Z-matrices have an abelian group structure on the index set of their rows and columns. The term "diagonal" is inspired by the case where Z is a finite cyclic group $\mathbb{Z}_N$, identified with the set of integers $\{0, 1, \dots, N-1\}$ with addition modulo N. In this case, labeling the rows and columns of M in order by $0, 1, \dots, N-1$, the Z-diagonals are indeed diagonal subsets of the matrix. For a general finite abelian group Z, we can fix an arbitrary ordering of the elements of Z (with 0 as the first element) and use this ordering for rows and columns. However, in general the Z-diagonals no longer look like diagonal subsets of the matrix (take, for instance, $\mathrm{Z} = \mathbb{Z}_2^2$).

We equip Z with its probability Haar measure (i.e., normalized counting measure). This induces the standard inner-product operation $\langle f, g\rangle := \mathbb{E}_{x\in\mathrm{Z}} f(x)\overline{g(x)}$ for any functions $f, g \in \mathbb{C}^{\mathrm{Z}}$, with the associated L^2-norm $\|f\|_{L^2} = \langle f, f\rangle^{1/2}$. In what follows, the notation $\|\cdot\|_2$ will always indicate this normalized L^2-norm when applied to functions in $\mathbb{C}^{\mathrm{Z}}$. When we wish to use instead the Euclidean norm of a vector without normalization, then we will signal this explicitly by writing $\|\cdot\|_{\ell^2}$ instead of $\|\cdot\|_2$ (thus $\|f\|_{\ell^2} = (\sum_{x\in\mathrm{Z}} |f(x)|^2)^{1/2}$).

It will be useful to view Z-matrices as kernels of integral linear operators on $L^2(\mathrm{Z})$ relative to the normalized Haar measure (in particular, in this sense the concept of Z-matrix extends naturally to compact abelian groups). However, this view entails an associated normalization of certain operations, which we shall use throughout the paper and which we underline as follows.

Definition 2.3. For any Z-matrix $M \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ and any function $f \in \mathbb{C}^{\mathrm{Z}}$, we define the function $Mf \in \mathbb{C}^{\mathrm{Z}}$ by $Mf(x) = \mathbb{E}_{y\in\mathrm{Z}} M(x, y) f(y)$. Accordingly, the product of two Z-matrices M, M' is the Z-matrix MM' defined by $(MM')(x, y) = \mathbb{E}_{z\in\mathrm{Z}} M(x, z)M'(z, y)$ (thus MM' is the kernel of the composition of the linear integral operators with kernels M, M'). The notions of eigenvector and eigenvalue of a Z-matrix are also normalized accordingly, thus a vector $v \in \mathbb{C}^{\mathrm{Z}}$ is an *eigenvector of M with eigenvalue* $\lambda \in \mathbb{C}$ if $\mathbb{E}_{y\in\mathrm{Z}} M(x, y)v(y) = \lambda v(x)$ for every $x \in \mathrm{Z}$. Finally, note that when we use the notation $\|\cdot\|_2$ with a Z-matrix, we always mean the version of the Euclidean norm (or Hilbert-Schmidt norm, or Schatten 2-norm) compatible with the normalized inner-product on $\mathbb{C}^{\mathrm{Z}}$ specified above, so $\|M\|_2 := (\mathbb{E}_{x,y\in\mathrm{Z}} |M(x, y)|^2)^{1/2}$.

We have seen above that from a Z-matrix we obtain a map $\mathfrak{D}_M : \mathrm{Z} \to L^2(\mathrm{Z})$ sending each element of Z to its corresponding Z-diagonal function in M. We can invert this map.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Definition 2.4. Given $F : \mathrm{Z} \to L^2(\mathrm{Z})$, we define the Z-matrix $\widetilde{M}(F)(x,y) := [F(x-y)](y)$.

In this notation $[F(x-y)](y)$, the square brackets enclose the *function* $F(x-y) \in L^2(\mathrm{Z})$, and the subsequent term (y) indicates the argument at which the function $F(x-y)$ is evaluated.

Lemma 2.5. *For any Z-matrix M we have $\widetilde{M}(\mathfrak{D}_M) = M$, and for any map $F : \mathrm{Z} \to L^2(\mathrm{Z})$ we have $\mathfrak{D}_{\widetilde{M}(F)} = F$.*

Proof. We have $\widetilde{M}(\mathfrak{D}_M)(x,y) = [\mathfrak{D}_M(x-y)](y) = [z \mapsto M(z+x-y,z)](y) = M(x,y)$. Conversely, for any $F : \mathrm{Z} \to L^2(\mathrm{Z})$ and $x, y \in \mathrm{Z}$, we have $[\mathfrak{D}_{\widetilde{M}(F)}(x)](y) = [z \mapsto \widetilde{M}(F)(z+x,z)](y) = \widetilde{M}(F)(x+y,y) = [F(x)](y)$. □

Next, we give an alternative description of the product of two Z-matrices M, M', consisting in writing the product in terms of the Z-diagonal functions of the matrices (rather than in terms of rows and columns as usual). To do so, let us view the Z-matrix M as the sum of the restrictions of M to its Z-diagonals: $M(x,y) = \sum_{t\in\mathrm{Z}} 1_{\mathrm{diag}(t)}(x,y)M(x,y)$.[5] Note that $1_{\mathrm{diag}(t)}(x,y)$ for $x, y \in \mathrm{Z}$ is a permutation matrix, so $1_{\mathrm{diag}(t)}(x,y)M(x,y)$ can be regarded as a weighted permutation Z-matrix, denoted M_t. We then have the following formula.

Lemma 2.6 (Product of weighted permutation Z-matrices). *Given two weighted permutation Z-matrices M_t and $M'_{t'}$, their product is a weighted permutation Z-matrix, supported on the Z-diagonal at level $t+t'$, with (x,y)-entry given by the following formula:*

$$M_t M'_{t'}(x,y) = \tfrac{1}{|\mathrm{Z}|} 1_{\mathrm{diag}(t+t')}(x,y) M(x,x-t)M'(y+t',y). \tag{2.2}$$

Proof. We have

$$M_t M'_{t'}(x,y) = \mathbb{E}_{v\in\mathrm{Z}} M_t(x,v)M'_{t'}(v,y) = \mathbb{E}_{v\in\mathrm{Z}} 1_{\mathrm{diag}(t)}(x,v)M(x,v)1_{\mathrm{diag}(t')}(v,y)M'(v,y).$$

Here $1_{\mathrm{diag}(t)}(x,v)1_{\mathrm{diag}(t')}(v,y) = 1(x=v+t)1(v=y+t'=x-t)$, so the last average above equals

$$\begin{aligned} &1_{\mathrm{diag}(t+t')}(x,y)\mathbb{E}_{v\in\mathrm{Z}} 1(x=v+t)M(x,v)M'(v,y) \\ &\qquad = \tfrac{1_{\mathrm{diag}(t+t')}(x,y)}{|\mathrm{Z}|} M(x,x-t)M'(x-t,y) = \tfrac{1_{\mathrm{diag}(t+t')}(x,y)}{|\mathrm{Z}|} M(x,x-t)M'(y+t',y), \end{aligned}$$

as claimed in (2.2). □

Now we can give the formula for multiplication of Z-matrices via Z-diagonals.

Proposition 2.7. *Let Z be a finite abelian group and let $M, M' \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$. Then for every $w \in \mathrm{Z}$ we have*

$$\mathfrak{D}_{MM',w}(z) = \mathbb{E}_{t\in\mathrm{Z}}\, \mathfrak{D}_{M,t}(z+w-t)\, \mathfrak{D}_{M',w-t}(z). \tag{2.3}$$

Proof. Using (2.2) we have that

$$\begin{aligned} MM'(z+w,z) &= \big(\textstyle\sum_{t\in\mathrm{Z}} M_t(z+w,z)\big)\big(\sum_{t'\in\mathrm{Z}} M'_{t'}(z+w,z)\big) = \sum_{t,t'\in\mathrm{Z}} M_t M'_{t'}(z+w,z) \\ &= \textstyle\sum_{t,t'\in\mathrm{Z}} \frac{1}{|\mathrm{Z}|} 1_{\mathrm{diag}(t+t')}(z+w,z)M(z+w,z+w-t)M'(z+t',z) \\ &= \tfrac{1}{|\mathrm{Z}|} \textstyle\sum_{t,t'\in\mathrm{Z}: t+t'=w} M(z+w,z+w-t)M'(z+t',z) \\ &= \mathbb{E}_{t\in\mathrm{Z}} M(z+w,z+w-t)M'(z+w-t,z), \end{aligned}$$

and this equals $\mathbb{E}_{t\in\mathrm{Z}}\, \mathfrak{D}_{M,t}(z+w-t)\, \mathfrak{D}_{M',w-t}(z)$, as claimed. □

[5] The function $1_{\mathrm{diag}(t)}(x,y)$ is the characteristic function of the t-th diagonal, i.e., $1_{\mathrm{diag}(t)}(x,y) = 1$ if $x-y=t$ and 0 otherwise.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Remark 2.8. Proposition 2.7 is useful in particular because it implies that if $\mathcal{A}$ is a shift-invariant algebra of functions (i.e., closed under addition, multiplication, and shifting) and every Z-diagonal of M and M' is a function in $\mathcal{A}$, then this also holds for the Z-matrix product MM'. (This observation is used for instance in Lemma 2.48.)

Remark 2.9. Among the notions introduced above and in what follows, many are easily extended to all compact abelian groups, and some even to locally compact abelian groups. Similarly, most proofs extend to compact abelian groups. We have nevertheless restricted our notation and treatment to finite abelian groups Z, because the main applications of algorithmic type concern primarily such groups, but also because this restriction simplifies the treatment markedly, avoiding various analytic issues related to measurability or convergence (especially once we start using Fourier analysis below).

2.1. Invariant operators

In what follows we shall often refer to maps $\mathbb{C}^{\mathrm{Z}} \to \mathbb{C}^{\mathrm{Z}}$ as *operators* (which will not necessarily be linear). We will apply such operators to Z-matrices along the Z-diagonals, and the main operators that we shall apply have a specific invariance property, defined as follows.

Definition 2.10. Let Z be a finite abelian group, let K be any operator $\mathbb{C}^{\mathrm{Z}} \to \mathbb{C}^{\mathrm{Z}}$, and let M be a Z-matrix. We denote by $\mathcal{K}(M)$ the Z-matrix $\widetilde{M}(K \circ \mathfrak{D}_M) = \widetilde{M}(K(\mathfrak{D}_M))$. For a function $f : \mathrm{Z} \to \mathbb{C}$ and $h \in \mathrm{Z}$, let $T^h f$ denote the function $x \mapsto f(x+h)$. We say that the operator K is *invariant* if $K(T^h f) = T^h K(f)$ and $K(\overline{f}) = \overline{K(f)}$, and we then say that the associated matrix operator $\mathcal{K}$ is also *invariant*.

Remark 2.11. Here and throughout the paper, note that given a map $K : \mathbb{C}^{\mathrm{Z}} \to \mathbb{C}^{\mathrm{Z}}$, for any Z-matrix M we use the *calligraphic* notation $\mathcal{K}(M)$ to denote the matrix obtained by applying K to each Z-diagonal of M.

We shall apply such invariant operators to self-adjoint matrices M, such as rank-1 matrices of the form $f \otimes \overline{f}$ for $f \in \mathbb{C}^{\mathrm{Z}}$. The following lemma shows that invariance of K suffices to ensure that $\mathcal{K}(M)$ is also self-adjoint.

Lemma 2.12. *If $M \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ is self-adjoint and K is invariant, then $\mathcal{K}(M)$ is self-adjoint.*

Proof. We need to show that for all $x, y \in \mathrm{Z}$ we have $\mathcal{K}(M)(x,y) = \overline{\mathcal{K}(M)(y,x)}$. By definition we have $\mathcal{K}(M)(x,y) = \widetilde{M}(K(\mathfrak{D}_M))(x,y) = [K(\mathfrak{D}_M)(x-y)](y) = [K(\mathfrak{D}_{M,x-y})](y)$. Recall that $\mathfrak{D}_{M,x-y}(z) = M(z+x-y,z) = \overline{M(z,z+x-y)}$, where in the last equality we used that M is self-adjoint. Let $t \in \mathrm{Z}$. Applying the shift operator T^t to this function, we have $T^t\mathfrak{D}_{M,x-y}(z) = \overline{M(z+t,z+t+x-y)}$. Letting $t := y-x$, we further obtain $T^{y-x}\mathfrak{D}_{M,x-y}(z) = \overline{\mathfrak{D}_{M,y-x}}$.

Since K is invariant, for any $t \in \mathrm{Z}$ and $f \in L^2(\mathrm{Z})$, we have $T^{-t}K(T^t f) = K(f)$. In particular, $\mathcal{K}(M)(x,y) = [K(\mathfrak{D}_{M,x-y})(z)](y) = [T^{-t}K(T^t\mathfrak{D}_{M,x-y}(z))](y) = [K(T^t\mathfrak{D}_{M,x-y}(z))](y-t)$. Letting $t := y-x$, we have $\mathcal{K}(M)(x,y) = [K(\overline{\mathfrak{D}_{M,y-x}})(z)](x)$. As K commutes with complex conjugation, we have $[K(\mathfrak{D}_{M,x-y})(z)](x) = \overline{[K(\mathfrak{D}_{M,y-x})(z)](x)}$, and this is by definition equal to $\overline{\mathcal{K}(M)(y,x)}$. □

Using various invariant operators $K : \mathbb{C}^{\mathrm{Z}} \to \mathbb{C}^{\mathrm{Z}}$ on the diagonals, and then taking the spectral decomposition of the resulting matrices, we obtain eigenvectors with various properties. Let us fix such an operator K and let $f : \mathrm{Z} \to \mathbb{C}$. As noted earlier, the matrix $f \otimes \overline{f}$ is a rank-1 self-adjoint matrix. Let $Q := \mathcal{K}(f \otimes \overline{f})$. Since Q is Hermitian by Lemma 2.12, we can apply the classical spectral theorem, obtaining the decomposition $Q = \sum_{i=1}^{n} \lambda_i\, v_i \otimes \overline{v_i}$, where $n = |\mathrm{Z}|$, the v_i are eigenvectors of Q forming an orthonormal basis of $\mathbb{C}^{\mathrm{Z}}$, and for each $i \in [n]$ the real scalar λ_i is the eigenvalue of Q corresponding to v_i. We can also assume (by permuting if necessary) that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$. The general idea of our spectral approach is that eigenvectors of Q corresponding to the largest eigenvalues are interesting "components" of the function f. The coefficients $c_i := \langle f, v_i \rangle$ are also important, and their relation with the eigenvalues λ_i is often nontrivial.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

In this section, we give various more detailed illustrations of this phenomenon, while keeping the involved machinery at a relatively simple level (in particular, not yet requiring deeper background such as nilspace theory).

2.2. A simple instance of the spectral approach: recovering classical Fourier analysis

As a first illustration of our use of invariant operators, let us detail here the remark made in the introduction, namely that if we apply to $f \otimes \overline{f}$ the operator consisting in simply averaging out each Z-diagonal, then the spectral decomposition of the resulting matrix (over $\mathbb{C}$) yields essentially the Fourier decomposition of f.

To see this in more detail, let K be the invariant operator consisting in averaging the function $f : \mathrm{Z} \to \mathbb{C}$ over the group Z, that is $K(f) = \mathbb{E}_{\mathrm{Z}}(f)$. Note that K is the orthogonal projection to the subspace of constant functions in $\mathbb{C}^{\mathrm{Z}}$. In this case, one can easily show that the classical Fourier characters $\chi \in \widehat{\mathrm{Z}}$ are eigenvectors of $\mathcal{K}(f \otimes \overline{f})$, and the eigenvalue corresponding to a character χ is the squared magnitude of the corresponding Fourier coefficient, that is, $|\widehat{f}(\chi)|^2$. Indeed, using the standard Fourier expansion $f = \sum_{\chi \in \widehat{\mathrm{Z}}} \langle f, \chi \rangle \chi$, we have

$$\begin{aligned}\mathcal{K}(f \otimes \overline{f})(x, y) &= \mathbb{E}_{z \in \mathrm{Z}} f(x+z)\overline{f(y+z)} = \mathbb{E}_{z \in \mathrm{Z}} \sum_{\chi \in \widehat{\mathrm{Z}}} \langle f, \chi \rangle \chi(x+z) \overline{\sum_{\chi' \in \widehat{\mathrm{Z}}} \langle f, \chi' \rangle \chi'(y+z)} \\ &= \sum_{\chi, \chi' \in \widehat{\mathrm{Z}}} \langle f, \chi \rangle \overline{\langle f, \chi' \rangle} \mathbb{E}_{z \in \mathrm{Z}} \chi(x+z) \overline{\chi'(y+z)} = \sum_{\chi \in \widehat{\mathrm{Z}}} |\langle f, \chi \rangle|^2 \chi(x) \overline{\chi(y)}.\end{aligned}$$

Remark 2.13. Note that even though the previous calculations tell us that Fourier characters are eigenvectors, we may not be able to recover all of them from the previous decomposition using only spectral analysis. To illustrate this, consider $f = \chi + \chi'$ for two Fourier characters $\chi \neq \chi'$ on Z. Then, looking at the spectrum of Q as described above, we will just find an eigenspace generated by χ and χ' for the eigenvalue 1, and another eigenspace of eigenvalue 0 generated by all other characters. Thus, if we want to recover the Fourier decomposition using this procedure, then we need a way of finding individual characters (such as χ, χ' here) inside eigenspaces of dimension greater than 1. The same phenomenon can occur in the higher-order versions of this spectral approach, and will be addressed in Subsection 7.2.

Remark 2.14. This example, involving the averaging operator K, enables us to highlight the connection between our approach and principal component analysis. More precisely, note that the operator $\mathcal{K}(f \otimes \overline{f})$ in this case is equal to the Gram matrix of the shifts of f, which can also be viewed as a covariance matrix. Thus, this operator is positive semidefinite, and the eigenvectors corresponding to dominant eigenvalues of this operator are the principal components of the set of shifts of f.

In what follows, we shall apply more subtle invariant operators than the averaging K in the above example. In particular, by applying an operator which keeps only the large Fourier coefficients of each Z-diagonal of $f \otimes \overline{f}$ (roughly speaking), we will see that the resulting dominant eigenvectors are important *quadratic*-Fourier-analytic components of f. This is the first nonclassical level of the general *order-increment principle* mentioned in the introduction. To detail all of this, we first need to discuss more precisely the notions of higher-order structures and characters that we will be able to capture in the spectral approach. We do this in the next two subsections.

2.3. Structured functions of order k

An important principle in higher-order Fourier analysis is the dichotomy between structure and randomness relative to the Gowers norms. A 1-bounded function $f : \mathrm{Z} \to \mathbb{C}$ is said to be *quasirandom of order* k if $\|f\|_{U^{k+1}} \leq \delta$ for some small $\delta > 0$, and a more precise aspect of the dichotomy is the decomposition of a function into a sum of a quasirandom part and a structured part, where *structure of order* k can be defined in various interrelated ways based on the idea of being orthogonal to all quasirandom functions of order k. In infinite settings, such as the ergodic theoretic one [43], or the ultraproduct setting [58], this

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

orthogonality can be defined exactly, and the notion of structure is captured in an exact way by certain subspaces of a Hilbert space, corresponding to certain σ-algebras or *factors* (see also Remark 2.28). However, in the setting of finite abelian groups, which we focus on in this paper, the notion of higher-order structure takes *approximate* and *quantitative* (or *parametric*) forms. In this subsection we give the precise definition of *structured function of order k* that we shall use in the rest of the paper. We discuss this first for order 1, and then generalize to higher orders.

For order 1 (i.e., in the setting of classical Fourier analysis), we use the following notion.

Definition 2.15 (Fourier-structured functions). Let Z be a finite abelian group. A function $f : \mathrm{Z} \to \mathbb{C}$ is (R, δ)-*Fourier-structured* if there is a function $g : \mathrm{Z} \to \mathbb{C}$ that is a linear combination of at most R Fourier characters on Z and that satisfies $\|f - g\|_2 \leq \delta$.

Remark 2.16. This property is almost identical to the notion of quantitative *almost-periodicity* from [66, Definition 10.34], but note that we do not require the coefficients in the linear combination to be 1-bounded. For functions f satisfying $\|f\|_2 \leq 1$, the two properties are equal. Indeed, if f is (R, δ)-Fourier-structured and $\|f\|_2 \leq 1$, then we can see that f is (R, δ)-almost-periodic as follows. Letting $g = \sum_{i=1}^{R} c_i \chi_i$ be the linear combination of characters χ_i with $\|f - g\|_2 \leq \delta$, we know that the orthogonal projection to the subspace spanned by these characters, namely $g' = \sum_{i=1}^{R} \widehat{f}(\chi_i)\chi_i$, satisfies $\|f - g'\|_2 \leq \delta$. Moreover, it follows from Parseval's identity that $|\widehat{f}(\chi_i)| \leq \|f\|_2 \leq 1$ for each i, and f is thus (R, δ)-almost-periodic as per [66]. Naming this property in terms of *structure* rather than almost-periodicity, as we do, is motivated by the developments below, which will introduce higher-order generalizations of this property that have a more direct connection with the notion of structure (as in the structure-randomness dichotomy) than with almost-periodicity.

Before we motivate Definition 2.15 further, let us recall other known tools to measure the Fourier structure of a function. These involve certain norms. A central example is the Wiener norm (also known as the algebra norm, or the spectral norm; see [30, 64]), which we recall here.

Definition 2.17. The *Wiener norm* of a function $f : \mathrm{Z} \to \mathbb{C}$ is $\|f\|_{A(\mathrm{Z})} := \|\widehat{f}\|_{\ell^1}$.

Being (R, δ)-Fourier-structured is closely related to being bounded in the Wiener norm. We first show that the latter property implies the former, using the next lemma.

Recall that for a metric space X and a function $f : X \to \mathbb{C}$, the *support* of f, denoted $\mathrm{Supp}(f)$, is the closure of the set $\{x \in X : f(x) \neq 0\}$.

Lemma 2.18. *Let S be a finite set, let $f : S \to \mathbb{C}$, and let $\delta > 0$. Let $g : S \to \mathbb{C}$ be the function $g(s) = f(s)\,\mathbf{1}(|f(s)| \geq \delta)$. Then $|\mathrm{Supp}(g)| \leq \delta^{-1}\|f\|_{\ell^1(S)}$ and $\|f - g\|_{\ell^2(S)} \leq \delta\|f\|_{\ell^1(S)}$.*

Proof. It is clear that $\delta\,|\mathrm{Supp}(g)| \leq \|g\|_{\ell^1} \leq \|f\|_{\ell^1}$, which proves the first claim. On the other hand, by Hölder's inequality we have $\|f - g\|_{\ell^2} \leq \|f - g\|_{\ell^\infty}\|f - g\|_{\ell^1} \leq \delta\|f\|_{\ell^1}$. □

Corollary 2.19. *For any $\delta > 0$, any function $f : \mathrm{Z} \to \mathbb{C}$ is $(\delta^{-1}\|f\|^2_{A(\mathrm{Z})}, \delta)$-Fourier-structured.*

Proof. Applying Lemma 2.18 with $\widehat{f}$ and $\delta > 0$, we obtain that there is a function $\widehat{g} : \widehat{\mathrm{Z}} \to \mathbb{C}$ whose support has size at most $\delta^{-1}\|f\|_{A(\mathrm{Z})}$ and such that $\|\widehat{f} - \widehat{g}\|_{\ell^2} \leq \delta$. The reverse Fourier transform g of $\widehat{g}$ is a linear combination of at most $\delta^{-1}\|f\|_{A(\mathrm{Z})}$ characters, and satisfies $\|f - g\|_2 = \|\widehat{f} - \widehat{g}\|_{\ell^2} \leq \delta$. Thus f is $(\delta^{-1}\|f\|_{A(\mathrm{Z})}, \delta\|f\|_{A(\mathrm{Z})})$-Fourier-structured, whence, changing $\delta\|f\|_{A(\mathrm{Z})}$ to δ, the result follows. □

Remark 2.20. More generally, if for some $\alpha \in (0, 1]$ the norm $\|\widehat{f}\|_{\ell^{2-\alpha}}$ is bounded above by some constant C, then for every $\delta > 0$ it follows similarly that f is $(C/\delta^{2-\alpha}, \delta^\alpha C)$-Fourier-structured. In particular, this is the case for the U^2-dual norm, since we have the well-known formula $\|f\|^*_{U^2} = \|\widehat{f}\|_{\ell^{4/3}}$.

Let us now consider the converse direction, that is, whether (R, δ)-Fourier-structure implies being bounded in the Wiener norm or in the U^2-dual norm. We first recall the following example, which rules out a naive converse to Corollary 2.19.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Example 2.21. Let $\mathbb{Z}_N$ be the finite cyclic group of order N, and let $f : \mathbb{Z}_N \to \{0,1\}$ be the indicator function of the interval $[1, \lfloor N/2 \rfloor]$. We have the well-known fact that $\|f\|_{A(\mathrm{Z})} = \Omega(\log N)$, whereas any norm $\|\widehat{f}\|_{\ell^p}$ with $p > 1$ is bounded independently of N. In particular $\|f\|_{U^2}^*$ is bounded, so by Remark 2.20 the function f is $(O(\delta^{-2/3}, O(\delta^{4/3}))$-Fourier-structured for every δ, independently of N.

Notwithstanding this example, we have the following equivalence result, providing an alternative expression for Fourier structure, using the U^2-dual norm instead of the Fourier transform.

Lemma 2.22. *Consider the following properties of functions $f : \mathrm{Z} \to \mathbb{C}$.*

(i) *For some $R > 0$ and $\delta \in [0,1]$, the function f is (R,δ)-Fourier-structured.*
(ii) *For some $R' > 0$ and $\delta' \in [0,1]$, there exists $g : \mathrm{Z} \to \mathbb{C}$ with $\|g\|_{U^2}^* \le R'$ and $\|f - g\|_2 \le \delta'$.*

These properties are equivalent in the following sense: if (i) holds with (R,δ) then (ii) holds with $\delta' = \delta$ and $R' = R^{3/4}\|f\|_2$; conversely, if (ii) holds with (R',δ') then, for every $\alpha > 0$, (i) holds with $R = \alpha^{-6}R'^4(\|f\|_2 + 1)^2$ and $\delta = \alpha + \delta'$.

Proof. If (i) holds, then, letting g be the projection of f to the linear span of the given R characters, we know that $\|f - g\|_2 \le \delta$, and g has Fourier coefficients of modulus at most $\|f\|_2$, so $\|g\|_{U^2}^* \le R^{3/4}\|f\|_2$, whence (ii) holds as claimed.

If (ii) holds, then we claim that for any $\gamma > 0$ the function g is $(\gamma^{-2}(\|f\|_2 + 1)^2, \gamma^{1/3}R'^{2/3})$-Fourier structured. Indeed, let $S = \{\chi \in \widehat{\mathrm{Z}} : |\widehat{g}(\chi)| > \gamma\}$, and define $h(x) := \sum_{\chi \in S} \widehat{g}(\chi)\chi(x)$. By Plancherel's theorem we have $\|g - h\|_2^2 = \sum_{\chi \notin S} |\widehat{g}(\chi)|^2 \le \gamma^{2/3}\,\|\widehat{g}\|_{\ell^{4/3}}^{4/3} \le \gamma^{2/3}R'^{4/3}$. We also have $\gamma^2|S| \le \|g\|_2^2 \le (\|f\|_2 + \delta')^2$. Our claim follows. Then the original function f is $(\gamma^{-2}(\|f\|_2+1)^2, \gamma^{1/3}R'^{2/3} + \delta')$-Fourier structured. Choosing γ so that $\gamma^{1/3}R'^{2/3} = \alpha$, we obtain that (i) holds as claimed. □

The statement of property (ii) in Lemma 2.22 clearly suggests the following natural generalization to higher orders. This is the main notion of higher-order structure used in this paper.

Definition 2.23 (Structured functions of order k). Let k be a positive integer. We say that a function $f : \mathrm{Z} \to \mathbb{C}$ is (R,δ)-*structured of order k* if there is a function $g : \mathrm{Z} \to \mathbb{C}$ satisfying $\|g\|_{U^{k+1}}^* \le R$ and $\|f - g\|_2 \le \delta$.

Remark 2.24. Since the U^k norms form an increasing sequence (i.e., $\|f\|_{U^k} \le \|f\|_{U^{k+1}}$ for all k), the U^k-dual norms form a decreasing sequence. It follows that (R,δ)-structure of order k becomes a weaker (more inclusive) property as k increases.

Remark 2.25. There are other higher-order parametric notions of structure, for instance the one based on the *uniform almost-periodicity norms* introduced by Tao in [64, Definition 5.2], which proceeds by generalizing the Wiener norm instead of the U^2-dual norm. The notion in Definition 2.23 turned out to be sufficiently natural and technically convenient for us.

We know, by Lemma 2.22, that structure of order 1 means proximity in L^2 to a combination of a few Fourier characters. Before we develop further aspects of structure of order k for general $k > 1$, let us observe that structure of order 0 also has a useful meaning, even though this case is peculiar in the sense that $\|f\|_{U^1} := |\mathbb{E}_{x \in \mathrm{Z}} f(x)|$ is a seminorm and not a norm. Indeed, let us say that a function $f : \mathrm{Z} \to \mathbb{C}$ is (R,δ)-*structured of order 0* if there exists $g : \mathrm{Z} \to \mathbb{C}$ such that $\|f - g\|_2 \le \delta$ and such that the seminorm $\|g\|_{U^1}^* := \sup_{h:\mathrm{Z}\to\mathbb{C},\, \|h\|_{U^1} \le 1} |\langle g, h\rangle|$ is at most R.

Lemma 2.26. *A function $f : \mathrm{Z} \to \mathbb{C}$ is (R,δ)-structured of order 0 for some R if and only if there is a constant function g such that $\|f - g\|_2 \le \delta$.*

Proof. For the forward implication, let g be such that $\|f - g\|_2 \le \delta$ and $\|g\|_{U^1}^* \le R$. Note that, for each nontrivial character $\chi \in \widehat{\mathrm{Z}} \setminus \{\mathrm{id}\}$, we have $\|\chi\|_{U^1} = 0$. Hence, for every integer $n \ge 1$, we have $\|n\chi\|_{U^1} \le 1$ and therefore $n|\widehat{g}(\chi)| = |\langle g, n\chi\rangle| \le \|g\|_{U^1}^* \le R$, whence $|\widehat{g}(\chi)| \le R/n$. Taking the limit as $n \to \infty$ we deduce that $\widehat{g}(\chi) = 0$ for all $\chi \ne \mathrm{id}$, so g is constant.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

For the converse, note that if g is constant and $\|f-g\|_2 \le \delta$, then $\|g\|_{U^1}^* = |g| \le \|f\|_2 + \delta$, so f is (R,δ)-structured of order 0 for $R = \|f\|_2 + \delta$. □

To elucidate further the *structure of order k* property from Definition 2.23, and for other uses later on, let us establish the following equivalent formulation.

Proposition 2.27. *Consider the following two properties of functions $f : \mathrm{Z} \to \mathbb{C}$.*

(i) *The function f is (R,δ)-structured of order k.*
(ii) *For every 1-bounded function $h : \mathrm{Z} \to \mathbb{C}$ satisfying $\|h\|_{U^{k+1}} \le a \le 1$, we have $|\langle f, h\rangle| \le b$.*

For 1-bounded functions $f : \mathrm{Z} \to \mathbb{C}$, these properties are equivalent in the following sense: if (i) holds with (R,δ) then (ii) holds with any a and $b = Ra + \delta$; conversely, if (ii) holds with (a,b), then (i) holds with some $R = O_a(1)$ and $\delta = 2(a^{1/2} + b^{1/2})$.

Proof. If (i) holds with (R,δ), then $|\langle f,h\rangle| \le |\langle g,h\rangle| + |\langle f-g,h\rangle| \le \|g\|_{U^{k+1}}^* \|h\|_{U^{k+1}} + \|f-g\|_2\|h\|_2 \le Ra + \delta$, so (ii) holds as claimed.

For the converse, suppose that (ii) holds with a, b. Let $\nu > 0$ (to be fixed later) and $\mathcal{G} : \mathbb{R}_+ \to \mathbb{R}_+$ be the function $\mathcal{G}(m) = \frac{\nu}{2(m+1)}$. By a standard regularity lemma, for example, [24, Corollary 5.2],[6] there is $R = R(\nu,\mathcal{G}) > 0$ such that we have a decomposition $f = f_s + f_e + f_r$ where $\|f_s\|_{U^{k+1}}^* \le R$, $\|f_e\|_2 \le \nu/2^{3/2}$, $\|f_r\|_{U^{k+1}} \le \mathcal{G}(\|f_s\|_{U^{k+1}}^*)$, and $|f_r| \le 2$. Note that $|\langle f, f_r\rangle| = |\langle f_s, f_r\rangle + \langle f_e, f_r\rangle + \|f_r\|_2^2| \ge \|f_r\|_2^2 - |\langle f_s, f_r\rangle| - |\langle f_e, f_r\rangle| \ge \|f_r\|_2^2 - \|f_s\|_{U^{k+1}}^* \frac{\nu}{2(\|f_s\|_{U^{k+1}}^*+1)} - 2^{1/2}\nu/2^{3/2} \ge \|f_r\|_2^2 - \nu$. Hence $\|f_r\|_2 \le (|\langle f, f_r\rangle| + \nu)^{1/2}$, so

$$\|f - f_s\|_2 \le \nu + (|\langle f, f_r\rangle| + \nu)^{1/2}. \tag{2.4}$$

Let $\nu = a$ (and so, $R = O_a(1)$) and note that, if $h = f_r/2$, we have $\|h\|_\infty \le 1$ and $\|h\|_{U^{k+1}} \le \frac{\nu}{4(\|f_s\|_{U^{k+1}}^*+1)} \le a$. Hence, applying (ii) with this function h, we have $|\langle f, f_r/2\rangle| \le b$, so by (2.4) we deduce that $\|f - f_s\|_2 \le a + (2b+a)^{1/2}$. Hence (i) holds with R and δ as claimed. □

Remark 2.28. Property (ii) in Proposition 2.27 is a formulation of k-th order structure which is strongly motivated from several viewpoints on higher-order Fourier analysis in infinite settings. Specifically, in the ergodic-theory setting from [43], as well as in the ultralimit setting of [58], or in the setting of cubic couplings [11], the notion of higher-order structure is captured by certain σ-algebras (the Host–Kra characteristic factors in [43] and their generalizations in [11], or the higher-order Fourier σ-algebras in [58]), where measurability relative to such a σ-algebra of order k can be defined as orthogonality to every function with vanishing U^{k+1}-seminorm (see [58, Theorem 1] or [11, Lemma 5.8]). Property (ii) above is a finite analogue of this.

2.4. Characters of order k

Here we use the notion of higher-order structured function from the previous subsection to define higher-order analogues of Fourier characters on abelian groups. In particular, quadratic (i.e., order 2) characters will play a key role in our algorithmic applications. By Z we continue to denote a finite abelian group, but note that many of the following notions still hold for more general compact abelian groups.

Definition 2.29. Let k be a positive integer. An (R,δ)-*character of order k* on a finite abelian group Z is a function $f : \mathrm{Z} \to \mathbb{C}$ with $\|f\|_2 \le 1$ such that for every $t \in \mathrm{Z}$, the multiplicative derivative $\Delta_t f$ is (R,δ)-structured of order $k-1$. We refer to R as the *complexity parameter*, and refer to δ as the *precision parameter*. When it is unnecessary to specify these parameters, we simply call f a *character of order k* on Z.

[6] Note that this result is stated for $\mathbb{R}$-valued functions. For $\mathbb{C}$-valued ones we can simply separate the real and imaginary parts and apply [24, Corollary 5.2] to each.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Many central examples of such characters are actually 1-bounded functions. The weaker bound $\|f\|_2 \le 1$ in the definition is useful from linear-algebraic and functional-analytic viewpoints.

Remark 2.30. Using that the U^k-dual norms decrease with k (see Remark 2.24), we see that for every function $f : \mathrm{Z} \to \mathbb{C}$ with $\|f\|_2 \le 1$, we have $\|f\|_{U^{k+1}}^* \le \|f\|_{U^2}^* \le |\mathrm{Z}|^{3/4}$, so any such function is a $(|\mathrm{Z}|^{3/4}, 0)$-character of order k. The notion of (R, δ)-character of order k is interesting only when Z is large, R is bounded and δ is small. In what follows, we will see examples and theorems involving such characters where R depends on δ but not on $|\mathrm{Z}|$.

Before we look at some basic examples of higher-order characters, let us briefly recall the notion of Gowers U^k-products [66, p. 419] and the Gowers-Cauchy-Schwarz inequality [66, (11.6)].

Definition 2.31 (Gowers U^k-product). Let Z be a finite abelian group, $k \ge 1$, and let $\{f_v \in \mathbb{C}^{\mathrm{Z}}\}_{v \in [\![k]\!]}$ be a set of functions (where $[\![k]\!]$ denotes $\{0,1\}^k$). The *Gowers U^k-product* $\langle (f_v)_{v \in [\![k]\!]} \rangle_{U^k}$ of these functions is defined as $\mathbb{E}_{x,t_1,\dots,t_k \in \mathrm{Z}} \prod_{v \in \{0,1\}^k} \mathcal{C}^{|v|} f_v(x + v(1)\, t_1 + \cdots + v(k)\, t_k)$. The *Gowers-Cauchy-Schwarz inequality* states that $|\langle (f_v)_{v \in [\![k]\!]} \rangle_{U^k}| \le \prod_{v \in [\![k]\!]} \|f_v\|_{U^k}$.

Example 2.32. The most basic example of a character of order k is a global polynomial phase function of degree k on Z, that is, a function $\phi : \mathrm{Z} \to \mathbb{C}$ such that $\Delta_{t_1} \cdots \Delta_{t_{k+1}} \phi(x) = 1$ for all $x, t_1, \dots, t_{k+1} \in \mathrm{Z}$. Indeed, such a function is in fact a $(1,0)$-character of order k. To see this, note that it follows from the definition of ϕ that $|\phi(x)| = 1$ for all $x \in \mathrm{Z}$. Then, for any fixed $t \in \mathrm{Z}$, note that $\Delta_t \phi$ is a polynomial phase function of degree $k-1$, so for every $x, t_1, \dots, t_k \in \mathrm{Z}$ we have $\overline{\Delta_t \phi(x)} = \prod_{v \in \{0,1\}^k \setminus 0^k} \mathcal{C}^{|v|} \Delta_t \phi(x + v(1)t_1 + \cdots + v(k)t_k)$ (where $\mathcal{C}$ denotes the complex-conjugation operator and $0^k = (0, \dots, 0)$). It follows that for every function $g : \mathrm{Z} \to \mathbb{C}$ with $\|g\|_{U^k} \le 1$, we have $\langle g, \Delta_t \phi \rangle = \langle g, \Delta_t \phi, \Delta_t \phi, \dots, \Delta_t \phi \rangle_{U^k} \le \|g\|_{U^k} \|\Delta_t \phi\|_{U^k}^{2^k - 1} \le 1$, whence $\|\Delta_t \phi\|_{U^k}^* \le 1$, which proves that $\Delta_t \phi$ is $(1,0)$-structured of order $k-1$, as required.

More complicated examples of higher-order characters involve nilpotent structures. In particular, nilsequences with so-called *vertical frequency*, called *nilcharacters* in [35, Definition 6.1], and their generalizations introduced in the present paper, called *nilspace characters*, provide important examples of higher-order characters (see Section 4).

To illustrate Definition 2.29 further, let us consider in more detail its special case for $k = 1$. This tells us that f is an (R, δ)-character of order 1 if for every $t \in \mathrm{Z}$ we have that $\Delta_t f$ is (R, δ)-structured of order 0. By Lemma 2.26, this implies that there is a constant function g_t such that $\|g_t - \Delta_t f\|_2 \le \delta$. If δ is sufficiently small, then this in turn implies indeed, as it should, that such a function f is close in L^2 to a scalar multiple of a Fourier character.

Proposition 2.33. *If $f : \mathrm{Z} \to \mathbb{C}$ is an (R, δ)-character of order 1 with $\delta \in (0, \|f\|_2/20]$, then there exists $\chi \in \widehat{\mathrm{Z}}$ such that $\|f - \widehat{f}(\chi)\chi\|_2 \le 4\delta^{1/2} \|f\|_2^{1/2}$.*

There are various ways to prove this by reducing to a simple instance of a so-called 99% *inverse theorem* for the U^2-norm (for more on this topic, see [16]). Let us give a short proof here.

Proof. By assumption, for every $t \in \mathrm{Z}$ there is a constant g_t such that $\|\Delta_t f - g_t\|_2 \le \delta$. Hence $\big| \|\Delta_t f\|_{U^1} - |g_t| \big| \le |\mathbb{E}(\Delta_t f - g_t)| \le \delta$. Then $\mathbb{E}_t |g_t| \le \|f\|_1^2 + \delta \le \|f\|_2^2 + \delta$, and, writing $h_t := \Delta_t f - g_t$, we have $\mathbb{E}_t |g_t|^2 = \mathbb{E}_t \mathbb{E}_x |g_t|^2 \ge \mathbb{E}_t \mathbb{E}_x (|\Delta_t f(x)|^2 - 2|h_t||\Delta_t f|) \ge \|f\|_2^4 - 2\delta \|f\|_2^2$. Hence $\|f\|_{U^2}^4 = \mathbb{E}_t \|\Delta_t f\|_{U^1}^2 \ge \mathbb{E}_t |g_t|^2 - 2\delta \mathbb{E}_t |g_t| + \delta^2 \ge \|f\|_2^4 - 4\delta \|f\|_2^2 - \delta^2$. Completing squares and rearranging, we deduce that $\|f\|_{U^2}^2 \ge \|f\|_2^2 - 5\delta$ (which is an assumption for a 99%-inverse-theorem). Let $x = \|\widehat{f}\|_\infty$. Using that $\|f\|_2 \ge x \ge \|f\|_{U^2}^2 / \|f\|_2$, we have $x(\|f\|_2 - x) \le \|f\|_2^2 - \|f\|_{U^2}^2 \le 5\delta$, so either $x \le \|f\|_2 \frac{1 - \sqrt{1 - 20\delta/\|f\|_2}}{2}$, or $x \ge \|f\|_2 \frac{1 + \sqrt{1 - 20\delta/\|f\|_2}}{2}$. The former case combined with $x \ge \|f\|_{U^2}^2 / \|f\|_2$ implies $\|f\|_{U^2}^2 \le 10\delta \|f\|_2$, and this combined with $\|f\|_2^2 \le \|f\|_{U^2}^2 + 5\delta$ implies $\|f\|_2 \le 5(\delta + \sqrt{\delta^2 + \delta/5}) \le 4\sqrt{\delta}$ (using that $\delta \le 1/20$), so the conclusion holds with any χ. The latter

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

case implies the conclusion with χ such that $|\widehat{f}(\chi)| = \|\widehat{f}\|_\infty$, as then $\|f - \widehat{f}(\chi)\chi\|_2^2 = \|f\|_2^2 - \|\widehat{f}\|_\infty^2 \leq 15\delta\|f\|_2$. □

This last result can be proved with better constants by a finer (but longer) argument using the Fourier coefficients of f. We omit this, as such an improvement is not needed in this paper.

It will be useful to introduce the following generalization of higher-order characters.

Definition 2.34 (Weak character of order k). An (R, δ_1, δ_2)-*weak character of order k* on a finite abelian group Z is a function $f : \mathrm{Z} \to \mathbb{C}$ with $\|f\|_2 \leq 1$ such that there is a set $S \subset \mathrm{Z}$ with $|S| \geq |\mathrm{Z}|(1-\delta_2)$ and such that for every $t \in S$ the function $\Delta_t f$ is (R, δ_1)-structured of order $k-1$. We refer to R as the *complexity parameter* and to δ_1, δ_2 as the *precision parameters* of f.

There is a fact that we must now establish, in order to connect properly the notions of higher-order structure and higher-order characters. Namely, we need to show that characters of order k are structured functions of order k. In fact, this holds more generally for *weak* characters of order k.

Lemma 2.35. *Let g be an (R, δ_1, δ_2)-weak-character of order k on a finite abelian group Z. Then for every 1-bounded function $f : \mathrm{Z} \to \mathbb{C}$ satisfying $\|f\|_{U^{k+1}} \leq a \leq 1$, we have $|\langle f, g\rangle| \leq (a^2 R + \delta_1 + \delta_2^{1/2})^{1/2}$. In particular g is an (R', δ')-structured function of order k, for $R' = O_a(1)$ and $\delta' = 2(a^{1/2}(R^{1/4}+1) + \delta_1^{1/4} + \delta_2^{1/8})$.*

Proof. By Definition 2.34, for a set $S \subset \mathrm{Z}$ of cardinality at least $(1-\delta_2)|\mathrm{Z}|$, for every $t \in S$ there is a function $\mathcal{E}_t : \mathrm{Z} \to \mathbb{C}$ with $\|\mathcal{E}_t\|_2 \leq \delta_1$ and a function $h_t : \mathrm{Z} \to \mathbb{C}$ with $\|h_t\|_{U^k}^* \leq R$, such that $\Delta_t g = h_t + \mathcal{E}_t$. Then, as $\|\Delta_t f\|_2 \leq 1$ and $(\mathbb{E}_t \|\Delta_t g\|_2^2)^{1/2} = \|g\|_2^2 \leq 1$, we have

$$\begin{aligned}
|\langle f, g\rangle|^2 &= \mathbb{E}_t \langle \Delta_t f, \Delta_t g\rangle = \mathbb{E}_t 1_S(t) \langle \Delta_t f, h_t\rangle + \mathbb{E}_t 1_S(t) \langle \Delta_t f, \mathcal{E}_t\rangle + \mathbb{E}_t 1_{S^c}(t) \langle \Delta_t f, \Delta_t g\rangle \\
&\leq \mathbb{E}_t 1_S(t) \|\Delta_t f\|_{U^k} \|h_t\|_{U^k}^* + \mathbb{E}_t 1_S(t) \|\Delta_t f\|_2 \|\mathcal{E}_t\|_2 + \mathbb{E}_t 1_{S^c}(t) \|\Delta_t f\|_2 \|\Delta_t g\|_2 \\
&\leq \mathbb{E}_t 1_S(t) \|\Delta_t f\|_{U^k} \|h_t\|_{U^k}^* + \delta_1 \mathbb{E}_t \|\Delta_t f\|_2 + (\mathbb{E}_t 1_{S^c}(t))^{1/2} (\mathbb{E}_t \|\Delta_t g\|_2^2)^{1/2}. \\
&\leq \mathbb{E}_t 1_S(t) \|\Delta_t f\|_{U^k} \|h_t\|_{U^k}^* + \delta_1 + \delta_2^{1/2}.
\end{aligned}$$

By Hölder's inequality, letting $q = 2^k/(2^k-1)$, we have that $\mathbb{E}_t 1_S(t) \|\Delta_t f\|_{U^k} \|h_t\|_{U^k}^*$ is at most $(\mathbb{E}_t \|\Delta_t f\|_{U^k}^{2^k})^{1/2^k} (\mathbb{E}_t 1_S(t) \|h_t\|_{U^k}^{*\,q})^{1/q} = \|f\|_{U^{k+1}}^2 (\mathbb{E}_t 1_S(t) R^q)^{1/q}$. Hence $|\langle f, g\rangle|^2 \leq a^2 R + \delta_1 + \delta_2^{1/2}$, as claimed. By Proposition 2.27, it follows that g is (R', δ')-structured of order k with $R' = O_a(1)$ and $\delta' = 2(a^{\frac{1}{2}} + b^{\frac{1}{2}})$, where $b = (a^2 R + \delta_1 + \delta_2^{\frac{1}{2}})^{\frac{1}{2}} \leq a R^{\frac{1}{2}} + \delta_1^{\frac{1}{2}} + \delta_2^{\frac{1}{4}}$. □

Remark 2.36. It is now natural to wonder how special the characters of order k are among the structured functions of order k. We have seen in Proposition 2.33 that characters of order 1 are essentially scalar multiples of Fourier characters (up to small L^2-errors), and these order-1 characters are indeed special among structured functions of order 1, we may even say that they are *fundamental* in the sense that every structured function of order 1 is a combination of a bounded number of characters of order 1 (up to a small L^2-error), by Fourier analysis. This picture generalizes to higher orders; we shall discuss this in Section 4 using the concept of *nilspace character*. Specifically, we will show that every structured function of order k is, up to a small L^2-error, a sum of a bounded number of k-step nilspace characters (this will follow from the regularity lemma Theorem 5.1). Moreover, focusing on the quadratic case, we will also show that 2-step nilspace characters are characters of order 2 as per Definition 2.29. We expect this latter fact to extend to general order k, but as explained in later sections, proving this is outside the scope of this paper.

As mentioned in the introduction, we aim to provide algorithms to obtain decompositions of functions into structured and quasirandom parts of order 2 (i.e., relative to the U^3-norm). A key step in our approach consists in modifying Z-matrices by applying an operator to the Z-diagonals which performs a *Fourier regularization* of these diagonals, that is, it isolates a suitable Fourier-structured part of these functions. We shall now introduce this operator and develop some of its main properties.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

2.5. The Fourier denoising operator K_ε

In this subsection, we define the main notion of Fourier regularization that we shall use to develop our spectral approach in quadratic Fourier analysis.

Informally speaking, the goal is to define a notion on finite abelian groups satisfying properties similar to those of the expectation operator corresponding to the first-order Fourier σ-algebra $\mathcal{F}_1$ in the ultraproduct setting, or the Kronecker factor in ergodic theory (see Remark 2.28). In other words, we want a tool that decomposes a function on a finite abelian group into a structured and a quasirandom part of order 1 (i.e., relative to classical Fourier analysis), and which satisfies certain additional properties that are useful from algorithmic viewpoints. In the infinite setting of ultraproducts (see [58], or later [12]), the regularization has the form $f = f_s + f_r$ where $f_s = \mathbb{E}(f|\mathcal{F}_1)$ while $f_r = f - f_s$. The function f_r is also the projection of f to the orthogonal complement of $L^2(\mathcal{F}_1)$. Thus, both operators $f \mapsto f_s$ and $f \mapsto f_r$ are projections in the infinite setting, so both operators are contractions relative to the L^2-norm. In particular, we have the continuity properties

$$\|f - f'\|_2 \geq \|f_s - f'_s\|_2 \text{ and } \|f - f'\|_2 \geq \|f_r - f'_r\|_2.$$

In the finite setting, a first candidate for the desired tool could be the Fourier cut-off operator mentioned in Subsection 1.1, namely the operator which, for a fixed $\varepsilon > 0$, simply eliminates the Fourier coefficients of f that have modulus less than ε. Unfortunately, this operator is too blunt and does not have suitable features such as the above contraction property. However, there is a variant of this operator that does have such properties. Before we describe it, let us mention another very useful fact that holds in infinite settings: the *translation invariance* whereby we have $(T^h f)_s = T^h(f_s)$ and $(T^h f)_r = T^h(f_r)$ for every group element h. To guarantee this property, we consider operators of the following form.

Definition 2.37. Let $r : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be an arbitrary function with $r(0) = 0$. For $z \in \mathbb{C}$ we define $q_r(z) := z\, r(|z|)/|z|$ for $z \neq 0$ and $q_r(0) = 0$. We then define an operator K_r on $\mathbb{C}^{\mathbb{Z}}$ as follows: for any function $f \in \mathbb{C}^{\mathbb{Z}}$, letting $f = \sum_{\chi \in \widehat{\mathbb{Z}}} \widehat{f}(\chi)\chi$ be its Fourier expansion, we define

$$K_r(f) := \sum_{\chi \in \widehat{\mathbb{Z}}} q_r(\widehat{f}(\chi))\, \chi. \tag{2.5}$$

Since translation only affects the phase of the Fourier coefficients, we have $T^h K_r(f) = K_r(T_h f)$. It is also clear that $\overline{K_r(f)} = K_r(\overline{f})$. Hence, K_r is invariant as per Definition 2.10.

We now want to identify a simple property of functions r which ensures that the associated operator K_r has the desired continuity property.

Recall that a function $g : \mathbb{C} \to \mathbb{C}$ is said to be *C-Lipschitz* if for all x and y in the domain of g we have $|g(x) - g(y)| \leq C|x - y|$. Similarly, given a map $K : L^2(X) \to L^2(X)$ (for some topological space X with a Borel probability measure), we define the Lipschitz constant

$$\|K\|_{\mathrm{Lip}} := \inf\{C \geq 0 : \text{for all } f, g \in L^2(X), \text{ we have } \|K(f) - K(g)\|_2 \leq C\|f - g\|_2\}. \tag{2.6}$$

When we need to specify the domain of K we write $\|K\|_{\mathrm{Lip}(L^2(X))}$. The following lemma identifies a useful property of the kind alluded to above.

Lemma 2.38. *Let $r : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ with $r(0) = 0$ be such that q_r is C-Lipschitz for some constant C. Then for any finite abelian group Z we have $\|K_r\|_{\mathrm{Lip}(\mathbb{C}^{\mathrm{Z}})} \leq C$.*

Proof. By Plancherel's theorem we have $\|K_r(f) - K_r(g)\|_2^2 = \sum_{\chi \in \widehat{G}} |q_r(\widehat{f}(\chi)) - q_r(\widehat{g}(\chi))|^2 \leq \sum_\chi C^2 |\widehat{f}(\chi) - \widehat{g}(\chi)|^2 = C^2\|f - g\|_2^2$. □

Now we focus on a specific function r and define the principal operator that we shall use.

Definition 2.39 (Fourier denoising operator). For $\varepsilon > 0$, we denote by K_ε the operator defined by applying (2.5) with the specific function $r = r_\varepsilon : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}, x \mapsto \max(x - \varepsilon, 0)$.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Lemma 2.40. *The function $q_\varepsilon := q_{r_\varepsilon}$ is a 1-Lipschitz function on $\mathbb{C}$, whence the operator K_ε is 1-Lipschitz on $\mathbb{C}^Z$.*

Proof. We prove the following inequality by a simple case distinction:

$$\forall x, y \in \mathbb{C},\ |q_\varepsilon(x) - q_\varepsilon(y)| \le |x - y|. \tag{2.7}$$

The first case is $\min(|x|, |y|) \ge \varepsilon$. In this case, let $u := x/|x|$ and $v = y/|y|$. Then $q_\varepsilon(x) = u(|x| - \varepsilon)$ and $q_\varepsilon(y) = u(|y| - \varepsilon)$. Note that $\mathrm{Re}(u\overline{v}) \le |u\overline{v}| \le 1$. Using that u, v have modulus 1, we have that $|u(|x| - \varepsilon) - v(|y| - \varepsilon)|^2$ equals

$$\begin{aligned}
&(|x| - \varepsilon)^2 + (|y| - \varepsilon)^2 - 2(|x| - \varepsilon)(|y| - \varepsilon)\,\mathrm{Re}(u\overline{v}) \\
&= |x|^2 + |y|^2 - 2\varepsilon(|x| + |y|) + 2\varepsilon^2 - 2\,\mathrm{Re}(x\overline{y}) + 2\varepsilon(|x| + |y|)\,\mathrm{Re}(u\overline{v}) - 2\varepsilon^2\,\mathrm{Re}(u\overline{v}) \\
&= |x - y|^2 - 2\varepsilon\,(|x| + |y| - \varepsilon)\,(1 - \mathrm{Re}(u\overline{v}))\ \le\ |x - y|^2.
\end{aligned}$$

The second case is $\min(|x|, |y|) < \varepsilon \le \max(|x|, |y|)$. In this case, suppose without loss of generality that $|x| < \varepsilon$, $|y| \ge \varepsilon$, and then observe that $q_\varepsilon(x) = 0$, so $|q_\varepsilon(x) - q_\varepsilon(y)| = |y| - \varepsilon$. The triangle inequality then yields $|x - y| \ge |y| - |x| \ge |y| - \varepsilon = |q_\varepsilon(x) - q_\varepsilon(y)|$, as required.

The third case is $\max(|x|, |y|) \le \varepsilon$. Here (2.7) holds trivially, since $q_\varepsilon(x) = q_\varepsilon(y) = 0$. □

Remark 2.41. Letting $q'_\varepsilon(z) := z - q_\varepsilon(z)$, a similar argument shows that q'_ε is also a contraction on $\mathbb{C}$, that is, we have $|x - y| \ge |q'_\varepsilon(x) - q'_\varepsilon(y)|$ for all $x, y \in \mathbb{C}$. This is less important for our applications, so we omit the proof. We do however record the following result, which will be used later in this paper.

Lemma 2.42. *We have $|q'_\varepsilon(z)| = \min(|z|, \varepsilon)$ for every $z \in \mathbb{C}$, and thus*

$$\forall x, y \in \mathbb{C},\ |q'_\varepsilon(x + y)| \le |q'_\varepsilon(x)| + |q'_\varepsilon(y)|. \tag{2.8}$$

Proof. We see that $|q'_\varepsilon(z)| = \min(|z|, \varepsilon)$ by a simple case distinction. Indeed for $|z| \le \varepsilon$ we have $|q'_\varepsilon(z)| = |z - q_\varepsilon(z)| = |z| = \min(|z|, \varepsilon)$, and for $|z| > \varepsilon$ we have $|q'_\varepsilon(z)| = |z - q_\varepsilon(z)| = |z(1 - \frac{|z| - \varepsilon}{|z|})| = \varepsilon = \min(|z|, \varepsilon)$.

To see that (2.8) holds, note first that $|q'_\varepsilon(x + y)| = \min(|x + y|, \varepsilon) \le \min(|x| + |y|, \varepsilon)$, so it suffices to prove that

$$\forall a, b \in \mathbb{R}_{\ge 0},\ \min(a + b, \varepsilon) \le \min(a, \varepsilon) + \min(b, \varepsilon). \tag{2.9}$$

This can also be proved by a case distinction.

The first case is $a + b \le \varepsilon$. Then $\min(a + b, \varepsilon) = a + b = \min(a, \varepsilon) + \min(b, \varepsilon)$ (where the last equality is since $a \le a + b \le \varepsilon$ so that $a = \min(a, \varepsilon)$, and similarly $b = \min(b, \varepsilon)$). So in this case we have equality.

The second case is $a + b > \varepsilon$. Then the left side of (2.9) is ε. There is then a first sub-case in which both a and b are less than ε, in which case the right side of (2.9) is $a + b$, which by assumption in this second case is $> \varepsilon$, so (2.9) holds in this sub-case. In the second sub-case at least one of a, b is at least ε, say wlog it is a. Then the right side of (2.9) is $\varepsilon + \min(b, \varepsilon) \ge \varepsilon$, so (2.9) holds. This completes this second case. □

We have thus proved that the operator K_ε has the desired continuity property mentioned at the beginning of this section.

It remains to verify that, for functions $f : Z \to \mathbb{C}$, the operator K_ε yields a decomposition of f into structured and quasirandom parts, in which $K_\varepsilon(f)$ is a valid Fourier-structured part. Note that if $\|f\|_2 \le 1$, then letting $K'_\varepsilon(f) = f - K_\varepsilon(f)$, we have the decomposition $f = K_\varepsilon(f) + K'_\varepsilon(f)$ where $\|K'_\varepsilon(f)\|_{U^2} \le \varepsilon^{1/2}$. We now prove that $K_\varepsilon(f)$ is indeed Fourier structured.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Lemma 2.43. *For every function $f : \mathrm{Z} \to \mathbb{C}$, the function $K_\varepsilon(f)$ is $(\varepsilon^{-2}\|f\|_2^2, 0)$-Fourier structured. In particular we have*

$$\|K_\varepsilon(f)\|_{A(\mathrm{Z})} \leq \varepsilon^{-1}\|f\|_2^2. \tag{2.10}$$

Proof. By Parseval's identity we have $|\{\chi \in \widehat{\mathrm{Z}} : |\widehat{f}(\chi)| \geq \varepsilon\}| \leq \|f\|_2^2/\varepsilon^2$, so $K_\varepsilon(f)$ is $(\varepsilon^{-2}\|f\|_2^2, 0)$-Fourier structured. To obtain (2.10), we use the Cauchy-Schwarz inequality and the fact that $\|K_\varepsilon(f)\|_2 \leq \|f\|_2$. □

Thus, if $\|f\|_2 \leq 1$ then $K_\varepsilon(f)$ is $(\varepsilon^{-2}, 0)$-Fourier structured. This completes the verification that $f = K_\varepsilon(f) + K'_\varepsilon(f)$ is a decomposition into Fourier-structured and quasirandom parts.

Remark 2.44. There are many possible choices of Fourier regularizing operators other than K_ε. In general, we want such an operator to leave the large Fourier coefficients roughly unchanged, to reduce the small ones to values close to 0, and to be invariant. For our proofs in this paper, it turned out that K_ε presented particularly convenient properties, such as Lemmas 2.40 and 2.42. Within the family of operators K_r defined in (2.5), some of these convenient properties of K_ε are available with other choices of r. For instance, it is not hard to show that if r is C-Lipschitz, then the operator K_r is $3C$-Lipschitz. One can then choose functions q_r that still reduce to 0 the small Fourier coefficients (like q_ε) while leaving the large Fourier coefficients unchanged (unlike q_ε). A specific example is $r(x) = x$ if $x \geq \varepsilon$, $r(x) = 2x - \varepsilon$ for $x \in [\varepsilon/2, \varepsilon]$, and $r(x) = 0$ otherwise. Looking into other possible choices of Fourier regularizing operators may be an interesting research direction, in particular regarding quantitative improvements of the resulting algorithms. In this direction, let us emphasize that the property of r being Lipschitz is quite convenient to avoid certain technical issues, as shown in the next example.

Example 2.45. Suppose that r is the blunt cut-off function $r(x) = x\,\mathbf{1}(x \geq \varepsilon)$ for a fixed $\varepsilon > 0$, and thus K_r is precisely the *Fourier cut-off operator* from Subsection 1.1. Let p be a prime number such that $\varepsilon > 1/\sqrt{p}$, and let $\delta > \varepsilon - 1/\sqrt{p}$. Then one can find a 1-bounded function $f : \mathbb{Z}_p \to \mathbb{C}$ and a δ-bounded function $g : \mathbb{Z}_p \to \mathbb{C}$ such that, while $\|(f+g) - f\|_2 = \delta$, we have $\|K_r(f+g) - K_r(f)\|_2 \geq 1$. Choosing carefully ε and p, we can make δ arbitrarily small, whence this choice of K_r fails to be a Lipschitz continuous operator in general, which is problematic for later calculations in this paper and also in specific applications. Let us give an explicit example of the latter. Let $f(x) := \exp(2\pi i x^2/p)$ and $g := \delta f$. Then, by Gauss-sum estimates, we know that all Fourier coefficients of f have absolute value $1/\sqrt{p}$. Thus $K_r(f) = 0$, whereas $K_r(f+g) = f + g$. Hence $\|K_r(f+g) - K_r(f)\|_2 = \|f+g\|_2 = 1 + \delta \geq 1$. This devolves into a noncontinuous behavior of $\mathcal{K}_r$ when we vary ε. For instance, let $h(x) = \exp(2\pi i x^3/p)$ for a large prime p. By similar arguments as above, it is easy to see that if we consider $\mathcal{K}_r(h \otimes \overline{h})$ for different ε, we get the following behavior: if $\varepsilon < 1/\sqrt{p}$ then $\mathcal{K}_r(h \otimes \overline{h}) = h \otimes \overline{h}$ and thus we see only one nonzero eigenvalue, namely 1, with corresponding eigenvector h. On the other hand, when $\varepsilon \geq 1/\sqrt{p}$, all diagonals $\mathcal{D}_{\mathcal{K}_r(h\otimes\overline{h}),t}$ for $t \in \mathrm{Z}\setminus\{0\}$ are 0, and our operator becomes the linear operator that simply multiplies by $1/p$. Hence, $\mathcal{K}_r$ exhibits a jump between two drastically different behaviors at $\varepsilon = 1/\sqrt{p}$.

Remark 2.46. In higher orders $k \geq 2$, an interesting choice of operator K to apply to the Z-diagonals is the dual-function operator $f \mapsto [f]_{U^k}$, where for $f : \mathrm{Z} \to \mathbb{C}$ we define the dual function $[f]_k$ (also denoted by $\mathcal{D}_k(f)$; see [66, Definition 11.13]) as follows:

$$[f]_k(x) := \mathbb{E}_{x,t_1,\dots,t_k \in \mathrm{Z}} \prod_{v \in \{0,1\}^k \setminus \{0^k\}} \mathcal{C}^{|v|+1} f(x + v(1)t_1 + \cdots + v(k)t_k). \tag{2.11}$$

We call the resulting Z-matrix $\mathcal{K}(f \otimes \overline{f})$ *the U^{k+1}-dual operator*, and denote it by $[[f]]_{U^{k+1}}$. It turns out that $[[f]]_{U^{k+1}}$ is a positive semi-definite operator, which has another definition very similar to the formula for the dual function (2.11) given by an average over cubes, but using an edge of the cube instead of a vertex. A useful property of dual functions is that $[f]_{U^k}$ is a structured function of order $k-1$ (in the sense of Definition 2.23) which satisfies $\langle f, [f]_{U^k}\rangle = \|f\|_{U^k}^{2^k}$, a correlation which indicates

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

that $[f]_{U^k}$ is nontrivially related to the structured part of f. Initial experiments with the eigenvectors of $[[f]]_{U^{k+1}}$ show promise in describing the quadratic structure of f, but further investigation is required.

To close this section, let us introduce some additional tools that will enable us, among other things, to detail more precisely the observation made in Remark 2.8, namely the convenient interaction of products of Z-matrices with the Fourier structure of their Z-diagonals. To do so, we introduce the following matrix norm which quantifies this Fourier structure.

Definition 2.47. For any matrix $M \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ we define

$$\|M\|_{MA} := \max_{t\in\mathrm{Z}} \|\mathcal{D}_{M,t}\|_{A(\mathrm{Z})}. \tag{2.12}$$

Lemma 2.48. *The function $\|\cdot\|_{MA}$ is a sub-multiplicative norm on the vector space $\mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ with respect to matrix multiplication.*

Proof. The fact that $\|\cdot\|_{MA}$ is a norm is easily checked. The sub-multiplicativity (namely that $\|XY\|_{MA} \le \|X\|_{MA}\|Y\|_{MA}$ for all $X, Y \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$) follows directly from Proposition 2.7. □

This matrix norm immediately enables the control of the norm of *polynomials* evaluated on such matrices, which will be very useful to isolate certain eigenvalues later on. Let us record this.

Definition 2.49. Given any polynomial $P(x) = \sum_{i=0}^n a_i x^i$ with coefficients $a_i \in \mathbb{C}$, let P^+ denote the polynomial defined by

$$P^+(x) := \sum_{i=0}^n |a_i| x^i. \tag{2.13}$$

Remark 2.50. Note that for any polynomials P, Q with coefficients in $\mathbb{C}$, for any $x \in \mathbb{R}_{\ge 0}$ we have $(PQ)^+(x) \le P^+(x)Q^+(x)$. In particular, for any integer $m \ge 1$ and $x \in \mathbb{R}_{\ge 0}$, we have $(P(x)^m)^+ \le P^+(x)^m$.

Lemma 2.48 has the following direct consequence.

Lemma 2.51. *For any matrix $M \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ and polynomial $P \in \mathbb{C}[x]$, we have $\|P(M)\|_{MA} \le P^+(\|M\|_{MA})$.*

Proof. Letting $P(x) = \sum_{i=0}^n a_i x^i$ we have $P(M) = \sum_{i=0}^n a_i M^i$, and by Lemma 2.48, we have

$$\|P(M)\|_{MA} \le \sum_{i=0}^n |a_i| \|M^i\|_{MA} \le \sum_{i=0}^n |a_i| \|M\|_{MA}^i = P^+(\|M\|_{MA}).$$ □

Let us record an important fact about weak quadratic characters, namely that these functions are precisely those which are approximately invariant under the operator $\mathcal{K}_\varepsilon$.

Proposition 2.52. *Let $f : \mathrm{Z} \to \mathbb{C}$ be a 1-bounded function. Then the following holds:*

(i) *For any $\gamma, \delta, \varepsilon \in (0,1]$, if $\|\mathcal{K}_\varepsilon(f\otimes\overline{f}) - f\otimes\overline{f}\|_2 \le \gamma$, then f is a $(\frac{4}{\varepsilon^2\delta^2}, (\frac{2\gamma}{\delta})^{\frac{1}{2}}, \delta)$-weak quadratic character.*

(ii) *For any $R > 0$ and $\delta_1, \delta_2 \in [0,1]$, if f is an (R, δ_1, δ_2)-weak quadratic character, then $\|\mathcal{K}_\varepsilon(f\otimes\overline{f}) - f\otimes\overline{f}\|_2 \le 6\varepsilon^{1/4}R^{1/2} + 4\delta_1 + 2\delta_2^{1/2}$.*

The proof will use the following basic property of K_ε.

Lemma 2.53. *Let $R > 0$, $\delta \in [0,1]$, and $g : \mathrm{Z} \to \mathbb{C}$ be (R,δ)-structured of order 1. Then $\|K_\varepsilon(g) - g\|_2 \le 2\delta + 4\varepsilon^{1/4}R^{1/2}(1+\|g\|_2)^{1/4}$.*

Proof. By Lemma 2.22, for any $\alpha > 0$ there is a function $h = \sum_{i\in[M]} c_i\chi_i$ with $M \le \alpha^{-6}R^4(1+\|g\|_2)^2$ and $\|g-h\|_2 \le \alpha+\delta$. Then, using Lemma 2.40, we have $\|K_\varepsilon(g)-g\|_2 \le \|K_\varepsilon(g)-K_\varepsilon(h)\|_2 + \|K_\varepsilon(h)-h\|_2 + \|h-g\|_2 \le \|K_\varepsilon(h)-h\|_2 + 2\|h-g\|_2 \le \|K_\varepsilon(h)-h\|_2 + 2(\alpha+\delta)$. By (2.5) and Lemma 2.42, we have $\|K_\varepsilon(h)-h\|_2 = (\sum_{i\in[M]} |q'_\varepsilon(c_i)|^2)^{1/2} \le \varepsilon M^{1/2}$. Thus $\|K_\varepsilon(g)-g\|_2 \le 2(\alpha+\delta) + \varepsilon\alpha^{-3}R^2(1+\|g\|_2)$, a quantity which we minimize by choosing $\alpha = (3/2)^{-1/4}\varepsilon^{1/4}R^{1/2}(1+\|g\|_2)^{1/4}$. The result follows. □

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Proof of Proposition 2.52. Since $\mathbb{E}_t\|K_\varepsilon(\Delta_t f)-\Delta_t f\|_2^2 \le \gamma^2$, it follows from Markov's inequality that $|\{t\in \mathrm{Z} : \|K_\varepsilon(\Delta_t f)-\Delta_t f\|_2 < \frac{\gamma}{(\delta/2)^{1/2}}\}| \ge |\mathrm{Z}|(1-\delta/2)$. By Lemma 2.43, we have that $K_\varepsilon(\Delta_t f)$ is $(\varepsilon^{-2}\|\Delta_t f\|_2^2, 0)$-Fourier structured. Note that $\mathbb{E}_t\|\Delta_t f\|_2^2 = \|f\|_2^4$, so $|\{t\in \mathrm{Z} : \|\Delta_t f\|_2^2 < \frac{\|f\|_2^4}{\delta/2}\}| \ge |\mathrm{Z}|(1-\delta/2)$. Hence, by the union bound, for at least $(1-\delta)|\mathrm{Z}|$ elements $t\in \mathrm{Z}$, we have that $\Delta_t f$ is $(\frac{2\|f\|_2^4}{\varepsilon^2\delta}, \frac{2^{1/2}\gamma}{\delta^{1/2}})$-Fourier structured, hence $(R', \frac{2^{1/2}\gamma}{\delta^{1/2}})$-structured of order 1 where, by Lemma 2.22 and what we know about t, we have $R' = \frac{2^{3/4}\|f\|_2^3}{\varepsilon^{3/2}\delta^{3/4}}\|\Delta_t f\|_2 \le \frac{2^{5/4}\|f\|_2^5}{\varepsilon^{3/2}\delta^{5/4}} \le \frac{4}{\varepsilon^2\delta^2}$. This proves (i).

To prove (ii), let $S\subset \mathrm{Z}$ with $|S|\ge(1-\delta_2)|\mathrm{Z}|$ and such that for every $t\in S$ we have that $\Delta_t f$ is (R,δ_1)-structured of order 1. Note that

$$\mathbb{E}_t\|K_\varepsilon(\Delta_t f)-\Delta_t f\|_2^2 = \mathbb{E}_t 1_S(t)\|K_\varepsilon(\Delta_t f)-\Delta_t f\|_2^2 + \mathbb{E}_t 1_{S^c}(t)\|K_\varepsilon(\Delta_t f)-\Delta_t f\|_2^2.$$

By the triangle inequality and Lemma 2.40, we have $\|K_\varepsilon(\Delta_t f)-\Delta_t f\|_2 \le 2\|\Delta_t f\|_2$, so $\mathbb{E}_t 1_{S^c}(t)\|K_\varepsilon(\Delta_t f)-\Delta_t f\|_2^2 \le 4\mathbb{E}_t 1_{S^c}(t)\|\Delta_t f\|_2^2 \le 4\delta_2$ (using that $\|\Delta_t f\|_2 \le \|f\|_\infty \le 1$). On the other hand, for every $t\in S$ we have $\Delta_t f = g_t + \mathcal{E}_t$ where $\|\mathcal{E}_t\|_2\le\delta_1$ (in particular $\|g_t\|_2\le 1+\delta_1$). Hence

$$\|K_\varepsilon(\Delta_t f)-\Delta_t f\|_2 = \|K_\varepsilon(g_t+\mathcal{E}_t)-K_\varepsilon(g_t)+K_\varepsilon(g_t)-g_t-\mathcal{E}_t\|_2.$$

By Lemma 2.40 we have $\|K_\varepsilon(g_t+\mathcal{E}_t)-K_\varepsilon(g_t)\|_2 \le \|\mathcal{E}_t\|_2 \le \delta_1$, and by Lemma 2.53 we have $\|K_\varepsilon(g_t)-g_t\|_2 \le 2\delta_1 + 4\varepsilon^{1/4}R^{1/2}(2+\delta_1)^{1/4} \le 2\delta_1 + 6\varepsilon^{1/4}R^{1/2}$. We conclude that for $t\in S$ we have $\|K_\varepsilon(\Delta_t f)-\Delta_t f\|_2 \le 4\delta_1 + 6\varepsilon^{1/4}R^{1/2}$, so $\mathbb{E}_t 1_S(t)\|K_\varepsilon(\Delta_t f)-\Delta_t f\|_2^2 \le (4\delta_1+6\varepsilon^{1/4}R^{1/2})^2$.

Hence $\|\mathcal{K}_\varepsilon(f\otimes\overline{f})-f\otimes\overline{f}\|_2 \le (4\delta_2 + [4\delta_1+6\varepsilon^{1/4}R^{1/2}]^2)^{1/2}$ and the result follows. □

Remark 2.54. Note that, to prove (i), we did not need the assumption $\|f\|_\infty\le 1$. Indeed, a bound $\|f\|_2\le 1$ was sufficient. Moreover, the proof of statement (ii) used the bound $\|f\|_\infty\le 1$ only to control the term involving S^c. Since for proper (i.e., nonweak) quadratic characters, there is no such term (as S^c is empty), it follows that for quadratic characters f, Proposition 2.52 holds in general (i.e., without assuming that f is 1-bounded). In particular, if f is an (R,δ)-quadratic character then $\|\mathcal{K}_\varepsilon(f\otimes\overline{f})-f\otimes\overline{f}\|_2 \le 6\varepsilon^{1/4}R^{1/2}+4\delta$.

Finally, we record the following fact, which is useful to control the behavior of weak quadratic characters under scalar multiplication.

Lemma 2.55. *For any function $f:\mathrm{Z}\to\mathbb{C}$ and $c\in\mathbb{C}$, we have*

$$\|\mathcal{K}_\varepsilon\big((cf)\otimes(\overline{cf})\big)-(cf)\otimes(\overline{cf})\|_2 = |c|^2\,\|\mathcal{K}_{\varepsilon/|c|^2}(f\otimes\overline{f})-f\otimes\overline{f}\|_2. \tag{2.14}$$

In particular,

$$\|\mathcal{K}_\varepsilon((cf)\otimes(\overline{cf}))-(cf)\otimes(\overline{cf})\|_2 \le \max\big(|c|^2\,\|\mathcal{K}_{\varepsilon^{1/2}}(f\otimes\overline{f})-f\otimes\overline{f}\|_2, 2\varepsilon^{1/2}\|f\|_2^2\big). \tag{2.15}$$

Proof. By (2.5) and Lemma 2.42, we have that $\|\mathcal{K}_\varepsilon((cf)\otimes(\overline{cf}))-(cf)\otimes(\overline{cf})\|_2^2$ equals

$$\begin{aligned}\mathbb{E}_t\|K_\varepsilon(|c|^2\Delta_t f)-|c|^2\Delta_t f\|_2^2 &= \mathbb{E}_t\sum_{\chi\in\widehat{\mathrm{Z}}}|\widehat{K_\varepsilon(|c|^2\Delta_t f)}(\chi)-|c|^2\widehat{\Delta_t f}(\chi)|^2\\ &= \mathbb{E}_t\sum_\chi\big|q'_\varepsilon\big(|c|^2\widehat{\Delta_t f}(\chi)\big)\big|^2 = \mathbb{E}_t\sum_\chi\min\big(|c|^2|\widehat{\Delta_t f}(\chi)|,\varepsilon\big)^2.\end{aligned}$$

Using that $\min(\lambda a,\lambda b)=\lambda\min(a,b)$ for any $a,b,\lambda\ge 0$, we see that the last average above equals $|c|^4\,\mathbb{E}_t\sum_\chi\min\big(|\widehat{\Delta_t f}(\chi)|,\varepsilon/|c|^2\big)^2 = |c|^4\,\|\mathcal{K}_{\varepsilon/|c|^2}(f\otimes\overline{f})-f\otimes\overline{f}\|_2^2$, and (2.14) follows.

To see (2.15), note first that the argument above yields the equalities

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

$$\forall\, c \in \mathbb{C}, \quad \|\mathcal{K}_\varepsilon((cf)\otimes(\overline{cf})) - (cf)\otimes(\overline{cf})\|_2^2 = \mathbb{E}_t \sum_\chi \min\left(|c|^2|\widehat{\Delta_t f}(\chi)|, \varepsilon\right)^2, \tag{2.16}$$

$$\forall\, c \in \mathbb{C}\setminus\{0\}, \quad \|\mathcal{K}_{\varepsilon/|c|^2}(f\otimes\overline{f}) - f\otimes\overline{f}\|_2^2 = \mathbb{E}_t \sum_\chi \min\left(|\widehat{\Delta_t f}(\chi)|, \varepsilon/|c|^2\right)^2, \tag{2.17}$$

and that (2.16) implies that $\|\mathcal{K}_\varepsilon((cf)\otimes(\overline{cf})) - (cf)\otimes(\overline{cf})\|_2$ increases as $|c|$ increases, while (2.17) implies that $\|\mathcal{K}_{\varepsilon/|c|^2}(f\otimes\overline{f}) - f\otimes\overline{f}\|_2$ decreases as $|c|$ increases.

Now $\|\mathcal{K}_{\varepsilon/|c|^2}(f\otimes\overline{f}) - f\otimes\overline{f}\|_2 \le 2\|f\otimes\overline{f}\|_2 = 2\|f\|_2^2$. Hence, if $|c|^2 \le \varepsilon^{1/2}$, then, by the increasing property, the left side of (2.15) is at most $2\varepsilon^{1/2}\|f\|_2^2$, while if $|c|^2 > \varepsilon^{1/2}$, then by the decreasing property, the left side of (2.15) is at most $|c|^2\,\|\mathcal{K}_{\varepsilon^{1/2}}(f\otimes\overline{f}) - f\otimes\overline{f}\|_2$, and (2.15) follows. □

3. Elementary motivation of the spectral approach in quadratic Fourier analysis

As discussed in the introduction, the spectral approach in quadratic Fourier analysis is based on the idea that, given a function $f : \mathrm{Z} \to \mathbb{C}$, the dominant eigenvectors of the Fourier-regularized matrix $\mathcal{K}_\varepsilon(f\otimes\overline{f})$ are interesting quadratic Fourier components of f. In this section, we present some initial nontrivial results illustrating this idea more precisely.

The main result is that certain eigenvectors of the matrix $\mathcal{K}_\varepsilon(f\otimes\overline{f})$, corresponding to large eigenvalues, are weak quadratic characters. We establish this formally in the first subsection below. In the second subsection, we gather further useful properties of the operator K_ε.

3.1. Spectrally isolated eigenvectors of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$ are (weak) quadratic characters

Definition 3.1. Let M be a self-adjoint linear operator on a finite-dimensional Hilbert space,[7] let λ be an eigenvalue of M and let $\theta > 0$. We say that λ is *θ-isolated* if the multiplicity of λ is 1 and every other eigenvalue λ' of M satisfies $|\lambda - \lambda'| > \theta$.

The main result in this subsection, Proposition 3.4, tells us that eigenvectors of the matrix $\mathcal{K}_\varepsilon(f\otimes\overline{f})$ corresponding to isolated large eigenvalues are weak quadratic characters, with parameters depending on how structured the diagonals of the matrix are. To establish this result, we will use the interaction of the matrix norm $\|\cdot\|_{MA}$ (recall Definition 2.47) with polynomial matrix expressions, described in Lemma 2.51. To apply this effectively on isolated eigenvalues, we shall use the following polynomial approximations of indicator functions.

Definition 3.2. For any $n \in \mathbb{N}$ and $\lambda \in \mathbb{R}\setminus\{0\}$, we define the polynomial

$$p_n(x,\lambda) := \lambda^{-1}x\big(1 - (x-\lambda)^2/4\big)^n.$$

Note that $p_n(\lambda,\lambda) = 1$. The idea is that if n is sufficiently large and x sufficiently separated from λ, then $p_n(x,\lambda)$ is small. We make this precise as follows.

Lemma 3.3. *For any $\theta > 0$ and $\lambda \in [-1,1]\setminus\{0\}$, if $n \ge 4\theta^{-2}\ln(\theta^{-1}\lambda^{-1})$, then $|p_n(x,\lambda)| \le \theta|x|$ for every $x \in [-1,1]$ satisfying $|x-\lambda| \ge \theta$.*

Proof. The inequality holds for $x = 0$, as $p_n(0,\lambda) = 0$, so we can assume that $x \ne 0$. Then

$$\begin{aligned}|p_n(x,\lambda)|/|x| &= |\lambda^{-1}\big(1-(x-\lambda)^2/4\big)^n| = \lambda^{-1}\big(1-(x-\lambda)^2/4\big)^n\\ &\le \lambda^{-1}\exp(-n(x-\lambda)^2/4) \le \lambda^{-1}\exp(-n\theta^2/4) \le \theta\end{aligned}$$

where the last inequality follows from the assumption on n. □

[7] For us, M will always be a Z-matrix and thus its eigenvalues are defined as per Definition 2.3.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Recall from (2.13) the definition of p_n^+. Since $p_n(x,\lambda) = \lambda^{-1}x(-\frac{x^2}{4} + \lambda\frac{x}{2} + (1-\frac{\lambda^2}{4}))^n$, it follows that for $0 < |\lambda| \le 1$ and $x \ge 0$ we have

$$p_n^+(x,\lambda) = |\lambda|^{-1}x\big(\tfrac{x^2}{4} + \tfrac{|\lambda|x}{2} + (1-\tfrac{\lambda^2}{4})\big)^n \le |\lambda|^{-1}x\big(\tfrac{x^2}{4} + x + 1\big)^n = |\lambda|^{-1}x\big(\tfrac{x}{2}+1\big)^{2n}. \tag{3.1}$$

Recall also that we normalize eigenvectors v of a matrix to have $\|v\|_2 = 1$. We can now present the main result of this subsection.

Proposition 3.4. *Let $M \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ be a self-adjoint* Z*-matrix with $\|M\|_2 \le 1$, and let $\theta \in (0,1]$. Suppose that $\lambda \ne 0$ is a θ-isolated eigenvalue of M, and let g be a corresponding eigenvector of M. Then, letting $n := \lceil 4\theta^{-2}\ln(\theta^{-1}\lambda^{-1})\rceil$, we have that for at least $(1-\theta^{1/2})|\mathrm{Z}|$ values of $t \in \mathrm{Z}$, the function $\Delta_t g$ is $(p_n^+(\|M\|_{MA},\lambda)^2\theta^{-1}, 2\theta^{1/2})$-Fourier structured, so g is a weak quadratic character with parameters $\big(2p_n^+(\|M\|_{MA},\lambda)^{3/2}\theta^{-5/4}, 3\theta^{1/2}, \theta\big)$. In particular, if $|\lambda| \ge \theta$, then g is a weak quadratic character with parameters*

$$\big(2\theta^{-3}\|M\|_{MA}(\|M\|_{MA}/2+1)^{2\lceil 8\theta^{-2}\ln(\theta^{-1})\rceil},\ 3\theta^{1/2},\ \theta\big). \tag{3.2}$$

Proof. As M is self-adjoint, let $\{g_i\}_{i=1}^{|\mathrm{Z}|-1}$ and $\{\lambda_i\}_{i=1}^{|\mathrm{Z}|-1}$ be eigenvectors and corresponding eigenvalues of M in a way that $\{g, g_1, \dots, g_{|\mathrm{Z}|-1}\}$ forms an orthonormal basis. Since $\|M\|_2 \le 1$, we have $\lambda^2 + \sum_{i=1}^{|\mathrm{Z}|-1}\lambda_i^2 \le 1$. In particular, all eigenvalues of M are of absolute value at most 1. Let $M' := p_n(M,\lambda)$. We have (using $p_n(\lambda,\lambda) = 1$) that $M' = g\otimes\overline{g} + \sum_{i=1}^{|\mathrm{Z}|-1} p_n(\lambda_i,\lambda)g_i\otimes\overline{g_i}$.

By Lemma 3.3, $|p_n(\lambda_i,\lambda)| \le \theta|\lambda_i|$ holds for $1 \le i \le |\mathrm{Z}|-1$, so $\sum_{i=1}^{|\mathrm{Z}|-1} p_n(\lambda_i,\lambda)^2 \le \theta^2\sum_{i=1}^{|\mathrm{Z}|-1}\lambda_i^2 \le \theta^2$. Therefore $\mathbb{E}_t\|\mathcal{D}_{M',t} - \Delta_t g\|_2^2 = \|g\otimes\overline{g} - M'\|_2^2 \le \theta^2$. This means that for a subset $S \in \mathrm{Z}$ with $|S| \ge (1-\theta/2)|\mathrm{Z}|$ we have that, if $t \in S$, then $\|\mathcal{D}_{M',t} - \Delta_t g\|_2 \le (2\theta)^{1/2}$.

As $\mathbb{E}_t\|\Delta_t g\|_2^2 = \|g\|_2^4 \le 1$, the set $S' := \{t \in \mathrm{Z} : \|\Delta_t g\|_2^2 \le \frac{2}{\theta}\}$ satisfies $|S'|/|\mathrm{Z}| \ge 1-\theta/2$.

By Lemma 2.51, $\|M'\|_{MA} \le p_n^+(\|M\|_{MA},\lambda)$. By Corollary 2.19, for every $t \in S\cap S'$, $\mathcal{D}_{M',t}$ is $(p_n^+(\|M\|_{MA},\lambda)^2\theta^{-1},\theta)$-Fourier structured, so $\Delta_t g$ is $(p_n^+(\|M\|_{MA},\lambda)^2\theta^{-1}, 3\theta^{1/2})$-Fourier structured and $\|\Delta_t g\|_2 \le (2/\theta)^{1/2}$. The parameters for g involving λ follow by Lemma 2.22, and the final parameters in (3.2) follow from the latter and (3.1). □

Proposition 3.4 is just one member of a family of similar results; we chose this one for purposes of illustration and motivation. In other variants one could want for instance the precision parameter of the character g to be independent of the isolation θ. Proving such variants can involve different choices of the polynomial p_n. Proposition 3.4 implies the next theorem.

Theorem 3.5. *Let* Z *be a finite abelian group, let $f : \mathrm{Z} \to \mathbb{C}$ be a 1-bounded function, and let $\theta, \varepsilon \in (0,1]$. Suppose that there exists a θ-isolated eigenvalue λ of $\mathcal{K}_\varepsilon\big(f\otimes\overline{f}\big)$ with $|\lambda| \ge \theta$, and let g be a corresponding eigenvector with $\|g\|_2 = 1$. Then g is a weak quadratic character with parameters $\big(2\theta^{-3}\varepsilon^{-1}(1/(2\varepsilon)+1)^{2n}, 3\theta^{1/2}, \theta\big)$, where $n = \lceil 8\theta^{-2}\ln(\theta^{-1})\rceil$.*

Proof. Let M be the self-adjoint matrix $\mathcal{K}_\varepsilon\big(f\otimes\overline{f}\big)$ and $M' := f\otimes\overline{f}$. Note that $\|\mathcal{D}_{M',t}\|_2 \le \|\mathcal{D}_{M',t}\|_\infty \le 1$ holds for every $t \in \mathrm{Z}$, so $\|\mathcal{D}_{M,t}\|_2 \le 1$ holds for every $t \in \mathrm{Z}$ by Lemma 2.40. By (2.10) we have $\|\mathcal{D}_{M,t}\|_{A(\mathrm{Z})} \le \varepsilon^{-1}$ for every $t \in \mathrm{Z}$. Note also that $\|M\|_2 = (\mathbb{E}_t\|\mathcal{D}_{M,t}\|_2^2)^{1/2} \le 1$. The result now follows by applying Proposition 3.4. □

There are various directions in which one can want to strengthen Theorem 3.5, and which provide good motivation for the upcoming sections of the paper.

One such direction seeks to ensure that the eigenvector g is not just a *weak* quadratic character, but rather a proper quadratic character. Another direction, which is important in connection with inverse theorems, is to add a nontrivial lower bound for the correlation between the original function f and the quadratic character g. In yet another direction, the goal is to obtain a correlating quadratic character g even when none of the largest eigenvalues is isolated.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

In the upcoming sections we will make progress in these directions. To this end, we prepare in the next subsection by obtaining further useful properties of the operator K_ε. Then, in Section 4, we will start using nilspace theory, and especially certain decompositions in terms of *nilspace characters*, which will be key ingredients for our results in the above directions. We will then be able to say more about the relations between weak quadratic characters and 2-step nilspace characters (in particular with Theorem 4.19; see also Question 4.21).

3.2. *Further properties of the operator* K_ε

Here, we gather several inequalities involving the operator K_ε from Definition 2.39.

Throughout this subsection we maintain the notation Z for a finite abelian group. Recall also the distinction between K_ε and $\mathcal{K}_\varepsilon$ from Remark 2.11.

The first two results capture continuity of $\mathcal{K}_\varepsilon$ relative to the L^2-norm.

Lemma 3.6. *Let $M, N \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ and let $\varepsilon \in (0,1]$. Then*

$$\|\mathcal{K}_\varepsilon(M) - \mathcal{K}_\varepsilon(N)\|_2 \le \|M - N\|_2. \tag{3.3}$$

Proof. Computing the L^2-norm via diagonals, note that the left side of (3.3) squared equals $\mathbb{E}_t \|K_\varepsilon(\mathcal{D}_{M,t}) - K_\varepsilon(\mathcal{D}_{N,t})\|_2^2$. For every $t \in \mathrm{Z}$, by Lemma 2.40, $\|K_\varepsilon(\mathcal{D}_{M,t}) - K_\varepsilon(\mathcal{D}_{N,t})\|_2^2 \le \|\mathcal{D}_{M,t} - \mathcal{D}_{N,t}\|_2^2 = \|\mathcal{D}_{M-N,t}\|_2^2$. Since $\mathbb{E}_t \|\mathcal{D}_{M-N,t}\|_2^2 = \|M - N\|_2^2$, the result follows. □

Corollary 3.7. *For any functions $f, g : \mathrm{Z} \to \mathbb{C}$ and $\varepsilon \in (0,1]$, we have*

$$\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon(g \otimes \overline{g})\|_2 \le \|f - g\|_2 \, (\|f\|_2 + \|g\|_2). \tag{3.4}$$

Proof. Using (3.3) we see that $\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon(g \otimes \overline{g})\|_2$ is at most

$$\begin{aligned}\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon(f \otimes \overline{g})\|_2 + \|\mathcal{K}_\varepsilon(f \otimes \overline{g}) - \mathcal{K}_\varepsilon(g \otimes \overline{g})\|_2 \\ \le \|f \otimes \overline{f} - f \otimes \overline{g}\|_2 + \|f \otimes \overline{g} - g \otimes \overline{g}\|_2 = \|f \otimes (\overline{f} - \overline{g})\|_2 + \|(f - g) \otimes \overline{g}\|_2,\end{aligned}$$

and this equals $\|f - g\|_2 \, (\|f\|_2 + \|g\|_2)$ as required. □

The next results build up towards a continuity property of the map $f \mapsto \mathcal{K}_\varepsilon(f \otimes \overline{f})$ relative to the U^3-norm. This property will eventually be established in Corollary 3.13 below. We begin with the following simple lemma.

Lemma 3.8. *For any function $f : \mathrm{Z} \to \mathbb{C}$ and any $\varepsilon \in (0,1]$, we have $\|K_\varepsilon(f)\|_2 \le \varepsilon^{-1}\|f\|_{U^2}^2$.*

Proof. Let $g = K_\varepsilon(f)$ and let $S := \mathrm{Supp}(\widehat{g}) = \{\chi \in \widehat{\mathrm{Z}} : |\widehat{f}(\chi)| > \varepsilon\}$. By Markov's inequality we then have $|S|\varepsilon^4 \le \|\widehat{f}\|_{\ell^4}^4 = \|f\|_{U^2}^4$. Furthermore, we have $|\widehat{f}| \ge |\widehat{g}|$. By Parseval's identity and the Cauchy-Schwarz inequality, we then have

$$\|g\|_2^2 = \textstyle\sum_{\chi\in S} |\widehat{g}(\chi)|^2 \le \sum_\chi |\widehat{f}(\chi)|^2 1_S(\chi) \le |S|^{1/2}\||\widehat{f}|^2\|_{\ell^2} = |S|^{1/2}\|f\|_{U^2}^2 \le \varepsilon^{-2}\|f\|_{U^2}^4.$$ □

A similar argument yields the following result.

Lemma 3.9. *For any functions $f, g : \mathrm{Z} \to \mathbb{C}$ and $\varepsilon \in (0,1]$, we have*

$$\|K_\varepsilon(f) - K_\varepsilon(g)\|_2 \le \varepsilon^{-1/2} \, \|f - g\|_{U^2} \, (\|f\|_2^2 + \|g\|_2^2)^{1/4}. \tag{3.5}$$

Proof. We have

$$\|K_\varepsilon(f) - K_\varepsilon(g)\|_2^2 = \textstyle\sum_{\chi\in\widehat{\mathrm{Z}}} |\widehat{K_\varepsilon(f)}(\chi) - \widehat{K_\varepsilon(g)}(\chi)|^2 = \sum_{\chi\in\widehat{\mathrm{Z}}} |q_\varepsilon(\widehat{f}(\chi)) - q_\varepsilon(\widehat{g}(\chi))|^2.$$

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Since q_ε eliminates Fourier coefficients of modulus at most ε, the last sum above gets nonzero contributions only from characters χ in the set $S := \{\chi : |\widehat{f}(\chi)| \geq \varepsilon\} \cup \{\chi : |\widehat{g}(\chi)| \geq \varepsilon\}$. We have $\varepsilon^2 |\{\chi : |\widehat{f}(\chi)| \geq \varepsilon\}| \leq \sum_\chi |\widehat{f}(\chi)|^2 = \|f\|_2^2$ and similarly $|\{\chi : |\widehat{g}(\chi)| \geq \varepsilon\}| \leq \varepsilon^{-2}\|g\|_2^2$. Hence $|S| \leq \varepsilon^{-2}(\|f\|_2^2 + \|g\|_2^2)$.

By Lemma 2.40 the function q_ε is a contraction, whence (using Cauchy-Schwarz) we have

$$\sum_{\chi\in S} |q_\varepsilon(\widehat{f}(\chi)) - q_\varepsilon(\widehat{g}(\chi))|^2 \leq \sum_{\chi\in S} |\widehat{f}(\chi) - \widehat{g}(\chi)|^2 = \sum_{\chi\in S} |\widehat{(f-g)}(\chi)|^2 \leq |S|^{1/2}\|\widehat{f-g}\|_{\ell^4}^2.$$

Since $\|\widehat{f-g}\|_{\ell^4}^2 = \|f-g\|_{U^2}^2$, the result follows. □

Applying Lemma 3.9 to diagonal functions of Z-matrices (recall (2.1)), we shall obtain below a result (Lemma 3.12) enabling us to control L^2-norms of the form $\|\mathcal{K}_\varepsilon(M) - \mathcal{K}_\varepsilon(N)\|_2$, for Z-matrices M and N, in terms of a norm which is the case $k = 2$ of the following notion.

Definition 3.10. For any integer $k \geq 2$, we define the following norm on $\mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$:

$$\forall\, M \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}, \quad \|M\|_{DU^k} := \left(\mathbb{E}_{t\in\mathrm{Z}}\|\mathfrak{D}_{M,t}\|_{U^k}^{2^k}\right)^{1/2^k}. \tag{3.6}$$

Lemma 3.11. *For every $k \geq 2$ we have that $\|\cdot\|_{DU^k}$ is a norm on the vector space $\mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$.*

Proof. The only nontrivial norm axiom is the sub-additivity. This follows from the sub-additivity of the U^k-norm and the L^{2^k}-norm, indeed

$$\begin{aligned}\|M+N\|_{DU^k} &= \left(\mathbb{E}_t\|\mathfrak{D}_{M,t} + \mathfrak{D}_{N,t}\|_{U^k}^{2^k}\right)^{1/2^k} \leq \left(\mathbb{E}_t(\|\mathfrak{D}_{M,t}\|_{U^k} + \|\mathfrak{D}_{N,t}\|_{U^k})^{2^k}\right)^{1/2^k}\\ &\leq \left(\mathbb{E}_t\|\mathfrak{D}_{M,t}\|_{U^k}^{2^k}\right)^{1/2^k} + \left(\mathbb{E}_t\|\mathfrak{D}_{N,t}\|_{U^k}^{2^k}\right)^{1/2^k} = \|M\|_{DU^k} + \|N\|_{DU^k}.\end{aligned}$$ □

We can now obtain the lemma announced above.

Lemma 3.12. *Let $M, N \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ and $\varepsilon \in (0,1]$. Then*

$$\|\mathcal{K}_\varepsilon(M) - \mathcal{K}_\varepsilon(N)\|_2 \leq \varepsilon^{-1/2}\,\|M-N\|_{DU^2}\,(\|M\|_2^2 + \|N\|_2^2)^{1/4}. \tag{3.7}$$

Proof. The left side of (3.7) squared equals $\mathbb{E}_t\|K_\varepsilon(\mathfrak{D}_{M,t}) - K_\varepsilon(\mathfrak{D}_{N,t})\|_2^2$. Applying Lemma 3.9 for every $t \in \mathrm{Z}$, this is seen to be at most $\varepsilon^{-1}\mathbb{E}_t\|\mathfrak{D}_{M,t} - \mathfrak{D}_{N,t}\|_{U^2}^2\,(\|\mathfrak{D}_{M,t})\|_2^2 + \|\mathfrak{D}_{N,t}\|_2^2)^{1/2}$. Using that $\mathfrak{D}_{M,t} - \mathfrak{D}_{N,t} = \mathfrak{D}_{M-N,t}$ and Cauchy-Schwarz, the last quantity is seen to be at most $\varepsilon^{-1}(\mathbb{E}_t\|\mathfrak{D}_{M-N,t}\|_{U^2}^4)^{1/2}\,(\mathbb{E}_t\|\mathfrak{D}_{M,t})\|_2^2 + \mathbb{E}_t\|\mathfrak{D}_{N,t}\|_2^2)^{1/2}$ and (3.7) follows. □

Similarly to how we deduced Corollary 3.7 from Lemma 3.6, we shall deduce the following from Lemma 3.12.

Corollary 3.13. *Let $\varepsilon \in (0,1]$ and let $f, g : \mathrm{Z} \to \mathbb{C}$ be functions. Then*

$$\left\|\mathcal{K}_\varepsilon(f\otimes\overline{f}) - \mathcal{K}_\varepsilon(g\otimes\overline{g})\right\|_2 \leq \varepsilon^{-1/2}\,\|f-g\|_{U^3}\,(\|f\|_{U^3}^{4/3} + \|g\|_{U^3}^{4/3})^{3/4}(\|f\|_2^2 + \|g\|_2^2)^{1/2}. \tag{3.8}$$

To prove Corollary 3.13, we introduce the following special case of a Gowers U^k-product.

Definition 3.14. For any integer $k \geq 2$ and any functions $f, g : \mathrm{Z} \to \mathbb{C}$, we denote by $\langle f, g\rangle_{U^k}$ the special case of the Gowers U^k-product[8] $\langle (f_\nu)_{\nu\in[\![k]\!]}\rangle_{U^k}$ in which $f_\nu = f$ for $\nu(k) = 0$ and $f_\nu = g$ for $\nu(k) = 1$. Note that

$$\langle f, g\rangle_{U^k} = \mathbb{E}_{t\in\mathrm{Z}}\|fT^t\overline{g}\|_{U^{k-1}}^{2^{k-1}}. \tag{3.9}$$

[8]See Definition 2.31.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

In particular, we always have $\langle f, g\rangle_{U^k} \geq 0$. Note also that from (3.9) it follows that

$$\langle f, g\rangle_{U^k} = \|f \otimes \overline{g}\|_{DU^{k-1}}^{2^{k-1}}. \tag{3.10}$$

Proof of Corollary 3.13. We have

$$\left\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon(g \otimes \overline{g})\right\|_2 \leq \left\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon(f \otimes \overline{g})\right\|_2 + \left\|\mathcal{K}_\varepsilon(f \otimes \overline{g}) - \mathcal{K}_\varepsilon(g \otimes \overline{g})\right\|_2.$$

To bound the first term on the right hand side, note that, by Lemma 3.12 and (3.10), we have

$$\begin{aligned}\left\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon(f \otimes \overline{g})\right\|_2^2 &\leq \varepsilon^{-1}\, \|f \otimes (\overline{f} - \overline{g})\|_{DU^2}^2\, (\|f \otimes \overline{f}\|_2^2 + \|f \otimes \overline{g}\|_2^2)^{1/2} \\ &= \varepsilon^{-1}\, \langle f, f - g\rangle_{U^3}^{1/2}\, \|f\|_2\, (\|f\|_2^2 + \|g\|_2^2)^{1/2}.\end{aligned}$$

Similarly, $\left\|\mathcal{K}_\varepsilon(f \otimes \overline{g}) - \mathcal{K}_\varepsilon(g \otimes \overline{g})\right\|_2^2 \leq \varepsilon^{-1}\, \langle f - g, g\rangle_{U^3}^{1/2}\, \|g\|_2\, (\|f\|_2^2 + \|g\|_2^2)^{1/2}$. In particular, the product $\left\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon(f \otimes \overline{g})\right\|_2 \left\|\mathcal{K}_\varepsilon(f \otimes \overline{g}) - \mathcal{K}_\varepsilon(g \otimes \overline{g})\right\|_2$ is at most

$$\varepsilon^{-1}\, \langle f, f - g\rangle_{U^3}^{1/4}\, \langle f - g, g\rangle_{U^3}^{1/4}\, \|f\|_2^{1/2}\, \|g\|_2^{1/2}\, (\|f\|_2^2 + \|g\|_2^2)^{1/2}.$$

Hence $\left\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon(g \otimes \overline{g})\right\|_2^2 \leq \varepsilon^{-1}\, (\|f\|_2^2 + \|g\|_2^2)^{\frac{1}{2}} (\langle f, f - g\rangle_{U^3}^{\frac{1}{4}} \|f\|_2^{\frac{1}{2}} + \langle f - g, g\rangle_{U^3}^{\frac{1}{4}} \|g\|_2^{\frac{1}{2}})^2$. By the Gowers-Cauchy-Schwarz inequality, we have $\langle f, f - g\rangle_{U^3}^{1/4} \leq \|f\|_{U^3} \|f - g\|_{U^3}$ and $\langle f - g, g\rangle_{U^3}^{1/4} \leq \|g\|_{U^3} \|f - g\|_{U^3}$. Hence

$$\left\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon(g \otimes \overline{g})\right\|_2^2 \leq \varepsilon^{-1}\, \|f - g\|_{U^3}^2\, (\|f\|_2^2 + \|g\|_2^2)^{\frac{1}{2}}\, (\|f\|_{U^3} \|f\|_2^{\frac{1}{2}} + \|g\|_{U^3} \|g\|_2^{\frac{1}{2}})^2.$$

By Hölder's inequality with indices 4 and 4/3 applied to the last factor, the last line is at most $\varepsilon^{-1}\, \|f - g\|_{U^3}^2\, (\|f\|_2^2 + \|g\|_2^2)^{\frac{1}{2}}\, (\|f\|_{U^3}^{\frac{4}{3}} + \|g\|_{U^3}^{\frac{4}{3}})^{\frac{3}{2}}\, (\|f\|_2^2 + \|g\|_2^2)^{\frac{1}{2}}$. The result follows. □

The third and final type of bounds concerns functions and Z-matrices that are stable under the Fourier denoising operator, particularly functions f such that $\|K_\varepsilon(f) - f\|_2$ is small. We now prove various versions of the fact that the sum of stable functions (or Z-matrices) is also stable.

Lemma 3.15. *For any two functions $f, g : \mathrm{Z} \to \mathbb{C}$ and any $\varepsilon > 0$ we have*

$$\|K_\varepsilon(f + g) - (f + g)\|_2 \leq \|K_\varepsilon(f) - f\|_2 + \|K_\varepsilon(g) - g\|_2. \tag{3.11}$$

Proof. By Parseval and the definitions of K_ε, q_ε, and q'_ε (recall Definition 2.39), we have

$$\begin{aligned}\|K_\varepsilon(f + g) - (f + g)\|_2 &= \big(\textstyle\sum_{\chi \in \widehat{\mathrm{Z}}} |\widehat{K_\varepsilon(f + g)}(\chi) - (\widehat{f}(\chi) + \widehat{g}(\chi))|^2\big)^{1/2} \\ &= \big(\textstyle\sum_{\chi \in \widehat{\mathrm{Z}}} |q_\varepsilon(\widehat{f}(\chi) + \widehat{g}(\chi)) - (\widehat{f}(\chi) + \widehat{g}(\chi))|^2\big)^{1/2} = \big(\textstyle\sum_{\chi \in \widehat{\mathrm{Z}}} \left|q'_\varepsilon(\widehat{f}(\chi) + \widehat{g}(\chi))\right|^2\big)^{1/2}.\end{aligned}$$

By Lemma 2.42, this is at most $\big(\sum_{\chi \in \widehat{\mathrm{Z}}} \left(\left|q'_\varepsilon(\widehat{f}(\chi))\right| + \left|q'_\varepsilon(\widehat{g}(\chi))\right|\right)^2\big)^{\frac{1}{2}} \leq \big(\sum_{\chi \in \widehat{\mathrm{Z}}} \left|q'_\varepsilon(\widehat{f}(\chi))\right|^2\big)^{\frac{1}{2}} + \big(\sum_{\chi \in \widehat{\mathrm{Z}}} \left|q'_\varepsilon(\widehat{g}(\chi))\right|^2\big)^{1/2} = \|K_\varepsilon(f) - f\|_2 + \|K_\varepsilon(g) - g\|_2$. □

Corollary 3.16. *For any Z-matrices $M_1, M_2, \ldots, M_n$, and any $\varepsilon \in (0, 1]$, we have*

$$\left\|\mathcal{K}_\varepsilon\Big(\textstyle\sum_{i \in [n]} M_i\Big) - \sum_{i \in [n]} M_i\right\|_2 \leq \sum_{i \in [n]} \|\mathcal{K}_\varepsilon(M_i) - M_i\|_2. \tag{3.12}$$

Proof. First let us note that for any two such Z-matrices M and N we have

$$\|\mathcal{K}_\varepsilon(M + N) - M - N\|_2 \leq \|\mathcal{K}_\varepsilon(M) - M\|_2 + \|\mathcal{K}_\varepsilon(N) - N\|_2. \tag{3.13}$$

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Indeed, this follows from Lemma 3.15 applied to the diagonals (calculating the L^2-norms via diagonals as in previous proofs above). Now (3.12) follows from (3.13) by a simple induction. □

We now combine the above three types of bounds to give estimates that will be useful to handle various error terms that occur in applications of the regularity lemma proved in the next section.

Proposition 3.17. *Let $\varepsilon \in (0,1]$, $n \in \mathbb{N}$, and $\alpha_i \geq 0$ for $i \in [5]$. Let $f : \mathrm{Z} \to \mathbb{C}$ be a function satisfying $\|f\|_2 \leq 1$, and suppose that we have a decomposition $f = \sum_{i=1}^n f_i + f_e + f_r$ where*

(i) $\|f_e\|_2 \leq \alpha_1 \leq 1$,
(ii) $\|f_r\|_2 \leq 1$ *and* $\|f_r\|_{U^3} \leq \alpha_2$,
(iii) $\|f_i\|_2 \leq \alpha_3$ *for every* $i \in [n]$,
(iv) $\langle f_i, f_j\rangle_{U^3} \leq \alpha_4$ *for all* $i \neq j$ *in* $[n]$,
(v) $\|\mathcal{K}_\varepsilon(f_i \otimes \overline{f_i}) - f_i \otimes \overline{f}_i\|_2 \leq \alpha_5$ *for every* $i \in [n]$.

Then

$$\big\| \textstyle\sum_{i=1}^n f_i \otimes \overline{f}_i - \mathcal{K}_\varepsilon(f \otimes \overline{f})\big\|_2 \leq 5\alpha_1 + 16\varepsilon^{-1/2}\alpha_2 + \varepsilon^{-1/2}n^2(n^{1/2}\alpha_3 + 3)\alpha_4 + n\alpha_5. \tag{3.14}$$

Proof. Let $F = \sum_{i=1}^n f_i$. We begin by eliminating f_e using the first kind of continuity (relative to L^2 errors). We have $f \otimes \overline{f} = (F + f_r) \otimes \overline{(F + f_r)} + (F + f_r) \otimes \overline{f_e} + f_e \otimes \overline{(F + f_r)} + f_e \otimes \overline{f_e}$.

By Lemma 3.6, we have that $\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon((F + f_r) \otimes \overline{(F + f_r)})\|_2$ is at most

$$\begin{aligned}&\|(F + f_r) \otimes \overline{f_e} + f_e \otimes \overline{(F + f_r)} + f_e \otimes \overline{f_e}\|_2\\ &\leq\ 2\|F + f_r\|_2\,\|f_e\|_2 + \|f_e\|_2^2 \ \leq\ 2(1 + \alpha_1)\alpha_1 + \alpha_1^2 \ \leq\ 2\alpha_1 + 3\alpha_1^2 \ \leq\ 5\alpha_1.\end{aligned}$$

Next, we eliminate f_r using the second kind of continuity (relative to U^3 errors). By the triangle inequality, we have that $\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon(F \otimes \overline{F})\|_2$ is at most

$$\begin{aligned}&\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon((F + f_r) \otimes \overline{(F + f_r)})\|_2 + \|\mathcal{K}_\varepsilon((F + f_r) \otimes \overline{(F + f_r)}) - \mathcal{K}_\varepsilon(F \otimes \overline{F})\|_2\\ &\qquad\leq 5\alpha_1 + \|\mathcal{K}_\varepsilon((F + f_r) \otimes \overline{(F + f_r)}) - \mathcal{K}_\varepsilon(F \otimes \overline{F})\|_2.\end{aligned}$$

By Corollary 3.13 with $f = F + f_r$ and $g = F$, and using $\|F + f_r\|_2 \leq 1 + \alpha_1$, $\|F\|_2 \leq 2 + \alpha_1$, and $\|\cdot\|_{U^3} \leq \|\cdot\|_{L^2}$, we see that $\|\mathcal{K}_\varepsilon((F + f_r) \otimes \overline{(F + f_r)}) - \mathcal{K}_\varepsilon(F \otimes \overline{F})\|_2$ is at most

$$\begin{aligned}&\varepsilon^{-1/2}\|f_r\|_{U^3}(\|F + f_r\|_{U^3}^{4/3} + \|F\|_{U^3}^{4/3})^{3/4}\ (\|F + f_r\|_2^2 + \|F\|_2^2)^{1/2}\\ &\leq \varepsilon^{-1/2}\alpha_2((1 + \alpha_1)^{4/3} + (2 + \alpha_1)^{4/3})^{3/4}\ ((1 + \alpha_1)^2 + (2 + \alpha_1)^2)^{1/2}\\ &\leq \varepsilon^{-1/2}\alpha_2(2^{4/3} + 3^{4/3})^{3/4}(4 + 9)^{1/2} \leq 16\varepsilon^{-1/2}\alpha_2.\end{aligned}$$

Thus, so far we have proved that

$$\|\mathcal{K}_\varepsilon(f \otimes \overline{f}) - \mathcal{K}_\varepsilon(F \otimes \overline{F})\|_2 \leq 5\alpha_1 + 16\varepsilon^{-1/2}\alpha_2. \tag{3.15}$$

Next, from the expansion $F \otimes \overline{F} = \sum_{i\in[n]} f_i \otimes \overline{f_i} + \sum_{i\neq j\in[n]} f_i \otimes \overline{f_j}$, we eliminate the latter cross-terms. By Lemma 3.12, we have that $\|\mathcal{K}_\varepsilon(F \otimes \overline{F}) - \mathcal{K}_\varepsilon(\sum_{i\in[n]} f_i \otimes \overline{f_i})\|_2$ is at most

$$\begin{aligned}&\varepsilon^{-1/2}\big\|\textstyle\sum_{i\neq j\in[n]} f_i \otimes \overline{f_j}\big\|_{DU^2}\big(\|F \otimes \overline{F}\|_2^2 + \big\|\sum_{i\in[n]} f_i \otimes \overline{f_i}\big\|_2^2\big)^{1/4}\\ &\qquad\leq \varepsilon^{-1/2}\textstyle\sum_{i\neq j\in[n]}\langle f_i, f_j\rangle_{U^3}\ \big((2 + \alpha_1)^4 + n^2\alpha_3^4\big)^{1/4}\\ &\qquad\leq \tfrac{\alpha_4}{\varepsilon^{1/2}}n^2(2 + \alpha_1 + n^{1/2}\alpha_3) \leq \tfrac{\alpha_4}{\varepsilon^{1/2}}n^2(3 + n^{1/2}\alpha_3).\end{aligned}$$

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Combining this with (3.15) and the triangle inequality, we deduce that

$$\Big\|\mathcal{K}_\varepsilon(f\otimes\overline{f}) - \mathcal{K}_\varepsilon\Big(\sum_{i\in[n]} f_i\otimes\overline{f_i}\Big)\Big\|_2 \le 5\alpha_1 + 16\varepsilon^{-1/2}\alpha_2 + \varepsilon^{-1/2}n^2(n^{1/2}\alpha_3+3)\alpha_4. \tag{3.16}$$

Now, by Corollary 3.16 and property (iv), we have $\|\sum_{i\in[n]} f_i\otimes\overline{f_i} - \mathcal{K}_\varepsilon(\sum_{i\in[n]} f_i\otimes\overline{f_i})\|_2 \le n\alpha_5$. Then, combining this with (3.16) via the triangle inequality, we deduce (3.14). □

4. Nilspace characters

In this section, we begin to use nilspace theory to progress towards a central goal of this paper, namely proving the validity of our spectral algorithms in quadratic Fourier analysis (Theorem 1.1). To do so, in this section, we introduce and study what will be the main nilspace-theoretic notion of a higher-order Fourier analytic component of a function: the notion of a *nilspace character*. In particular, we shall establish in Theorem 4.19 a fact that will be important for us, namely that a 2-step nilspace character of bounded complexity on a finite abelian group Z is a *character of order 2* on Z (as per Definition 2.29), with a bounded complexity parameter depending on the precision parameter but not on the cardinality of Z. To explain this in more detail, we need to recall certain basic notions from nilspace theory.

In addition to the original paper [5], there are various subsequent sources treating nilspace theory in depth, including [6, 7, 37, 38, 39]. Therefore, rather than stating the full definitions here, let us mention the basic concepts of this theory and give references for their definitions. The basic concepts that we shall use include that of a *nilspace* [6, Definition 1.2.1], a *compact nilspace* [7, Definition 1.0.2] and *compact finite-rank (CFR) nilspaces* [7, Definition 1.0.2], as well as the *morphisms of nilspaces* [6, Definition 2.2.11]. We shall also use the *structure groups* and *bundle structure* of nilspaces [6, Definitions 3.2.17 and 3.2.18, and Theorem 3.2.19], as well as the existence of a unique probability Haar measure on each compact nilspace [7, Proposition 2.2.5]. It is worth recalling the notion of *nilspace polynomial*, see [12, Definition 1.3].

Definition 4.1. A k-step *nilspace polynomial* on a finite abelian group Z is a function of the form $F\circ\phi : \mathrm{Z}\to\mathbb{C}$ constructed as follows: there is a k-step CFR nilspace X such that the map ϕ is a nilspace morphism from Z (viewed as a 1-step nilspace) into X and $F : \mathrm{X}\to\mathbb{C}$ is a 1-bounded continuous function. Usually, additional parameters are specified, such as a parameter bounding the complexity of the nilspace X as well as the Lipschitz norm of F (see [12]), and in fact an adequate bound on some of these parameters is required for this notion of nilspace polynomial to be nontrivial; see Remark 4.6.

Nilspace characters are a special type of nilspace polynomials. To define nilspace characters, we first need to extend, to the setting of compact nilspaces, the notion (well-known in the setting of nilmanifolds) of a function having a *vertical frequency* (see [35, Definition 6.1] and more recently [51, Definition 2.15]).

Definition 4.2 (Functions with vertical frequency on a compact nilspace). Let X be a k-step compact nilspace, let $\mathrm{Z}_k(\mathrm{X})$ be the k-th structure group of X (a compact abelian group), and let χ be a character in $\widehat{\mathrm{Z}_k(\mathrm{X})}$. We say that a function $F : \mathrm{X}\to\mathbb{C}$ has *vertical frequency* χ if for every $x\in\mathrm{X}$ and every $z\in\mathrm{Z}_k(\mathrm{X})$ we have $F(x+z) = \chi(z)F(x)$. Given a bounded Borel function $F : \mathrm{X}\to\mathbb{C}$ and $\chi\in\widehat{\mathrm{Z}_k(\mathrm{X})}$, we denote by F_χ the *component of F with vertical frequency χ*, defined as follows (where μ is the probability Haar measure on $\mathrm{Z}_k(\mathrm{X})$):

$$F_\chi(x) = \int_{\mathrm{Z}_k(\mathrm{X})} F(x+z)\overline{\chi(z)}\,\mathrm{d}\mu(z). \tag{4.1}$$

It is easy to see that functions with vertical frequency χ form a closed subspace of the Hilbert space $L^2(\mathrm{X})$. We thus obtain the following notion (see [11, Definition 3.50]).

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Definition 4.3. Let X be a k-step compact nilspace, and let χ be a character in $\widehat{\mathrm{Z}_k(\mathrm{X})}$. We denote by $W(\chi, \mathrm{X})$ the Hilbert space of functions in $L^2(\mathrm{X})$ that have vertical frequency χ.

Remark 4.4. Let X be a k-step CFR nilspace, let $\chi \in \widehat{\mathrm{Z}_k(\mathrm{X})}$, and let $p \in [1, \infty]$. Then equation (4.1) defines a continuous linear operator $(\cdot)_\chi : L^p(\mathrm{X}) \to L^p(\mathrm{X})$ of norm 1. The space $W^p(\chi, \mathrm{X}) := \{f \in L^p(\mathrm{X}) : f \text{ has vertical frequency } \chi\}$ is a closed subspace of $L^p(\mathrm{X})$. Moreover, for $p = 2$, we have the decomposition $L^2(\mathrm{X}) = \bigoplus_{\chi \in \widehat{\mathrm{Z}_k(\mathrm{X})}} W^2(\chi, \mathrm{X})$ where the subspaces $W^2(\chi, \mathrm{X})$ are pairwise orthogonal. We will not need these results in this paper, as we will use mainly the components F_χ of much more specific functions F (Lipschitz continuous functions), so we omit the proofs.

We can now formulate the main concept in this section.

Definition 4.5 (Nilspace character). A k-step *nilspace character* on a finite abelian group Z is a 1-bounded function g of the following form: there is a k-step CFR nilspace X, a nilspace morphism $\phi : \mathcal{D}_1(\mathrm{Z}) \to \mathrm{X}$, and a 1-bounded continuous function $F : \mathrm{X} \to \mathbb{C}$ having vertical frequency $\chi \in \widehat{\mathrm{Z}_k(\mathrm{X})}$, such that $g = F \circ \phi$.

Various additional properties can be required for such a nilspace character $g = F \circ \phi$ (as for nilspace polynomials more generally): we can require F to have bounded Lipschitz constant C (we then say that g is a *C-Lipschitz nilspace character*); we can require the complexity[9] of the underlying nilspace to be at most m (we then say that g is a *complexity-m nilspace character*); we can also add an equidistribution requirement, namely that the morphism ϕ should be b-balanced[10] (we then say that g is a *b-balanced nilspace character*). Furthermore, one can require that the modulus of the character be equal to 1 everywhere. Note that with the latter requirement, when X is a nilmanifold, we recover the notion of a *nilcharacter* introduced by Green, Tao, and Ziegler in [35, Definition 6.1]. Note also that, on the other hand, the modulus-1 requirement can often clash with the continuity requirement for F (see the discussion of topological obstructions in [35, p. 1254]), in which case one can either weaken the continuity to piecewise continuity, or work with multidimensional generalizations (as in [35]). Finally, note that the concept of nilspace character can be extended to compact abelian groups, by requiring that ϕ be a *continuous* morphism. As usual in this paper, we focus on finite abelian groups Z.

Remark 4.6. In main results involving nilspace polynomials or characters, especially in inverse theorems or regularity results, some of the above-mentioned additional parameters actually *must* be adequately bounded in order for these results to be nontrivial. Indeed, concerning nilspace polynomials, note that if in Definition 4.1 we do not ensure a Lipschitz bound on F independent of $|\mathrm{Z}|$, then *any* 1-bounded function on Z could be a valid nilspace polynomial, as one could take X to be a copy of Z embedded in some torus (equipped with its standard metric), take φ the identity map, and $F = f$. For nilspace characters, the role of Lipschitz bounds is more subtle, but it is still important in our results. For instance, such a Lipschitz bound is involved in our expression of 2-step nilspace characters as quadratic characters of bounded complexity (see Theorem 4.19 and Remark 4.20).

In the following subsections, we develop various properties of nilspace characters, especially to prove that they are adequate higher-order generalizations of classical Fourier characters.

4.1. Uniform approximation of Lipschitz functions by vertical Fourier sums

The structure theorem (or regularity lemma) from [12, Theorem 1.5] yields, as a key ingredient of the structured part in the resulting decomposition, a function on a compact nilspace that has a bounded Lipschitz norm. It will be very useful for later calculations to be able to approximate such a function by a sum of boundedly many functions with vertical frequency. That is the purpose of the main result in this subsection. To obtain this, we first gather some tools related to Lipschitz functions. We shall recall some basic definitions and notation following [14, §8.4.2].

[9] See [12, Definition 1.2].
[10] We recall this notion in Definition 4.17 below.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Given a metric space X, we say that a metric d on X is a *compatible metric* if d generates the given topology on X.

Definition 4.7. Let X be a compact metric space with a compatible metric d. For any continuous function $f : X \to \mathbb{C}$, we write $\|f\|_L$ for the quantity $\sup_{x\neq y} \frac{|f(x)-f(y)|}{d(x,y)}$. We say that f is *Lipschitz* if $\|f\|_L$ is finite, and we say that f is *C-Lipschitz* if $\|f\|_L \leq C$. We define $\|f\|_{\mathrm{sum}} := \|f\|_L + \|f\|_\infty$.

Note that $\|\cdot\|_{\mathrm{sum}}$ is a norm whereas $\|\cdot\|_L$ is only a seminorm.

Given a k-step compact nilspace X, we say a compatible metric d on X is $\mathrm{Z}_k(\mathrm{X})$-*invariant* if $d(x+z, y+z) = d(x,y)$ for all $x, y \in \mathrm{X}$ and $z \in \mathrm{Z}_k(\mathrm{X})$.

Lemma 4.8. *Let* X *be a k-step compact nilspace and let d be a* $\mathrm{Z}_k(\mathrm{X})$*-invariant compatible metric on* X*. Then, for every Lipschitz function* $F : \mathrm{X} \to \mathbb{C}$ *(relative to d) and every* $\chi \in \widehat{\mathrm{Z}_k(\mathrm{X})}$*, the function* $F_\chi : \mathrm{X} \to \mathbb{C}$, $x \mapsto \int_{\mathrm{Z}_k(\mathrm{X})} F(x+z)\overline{\chi(z)}\,\mathrm{d}z$ *satisfies* $\|F_\chi\|_L \leq \|F\|_L$ *and* $\|F_\chi\|_\infty \leq \|F\|_\infty$.

Proof. For all $x, y \in \mathrm{X}$, using the $\mathrm{Z}_k(\mathrm{X})$-invariance we have

$$|F_\chi(x) - F_\chi(y)| \leq \int_{\mathrm{Z}_k(\mathrm{X})} |F(x+z) - F(y+z)|\,\mathrm{d}z \leq \int_{\mathrm{Z}_k(\mathrm{X})} Cd(x+z, y+z)\,\mathrm{d}z = Cd(x,y),$$

so $\|F_\chi\|_L \leq \|F\|_L$. We also have $|F_\chi(x)| \leq \int_{\mathrm{Z}_k(\mathrm{X})} |F(x+z)|\,\mathrm{d}z \leq \|F\|_\infty$. □

Remark 4.9. Note that for Lemma 4.8 to work, we needed the metric d on X to be $\mathrm{Z}_k(\mathrm{X})$-invariant. By [7, Lemma 2.1.11], one can equip every compact k-step nilspace with such a metric.[11]

Remark 4.10. The Lipschitz constant of a function is not a topological property of the space, in the sense that given a compact metric space S and a continuous function $F : S \to \mathbb{C}$, different compatible metrics can make F be Lipschitz or not. This holds even if we assume that the space is a k-step compact nilspace X and all the considered metrics are $\mathrm{Z}_k(\mathrm{X})$-invariant. For instance, let $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ and let $|x|_{\mathbb{T}}$ be the usual distance from x to the nearest integer, under the usual identification of x with a real number in $[0,1)$. Clearly $d(x,y) := |x-y|_{\mathbb{T}}$ is a $\mathbb{T}$-invariant metric on $\mathbb{T}$. Moreover, so is $d_\alpha(x,y) := |x-y|_{\mathbb{T}}^\alpha$ for any $\alpha \in (0,1]$. Let $F : \mathbb{T} \to \mathbb{R}$ be the function defined by $F(x) = \sqrt{x}$ for $x \in [0, 1/2)$ and $F(x) = \sqrt{1-x}$ otherwise. It can be checked that F is a 1-Lipschitz function relative to the metric $d_{1/2}$ on $\mathbb{T}$, but it is not Lipschitz relative to d.

When working on CFR nilspaces, problematic metrics, such as in the above example, can be conveniently ruled out by focusing on metrics induced by the manifold structure of these spaces, that is, by *Riemannian metrics*. This generalizes a standard approach in the setting of nilmanifolds, for example, in [35]. We gather some basic facts on Riemannian metrics in Appendix A (see in particular Definition A.11) and use these to get the following approximation result announced above. For simplicity, compact abelian Lie groups will always be assumed to be endowed with the *flat metric*, see Definition A.12.[12]

Lemma 4.11. *Let* X *be a k-step* CFR *nilspace endowed with a Riemannian metric. Then there is a function* $\mathrm{Q}_{\mathrm{X}} : \mathbb{R}_{>0} \times \mathbb{R}_{\geq 0} \to \mathbb{N}$ *such that the following holds. For any* $C, \delta > 0$*, there is a set* $S \subset \widehat{\mathrm{Z}_k(\mathrm{X})}$ *with* $|S| \leq \mathrm{Q}_{\mathrm{X}}(\delta, C)$ *such that, for any C-Lipschitz function* $F : \mathrm{X} \to \mathbb{C}$ *(relative to the given metric), we have* $\|F - \sum_{\chi\in S} F_\chi\|_\infty \leq \delta$.

This includes, as a special case, a uniform approximation result for Lipschitz functions on nilmanifolds. Uniform approximation results by sums of functions with vertical frequency have been used in

[11] Later, we will equip k-step CFR nilspaces with metrics coming from Riemannian tensors, see Appendix A. For those we can also guarantee $\mathrm{Z}_k(\mathrm{X})$-invariance, see Lemma A.10.

[12] Modulo minor changes in the constants, all results work the same way if we endow compact abelian Lie groups with any Riemannian metric. This is because by Proposition A.13, the identity map $\mathrm{id} : \mathrm{X} \to \mathrm{X}$ is bi-Lipschitz when we consider any two Riemannian metrics on X. Hence, if a map $F : \mathrm{X} \to \mathbb{C}$ is Lipschitz with respect to a Riemannian metric, then it is also Lipschitz relative to any other Riemannian metric. We fix one such metric in Definition A.12 and use this one unless otherwise stated.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

this setting previously (e.g., [35, Lemma 6.4]). Note, however, that Lemma 4.11 yields an approximation by a partial Fourier series of F relative to the last structure group, which is more precise than an approximation in terms of *some* functions with vertical frequencies.

The proof of Lemma 4.11 will use the following special case for compact abelian Lie groups.

Lemma 4.12. *Let G be a compact abelian Lie group. Then there exists $\mathrm{Q}_G : \mathbb{R}_{>0} \times \mathbb{R}_{\geq 0} \to \mathbb{N}$ such that the following holds. For any $C, \delta > 0$, there is a set $S \subset \widehat{G}$ with $|S| \leq \mathrm{Q}_G(\delta, C)$ such that for any C-Lipschitz function $f : G \to \mathbb{C}$ we have $\|f - \sum_{\chi \in S} \widehat{f}(\chi)\chi\|_\infty < \delta$.*

Proof. We know that $G \cong \mathrm{Z} \times \mathbb{T}^n$ for some finite abelian group Z and some integer $n \geq 0$.

First, suppose that Z is trivial, so that $G \cong \mathbb{T}^n$. Then, it follows from known results (due to Lebesgue for $n = 1$, and to Zhizhiashvili for $n > 1$, see [49, p. 201] and [21, p. 642] respectively) that if f is C-Lipschitz on $\mathbb{T}^n$ then the Fourier sums $S_M(f) := \sum_{r \in \mathbb{Z}^n : \|r\|_{\ell^\infty} \leq M} \widehat{f}(\chi_r)\chi_r$ on $\mathbb{T}^n$ satisfy $\|S_M(f) - f\|_\infty \leq K_{C,n} \frac{\log^n M}{M}$, where the constant $K_{C,n}$ depends only on C and n.[13] Hence the claim holds for M sufficiently large in terms of δ, C, and n.

Now if $G = \mathrm{Z} \times \mathbb{T}^n$ for a general finite abelian group Z, then for every $a \in \mathrm{Z}$ and $x \in \mathbb{T}^n$ we have $f(a, x) = \sum_{\xi \in \widehat{\mathrm{Z}}} \left(\mathbb{E}_{z \in \mathrm{Z}} f(z, x)\overline{\xi(z)}\right) \xi(a)$. For every $z \in \mathrm{Z}$ the function $f(z, \cdot)$ is C-Lipschitz on $\mathbb{T}^n$, so we can apply the previous paragraph with M such that $\|S_M(f(z, \cdot)) - f(z, \cdot)\|_\infty \leq \delta/|\mathrm{Z}|$. The result follows with $S = \{(\xi, \chi_r) \in \widehat{\mathrm{Z}} \times \widehat{\mathbb{T}^n} = \widehat{G} : \|r\|_{\ell^\infty} \leq M\}$. □

Proof of Lemma 4.11. Let $\mathrm{X} \times \mathrm{Z}_k(\mathrm{X}) \to \mathrm{X}$, $(x, z) \mapsto x + z$ be the addition map. On $\mathrm{X} \times \mathrm{Z}_k(\mathrm{X})$, we consider the usual Riemannian metric tensor which is the product of the Riemannian tensor of X with the usual Riemannian tensor on the compact abelian group $\mathrm{Z}_k(\mathrm{X})$. Note that, in particular, $(x, z) \mapsto x+z$ is a morphism between k-step CFR nilspaces (regarding $\mathrm{Z}_k(\mathrm{X})$ as the group nilspace $\mathcal{D}_k(\mathrm{Z}_k(\mathrm{X}))$) and by Proposition A.13 we know that it is Lipschitz. Hence the map $(x, z) \mapsto F(x+z)$ (being the composition of two Lipschitz maps) is C'-Lipschitz with C' depending only on X and C. We claim that, for any fixed $x_0 \in \mathrm{X}$, the map $z \mapsto F(x_0 + z)$ is also C'-Lipschitz.

Indeed, we have $|F(x_0 + z) - F(x_0 + z')| \leq C' d_{\mathrm{X} \times \mathrm{Z}_k(\mathrm{X})}((x_0, z), (x_0, z'))$. We claim that, for all $z, z' \in \mathrm{Z}_k(\mathrm{X})$ and $x_0 \in \mathrm{X}$, we have $d_{\mathrm{Z}_k(\mathrm{X})}(z, z') = d_{\mathrm{X} \times \mathrm{Z}_k(\mathrm{X})}((x_0, z), (x_0, z'))$. Note that (x_0, z) and (x_0, z') are in different connected components if and only if z and z' are in different connected components. Hence, in this case we have $d_{\mathrm{Z}_k(\mathrm{X})}(z, z') = 1 = d_{\mathrm{X} \times \mathrm{Z}_k(\mathrm{X})}((x_0, z), (x_0, z'))$ by definition. If z and z' lie in the same connected component, let $\gamma : [0, 1] \to \mathrm{X} \times \mathrm{Z}_k(\mathrm{X})$, $t \mapsto (x(t), z(t))$ be a path connecting (x_0, z) with (x_0, z'). Then $L_\gamma = \int_0^1 \sqrt{g_{\mathrm{X} \times \mathrm{Z}_k(\mathrm{X})}(\nabla\gamma, \nabla\gamma)}$ where $g_{\mathrm{X} \times \mathrm{Z}_k(\mathrm{X})}$ is the Riemannian metric tensor on $\mathrm{X} \times \mathrm{Z}_k(\mathrm{X})$. As $g_{\mathrm{X} \times \mathrm{Z}_k(\mathrm{X})}(\nabla\gamma, \nabla\gamma) = g_{\mathrm{X}}(\nabla x(t), \nabla x(t)) + g_{\mathrm{Z}_k(\mathrm{X})}(\nabla z(t), \nabla z(t))$, if we let $\gamma' : [0, 1] \to \mathrm{Z}_k(\mathrm{X})$, $t \mapsto z(t)$, we have that $L_\gamma \geq L_{\gamma'}$. Thus, $d_{\mathrm{X} \times \mathrm{Z}_k(\mathrm{X})}((x_0, z), (x_0, z')) \geq d_{\mathrm{Z}_k(\mathrm{X})}(z, z')$. By a similar argument we obtain the reversed inequality, and the claim follows.

Therefore, since the map $z \mapsto F(x_0 + z)$ is C'-Lipschitz, we can apply Lemma 4.12 to it and complete the proof (noting that the Lipschitz constant C' does not depend on $x_0 \in \mathrm{X}$). □

Remark 4.13. Note that the proof of Lemma 4.12 shows that, in Lemma 4.11, letting Z be a finite abelian group and $n \geq 0$ be an integer such that $\mathrm{Z}_k(\mathrm{X}) \cong \mathrm{Z} \times \mathbb{T}^n$, we can add the information that the set S is of the concrete form $S = \{(\xi, \chi_r) \in \widehat{\mathrm{Z}} \times \widehat{\mathbb{T}^n} : r \in \mathbb{Z}^n, \|r\|_{\ell^\infty} \leq M\}$ for some sufficiently large integer M.

We close this section by discussing another advantage that functions with a uniform bound on their Lipschitz constant present for our later calculations, namely, that we can then have uniform control on the convergence of integrals. Before we go into this, let us recall a standard example showing that such uniform control fails without a uniform Lipschitz bound.

Example 4.14. Let $I = [0, 1]$ with the usual topology and for each positive integer n consider the Dirac measure $\delta_{1/n}$. These measures converge weakly to the Dirac measure δ_0, that is, for every continuous function $f : I \to \mathbb{C}$ we have $\int f \, \mathrm{d}\delta_{1/n} \to \int f \, \mathrm{d}\delta_0$ as $n \to \infty$. However, this convergence is not uniform

[13] In [21, p. 642] the metric used on $\mathbb{T}^n$ is the Euclidean one; but as noted earlier, by Proposition A.13, this is equivalent to the fixed Riemannian metric on $\mathbb{T}^n$, and the associated constants are taken into account in $K_{C,n}$.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

in f, as can be seen with the sequence of continuous functions $f_n(x) = (1-nx)\mathbf{1}_{[0,1/n]}(x)$, $n \in \mathbb{N}$, which satisfies $\int f_n \, d\delta_{1/n} - \int f_n \, d\delta_0 = -1$, and so $\sup_{\|f\|_\infty \le 1} \left| \int f \, d\delta_{1/n} - \int f \, d\delta_0 \right| \not\to 0$ as $n \to \infty$.

Such examples could a priori be problematic for our proofs later, because the *balance* property (recalled below) is defined in terms of weak convergence of measures, and we then do not have uniform convergence of integrals. However, restricting our attention to Lipschitz functions with uniform Lipschitz constant will enable us to avoid such issues. A convenient way to implement such a uniform control is to use the Kantorovich-Rubinstein (or Wasserstein) norm.

Definition 4.15 (Kantorovich-Rubinstein norm). Let X be a compact metric space, and let $\mathcal{P}(X)$ denote the set of Borel probability measures on X. The *Kantorovich-Rubinstein norm* on $\mathcal{P}(X)$ is defined as follows:

$$\|\mu\|_{KR} := \sup_{f: X \to \mathbb{C} \text{ Lipschitz with } \|f\|_{\text{sum}} \le 1} \left| \int f \, d\mu \right|.$$

The full definition of $\|\mu\|_{KR}$ is as a norm on the vector space of Borel measures with bounded variation (see [14, Theorem 8.4.5]). A convenient fact about this norm is its direct relation with weak convergence of measures, as the following result indicates (this is a direct consequence of [14, Theorem 8.4.13]).

Theorem 4.16. *Let X be a compact metric space and let $\mu, \mu_1, \mu_2, \dots$ be measures in $\mathcal{P}(X)$. Then $\mu_j \to \mu$ weakly if and only if $\|\mu_j - \mu\|_{KR} \to 0$.*

We will use Theorem 4.16 to transform the *balance* property of the regularity lemma [12, Theorem 1.5] into a bound for the Kantorovich-Rubinstein norm. Let us recall the definition of balance (see [12, Definition 5.1]).

Definition 4.17. Let Y be a k-step CFR nilspace. For each $n \in \mathbb{N}$ fix a metric d_n on $\mathcal{P}(\mathrm{C}^n(\mathrm{Y}))$. Let X be a k-step compact nilspace and $\phi : \mathrm{X} \to \mathrm{Y}$ be a continuous morphism. For $b > 0$, we say that ϕ is *b-balanced* if for every $n \le 1/b$ we have $d_n(\mu_{\mathrm{C}^n(\mathrm{X})} \circ (\phi^{\llbracket n \rrbracket})^{-1}, \mu_{\mathrm{C}^n(\mathrm{Y})}) \le b$.

Corollary 4.18. *Let Y be a k-step CFR nilspace. For each $n \in \mathbb{N}$ fix a metric d_n on $\mathcal{P}(\mathrm{C}^n(\mathrm{Y}))$ and a metric d'_n on $\mathrm{C}^n(\mathrm{Y})$. Then the following properties hold:*

- *For every $\delta > 0$ there exists $b = b(\delta) > 0$ such that if $\phi : \mathrm{X} \to \mathrm{Y}$ is a b-balanced morphism, then for all $n \le 1/\delta$ we have $\|\mu_{\mathrm{C}^n(\mathrm{X})} \circ (\phi^{\llbracket n \rrbracket})^{-1} - \mu_{\mathrm{C}^n(\mathrm{Y})}\|_{KR} \le \delta$.*
- *For every $b > 0$ there exists $\delta = \delta(b) > 0$ such that if for every $n \le 1/\delta$ we have $\|\mu_{\mathrm{C}^n(\mathrm{X})} \circ (\phi^{\llbracket n \rrbracket})^{-1} - \mu_{\mathrm{C}^n(\mathrm{Y})}\|_{KR} \le \delta$, then ϕ is b-balanced.*

Proof. For any fixed $n \ge 0$, the space $\mathcal{P}(\mathrm{C}^n(\mathrm{Y}))$ is endowed with the weak topology, which is metrizable. The balance property is using by definition one such possible metric d_n, see [12, Definition 5.1]. By Theorem 4.16, the Kantorovich-Rubinstein norm also defines a metric for the weak topology on $\mathcal{P}(\mathrm{C}^n(\mathrm{Y}))$. The result follows, as by Theorem 4.16 convergence of a sequence $\mu_i \in \mathcal{P}(\mathrm{C}^n(\mathrm{Y}))$ relative to d_n is equivalent to convergence relative to $\|\cdot\|_{KR}$. □

4.2. 2-step nilspace characters are quadratic characters

In this section we prove the following result.

Theorem 4.19. *Let X be a 2-step CFR nilspace and let $\sigma \in (0, 1/2)$. Let Z be a finite abelian group, let $\phi : \mathrm{Z} \to \mathrm{X}$ be a morphism, and let $F \circ \phi$ be a 2-step nilspace character on Z with vertical frequency χ. Then, for every $h \in \mathrm{Z}$, we have $\Delta_h(F \circ \phi) = \sum_{\gamma \in S_h} \lambda_{h,\gamma} \gamma(z) + \mathcal{E}_h(z)$ where $S_h \subset \widehat{\mathrm{Z}}$ satisfies $|S_h| = O_{\mathrm{X}, \|F\|_{sum}}(\sigma^{-O_{\mathrm{X}}(1)})$, $\mathcal{E}_h : \mathrm{Z} \to \mathbb{C}$ satisfies $\|\mathcal{E}_h\|_\infty \le \sigma$, and for all $h \in \mathrm{Z}$ and $\gamma \in S_h$ we have $|\lambda_{h,\gamma}| \le \|F\|_\infty^2$.*

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Remark 4.20. In particular, this establishes that any such 2-step nilspace character $F \circ \phi$ is a 1-bounded (R, σ)-character of order 2 on Z, in the sense of Definition 2.29, with precision parameter σ and complexity parameter $R = O_{\mathrm{X}, \|F\|_{\mathrm{sum}}}(\sigma^{-O_{\mathrm{X}}(1)})$. Note the additional strength (compared with Definition 2.29) in that the error $\mathcal{E}_h$ here is small in $\|\cdot\|_\infty$, not just in $\|\cdot\|_2$.

Proof. Since X is 2-step and of finite rank, the 1-step nilspace factor X_1 is isomorphic to a compact abelian Lie group, which is therefore of the form $A \times \mathbb{T}^\ell$ for some finite abelian group A and some $\ell \in \mathbb{Z}_{\geq 0}$. In particular, X_1 has exactly $|A|$ connected components.

Now observe that $\pi_1 \circ \phi$ is a morphism between 1-step nilspaces (from $\mathcal{D}_1(\mathrm{Z})$ to X_1), so it is an affine homomorphism of the corresponding abelian groups, and we therefore have $\pi_1 \circ \phi = \psi + \pi_1(\phi(0))$ where $\psi : \mathrm{Z} \to A \times \mathbb{T}^\ell$ is a continuous group homomorphism.

Let $\Theta(\mathrm{X}_1)^0$ be the connected component of the identity of the translation group of X_1. Note that this is isomorphic to the subgroup $\{0\} \times \mathbb{T}^\ell$ of $A \times \mathbb{T}^\ell$ (since $\Theta(\mathrm{X}_1) \cong A \times \mathbb{T}^\ell = \mathrm{Z}_1(\mathrm{X}_1)$). Thus, for every element $(0, t) \in A \times \mathbb{T}^\ell$ there is an element $\beta_t^{(1)} = \beta_{(0,t)}^{(1)} \in \Theta(\mathrm{X}_1)^0$ such that, for every $x \in \mathrm{X}$, $\beta_t^{(1)} \pi_1(x) = \pi_1(x) + (0, t)$. Let $h : \Theta(\mathrm{X}) \to \Theta(\mathrm{X}_1)$ be the continuous homomorphism defined as in [7, Lemma 2.9.3], that is, for $\alpha \in \Theta(\mathrm{X})$ we let $h(\alpha)$ be the unique map such that, for every $x \in \mathrm{X}$, we have $h(\alpha)(\pi_1(x)) = \pi_1(\alpha(x))$. By [7, Theorem 2.9.10 (*ii*)], the restriction of h to $\Theta(\mathrm{X})^0$ defines a surjective homomorphism between Lie groups $h|_{\Theta(\mathrm{X})^0} : \Theta(\mathrm{X})^0 \to \Theta(\mathrm{X}_1)^0$. By the Open Mapping Theorem for topological groups, $h|_{\Theta(\mathrm{X})^0}$ is an open map and in particular $\Theta(\mathrm{X}_1)^0 \cong \Theta(\mathrm{X})^0 / \ker(h|_{\Theta(\mathrm{X})^0})$. Hence, the map $h|_{\Theta(\mathrm{X})^0}$ induces an open Hausdorff equivalence relation, see [4, Chapter 1, §5, 2 and §8, 3]. By [4, p. 107, Proposition 10] there exists a compact $B \subset \Theta(\mathrm{X})^0$ such that $h|_{\Theta(\mathrm{X})^0}(B) = \Theta(\mathrm{X}_1)^0$. Thus, for every $t \in \mathbb{T}^\ell$, we let $\beta_t \in B$ be such that $h|_{\Theta(\mathrm{X})^0}(\beta_t) = \beta_1^{(1)}$.[14] Note that B can be chosen depending only on X. Thus, in what follows, we assume that for each X, we make one such choice of a set B. Hence, if some quantity or object depends on B, we will simply say that it depends on X.

Now, for each $a \in A$, consider the *translation bundle* corresponding to $\beta_a^{(1)}$, namely the following standard construction in nilspace theory (see [6, Definition 3.3.34]):

$$\mathrm{Y}_a := \mathcal{T}(\mathrm{X}, \beta_a^{(1)}, 1) = \{(x_0, x_1) \in \mathrm{X}^2 : \pi_1(x_0) + (a, 0) = \pi_1(x_1)\}.$$

By [6, Lemma 3.3.35] we know that Y_a is a 2-step compact nilspace that is a subnilspace of the arrow space $\mathrm{X} \bowtie_1 \mathrm{X}$. Abusing the notation we will identify $\mathrm{Y}_{(a,0)} = \mathrm{Y}_a$.

By [6, Lemma 3.3.38], the 1-step canonical factor $\pi_1(\mathrm{Y}_a)$ (denoted $\mathcal{T}^*$ in [6]) is, from the purely algebraic viewpoint, a degree- 1 extension of X_1 by $\mathrm{Z}_2(\mathrm{X})$ and, by the results from [7, Section 2.1], this extension is a continuous $\mathrm{Z}_2(\mathrm{X})$-bundle over X_1. In particular Y_1 is a compact 1-step nilspace, that is, an affine *compact abelian group*, which is an extension of the affine compact abelian group X_1 by $\mathrm{Z}_2(\mathrm{X})$.

Recall also that, by [6, Definition 3.3.34 and Proposition 3.3.36], the nilspace Y_a is a continuous Z-bundle over $\pi_1(\mathrm{Y}_a)$, for the Polish group $\mathrm{Z} = \mathrm{Z}_2(\mathrm{X})^\Delta = \{(z, z) : z \in \mathrm{Z}_2(\mathrm{X})\} \leq \mathrm{Z}_2(\mathrm{X}) \times \mathrm{Z}_2(\mathrm{X})$. In other words, the nilspace factor map $\pi_1 : \mathrm{Y}_a \to \pi_1(\mathrm{Y}_a)$ is the map sending any $(x_0, x_1) \in \mathrm{Y}_a$ to its orbit under the Z-action, that is, the action of the diagonal subgroup $\mathrm{Z}_2(\mathrm{X})^\Delta$.

Let p denote the projection homomorphism $(a, t) \mapsto (a, 0)$ on $A \times \mathbb{T}^\ell$.

Our aim is to show that the multiplicative derivative $\Delta_h(F \circ \phi) : z \mapsto \overline{F \circ \phi(z)} F(\phi(z + h))$ factors through $\mathrm{Y}_{p(\psi(h))}$. To this end, we note the decomposition $\Delta_h(F \circ \phi)(z) = V_h \circ \delta_h(z)$, where

$$\begin{array}{rccc} \delta_h : & \mathcal{D}_1(\mathrm{Z}) & \to & \mathrm{Y}_{p(\psi(h))} \\ & z & \mapsto & \left(\phi(z), \beta_{p(\psi(h)) - \psi(h)}\, \phi(z + h)\right) \end{array}$$

[14] There may be different elements $\beta_t, \beta_t' \in B$ such that $h|_{\Theta(\mathrm{X})^0}(\beta_t) = h|_{\Theta(\mathrm{X})^0}(\beta_t') = \beta_1^{(1)}$. For each t we simply choose one such element. This choice can be made in a measurable way, but we omit the proof as we shall not use this fact.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

and

$$\begin{aligned} V_h : \mathrm{Y}_{p(\psi(h))} &\to \mathbb{C} \\ (x_0, x_1) &\mapsto \overline{F(x_0)} F(\beta^{-1}_{p(\psi(h))-\psi(h)} x_1). \end{aligned}$$

Let us check first that these maps are well-defined. The first thing to note is that $p(\psi(h)) - \psi(h)$ is in $\{0\} \times \mathbb{T}^\ell$, so the lift $\beta_{p(\psi(h))-\psi(h)}$ indeed exists as defined above. Next, we need to prove that $\delta_h(z) \in \mathrm{Y}_{p(\psi(h))}$. That is, we need to check that

$$\pi_1(\phi(z)) + p(\psi(h)) = \pi_1\big(\beta_{p(\psi(h))-\psi(h)}\, \phi(z+h)\big).$$

But the right hand side here equals by definition $p(\psi(h)) - \psi(h) + \pi_1(\phi(z+h)) = p(\psi(h)) - \psi(h) + \pi_1(\phi(0)) + \psi(z+h) = \pi_1(\phi(z)) + p(\psi(h))$, as required.

We claim that δ_h is a nilspace morphism. To see this, note that, given an n-cube c on Z, it suffices to check that $(\phi \circ \mathrm{c}, \phi \circ (\mathrm{c}+h))$ is an n-cube on $\mathrm{X} \bowtie_1 \mathrm{X}$, that is, that the 1-arrow $\langle \phi \circ \mathrm{c}, \phi \circ (\mathrm{c}+h)\rangle_1$ is in $\mathrm{C}^{n+1}(\mathrm{X})$. Indeed, if we show this then it will follow that $\delta_h \circ \mathrm{c}$ is also a cube, since the additional application of the translation $\beta_{p(\psi(h))-\psi(h)}$ in the second coordinate preserves cubes by definition of translations. But the 1-arrow $\langle \phi \circ \mathrm{c}, \phi \circ (\mathrm{c}+h)\rangle_1$ is clearly a cube because it is just $\phi \circ \langle \mathrm{c}, \mathrm{c}+h\rangle_1$, and this is indeed an $(n+1)$-cube on X since $\langle \mathrm{c}, \mathrm{c}+h\rangle_1$ is an $(n+1)$-cube on Z and ϕ is a morphism.

Now note that the space $\mathrm{Y}_{p(\psi(h))}$ is among the finitely many spaces Y_a, $a \in A$. On the other hand, by Corollary A.15, V_h is a C-Lipschitz function with C bounded in terms of $\|F\|_L$, X, and $\mathrm{Y}_{p(\psi(h))}$.[15] Hence, we can take this bound C to be uniform in h, since $\mathrm{Y}_{p(\psi(h))}$ can take at most $|A|$ different values, and this number $|A|$ depends only on X_1. Therefore $\|V_h\|_L = O_{\mathrm{X},\|F\|_{\mathrm{sum}}}(1)$. Moreover V_h is invariant under the action of the structure group $\mathrm{Z}_2(\mathrm{X})$ acting diagonally on $\mathrm{Y}_{p(\psi(h))}$, so there exists $W_h : \pi_1(\mathrm{Y}_{p(\psi(h))}) \to \mathbb{C}$ such that $W_h \circ \pi_1 = V_h$. Hence, by Lemma A.17 we have $\|W_h\|_L = O_{\mathrm{X},\|F\|_{\mathrm{sum}}}(1)$ as well.

Finally, by Proposition A.13 the metric on $\pi_1(\mathrm{Y}_{p(\psi(h))})$ is Lipschitz equivalent to the usual metric on any compact abelian Lie group. Hence, by Lemma 4.12 we can approximate the function W_h in L^∞ by a linear combination of boundedly many characters:

$$\forall\, (x_0, x_1) \in \mathrm{Y}_{p(\psi(h))}, \quad W_h(\pi_1(x_0, x_1)) = \textstyle\sum_{\gamma \in S} \lambda_{h,\gamma} \gamma(\pi_1(x_0, x_1)) + \mathcal{E}_h(\pi_1(x_0, x_1))$$

where $\|\mathcal{E}_h\|_\infty \le \sigma$, $\lambda_{h,\gamma} = \widehat{W_h}(\gamma)$ (so, in particular $|\lambda_{h,\gamma}| \le \|F\|_\infty^2$), and $|S| = O_{\sigma,\mathrm{X},\|F\|_{\mathrm{sum}}}(1)$. Moreover, from the proof of Lemma 4.12 we can give a more precise estimate of the size of S. Assume that the 1-step nilspace $\mathrm{Y}_{p(\psi(h))} \cong \mathcal{D}_1(A' \times \mathbb{T}^{\ell'})$ where A' and ℓ' depend solely on X (and $p(\psi(h))$, but as this is a finite set depending only on X we regard such a dependence as a dependence in X). By Remark 4.13, note that letting $S = \{(\widetilde{\chi}, \chi_r) \in \widehat{A'} \times \widehat{\mathbb{T}^{\ell'}} : r \in \mathbb{Z}^{\ell'}, \|r\|_\infty \le M\}$ then for the given σ we need to ensure that the upper bound $O_{\mathrm{X},\|F\|_{\mathrm{sum}}}(\frac{\log^{\ell'} M}{M})$ is at most σ. Since $\frac{\log^{\ell'} M}{M} \le \frac{O_{\ell'}(1)}{M^{1/2}}$, it suffices to assume that $M \ge \Omega_{\mathrm{X},\|F\|_{\mathrm{sum}}}(\sigma^{-2})$. The estimate on the size of S follows by taking the ℓ'-th power of such a quantity and multiplying by $|A'|$.

Hence, we conclude that $\overline{F \circ \phi(z)}\, T^h(F \circ \phi)(z) = \sum_{\gamma \in S} \lambda_{h,\gamma} \gamma(\pi_1(\delta_h(z))) + \mathcal{E}_h(\pi_1(\delta_h(z)))$. Now note that each function $\gamma(\pi_1(\delta_h(z)))$ is a Fourier character on Z multiplied by a complex number of modulus 1, so the result follow by relabelling said character as γ. □

Theorem 4.19 tells us that 2-step nilspace characters (Definition 4.5) are quadratic characters (Definition 2.29), and these are clearly *weak* quadratic characters (Definition 2.34). The converse direction

[15] To be more precise, note that by standard arguments and Corollary A.15 we have that there exists some $L > 0$ depending only on $\|F\|_{\mathrm{sum}}$ and on the compact set B such that for any $(x_0, x_1), (y_0, y_1) \in \mathrm{Y}_{p(\psi(h))}$ we have $|V_h(x_0, x_1) - V_h(y_0, y_1)| \le L(d_{\mathrm{X}}(x_0, y_0) + d_{\mathrm{X}}(y_0, y_1))$. But the map $\mathrm{Y}_{p(\psi(h))} \to \mathrm{X}$ defined as $(x_0, x_1) \mapsto x_0$ is a continuous morphism (and similarly the projection to x_1). Thus, by Proposition A.13, it is Lipscshitz, and thus $d(x_0, y_0) \le L' d_{\mathrm{Y}_{p(\psi(h))}}((x_0, x_1), (y_0, y_1))$ for some $L' \ge 0$. The result follows.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

is also interesting but not needed for this paper, so we do not pursue it here and instead we pose the following question.

Question 4.21. Is every weak quadratic character of bounded complexity close (in the L^2-norm, for instance) to a 2-step nilspace character of bounded complexity?

4.3. Approximate orthogonality of balanced nilspace characters

In the structure theorem (or regularity lemma) for the U^{k+1}-norm obtained in [12, Theorem 1.5], the structured part of the original function $f : \mathrm{Z} \to \mathbb{C}$ is of the form $F \circ \phi$ for some highly balanced morphism ϕ from Z into some k-step CFR nilspace X, and some Lipschitz function $F : \mathrm{X} \to \mathbb{C}$. It is then natural to apply Lemma 4.11 to the function F and examine the properties of the resulting Fourier components F_χ.

Recall the notation from Definition 3.14 for the special type of Gowers U^k products (Definition 2.31) involving only two different functions. The first main result of this subsection establishes the following: the more balanced the morphism ϕ is, the smaller the U^{k+1}-products $\langle F_\chi \circ \phi, F_{\chi'} \circ \phi \rangle_{U^{k+1}}$ will be for distinct characters χ, χ'. This will be used in the proof of the validity of our algorithms.

Proposition 4.22. *Let $\delta, C > 0$, and let X be a k-step compact nilspace endowed with a $\mathrm{Z}_k(\mathrm{X})$-invariant metric. There exists $b = b(\delta, C, \mathrm{X}) > 0$ such that if $\phi : \mathrm{Z} \to \mathrm{X}$ is a b-balanced morphism, then for every 1-bounded C-Lipschitz function $F : \mathrm{X} \to \mathbb{C}$ and every $\chi \neq \chi'$ in $\widehat{\mathrm{Z}_k(\mathrm{X})}$, we have $\langle F_\chi \circ \phi, F_{\chi'} \circ \phi \rangle_{U^{k+1}} \leq \delta$.*

Proof. Let $\mathbf{F}$ denote the function on $\mathrm{X}^{[\![k+1]\!]}$ sending $x = (x_v)_{v \in [\![k+1]\!]}$ to $\prod_{v \in [\![k+1]\!]} y_v \in \mathbb{C}$ with $y_v = F_\chi(x_v)$ if $v_{k+1} = 0$ and $y_v = F_{\chi'}(x_v)$ otherwise. Note that

$$\langle F_\chi \circ \phi, F_{\chi'} \circ \phi \rangle_{U^{k+1}} = \mathbb{E}_{\mathrm{c} \in \mathrm{C}^{k+1}(\mathrm{Z})} \mathbf{F} \circ \phi^{[\![k+1]\!]}(\mathrm{c}).$$

By Lemma 4.8, the functions F_χ and $F_{\chi'}$ are both C-Lipschitz. For every $n \in \mathbb{N}$ we can endow $\mathrm{X}^{[\![k+1]\!]}$ with the metric $d'_n(\mathrm{c}, \mathrm{c}') := \sum_{v \in [\![n]\!]} d_{\mathrm{X}}(\mathrm{c}(v), \mathrm{c}'(v))$ (and $\mathrm{C}^n(\mathrm{X})$ with the restriction of this metric) which ensures that $\mathbf{F}$ has $\|\mathbf{F}\|_{\mathrm{sum}} \leq 1 + C$ relative to that metric.[16]

Let $\eta > 0$ be a parameter to be fixed later. By Corollary 4.18 (applied with any fixed metrics d_n on $\mathcal{P}(\mathrm{C}^n(\mathrm{X}))$ and the metrics d'_n on $\mathrm{C}^n(\mathrm{X})$ for $n \in \mathbb{N}$ that we have just defined), we know that there exists $b = b(\eta, \mathrm{X}) > 0$ sufficiently small such that, if $\phi : \mathrm{Z} \to \mathrm{X}$ is b-balanced, then $\|\mu_{\mathrm{C}^{k+1}(\mathrm{Z})} \circ (\phi^{[\![k+1]\!]})^{-1} - \mu_{\mathrm{C}^{k+1}(\mathrm{X})}\|_{KR} \leq \eta$ (here we assume that the metrics d, d_n, d'_n are chosen *a priori* depending only on X, hence we only record such a dependence). By definition of $\|\cdot\|_{KR}$, we then have

$$\left| \int_{\mathrm{C}^{k+1}(\mathrm{X})} \mathbf{F}(\mathrm{c}) \, \mathrm{d}\mu_{\mathrm{C}^{k+1}(\mathrm{X})}(\mathrm{c}) - \int_{\mathrm{C}^{k+1}(\mathrm{Z})} \mathbf{F} \circ \phi^{[\![k+1]\!]}(\mathrm{c}) \, \mathrm{d}\mu_{\mathrm{C}^{k+1}(\mathrm{Z})}(\mathrm{c}) \right| \leq \eta \|\mathbf{F}\|_{\mathrm{sum}} \leq \delta$$

where η is chosen to be $\delta/(C+1)$.

We now claim that $\int_{\mathrm{C}^{k+1}(\mathrm{X})} \mathbf{F}(\mathrm{c}) \, \mathrm{d}\mu_{\mathrm{C}^{k+1}(\mathrm{X})}(\mathrm{c}) = 0$. Confirming this will complete the proof. Recall that the Haar measure on $\mathrm{C}^{k+1}(\mathrm{X})$ is preserved by the action of the cube group $G := \mathrm{C}^{k+1}(\mathcal{D}_k(\mathrm{Z}_k(\mathrm{X})))$, by [7, Lemma 2.2.6]. Hence

$$\int_{\mathrm{C}^{k+1}(\mathrm{X})} \mathbf{F}(\mathrm{c}) \, \mathrm{d}\mu_{\mathrm{C}^{k+1}(\mathrm{X})}(\mathrm{c}) = \int_{\mathrm{C}^{k+1}(\mathrm{X})} \int_G \mathbf{F}(\mathrm{c} + t) \, \mathrm{d}\mu_G(t) \, \mathrm{d}\mu_{\mathrm{C}^{k+1}(\mathrm{X})}(\mathrm{c}).$$

[16] To prove this, note that as $\|F\|_\infty \leq 1$ we have $\|\mathbf{F}\|_\infty \leq 1$. For the Lipschitz constant, note that, by our choice of metric on $\mathrm{C}^n(\mathrm{X})$, we can use a standard telescoping-sum argument. More precisely, let $(v_i)_{i=1,\dots,2^{k+1}} = [\![k+1]\!]$ be an arbitrary enumeration of the elements of $[\![k+1]\!]$. Then, for $i = 1, \dots, 2^{k+1}$, let $\mathrm{c}''_i \in \mathrm{X}^{[\![k+1]\!]}$ be defined as $\mathrm{c}''_i(v_j) = \mathrm{c}(v_j)$ for $j \leq i$ and $\mathrm{c}'(v)$ otherwise. Then $|\mathbf{F}(\mathrm{c}) - \mathbf{F}(\mathrm{c}')| \leq \sum_{i=1}^{2^{k+1}} |\mathbf{F}(\mathrm{c}''_i) - \mathbf{F}(\mathrm{c}''_{i+1})| \leq \sum_{i=1}^{2^{k+1}} C d_{\mathrm{X}}(\mathrm{c}(v_i), \mathrm{c}'(v_i))$. Hence $\|\mathbf{F}\|_L \leq C$.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

But by definition of $\mathbf{F}$, for every c we have

$$\int_G \mathbf{F}(\mathrm{c}+t)\,\mathrm{d}\mu_G(t) = \mathbf{F}(\mathrm{c})\int_G \prod_{v\in[\![n]\!]}\chi_v(t(v))\,\mathrm{d}\mu_G(t)$$

where $\chi_v = \chi$ if $v_{k+1} = 0$ and $\chi_v = \chi'$ otherwise. By [11, Lemma 3.52] the last integral above is 0 and the claim follows. □

The second main result is that ϕ being highly balanced also ensures that distinct nilspace characters $F_\chi \circ \phi$ and $F_{\chi'} \circ \phi$ are quasiorthogonal in the following sense.

Definition 4.23. Given an inner-product space H and $\delta > 0$, we say that two elements $f, g \in H$ are *δ-quasiorthogonal* if $|\langle f, g\rangle| \le \delta$.

Theorem 4.24. *Let $\delta, C > 0$ and let X be a k-step compact nilspace endowed with a $\mathrm{Z}_k(\mathrm{X})$-invariant metric. There exists $b = b(\delta, C, \mathrm{X}) > 0$ such that if $\phi : \mathrm{Z} \to \mathrm{X}$ is a b-balanced morphism, then for every 1-bounded C-Lipschitz function $F : \mathrm{X} \to \mathbb{C}$ and $\chi \neq \chi'$ in $\widehat{\mathrm{Z}_k(\mathrm{X})}$, the functions $F_\chi \circ \phi$, $F_{\chi'} \circ \phi$ are δ-quasiorthogonal.*

Proof. Let $\mathbf{F} : \mathrm{X} \to \mathbb{C}$ be the function $x \mapsto \overline{F_\chi(x)}F_{\chi'}(x)$ and note that $\langle F_{\chi'}, F_\chi\rangle = \mathbb{E}_{z\in\mathrm{Z}}\mathbf{F}\circ\phi(z)$. By Lemma 4.8, the functions F_χ and $F_{\chi'}$ are both C-Lipschitz and 1-bounded. Hence $\|\mathbf{F}\|_\infty \le 1$. For $x, y \in \mathrm{X}$, $|\overline{F_\chi(x)}F_{\chi'}(x) - \overline{F_\chi(y)}F_{\chi'}(y)| \le |\overline{F_\chi(x)}F_{\chi'}(x) - \overline{F_\chi(x)}F_{\chi'}(y)| + |\overline{F_\chi(x)}F_{\chi'}(y) - \overline{F_\chi(y)}F_{\chi'}(y)| \le 2\|F\|_\infty C d_{\mathrm{X}}(x, y)$, so $\|\mathbf{F}\|_{\mathrm{sum}} \le 1 + 2C$.

Let $\eta > 0$ be a parameter to be fixed later. By Corollary 4.18, we know that there exists $b = b(\eta, \mathrm{X}) > 0$ sufficiently small such that, if $\phi : \mathrm{Z} \to \mathrm{X}$ is b-balanced, then $\|\mu_\mathrm{Z}\circ\phi^{-1} - \mu_\mathrm{X}\|_{KR} \le \eta$. By definition of $\|\cdot\|_{KR}$, we then have

$$\left|\int_\mathrm{X} \mathbf{F}(x)\,\mathrm{d}\mu_\mathrm{X}(x) - \int_\mathrm{Z} \mathbf{F}\circ\phi(z)\,\mathrm{d}\mu_\mathrm{Z}(z)\right| \le \eta\|\mathbf{F}\|_{\mathrm{sum}} \le \delta$$

where η is chosen to be $\delta/(2C+1)$.

We now claim that $\int_\mathrm{X} \mathbf{F}(x)\,\mathrm{d}\mu_\mathrm{X}(x) = 0$. Confirming this will complete the proof. Since the Haar measure on X is preserved by the action of $\mathrm{Z}_k(\mathrm{X})$ (by [7, Lemma 2.2.6]), we have

$$\int_\mathrm{X} \mathbf{F}(x)\,\mathrm{d}\mu_\mathrm{X}(x) = \int_\mathrm{X}\int_{\mathrm{Z}_k(\mathrm{X})} \mathbf{F}(x+t)\,\mathrm{d}\mu_{\mathrm{Z}_k(\mathrm{X})}(t)\,\mathrm{d}\mu_\mathrm{X}(x).$$

By definition of $\mathbf{F}$, for every x, using that $\chi \neq \chi'$, we have

$$\int_{\mathrm{Z}_k(\mathrm{X})} \mathbf{F}(x+t)\,\mathrm{d}\mu_{\mathrm{Z}_k(\mathrm{X})}(t) = \mathbf{F}(x)\int_{\mathrm{Z}_k(\mathrm{X})} (\overline{\chi}\chi')(t)\,\mathrm{d}\mu_{\mathrm{Z}_k(\mathrm{X})}(t) = 0,$$

which concludes the proof. □

We strongly believe that the two main results of this subsection are in fact consequences of a single stronger result, which would constitute in various ways a more natural form of approximate orthogonality. As this is not needed in this paper, we leave it as the following question.

Question 4.25. Can the smallness of $\max_{\chi\neq\chi'}\{\langle F_\chi\circ\phi, F_{\chi'}\circ\phi\rangle_{U^{k+1}}\}$ in Proposition 4.22, and the smallness of $\max_{\chi\neq\chi'}\{\langle F_\chi\circ\phi, F_{\chi'}\circ\phi\rangle\}$ in Theorem 4.24, both be replaced with smallness of the larger quantity $\max_{\chi\neq\chi'}\sup_{h\in\mathrm{Z}}\|F_\chi\circ\phi\,\overline{T^h F_{\chi'}\circ\phi}\|_{U^k}$?

To motivate this question further, let us give some examples clarifying the relations between the three quantities in the question.

Example 4.26. Let us illustrate the aforementioned relations just in the case of the U^2-norm.

First we give examples of functions f, g showing that the product $\langle f, g\rangle_{U^3} = \mathbb{E}_h\|f\overline{T^h g}\|_{U^2}^4$ can be arbitrarily small while $\sup_{h\in\mathrm{Z}}\|f\overline{T^h g}\|_{U^2}$ is bounded away from 0. Let $f(x) = e(x^3/p)$ and

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

$g(x) = e((x^3 + bx^2)/p)$ for some $b \neq 0$ in $\mathbb{Z}_p$, with $p \geq 5$ prime. Then $\mathbb{E}_h \|f\overline{T^h g}\|_{U^2}^4 = o(1)_{p\to\infty}$, because $\|f\overline{T^h g}\|_{U^2} = o(1)_{p\to\infty}$ when $3h + b \neq 0 \mod p$. However, note that

$$\|f\overline{T^h g}\|_{U^2} = \|e(-(3h + b)x^2 - (3h^2 + 2bh)x - (h^3 + bh^2))/p)\|_{U^2},$$

so for $h = -3^{-1}b \mod p$, the quadratic term above cancels and we obtain

$$\|f\overline{T^h g}\|_{U^2} = \|e(-(3^{-1}b^2 - 3^{-1}2b^2)x - (h^3 + bh^2))/p)\|_{U^2} = \|e(3^{-1}b^2x)/p)\|_{U^2}$$

which equals 1 since $e(3^{-1}b^2x)/p)$ is a Fourier character.

Next let us give examples showing that $\langle f, g\rangle_{U^3}$ can be arbitrarily small while the usual inner product $\langle f, g\rangle$ is large (so that, in particular, the smallness of $\langle f, g\rangle_{U^3}$ obtained in Proposition 4.22 does not imply the quasiorthogonality obtained in Theorem 4.24). Taking simply $g = f = e(x^3/p)$, we have $\langle f, g\rangle_{U^3} = \|f\|_{U^3}^8 = o(1)_{p\to\infty}$, yet $\langle f, g\rangle = 1$.

4.4. k-step nilspace polynomials are structured functions of order k

For general theoretical reasons in the nilspace approach to higher-order Fourier analysis, and especially for our main proofs in Section 5, it is useful to establish that k-step nilspace polynomials of bounded complexity are structured functions of order k in the sense of Definition 2.23. We prove this with the following result, with the added strength that the L^2-error in the latter definition is here reduced to an L^∞-error.

Theorem 4.27. *Let $m \in \mathbb{N}$ and $C \geq 0$. Then for any $\delta > 0$ there exists a constant $N = N_{\delta,m,C} \geq 0$ such that the following holds. Let X be a k-step CFR nilspace of complexity at most m, let Z be a finite abelian group, let $\phi : \mathrm{Z} \to \mathrm{X}$ be a nilspace morphism, and let $F : \mathrm{X} \to \mathbb{C}$ be a Lipschitz function with $\|F\|_{sum} \leq C$. Then there is a function $h : \mathrm{Z} \to \mathbb{C}$ such that $\|F \circ \phi - h\|_\infty \leq \delta$ and $\|h\|_{U^{k+1}}^* \leq N$.*

To prove this, we shall use a quantitative form of the Stone-Weierstrass theorem for functions $F : \mathrm{X} \to \mathbb{C}$ such that the complexity of X and $\|F\|_{\mathrm{sum}}$ are bounded (see Proposition 4.29). As an ingredient we use the following result on $[0, 1/2]^n$ (equipped with the Euclidean metric).

Lemma 4.28. *Let $n \in \mathbb{N}$ and $C \geq 0$. Then, for any $\delta > 0$ there exists $N = N_{\delta,n,C} \geq 0$ such that the following holds. Let $g : [0, 1/2]^n \to \mathbb{C}$ be a Lipschitz function such that $\|g\|_{sum} \leq C$. Then, for every $i \in [N]$ and $j \in [n]$, there exist a complex number $\lambda_i = O_C(1)$ and a Lipschitz function $h_{i,j} : [0, 1/2] \to \mathbb{C}$ with $\|h_{i,j}\|_{sum} = O_C(1)$, and there exists $\mathcal{E} : [0, 1/2]^n \to \mathbb{C}$ with $\|\mathcal{E}\|_\infty \leq \delta$, such that $g(x_1, \dots, x_n) = \sum_{i=1}^N \lambda_i \prod_{j=1}^n h_{i,j}(x_j) + \mathcal{E}(x_1, \dots, x_n)$.*

Proof. We embed isometrically $[0, 1/2]^n$ into $\mathbb{T}^n$ by identifying $\mathbb{T}^n \cong [0, 1)^n$ (using the Euclidean metric on $[0, 1/2]^n$ and the flat metric on $\mathbb{T}^n$). By the Tietze extension theorem, we extend g to a continuous function $g' : \mathbb{T}^n \to \mathbb{C}$ with $\|g'\|_{\mathrm{sum}} = O(C)$. The partial Fourier series $S_M(g') := \sum_{r\in\mathbb{Z}^n : \|r\|_{\ell^\infty} \leq M} \widehat{g'}(\chi_r)\chi_r$ on $\mathbb{T}^n$ satisfy $\|S_M(g') - g'\|_\infty \leq K_{C,n}\frac{\log^n M}{M}$, where $K_{C,n}$ depends only on C and n (see [49, p. 201] for $n = 1$ and [21, p. 642] for $n > 1$). Letting M be large enough so that $K_{C,n}\frac{\log^n M}{M} \leq \delta$, we have $\|S_M(g') - g'\|_\infty \leq \delta$. Note that each of the terms $\widehat{g'}(\chi_r)\chi_r$ has the desired form. The result follows. □

Proposition 4.29. *Let X be a k-step CFR nilspace and let $C \geq 0$. Let $K : \mathrm{Cor}^{k+1}(\mathrm{X}) \to \mathrm{X}$ be the corner-completion function (see [6, Lemma 2.1.12]). Then, for any $\delta > 0$, there exists $N = N_{\delta,\mathrm{X},C} \geq 0$ such that the following holds. For any $F : \mathrm{X} \to \mathbb{C}$ with $\|F\|_{sum} \leq C$ there exist Lipschitz functions $(h_{i,v} : \mathrm{X} \to \mathbb{C})_{i\in[N], v\in[\![k+1]\!]\setminus\{0^{k+1}\}}$ with $\|h_{i,v}\|_{sum} = O_{\mathrm{X},C}(1)$ such that $\|F \circ K - \sum_{i\in[N]} \prod_{v\in[\![k+1]\!]\setminus\{0^{k+1}\}} h_{i,v}(x_v)\|_\infty \leq \delta$.*

Proof. By Lemma A.9 and the fact that X is compact, the function $F \circ K$ is $O_{C,\mathrm{X}}(1)$-Lipschitz. By the Tietze extension theorem, we can extend $F \circ K$ to a continuous function on $\mathrm{X}^{[\![k+1]\!]\setminus\{0^{k+1}\}}$, which we

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

denote by F', and for which we can assume that $\|F'\|_{\mathrm{sum}} = O_{C,\mathrm{X}}(1)$. As X is a compact manifold by Theorem A.2, we can find a finite C^∞ partition of unity $(\rho_i)_{i\in\mathcal{I}}$ such that, for all $i \in \mathcal{I}$, $\mathrm{Supp}(\rho_i) \subset U_i$ is compact and there exists a diffeomorphism $\phi_i : U_i \to (0,1/2)^\ell$ (for some $\ell = \ell_{\mathrm{X}}$). This then induces a partition of unity in $\mathrm{X}^{[\![k+1]\!]\setminus\{0^{k+1}\}}$ simply by considering products of different ρ_i, namely $\{\prod_{v\in[\![k+1]\!]\setminus\{0^{k+1}\}} \rho_{I_v} : I \in \mathcal{I}^{[\![k+1]\!]\setminus\{0^{k+1}\}}\}$. For any $I \in \mathcal{I}^{[\![k+1]\!]\setminus\{0^{k+1}\}}$, let ρ_I^* be the function mapping $\mathrm{c}^* \in \mathrm{X}^{[\![k+1]\!]\setminus\{0^{k+1}\}}$ to $\prod_{v\in[\![k+1]\!]\setminus\{0^{k+1}\}} \rho_{I_v}(\mathrm{c}^*(x_v))$.

Hence $F' = \sum_{I\in\mathcal{I}^{[\![k+1]\!]\setminus\{0^{k+1}\}}} F'\rho_I$. Note that $|\mathcal{I}^{[\![k+1]\!]\setminus\{0^{k+1}\}}| = O_{\mathrm{X}}(1)$ so if we prove the desired decomposition for each function $F'\rho_I$ then the result will follow. Note also that by construction $\|F'\rho_I\|_{\mathrm{sum}} = O_{\mathrm{X},C}(1)$. For $I \in \mathcal{I}^{[\![k+1]\!]\setminus\{0^{k+1}\}}$, let ϕ_I denote the map $\prod_{v\in[\![k+1]\!]\setminus\{0^{k+1}\}} \phi_{I_v} : \prod_{v\in[\![k+1]\!]\setminus\{0^{k+1}\}} U_{I_v} \to (0,1/2)^{\ell(2^{k+1}-1)}$. Then $F'\rho_I = (F'\rho_I)\circ\phi_I^{-1}\circ\phi_I$. Note that $\|\phi_I\|_{\mathrm{Lip}}$ and $\|\phi_I^{-1}\|_{\mathrm{Lip}}$ are both $O_{\mathrm{X}}(1)$, which implies that $\|(F'\rho_I)\circ\phi_I^{-1}\|_{\mathrm{sum}} = O_{C,\mathrm{X}}(1)$. Thus, we can apply Lemma 4.28 to $(F'\rho_I)\circ\phi_I^{-1}$ with $\delta_I = \delta/|\mathcal{I}^{[\![k+1]\!]\setminus\{0^{k+1}\}}|$. Hence $\big\|(F'\rho_I)\circ\phi_I^{-1} - \sum_{i\in[N]}\prod_{v\in[\![k+1]\!]\setminus\{0^{k+1}\}} h_{i,v}^I\big\|_\infty \le \delta_I$, where $h_{i,v}^I : (0,1/2)^\ell \to \mathbb{C}$.[17] The desired decomposition follows simply by considering $\prod_{v\in[\![k+1]\!]\setminus\{0^{k+1}\}} h_{i,v}^I\circ\phi_I$ for $i \in [N]$ and $I \in \mathcal{I}^{[\![k+1]\!]\setminus\{0^{k+1}\}}$. □

Proof of Theorem 4.27. Let $K_{k+1} := \{0,1\}^{k+1}\setminus\{0^{k+1}\}$. For every $x, t_1, \dots, t_{k+1} \in \mathrm{Z}$ we have

$$F\circ\phi(x) = F\circ K\big((\phi(x + v_1t_1 + \cdots + v_{k+1}t_{k+1}))_{v\in K_{k+1}}\big)$$

where $K : \mathrm{Cor}^{k+1}(\mathrm{X}) \to \mathrm{X}$ is the continuous corner completion function.

By Proposition 4.29 applied to F and K there exists $N = N_{\delta,m,C} \ge 0$ (recall that X has complexity at most m) and functions $(h_{i,v} : \mathrm{X} \to \mathbb{C})_{i\in[N],v\in K_{k+1}}$ with $\|h_{i,v}\|_{\mathrm{sum}} = O_{m,C}(1)$ such that $\|F\circ K - \sum_{i\in[N]}\prod_{v\in K_{k+1}} h_{i,v}(x_v)\|_\infty \le \delta$.

Let $h : \mathrm{Z} \to \mathbb{C}$, $x \mapsto \mathbb{E}_{t_1,\dots,t_{k+1}\in\mathrm{Z}}\sum_{i\in[N]}\prod_{v\in K_{k+1}} h_{i,v}(\phi(x + v\cdot t))$, where $t = (t_1,\dots,t_{k+1})$. We have for every $x \in \mathrm{Z}$,

$$\begin{aligned}|F\circ\phi(x) - h(x)| &= \big|\mathbb{E}_{t_1,\dots,t_{k+1}\in\mathrm{Z}}(F\circ K - \textstyle\sum_{i\in[N]}\prod_{v\in K_{k+1}} h_{i,v})\big((\phi(x+v\cdot t)_{v\in K_{k+1}}\big)\big| \\ &\le \big\|F\circ K - \textstyle\sum_{i\in[N]}\prod_{v\in K_{k+1}} h_{i,v})\big\|_\infty \;\le\; \delta.\end{aligned}$$

To bound $\|h\|_{U^{k+1}}^*$, let $g : \mathrm{Z} \to \mathbb{C}$ with $\|g\|_{U^{k+1}} \le 1$, and note that $\mathbb{E}_x h(x)\overline{g(x)}$ equals

$$\begin{aligned}&= \mathbb{E}_{x,t_1,\dots,t_{k+1}\in\mathrm{Z}}\,\overline{g(x)}\sum_{i\in[N]}\prod_{v\in K_{k+1}} h_{i,v}(\phi(x+v\cdot t)) = \sum_{i\in[N]}\langle g, (h_{i,v}\circ\phi)_{v\in K_{k+1}}\rangle_{U^{k+1}} \\ &\le \sum_{i\in[N]}\|g\|_{U^{k+1}}\prod_{v\in K_{k+1}}\|h_{i,v}\circ\phi\|_{U^{k+1}} \le \sum_{i\in[N]}\prod_{v\in K_{k+1}}\|h_{i,v}\|_\infty,\end{aligned}$$

where we used the Gowers-Cauchy-Schwarz inequality. Hence $\|h\|_{U^{k+1}}^* = O_{\delta,m,C}(1)$. □

5. Refined regularity in terms of nilspace characters

5.1. The refined regularity lemma

Recall that a complexity notion for compact nilspaces is a bijection between the countable set of isomorphism classes of CFR nilspaces and the natural numbers (see [12, Definition 1.2]). Throughout this section, we assume that we have fixed an arbitrary complexity notion, denoting the resulting list of nilspaces by $\{\mathrm{Y}_1, \mathrm{Y}_2, \dots\}$. For each Y_i in this fixed complexity notion, let $\mathrm{Q}_{\mathrm{Y}_i}(\eta, C)$ be the function given by Lemma 4.11. For each $m \in \mathbb{N}$ and $\eta > 0$, we then define

[17] Lemma 4.28 requires $(F'\rho_I)\circ\phi_I^{-1}$ to be defined and Lipschitz on $[0,1/2]^n$ and we have it defined on $(0,1/2)^n$ where $n = \ell(2^{k+1}-1)$. But since $\mathrm{Supp}(\rho_i) \subset U_i$ is compact, we can define $(F'\rho_I)\circ\phi_I^{-1}$ to be 0 on $[0,1/2]^n\setminus(0,1/2)^n$, and the resulting function is still Lipschitz as required, by [14, Theorem 2.5.6 and Remark 2.5.5].

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

$$\mathrm{Q}(\eta, m) := \max_{i \le m} \mathrm{Q}_{\mathrm{Y}_i}(\eta, m). \tag{5.1}$$

We can now prove the following regularity theorem, refining [12, Theorem 1.5] for finite abelian groups. (The proof could be extended to compact abelian groups as enabled by [12, Theorem 1.5], but we do not need this.)

Theorem 5.1 (Regularity). *Let $k \in \mathbb{N}$ and let $\mathcal{D} : \mathbb{R}_{>0} \times \mathbb{N} \to \mathbb{R}_{>0}$ be an arbitrary function. For every $\eta > 0$ there exists $N = N(\eta, \mathcal{D}) > 0$ such that the following holds. For every finite abelian group Z and every 1-bounded function $f : \mathrm{Z} \to \mathbb{C}$, for some $m \in [1, N]$ there exists a k-step CFR nilspace Y of complexity at most m, a $\mathcal{D}(\eta, m)$-balanced morphism $\phi : \mathrm{Z} \to \mathrm{Y}$, and a 1-bounded m-Lipschitz function $F : \mathrm{Y} \to \mathbb{C}$, such that*

$$f = f_s + f_e + f_r, \tag{5.2}$$

where $f_s = F \circ \phi$, $\|f_e\|_1 \le \eta$, $\|f_r\|_\infty \le 1$, and $\max(\|f_r\|_{U^{k+1}}, |\langle f_r, f_s\rangle|, |\langle f_r, f_e\rangle|) \le \mathcal{D}(\eta, m)$.

Moreover, there is a set $S \subset \widehat{\mathrm{Z}_k(\mathrm{Y})}$ with $|S| \le \mathrm{Q}(\eta, m)$, such that, letting $g_\chi := F_\chi \circ \phi$ for each $\chi \in S$, we have

$$f_s = \textstyle\sum_{\chi \in S} g_\chi \; + \; g_e, \tag{5.3}$$

where $\|g_e\|_\infty \le \eta$, and for all $\chi \ne \chi'$ in S, we have $\max\left(|\langle g_\chi, g_{\chi'}\rangle|, \langle g_\chi, g_{\chi'}\rangle_{U^{k+1}}\right) \le \mathcal{D}(\eta, m)$.

Proof. Having fixed $\mathcal{D}$ and η, we apply [12, Theorem 1.5] with a function $\mathcal{D}'$ for which we shall specify constraints throughout the proof. Firstly we require that $\mathcal{D}'(\eta, m) \le \mathcal{D}(\eta, m)$. For this function $\mathcal{D}'$ (to be fully specified later), let $f = f_s + f_e + f_r$ be the decomposition provided by [12, Theorem 1.5]. It then remains to decompose $f_s = F \circ \phi$ further as claimed in (5.3). By [12, Theorem 1.5] we know that $F : \mathrm{Y} \to \mathbb{C}$ is m-Lipschitz, 1-bounded, and the complexity of Y is at most m. Also, the morphism $\phi : \mathrm{Z} \to \mathrm{Y}$ is $\mathcal{D}'(\eta, m)$-balanced.

As Y has complexity at most m, Lemma 4.11 applied with $\delta = \eta$ yields a set $S \subset \widehat{\mathrm{Z}_k(\mathrm{Y})}$ with $|S| \le \mathrm{Q}(\eta, m)$ such that $F = \sum_{\chi \in S} F_\chi + F_e$ where $\|F_e\|_\infty \le \eta$. Letting $g_e = F_e \circ \phi$, we obtain (5.3). By Lemma 4.8, we have $\|F_\chi\|_L \le \|F\|_L$ and $\|F_\chi\|_\infty \le \|F\|_\infty$, so $\|F_\chi\|_{\mathrm{sum}} \le m + 1$.

Finally, we apply Proposition 4.22 and Theorem 4.24 with $\delta = \mathcal{D}(\eta, m)$, thus obtaining that there exists $b = b(\mathcal{D}(\eta, m), m, \mathrm{Y}) > 0$ such that if ϕ is b-balanced, then the conclusions of Proposition 4.22 and Theorem 4.24 hold, giving us the claimed smallness of U^{k+1}-products and quasiorthogonality, namely $\max\left(|\langle g_\chi, g_{\chi'}\rangle|, \langle g_\chi, g_{\chi'}\rangle_{U^{k+1}}\right) \le \mathcal{D}(\eta, m)$ for all $\chi \ne \chi'$ in S. Since Y has complexity at most m, this function $b(\mathcal{D}(\eta, m), m, \mathrm{Y})$ can be bounded by a function $b'(\mathcal{D}(\eta, m), m)$. We require that $\mathcal{D}'(\eta, m) < b'(\mathcal{D}(\eta, m), m)$. □

5.2. *Correlation, approximate diagonalization of the U^k norm, and regularization*

This subsection presents three main consequences of the refined regularity theorem that follow by elementary arguments once we have Theorem 5.1.

We begin with a simple refinement of the inverse theorem [12, Theorem 5.2] (in the case of finite abelian groups), which shows that a 1-bounded function $f : \mathrm{Z} \to \mathbb{C}$ with large U^{k+1}-norm correlates nontrivially not just with the structured part f_s (which is a nilspace polynomial in the sense of [12]) but rather with one of the nilspace characters in the decomposition of f_s.

Theorem 5.2 (Inverse theorem with nilspace characters). *Under the assumptions and notation of Theorem 5.1, suppose in addition that $\|f\|_{U^{k+1}} \ge \delta > 0$. Then*

$$\max_{\chi \in S} |\langle f, g_\chi\rangle| \ge \tfrac{\delta^{2^{k+1}}}{4|S|} = \Omega_\delta(1). \tag{5.4}$$

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Proof. Following the proof of [12, Theorem 5.2], we have that if η and $\mathcal{D}(\eta, m)$ are sufficiently small in terms of δ, then $|\langle f, f_s\rangle| \geq \delta^{2^{k+1}}/2$. Substituting here $f_s = \sum_{\chi\in S} g_\chi + g_e$, and noting that we can also take η to be at most $\delta^{2^{k+1}}/4$, we see that there exists $\chi \in S$ such that $|\langle f, g_\chi\rangle| \geq \delta^{2^{k+1}}/(4|S|)$ where $|S| = O_\delta(1)$. □

We can now prove a new, more detailed correlation estimate, which gives information not just on the maximal correlation but also on the correlation of f with each nilspace character in (5.3) having nonnegligible L^2-norm. The information in question is that every such inner product is close to the squared L^2-norm of the nilspace character.

Lemma 5.3. *Under the assumptions and notation of Theorem 5.1, for each $\chi \in S$ we have*

$$\big|\langle f, g_\chi\rangle - \|g_\chi\|_2^2\big| \leq \mathcal{D}(\eta, m)\big(|S| + O_{\eta,m}(1)\big) + 3\eta. \tag{5.5}$$

Proof. We have $\langle f, g_\chi\rangle = \sum_{\chi'\in S}\langle g_{\chi'}, g_\chi\rangle + \langle g_e + f_e, g_\chi\rangle + \langle f_r, g_\chi\rangle$. Hence

$$\big|\langle f, g_\chi\rangle - \|g_\chi\|_2^2\big| \leq \sum_{\chi'\in S\setminus\{\chi\}} |\langle g_{\chi'}, g_\chi\rangle| + |\langle g_e + f_e, g_\chi\rangle| + |\langle f_r, g_\chi\rangle|.$$

Using the quasiorthogonality, the sum over $S \setminus \{\chi\}$ here is at most $|S|\mathcal{D}(\eta, m)$.

Next, using that $\|g_\chi\|_\infty \leq 1$, we have $|\langle g_e + f_e, g_\chi\rangle| \leq \|g_e\|_\infty + \|f_e\|_1 \leq 2\eta$.

Finally, we bound $|\langle f_r, g_\chi\rangle|$. We have $g_\chi = F_\chi \circ \phi$ where, as explained in the proof of Theorem 5.1, we have $\|F_\chi\|_{\text{sum}} \leq m + 1$. By Theorem 4.27 with $\delta = \eta$, there is $h' : \mathrm{Z} \to \mathbb{C}$ with $\|h'\|^*_{U^{k+1}} = O_{\eta,m}(1)$ such that $\|g_\chi - h'\|_\infty \leq \eta$. Hence g_χ is an (R, η)-structured function of order k with $R = O_{\eta,m}(1)$. Since f_r is 1-bounded and $\|f_r\|_{U^{k+1}} \leq \mathcal{D}(\eta, m)$, by Proposition 2.27 with $h = f_r$, we have $|\langle f_r, g_\chi\rangle| \leq \mathcal{D}(\eta, m)O_{\eta,m}(1) + \eta$. The result follows. □

The next main result in this subsection is an approximate diagonalization of the U^{k+1}-norm. To motivate this, recall the well-known formula for the U^2-norm in terms of Fourier coefficients:

$$\|f\|_{U^2}^4 = \sum_{\chi\in\widehat{\mathrm{Z}}} |\widehat{f}(\chi)|^4. \tag{5.6}$$

This formula can be viewed as a diagonalization of the U^2-norm in the sense that when we write $\|f\|_{U^2}^4$ as $\langle f, f, f, f\rangle_{U^2}$, substitute here f by its Fourier expansion $f = \sum_{\chi\in\widehat{\mathrm{Z}}} \widehat{f}(\chi)\chi$, and expand the U^2-product, the orthonormality of characters yields cancellation of all nondiagonal terms in the expansion, resulting in the right side of (5.6). Note that this right side can be rewritten as $\sum_{\chi\in\widehat{\mathrm{Z}}} \|\widehat{f}(\chi)\chi\|_{U^2}^4$.

As a consequence of Theorem 5.1, we can obtain the following generalization of (5.6). This type of result will help in the final proof of this section, but it is also of independent interest. Indeed, formulas of this type, which can be viewed as approximate "Pythagorean theorems" for the Gowers norms, have been sought as a desirable feature of higher-order Fourier analysis; see for instance a special case of such a result for the U^3-norm in [29, Theorem 4.1].

Theorem 5.4 (Approximate diagonalization of $\|\cdot\|_{U^{k+1}}$). *Let $k \in \mathbb{N}$, let $\mathcal{D} : \mathbb{R}_{>0} \times \mathbb{N} \to \mathbb{R}_{>0}$ be an arbitrary function, and let $\eta > 0$. Let Z be a finite abelian group and let $f : \mathrm{Z} \to \mathbb{C}$ be a 1-bounded function. Let $f = \sum_{\chi\in S} g_\chi + g_e + f_r$ and $m \leq N(\eta, \mathcal{D})$ be the decomposition and complexity bound given by Theorem 5.1. Then*

$$\|f\|_{U^{k+1}}^{2^{k+1}} = \sum_{\chi\in S} \|g_\chi\|_{U^{k+1}}^{2^{k+1}} + \mathcal{E}, \tag{5.7}$$

where $|\mathcal{E}| \leq 2^{2^{k+1}}\big(4\eta^{\frac{k+2}{2^{k+1}}} + \mathcal{D}(\eta, m)\big) + (|S|^{2^{k+1}} - 1)\mathcal{D}(\eta, m)^{1/2^k}$.

We extract the final part of the proof of this theorem as the following separate lemma, which will be used in other results below. This lemma gives more properly an approximate "Pythagorean formula"

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

for Gowers norms (as it holds not just for the nilspace characters in Theorem 5.4, but for any system of 1-bounded functions with pairwise small Gowers inner-products).

Lemma 5.5 (Approximate Pythagorean formula for $\|\cdot\|_{U^s}$). *Let $s \geq 2$ be an integer, let G be a compact abelian group, let B be a finite set, and let $(g_i)_{i\in B}$ be a collection of 1-bounded Borel functions on G such that for any $i \neq j$ in B we have $\langle g_i, g_j\rangle_{U^s} \leq \delta$. Then*

$$\Big| \Big\| \sum_{i\in B} g_i \Big\|_{U^s}^{2^s} - \sum_{i\in B} \|g_i\|_{U^s}^{2^s} \Big| \leq (|B|^{2^s} - 1)\delta^{2^{1-s}}. \tag{5.8}$$

Proof. We expand $\| \sum_{i\in B} g_i\|_{U^s}^{2^s}$ into $|B|^{2^s}$ Gowers U^s-products, obtaining

$$\Big\| \sum_{i\in B} g_i \Big\|_{U^s}^{2^s} = \sum_{(i_v)_{v\in[\![s]\!]}\in B^{[\![s]\!]}} \langle (g_{i_v})_{v\in[\![s]\!]}\rangle_{U^s}.$$

In the right side here, the summands for which the corresponding vector $(i_v)_{v\in[\![s]\!]}$ is constant add up to $\sum_{i\in B}\|g_i\|_{U^s}^{2^s}$. The main task is thus to show that every other term is small. Every such term is a product $\langle (g_{i_v})_{v\in[\![s]\!]}\rangle_{U^s}$ such that for some pair of vertices $v, w \in [\![s]\!]$ we have $i_v \neq i_w$. We claim that we can then assume additionally that these vertices v, w are *adjacent* in the discrete cube $[\![s]\!]$ (i.e., only one of their coordinates differs); this follows from the fact that starting from any vertex of $[\![s]\!]$ there is a Hamiltonian path in the graph of 1-faces (or "edges") of the discrete cube $[\![s]\!]$, so if our claim was false then starting from any vertex and following such a path, we would deduce that $(i_v)_{v\in[\![s]\!]}$ is constant, a contradiction. Assuming thus that v, w are adjacent with $i_v \neq i_w$, now note that by $s-1$ applications of the Cauchy-Schwarz inequality, similar to the applications involved in the proof of the U^s-Gowers-Cauchy-Schwarz inequality (see [66]), and using that $\|g_i\|_{U^s} \leq 1$ for every $i \in B$, we obtain that the U^s-product in question is at most $\langle g_{i_v}, g_{i_w}\rangle_{U^s}^{1/2^{s-1}} \leq \delta^{1/2^{s-1}}$. The result follows. □

We shall also use the following upper bound for the Gowers norms.

Lemma 5.6. *Let $s \geq 2$ be an integer, let G be a compact abelian group, and let $f : G \to \mathbb{C}$ be a bounded Borel function. Then*

$$\|f\|_{U^s}^{2^s} \leq \min\big(\|f\|_2^{2s+2}\|f\|_{L^\infty}^{2^s-2s-2},\ \|f\|_1^{s+1}\|f\|_{L^\infty}^{2^s-s-1}\big). \tag{5.9}$$

Proof. To see that $\|f\|_{U^s}^{2^s} \leq \|f\|_2^{2s+2}\|f\|_{L^\infty}^{2^s-2s-2}$, note that the graph of 1-faces on the discrete cube $[\![s]\!]$ contains two disjoint stars S_1, S_2 of order $s+1$, namely the star with center 0^s and the star with center 1^s (note that for $s = 3$ these stars partition $[\![s]\!]$, and for $s > 3$ their union leaves $2^s - 2s - 2$ vertices in the complement). Using the Cauchy-Schwarz inequality, we then have

$$\begin{aligned}\|f\|_{U^s}^{2^s} \leq& \|f\|_{L^\infty}^{2^s-2s-2} \int_{\mathrm{c}\in\mathrm{C}^s(G)} \prod_{v\in S_1} |f(\mathrm{c}(v))| \prod_{w\in S_2} |f(\mathrm{c}(w))| \\ \leq& \|f\|_{L^\infty}^{2^s-2s-2} \Big(\int_{\mathrm{c}\in\mathrm{C}^s(G)} \prod_{v\in S_1} |f(\mathrm{c}(v))|^2\Big)^{1/2} \Big(\int_{\mathrm{c}\in\mathrm{C}^s(G)} \prod_{w\in S_2} |f(\mathrm{c}(w))|^2\Big)^{1/2}.\end{aligned}$$

The variables in each star S_i can be changed into $s+1$ independent variables $x_1, \dots, x_{s+1} \in G$, and it follows that the last two integrals both equal $\|f\|_2^{2s+2}$.

The proof of the other inequality $\|f\|_{U^s}^{2^s} \leq \|f\|_1^{s+1}\|f\|_{L^\infty}^{2^s-s-1}$ is similar. □

Proof of Theorem 5.4. Letting $f_c = \sum_{\chi\in S} F_\chi \circ \phi$, we have that $\big|\|f\|_{U^{k+1}} - \|f_c\|_{U^{k+1}}\big|$ is at most

$$\|g_e + f_e + f_r\|_{U^{k+1}} \leq \|g_e\|_\infty + 3^{1-\frac{k+2}{2^{k+1}}}\|f_e\|_{L^1}^{\frac{k+2}{2^{k+1}}} + \|f_r\|_{U^{k+1}} \leq \eta + 3^{1-\frac{k+2}{2^{k+1}}}\eta^{\frac{k+2}{2^{k+1}}} + \mathcal{D}(\eta, m),$$

where we used that $\|f_e\|_\infty \leq 3$ and thus $\|f_e\|_{U^{k+1}}^{2^{k+1}} \leq 3^{2^{k+1}-(k+2)}\|f_e\|_1^{k+2}$ (by Lemma 5.6). Hence, using $\|f\|_{U^{k+1}} \leq \|f\|_{L^\infty} \leq 1$ and $\|f_c\|_{U^{k+1}} \leq \|f_c\|_{L^\infty} \leq 1+\eta$, we obtain that $|\|f\|_{U^{k+1}}^{2^{k+1}} - \|f_c\|_{U^{k+1}}^{2^{k+1}}| \leq$

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

$\big|\|f\|_{U^{k+1}}-\|f_c\|_{U^{k+1}}\big|\,\big|\|f\|_{U^{k+1}}^{2^{k+1}-1}+\|f\|_{U^{k+1}}^{2^{k+1}-2}\|f_c\|_{U^{k+1}}+\cdots+\|f_c\|_{U^{k+1}}^{2^{k+1}-1}\big) \le (\eta+3^{1-\frac{k+2}{2^{k+1}}}\eta^{\frac{k+2}{2^{k+1}}}+\mathcal{D}(\eta,m))\,\big(1+(1+\eta)+(1+\eta)^2+\cdots+(1+\eta)^{2^{k+1}-1}\big)$, which is at most $2^{2^{k+1}}(\eta+3^{1-\frac{k+2}{2^{k+1}}}\eta^{\frac{k+2}{2^{k+1}}}+\mathcal{D}(\eta,m))$.

Hence $\big|\|f\|_{U^{k+1}}^{2^{k+1}}-\|f_c\|_{U^{k+1}}^{2^{k+1}}\big| \le 2^{2^{k+1}}(4\eta^{\frac{k+2}{2^{k+1}}}+\mathcal{D}(\eta,m))$. Combining this with Lemma 5.5 applied to f_c, the result follows. □

In a similar spirit, we can also prove the following approximate Parseval identity, to be used in the proof of Proposition 5.8 below. Higher-order analogues of Parseval's identity have also been sought as desirable features of higher-order Fourier analysis (see [25, Section 16]).

Theorem 5.7 (Approximate higher-order Parseval identity). *Let $k \in \mathbb{N}$, let $\mathcal{D}:\mathbb{R}_{>0}\times\mathbb{N}\to\mathbb{R}_{>0}$ be an arbitrary function, and let $\eta>0$. Let Z be a finite abelian group, and let $f:\mathrm{Z}\to\mathbb{C}$ be a 1-bounded function. Let $f=\sum_{\chi\in S}g_\chi+g_e+f_e+f_r$ and $m\le N(\eta,\mathcal{D})$ be the decomposition and complexity bound given by Theorem 5.1. Then*

$$\Big|\|f\|_2^2-\sum_{\chi\in S}\|g_\chi\|_2^2-\|f_r\|_2^2\Big| \le (|S|^2+4)\mathcal{D}(\eta,m)+8\eta. \tag{5.10}$$

Proof. Letting again $f_c=\sum_{\chi\in S}g_\chi$, and recalling that $f_s=f_c+g_e$, we have, using the estimates given in Theorem 5.1, that

$$|\|f\|_2^2-\|f_s\|_2^2| \le 2|\langle f_s,f_r\rangle|+2|\langle f_e,f_r\rangle|+2|\langle f_s,f_e\rangle|+\|f_e\|_2^2+\|f_r\|_2^2 \le 4\mathcal{D}(\eta,m)+5\eta+\|f_r\|_2^2.$$

As $\|f_s\|_\infty\le 1$ and $\|g_e\|_\infty\le\eta$, we have $|\|f_s\|_2^2-\|f_c\|_2^2|\le 2|\langle f_c,g_e\rangle|+\|g_e\|_2^2\le 2\eta+\eta^2\le 3\eta$, and so $|\|f\|_2^2-\|f_c\|_2^2|\le|\|f\|_2^2-\|f_s\|_2^2|+|\|f_s\|_2^2-\|f_c\|_2^2|\le 4\mathcal{D}(\eta,m)+8\eta+\|f_r\|_2^2$. We finish by using that $|\|f_c\|_2^2-\sum_{\chi\in S}\|g_\chi\|_2^2|\le|S|^2\max_{\chi\neq\chi'\in S}|\langle g_\chi,g_{\chi'}\rangle|\le|S|^2\mathcal{D}(\eta,m)$. □

The final result of this subsection is a key ingredient for the proof of our main result Theorem 1.1, to be completed in the last section (see Theorem 7.12).

Proposition 5.8. *Let $k\in\mathbb{N}$, let $\mathcal{D}:\mathbb{R}_{>0}\times\mathbb{N}\to\mathbb{R}_{>0}$ be an arbitrary function, and let $\eta>0$. Let Z be a finite abelian group, and let $f:\mathrm{Z}\to\mathbb{C}$ be a 1-bounded function. Let*

$$f=\textstyle\sum_{\chi\in S}g_\chi+g_e+f_e+f_r$$

be the decomposition given by Theorem 5.1, with associated complexity bound $m\le N(\eta,\mathcal{D})$. For any $\rho>0$, let $S_\rho:=\{\chi\in S:\|g_\chi\|_2^2\ge\rho\}$. Then $|S_\rho|\le\rho^{-1}(1+c_0)$ and

$$\big\|f-\textstyle\sum_{\chi\in S_\rho}g_\chi\big\|_{U^{k+1}} \le \rho^{(k+1)/2^{k+1}}(1+c_0)^{1/2^{k+1}}+2|S|\mathcal{D}(\eta,m)^{1/4^{k+1}}+4\eta^{(k+2)/2^{k+1}}, \tag{5.11}$$

where $c_0=(|S|^2+4)\mathcal{D}(\eta,m)+8\eta$. Moreover, for any $\sigma>0$, there exists $h:\mathrm{Z}\to\mathbb{C}$ with $\|h\|_{U^{k+1}}^*=O_{\sigma,\rho,\eta,m}(1)$ and $\|h-\sum_{\chi\in S_\rho}g_\chi\|_\infty\le\sigma$.

The last sentence of Proposition 5.8 tells us that $\sum_{\chi\in S_\rho}g_\chi$ is a valid k-th order structured part of f, since it is an (R,δ)-structured function of order k, with the additional strength that the error is small uniformly (not just in L^2 as in Definition 2.23).

Proof. By (5.10), and using that $\|f_r\|_2\ge 0$, we have $\sum_{\chi\in S}\|g_\chi\|_2^2\le\|f\|_2^2+c_0\le 1+c_0$ where $c_0=(|S|^2+4)\mathcal{D}(\eta,m)+8\eta$, so $|S_\rho|\le\rho^{-1}(1+c_0)$. Using that $\|g_e\|_\infty\le\eta$, $\|f_e\|_\infty\le 3$, $\|f_e\|_1\le\eta$, and (5.9) applied with $s=k+1$, we have

$$\begin{aligned}\big|\big\|f-\textstyle\sum_{\chi\in S_\rho}g_\chi\big\|_{U^{k+1}}-\big\|\sum_{\chi\in S\setminus S_\rho}g_\chi\big\|_{U^{k+1}}\big| &\le \|g_e\|_{U^{k+1}}+\|f_e\|_{U^{k+1}}+\|f_r\|_{U^{k+1}}\\ &\le \|g_e\|_\infty+(\|f_e\|_\infty^{2^{k+1}-k-2}\|f_e\|_1^{k+2})^{1/2^{k+1}}+\|f_r\|_{U^{k+1}} \le 4\eta^{(k+2)/2^{k+1}}+\mathcal{D}(\eta,m)=:c_1.\end{aligned}$$

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Hence $\|f - \sum_{\chi\in S_\rho} g_\chi\|_{U^{k+1}} \le \|\sum_{\chi\in S\setminus S_\rho} g_\chi\|_{U^{k+1}} + c_1$. By (5.8) applied with $B = S\setminus S_\rho$, we have $\|\sum_{\chi\in S\setminus S_\rho} g_\chi\|_{U^{k+1}}^{2^{k+1}} \le \sum_{\chi\in S\setminus S_\rho} \|g_\chi\|_{U^{k+1}}^{2^{k+1}} + (|S\setminus S_\rho|^{2^{k+1}} - 1)\mathcal{D}(\eta,m)^{1/2^k}$. Moreover, by (5.9) and the fact that each g_χ is 1-bounded, we have

$$\sum_{\chi\in S\setminus S_\rho} \|g_\chi\|_{U^{k+1}}^{2^{k+1}} \le \sum_{\chi\in S\setminus S_\rho} \|g_\chi\|_2^{2k+4} < \rho^{k+1} \sum_{\chi\in S\setminus S_\rho} \|g_\chi\|_2^2 \le \rho^{k+1}(1+c_0).$$

Putting all this together, we conclude that

$$\begin{aligned}\big\|f - \textstyle\sum_{\chi\in S_\rho} g_\chi\big\|_{U^{k+1}} &\le \big(\rho^{k+1}(1+c_0) + |S|^{2^{k+1}}\mathcal{D}(\eta,m)^{1/2^k}\big)^{1/2^{k+1}} + c_1\\ &\le \rho^{(k+1)/2^{k+1}}(1+c_0)^{1/2^{k+1}} + |S|\mathcal{D}(\eta,m)^{1/2^{k+1}} + c_1,\end{aligned}$$

where in the last inequality we used that $\|\cdot\|_{\ell^{2^{k+1}}} \le \|\cdot\|_{\ell^1}$.

To prove the last claim, we apply Theorem 4.27. Recall that for each $\chi \in S_\rho$ we have $g_\chi = F_\chi \circ \phi$ with $\|F_\chi\|_{\mathrm{sum}} \le m+1$. Thus $\sum_{\chi\in S_\rho} g_\chi = F_\rho \circ \phi$ where $F_\rho := \sum_{\chi\in S_\rho} F_\chi : \mathrm{X} \to \mathbb{C}$ satisfies $\|F_\rho\|_{\mathrm{sum}} \le |S_\rho|(m+1)$. It then follows from Theorem 4.27, applied with $\delta = \sigma$, that there is a function $h : \mathrm{Z} \to \mathbb{C}$ with $\|F_\rho \circ \phi - h\|_\infty \le \sigma$ and $\|h\|_{U^{k+1}}^* = O_{\sigma,\rho,\eta,m}(1)$. □

6. A structure theorem for $\mathcal{K}_\varepsilon(f\otimes\overline{f})$ using 2-step nilspace characters

In this section, we restrict the general k-th order treatment of Section 5 and focus on the quadratic case (second order) to obtain a key ingredient for our algorithmic applications: Theorem 6.1 below. The idea is to apply the Fourier denoising operator K_ε (recall Definition 2.39) to the decomposition resulting from our upgraded regularity lemma (Theorem 5.1) in the case $k = 2$. The resulting structure theorem for $\mathcal{K}_\varepsilon(f\otimes\overline{f})$ decomposes this matrix into a sum of a bounded number of rank-1 matrices corresponding to nonnegligible 2-step nilspace characters.

Theorem 6.1. *For any function $\mathcal{D} : \mathbb{R}_{>0}\times\mathbb{N} \to \mathbb{R}_{>0}$ and any $\eta \in (0,1)$, there exists $M = M(\eta,\mathcal{D}) > 0$ such that the following holds. Let Z be a finite abelian group, and let $f : \mathrm{Z} \to \mathbb{C}$ be a 1-bounded function. Then, for some $m \le M$, there is a 2-step CFR nilspace X of complexity at most m, a $\mathcal{D}(\eta,m)$-balanced morphism $\phi : \mathrm{Z} \to \mathrm{X}$, a 1-bounded m-Lipschitz function $F : \mathrm{X} \to \mathbb{C}$, and a set $S \subset \widehat{\mathrm{Z}_k(\mathrm{X})}$ with $|S| \le \mathrm{Q}(\eta,m)$, such that for every $\varepsilon > 0$ we have*

$$\mathcal{K}_\varepsilon(f\otimes\overline{f}) = \sum_{\chi\in S} g_\chi \otimes \overline{g_\chi} + E, \tag{6.1}$$

where for every $\chi \in S$ we have $g_\chi := F_\chi \circ \phi$ with $\|g_\chi\|_2 \ge \eta^2/\mathrm{Q}(\eta,m)$, and $E \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ satisfies $\|E\|_2 \le 21\sqrt{\eta} + 20\mathcal{D}(\eta,m)\varepsilon^{-1/2}\,\mathrm{Q}(\eta,m)^{5/2} + O_{m,\eta}(\varepsilon^{1/4})$.

Remark 6.2. Note that we can ensure that $\|E\|_2 \le 23\sqrt{\eta}$ by letting $\varepsilon = \varepsilon_{\eta,m} > 0$ be such that $O_{m,\eta}(\varepsilon^{1/4}) \le \sqrt{\eta}$ and then setting $\mathcal{D}(\eta,m)$ so that $20\mathcal{D}(\eta,m)\varepsilon^{-1/2}\,\mathrm{Q}(\eta,m)^{5/2} \le \sqrt{\eta}$.

Proof. We apply Theorem 5.1 with η and $\mathcal{D}$. This yields that for some $M = M(\eta,\mathcal{D}) > 0$, there exists $m \le M$ such that we have the decomposition $f = \sum_{\chi\in S} g_\chi + g_e + f_e + f_r$, with the properties specified in Theorem 5.1. Relabeling this set S (resulting from Theorem 5.1) as S', we now restrict the last sum to the following subset of S', to ensure that each $\|g_\chi\|_2$ is large:

$$S := \{\chi \in S' : \|g_\chi\|_2 \ge \eta^2/|S'|\}.$$

Note that $\|\sum_{\chi\in S'\setminus S} g_\chi\|_2 \le |S'\setminus S|\frac{\eta^2}{|S'|} \le \eta^2$. Therefore, letting $f_2 := \sum_{\chi\in S'\setminus S} g_\chi + g_e + f_e$, we have $f = \sum_{\chi\in S} g_\chi + f_2 + f_r$, where $\|f_2\|_2 \le \eta^2 + \eta + \sqrt{3\eta} \le 4\sqrt{\eta}$ (where we used that $\eta \in (0,1)$) and $\|f_r\|_{U^3} \le \mathcal{D}(\eta,m)$. Note that since $|S'| \le \mathrm{Q}(\eta,m)$, for every $\chi \in S$ we have $\|g_\chi\|_2 \ge \eta^2/\mathrm{Q}(\eta,m)$.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Now the result will follow from Proposition 3.17 by setting the parameters α_i in that proposition according to the above data. More precisely, in our present application of Proposition 3.17, we have that condition (i) in that proposition holds with $\alpha_1 = 4\sqrt{\eta}$, condition (ii) holds with $\alpha_2 = \mathcal{D}(\eta, m)$, condition (iii) holds with $\alpha_3 = 1$ and $n = |S|$ (letting the f_i in that proposition be the nilspace characters g_χ). Furthermore, condition (iv) holds with $\alpha_4 = \mathcal{D}(\eta, m)$ (by the last sentence of Theorem 5.1). Finally, by Theorem 4.19 we have that each g_χ is a quadratic character on Z with parameters (R, σ) where $R = O_m(\sigma^{-O_m(1)})$, so by Proposition 2.52, we have $\|\mathcal{K}_\varepsilon(g_\chi \otimes \overline{g_\chi}) - g_\chi \otimes \overline{g_\chi}\|_2 \leq 4\sigma + O_m(\varepsilon^{1/4}\sigma^{-O_m(1)})$. Setting $\sigma = \eta/(4\,\mathrm{Q}(\eta, m))$, we have that condition (v) in Proposition 3.17 holds with $\alpha_5 = \eta/\mathrm{Q}(\eta, m) + O_{m,\eta}(\varepsilon^{1/4})$.

Substituting these quantities into the bound $5\alpha_1 + 16\varepsilon^{-1/2}\alpha_2 + \varepsilon^{-1/2}n^2(n^{1/2}\alpha_3 + 3)\alpha_4 + n\alpha_5$ from Proposition 3.17, and using that $|S| \leq \mathrm{Q}(\eta, m)$, the result follows. □

7. Algorithmic consequences of the structure theorem for $\mathcal{K}_\varepsilon(f \otimes \overline{f})$

In this section we use the notions of pseudoeigenvalues and pseudoeigenvectors, well-known in matrix analysis (see [69, p. 16]). However, we adapt these notions to Z-matrices as follows, coherently with the normalization used in this paper (recall Definition 2.3). We shall use these notions only in this normalized form throughout this section.

Definition 7.1. Let Z be a finite abelian group, let $M \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ be a Z-matrix, and let $\beta > 0$. A *β-pseudoeigenvalue* of M is a number $\lambda \in \mathbb{C}$ such that for some $v \in \mathbb{C}^{\mathrm{Z}}$ with $\|v\|_2 = 1$ we have $\|Mv - \lambda v\|_2 < \beta$. We then say that v is a *β-pseudoeigenvector* corresponding to λ, or *(λ, β)-pseudoeigenvector*. The set of β-pseudoeigenvalues of M is denoted by $\sigma_\beta(M)$.

Given a self-adjoint Z-matrix $M \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$, we denote by $\mathrm{Spec}(M)$ the spectrum of M, that is, the multiset of (real) eigenvalues of M, where the multiplicity of each eigenvalue equals the dimension of its eigenspace. For any $\rho > 0$ we define the multiset

$$\mathrm{Spec}_\rho(M) := \{\lambda \in \mathrm{Spec}(M) : \lambda \geq \rho\}. \tag{7.1}$$

By Theorem 6.1, we know that the self-adjoint matrix $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ is (up to a small L^2-error) a sum of rank-1 matrices $g_\chi \otimes \overline{g_\chi}$ where the g_χ are quasiorthogonal nilspace characters with nonnegligible L^2-norm. Then every g_χ yields a pseudoeigenvector of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$, as follows.

Lemma 7.2. *Under the assumptions of Theorem 6.1, let $\mathcal{K}_\varepsilon(f \otimes \overline{f}) = \sum_{\chi\in S} g_\chi \otimes \overline{g_\chi} + E$ be the decomposition, obtained in (6.1), for the 1-bounded function $f : \mathrm{Z} \to \mathbb{C}$. Then for any $\chi \in S$, the function $u_\chi := g_\chi/\|g_\chi\|_2$ is a β-pseudoeigenvector of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ with pseudoeigenvalue $\lambda_\chi := \|g_\chi\|_2^2$, where $\beta = 21\sqrt{\eta} + \mathcal{D}(\eta, m)(20\varepsilon^{-1/2}\,\mathrm{Q}(\eta, m)^{5/2} + \mathrm{Q}(\eta, m)^2/\eta^2) + O_{m,\eta}(\varepsilon^{1/4})$.*

Remark 7.3. Note that, choosing $\varepsilon > 0$ sufficiently small so that the term $O_{m,\eta}(\varepsilon^{1/4})$ is at most $\eta^{1/2}$, we can then prescribe the function $\mathcal{D}$ to be sufficiently fast-decreasing so that $\beta = O(\eta^{1/2})$.

Proof. Recall from Theorem 6.1 that $\|g_\chi\|_2 \geq \eta^2/\mathrm{Q}(\eta, m)$ for each $\chi \in S$. Letting $M = \mathcal{K}_\varepsilon(f \otimes \overline{f})$, we have $Mu_\chi = \lambda_\chi u_\chi + \sum_{\chi'\in S\setminus\{\chi\}} \lambda_{\chi'} u_{\chi'} \langle u_{\chi'}, u_\chi\rangle + Eu_\chi$. Hence

$$\|Mu_\chi - \lambda_\chi u_\chi\|_2 \leq \sum_{\chi'\in S\setminus\{\chi\}} \|g_{\chi'}\|_2^2 \frac{|\langle g_{\chi'}, g_\chi\rangle|}{\|g_{\chi'}\|_2\|g_\chi\|_2} + \|Eu_\chi\|_2 \leq \sum_{\chi'\in S\setminus\{\chi\}} \frac{\|g_{\chi'}\|_2}{\|g_\chi\|_2} |\langle g_{\chi'}, g_\chi\rangle| + \|E\|_2.$$

Using that $|\langle g_{\chi'}, g_\chi\rangle| \leq \mathcal{D}(\eta, m)$ (by Theorem 5.1), we conclude that $\|Mu_\chi - \lambda_\chi u_\chi\|_2 \leq |S|\frac{\mathrm{Q}(\eta,m)}{\eta^2}\mathcal{D}(\eta, m) + \|E\|_2$, and the result follows. □

We now face the question of how to recover, algorithmically, the pseudo-eigenvectors (normalized nilspace characters) u_χ from the eigenvalues and eigenvectors of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Firstly, in Subsection 7.1 we obtain a decomposition of f into a quadratically structured part $\sum_{\chi\in S_\rho} g_\chi$ and a U^3-noise part, using the leading eigenvalues and eigenvectors of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$ (see Theorem 7.12). This will enable us to obtain such a decomposition algorithmically.

Secondly, in Subsection 7.2 we focus on this structured part $\sum_{\chi\in S_\rho} g_\chi$ to describe how one can recover the individual nilspace characters g_χ using dominant eigenvalues of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$.

Before we get started on the first task, let us record the following simple result, that will enable us to express the pseudoeigenvectors u_χ in terms of the eigenvectors of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$.

Lemma 7.4. *Let Z be a finite abelian group and let $M\in\mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ be a self-adjoint Z-matrix, with real eigenvalues λ_i for $i\in[|\mathrm{Z}|]$. Let λ be a β-pseudoeigenvalue of M and let u be a corresponding (λ,β)-pseudoeigenvector. Let $\{v_i : i\in[|\mathrm{Z}|]\}$ be an orthonormal basis of eigenvectors of M, let $u=\sum_{i\in[|\mathrm{Z}|]}\mu_i v_i$ be the expansion of u in this basis, and consider the eigenvalue cluster $C_\delta(\lambda):=\{i\in[|\mathrm{Z}|] : |\lambda_i-\lambda|\le\delta\}$. Then*

$$\left\|u-\sum_{i\in C_\delta(\lambda)}\mu_i v_i\right\|_2\le\beta/\delta. \tag{7.2}$$

Proof. Suppose that the eigenvalues λ_i of M satisfy $\lambda_i\ge\lambda_j$ if $i<j$. By assumption we have $Mu-\lambda u=\sum_{i=1}^{|\mathrm{Z}|}\mu_i(\lambda_i-\lambda)v_i$, so $\|Mu-\lambda u\|_2^2=\sum_{i=1}^{|\mathrm{Z}|}\mu_i^2(\lambda_i-\lambda)^2$. Since u is a β-pseudoeigenvector of M, the left side of the last equation is at most β^2, so $\sum_{i=1}^{|\mathrm{Z}|}\mu_i^2(\lambda_i-\lambda)^2\le\beta^2$. Let $I_\delta(\lambda):=\{i\in[|\mathrm{Z}|] : |\lambda_i-\lambda|>\delta\}$ be the complement of $C_\delta(\lambda)$. Then $\delta^2\sum_{i\in I_\delta}\mu_i^2\le\sum_{i=1}^{|\mathrm{Z}|}\mu_i^2(\lambda_i-\lambda)^2\le\beta^2$, whence $\left\|u-\sum_{i\in C_\delta(\lambda)}\mu_i v_i\right\|_2^2=\sum_{i\in I_\delta}\mu_i^2\le\beta^2/\delta^2$. □

7.1. Obtaining a quadratically structured part of f from the spectrum of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$

The case $k=2$ of Proposition 5.8 establishes the sum of 2-step nilspace characters $\sum_{\chi\in S_\rho} g_\chi$ as a valid quadratic (or second-order) structured part of f. The main result in this subsection, Theorem 7.12, tells us that this structured part can be recovered, up to a small error, as the projection of the function f to the linear span of the eigenvectors of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$ with large eigenvalues. This will prove Theorem 1.1, and thus the validity of our spectral algorithm to obtain the order-2 structured part of a bounded function (i.e., Algorithm 1).

To prove this result, we shall need to compare two orthogonal projections. On the one hand, the projection of f to the space spanned by its dominant 2-step nilspace characters g_χ (provided by Proposition 5.8). On the other hand, the projection of f to the linear span of the leading eigenvectors of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$. We are thus naturally led to standard tools in the theory of orthogonal projections in Hilbert space, which we now gather before proving the main result.

The first such tool is the following metric on the set of orthogonal projections.

Definition 7.5. Let W be a real or complex Hilbert space with inner product $\langle\cdot,\cdot\rangle$ and norm $\|v\|_W:=\langle v,v\rangle^{1/2}$. For any two orthogonal projections P,Q on W, we denote by $d(P,Q)$ the operator norm of $P-Q$, i.e.

$$d(P,Q):=\|P-Q\|=\sup_{v\in W,\|v\|_W\le 1}\|Pv-Qv\|_W. \tag{7.3}$$

We shall apply this in the Hilbert space $\mathbb{C}^{\mathrm{Z}}$ for a finite abelian group Z, equipped with the inner product and L^2-norm with the normalization used throughout this paper. Recall that any such orthogonal projection P is uniquely associated with the vector subspace that is its image, which we shall denote by U_P. Thus, the above metric gives a notion of distance between linear subspaces of W, which is well-known in the theory of Grassmannians (see, for instance, [55, (3)] or the book [2]). Recall the following basic fact (see [2, §39], where the metric is called the *aperture* of the subspaces U_P, U_Q).

Lemma 7.6. *For any pair of orthogonal projections P,Q on a Hilbert space W, we have $d(P,Q)\le 1$, and if $d(P,Q)<1$ then the dimensions of the subspaces U_P, U_Q are equal.*

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Proof. To see that $\|P-Q\| \le 1$, let us recall here a short standard argument: note that $P-Q = P(1-Q)-(1-P)Q$, so that for every v we have $Pv-Qv = P(1-Q)v-(1-P)Qv$, and since $P(1-Q)v$ and $(1-P)Qv$ are clearly orthogonal, it follows that $\|(P-Q)v\|_W^2 = \|P(1-Q)v\|_W^2 + \|(1-P)Qv\|_W^2 \le \|(1-Q)v\|_W^2 + \|Qv\|_W^2 = \|v\|_W^2$, as required.

The second claim in the proof is the main theorem in [2, §39]. □

The following alternative expression for $d(P,Q)$ will be useful (see [2, p. 111, equation (2)]):

$$d(P,Q) = \max\Big(\sup_{v\in U_P:\|v\|_W\le 1} \operatorname{dist}(v,U_Q), \sup_{v\in U_Q:\|v\|_W\le 1} \operatorname{dist}(v,U_P)\Big), \tag{7.4}$$

where $\operatorname{dist}(v,U_Q) := \|v-Q(v)\|_W$. We shall also use the following estimate.

Lemma 7.7. *Let P,Q be orthogonal projections on a Hilbert space W. Let $b_1,b_2,\dots,b_\ell$ be an orthonormal basis of U_P (the image of P). Then*

$$\sup_{v\in U_P:\|v\|_W\le 1} \operatorname{dist}(v,U_Q) \le \ell^{1/2} \max_{1\le i\le \ell} \|b_i - Q(b_i)\|_W. \tag{7.5}$$

Proof. Let $v = \sum_{i=1}^{\ell} \lambda_i b_i$ with $\|v\|_W \le 1$. By the Cauchy-Schwarz inequality, we have that $\|v-Q(v)\|_W$ is at most $\sum_{i=1}^{\ell} |\lambda_i| \|b_i - Q(b_i)\|_W \le \big(\sum_{i=1}^{\ell} |\lambda_i|\big) \max_{1\le i\le \ell} \|b_i - Q(b_i)\|_W \le \ell^{1/2} \max_{1\le i\le \ell} \|b_i - Q(b_i)\|_W$ (where we used that $\|v\|_W^2 = \sum_{i=1}^{\ell} |\lambda_i|^2$). □

Definition 7.8 (Projection to the span of spectrally dominant eigenvectors). Let T be a self-adjoint linear operator on a Hilbert space W, and let $\rho > 0$. We denote by $\operatorname{Eigen}_\rho(T)$ the subspace of W spanned by eigenvectors with corresponding eigenvalues in $\operatorname{Spec}_\rho(T)$. We denote by $\mathcal{P}_{T,\rho}$ the orthogonal projection to $\operatorname{Eigen}_\rho(T)$ in W.

Another ingredient is the following refinement of the Gram-Schmidt process.

Lemma 7.9 (Quantitative Gram-Schmidt process). *Let W be a Hilbert space. For each positive integer $s \ge 2$ there is a constant $C_s > 0$ such that the following holds. Let $u_1,\dots,u_s \in W$ such that $\|u_i\|_W = 1$ and such that for some $\tau \in (0,1)$ we have $|\langle u_i,u_j\rangle| \le \tau/C_s$ for all pairs $i \ne j$. Then there exist orthonormal vectors $w_1,\dots,w_s \in W$ such that for every $i \in [s]$ we have $\|u_i - w_i\|_2 \le C_s\tau$ and such that, for any $i \in [s]$, $\operatorname{Span}(w_1,\dots,w_i) = \operatorname{Span}(u_1,\dots,u_i)$.*

Estimates for C_s can be obtained from the proof, for example, $C_2 = 1$ and $C_s = 13\cdot 5^{s-3} - 1$ for $s \ge 3$.

Proof. We apply the Gram-Schmidt orthonormalization to the set $\{u_1,\dots,u_s\}$, and keep track of the changes performed by this process, to ensure the L^2-smallness of these changes.

Following the formulas used in the process (see for instance [41, §0.6.4]), we have $w_1 = u_1$, then $w_2' = u_2 - \langle u_2,w_1\rangle w_1$ and $w_2 := w_2'/\|w_2'\|_W$. Hence $\|u_2 - w_2'\|_W = |\langle u_2,w_1\rangle| \le \tau$. We can therefore let $C_2 = 1$.

Now, for general $s > 2$, the formulae are $w_s' = u_s - \sum_{j=1}^{s-1} \langle u_s,w_j\rangle w_j$ and $w_s = w_s'/\|w_s'\|_W$. For $j \in [s-1]$, we have $|\langle u_s,w_j\rangle| \le |\langle u_s,u_j\rangle| + \|u_j - w_j\|_W$ by the triangle and Cauchy-Schwarz inequalities, and the fact that u_j is a unit vector. Hence by induction we can assume that for every $j \in [s-1]$ we have $|\langle u_s,w_j\rangle| \le (1+C_j)\tau$. It follows that $\|u_s - w_s'\|_W \le \sum_{j\in[s-1]} |\langle u_s,w_j\rangle| \le \tau \sum_{j\in[s-1]} (1+C_j) = \tau\big(s + \sum_{j=3}^{s-1} C_j\big)$. This in turn implies, by the triangle inequality, that $|1 - \|w_s'\|_W| \le \big(s + \sum_{j=3}^{s-1} C_j\big)\tau$. Therefore, letting $w_s := w_s'/\|w_s'\|_W$, we have $\|u_s - w_s\|_W$ equals

$$\frac{\|u_s\|w_s'\|_W - w_s'\|_W}{\|w_s'\|_W} \le \frac{\|u_s(\|w_s'\|_W - 1)\|_W + \|u_s - w_s'\|_W}{1-(1-\|w_s'\|_W)} \le \frac{2\tau\big(s+\sum_{j=3}^{s-1} C_j\big)}{1-\tau\big(s+\sum_{j=3}^{s-1} C_j\big)}.$$

If $\tau \le 1/\big(2(s+\sum_{j=3}^{s-1} C_j)\big)$, then $\|u_s - w_s\|_W \le 4(s+\sum_{j=3}^{s-1} C_j)\tau$. Letting $C_s := 4(s+\sum_{j=3}^{s-1} C_j)$, we have that if $\tau \le 1/C_s$ then $\|u_s - w_s\|_W \le C_s\tau$. This completes the induction. □

We also use the following version of the Hoffman-Wielandt theorem [41, Corollary 6.3.8].

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Theorem 7.10. *Let* Z *be a finite abelian group and let* $A, B \in \mathbb{C}^{\mathrm{Z}\times\mathrm{Z}}$ *be self-adjoint* Z*-matrices. Let* $\lambda_1 \geq \cdots \geq \lambda_n$ *be the eigenvalues of* A *arranged in nonincreasing order, and similarly let* $\lambda'_1 \geq \cdots \geq \lambda'_n$ *be the eigenvalues of* B *arranged in nonincreasing order. Then*

$$\sum_{i=1}^{n} |\lambda'_i - \lambda_i|^2 \leq \|A - B\|_2^2. \tag{7.6}$$

Note that in the standard matrix-analysis setup (such as in [41]), the norm on the right side in (7.6) is the Frobenius norm (involving summation over all entries, rather than normalized summation as here), and the eigenvalues are also defined relative to summation. In (7.6), we give the equivalent *normalized* version of the Hoffman-Wielandt inequality, where we divide the original inequality by $n^2 = |\mathrm{Z}|^2$. On the right side, this division goes into having the normalized L^2-norm, that is, the Hilbert-Schmidt norm, while on the left, the division is compensated by the fact that our eigenvalues are $1/n$ times the usually defined eigenvalues (recall Definition 2.3).

The following result enables us, given proximal self-adjoint Z-matrices, to find proximal orthogonal projections onto dominant eigenspaces of these Z-matrices.

Theorem 7.11. *Let* $\rho_1, \rho_2 \in (0,1)$. *Let* T_1, T_2 *be self-adjoint* Z*-matrices such that* $\|T_2\|_2 \leq 1$ *and* $\|T_1 - T_2\|_2 \leq \frac{\rho_2\rho_1^4}{120}$. *Then there is* $\rho \in [\rho_1/2, \rho_1]$ *such that* $d(\mathcal{P}_{T_1,\rho}, \mathcal{P}_{T_2,\rho}) \leq \rho_2$, *whence in particular* $|\mathrm{Spec}_\rho(T_1)| = |\mathrm{Spec}_\rho(T_2)|$. *Moreover* T_2 *has no eigenvalue in* $[\rho - \frac{\rho_1^3}{30}, \rho + \frac{\rho_1^3}{30}]$.

Proof. Let $\delta = \|T_1 - T_2\|_2$. Let $S = \{\lambda_1 \geq \ldots \geq \lambda_t\} = \mathrm{Spec}_\rho(T_2) \cap [\rho_1/2, \rho_1]$ be the (multi)set of eigenvalues of T_2 which also lie in $[\rho_1/2, \rho_1]$. The assumption $\|T_2\|_2 \leq 1$, together with the fact that (by our general choice of normalizations) the norm $\|T_2\|_2$ is the sum of squares of the eigenvalues of T_2, implies that $\sum_{i\in[t]} \lambda_i^2 \leq \|T_2\|_2^2 \leq 1$, so $t \leq 4/\rho_1^2$. Note that there exists $\rho' \in [\rho_1/2 + \frac{\rho_1^3}{10}, \rho_1]$ such that $S \cap [\rho' - \frac{\rho_1^3}{10}, \rho') = \emptyset$. Indeed, otherwise, letting $S' = S \cup \{\rho_1/2, \rho_1\}$, we would have that any two consecutive elements of S' differ by at most $\frac{\rho_1^3}{10}$, implying that we could partition $[\rho_1/2, \rho_1]$ into $t+1$ consecutive intervals of length at most $\frac{\rho_1^3}{10}$, which yields the contradiction that $\frac{\rho_1^3}{10}(t+1) \leq \frac{\rho_1^3}{10}(\frac{4}{\rho_1^2}+1) < \frac{\rho_1}{2}$.

We claim that $(\mathrm{Spec}(T_1) \cup \mathrm{Spec}(T_2)) \cap [\rho' - 2\frac{\rho_1^3}{30}, \rho' - \frac{\rho_1^3}{30}) = \emptyset$. Indeed, Theorem 7.10 implies that each eigenvalue of T_1 is at distance at most δ from an eigenvalue of T_2. As the interval $[\rho' - \frac{\rho_1^3}{10}, \rho')$ contains no eigenvalue of T_2, the interval $[\rho' - \frac{\rho_1^3}{10} + \delta, \rho' - \delta)$ cannot contain eigenvalues of T_1, and since $\delta = \|T_1 - T_2\|_2 \leq \frac{\rho_2\rho_1^4}{120} \leq \frac{\rho_1^3}{30}$, the claim follows. Let $\rho := \rho' - \frac{\rho_1^3}{30}$, $P_1 := \mathcal{P}_{T_1,\rho}$, and $P_2 := \mathcal{P}_{T_2,\rho}$.

Let v be an eigenvector of T_1 with $\|v\|_2 = 1$ and eigenvalue κ in $\mathrm{Spec}_\rho(T_1)$. We have $\|T_2 v - \kappa v\|_2 = \|T_1 v + (T_2 - T_1)v - \kappa v\|_2 = \|(T_2 - T_1)v\|_2 \leq \delta$.

By Lemma 7.4, there exists $w \in \mathrm{Eigen}_\rho(T_2)$ such that $\|v - w\|_2 \leq \delta 30\rho_1^{-3}$ (more precisely, w is the combination of eigenvectors corresponding to the relevant eigenvalue cluster, as given by Lemma 7.4, noting in addition that, by our choice of ρ, there is no such eigenvalue less than ρ in this cluster). It follows that $\|v - P_2(v)\|_2 \leq \|v - w\|_2 \leq 30\delta\rho_1^{-3}$. The image of P_1 is $\mathrm{Eigen}_\rho(T_1)$, which has dimension $|\mathrm{Spec}_\rho(T_1)| \leq \rho^{-2}\sum_{\lambda\in\mathrm{Spec}(T_1)} \lambda^2 = \rho^{-2}\|T_1\|_2^2 \leq (1+\delta)^2/\rho^2 \leq 16/\rho_1^2$. Hence, by Lemma 7.7 applied with $P = P_1, Q = P_2$, we have $\sup_{v\in U_{P_1}:\|v\|\leq 1} \mathrm{dist}(v, U_{P_2}) \leq 120\delta\rho_1^{-4}$. Similarly, we obtain $\sup_{v\in U_{P_2}:\|v\|\leq 1} \mathrm{dist}(v, U_{P_1}) \leq 120\delta\rho_1^{-4}$. Hence, by (7.4), we have $d(P_1, P_2) \leq 120\delta\rho_1^{-4}$, so if $\delta \leq \rho_2\rho_1^4/120$, then $d(P_1, P_2) \leq \rho_2$. As $\rho_2 < 1$, by the last sentence of Lemma7.6 we have $|\mathrm{Spec}_\rho(T_1)| = |\mathrm{Spec}_\rho(T_2)|$. □

We can now prove the main result of this subsection.

Theorem 7.12. *Let* $\mathcal{H} : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ *and* $\mathcal{B} : \mathbb{R}_{>0} \times \mathbb{N} \to \mathbb{R}_{>0}$ *be arbitrary functions, and let* $\rho_0 \in (0,1)$. *Then there exists* $M = M(\rho_0, \mathcal{B}, \mathcal{H}) > 0$ *and* $\varepsilon_0 = \varepsilon_0(\rho_0, \mathcal{B}, \mathcal{H}) > 0$ *such that the following holds. For any finite abelian group* Z *and any* 1*-bounded function* $f : \mathrm{Z} \to \mathbb{C}$, *there exists* $m \leq M$, $\varepsilon = \varepsilon_{\mathcal{H},\rho_0,m} \in [\varepsilon_0, 1]$, *and* $\rho = \rho_{\rho_0,f} \in [\rho_0/2, \rho_0]$ *satisfying the following.*

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Letting P_ρ denote the orthogonal projection to $\mathrm{Eigen}_\rho\big(\mathcal{K}_\varepsilon(f\otimes\overline{f})\big)$ in $\mathbb{C}^{\mathrm{Z}}$, we have

$$\|f - P_\rho(f)\|_{U^3} \le 2\rho^{3/8}. \tag{7.7}$$

Moreover, there is a 2-step cfr nilspace X *of complexity $\le m$, a $\mathcal{B}(\rho_0,m)$-balanced morphism $\phi : \mathrm{Z}\to\mathrm{X}$, a 1-bounded m-Lipschitz function $F:\mathrm{X}\to\mathbb{C}$, and a set $S_\rho = S_{\rho,m,f}\subset\widehat{\mathrm{Z}_k(\mathrm{X})}$ with $|S_\rho| = |\mathrm{Spec}_\rho\big(\mathcal{K}_\varepsilon(f\otimes\overline{f})\big)| \le 10/\rho$, such that for every $\chi\in S_\rho$, letting $g_\chi = F_\chi\circ\phi$, we have $\|g_\chi\|_2^2\ge\rho$ and*

$$\|P_\rho(f) - \sum_{\chi\in S_\rho} g_\chi\|_2 \le \mathcal{H}(\rho_0), \tag{7.8}$$

and there exists $h:\mathrm{Z}\to\mathbb{C}$ such that $\|h-\sum_{\chi\in S_\rho} g_\chi\|_\infty\le\mathcal{H}(\rho_0)$ and $\|h\|_{U^3}^* = O_{\rho_0,\mathcal{B},\mathcal{H}}(1)$.

Remark 7.13. This result can be viewed as a spectral structure-randomness decomposition of order 2. In particular, note that by (7.8) and the properties of h, we have that $P_\rho(f)$ is a structured function of order 2 in the sense of Definition 2.23, with parameters $\big(O_{\rho_0,\mathcal{B},\mathcal{H}}(1), 2\mathcal{H}(\rho_0)\big)$.

Remark 7.14. It will follow from the proof that ρ can be chosen so that there are no eigenvalues of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$ in the interval $[\rho-\rho_0^3/30, \rho+\rho_0^3/30)$.

Proof. Let C_s be the constant defined in Lemma 7.9. Let

$$\Upsilon(\rho_0) := \min\left(1, \frac{\mathcal{H}(\rho_0)}{\rho_0^{5/2}/26+1}, \frac{(\rho_0/2)^{3/8}/2}{\rho_0^4/600+\rho_0^{5/2}/26+1}\right), \tag{7.9}$$

and let $\eta = \rho_0^8\Upsilon(\rho_0)^2/(9\cdot 10^6)\in[0,1]$. Let $\mathcal{D}:\mathbb{R}_{>0}\times\mathbb{N}\to\mathbb{R}_{>0}$ and $\varepsilon>0$ be chosen so that the following conditions hold:

(i) The number $\tau := C_{\mathrm{Q}(\eta,m)}\,\mathrm{Q}(\eta,m)^2\eta^{-4}\mathcal{D}(\eta,m)$ is at most 1,
(ii) $2\,\mathrm{Q}(\eta,m)C_{\mathrm{Q}(\eta,m)}\tau\le\sqrt{\eta}$,
(iii) $(\mathrm{Q}(\eta,m)^2+4)\mathcal{D}(\eta,m)\le\eta$,
(iv) $\varepsilon=\varepsilon_{m,\eta}>0$ is chosen following Remark 6.2 (which also imposes another condition on $\mathcal{D}$ being small enough), so that $\|E\|_2\le 23\sqrt{\eta}$.
(v) $\mathcal{D}(\eta,m)(\mathrm{Q}(\eta,m)+O_{\eta,m}(1))\le\eta^{1/2}$, where the implicit constant is defined as per Lemma 5.3,
(vi) $2\,\mathrm{Q}(\eta,m)\mathcal{D}(\eta,m)^{1/4^3}\le\eta^{1/2}$, and
(vii) $\mathcal{D}(\eta,m)\le\mathcal{B}(\rho_0,m)$.

We apply Theorem 6.1 to f with parameters η and $\mathcal{D}$. Thus, we obtain the decomposition $\mathcal{K}_\varepsilon(f\otimes\overline{f}) = \sum_{\chi\in S} g_\chi\otimes\overline{g_\chi}+E$. Note that letting $\varepsilon_0 := \min_{m\le M}(\varepsilon_{\eta,m})$ we have that indeed $\varepsilon\ge\varepsilon_0>0$ and $\varepsilon_0=\varepsilon_0(\rho_0,\mathcal{B})$.

Next, note that by our choices above, the unit vectors $u_\chi := g_\chi/\|g_\chi\|_2$ satisfy, for any $\chi\ne\chi'$,

$$|\langle u_\chi,u_{\chi'}\rangle| \le |\langle g_\chi,g_{\chi'}\rangle|/(\|g_\chi\|_2\|g_{\chi'}\|_2) \le \mathrm{Q}(\eta,m)^2\eta^{-4}\mathcal{D}(\eta,m)\le\tau/C_s, \tag{7.10}$$

with $s=|S|$ and τ defined in (i) (in particular $\tau\le 1$).

We apply Lemma 7.9 to $\{u_\chi\}_{\chi\in S}$ ordered decreasingly according to $\|g_\chi\|_2^2$. (7.11)

Thus, we obtain the orthonormal vectors w_χ. We then define $g'_\chi := \|g_\chi\|_2 w_\chi$ for each $\chi\in S$. We thus have the orthogonal set $\{g'_\chi:\chi\in S\}$ which satisfies

$$\|g'_\chi\|_2 = \|g_\chi\|_2 \text{ and } \|g_\chi - g'_\chi\|_2 = \|g_\chi\|_2\|u_\chi - w_\chi\|_2 \le C_s\tau. \tag{7.12}$$

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Now we consider the following Z-matrices:

$$T_1 := \sum_{\chi\in S} g'_\chi \otimes \overline{g'_\chi} \quad \text{and} \quad T_2 := \mathcal{K}_\varepsilon(f\otimes\overline{f}) = \sum_{\chi\in S} g_\chi \otimes \overline{g_\chi} + E. \tag{7.13}$$

Note that by Lemma 3.6 with $M = \mathcal{K}_\varepsilon(f\otimes\overline{f})$ and $N = 0$, we have $\|T_2\|_2 \le 1$. We also have $\|g'_\chi\otimes\overline{g'_\chi} - g_\chi\otimes\overline{g_\chi}\|_2 \le \|g'_\chi - g_\chi\|_2(\|g'_\chi\|_2 + \|g_\chi\|_2) \le 2C_s\tau$. Hence $\|T_1 - T_2\|_2 \le |S|2C_s\tau + \|E\|_2$, and by (ii) and (iv), we have $\|T_1 - T_2\|_2 \le 24\sqrt{\eta}$.

We want to apply Theorem 7.11 to T_1 and T_2 with $\rho_1 = \rho_0$, and $\rho_2 = \Upsilon(\rho_0)$. Hence, we want to ensure that $\|T_1 - T_2\|_2 \le 24\sqrt{\eta} \le \rho_0^4\Upsilon(\rho_0)/120$. But note that this holds by our choice of η. Hence, there exists $\rho\in[\rho_0/2,\rho_0]$ such that $d(\mathcal{P}_{T_1,\rho},\mathcal{P}_{T_2,\rho}) \le \Upsilon(\rho_0)$. By definition of the metric d (recall (7.3)) and the fact that $\|f\|_2 \le 1$, this implies that $\|\mathcal{P}_{T_2,\rho}(f) - \mathcal{P}_{T_1,\rho}(f)\|_2 \le \Upsilon(\rho_0)$.

Note that $\mathcal{P}_{T_2,\rho}(f)$ is the desired projection $P_\rho(f)$. Let $S_\rho := \{\chi\in S : \|g_\chi\|_2^2 \ge \rho\}$. Note that by (7.12) and Theorem 7.11 we know that $|S_\rho| = |\{\chi\in S : \|g'_\chi\|_2^2 \ge \rho\}| = |\mathrm{Spec}_\rho(T_2)|$ and that $\rho\in[\rho_0/2,\rho_0]$ can be chosen explicitly depending on T_2. Note also that by Proposition 5.8 and (iii) we have $|S_\rho| \le 10/\rho$.

We claim that $\mathcal{P}_{T_1,\rho}(f)$ is close in L^2 to $\sum_{\chi\in S_\rho} g_\chi$. Indeed, the vectors $w_\chi = \frac{g'_\chi}{\|g'_\chi\|_2}$ form an orthonormal set of eigenvectors of T_1, with corresponding eigenvalues $\|g'_\chi\|_2^2$. This together with $\|g'_\chi\|_2 = \|g_\chi\|_2$ implies the following expression for the projection:

$$\mathcal{P}_{T_1,\rho}(f) = \textstyle\sum_{\chi\in S_\rho}\left\langle f, \frac{g'_\chi}{\|g'_\chi\|_2}\right\rangle \frac{g'_\chi}{\|g'_\chi\|_2} = \sum_{\chi\in S_\rho}\left\langle f, \frac{g'_\chi}{\|g_\chi\|_2}\right\rangle \frac{g'_\chi}{\|g_\chi\|_2}.$$

By (7.12) we have $\|g_\chi - g'_\chi\|_2 \le C_s\tau$, so $\|\langle f, g'_\chi\rangle g'_\chi - \langle f, g_\chi\rangle g_\chi\|_2 \le 2C_s\tau$, and so

$$\begin{aligned}\Big\|\mathcal{P}_{T_1,\rho}(f) - \textstyle\sum_{\chi\in S_\rho}\left\langle f, \frac{g_\chi}{\|g_\chi\|_2}\right\rangle \frac{g_\chi}{\|g_\chi\|_2}\Big\|_2 &\le \textstyle\sum_{\chi\in S_\rho}\Big\|\left\langle f, \frac{g'_\chi}{\|g_\chi\|_2}\right\rangle \frac{g'_\chi}{\|g_\chi\|_2} - \left\langle f, \frac{g_\chi}{\|g_\chi\|_2}\right\rangle \frac{g_\chi}{\|g_\chi\|_2}\Big\|_2 \\ &\le 2\rho^{-1}|S_\rho|C_s\tau \ \le\ 2\rho^{-1}|S|C_s\tau \ \le\ \rho_0^3\Upsilon(\rho_0)/1500,\end{aligned}$$

where we used that $|S_\rho| \le |S|$, that $2|S|C_s\tau \le \sqrt{\eta}$ (by (ii)), and that $\rho \ge \rho_0/2$.

Finally, by (5.5), for every $\chi\in S_\rho$ we have

$$|\langle f, g_\chi/\|g_\chi\|_2\rangle - \|g_\chi\|_2| \le \rho^{-1/2}\big(\mathcal{D}(\eta,m)(\mathrm{Q}(\eta,m) + O_{\eta,m}(1)) + 3\eta^{1/2}\big).$$

Recalling the definition of η and using (v), we have $|\langle f, \frac{g_\chi}{\|g_\chi\|_2}\rangle - \|g_\chi\|_2| \le \rho_0^{7/2}\Upsilon(\rho_0)/530$. Hence, using $|S_\rho| \le 10/\rho$ and $\rho \ge \rho_0/2$, we conclude that

$$\Big\|\textstyle\sum_{\chi\in S_\rho} g_\chi - \sum_{\chi\in S_\rho}\left\langle f, \frac{g_\chi}{\|g_\chi\|_2}\right\rangle \frac{g_\chi}{\|g_\chi\|_2}\Big\|_2 \le |S_\rho|\rho_0^{7/2}\Upsilon(\rho_0)/530 \le 2\rho_0^{5/2}\Upsilon(\rho_0)/53.$$

Hence $\|\mathcal{P}_{T_1,\rho}(f) - \sum_{\chi\in S_\rho} g_\chi\|_2 \le \rho_0^{5/2}\Upsilon(\rho_0)/26$, which proves our claim.

Using this last estimate we can prove (7.8), indeed $\|P_\rho(f) - \sum_{\chi\in S_\rho} g_\chi\|_2$ is at most

$$\|\textstyle\sum_{\chi\in S_\rho} g_\chi - \mathcal{P}_{T_1,\rho}(f)\|_2 + \|\mathcal{P}_{T_1,\rho}(f) - \mathcal{P}_{T_2,\rho}(f)\|_2 \le \rho_0^{5/2}\Upsilon(\rho_0)/26 + \Upsilon(\rho_0) \le \mathcal{H}(\rho_0).$$

Now, using the above estimates, the fact that the L^2-norm dominates the U^3-norm, and the bound (5.11), we conclude that

$$\begin{aligned}\|f - P_\rho(f)\|_{U^3} &\le \|f - \textstyle\sum_{\chi\in S_\rho} g_\chi\|_{U^3} + \|P_\rho(f) - \sum_{\chi\in S_\rho} g_\chi\|_2 \\ &\le \Big[\rho^{3/8}(1+c_0)^{1/8} + 2|S|\mathcal{D}(\eta,m)^{1/4^3} + 4\eta^{1/2}\Big] + \rho_0^{5/2}\Upsilon(\rho_0)/26 + \Upsilon(\rho_0).\end{aligned}$$

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Now using that $1+c_0 \le 1+(|S|^2+4)\mathcal{D}(\eta,m)+8\eta \le 1+9\eta \le 10$ (by (iii)), and that $2|S|\mathcal{D}(\eta,m)^{1/4^3} \le \eta^{1/2}$ by (vi), we deduce that the last line above is at most $10^{1/8}\rho^{3/8} + 5\eta^{1/2} + \Upsilon(\rho_0)\left(\rho_0^{5/2}/26+1\right) \le 10^{1/8}\rho^{3/8} + \Upsilon(\rho_0)\left(\rho_0^4/600 + \rho_0^{5/2}/26+1\right)$, and this is at most $2\rho^{3/8}$ by our choice of Υ.

To prove the final statement of the theorem, we apply Proposition 5.8 for $\sigma := \mathcal{H}(\rho_0)$. Thus, we obtain some $h : \mathrm{Z} \to \mathbb{C}$ such that $\|h - \sum_{\chi\in S_\rho} g_\chi\|_\infty \le \mathcal{H}(\rho_0)$ and $\|h\|_{U^3}^* = O_{\rho_0,\mathcal{B},\mathcal{H}}(1)$. □

We can now complete the proof of the validity of our spectral U^3-regularization algorithm.

Proof of Theorem 1.1. Note that this theorem is just a version of Theorem 7.12 where we fix an arbitrary function $\mathcal{B}$ and we let $\mathcal{H}(\rho_0) = \rho_0/2$. □

7.2. *Obtaining dominant nilspace characters of f from the spectrum of* $\mathcal{K}_\varepsilon(f\otimes\overline{f})$

Now that we have a method for recovering algorithmically the structured part $\sum_{\chi\in S_\rho} g_\chi$ using the eigenvectors of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$, it is natural to try to go further and recover as much as possible the individual nilspace characters $g_\chi, \chi \in S_\rho$. Lemma 7.4 guarantees that every such function g_χ is (up to an error with small L^2-norm) a linear combination of eigenvectors of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$ that have corresponding eigenvalues close to the number $\|g_\chi\|_2^2$. The main problem that we may face is that there could be a cluster of many such eigenvalues, with corresponding eigenvectors thus spanning a multidimensional space, in which case we would have to find how to recover g_χ algorithmically inside this space.

A simple observation is that if the dominant eigenvalues of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$ are sufficiently isolated, then the above clustering cannot happen and we are therefore able to recover g_χ as a scalar multiple of precisely one of the eigenvectors (up to a small L^2-error). In fact, this holds more generally in a sense that will be very useful for us. Namely, for any unit vector h in the linear span of the leading eigenvectors of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$, if the leading eigenvalues of the *new* matrix $\mathcal{K}_\varepsilon(h\otimes\overline{h})$ are well-separated, then there is a bijection from the set of leading eigenvectors of the latter matrix to the set of leading nilspace character g_χ of f, such that g_χ is close in L^2 to a scalar multiple of its corresponding eigenvector under this bijection. This is formalized in the following main theorem of this subsection.

We say that a multiset X of complex numbers is *δ-separated* if $|x-y| > \delta$ for every $x \ne y$ in X. Note that, in particular, this implies that every element of X has multiplicity 1.

Theorem 7.15. *For any $q > 0$ and $r > q+1$, there exists $B_{q,r} \in (0,1)$ such that the following holds. Let $\mathcal{B}$ be a function $\mathbb{R}_{>0}\times\mathbb{N} \to \mathbb{R}_{>0}$, let $\mathcal{H} : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$, $x \mapsto (x/2)^r$, and let $\rho_0 \in (0, B_{q,r}]$. Let $M(\rho_0,\mathcal{B},\mathcal{H}), \varepsilon_0(\rho_0,\mathcal{B},\mathcal{H})$ be the resulting numbers given by Theorem 7.12. Let Z be a finite abelian group, let $f : \mathrm{Z} \to \mathbb{C}$ be 1-bounded, and, for an appropriate $\rho \in [\rho_0/2, \rho_0]$, let $f = \sum_{\chi\in S_\rho} g_\chi + f_r + f_e$ be the corresponding decomposition given by Theorem 7.12, where $\|f_r\|_{U^3} \le 2\rho^{3/8}$ and $\|f_e\|_2 \le (\rho_0/2)^r$. Let h be a function in $\mathrm{Eigen}_\rho\big(\mathcal{K}_\varepsilon(f\otimes\overline{f})\big) \subset \mathbb{C}^{\mathrm{Z}}$, with $\|h\|_2 \le 1$, such that $A := \mathcal{K}_\varepsilon(h\otimes\overline{h})$ has spectrum $\mathrm{Spec}_{\rho^q}(A)$ being δ-separated for some $\delta > \rho^q$, and $|\mathrm{Spec}_{\rho^q}(A)| = |S_\rho|$. Then there is a bijection $J : \mathrm{Spec}_{\rho^q}(A) \to S_\rho$ such that, for every $\lambda \in \mathrm{Spec}_{\rho^q}(A)$, letting $\chi = J(\lambda)$ and v be an eigenvector of A (with $\|v\|_2 = 1$) corresponding to λ, we have $\|\langle f, v\rangle v - g_\chi\|_2 \le 28\rho^{r-q}$ and $|\langle f, v\rangle| \ge \sqrt{\rho/2}$.*

We can let $B_{q,r} := \min\big(1, (1/14)^{1/(r-1)}, (1/14)^{1/(r-q)}, (1/56)^{1/(r-q)}, (1/452)^{1/(r-q-1)}\big)$. Let us also mention that the specific choice of function $(x/2)^r$ in the theorem is sufficient for our purposes in what follows, but that the result could be extended to a more general function $\mathcal{H}$.

Remark 7.16. The structured part $P_\rho(f)$ of f obtained in Theorem 7.12 has L^2-norm at most 1 (being an orthogonal projection of the 1-bounded function f), and it can then be seen that, setting $h = P_\rho(f)$, the spectra $\mathrm{Spec}_\rho(\mathcal{K}_\varepsilon(h\otimes\overline{h}))$ and $\mathrm{Spec}_\rho(\mathcal{K}_\varepsilon(f\otimes\overline{f}))$ are roughly equal (this is proved in Lemma 7.18 below). Therefore, using this choice of h, Theorem 7.15 implies that, if $\mathrm{Spec}_\rho\big(\mathcal{K}_\varepsilon(f\otimes\overline{f})\big)$ is sufficiently separated, then every nilspace character g_χ in the decomposition of f with $\chi \in S_\rho$ can be recovered directly (up to a small L^2-error) as an eigenvector in $\mathrm{Eigen}_\rho\big(\mathcal{K}_\varepsilon(h\otimes\overline{h})\big)$. The purpose of the

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

greater generality of Theorem 7.15 (enabling us to consider more general functions h in the unit ball of $\mathrm{Eigen}_\rho(\mathcal{K}_\varepsilon(f\otimes\overline{f}))$) is to give us more options for recovering the individual quadratic characters of f, even when the eigenvalues of $\mathcal{K}_\varepsilon(f\otimes\overline{f})$ are *not* sufficiently separated (a situation that can occur even in the classical Fourier-analytic setting, as we mentioned in Remark 2.13). Indeed, the subspace $\mathrm{Eigen}_\rho(\mathcal{K}_\varepsilon(f\otimes\overline{f}))$ has a dimension much smaller than $|\mathrm{Z}|$, and the theorem tells us that it suffices to find a vector (or function) h in this subspace such that the dominant eigenvalues of $\mathcal{K}_\varepsilon(h\otimes\overline{h})$ are sufficiently separated. In the last part of this subsection, we will see that it is useful to consider a *random* vector h in the unit ball of this subspace, because we can then guarantee that, with high probability, the spectrum $\mathrm{Spec}_{\rho^q}(\mathcal{K}_\varepsilon(h\otimes\overline{h}))$ for such a random h is adequately separated.

The proof of Theorem 7.15 relies on the following result. This tells us essentially that if h is *any* unit vector that is close to a linear combination of a few quasiorthogonal quadratic characters of unit norm, then the leading eigenvalues of $\mathcal{K}_\varepsilon(h\otimes\overline{h})$ are close to the squared moduli of the coefficients in the linear combination.

Proposition 7.17. *Let $\varepsilon\in(0,1]$, let Z be a finite abelian group, let D be a finite set, and let g_χ, $\chi\in D$ be functions in $\mathbb{C}^{\mathrm{Z}}$ satisfying $\|g_\chi\|_2\in[\tau_1,1]$, $\|\mathcal{K}_{\varepsilon^{1/2}}(g_\chi\otimes\overline{g_\chi})-g_\chi\otimes\overline{g_\chi}\|_2\le\tau_2$, and $\max(|\langle g_\chi,g_{\chi'}\rangle|,\langle g_\chi,g_{\chi'}\rangle_{U^3})\le\tau_3\le\tau_1^2/|D|$ for every $\chi'\in D\setminus\{\chi\}$. Let $u_\chi=g_\chi/\|g_\chi\|_2$ for each $\chi\in D$, and suppose that $h\in\mathbb{C}^{\mathrm{Z}}$ is a function with $\|h\|_2\le1$ such that, for some function $h'=\sum_{\chi\in D}c_\chi u_\chi$ with $|c_\chi|\le1$, we have $\|h-h'\|_2\le\tau_4$. Then*

$$\mathcal{K}_\varepsilon(h\otimes\overline{h})=\textstyle\sum_{\chi\in D}|c_\chi|^2u_\chi\otimes\overline{u_\chi}+E, \tag{7.14}$$

where $\|E\|_2\le5\tau_4+\varepsilon^{-1/2}|D|^2(|D|^{1/2}+3)\tau_3+|D|\max\big(\tau_2/\tau_1^2,2\varepsilon^{1/2}\big)$. In particular, letting $\lambda_1\ge\lambda_2\ge\cdots\ge\lambda_{|\mathrm{Z}|}$ be the eigenvalues of the matrix $\mathcal{K}_\varepsilon(h\otimes\overline{h})$, and $\chi_1,\chi_2,\dots,\chi_{|D|}$ be an ordering such that $c_{\chi_i}\ge c_{\chi_{i+1}}$ for all i, and setting $c_{\chi_i}=0$ for $i>|D|$, we have

$$\big(\textstyle\sum_{i=1}^{|\mathrm{Z}|}\big||c_{\chi_i}|^2-\lambda_i\big|^2\big)^{\frac12}\le2|D|C_{|D|}^2\tau_3/\tau_1^2+\|E\|_2. \tag{7.15}$$

To prove Proposition 7.17, we will combine the refined Gram-Schmidt process (Lemma 7.9) with Theorem 7.10.

Proof of Proposition 7.17. We first apply Proposition 3.17 with the following elements and parameters α_i: we take $f=h=\sum_{\chi\in D}c_\chi u_\chi+f_e$, where $f_e:=h-h'$ (so $\alpha_1=\tau_4$), $f_r=0$ (so $\alpha_2=0$), the f_i are the functions $c_\chi u_\chi$, so $\alpha_3=1$ and $\alpha_4=\tau_3$. Finally, we have $\|\mathcal{K}_\varepsilon\big((c_\chi u_\chi)\otimes\overline{(c_\chi u_\chi)}\big)-(c_\chi u_\chi)\otimes\overline{(c_\chi u_\chi)}\|_2\le\alpha_5$ where we shall now estimate α_5 using Lemma 2.55. Indeed, since $\|\mathcal{K}_{\varepsilon^{1/2}}(g_\chi\otimes\overline{g_\chi})-g_\chi\otimes\overline{g_\chi}\|_2\le\tau_2$, we deduce by Lemma 2.55 applied with $c=c_\chi/\|g_\chi\|_2$ (and using that $\|g_\chi\|_2\le1$) that

$$\|\mathcal{K}_\varepsilon\big((c_\chi u_\chi)\otimes\overline{(c_\chi u_\chi)}\big)-(c_\chi u_\chi)\otimes\overline{(c_\chi u_\chi)}\|_2\le\max\big(|c_\chi|^2\tau_2/\|g_\chi\|_2^2,2\varepsilon^{1/2}\big)\le\max\big(\tau_2/\tau_1^2,2\varepsilon^{1/2}\big),$$

so we can let $\alpha_5=\max\big(\tau_2/\tau_1^2,2\varepsilon^{1/2}\big)$.

Thus we obtain

$$\mathcal{K}_\varepsilon(h\otimes\overline{h})=\textstyle\sum_{\chi\in D}|c_\chi|^2u_\chi\otimes\overline{u_\chi}+E,$$

where $\|E\|_2\le5\tau_4+\varepsilon^{-1/2}|D|^2(|D|^{1/2}+3)\tau_3+|D|\max\big(\tau_2/\tau_1^2,2\varepsilon^{1/2}\big)$.

Now, we apply Theorem 7.10 to deduce that, letting λ_i' be the eigenvalues of the self-adjoint matrix $M_u:=\sum_{\chi\in D}|c_\chi|^2u_\chi\otimes\overline{u_\chi}$ in descending order, and letting $n=|\mathrm{Z}|$, we have

$$\textstyle\sum_{i\in[n]}|\lambda_i'-\lambda_i|^2\le\big\|\mathcal{K}_\varepsilon(h\otimes\overline{h})-\sum_{\chi\in D}|c_\chi|^2u_\chi\otimes\overline{u_\chi}\big\|_2^2=\|E\|_2^2.$$

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

We can thus focus on approximating the eigenvalues λ_i'. For this we apply Lemma 7.9 to the quasi-orthogonal set $\{u_\chi : \chi \in D\}$ with $\tau = C_{|D|}\tau_3/\tau_1^2 \le 1$ (note that, for $\chi \ne \chi'$, $|\langle u_\chi, u_{\chi'}\rangle| = |\langle \frac{g_\chi}{\|g_\chi\|_2}, \frac{g_{\chi'}}{\|g_{\chi'}\|_2}\rangle| \le \tau_3/\tau_1^2$), obtaining an orthonormal set $\{w_\chi : \chi \in D\}$ with $\|u_\chi - w_\chi\|_2 \le C_{|D|}^2\tau_3/\tau_1^2$ for all χ.

It follows that the matrices M_u and $M_w := \sum_{\chi\in D} |c_\chi|^2 w_\chi \otimes \overline{w_\chi}$ satisfy

$$\|M_u - M_w\|_2 \le \textstyle\sum_{\chi\in D} |c_\chi|^2 \, \|(u_\chi - w_\chi)\otimes\overline{u_\chi} - w_\chi\otimes(\overline{w_\chi} - \overline{u_\chi})\|_2 \le 2|D|C_{|D|}^2\tau_3/\tau_1^2.$$

Moreover, the eigenvalues of the matrix M_w are clearly determined: since the w_χ are orthonormal, it follows that for each $\chi \in D$ we have $M_w w_\chi = |c_\chi|^2 w_\chi$, so each w_χ is an eigenvector with eigenvalue $|c_\chi|^2$. Applying again Theorem 7.10 to M_u, M_w, we deduce that $\big(\sum_{i=1}^{|R|} |\lambda_i' - |c_{\chi_i}|^2|^2\big)^{1/2} \le 2|D|C_{|D|}^2\tau_3/\tau_1^2$. Hence for each i we have $|\lambda_i - |c_{\chi_i}|^2| \le 2|D|C_{|D|}^2\tau_3/\tau_1^2 + \|E\|_2$ and the result follows. □

We can now prove the main result of this section.

Proof of Theorem 7.15. Recall that our assumptions include an application of Theorem 7.12. In particular, we have applied Theorem 6.1 with some particular choices of η and $\mathcal{D}$. For simplicity, instead of applying again Theorem 6.1 with a new list of constants, we are going to assume that we have repeated the argument of Theorem 7.12 with two additional constraints specified below, (viii) and (ix), on ε and $\mathcal{D}$. Recall also that $\mathcal{H}(\rho_0) = (\rho_0/2)^r$, so in particular $\Upsilon(\rho_0) \le (\rho_0/2)^r \le \rho^r$, and thus $\eta = \rho_0^8\Upsilon(\rho_0)/(9\cdot 10^6) \le \rho^{8+2r}/5^6$.

Let T_1 and T_2 be defined as in (7.13):

$$T_1 := \textstyle\sum_{\chi\in S} g_\chi' \otimes \overline{g_\chi'} \quad \text{and} \quad T_2 := \mathcal{K}_\varepsilon(f\otimes\overline{f}) = \sum_{\chi\in S} g_\chi\otimes\overline{g_\chi} + E. \tag{7.16}$$

Then, arguing as in the proof of Theorem 7.12, we have $d(\mathcal{P}_{T_1,\rho}, \mathcal{P}_{T_2,\rho}) \le \Upsilon(\rho_0) \le \rho^r$. Since $\|h\|_2 \le 1$ and $h = \mathcal{P}_{T_2,\rho}(h)$, it follows that $\|\mathcal{P}_{T_1,\rho}(h) - h\|_2 \le \rho^r$. Moreover $\mathcal{P}_{T_1,\rho}(h) = \sum_{\chi\in S_\rho}\langle h, w_\chi\rangle w_\chi$, where $w_\chi = g_\chi'/\|g_\chi'\|_2$ (indeed the orthonormality of the w_χ implies that the eigenvalues of T_1 are precisely the numbers $\|g_\chi'\|_2^2 = \|g_\chi\|_2^2$, $\chi \in S$, so $\mathrm{Eigen}_\rho(T_1)$ is the linear span of $\{w_\chi : \chi \in S_\rho\}$). By Lemma 7.9, we have $\|u_\chi - w_\chi\|_2 \le C_s\tau$ where τ is defined as per (i) for $\chi \in S_\rho$ (where $s = |S|$). Combining this with condition (ii) from the proof of Theorem 7.12, and with the fact that $S_\rho \subset S$, we conclude that $\big\|\mathcal{P}_{T_1,\rho}(h) - \sum_{\chi\in S_\rho}\langle h, w_\chi\rangle u_\chi\big\|_2 \le C_s\tau\,|S_\rho| \le \mathrm{Q}(\eta,m)C_s\tau \le \eta^{1/2}/2 \le \rho^{4+r}/250$. Thus, letting $c_\chi = \langle h, w_\chi\rangle$, we have

$$\big\|h - \textstyle\sum_{\chi\in S_\rho} c_\chi u_\chi\big\|_2 \le 2\rho^r. \tag{7.17}$$

We now want to apply Proposition 7.17 to h, with $D = S_\rho$ and $h' = \sum_{\chi\in S_\rho} c_\chi u_\chi$. By definition of S_ρ (recall Proposition 5.8), we can let $\tau_1 = (\rho_0/2)^{1/2}$. To estimate τ_2, we need to bound $\|\mathcal{K}_{\varepsilon^{1/2}}(g_\chi\otimes\overline{g_\chi}) - (g_\chi\otimes\overline{g_\chi})\|_2$. Note that $g_\chi = F_\chi\circ\phi$ where $F_\chi : \mathrm{X} \to \mathbb{C}$ is a Lipschitz function with $\|F_\chi\|_{\mathrm{sum}} \le 1+m$. By Theorem 4.19, for any $\sigma \in (0,1/2)$ we have that g_χ is a quadratic character with parameters $(O_m(\sigma^{-O_m(1)}),\sigma)$. By Proposition 2.52, we have $\|\mathcal{K}_{\varepsilon^{1/2}}\big((g_\chi)\otimes\overline{(g_\chi)}\big) - (g_\chi)\otimes\overline{(g_\chi)}\|_2 \le 4\sigma + O_m(\varepsilon^{1/8}\sigma^{-O_m(1)})$. Setting $\sigma = \eta/4$, we have $\|\mathcal{K}_{\varepsilon^{1/2}}\big((g_\chi)\otimes\overline{(g_\chi)}\big) - (g_\chi)\otimes\overline{(g_\chi)}\|_2 \le \eta + O_{m,\eta}(\varepsilon^{1/8}) =: \tau_2$. Finally, recall that by Theorem 5.1 we can let $\tau_3 = \mathcal{D}(\eta,m)$, and since $|S_\rho| \le |S| \le \mathrm{Q}(\eta,m)$, we have that $\tau_3|D| \le \tau_3\,\mathrm{Q}(\eta,m) \le \mathcal{D}(\eta,m)\,\mathrm{Q}(\eta,m)$, which is at most $\tau_1^2 = \rho_0/2$ by assumption (i) from the proof of Theorem 7.12 and the above bound on η. Finally, by (7.17) we can let $\tau_4 = 2\rho^r$. Applying now Proposition 7.17 with these parameters τ_i, we obtain (7.14) with

$$\begin{aligned}\|E\|_2 &\le 5\tau_4 + \varepsilon^{-1/2}|D|^2(|D|^{1/2}+3)\tau_3 + |D|\max\big(\tau_2/\tau_1^2, 2\varepsilon^{1/2}\big)\\ &\le 10\rho^r + \varepsilon^{-1/2}|D|^2(|D|^{1/2}+3)\mathcal{D}(\eta,m) + |D|\max\big(2(\eta + O_{m,\eta}(\varepsilon^{1/8}))/\rho_0, 2\varepsilon^{1/2}\big)\\ &\le 11\rho^r + O_{m,\eta}(\varepsilon^{-1/2}\mathcal{D}(\eta,m)) + O_{m,\eta}(\varepsilon^{1/8}),\end{aligned} \tag{7.18}$$

where for this last inequality we used that $2|D|\eta \le (20/\rho)\eta \le \rho^r$.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

To simplify the upper bound in (7.18) further, we impose additional constraints on both $\varepsilon = \varepsilon_{\eta,m}$ and $\mathcal{D}$. These assumptions are to be thought of as added to the list at the beginning of the proof Theorem 7.12, as the following eighth condition:

(viii) We let $\varepsilon = \varepsilon_{\eta,m} > 0$ be such that the term $O_{m,\eta}(\varepsilon^{1/8})$ in (7.18) is at most $(\rho_0/2)^r$. Furthermore, we then let $\mathcal{D}$ be such that the term $O_{m,\eta}(\varepsilon^{-1/2}\mathcal{D}(\eta,m))$ in (7.18) is at most $(\rho_0/2)^r$.

Hence $\|E\|_2 \leq 13\rho^r$. Let $\{\lambda_1 \geq \lambda_2 \geq \cdots\}$ be the eigenvalues of $\mathcal{K}_\varepsilon(h \otimes \overline{h})$ ordered decreasingly, let the elements of S_ρ be ordered $\chi_1, \chi_2, \ldots$ so that $|c_{\chi_1}|^2 \geq |c_{\chi_2}|^2 \geq \cdots$, and also let $c_{\chi_i} := 0$ for $|S_\rho| < i \leq |\mathrm{Z}|$. We then have (7.15), and thus, for all $i \in [|\mathrm{Z}|]$, it follows that $||c_{\chi_i}|^2 - \lambda_i| \leq \gamma := 2|S_\rho|C^2_{|S_\rho|}\tau_3/\tau_1^2 + \|E\|_2$. Similarly, we would like to have a simpler bound for this quantity, so we assume the following additional condition:

(ix) The function $\mathcal{D}$ is chosen so that $2\,\mathrm{Q}(\eta,m)C^2_{\mathrm{Q}(\eta,m)}\mathcal{D}(\eta,m)(\rho_0/2)^{-1} \leq (\rho_0/2)^r$.

Thus we obtain that

$$\forall i \in [|S_\rho|],\ \text{we have } ||c_{\chi_i}|^2 - \lambda_i| \leq 14\rho^r \text{ and } \forall i \in [|S_\rho|+1, |\mathrm{Z}|],\ \text{we have } |\lambda_i| \leq 14\rho^r. \tag{7.19}$$

Our aim now is to apply Lemma 7.4 to $A := \mathcal{K}_\varepsilon(h \otimes \overline{h})$. Let $\{v_i : i \in |\mathrm{Z}|\}$ be an orthonormal basis of eigenvectors of A, where v_i has corresponding eigenvalue λ_i.

We claim that each u_χ with $\chi \in S_\rho$ is a β-pseudoeigenvector of A with pseudoeigenvalue $|c_\chi|^2$, for $\beta = 14\rho^r$. Indeed (7.14) implies $\mathcal{K}_\varepsilon(h \otimes \overline{h})u_\chi = \sum_{\chi' \in S_\rho} |c_{\chi'}|^2 u_{\chi'}\langle u_\chi, u_{\chi'}\rangle + Eu_\chi$, so

$$\begin{aligned}\|\mathcal{K}_\varepsilon(h \otimes \overline{h})u_\chi - |c_\chi|^2 u_\chi\|_2 &\leq \textstyle\sum_{\chi' \in S_\rho \setminus \{\chi\}} |c_{\chi'}|^2 |\langle u_\chi, u_{\chi'}\rangle| + \|E\|_2 \\ &\leq \tau C_s\,\mathrm{Q}(\eta,m) + \|E\|_2 \leq 14\rho^r,\end{aligned}$$

the second inequality holding by (7.10) and the third by (ii), using $\eta \leq \rho^r$ and our bound on $\|E\|_2$.

Fix any $\lambda \in \mathrm{Spec}_{\rho^q}(A)$ and let v be a corresponding unit eigenvector. Since $\mathrm{Spec}_{\rho^q}(A)$ is δ-separated and $\lambda \geq \rho^q > 14\rho^r$, we have $||c_\chi|^2 - \lambda| \leq 14\rho^r$ for some $\chi \in S_\rho$. Let $C_{\rho^q/2}(|c_\chi|^2) := \{i \in [|\mathrm{Z}|] : |\lambda_i - |c_\chi|^2| \leq \rho^q/2\}$. We claim that $C_{\rho^q/2}(|c_\chi|^2) = \{\lambda\}$. First, note that if $i > |S_\rho|$, the second part of (7.19) tells us that $\lambda_i \leq 14\rho^r$. Since $\rho^q - \rho^q/2 - 14\rho^r > 14\rho^r$, it follows that for $i > |S_\rho|$ we have $|\lambda_i - |c_\chi|^2| > \rho^q/2$. Now if $i \leq |S_\rho|$, then, since by assumption $|\mathrm{Spec}_{\rho^q}(A)| = |S_\rho|$, it follows that $\lambda_i \geq \rho^q$ for all $i \leq |S_\rho|$. Let $\lambda_i \neq \lambda$ for $i \geq |S_\rho|$. We want to prove that $|\lambda_i - |c_\chi|^2| > \rho^q$. As $\mathrm{Spec}_{\rho^q}(A)$ is δ-separated, we have $|\lambda - \lambda_i| \geq \delta \geq \rho^q$. Thus, $|\lambda_i - |c_\chi|^2| \geq |\lambda_i - \lambda| - |\lambda - |c_\chi|^2| \geq \rho^q - 14\rho^r > \rho^q/2$. Therefore, we can define a map $J : \mathrm{Spec}_{\rho^q}(A) \to S_\rho$ that associates λ with the unique $\chi \in S_\rho$ such that $|\lambda - |c_\chi|^2| \leq \rho^q/2$. As $|\mathrm{Spec}_{\rho^q}(A)| = |S_\rho|$, J is a bijection. By Lemma 7.4, we have $\|u_\chi - \langle u_\chi, v\rangle v\|_2 \leq 28\rho^{r-q}$.

In particular $|1 - |\langle u_\chi, v\rangle|\ | \leq 28\rho^{r-q} \leq 1/2$. Hence $\|u_\chi/\langle u_\chi, v\rangle - v\|_2 \leq \frac{1}{\langle u_\chi, v\rangle}\beta/\delta \leq 56\rho^{r-q}$. Note that $\left|\frac{1}{\langle u_\chi, v\rangle} - \frac{\overline{\langle u_\chi, v\rangle}}{|\langle u_\chi, v\rangle|}\right| = \left|\frac{1}{|\langle u_\chi, v\rangle|} - 1\right| = \left|\frac{|\langle u_\chi, v\rangle| - 1}{\langle u_\chi, v\rangle}\right| \leq 56\rho^{r-q}$ and thus, letting $z := \frac{\overline{\langle u_\chi, v\rangle}}{|\langle u_\chi, v\rangle|}$ we have that $\left\|v - \frac{zg_\chi}{\|g_\chi\|_2}\right\|_2 \leq 112\rho^{r-q}$.

With this, we can now recover the component g_χ of f. Indeed, note that

$$\|g_\chi - \langle f, v\rangle v\|_2 \leq \|g_\chi - \langle f, \tfrac{zg_\chi}{\|g_\chi\|_2}\rangle \tfrac{zg_\chi}{\|g_\chi\|_2}\|_2 + \|\langle f, v - \tfrac{zg_\chi}{\|g_\chi\|_2}\rangle v\|_2 + \|\langle f, \tfrac{zg_\chi}{\|g_\chi\|_2}\rangle (v - \tfrac{zg_\chi}{\|g_\chi\|_2})\|_2.$$

Thus $\|g_\chi - \langle f, v\rangle v\|_2 \leq \|g_\chi - \langle f, \frac{zg_\chi}{\|g_\chi\|_2}\rangle \frac{zg_\chi}{\|g_\chi\|_2}\|_2 + 224\rho^{r-q}$. Moreover $\|g_\chi - \langle f, \frac{zg_\chi}{\|g_\chi\|_2}\rangle \frac{zg_\chi}{\|g_\chi\|_2}\|_2 = \frac{1}{\|g_\chi\|_2}|\|g_\chi\|_2^2 - \langle f, g_\chi\rangle g_\chi|$. As $\chi \in S_\rho$, we have $\|g_\chi\|_2 \geq \sqrt{\rho}$, and by Lemma 5.3 combined with (v) and $4\sqrt{\eta}/\sqrt{\rho} \leq \rho^r$, we have that $\|g_\chi - \langle f, v\rangle v\|_2 \leq \rho^r + 224\rho^{r-q} \leq 225\rho^{r-q}$.

Finally, to see that v indeed correlates with f, note that by Lemma 5.3 combined with (v), we have $|\langle f, g_\chi\rangle - \|g_\chi\|_2^2| \leq \rho^r$ (similarly as before). Hence $|\langle f, \langle f, v\rangle v\rangle| \geq |\langle f, g_\chi\rangle| - \|g_\chi - \langle f, v\rangle v\|_2 \geq \rho - 226\rho^{r-q} \geq \rho/2$ where we used that $\|g_\chi\|_2^2 \geq \rho$. Therefore $|\langle f, v\rangle| \geq \sqrt{\rho/2}$. □

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

The second consequence of Proposition 7.17 is the following confirmation of the fact mentioned at the beginning of Remark 7.16.

Lemma 7.18. *Under the same assumptions as Theorems 7.12 and 7.15, assume further that* $r > 3$ *and* $\rho < (1/450)^{1/(r-3)}$. *Let* $h = P_\rho(f)$. *Then there is a bijection* $B : \mathrm{Spec}_\rho(\mathcal{K}_\varepsilon(h \otimes \overline{h})) \to \mathrm{Spec}_\rho(\mathcal{K}_\varepsilon(f \otimes \overline{f}))$ *which preserves multiplicity of eigenvalues and satisfies* $|B(\lambda) - \lambda| \leq 15\rho^r$.

Proof. Recall from (7.13) the definition of the Z-matrices T_1 and T_2 and recall that $\|T_1 - T_2\|_2 \leq 24\sqrt{\eta}$. By definition, the $|\mathrm{Z}|$ eigenvalues of T_1 are precisely $\|g_{\chi_1}\|_2^2 \geq \|g_{\chi_2}\|_2^2 \geq \cdots$ (where we complete the list with zeroes if necessary), and the $|\mathrm{Z}|$ eigenvalues of T_2 are $\lambda_1 \geq \lambda_2 \geq \cdots$. By Theorem 7.10 we have $\sum_{i=1}^{|S_\rho|} |\|g_{\chi_i}\|_2^2 - \lambda_i|^2 \leq \sum_{i=1}^{|\mathrm{Z}|} |\|g_{\chi_i}\|_2^2 - \lambda_i|^2 \leq 24^2\eta \leq \rho^{2r}$.

Now we apply Proposition 7.17 with $h = P_\rho(f)$ and $h' = \sum_{\chi \in S_\rho} g_\chi$, using that $\|P_\rho(f)\|_2 \leq 1$ and that, by (7.8), we have $\|h - h'\|_2 = \rho^r$. Let $A := \mathcal{K}_\varepsilon(h \otimes \overline{h})$ and let $\mathrm{Spec}(A) = \{\mu_1 \geq \mu_2 \geq \cdots\}$. Reusing the bound (7.19) from Theorem 7.15 we have $\big(\sum_{i=1}^{|\mathrm{Z}|} ||c_{\chi_i}|^2 - \mu_i|^2\big)^{1/2} \leq 14\rho^r$ where crucially here $c_\chi = \|g_\chi\|_2$ if $\chi \in S_\rho$ and $c_\chi = 0$ otherwise. In particular, $\big(\sum_{i=1}^{|S_\rho|} |\|g_{\chi_i}\|_2^2 - \mu_i|^2\big)^{1/2} \leq 14\rho^r$. By the triangle inequality, we deduce that $\big(\sum_{i=1}^{|S_\rho|} |\mu_i - \lambda_i|^2\big)^{1/2} \leq 15\rho^r$.

To conclude that this gives the desired bijection, we want to ensure that all the μ_i are larger than ρ. By Remark 7.14, we have $\lambda_{|S_\rho|} \geq \rho + \rho_0^3/30$. By our assumption on r and ρ, we know that $\rho_0^3/30 \geq \rho^3/30 > 15\rho^r$. Therefore $\mu_{|S_\rho|} \geq \rho$ and the result follows. □

Remark 7.19. It follows from Lemma 7.2 combined with Lemma 7.4 that if an eigenvalue of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ is large and sufficiently separated from the rest, then the projection of f to the corresponding eigenspace is very close to one of the nilspace characters g_χ. Moreover, Theorem 7.15 could be generalized to be applicable in clusters of eigenvalues. That is, if C is a submultiset of $\mathrm{Spec}_\rho\big(\mathcal{K}_\varepsilon(f \otimes \overline{f})\big)$ such that C is included in a ball of small radius separated from the rest of the eigenvalues, then the mentioned generalization of Theorem 7.15 would involve a vector h chosen in the eigenspace corresponding to C, instead of the whole space $\mathrm{Eigen}_\rho\big(\mathcal{K}_\varepsilon(f \otimes \overline{f})\big)$. Another possible extension of Theorem 7.15 involves not assuming that $|\mathrm{Spec}_{\rho^q}(A)| = |S_\rho|$. In this case, we could prove an analogous result but where the map $J : \mathrm{Spec}_{\rho^q}(A) \to S_\rho$ is only an injection. As we shall not use these extensions in the paper, we omit their proofs.

Finally, note that it follows from (7.19) that, for any $h \in \mathrm{Eigen}_\rho\big(\mathcal{K}_\varepsilon(f \otimes \overline{f})\big)$ with $\|h\|_2 \leq 1$, we have $|\mathrm{Spec}_{\rho^q}\big(\mathcal{K}_\varepsilon(h \otimes \overline{h})\big)| \leq |S_\rho|$. Indeed, choosing h in $\mathrm{Eigen}_\rho\big(\mathcal{K}_\varepsilon(f \otimes \overline{f})\big)$ generates a gap in the set of the eigenvalues of $\mathcal{K}_\varepsilon(h \otimes \overline{h})$, that is, there are at most $|S_\rho|$ such eigenvalues which can be larger than $14\rho^r$ and all the others must be small. In particular, we can combine this fact with Lemma 7.18 to deduce that for $h = P_\rho(f)$, we have $\mathrm{Spec}_{\rho^q}\big(\mathcal{K}_\varepsilon(h \otimes \overline{h})\big) = \mathrm{Spec}_\rho\big(\mathcal{K}_\varepsilon(h \otimes \overline{h})\big)$.

Proof of Theorem 1.6. This follows from Theorem 7.15 applied with $q = 7$ and $r = 11$. Indeed, this application gives us that each unit eigenvector v has a corresponding 2-step nilspace character g_χ such that $\|v - g_\chi/\langle f, v\rangle\|_2 \leq 56\rho^{7/2}$. The claimed quadratic character is then $g := g_\chi/\langle f, v\rangle$. To obtain the complexity and precision bounds for g, note that, by Remark 4.20, the function g_χ is an (R, σ)-character of order 2 with $R = O_{\mathrm{X}, \|F\|_{\mathrm{sum}}}(\sigma^{-O_{\mathrm{X}}(1)})$. Letting $\sigma = c\rho^2$ for a small enough constant $c > 0$, we have that g is a $(O_\rho(1), \rho)$-quadratic character and $|\langle f, g\rangle| \geq |\langle f, v\rangle| - \|v - g\|_2 \geq \sqrt{\rho/2} - 56\rho^{7/2} \geq \sqrt{\rho/4}$. □

7.2.1. Handling clustered eigenvalues

Suppose that we are working under the assumptions of Theorem 7.15, and for the given 1-bounded function $f : \mathrm{Z} \to \mathbb{C}$ we want to recover the individual g_χ, $\chi \in S_\rho$ in the decomposition of f involved in that theorem. As explained earlier (in particular, combining Remark 7.16 and Lemma 7.18), if the large spectrum of $\mathcal{K}_\varepsilon(f \otimes \overline{f})$ is sufficiently separated, then the desired recovery is straightforward using this spectrum, by Theorem 7.15. Hence, the last difficulty that we face for this recovery is the possibility that, for $h = P_\rho(f)$, the separation assumptions for the spectrum in Theorem 7.15 do not hold. Our goal now is to show that, in fact, if we select another function h *randomly* in $\mathrm{Eigen}_\rho\big(\mathcal{K}_\varepsilon(f \otimes \overline{f})\big)$ according

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

to a simple normal distribution (specified below), then with high probability the resulting spectrum $\mathrm{Spec}_{\rho^q}\big(\mathcal{K}_\varepsilon(h\otimes\overline{h})\big)$ will satisfy the conditions for Theorem 7.15 to be applicable, and thus enable the recovery of the g_χ.

Let us now detail the probabilistic argument. Let $\gamma\in[0,1]$ be some fixed parameter to be determined later. Let $A_1,B_1,A_2,B_2,\dots,A_t,B_t$ be i.i.d. real normal random variables with mean 0 and variance $\gamma/(2t)$. Then, for every $j\in[t]$, let X_j be the complex random variable A_j+iB_j, and let

$$X:=(X_1,\dots,X_t) \tag{7.20}$$

be the resulting complex Gaussian random vector in $\mathbb{C}^t$ (for background on such vectors see, e.g., [20]). Note that $\mathbb{E}\|X\|_{\ell^2}=\mathbb{E}(|X_1|^2+\cdots+|X_t|^2)=\mathbb{E}(A_1^2+B_1^2+\cdots+A_t^2+B_t^2)=\gamma$.

Letting $v_1,\dots,v_t$ be an orthonormal basis of complex eigenvectors for $\mathrm{Eigen}_\rho\big(\mathcal{K}_\varepsilon(f\otimes\overline{f})\big)$, we then define the random vector

$$h:=\textstyle\sum_{j\in[t]}X_jv_j. \tag{7.21}$$

Note that $\mathbb{E}\|h\|_2^2=\gamma$. We want to prove that, with high probability, the large spectrum of $\mathcal{K}_\varepsilon(h\otimes\overline{h})$ satisfies the properties required in Theorem 7.15. To do so, we plan to use Proposition 7.17, which requires us to prove that h is sufficiently close to a vector of the form $\sum_{\chi\in S_\rho}c_\chi u_\chi$ such that the set $\{|c_\chi|^2\}_{\chi\in S_\rho}$ is sufficiently separated. Hence, we would like to express h in terms of the u_{χ_i} rather than the v_j (note that here we are using that $t=|S_\rho|=|\mathrm{Spec}_\rho(f\otimes\overline{f})|$). However, the subspace of $\mathbb{C}^{\mathbb{Z}}$ generated by $\{u_{\chi_i}\}_{i\in[t]}$ may not be the same as the one generated by the $\{v_j\}_{j\in[t]}$. Fortunately, these two subspaces are *close* to each other, as we detail in the proof of Theorem 7.24 below. Let P denote the orthogonal projection in $\mathbb{C}^{\mathbb{Z}}$ onto the subspace spanned by $\{u_\chi:\chi\in S_\rho\}$. Let $W\in\mathbb{C}^{t\times t}$ be the matrix defined by

$$\forall\, j\in[t],\ P(v_j)=\sum_{i\in[t]}W_{i,j}u_{\chi_i}. \tag{7.22}$$

As we shall establish in the proof of Theorem 7.24, we then have

$$\forall j\in[t],\ v_j=\textstyle\sum_{i\in[t]}W_{i,j}u_{\chi_i}+d_j \text{ where } d_j:=v_j-P(v_j) \text{ satisfies } \|d_j\|_2\le\rho^r. \tag{7.23}$$

Remark 7.20. Note that the matrix W here is not a Z-matrix and we do not use the normalization for it as we do for Z-matrices (recall Definition 2.3). In particular, given a vector $x\in\mathbb{C}^t$, we use the standard definition $(Wx)_i=\sum_{j\in[t]}W_{i,j}x_j$. Recall that $\|x\|_{\ell^2}^2$ denotes the Euclidean norm of $x\in\mathbb{C}^t$, that is $\|x\|_{\ell^2}^2=\sum_{i\in[t]}|x_i|^2$. In particular, by (7.21) we have $\|h\|_2=\|X\|_{\ell^2}$.

Now, by definition of the matrix W above, and (7.23), for any vector $x\in\mathbb{C}^t$ we have

$$\textstyle\sum_j x_jv_j=\sum_i(\sum_j W_{i,j}x_j)u_{\chi_i}+\sum_j x_jd_j=\sum_i(Wx)_iu_{\chi_i}+\sum_j x_jd_j, \tag{7.24}$$

where all sums and averages here are for i and j in $[t]$. This will imply that the matrix W is *nearly unitary*. To justify this we will use the following independent result.

Lemma 7.21. *Let $N\in\mathbb{N}$, let $t\in[N]$, let $\kappa\in(0,1]$, and $\tau\in(0,1/(2t))$. Let $\{v_i\}_{i\in[t]}\subset\mathbb{C}^N$ be an orthonormal set of vectors, let $\{u_i\}_{i\in[t]}\subset\mathbb{C}^N$ be a set of unit vectors such that for every $i\ne j$ we have $|\langle u_i,u_j\rangle|\le\tau$, and let $\{d_i\}_{i\in[t]}\subset\mathbb{C}^N$ be such that for every $i\in[t]$ we have $\|d_i\|_2\le\kappa$. Let W be a matrix in $\mathbb{C}^{t\times t}$ such that for every $j\in[t]$ we have $v_j=\sum_{i\in[t]}W_{i,j}u_i+d_j$. Then, for every $x\in\mathbb{C}^t$ we have*

$$\big|\|Wx\|_{\ell^2}^2-\|x\|_{\ell^2}^2\big|\le(3\kappa t+12\tau t^2)\|x\|_{\ell^2}^2. \tag{7.25}$$

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Proof. By the analogue of (7.24) for v_i, u_i, we have

$$\|\sum_i (Wx)_i u_i\|_{\ell^2}^2 = \|\sum_j x_j (v_j - d_j)\|_{\ell^2}^2. \tag{7.26}$$

We now expand both sides.

The right side of (7.26) equals $\langle \sum_j x_j(v_j - d_j), \sum_{j'} x_{j'}(v_{j'} - d_{j'})\rangle = \sum_{j,j'} x_j \overline{x_{j'}}(\langle v_j, v_{j'}\rangle - \langle v_j, d_{j'}\rangle - \langle d_j, v_{j'}\rangle + \langle d_j, d_{j'}\rangle)$. As the v_i are orthonormal, we have $\sum_{j,j'} x_j \overline{x_{j'}}\langle v_j, v_{j'}\rangle = \|x\|_{\ell^2}^2$. By Cauchy-Schwarz, and recalling that $\kappa \le 1$, we have $\big|\|\sum_j x_j(v_j - d_j)\|_2^2 - \|x\|_{\ell^2}^2\big| \le \sum_{j,j'} |x_j x_{j'}| 3\kappa = 3\kappa\big(\sum_j |x_j|\big)^2 \le 3\kappa t \|x\|_{\ell^2}^2$.

The left-hand side of (7.26) is $\|\sum_i (Wx)_i u_i\|_2^2 = \sum_{i,i'} (Wx)_i \overline{(Wx)_{i'}}\langle u_i, u_{i'}\rangle$, which equals $\|Wx\|_{\ell^2}^2 + \sum_{i\neq i'} (Wx)_i \overline{(Wx)_{i'}}\langle u_i, u_{i'}\rangle$. Hence $\big|\|Wx\|_{\ell^2}^2 - \|\sum_i (Wx)_i u_i\|_2^2\big| \le \tau(\sum_i |(Wx)_i|)^2$, which is at most $\tau t \sum_i |(Wx)_i|^2 \le \tau t \|W\|_{\ell^2}^2 \|x\|_{\ell^2}^2$.

Finally, let us prove that $\|W\|_{\ell^2}^2 \le 12t$. Note that for every $j \in [t]$ we have

$$1 = \langle v_j, v_j\rangle = \sum_{i,i'} W_{i,j}\overline{W_{i',j}}\langle u_i, u_{i'}\rangle + \langle \sum_i W_{i,j} u_i, d_j\rangle + \langle d_j, \sum_i W_{i,j} u_i\rangle + \langle d_j, d_j\rangle.$$

As $\|\sum_i W_{i,j} u_i\|_2 \le \|v_j\|_2 + \|d_j\|_2 \le 1 + \kappa \le 2$, by Cauchy-Schwarz we have $\big|\langle \sum_i W_{i,j} u_i, d_j\rangle + \langle d_j, \sum_i W_{i,j} u_i\rangle + \langle d_j, d_j\rangle\big| \le 4\kappa + \kappa^2 \le 5\kappa$. To control the other term, note that

$$\sum_{i,i'} W_{i,j}\overline{W_{i',j}}\langle u_i, u_{i'}\rangle = \sum_i |W_{i,j}|^2 + \sum_{i\neq i'} W_{i,j}\overline{W_{i',j}}\langle u_i, u_{i'}\rangle,$$

so $|1 - \sum_i |W_{i,j}|^2| \le 5\kappa + \tau(\sum_i |W_{i,j}|)^2 \le 5\kappa + \tau t \sum_i |W_{i,j}|^2$. We conclude that for every $j \in [t]$ we have $(1 - \tau t)\sum_i |W_{i,j}|^2 \le 1 + 5\kappa$. Summing this over j, we deduce $(1 - \tau t)\|W\|_{\ell^2}^2 \le (1 + 5\kappa)t$, whence $\|W\|_{\ell^2}^2 \le \frac{t}{1-\tau t}(1 + 5\kappa) \le 12t$, as claimed.

Combining the bounds above, we deduce $\big|\|x\|_{\ell^2}^2 - \|Wx\|_{\ell^2}^2\big| \le \big|\|\sum_j x_j(v_j - d_j)\|_2^2 - \|x\|_{\ell^2}^2\big| + \big|\|Wx\|_{\ell^2}^2 - \|\sum_i (Wx)_i u_i\|_2^2\big| \le 3\kappa t\|x\|_{\ell^2}^2 + 12\tau t^2 \|x\|_{\ell^2}^2$ and (7.25) follows. □

Recall from (7.20) the definition of the complex Gaussian vector $X = (X_1, \dots, X_t)$, and from (7.21) the random function h. Our aim is to apply Theorem 7.15 to this h. To do so, we need to ensure that the leading eigenvalues of $\mathcal{K}_\varepsilon(h \otimes \overline{h})$ are sufficiently separated, and for this we need more information on these eigenvalues. Now Proposition 7.17 tells us that if we can write $h = \sum_{\chi \in D} c_\chi u_\chi + \mathcal{E}$ where $\|\mathcal{E}\|_2$ is small, then the eigenvalues of $\mathcal{K}_\varepsilon(h \otimes \overline{h})$ will be approximately the numbers $|c_\chi|^2$. But we *can* write h this way, indeed we have $h = \sum_j X_j v_j = \sum_i (WX)_i u_{\chi_i} + (h - P(h))$ where $\|h - P(h)\|_2$ will be shown to be small in the proof of Theorem 7.24 (as mentioned already in (7.23)). Hence the coefficients c_{χ_i} that we thus obtain are the coordinates of the image of X under the matrix W. Let $Y = (Y_1, \dots, Y_s) := WX$ be this image, so that the aforementioned eigenvalues are approximately the numbers $|Y_i|^2$. Now, if W were unitary, then, by standard results, the vector Y would also be complex Gaussian with the same distribution[18] as X, and this would imply (as we will show below in Lemma 7.23) that with high probability the eigenvalue separation that we need does indeed hold. However, Lemma 7.21 only gives that W is *nearly unitary*. Fortunately, any such matrix can be efficiently approximated by a unitary matrix, as the next result shows.

Lemma 7.22. *Let $t \ge 2$, $\delta \le 1/(24t)$, and let W be a matrix in $\mathbb{C}^{t\times t}$ satisfying*

$$\forall\, x \in \mathbb{C}^t, \quad \big|\|Wx\|_{\ell^2}^2 - \|x\|_{\ell^2}^2\big| \le \delta \|x\|_{\ell^2}^2. \tag{7.27}$$

Then there is a unitary matrix $M \in \mathbb{C}^{t\times t}$ such that $W = M + E$ where $\|E\|_\infty \le 30\delta t^{1/2}$.

[18] To see this, note that we can view the distribution of the (complex) random variable X as a $2t$-dimensional (real) random multivariate normal $X' = (\mathrm{Re}(X_i), \mathrm{Im}(X_i))_i$. Moreover, the action of a unitary W on the complex vector X can be viewed as the action of an orthonormal (real) matrix W' on X' by [20, Theorem 2.1 and Theorem 2.4]. The fact that the covariance matrix of X' is a multiple of the identity combined with standard results, for example, [36, Chapter 5, Theorem 2.2], implies that the distribution of $W'X'$ is the same as that of X'.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Here for a matrix $E \in \mathbb{C}^{t\times t}$ recall that $\|E\|_\infty := \max_{i,j} |E_{i,j}|$.

Proof. Let $\mathcal{I}_t$ denote the identity $t \times t$ matrix. Firstly we claim that (7.27) implies

$$\|W^*W - \mathcal{I}_t\|_\infty \le 6\delta. \tag{7.28}$$

This can be proved by following a standard argument proving that $x \mapsto Ax$ is a Euclidean isometry only if the matrix A is unitary (e.g., the proof of $(g) \Rightarrow (a)$ in [41, Theorem 2.1.4]). More precisely, let $x = z + w$ for any z, w, let $y = Wx$, let $A = W^*W$, and note that[19]

$$\begin{aligned} x^*x &= z^*z + w^*w + 2\,\mathrm{Re}(z^*w) \\ y^*y &= z^*Az + w^*Aw + 2\,\mathrm{Re}(z^*Aw). \end{aligned}$$

By assumption we have $|x^*x - y^*y| \le \delta x^*x$, $|z^*z - z^*Az| \le \delta z^*z$, and $|w^*w - w^*Aw| \le \delta w^*w$. Using this, we deduce that for every w, z we have

$$\begin{aligned} |\,\mathrm{Re}(z^*Aw) - \mathrm{Re}(z^*w)| &\le \tfrac{1}{2}(|x^*x - y^*y| + |z^*z - z^*Az| + |w^*w - w^*Aw|) \\ &\le \tfrac{\delta}{2}\big(w^*w + x^*x + z^*z\big) = \delta\big(w^*w + z^*z + \mathrm{Re}(w^*z)\big). \end{aligned} \tag{7.29}$$

Now, applying (7.29) with $z = e_i$ and $w = e_j$, we deduce that $|\,\mathrm{Re}(A_{i,j}) - \mathbf{1}(i = j)| \le 3\delta$. Moreover, applying (7.29) with $z = e_i$ and $w = ie_j$, we deduce that $|\,\mathrm{Re}(iA_{i,j})| = |\,\mathrm{Im}(A_{i,j})| \le 3\delta$. The last two inequalities imply (7.28). We can now apply [1, Lemma 7.2] using that $\delta < 1/(24t)$, thus obtaining a unitary matrix M such that $\|W - M\|_\infty \le 30\delta t^{1/2}$, as claimed. □

Recall that $Y = WX$ for X defined in (7.20). Our goal now is to prove that with high probability the coordinates $|Y_i|^2$ are sufficiently separated and large. To this end, the decomposition $W = M + E$ provided by Lemma 7.22 is useful, because we then have $Y = MX + EX$, where MX has the same distribution as X thanks to the fact that M is unitary, and EX is small. Thus, we have reduced our task to proving the following result, which we will apply to $V = MX$.

Lemma 7.23. *There exists an absolute constant $C > 0$ such that the following holds. Let $t \in \mathbb{N}$ and $\delta, \sigma, \gamma > 0$. Let $A_1, B_1, \dots, A_t, B_t$ be i.i.d. real normal variables with mean 0 and variance $\gamma/(2t)$, and let V be the complex Gaussian vector in $\mathbb{C}^t$ with j-th coordinate $V_j := A_j + iB_j$, $j \in [t]$. Let $Q(\delta, \gamma, \sigma, t)$ denote the following event: we have $\|V\|_{\ell^2} \le 1$ and for every pair $j \ne k$ in $[t]$ we have $||V_j|^2 - |V_k|^2| > 2\delta$ and $|V_j|^2 > \sigma$. Then*

$$\mathbb{P}(Q(\delta, \gamma, \sigma, t)) \ge 1 - t^2\big(1 - e^{-4\delta t/(\gamma\pi)}\big)^{1/2} - t\big(1 - e^{-t\sigma/\gamma}\big) - 2e^{-C(1-\sqrt{\gamma})^2 t/\gamma}. \tag{7.30}$$

Proof. Firstly we obtain a small upper bound on the probability that for some pair $j \ne k$ we have $||V_j|^2 - |V_k|^2| \le 2\delta$. By the union bound, this probability is at most $\binom{s}{2}$ times the probability, for any *fixed* pair j, k, that $||V_j|^2 - |V_k|^2| \le 2\delta$. Without loss of generality we fix $j = 1$, $k = 2$.

Since $\big||V_1|^2 - |V_2|^2\big| = |A_1^2 - A_2^2 + B_1^2 - B_2^2|$, we want a small upper bound for the probability of the event $E : |A_1^2 - A_2^2 + B_1^2 - B_2^2| \le 2\delta$. Let $A = A_1^2 - A_2^2$ and $B = B_1^2 - B_2^2$, and denote by f_A, f_B the probability density functions of A, B respectively. Since the probability density function f_{A+B} is the convolution $f_A * f_B$, we have

$$\mathbb{P}(E) = \int_{-2\delta}^{2\delta} f_A * f_B(z)\,\mathrm{d}z = \int_{-2\delta}^{2\delta}\int_{-\infty}^{\infty} f_A(x) f_B(z - x)\,\mathrm{d}x\,\mathrm{d}z = \int_{-\infty}^{\infty} f_A(x)\int_{[-2\delta, 2\delta]-x} f_B(y)\,\mathrm{d}y,$$

where the last equality follows from the Fubini–Tonelli theorem [15, Proposition 5.2.1] and the change of variables $y = z - x$.

[19] For vectors $x, y \in \mathbb{C}^t$ we denote $\sum_{i=1}^t \overline{x_i} y_i$ by x^*y. Note that this agrees with the usual convention of regarding vectors $x \in \mathbb{C}^t$ as column vectors and thus x^* is the row vector whose entries are $(\overline{x_i})_{i\in[t]}$. A similar convention applies to the expression x^*Ay where $A \in \mathbb{C}^{t\times t}$.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Note that $B = (B_1 - B_2)(B_1 + B_2)$ and that each of $B_1 - B_2$, $B_1 + B_2$ is a normal variable with mean 0. Hence B follows a *product distribution*, with $f_B(y)$ being the value at $|y|$ of a modified Bessel function of the second kind. From this we just use the consequence that, for every x we have $\int_{[-2\delta,2\delta]-x} f_B(y)\,\mathrm{d}y \le \int_{[-2\delta,2\delta]} f_B(y)\,\mathrm{d}y = \mathbb{P}(|B_1^2 - B_2^2| \le 2\delta)$. As $|B_1^2 - B_2^2| = |B_1 - B_2|\,|B_1 + B_2|$, this product being at most 2δ implies that either $|B_1 - B_2| \le (2\delta)^{1/2}$ or $|B_1 + B_2| \le (2\delta)^{1/2}$. Let us bound $\mathbb{P}(|B_1 - B_2| \le (2\delta)^{1/2})$ (the other event will be treated similarly). Since $B_1 - B_2$ is a normal random variable with mean 0 and variance γ/s, the variable $Z := (B_1 - B_2)\sqrt{s/\gamma}$ is standard normal, whence $\mathbb{P}(|B_1 - B_2| \le \sqrt{2\delta}) = \mathbb{P}(|Z| \le \sqrt{2\delta t/\gamma}) = \mathrm{erf}(\sqrt{\delta t/\gamma})$. By standard estimates for the *error function* erf, we have $\mathrm{erf}(\sqrt{\delta t/\gamma}) \le \big(1 - e^{-4\delta t/(\gamma\pi)}\big)^{1/2}$. The same upper bound applies to $\mathbb{P}(|B_1 + B_2| \le \sqrt{2\delta})$.

We conclude that the probability that some pair V_j, V_k satisfies $\big||V_j|^2 - |V_k|^2\big| \le 2\delta$ is at most

$$\tbinom{t}{2}\mathbb{P}(E) \le \tbinom{t}{2}\int_{-\infty}^{\infty} f_A(x)\ \mathbb{P}(|B_1^2 - B_2^2| \le 2\delta)\ \mathrm{d}x \le t^2\big(1 - e^{-4\delta t/(\gamma\pi)}\big)^{1/2}. \tag{7.31}$$

Next, note that for every $j \in [t]$ we have $\mathbb{P}(|V_j|^2 \le \sigma) = \mathbb{P}(A_j^2 + B_j^2 \in [0,\sigma])$. Since $A_j\sqrt{2t/\gamma}, B_j\sqrt{2t/\gamma}$ are both standard normal, the variable $2t(A_j^2 + B_j^2)/\gamma$ follows a χ^2-distribution with 2 degrees of freedom, so by standard results the cumulative distribution function of this variable is $\mathbb{P}(2t(A_j^2 + B_j^2)/\gamma \le x) = 1 - e^{-x/2}$. Therefore we have $\mathbb{P}(|V_j|^2 \le \sigma) = \mathbb{P}(2t(A_j^2 + B_j^2)/\gamma \le 2t\sigma/\gamma) = 1 - e^{-t\sigma/\gamma}$, so $\mathbb{P}(\exists\, j \in [t] \text{ s.t. } |V_j^2| \le \sigma) \le t(1 - e^{-t\sigma/\gamma})$.

Finally, we need to bound the probability that $\|V\|_{\ell^2} > 1$. By [72, Theorem 3.1.1] we have that $\big\|\|X/\sqrt{\gamma/2t}\|_{\ell^2} - \sqrt{2t}\big\|_{\psi_2} \le C$ for some absolute constant C (see [72, Definition 2.5.6] for the definition of the sub-Gaussian norm $\|\cdot\|_{\psi_2}$). Hence, $\big\|\|X\|_{\ell^2} - \sqrt{\gamma}\big\|_{\psi_2} \le C\sqrt{\gamma/2t}$. By [72, Proposition 2.5.2] there exists another absolute constant $C' > 0$ such that for all $x \ge 0$, $\mathbb{P}\big(|\|X\|_{\ell^2} - \sqrt{\gamma}| \ge x\big) \le 2e^{-C'x^2t/\gamma}$. Letting $x := 1 - \sqrt{\gamma}$ we conclude that $\mathbb{P}\big(\|X\|_{\ell^2} \ge 1\big) \le 2e^{-C'(1-\sqrt{\gamma})^2 t/\gamma}$. □

Now we have all the necessary ingredients to prove the following result.

Theorem 7.24. *For any $q > 6$ and $r > q + 3$, there exists $C_{q,r} > 0$ such that the following holds.*[20] *Let $\rho_0 \in (0, C_{q,r}]$. Under the assumptions of Theorems 7.12 and 7.15, recall that, for a 1-bounded function $f : \mathrm{Z} \to \mathbb{C}$ there exists $\rho \in [\rho_0/2, \rho_0]$ such that we have a decomposition $f = \sum_{\chi\in S_\rho} g_\chi + f_r + f_e$ where $\|f_r\|_{U^3} \le 2\rho^{3/8}$ and $\|f_e\|_2 \le (\rho_0/2)^r$. Let $t := |S_\rho|$ and, for $i \in [t]$, let A_i, B_i be independent normal random variables with mean 0 and variance $\rho/(2t)$. Let X be the complex random vector $X := (A_1 + iB_1, \dots, A_t + iB_t)$. Let $\{v_j : j \in [t]\}$ be an orthonormal basis of $\mathrm{Eigen}_\rho\big(\mathcal{K}_\varepsilon(f \otimes \overline{f})\big)$ and let h be the random function $\sum_{j\in[t]} X_j v_j$. Then, with probability $1 - o_{\rho\to 0}(1)$, we have that $\|h\|_2 \le 1$, that $\mathrm{Spec}_{\rho^q}\big(\mathcal{K}_\varepsilon(h \otimes \overline{h})\big)$ is ρ^q-separated, and that $|\mathrm{Spec}_{\rho^q}\big(\mathcal{K}_\varepsilon(h \otimes \overline{h})\big)| = |S_\rho|$.*

Proof. Recall that $t := |S_\rho|$ is equal to $|\mathrm{Spec}_\rho\big(\mathcal{K}_\varepsilon(f \otimes \overline{f})\big)|$. We first prove the following claim:

$$\begin{aligned}&\text{There is a unitary matrix } M \in \mathbb{C}^{t\times t} \text{ such that } \forall\, x = (x_1, \dots, x_t) \in \mathbb{C}^t, \text{ we have}\\ &\textstyle\sum_i x_i v_i = \sum_i (Mx)_i u_{\chi_i} + d \text{ where } \|d\|_2 \le 4 \cdot 10^6 \rho^{r-3}\|x\|_{\ell^2}.\end{aligned} \tag{7.32}$$

To prove this, let $w = \sum_i x_i v_i$ and note that $\|w\|_2 = \|x\|_{\ell^2}$. Let T_1 and T_2 be as defined in (7.13). By (7.11), we know that $\mathcal{P}_{T_1,\rho}$ is precisely the projection to the space spanned by the u_χ for $\chi \in S_\rho$. By our choice of parameters in Theorems 7.12 and 7.15, we have that $\|\mathcal{P}_{T_1,\rho} - \mathcal{P}_{T_2,\rho}\| \le \rho^r$. As $w \in \mathrm{Eigen}_\rho\big(\mathcal{K}_\varepsilon(f \otimes \overline{f})\big)$, we have $\mathcal{P}_{T_2,\rho}(w) = w$, whence $\|\mathcal{P}_{T_1,\rho}(w) - w\|_2 = \|\mathcal{P}_{T_1,\rho}(w) - \mathcal{P}_{T_2,\rho}(w)\|_2 \le \rho^r\|w\|_2 = \rho^r\|x\|_{\ell^2}$. Let $W \in \mathbb{C}^{t\times t}$ be the matrix defined in (7.22), thus for every $j \in [t]$ we have $\mathcal{P}_{T_1,\rho}(v_j) = \sum_{i\in[t]} W_{i,j}u_{\chi_i}$. Then, for every $j \in [t]$, we have $v_j = \sum_{i\in[t]} W_{i,j}u_{\chi_i} + d_j$ where $\|d_j\|_2 \le \rho^r$ (as we announced in (7.23)). Now we want to apply Lemma 7.21 to $\{v_i\}_{i\in[t]}$ and $\{u_\chi\}_{\chi\in S_\rho}$ with $\kappa := \rho^r$, $\tau := \rho^{4+r}$, and $t = |S_\rho|$. Since $|S_\rho| \le |S| \le \mathrm{Q}(\eta, m)$, using condition (ii) from Theorem 7.12, we have that

[20] We can let $C_{q,r} = \min(B_{q,r}, (1/295200)^{1/(r-2)}, (1/(8 \cdot 10^7))^{1/(r-q-3)})$, for $B_{q,r}$ as defined in Theorem 7.15.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

$\tau \le 1/(2t)$. Hence, Lemma 7.21 combined with the bound $t \le 10/\rho$ (given by Theorem 7.12) tells us that $\left|\|Wx\|_{\ell^2}^2 - \|x\|_{\ell^2}^2\right| \le 1230\rho^{r-1}\|x\|_{\ell^2}^2$. Next, by Lemma 7.22 (note that we can apply this because by assumption on ρ_0 it follows that $24t\delta \le 1$) we have that $W = M + E$ where M is unitary and $\|E\|_\infty \le 30\delta t^{1/2}$ where $\delta = 1230\rho^{r-1}$. In particular, note that

$$\|Ex\|_\infty \le \max_i \|E_i\|_{\ell^2}\|x\|_{\ell^2} \le t^{1/2}\|E\|_\infty\|x\|_{\ell^2} \le 30\delta t\|x\|_{\ell^2}.$$

Thus $\|\sum_i (Ex)_i u_{\chi_i}\|_2 \le t\|Ex\|_\infty \le 30\delta t^2\|x\|_{\ell^2}$. Letting $d := (\mathcal{P}_{T_1,\rho}(w) - \mathcal{P}_{T_2,\rho}(w)) + \sum_i (Ex)_i u_{\chi_i}$ the claim (7.32) follows.

Now let h be the random function $\sum_{i\in[t]} X_i v_i$ in the theorem. By (7.32), we have that $h = \sum_i (MX)_i u_{\chi_i} + d$. Let V denote the vector MX. We apply now Lemma 7.23 to V with parameters $t = |S_\rho| \le 10/\rho$, $\delta := \rho^q$, $\sigma := 2\rho^q$, $\gamma := \rho$. Let us check that under these conditions, the right side of (7.30) is $1 - o_{\rho\to 0}(1)$. The first term is $t^2\left(1 - e^{-4\delta t/(\gamma\pi)}\right)^{1/2} \le O\left(\frac{1}{\rho^4}(1 - e^{-40\rho^q/(\pi\rho^2)}\right)^{1/2}$. By L'Hôpital's rule it is easily checked that this quantity tends to 0 as $\rho \to 0$ if $q > 6$. The second term is $t\left(1 - e^{-t\sigma/\gamma}\right) = O\left(\frac{1}{\rho}(1 - e^{-20\rho^q/\rho^2})\right)$, which tends to 0 as $\rho \to 0$ if $q > 3$ (also by L'Hôpital's rule). The third and final term is $2e^{-C(1-\sqrt{\gamma})^2 t/\gamma}$. Note that we can assume that $t \ge 1$, as otherwise the result is trivial, and we can also assume that $\sqrt{\gamma} = \sqrt{\rho} \le 1/2$. Thus $2e^{-C(1-\sqrt{\gamma})^2 t/\gamma} \le O\left(e^{-C/(4\rho)}\right) \to 0$ as $\rho \to 0$.

In particular, as $\|V\|_{\ell^2} \le 1$ and M is unitary, we have $\|X\|_{\ell^2} \le 1$, and $h = \sum_i V_i u_{\chi_i} + d$ where $\|d\|_2 \le 4 \cdot 10^6 \rho^{r-3}$.

Finally, we apply Proposition 7.17 to h. Note that the constants τ_1, τ_2, and τ_3 are bounded as in the proof of Theorem 7.15, and we have just proved that we can let $\tau_4 = 4 \cdot 10^6 \rho^{r-3}$. Hence, we have (7.14) with $\|E\|_2 \le 3\rho^r + 20 \cdot 10^6 \rho^{r-3} \le 3 \cdot 10^7 \rho^{r-3}$. Let λ_i be the eigenvalues of $\mathcal{K}_\varepsilon(h \otimes \overline{h})$ in descending order. Assume also that $|V_1|^2 \ge |V_2|^2 \ge \cdots$ (by relabeling if necessary). Then, by (7.15) combined with (ix), we have $\left(\sum_{i=1}^t \left||V_i|^2 - \lambda_i\right|^2\right)^{\frac{1}{2}} \le 4 \cdot 10^7 \rho^{r-3}$. In particular, note that, for all $i, j \in [t]$ with $i \ne j$ we have $\left||V_i|^2 - |V_j|^2\right| \ge 2\rho^q$. As $4 \cdot 10^7 \rho^{r-3} \le \rho^q/2$, we have $|\lambda_i - \lambda_j| \ge \rho^q$ for all $i, j \in [t]$, $i \ne j$. Similarly, since for all $i \in [t]$ we have $|V_i|^2 \ge 2\rho^q$, it follows that $\lambda_i \ge \rho^q$ for all $i \in [t]$. By (7.15), we have $\lambda_i \le 4 \cdot 10^7 \rho^{r-3} < \rho^q$ for all $i > t$, and the result follows. □

Remark 7.25 (Uniform spherical distribution). In Theorem 7.24 we chose the random vector $X \in \mathbb{C}^t$ to follow a multivariate (complex) normal distribution and we showed that, with high probability, the function $h = \sum_i X_i v_i$ satisfies the desired properties. However, note that if instead we let $X' := X/\|X\|_{\ell^2}$, then X' follows a uniform distribution in the unit sphere of $\mathbb{C}^t$ (e.g., by [72, Exercise 3.3.7]). Moreover, if we let $h' := \sum_i X_i' v_i$ and the event $Q(\rho^q, \rho, 2\rho^q, t)$ from Lemma 7.23 holds for X, then we also have that $\|h'\|_2 \le 1$, $\mathrm{Spec}_{\rho^q}(\mathcal{K}_\varepsilon(h' \otimes \overline{h'}))$ is ρ^q-separated, and $|\mathrm{Spec}_{\rho^q}(\mathcal{K}_\varepsilon(h' \otimes \overline{h'}))| = |S_\rho|$. To see this, let $V' := MX'$ and note that $V' = V/\|X\|_{\ell^2}$ where $V = MX$ (as in the proof of Theorem 7.24). Thus, if $h' = \sum_i V_i' u_{\chi_i} + d'$ we have that h' satisfies the same conditions as h, that is, we have $\|d'\|_2 \le 4 \cdot 10^6 \rho^{r-3}$ (by (7.32)), for all $i \in [t]$ we have $|V_i'|^2 = \frac{|V_i|^2}{\|X\|_{\ell^2}^2} \ge 2\rho^q$, and for all $i \ne j$ we have $\left||V_i'|^2 - |V_j'|^2\right| = \frac{||V_i|^2 - |V_j|^2|}{\|X\|_{\ell^2}^2} \ge \rho^q$. Therefore, the function h' also satisfies the conclusion of Theorem 7.24 and we could have directly chosen X to be uniformly distributed in the unit sphere.

A. Metrics on CFR nilspaces derived from the manifold structure

Every k-step CFR nilspace X is a metric space [7, Remark 2.1.3 (*ii*)]. However, there can be different metrics on X that generate the same original topology, and some of these metrics are better than others for our purposes in this paper. In particular, we need to work with a family of metrics on k-step CFR nilspaces X which ensure sufficient control on Lipschitz constants of various functions on X (including complex-valued functions but also nilspace translations on X), and also on Lipschitz constants of compositions of such functions with morphisms between different nilspaces. Such properties are used for instance in the proof of Theorem 4.19.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Below we describe metrics with such properties, eventually obtaining as main results Proposition A.13, Corollary A.15 and Lemma A.17, giving the desired control on Lipschitz constants.

Let us first give an example showing some metrics that we want to avoid in this respect.

Example A.1. Consider the circle $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, which can be regarded as a 1-step CFR nilspace. If for $x \in \mathbb{R}$ we let $|x|_{\mathbb{T}} := \min_{n\in\mathbb{Z}}\{|x-n|\}$, then we obtain the well-known metric $d(x+\mathbb{Z}, y+\mathbb{Z}) := |x-y|_{\mathbb{T}}$ that generates the usual topology on $\mathbb{T}$. There are other metrics, like $d_\alpha(x+\mathbb{Z}, y+\mathbb{Z}) := |x-y|_{\mathbb{T}}^\alpha$ for $\alpha \in (0,1)$, that generate the same topology. The problem with these metrics is that, for example, the identity map $(\mathbb{T}, d) \to (\mathbb{T}, d_\alpha)$ is not Lipschitz. Hence, if we let $\mathrm{X} = \mathcal{D}_1(\mathbb{Z}_2 \times \mathbb{T})$ we can define a (somewhat artificial, yet valid) metric $d' : \mathrm{X}\times\mathrm{X} \to \mathbb{R}_{\geq 0}$ by setting $d'((a_0,a_1),(b_0,b_1))$ to have value 1 if $a_0 \neq b_0$, value $d(a_1,b_1)$ if $a_0 = b_0 = 0$, and value $d_{1/2}(a_1,b_1)$ if $a_0 = b_0 = 1$. Then, by the previous observation, it is easy to see that the nilspace translation $\alpha \in \Theta(\mathrm{X})$ defined by $\alpha(a_0,a_1) := (a_0+1, a_1)$ is not Lipschitz.

There is a very convenient family of metrics on CFR nilspaces which avoids issues such as those in the above example. This family is naturally suggested by what is already usual practice when working in the more classical setting of compact nilmanifolds (which are particular cases of CFR nilspaces). Indeed, on nilmanifolds we can use *Riemannian metrics*, that is, metrics induced by Riemannian metric tensors defined on the nilmanifolds. It turns out that we can use a similar construction on CFR nilspaces more generally, because we can endow these spaces with a differentiable structure that makes them *smooth* (i.e., C^∞) *manifolds*. In fact, this smooth structure is available, without much added difficulty, on the more general class of (not necessarily compact) *Lie-fibered nilspaces*, that is, topological nilspaces whose structure groups are Lie groups (for the precise definition see [10, Definition 2.19]).

Theorem A.2. *Let* X *be a k-step Lie-fibered (resp. CFR) nilspace. Then* X *can be endowed with a* C^∞ *manifold (resp. compact manifold) structure such that every cube set* $\mathrm{C}^n(\mathrm{X})$ *(*$n \in \mathbb{N}$*) is also a* C^∞ *manifold.*

Proof. Our main tool here is a structure theorem for nilspaces, namely [10, Theorem 1.8], which tells us that any k-step CFR nilspace is isomorphic to a quotient space F/H where F is a k-step free nilspace (see [10, Definition 1.2]) and $H \subset \Theta(F)$ is a k-th order lattice action (see [10, Definition 1.7]). Note that this theorem is stated for CFR nilspaces, but it can be checked that its proof works exactly the same for Lie-fibered nilspaces, see [10, §6] (where, of course, we have to drop the fiber-cocompact property).

The advantage of seeing X as F/H is that F has now a natural C^∞ structure which induces a uniquely defined C^∞ structure on the quotient F/H (by definition $F = \prod_{i=1}^k \mathcal{D}_i(\mathbb{Z}^{a_i} \times \mathbb{R}^{b_i})$). Let $\pi_H : F \to F/H$ be the quotient morphism. By [57, p. 51, 5.9.5] it suffices to prove that the set $R := \{(x,gx) \in F^2 : x \in F, g \in H\}$ is an embedded closed submanifold of F^2 and the projection $p_1 : R \to F$, $(x,gx) \mapsto x$ is a submersion. Note that the fact that R is closed follows from [10, Lemma 5.25]. To prove that R is an embedded submanifold, we are going to prove the following claim:

$$\begin{aligned}&\text{For any } (x_0, g_0x_0) \in R, \text{ there exists a neighborhood } U \subset F^2 \text{ of } (x_0, g_0x_0)\\ &\text{such that } U \cap R = U \cap \{(x,gx) \in F^2 : x \in F\} \subset F^2.\end{aligned} \tag{A.1}$$

We prove this by induction on the step k.

For $k = 0$ the claim is trivially true (0-step nilspaces are 1-point spaces).

Now assume that (A.1) holds for step at most $k-1$. Let $\eta_{k-1} : \Theta(F) \to \Theta(F_{k-1})$ be the homomorphism that forgets the action on the k-th order component (see [10, paragraph before Definition 1.5]). Let $U^{k-1} \subset F_{k-1}^2$ be a neighborhood of $(\pi_{k-1}(x_0), \eta_{k-1}(g_0)(\pi_{k-1}(x_0)))$ satisfying the claim. If $(x,gx) \in R \cap (\pi_{k-1}^2)^{-1}(U^{k-1})$, we have

$$(\pi_{k-1}(x), \eta_{k-1}(g)(\pi_{k-1}(x))) = (\pi_{k-1}(x), \eta_{k-1}(g_0)(\pi_{k-1}(x))).$$

In particular, by the fiber-transitive property (see [10, Definition 1.4]), there exists $\gamma \in H_k$ such that $gx = g_0x + \gamma$. Fix an arbitrary translation-invariant distance in F (e.g., the usual Euclidean distance) and

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

let $\delta > 0$ be a constant to be fixed later. For any $r \geq 0$ let $B_r(x_0) := \{y \in F : d(y, x_0) < r\}$. As H_k is a discrete lattice in the k-th structure group of F, there exists $\epsilon > 0$ such that if $\gamma' \in B_\epsilon(\underline{0}) \cap H_k$ then $\gamma' = 0$ (where $\underline{0} \in F$ is the constant 0 in all coordinates). As g_0 is continuous, if δ is sufficiently small and $x \in B_\delta(x_0)$ then $d(g_0x, gx) < \epsilon/3$. Now in F^2 the set $U := B_\delta(x_0) \times B_{\epsilon/3}(g_0x_0) \cap (\pi^2_{k-1})^{-1}(U^{k-1})$ is an open neighborhood of (x_0, g_0x_0). Note that if $(x, gx) \in R \cap U$ then there exists $\gamma \in H_k$ such that $(x, gx) = (x, g_0x + \gamma)$. As $d(x, x_0) < \delta$ we have that $d(g_0x, g_0x_0) < \epsilon/3$. Also, by definition $d(g_0x_0, g_0x + \gamma) < \epsilon/3$. Hence $d(\gamma, \underline{0}) < 2\epsilon/3$ and we have shown that in this case $\gamma = 0$. The claim (A.1) follows.

Using (A.1), we can prove that $R \subset F^2$ is an embedded submanifold. For any $(x_0, g_0x_0) \in R$ on the neighborhood U given by (A.1), we define the chart $\varphi : (x, y) \in U \mapsto (x, y - g_0x) \in F^2$. As g_0 has an expression in terms of polynomials, it is clear that this map is a diffeomorphism. Moreover, by construction of U, the elements of $R \cap U$ are exactly those whose images under φ have $0 \in F$ in the second coordinate. Hence R is a closed embedded submanifold of F^2.

The fact that $p_1 : R \to F$ is a submersion follows easily using the charts φ described above. In fact, it can be checked that the differential is the identity at every point.

A similar argument shows that the sets $\mathrm{C}^n(\mathrm{X})$ are manifolds as well. The key observation is that the fiber-transitive property implies that if $\mathrm{c} \in \mathrm{C}^n(F)$ and $d \in \mathrm{C}^n(H_\bullet)$ is such that $\pi^{[\![n]\!]}_{k-1} \circ d \circ \mathrm{c} = \pi^{[\![n]\!]}_{k-1} \circ \mathrm{c}$, then there exists $d' \in \mathrm{C}^n(\mathcal{D}_k(H_k))$ such that $d \circ \mathrm{c} = \mathrm{c} + d'$ and $\mathrm{C}^n(\mathcal{D}_k(H_k))$ is discrete. □

Remark A.3. It also follows from [57, p. 51, 5.9.5] that the quotient map $\pi_H : F \to F/H$ is a submersion.

Corollary A.4. *Let* X *be a Lie-fibered nilspace, $n \geq 0$, and $S \subset [\![n]\!]$ be a subset with the extension property (see [6, Definition 3.1.3]). Then* $\hom(S, \mathrm{X})$ *is a C^∞ manifold.*

Proof. The proof is essentially the same as the one given above for the cube sets $\mathrm{C}^n(\mathrm{X})$, as for a fiber-transitive action we have that if $\mathrm{c} \in \hom(S, F)$ and $d \in \hom(S, H_\bullet)$ is such that $\pi^S_{k-1} \circ d \circ \mathrm{c} = \pi^S_{k-1} \circ \mathrm{c}$ then there exists $d' \in \hom(S, \mathcal{D}_k(H_k))$ such that $d \circ \mathrm{c} = \mathrm{c} + d'$ and $\hom(S, \mathcal{D}_k(H_k))$ is discrete. □

Equipping k-step nilspaces with this C^∞ manifold structure, we also obtain that the canonical factor maps are C^∞.

Proposition A.5. *Let* X *be a k-step Lie-fibered nilspace. Then the action of the last structure group* $\mathrm{Z}_k(\mathrm{X})$ *on* X *is differentiable. Moreover, this makes* X *a fiber-bundle over* X_{k-1} *in the category of C^∞ manifolds.*

Proof. Note that if $\mathrm{X} \cong F/H$ for a fiber-transitive group H, we can describe the action of the last structure group as $F/H \times \mathcal{D}_k(F)/H_k \to F/H$, $(\pi_H(x), z \mod H_k) \mapsto (\pi_H(x + z))$. By [50, Theorem 4.29] to prove that this map is C^∞ it suffices to check that the action $F \times \mathcal{D}_k(F) \to F$ is C^∞, which is trivial as this is just addition on the k-th factor of F. As the action is proper (note that the map $F/H \times \mathcal{D}_k(F) \to F/H \times F/H$ is a homeomorphism onto its image) we have that by standard results (e.g., [50, Exercise 21-6]), $\mathrm{X} \cong F/H$ is a fiber-bundle over X_{k-1}. □

Moreover, morphisms are automatically C^∞ maps.

Theorem A.6 (Automatic differentiability of morphisms). *Let* X *and* Y *be k-step Lie-fibered nilspaces and let $\varphi : \mathrm{X} \to \mathrm{Y}$ be a continuous morphism. Then φ is C^∞. Moreover, if φ is a fibration then it is a submersion.*

To prove Theorem A.6 we shall use the following result.

Proposition A.7 (Free nilspaces are projective in the category of Lie-fibered nilspaces). *Let* X, Y *be k-step Lie-fibered nilspaces, and let F be a free nilspace. Suppose that we have a morphism $\phi : F \to \mathrm{X}$ and a fibration $\varphi : \mathrm{Y} \to \mathrm{X}$. Then there exists a morphism $\psi : F \to \mathrm{Y}$ such that $\phi = \varphi \circ \psi$.*

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Proof. By [10, Theorem 4.4] there exists a free nilspace F' and a fibration $\varphi : F' \to \mathrm{Y}$. Thus, without loss of generality we can assume that Y is also a free nilspace.

Consider now the fiber product $F \times_{\mathrm{X}} \mathrm{Y} := \{(f, y) \in F \times \mathrm{Y} : \phi(f) = \varphi(y)\}$ with cube sets $\mathrm{C}^n(F \times_{\mathrm{X}} \mathrm{Y}) := \{(\mathrm{c}, \mathrm{c}') \in \mathrm{C}^n(F) \times \mathrm{C}^n(\mathrm{Y}) : \phi \circ \mathrm{c} = \varphi \circ \mathrm{c}'\}$. This construction has appeared many times in nilspace theory, see [8, Lemma 4.2] and [10, Lemma 2.18]. We leave as an exercise for the reader to check that, even though in such results both ϕ and φ are required to be fibrations, it suffices that one of them (in this case φ) is a fibration for all the arguments to work. Using these results we have the following commutative diagram:

$$\begin{array}{ccc} F \times_{\mathrm{X}} \mathrm{Y} & \xrightarrow{p_2} & \mathrm{Y} \\ \downarrow{\scriptstyle p_1} & & \downarrow{\scriptstyle \varphi} \\ F & \xrightarrow{\phi} & \mathrm{X}, \end{array}$$

where it is not difficult to prove that $p_1 : F \times_{\mathrm{X}} \mathrm{Y} \to F$ is a fibration.

Next, we claim that $F \times_{\mathrm{X}} \mathrm{Y}$ is a free nilspace. We already know by [8, Lemma 4.2] and [10, Lemma 2.18] that it is a LCH k-step nilspace, so we just need to check that the structure groups are Lie groups. We prove this just for the k-th structure group, as the claim for the other structure groups follows similarly. By [9, Proposition A.20] the k-structure group of $F \times_{\mathrm{X}} \mathrm{Y}$, that is, $\mathrm{Z}_k(F \times_{\mathrm{X}} \mathrm{Y})$, is isomorphic to $\{(z_1, z_2) \in \mathrm{Z}_k(F) \times \mathrm{Z}_k(\mathrm{Y}) : \gamma_k(z_1) = \gamma'_k(z_2)\}$ where γ_k and γ'_k are the k-th structure morphisms of ϕ and φ respectively. Note that we have the following short exact sequence $0 \to \ker(\varphi) \to \mathrm{Z}_k(F \times_{\mathrm{X}} \mathrm{Y}) \to \mathrm{Z}_k(F) \to 0$. As both $\ker(\varphi)$ and $\mathrm{Z}_k(F)$ are Lie groups, so is $\mathrm{Z}_k(F \times_{\mathrm{X}} \mathrm{Y})$. Moreover, as F is a free nilspace, we have $\mathrm{Z}_k(F) \cong \mathbb{R}^a \times \mathbb{Z}^b$ for some $a, b \in \mathbb{Z}_{\geq 0}$, and since this group is projective in the category of abelian Lie groups, we have $\mathrm{Z}_k(F \times_{\mathrm{X}} \mathrm{Y}) \cong \mathrm{Z}_k(F) \times \ker(\varphi)$. Moreover, as $\ker(\varphi)$ is a closed subgroup of $\mathrm{Z}_k(\mathrm{Y}) \cong \mathbb{R}^{a'} \times \mathbb{Z}^{b'}$ for some $a', b' \in \mathbb{Z}_{\geq 0}$, we have that $\ker(\varphi)$ is also itself a product of the form $\mathbb{R}^{a''} \times \mathbb{Z}^{b''}$ for some $a'', b'' \in \mathbb{Z}_{\geq 0}$. Hence $\mathrm{Z}_k(F \times_{\mathrm{X}} \mathrm{Y})$ is of the desired form.

To conclude, we apply [10, Lemma 8.9] to the fibration $p_1 : F \times_{\mathrm{X}} \mathrm{Y} \to F$. Hence, there exists a cross-section $s : F \to F \times_{\mathrm{X}} \mathrm{Y}$ (note that such lemma gives a free nilspace F' and a nilspace isomorphism $\tau : F \times F' \to F \times_{\mathrm{X}} \mathrm{Y}$ and thus we may let s be τ composed with the obvious inclusion of F into $F \times F'$) with $p_1 \circ s = \mathrm{id}_F$, which is a morphism. Thus, we may take $\psi = p_2 \circ s$ and the result follows. □

Proof of Theorem A.6. Let X and Y be Lie-fibered k-step nilspaces, and let $\varphi : \mathrm{X} \to \mathrm{Y}$ be a morphism. By [10, Lemma 4.4] there exist free nilspaces F, F' and fibrations $\phi : F \to \mathrm{X}$ and $\phi' : F' \to \mathrm{Y}$. We now apply Proposition A.7 to $\varphi \circ \phi$ and ϕ', obtaining a morphism $\psi : F \to F'$ which makes the following diagram commutative:

$$\begin{array}{ccc} F & \overset{\psi}{\dashrightarrow} & F' \\ \downarrow{\scriptstyle \phi} & & \downarrow{\scriptstyle \phi'} \\ \mathrm{X} & \xrightarrow{\varphi} & \mathrm{Y}. \end{array}$$

By [10, §3.1] morphisms between k-step free nilspaces are basically polynomials of degree at most k, so they are C^∞ maps. By [50, Theorem 4.29] it then follows that φ is also C^∞.

Now if φ is a fibration, then firstly, similarly as in the proof of Proposition A.7, we can construct the fiber-product of $F \times_{\mathrm{Y}} F'$. Then, the fact that φ is a fibration implies that $p_1 : F \times_{\mathrm{Y}} F' \to F$ and $p_2 : F \times_{\mathrm{Y}} F' \to F'$ are also fibrations. By [10, Lemma 8.9], such maps are basically projections, that is, $F \times_{\mathrm{Y}} F' \cong F \times N$ and $F \times_{\mathrm{Y}} F' \cong F' \times N'$ for some free nilspaces N, N'. In particular, p_1, p_2 are submersions. Hence $\varphi \circ \phi \circ p_1 = \phi' \circ p_2$ where both $\phi' \circ p_2$ and $\phi \circ p_1$ are submersions. Hence, as $\phi \circ p_1$ is also surjective, φ is also a submersion. □

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

Corollary A.8. *Let* X *be a k-step* CFR *nilspace. Then the action of the translation group* X*, that is,* $\mathrm{X} \times \Theta(\mathrm{X}) \to \mathrm{X}$*,* $(x, \alpha) \mapsto \alpha(x)$*, is* C^∞*.*

Proof. Recall that, by [7, Theorem 2.9.10], the translation group $\Theta(\mathrm{X})$ is a k-step nilpotent filtered Lie group, so in particular it is a k-step Lie-fibered nilspace. Then the map $\mathrm{X} \times \Theta(\mathrm{X}) \to \mathrm{X}$, $(x, \alpha) \mapsto \alpha(x)$ is a continuous morphism, so the result follows by Theorem A.6. □

We can now also prove the following result, used for instance in the proof of Proposition 4.29.

Lemma A.9. *Let* X *be a k-step Lie-fibered nilspace. Then the corner completion map* $K : \mathrm{Cor}^{k+1}(\mathrm{X}) \to \mathrm{X}$ *is* C^∞*.*

Proof. By [10, Theorem 1.8], there is a fibration $\varphi : F \to \mathrm{X}$ where F is a k-step free nilspace and $\mathrm{X} \cong F/H$ for some fiber-transitive fiber-discrete group H. By Corollary A.4, we have that $\mathrm{Cor}^{k+1}(\mathrm{X})$ is a C^∞ manifold and we have the following commutative diagram:

$$\begin{array}{ccc} \mathrm{Cor}^{k+1}(F) & \xrightarrow{K'} & F \\ \downarrow{\scriptstyle \varphi^{[\![k+1]\!]\setminus\{0^{k+1}\}}} & & \downarrow{\scriptstyle \varphi} \\ \mathrm{Cor}^{k+1}(\mathrm{X}) & \xrightarrow{K} & \mathrm{X}. \end{array}$$

Here K' is the completion function of F. Note that φ and $\varphi^{[\![k+1]\!]\setminus\{0^{k+1}\}}$ are submersions by Remark A.3, and K' is C^∞ (this last claim can be checked by hand, as K' has an explicit polynomial expression). Hence K is also a C^∞ map. □

We can now define the desired metrics that we shall work with.

Lemma A.10. *Let* X *be a k-step* CFR *nilspace endowed with the* C^∞ *structure given by Theorem A.2. Then there exists a Riemannian metric tensor on* X *which is* $\mathrm{Z}_k(\mathrm{X})$*-invariant.*

Proof. Any smooth manifold admits a Riemannian metric tensor (see [50, Proposition 13.3]). If g is such a tensor, it is easy to see that $\int_{\mathrm{Z}_k(\mathrm{X})} g(\cdot + z, \cdot + z)\, \mathrm{d}_{\mathrm{Z}_k(\mathrm{X})}(z)$ is also a Riemannian metric tensor which is $\mathrm{Z}_k(\mathrm{X})$-invariant. □

Note that there may be many such Riemannian tensors, as the previous result does not guarantee uniqueness. However, this will not be a problem in what follows.

Recall that in a connected Riemannian manifold, the geodesic distance is defined as the length of the shortest path between two points, and this distance agrees with the topology of the manifold, see [50, Proposition 8.19]. This is the distance function that we shall use. Let us record it as follows.

Definition A.11 (Riemannian metric on a CFRk-step nilspace). Let X be a CFR k-step nilspace endowed with a Riemannian metric tensor g. We define the distance $d(x, y)$ for $x, y \in \mathrm{X}$ to be 1 if x and y are in different connected components, and $d(x, y) = \inf_{\gamma_{x,y}} \int \sqrt{g(\nabla\gamma, \nabla\gamma)}$ otherwise, where the infimum is on the C^∞ maps $\gamma : [0, 1] \to \mathrm{X}$ with $\gamma(0) = x$ and $\gamma(1) = y$.

A particular case is that of 1-step CFR nilspaces. These are in fact compact abelian Lie groups. For those, we will always use a metric defined as follows.

Definition A.12 (Riemannain metric on compact abelian Lie groups). Let G be a compact abelian Lie group, thus $G \cong \mathrm{Z} \times \mathbb{T}^n$ where Z is a finite abelian group and $n \in \mathbb{Z}_{\geq 0}$. We define a Riemannian metric tensor on G by giving a Riemannian metric tensor on each of its connected components: on $\{a\} \times \mathbb{T}^n$ for $a \in \mathrm{Z}$, we define a Riemannian metric tensor as the n-fold product of the usual metric tensor on $\mathbb{T}$ obtained as the induced Riemannian metric when $\mathbb{T}$ is embedded in $\mathbb{R}^2$ the usual way, see [13, Chapter 1, Example 2.7]. On G, we define the distance $d_G(x, y) = 1$ if $x, y \in G$ are in different connected components, and as the usual Riemannian distance otherwise (see Definition A.11). We call this metric on G the *usual* or *flat* metric on G.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

We can now reap the benefits of using this Riemannian distance, by proving the main results concerning Lipschitz properties, announced at the beginning of the appendix.

We start with the following.

Proposition A.13. *Let* X, Y *be k-step* CFR *nilspaces, and let* $\varphi : \mathrm{X} \to \mathrm{Y}$ *be a (continuous) morphism. Then, relative to any two Riemannian metrics on* X *and* Y*, the map* φ *is Lipschitz.*

Proof. Let g_{X} and g_{Y} be the metric tensors of X and Y respectively. Consider the unit tangent bundle $\mathrm{UT}(\mathrm{X})$ on X and the map $D : \mathrm{UT}(\mathrm{X}) \to \mathbb{R}_{\geq 0}$ defined as

$$(x, W_x) \mapsto \frac{g_{\mathrm{Y}}(\nabla\varphi(W_x), \nabla\varphi(W_x))}{g_{\mathrm{X}}(W_x, W_x)}.$$

This map is always well-defined as the denominator is constantly 1. Moreover, it is a continuous map and $\mathrm{UT}(\mathrm{X})$ is compact. Thus, it attains a maximum M. Hence $g_{\mathrm{Y}}(\nabla\varphi(W), \nabla\varphi(W)) \leq M g_{\mathrm{X}}(W, W)$ for any tangent vector W.

Then, given any two points $x, y \in \mathrm{X}$, if they are in different connected components then we always have $d_{\mathrm{Y}}(\varphi(x), \varphi(y)) \leq L d_{\mathrm{X}}(x, y)$ for any constant $L \geq \mathrm{diam}(\mathrm{Y})$ so there is nothing to prove. On the other hand, if they are in the same connected component, let $\gamma : [0, 1] \to \mathrm{X}$ be such that $\gamma(0) = x$, $\gamma(1) = y$, and $d(x, y) = \int \sqrt{g_{\mathrm{X}}(\nabla\gamma, \nabla\gamma)}$. Then $d(\varphi(x), \varphi(y)) \leq \int \sqrt{g_{\mathrm{Y}}(\nabla\varphi \circ \nabla\gamma, \nabla\varphi \circ \nabla\gamma)} \leq M \int \sqrt{g_{\mathrm{X}}(\nabla\gamma, \nabla\gamma)} = M d(x, y)$.Hence, the morphism φ is $\max(\mathrm{diam}(\mathrm{Y}), M)$-Lipschitz. □

Remark A.14. It is important to note that the Lipschitz constant depends on the metrics chosen. But once we fix such metrics, all continuous morphisms are Lipschitz maps. We will endow every k-step CFR nilspace with one such Riemannian metric tensor and a corresponding metric.

The idea of the previous result can be extended slightly, to show that if we have a compact family of translations acting on a nilspace, then their Lipschitz constants are bounded uniformly.

Corollary A.15. *Let* X *be a* CFR *k-step nilspace where we have fixed some Riemannian metric* g_{X}*. Let* $C \subset \Theta(\mathrm{X})$ *be compact. Then* $\sup_{\alpha \in C} \|\alpha\|_L < \infty$.

Proof. The idea of the proof is the same as before. The only difference is to consider the map $\mathrm{UT}(\mathrm{X}) \times C \to \mathbb{R}$ defined as $(x, W_x, \alpha) \mapsto \frac{g_{\mathrm{X}}(\nabla\alpha(W_x), \nabla\alpha(W_x))}{g_{\mathrm{X}}(W_x, W_x)}$, which is continuous by Corollary A.8 (in fact, C^∞). In particular, it attains a maximum and thus, repeating the same argument from the proof of Proposition A.13, the result follows. □

We need the next technical result that follows from the fact that $\pi_{k-1} : \mathrm{X} \to \mathrm{X}_{k-1}$ is a submersion by Theorem A.6 (we omit its proof).

Lemma A.16. *Let* X *be a* CFR *k-step nilspace and assume that we have fixed Riemannian metrics on* X *and* X_{k-1}*. Then for every point* $\pi_{k-1}(x) \in \mathrm{X}_{k-1}$ *there exists a neighborhood* $U_x \subset \mathrm{X}_{k-1}$ *and a Lipschitz cross-section* $s_x : U_x \to \mathrm{X}$*. In particular, by compactness, there exists a finite set* $\{\pi_{k-1}(x_1), \dots, \pi_{k-1}(x_s)\} \subset \mathrm{X}_{k-1}$ *such that* $\mathrm{X}_{k-1} = \bigcup_{i=1}^{s} U_{x_i}$.

Lemma A.17. *Let* Y *be a k-step* CFR *nilspace and assume that we have fixed Riemannian metrics on* Y *and* Y_{k-1}*. Let* $F : \mathrm{Y} \to \mathbb{C}$ *be Lipschitz with respect to the metric on* Y *and suppose that* $F = F' \circ \pi_{k-1}$ *for some function* $F' : \mathrm{Y}_1 \to \mathbb{C}$*. Then* $F' : \mathrm{Y}_{k-1} \to \mathbb{C}$ *is also Lipschitz, with* $\|F'\|_{\max}$ *bounded in terms of* $\|F\|_{\max}$ *and the chosen metrics.*

Proof. Consider a covering of X_{k-1} by open sets such as the one given in Lemma A.16. By the Lebesgue number lemma there exists $\delta > 0$ such that for any $\pi_{k-1}(x), \pi_{k-1}(y) \in \mathrm{X}_{k-1}$ if $d(\pi_{k-1}(x), \pi_{k-1}(y)) < \delta$ then $\pi_{k-1}(x), \pi_{k-1}(y) \in U_{x_i}$ for some $i \in [s]$.

Let $\pi_{k-1}(x), \pi_{k-1}(y)$ be any two points in X_{k-1}. If $d(\pi_{k-1}(x), \pi_{k-1}(y)) \geq \delta$ then clearly $|F'(\pi_{k-1}(x)) - F'(\pi_{k-1}(y))| \leq 2\|F\|_\infty \leq \frac{2\|F\|_\infty}{\delta} d(\pi_{k-1}(x), \pi_{k-1}(y))$. On the other hand, if $d(\pi_{k-1}(x), \pi_{k-1}(y)) < \delta$ then we have $|F'(\pi_{k-1}(x)) - F'(\pi_{k-1}(y))| = |F(s_i(\pi_{k-1}(x))) -$

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

$F(s_i(\pi_{k-1}(x)))| \leq \|F\|_L d(s_i(\pi_{k-1}(x)), s_i(\pi_{k-1}(x))) \leq \|F\|_L \|s_i\|_L d(\pi_{k-1}(x), \pi_{k-1}(y))$. Thus, we have $\|F\|_L \leq \max\left(\frac{2\|F\|_\infty}{\delta}, \max_{i\in[s]} \|F\|_L \|s_i\|_L\right)$. □

Acknowledgments. The authors thank R. Giménez Conejero and Alex Hof for valuable discussions regarding the appendix of this paper. We also thank the anonymous referee for valuable comments that improved the paper. For a majority of the authors' work on this article, the second-named author was affiliated with HUN-REN Alfréd Rényi Institute of Mathematics, and for the final part of the work he was affiliated with Université Paris Cité, Sorbonne Université, CNRS, IMJ-PRG.

Competing interests. The authors have no competing interests to declare.

Financial support. All authors received funding from projects PID2020-113350GB-I00 and PID2024-156180NB-I00 funded by Spain's MICIU/AEI and the EU. The second-named author received funding from project Momentum (Lendület) 30003 of the Hungarian Government and from HORIZON-MSCA-2024-PF-01, AlgHOF 101202161, funded by the European Union.[21] The second and third-named authors were also supported partially by the NKFIH "Élvonal" KKP 133921 grant and partially by the Hungarian Ministry of Innovation and Technology NRDI Office within the framework of the Artificial Intelligence National Laboratory Program (MILAB, RRF-2.3.1- 21-2022-00004).

References

[1] S. Aaronson, 'Algorithms for Boolean function query properties', *SIAM J. Comput.* **32**(5) (2003), 1140–1157.
[2] N. I. Akhiezer and I.M. Glazman, *Theory of Linear Operators in Hilbert Space. Vol. I., translated from the third Russian edition by E. R. Dawson.* Translation edited by W. N. Everitt Monogr. Stud. Math., 9 Pitman (Advanced Publishing Program) (Boston, MA-London, 1981), xxxii+312 pp.
[3] N. Alon, E. Fischer, M. Krivelevich and M. Szegedy, 'Efficient testing of large graphs', *Combinatorica* **20**(4) (2000), 451–476.
[4] N. Bourbaki, *Elements of Mathematics. General Topology. Part 1.* (Hermann, Paris; Addison-Wesley Publishing Co., Reading, MA-London-Don Mills, Ont. 1966).
[5] O. A. Camarena and B. Szegedy, *Nilspaces, Nilmanifolds and Their Morphisms*, 2010, arXiv:1009.3825.
[6] P. Candela, 'Notes on nilspaces: algebraic aspects', *Discrete Anal.* 2017, Paper No. 15, 59 pp.
[7] P. Candela, 'Notes on compact nilspaces', *Discrete Anal.* 2017, Paper No. 16, 57 pp.
[8] P. Candela, D. Gonzalez-Sanchez and B. Szegedy, 'On nilspace systems and their morphisms', *Ergodic Theory Dynam. Systems* **40**(11) (2020), 3015–3029.
[9] P. Candela, D. Gonzalez-Sanchez and B. Szegedy, 'On higher-order Fourier analysis in characteristic p', *Ergodic Theory Dynam. Systems* **43**(12) (2023), 3971–4040.
[10] P. Candela, D. Gonzalez-Sanchez and B. Szegedy, 'Free nilspaces, double coset nilspaces, and Gowers norms', *Math. Ann.* **395**, 23 (2026).
[11] P. Candela and B. Szegedy, 'Nilspace factors for general uniformity seminorms, cubic exchangeability and limits', *Mem. Amer. Math. Soc.* **287**(1425) (2023), v+101 pp.
[12] P. Candela, B. Szegedy, 'Regularity and inverse theorems for uniformity norms on compact abelian groups and nilmanifolds', *J. Reine Angew. Math.* **789** (2022), 1–42.
[13] M. P. do Carmo, *Geometria riemanniana [Riemannian Geometry]*, Projeto Euclides, 10 [Euclid Project] (Instituto de Matemática Pura e Aplicada, Rio de Janeiro, 1979), ix+238 pp.
[14] Ş. Cobzaş, R. Miculescu and A. Nicolae, *Lipschitz functions*, Lecture Notes in Math., 2241 Springer, Cham, 2019. xiv+591 pp.
[15] D. L. Cohn, *Measure Theory*, Second edition (Birkhäuser/Springer, New York, 2013). xxi+457 pp.
[16] T. Eisner and T. Tao, 'Large values of the Gowers-Host-Kra seminorms', *J. Anal. Math.* **117** (2012), 133–186.
[17] A. M. Frieze and R. Kannan, 'Quick approximation to matrices and applications', *Combinatorica* **19** (1999), 175–220.
[18] A. M. Frieze and R. Kannan, 'A simple algorithm for constructing Szemerédi's regularity partition', *Electron. J. Combin.* **6** (1999), #R17, 7 pp.
[19] H. Furstenberg, 'Ergodic behavior of diagonal measures and a theorem of Szemerédi on arithmetic progressions', *J. Analyse Math.* **31** (1977), 204–256.
[20] N. R. Goodman, 'Statistical analysis based on a certain multivariate complex Gaussian distribution (an introduction)', *Ann. Math. Statist.* **34** (1963), 152–177.
[21] B. I. Golubov, *Multiple series and Fourier integrals.* (Russian) Mathematical analysis, Vol. 19, pp. 3–54, 232 Itogi Nauki i Tekhniki [Progress in Science and Technology] Akad. Nauk SSSR, Vsesoyuz. Inst. Nauchn. i Tekhn. Inform., Moscow, 1982
[22] W. T. Gowers, 'A new proof of Szemerédi's theorem', *Geom. Funct. Anal.* **11** (2001), 465–588.

[21] Views and opinions expressed are those of the author(s) only and do not reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

[23] W. T. Gowers, 'Hypergraph regularity and the multidimensional Szemerédi theorem', *Ann. of Math. (2)* **166**(3) (2007), 897–946.
[24] W. T. Gowers, 'Decompositions, approximate structure, transference, and the Hahn-Banach theorem', *Bull. Lond. Math. Soc.* **42**(4) (2010), 573–606.
[25] W. T. Gowers, 'Generalizations of Fourier analysis, and how to apply them', *Bull. Amer. Math. Soc. (N.S.)* **54**(1) (2017), 1–44.
[26] W. T. Gowers and L. Milićević, 'A quantitative inverse theorem for the U^4 norm over finite fields', 2020, arXiv:1712.00241.
[27] W. T. Gowers and L. Milićević, 'An inverse theorem for Freiman multi-homomorphisms', 2020, arXiv:2002.11667.
[28] W.T. Gowers, J. Wolf, 'Linear forms and higher-degree uniformity for functions on F_p^n', *Geom. Funct. Anal.* **21**(1) (2011), 36–69.
[29] W. T. Gowers and J. Wolf, 'The true complexity of a system of linear equations', *Proc. Lond. Math. Soc., (1)* **100** (2010), 155–176.
[30] B. Green and T. Sanders, 'Boolean functions with small spectral norm', *Geom. Funct. Anal.* **18**(1) (2008), 144–162.
[31] B. Green and T. Tao, *An arithmetic regularity lemma, an associated counting lemma, and applications*, in *An Irregular Mind*, Bolyai Soc. Math. Stud. 21, János Bolyai Mathematical Society, Budapest, 2010, pp. 261–334.
[32] B. Green and T. Tao, 'An inverse theorem for the Gowers U^3-norm', *Proc. Edinburgh Math. Soc. (1)* **51** (2008), 73–153.
[33] B. Green and T. Tao, 'The primes contain arbitrarily long arithmetic progressions', *Ann. of Math. (2)* **167** (2008), 481–547.
[34] B. Green and T. Tao, 'The quantitative behaviour of polynomial orbits on nilmanifolds', *Ann. of Math. (2)* **175**(2) (2012), 465–540.
[35] B. Green, T. Tao and T. Ziegler, 'An inverse theorem for the Gowers $U^{s+1}[N]$-norm', *Ann. of Math. (2)* **176**(2) (2012), 1231–1372.
[36] A. Gut, *An Intermediate Course in Probability*, Second edition, Springer Texts Statist. (Springer, New York, 2009), xvi+303 pp.
[37] Y. Gutman, F. Manners and P. P. Varjú, 'The structure theory of nilspaces I', *J. Anal. Math.* **140**(1) (2020), 299–369.
[38] Y. Gutman, F. Manners and P. P. Varjú, 'The structure theory of nilspaces II: Representation as nilmanifolds', *Trans. Amer. Math. Soc.* **371**(7) (2019), 4951–4992.
[39] Y. Gutman, F. Manners and P. P. Varjú, 'The structure theory of nilspaces III: Inverse limit representations and topological dynamics', *Adv. Math.* **365** (2020), 107059, 53 pp.
[40] H. Hatami, P. Hatami and S. Lovett, 'Higher-order Fourier analysis and applications', *Found. Trends Theor. Comput. Sci.* **13**(4) (2019), front matter, 247–448
[41] R. A. Horn and C. R. Johnson, *Matrix Analysis , Second edition* (Cambridge University Press, Cambridge, 2013), xviii+643 pp.
[42] B. Host and B. Kra, 'A point of view on Gowers uniformity norms', *New York J. Math.* **18** (2012), 213–248.
[43] B. Host and B. Kra, 'Nonconventional ergodic averages and nilmanifolds', *Ann. of Math. (2)* **161**(1) (2005), 397–488.
[44] B. Host and B. Kra, 'Parallelepipeds, nilpotent groups, and Gowers norms', *Bull. Soc. Math. France* **136**(3) (2008), 405–437.
[45] B. Host and B. Kra, 'Uniformity seminorms on ℓ^∞ and applications', *J. Anal. Math.* **108** (2009), 219–276.
[46] A. Jamneshan and T. Tao, 'The inverse theorem for the U^3 Gowers uniformity norm on arbitrary finite abelian groups: Fourier-analytic and ergodic approaches' , *Discrete Anal.* 2023, Paper No. 11, 48 pp.
[47] A. Jamneshan, O. Shalom and T. Tao, *The structure of totally disconnected Host–Kra–Ziegler factors, and the inverse theorem for the U^k Gowers uniformity norms on finite abelian groups of bounded torsion*, 2023, arXiv:2303.04860.
[48] D. Kim, A. Li and J. Tidor, *Cubic Goldreich-Levin*, Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 4846–4892.
[49] H. Lebesgue, 'Sur la représentation trigonométrique approchée des fonctions satisfaisant à une condition de Lipschitz', *Bull. Soc. Math. France* **38** (1910), 184–210.
[50] J. M. Lee, *Introduction to Smooth Manifolds, Second edition*, Grad. Texts in Math., 218 Springer, New York, 2013. xvi+708 pp.
[51] J. Leng, A. Sah and M. Sawhney, *Quasipolynomial bounds on the inverse theorem for the Gowers $U^{s+1}[N]$-norm*, 2024, arXiv:2402.17994.
[52] F. Manners, *Quantitative bounds in the inverse theorem for the Gowers U^{s+1}-norms over cyclic groups*, 2018, arXiv:1811.00718.
[53] L. Milićević, Quantitative inverse theorem for Gowers uniformity norms U^5 and U^6 in F_2^n, to appear in *Canadian Journal of Mathematics*.
[54] L. Milićević, 'An inverse theorem for certain directional Gowers uniformity norms', *Publ. Inst. Math. (Beograd) (N.S.)* **113**(127) (2023), 1–56.
[55] I.D. Morris, 'A rapidly-converging lower bound for the joint spectral radius via multiplicative ergodic theory', *Adv. Math.* **225**(6) (2010), 3425–3445.
[56] V. Rödl, B. Nagle, J. Skokan, M. Schacht and Y. Kohayakawa, 'The hypergraph regularity method and its applications', *Proc. Natl. Acad. Sci. USA.* (2005) *Jun* **7; 102**(23) pp. 8109–13.
[57] K. Srinivasacharyulu, 'Variétés Différentielles et analytiques by N. Bourbaki. Hermann, Paris, 1967. 97 pages', *Canad. Math. Bull.* **11**(3), (1968), 514–514.
[58] B. Szegedy, *Higher order fourier analysis as an algebraic theory I*, 2009, arXiv:0903.0897.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press

[59] B. Szegedy, *Higher order fourier analysis as an algebraic theory II*, 2009, arXiv:0911.1157.
[60] B. Szegedy, *Higher order fourier analysis as an algebraic theory III*, 2010, arXiv:1001.4282.
[61] B. Szegedy, 'Limits of kernel operators and the spectral regularity lemma', *European J. of Comb.* **32** (7) (2011), 1156–1167
[62] E. Szemerédi, 'On sets of integers containing no k elements in arithmetic progression', *Acta Arith.* **27** (1975), 199–245.
[63] E. Szemerédi, 'Regular partitions of graphs', *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976), pp.* 399 – 401 *Colloq. Internat. CNRS*, **260**, (1978).
[64] T. Tao, 'A quantitative ergodic theory proof of Szemerédi's theorem', *Electron. J. Combin.* **13**(1) (2006), Research Paper 99, 49 pp.
[65] T. Tao, 'A variant of the hypergraph removal lemma', *J. Combin. Theory Ser. A*, **113** (7), 1257–1280.
[66] T. Tao and V. Vu, *Additive Combinatorics*, Cambridge Stud. Adv. Math., 105 (Cambridge University Press, Cambridge, 2010), xviii+512 pp.
[67] T. Tao and T. Ziegler, 'The inverse conjecture for the Gowers norm over finite fields via the correspondence principle', *Anal. PDE* **3**(1) (2010), 1–20.
[68] T. Tao and T. Ziegler, 'The inverse conjecture for the Gowers norm over finite fields in low characteristic', *Ann. Comb.* **16** (2012), 121–188.
[69] L. N. Trefethen and M. Embree, *Spectra and Pseudospectra. The Behavior of Nonnormal Matrices and Operators* (Princeton University Press, Princeton, NJ, 2005), xviii+606 pp.
[70] L. Trevisan, 'Additive combinatorics and theoretical computer science', *ACM SIGACT News (2)* **40** (2009), 50–66.
[71] M. Tulsiani and J. Wolf, 'Quadratic Goldreich–Levin Theorems', *SIAM J. Comput.* **43**(2) (2014), 730–766.
[72] R. Vershynin, *High-Dimensional Probability: An Introduction with Applications in Data Science*, Camb. Ser. Stat. Probab. Math., 47 (Cambridge University Press, Cambridge, 2018), xiv+284 pp.
[73] T. Ziegler, 'Universal Characteristic Factors and Furstenberg Averages', *J. Amer. Math. Soc.* **20** (2007), 53–97.

https://doi.org/10.1017/fms.2026.10238 Published online by Cambridge University Press