

Sparse Harmonic Transforms II: Best s -Term Approximation Guarantees for Bounded Orthonormal Product Bases in Sublinear-Time

Bosu Choi* Mark Iwen† Toni Volkmer‡

Abstract

In this paper we develop a sublinear-time compressive sensing algorithm for approximating functions of many variables which are compressible in a given Bounded Orthonormal Product Basis (BOPB). The resulting algorithm is shown to both have an associated best s -term recovery guarantee in the given BOPB, and also to work well numerically for solving sparse approximation problems involving functions contained in the span of fairly general sets of as many as $\sim 10^{230}$ orthonormal basis functions. All code is made publicly available.

As part of the proof of the main recovery guarantee new variants of the well known CoSaMP algorithm are proposed which can utilize any sufficiently accurate support identification procedure satisfying a Support Identification Property (SIP) in order to obtain strong sparse approximation guarantees. These new CoSaMP variants are then proven to have both runtime and recovery error behavior which are largely determined by the associated runtime and error behavior of the chosen support identification method. The main theoretical results of the paper are then shown by developing a sublinear-time support identification algorithm for general BOPB sets which is robust to arbitrary additive errors. Using this new support identification method to create a new CoSaMP variant then results in a new robust sublinear-time compressive sensing algorithm for BOPB-compressible functions of many variables.

Keywords High-dimensional function approximation · Sublinear-time algorithms · Function learning · Sparse approximation · Compressive Sensing · Sparse Fourier transforms (SFT)

Mathematics subject classification 65T40 · 68W25

*University of Texas at Austin, Oden Institute for Computational Engineering and Sciences, choibosu@utexas.edu.

†Michigan State University, Department of Mathematics, and the Department of Computational Mathematics, Science and Engineering (CMSE), markiwen@math.msu.edu.

‡Chemnitz University of Technology, Faculty of Mathematics, toni.volkmer@math.tu-chemnitz.de.

1 Introduction

In this paper we focus on rapidly computing best s -term approximations in the sense of compressive sensing [13, 17] for functions of many variables $f : \mathcal{D} \subseteq \mathbb{R}^D \rightarrow \mathbb{C}$. More specifically, we develop a numerical method that aims to very quickly approximate any given function f using a near optimal s -sparse linear combination of N' fixed basis functions $\mathcal{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_{N'}\}$ chosen in advance. The developed method has two basic components: (i) a low-cardinality grid of evaluation points, $\mathcal{G} \subset \mathcal{D}$, and (ii) a fast deterministic algorithm $\mathcal{H} : \mathbb{C}^{|\mathcal{G}|} \rightarrow \text{span}(\mathcal{B})$ which takes f evaluated on \mathcal{G} , $f(\mathcal{G}) \in \mathbb{C}^{|\mathcal{G}|}$, as input, and then outputs an accurate s -sparse approximation to f of the form $\sum_{j=1}^s a_{\ell_j} \mathbf{b}_{\ell_j} \in \text{span}(\mathcal{B})$. In particular, we require that \mathcal{H} can approximate all functions f near-optimally based only on the evaluations of f on \mathcal{G} so that

$$\|f - \mathcal{H}(f(\mathcal{G}))\| \lesssim C \inf_{z \in \mathbb{C}^{N'}, \|z\|_0 \leq s} \left\| f - \sum_{j=0}^{N'} z_j \mathbf{b}_j \right\| \quad (1.1)$$

holds with respect to suitable norms for all functions $f : \mathcal{D} \rightarrow \mathbb{C}$ in a sufficiently general function class.

Note that we are requiring several strong properties of both \mathcal{G} and \mathcal{H} above. First, we want the approximation algorithm \mathcal{H} to succeed *for all* functions f in a suitably large class when only given access to function evaluations of each f on the same fixed and nonadaptive grid \mathcal{G} . Second, we require that \mathcal{H} is *fast*, which will mean in this paper that we require it to use a total number of scalar arithmetic and read/write operations that scales *sublinearly* with respect to the basis size N' (e.g., herein we will focus on methods with runtimes that scale like $\mathcal{O}(\log^c(N'))$ for all sufficiently small sparsities s). This second requirement has several other beneficial repercussions beyond computational speed. Principally among them is the fact that the deterministic procedure \mathcal{H} can use at most $o(N')$ function evaluations since its fast runtime constrains the number of function evaluations \mathcal{H} can use. This effectively constrains the size of the nonadaptive grid $\mathcal{G} \subset \mathcal{D}$ that it makes sense to use in the first place. Similarly, any such \mathcal{H} must also have low $o(N')$ memory requirements given that it only has time to perform $o(N')$ -total scalar operations involving memory accesses.

The first work on sublinear-time algorithms \mathcal{H} of this kind focused almost exclusively on the one-dimensional Fourier basis where, e.g., $\mathcal{B} = \{e^{2\pi i \omega x} \mid \omega \in ([N'/2], \lfloor N'/2 \rfloor) \cap \mathbb{Z}\}$. The first of these [30, 21, 19, 22] were randomized algorithms which used grids \mathcal{G} that varied from function to function and that failed with some nonzero probability for each given f . All of these methods have runtimes that scale like $\mathcal{O}(s \log^c N')$ and achieve approximation errors along the lines of (1.1) with high probability (w.h.p.) for each given $f : [0, 1] \rightarrow \mathbb{C}$. Later on, entirely deterministic and explicit $\mathcal{O}(s^2 \log^c N')$ -time methods

\mathcal{H} were then devised which use one fixed and nonadaptive grid \mathcal{G} in order to guarantee approximation errors of the form (1.1) for all sufficiently smooth and periodic functions $f : [0, 1] \rightarrow \mathbb{C}$ (see [24, 25, 3, 39]). These deterministic methods were then randomized in [31] to achieve highly efficient $\mathcal{O}(s \log^c N')$ -time randomized discrete Fourier transform methods (generally known as “sparse Fourier transforms”) with high probability best s -term approximation guarantees (1.1) along the lines of the first methods mentioned above, as well as sped up to produce entirely deterministic methods that are significantly faster than the generic $\mathcal{O}(s^2 \log^c N')$ -time deterministic algorithms for periodic functions $f : [0, 1] \rightarrow \mathbb{C}$ which exhibit structured sparsity in the Fourier domain [5]. Code for many of these methods is publicly available¹, and a nice survey article covering the standard techniques used to construct many of these first sublinear-time Fourier methods appeared in 2014 [20].

As sublinear-time methods for the one-dimensional Fourier basis started to mature, similar algorithms began to be developed for other one-dimensional bases \mathcal{B} as well, including for the cosine, Chebyshev, and Legendre polynomial bases [23, 4] (see also [37] for traditional compressive sensing methods which focus on the Legendre polynomial basis). Recently these ideas have been extended yet further to produce sublinear-time algorithms with reconstruction guarantees for restricted classes of signals exhibiting approximate sparsity in any given one-dimensional Jacobi polynomial basis [18]. Another direction of research has focused on extending the types of sparse approximation algorithms discussed above to higher dimensional settings in order to approximate, e.g., functions $f : [0, 1]^D \rightarrow \mathbb{C}$ with respect to either multidimensional Fourier [36, 26, 34, 10, 28, 32, 27, 11, 29] or Chebyshev [35] bases \mathcal{B} of cardinality $N' = N^D$. In these cases achieving fast algorithms \mathcal{H} that run in $o(N')$ -time becomes increasingly important as D grows.

As in the one-dimensional setting, sublinear-time methods \mathcal{H} for approximating functions of D variables are most well developed in the case of the multidimensional Fourier basis where, e.g., $\mathcal{B} = \{e^{2\pi i \boldsymbol{\omega} \cdot \mathbf{x}} \mid \boldsymbol{\omega} \in ([N'/2], [N'/2]]^D \cap \mathbb{Z}^D\}$. For example, see Theorem 8 in [26] and Theorems 10 and 12 in [32] for explicit and deterministic $\mathcal{O}(s^2 \log^c(N'))$ -time methods that use function evaluations on a single fixed and nonadaptive grid $\mathcal{G} \subset [0, 1]^D$ in order to guarantee approximation errors of the form (1.1) for all sufficiently smooth and periodic functions $f : [0, 1]^D \rightarrow \mathbb{C}$. When it comes to approximating functions of many variables with respect to non-Fourier bases \mathcal{B} in $o(N')$ -time, however, very little is currently known. The first result in this direction [12] provided sublinear-time recovery guarantees for all functions that are *exactly s -sparse*² in any tensor product basis \mathcal{B} of one-

¹The code for an implementation of [31] is available at <https://sourceforge.net/projects/aafftannarborfa/>. The code for an implementation of [5] is available at https://www.math.msu.edu/~markiwen/Code/FAST_block_sparse.zip.

²A function is exactly s -sparse in \mathcal{B} if it is a linear combination of $\leq s$ unknown elements of \mathcal{B} .

dimensional bounded orthonormal bases. The aim of this paper is to augment this first general result with best s -term approximation guarantees along the lines of (1.1) while maintaining its fast runtime and small fixed and nonadaptive grid size. In doing so the authors aim to complement existing compressive sensing approaches for uncertainty quantification and function approximation [38, 9, 1, 7, 2] with a new class of methods whose runtimes scale sublinearly in the basis size used for approximation. These new methods will then hopefully allow for the extension of such techniques to, e.g., functions of hundreds or even thousands of variables in a more computationally feasible fashion.

1.1 Setup and Main Results

Let $L^2(\mathcal{D}, \mu)$ for $\mathcal{D} := \times_{j \in [D]} \mathcal{D}_j \subset \mathbb{R}^D$ denote all functions $f : \mathcal{D} \rightarrow \mathbb{C}$ that are square-integrable with respect to a given product of probability measures $\mu := \times_{j \in [D]} \mu_j$ over \mathcal{D} , and suppose that you are given a countable orthonormal (with respect to μ) basis,

$$\mathcal{B}' := \{T_{\mathbf{n}} : \mathcal{D} \rightarrow \mathbb{C} \mid \mathbf{n} \in \mathbb{N}^D\}, \quad (1.2)$$

of $L^2(\mathcal{D}, \mu)$ so that

$$\langle T_{\mathbf{k}}, T_{\mathbf{l}} \rangle_{(\mathcal{D}, \mu)} := \int_{\mathcal{D}} T_{\mathbf{k}}(\boldsymbol{\xi}) \overline{T_{\mathbf{l}}(\boldsymbol{\xi})} d\mu(\boldsymbol{\xi}) = \delta_{\mathbf{k}, \mathbf{l}} = \begin{cases} 1 & \text{if } \mathbf{k} = \mathbf{l} \\ 0 & \text{if } \mathbf{k} \neq \mathbf{l} \end{cases}.$$

Furthermore, suppose that \mathcal{B}' is a tensor product basis so that

$$T_{\mathbf{n}}(\boldsymbol{\xi}) := \prod_{j \in [D]} T_{j; n_j}(\xi_j) \quad (1.3)$$

holds for all $\mathbf{n} \in \mathbb{N}^D$ and $\boldsymbol{\xi} \in \mathcal{D}$, where each set $\mathcal{B}'_j := \{T_{j; n_j} : \mathcal{D}_j \rightarrow \mathbb{C} \mid n_j \in \mathbb{N}\}$ with $j \in [D] := \{0, \dots, D-1\}$ is itself an orthonormal (with respect to the probability measure μ_j over $\mathcal{D}_j \subset \mathbb{R}$) basis of $L^2(\mathcal{D}_j, \mu_j)$. We will call any such basis \mathcal{B}' an Orthonormal Product Basis (OPB) with respect to μ .

Our objective in this paper is to approximate smooth functions $f \in L^2(\mathcal{D}, \mu)$ as rapidly as possible using just a few point evaluations. Toward this end we will take the traditional approach of considering only a finite subset $\mathcal{B}_{N,d}$ of \mathcal{B}' , and then approximating f by approximating its projection \tilde{f} onto the span of $\mathcal{B}_{N,d}$ (consider, e.g., hyperbolic cross/sparse grid methods for approximating functions of several variables [40, 16, 8]). The potential improvement that the sublinear-time methods considered herein will then potentially provide over such standard methods in some cases will come from the fact that the finite basis $\mathcal{B}_{N,d}$ can be chosen to be extremely large herein (e.g., experiments were performed for Section 5 on a standard workstation using bases of cardinality 200^{100}).

More specifically, herein we will consider two different types of bases \mathcal{B}' , each of which will allow us to demonstrate that the finite basis $\mathcal{B}_{N,d} \subset \mathcal{B}'$ we select below for approximation purposes also promote computational efficiency.

We will characterize \mathcal{B}' below based on the behavior of its lowest order elements

$$\mathcal{B}_N := \{T_{\mathbf{n}} \mid \|\mathbf{n}\|_{\infty} < N\} \subset \mathcal{B}'$$

which we will assume throughout this paper is a finite Bounded Orthonormal System (BOS) with respect to the probability measure μ over \mathcal{D} with a finite BOS constant

$$K' := \max_{\mathbf{n} \in [N]^D} \|T_{\mathbf{n}}\|_{\infty} := \max_{\mathbf{n} \in [N]^D} \sup_{\boldsymbol{\xi} \in \mathcal{D}} |T_{\mathbf{n}}(\boldsymbol{\xi})| \in [1, \infty).$$

Note that this implies that each set $\mathcal{B}_{j;N} := \{T_{j;n_j} : \mathcal{D}_j \rightarrow \mathbb{C} \mid n_j \in [N]\} \subset \mathcal{B}'_j$ with $j \in [D]$ is itself also a BOS with respect to the probability measure μ_j over $\mathcal{D}_j \subset \mathbb{R}$ with BOS constant

$$K_j := \max_{n_j \in [N]} \|T_{j;n_j}\|_{\infty} \in [1, \infty). \quad (1.4)$$

Finally, we will further define K_j^0 to be

$$1 \leq K_j^0 := \|T_{j;0}\|_{\infty} \leq K_j \quad (1.5)$$

for each $j \in [D]$. Note that K_j^0 is strictly smaller than K_j for many BOSs of interest (e.g., the cosine and Chebyshev polynomial bases as \mathcal{B}'_j both have $K_j^0 = 1 < K_j = \sqrt{2}$). From these two definitions we can also see, e.g., that $\prod_{j \in [D]} K_j^0 \leq K' = \prod_{j \in [D]} K_j$ always holds. Due to the boundedness of K' assumed throughout the remainder of this paper we will always refer to \mathcal{B}_N (as well as \mathcal{B}' with slight abuse) as a Bounded Orthonormal Product Basis (BOPB) going forward.

We will approximate any given smooth $f \in L^2(\mathcal{D}, \mu)$ by approximating its projection \tilde{f} onto the span of the finite BOS set

$$\mathcal{B}_{N,d} := \{T_{\mathbf{n}} \mid \mathbf{n} \in [N]^D \text{ and } \|\mathbf{n}\|_0 \leq d\} \subseteq \mathcal{B}_N \subset \mathcal{B}' \quad (1.6)$$

for some $d \in [D + 1]$, where d is used to help constrain the BOS constant. The BOS constant $K \leq K'$ of $\mathcal{B}_{N,d}$ will be referred to as the *effective BOS constant* below. As is usually the case in compressive sensing scenarios involving BOSs, its size will be a significant consideration with respect to sampling and computational efficiency. In order to limit K 's size we will concentrate on the following two types of BOPBs going forward:

- **BOPBs of TYPE I:** We will say a BOPB is of type I if the BOS constants K_j are 1 for all but at most $\tilde{d} \in \mathbb{Z} \cap [0, D]$ BOS basis sets $\mathcal{B}_{j;N}$. In this case we let

$K_\infty := \max_{j \in [D]} K_j$ and note that $1 \leq K \leq K' \leq K_\infty^{\tilde{d}}$ so that K will scale sub-exponentially in D when $\tilde{d} \ll D$ independently of our choice of d in (1.6). We note that this type of BOPB includes several interesting examples of bases including the multidimensional Fourier basis (for which $\tilde{d} = 0$), and mixed BOPBs \mathcal{B}' that have one-dimensional Fourier bases used for all but \tilde{d} of their \mathcal{B}'_j component bases.

- **BOPBs of TYPE II:** We will say a BOPB is of type II if $K_0 := \max_{j \in [D]} K_j^0 = 1$. This type of BOPB includes many bases where having a small number of interacting dimensions, d , helps to limit the effective BOS constant K involved in the underlying sparse approximation problem. Examples include the multivariate cosine, Chebyshev, and Legendre polynomial bases, as well as mixed polynomial bases where each one-dimensional component basis \mathcal{B}'_j is, e.g., a potentially different Jacobi polynomial basis.

In either case above one can see that $K \leq K_\infty^d K_0^{D-d}$ will always hold. In particular, $K = 2^{d/2}$ always holds if \mathcal{B}' is either the multivariate cosine or Chebyshev basis in the type II case. In the type I case we note that $1 \leq K \leq K_\infty^{\min(d, \tilde{d})} K_0^{\tilde{d} - \min(d, \tilde{d})} \leq K_\infty^{\tilde{d}}$ will always hold so that d can be set to D without causing K to become too large if, e.g., $\tilde{d} \ll D$. This is certainly the case if \mathcal{B}' is the multidimensional Fourier basis where $\tilde{d} = 0$.

Let $f \in L^2(\mathcal{D}, \mu)$ be smooth enough³ that there exists a sequence $\{c_n\}_{n \in \mathbb{N}^D}$ such that

$$f(\boldsymbol{\xi}) = \sum_{\mathbf{n} \in \mathbb{N}^D} c_n T_{\mathbf{n}}(\boldsymbol{\xi}) \quad (1.7)$$

holds pointwise for all $\boldsymbol{\xi} \in \mathcal{D}$. Given such an $f : \mathcal{D} \rightarrow \mathbb{C}$, we will denote its orthogonal projection onto the span of $\mathcal{B}_{N,d}$ by $\tilde{f} : \mathcal{D} \rightarrow \mathbb{C}$. Let

$$\mathcal{I}_{N,d} := \{\mathbf{n} \in [N]^D \mid \|\mathbf{n}\|_0 \leq d \leq D\} \subset \mathbb{N}^D$$

be the set of indices corresponding to the basis elements in $\mathcal{B}_{N,d}$. We then have that

$$\tilde{f}(\boldsymbol{\xi}) := \sum_{\mathbf{n} \in \mathcal{I}_{N,d}} \tilde{c}_n T_{\mathbf{n}}(\boldsymbol{\xi}) \quad (1.8)$$

for all $\boldsymbol{\xi} \in \mathcal{D}$, where $\tilde{\mathbf{c}}$ will be considered to be a vector in $\mathbb{C}^{|\mathcal{I}_{N,d}|}$ indexed by $\mathcal{I}_{N,d}$. Note further that the entries of $\tilde{\mathbf{c}}$ will satisfy $\tilde{c}_n = c_n$ for all $\mathbf{n} \in \mathcal{I}_{N,d}$.

As mentioned above, we will ultimately approximate f by producing a sparse approximation in $\mathcal{B}_{N,d}$ to \tilde{f} . The best possible s -term approximation to \tilde{f} in $\mathcal{B}_{N,d}$ will be denoted

³Given that we will be recovering f based on point samples we will require at least enough smoothness to guarantee that any particular point sample we might possibly utilize actually contains information about the given function's basis coefficients $\{c_n\}_{n \in \mathbb{N}^D}$. Of course, the details regarding this smoothness requirement will vary with the choice of basis \mathcal{B}' .

by $\tilde{f}_s^{\text{opt}} : \mathcal{D} \rightarrow \mathbb{C}$, and will be defined as follows: Order the basis coefficients $\tilde{\mathbf{c}} \in \mathbb{C}^{|\mathcal{I}_{N,d}|}$ of \tilde{f} by their magnitudes so that

$$|\tilde{c}_{\mathbf{n}_1}| \geq |\tilde{c}_{\mathbf{n}_2}| \geq |\tilde{c}_{\mathbf{n}_3}| \geq \cdots \geq |\tilde{c}_{\mathbf{n}_{|\mathcal{I}_{N,d}|}}|,$$

where ties are broken lexicographically using the entries' indices in $\mathcal{I}_{N,d}$. Then \tilde{f}_s^{opt} will be defined to be

$$\tilde{f}_s^{\text{opt}}(\boldsymbol{\xi}) := \sum_{j=1}^s \tilde{c}_{\mathbf{n}_j} T_{\mathbf{n}_j}(\boldsymbol{\xi})$$

for all $\boldsymbol{\xi} \in \mathcal{D}$, and its (potentially) nonzero coefficients' indices will be denoted by

$$\Omega_{\tilde{f},s}^{\text{opt}} := \{\mathbf{n}_1, \dots, \mathbf{n}_s\} \subset \mathcal{I}_{N,d}.$$

Note that \tilde{f}_s^{opt} will indeed have the property that

$$\left\| \tilde{f} - \tilde{f}_s^{\text{opt}} \right\|_{L^2(\mathcal{D}, \mu)} = \inf_{\mathbf{z} \in \mathbb{C}^{|\mathcal{I}_{N,d}|}, \|\mathbf{z}\|_0 \leq s} \left\| \tilde{f} - \sum_{\mathbf{n} \in \mathcal{I}_{N,d}} z_{\mathbf{n}} T_{\mathbf{n}} \right\|_{L^2(\mathcal{D}, \mu)}.$$

Furthermore, if we let $\tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \in \mathbb{C}^{|\mathcal{I}_{N,d}|}$ denote the $\mathcal{B}_{N,d}$ basis coefficients of \tilde{f}_s^{opt} then we can see that both $\left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_0 \leq s$ and $\left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 = \left\| \tilde{f} - \tilde{f}_s^{\text{opt}} \right\|_{L^2(\mathcal{D}, \mu)}$ will hold. As a result, norms involving the vector $\tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \in \mathbb{C}^{|\mathcal{I}_{N,d}|}$ can be interpreted as best s -term approximation errors of \tilde{f} in a natural way.

Finally, to prove our main result below we will effectively be considering the point samples we take from f to instead be point samples taken from \tilde{f} that are contaminated with evaluation errors of size $(f - \tilde{f})(\boldsymbol{\xi})$ at each evaluation point $\boldsymbol{\xi} \in \mathcal{G}$. To bound all of these errors in a uniform fashion we will define

$$\gamma := \left\| f - \tilde{f} \right\|_{\infty} = \sup_{\boldsymbol{\xi} \in \mathcal{D}} \left| (f - \tilde{f})(\boldsymbol{\xi}) \right|. \quad (1.9)$$

The following theorem is proven in Section 3.

Theorem 1. (Main Result). *Let $\eta \in (0, \infty)$ and $s, d, N \in \mathbb{N} \setminus \{1\}$ with $d \leq D$ and $s < |\mathcal{I}_{N,d}|/2$. There exists a finite set of grid points $\mathcal{G} \subset \mathcal{D}$, an algorithm $\mathcal{H} : \mathbb{C}^{|\mathcal{G}|} \rightarrow (\mathcal{I}_{N,d} \times \mathbb{C})^s$, and an absolute universal constant $C' \in \mathbb{R}^+$ such that the function $a : \mathcal{D} \rightarrow \mathbb{C}$ defined by $a(\boldsymbol{\xi}) := \sum_{(\mathbf{n}, a_{\mathbf{n}}) \in \mathcal{H}(f(\mathcal{G}))} a_{\mathbf{n}} T_{\mathbf{n}}(\boldsymbol{\xi})$ satisfies*

$$\|f - a\|_{L^2(\mathcal{D}, \mu)} \leq \left\| f - \tilde{f} \right\|_{L^2(\mathcal{D}, \mu)} + C' \left(\sqrt{s} \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_1 + \gamma \sqrt{s} \right) + \eta \quad (1.10)$$

for all $f = \sum_{\mathbf{n} \in \mathbb{N}^D} c_{\mathbf{n}} T_{\mathbf{n}} \in L^2(\mathcal{D}, \mu)$ with $\gamma := \left\| f - \tilde{f} \right\|_{\infty} = \sup_{\boldsymbol{\xi} \in \mathcal{D}} \left| (f - \tilde{f})(\boldsymbol{\xi}) \right| < \infty$, where $\tilde{f} : \mathcal{D} \rightarrow \mathbb{C}$ is the finite dimensional approximation to f defined as per (1.8).

If the BOPB $\mathcal{B}_{N,d}$ is of type I so that the BOS constants K_j are 1 for all but at most $\tilde{d} \in \mathbb{Z} \cap [0, D]$ BOS basis sets $\mathcal{B}_{j;N}$, then

$$|\mathcal{G}| = \mathcal{O} \left(s^3 D K_{\infty}^{4\tilde{d}} d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right),$$

and the algorithm \mathcal{H} will have runtime complexity

$$\mathcal{O} \left((s^5 + s^3 N) D^2 K_{\infty}^{4\tilde{d}} d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \log \left(\left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 / \eta \right) \right).$$

If the BOPB $\mathcal{B}_{N,d}$ is of type II so that $K_0 = 1$, then

$$|\mathcal{G}| = \mathcal{O} \left(s^3 D K_{\infty}^{4d} d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right),$$

and the algorithm \mathcal{H} will have runtime complexity

$$\mathcal{O} \left((s^5 + s^3 N) D^2 K_{\infty}^{4d} d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \log \left(\left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 / \eta \right) \right).$$

Here we have assumed that any desired basis function $T_{\mathbf{n}} \in \mathcal{B}_{N,d}$ can be evaluated at any desired point in \mathcal{D} in $\mathcal{O}(ND)$ -time (which will be the case, e.g., for polynomial product bases of degree $\leq ND$).

Proof. This is a restatement of Corollary 1. □

Looking at Theorem 1 we can see that it effectively subsumes the theoretical recovery results of [12]. Consider, for example, the case where $f = \tilde{f} = \tilde{f}_s^{\text{opt}}$ (so that f is exactly s -sparse in $\mathcal{B}_{N,d}$). In this setting we will have both $\left\| f - \tilde{f} \right\|_{L^2(\mathcal{D}, \mu)} = \gamma = 0$ and $\tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} = \mathbf{0}$ hold true so that (1.10) implies that f is recovered exactly (up to any chosen tolerance η). Unlike the results in [12], however, Theorem 1 also guarantees that the method \mathcal{H} will work well for functions f which have $\|f - \tilde{f}_s^{\text{opt}}\|_{L^2(\mathcal{D}, \mu)}$ relatively small, but nonzero.

The authors would also like to emphasize the generality of Theorem 1, which is unique to the best of their knowledge in the literature related to sublinear-time sparse approximation methods. If, for example, one chooses $\mathcal{B}_{N,d}$ to be the multidimensional Fourier basis with $d = D$ and $\tilde{d} = 0$ then one immediately obtains a new sparse Fourier transform result for functions of many variables whose sampling and runtime requirements scale only

polylogarithmically in the total basis size $|\mathcal{B}_{N,d}| = N^D$. Though this new Fourier result does not compare favorably to the best deterministic multidimensional Fourier results of this kind [26, 32] with respect to achievable runtimes or error guarantees, it also does not use any of the specific algebraic structure of the Fourier basis. This allows Theorem 1 to be significantly more flexible than these older Fourier results in that it can apply to situations which they don't cover. For example, it can generate entirely discrete Fourier results where $\mathcal{G} \subset \{\frac{j}{N} \mid j \in [N]\}^D$ (unlike [26]) by using a discrete and finite multidimensional Fourier BOBP with $f = \tilde{f}$ which will work for any choice of N (unlike [32], which requires N to be prime).

Finally, the astute reader has likely noticed that Theorem 1 is phrased in the form of an existence result, which may be troubling to the practical numerical analyst who actually wants to know how to compute an accurate solution. Let us allay any anxieties that this choice of theoretical statement may have birthed – the algorithm \mathcal{H} referred to above is a modified version of the well known CoSaMP algorithm [33] (see Algorithm 1 in Section 3). It has been implemented and evaluated in Section 5, and the code made publicly available.⁴ In short, the result is entirely explicit and constructive with respect to the algorithm \mathcal{H} . The grid $\mathcal{G} \subset \mathcal{D}$, which ultimately responsible for the form of Theorem 1 as an existence result, on the other hand, is a bit more nuanced with respect to its practical construction.

As we shall see below, the grid $\mathcal{G} \subset \mathcal{D}$ is constructed by randomly selecting points from \mathcal{D} according to several prescribed probability distributions that are ultimately derived from the orthogonality measure μ (see Theorems 3 and 6 and their proofs for details). It is then proven that this randomly constructed grid will allow \mathcal{H} to satisfy the error guarantee (1.10) for all functions f as per (1.7) with high probability while simultaneously satisfying the stated upper bounds on its cardinality. The runtime complexity of \mathcal{H} follows from the boundedness of $|\mathcal{G}|$. Hence, the existence result is proven by randomly constructing a grid \mathcal{G} which is guaranteed to satisfy the conclusions of Theorem 1 with high probability.

In fact, this is completely analogous to the role of random sampling matrices in standard compressive sensing results involving the Restricted Isometry Property (RIP). Many compressive sensing methods are guaranteed to be accurate if they are used in combination with a random sampling matrix that has the RIP, a condition which can only be achieved near-optimally with high probability. Herein, the conclusions of Theorem 1 will hold for any grid \mathcal{G} that can be used to form two associated random sampling matrices: one with the RIP, and another with a property known as the *Support Identification Property (SIP)* which will be defined in Section 1.3. The conclusions of Theorem 1 will hold whenever these two conditions are satisfied by \mathcal{G} , and it will be shown that a randomly constructed

⁴See “SHT II: Best s-Term Approximation Guarantees for Bounded Orthonormal Product Bases in Sublinear-Time” on Mark Iwen’s code page <https://www.math.msu.edu/~markiwen/Code.html>.

grid \mathcal{G} will satisfy both conditions with high probability.

1.2 An Outline of the Paper and of the Proof of Theorem 1

After reviewing some relevant compressive sensing results and establishing necessary notation in Section 2, we will begin proving Theorem 1 in Section 3. The first step in that process will be to prove a compressive sensing recovery guarantee for a generalized version of the well known CoSaMP method [33]. This new theorem, Theorem 3, will establish a best s -term recovery guarantee for the CoSaMP algorithm where the support identification step is performed by any algorithm \mathcal{A} and grid \mathcal{G} pair which has the SIP (see Section 1.3 below for details on the SIP). With Theorem 3 in hand we will then turn our attention to constructing a sublinear-time algorithm \mathcal{A} and grid \mathcal{G} pair that have the SIP, an effort whose results are summarized by Theorem 4 (see also Proposition 1). Combining Theorems 3 and 4 then quickly establishes our main result above which appears in the form of Theorem 5 and Corollary 1 in Section 3.

The vast majority of the effort in the paper will be focussed on proving Theorem 4 in Section 4. That is, to demonstrate that Algorithm 2 therein can be used together with a randomly constructed grid \mathcal{G} in order to effectively achieve the SIP with high probability. This is done by Theorem 6 (a specialized version of Theorem 9 in Section 4.1) which formalizes the random sampling strategy one must use in order to construct \mathcal{G} so that the SIP is achieved with high probability, and by Theorem 7 which translates the conclusions of Theorem 6 into a SIP-type statement. Theorem 9, in turn, follows from Theorem 8 which is proven in Section 4.2.

Finally, the authors would like to note that the reader who is interested in seeing the proof of Theorem 1 unfold from basic compressive sensing principals in a more direct fashion (though without the benefit of waypoints explaining the relevance of each result to the final goal) might consider the following alternate reading order for the sections below: Such readers can begin with Section 4.2 after reviewing Section 2, followed by the first 4 paragraphs of Section 4, then Section 4.1, and finally the remainder of Section 4 after which Theorem 4 will have been proven. Reading Section 4 in this bottom up fashion first will then allow Section 3 to be read without having to temporarily take any of the theoretical statements therein for granted along the way. For readers who are mostly interested in the numerical ramifications of the methods developed herein, we suggest skipping down to conduct a careful review of Algorithms 1 and 2 (together with the equations referred to therein) after reading Section 2, after which the careful numerical evaluation conducted in Section 5 should be understandable.

Before moving on to establish some additional required notation, however, we will

first discuss the SIP in the next subsection. This is crucial as the notion of the SIP will allow for easier sublinear-time methods to be developed in the future. To emphasize this last point: Any basis for which the SIP can be established via a sublinear-time algorithm can be combined with Theorem 3 below in order to produce a new sublinear-time compressive sensing method for that basis. We expect that this new pathway for developing future sublinear-time algorithms will help to stimulate the further improvement and generalization of sparse Fourier transform techniques to other bases of interest going forward.

1.3 The Support Identification Property (SIP)

As above, let $[N] := \{0, \dots, N - 1\}$ for all $N \in \mathbb{N}$ and further define $\mathcal{P}([N])$ to be the power set of any such set $[N]$. In Section 3 we will prove that CoSaMP will still produce accurate sparse approximations as long as its support identification step employs a triple with the *support identification property*.

Definition 1 (The Support Identification Property (SIP)). *Let $s \in [N]$, $\beta \in (0, 1)$, $\Phi \in \mathbb{C}^{m \times N}$, $\mathcal{A} : \mathbb{C}^m \rightarrow \mathcal{P}([N])$, and $\Gamma : \mathbb{C}^m \rightarrow [0, \infty)$ with $\Gamma(\mathbf{0}) = 0$. The triple $(\Phi, \mathcal{A}, \Gamma)$ is said to have the Support Identification Property (SIP) of order (s, β) if*

$$\|\mathbf{v}_{\mathcal{A}(\Phi\mathbf{v}+\mathbf{e})^c}\|_2 \leq \beta\|\mathbf{v}\|_2$$

holds for all $\mathbf{e} \in \mathbb{C}^m$ and $\mathbf{v} \in \mathbb{C}^N$ with $\|\mathbf{v}\|_0 \leq s$ that also satisfy $\|\mathbf{v}\|_2 > \Gamma(\mathbf{e})$.

Note that many triples with the SIP exist. One completely trivial example is the triple consisting of the $N \times N$ identity matrix I , the function \mathcal{A} which always outputs $[N]$, and the zero function Γ . Of course this example is extremely unsatisfying – generally for compressive sensing applications we prefer that the any SIP triple $(\Phi \in \mathbb{C}^{m \times N}, \mathcal{A} : \mathbb{C}^m \rightarrow \mathcal{P}([N]), \Gamma)$ has $m \ll N$ and an efficient computational complexity for \mathcal{A} (preferably sublinear-in- N herein). Thankfully these types of SIP triples also exist – in fact it is easy to see that any fast and error-robust compressive sensing algorithm \mathcal{A} must in fact be a member of such a triple.

Lemma 1. *Let $\Gamma' : \mathbb{C}^m \rightarrow \mathbb{R}^+$ be such that $\Gamma'(\mathbf{0}) = 0$, and let $\mathcal{A} : \mathbb{C}^m \rightarrow \mathbb{C}^N$ be a compressive sensing algorithm with an associated measurement matrix $\Phi \in \mathbb{C}^{m \times N}$ that satisfies*

$$\|\mathbf{x} - \mathcal{A}(\Phi\mathbf{x} + \mathbf{e})\|_2 \leq \Gamma'(\mathbf{e})$$

for all $\mathbf{e} \in \mathbb{C}^m$ and $\mathbf{x} \in \mathbb{C}^N$ with $\|\mathbf{x}\|_0 \leq s$. Furthermore, let $\text{supp} : \mathbb{C}^N \rightarrow \mathcal{P}([N])$ output the indices of the nonzero entries of any given input vector. Then, the triple $(\Phi, \text{supp} \circ \mathcal{A}, (1/\beta')\Gamma')$ will have the SIP of order (s, β') for all $\beta' < 1$.

Proof. Let $\mathbf{e} \in \mathbb{C}^m$ and note that

$$\begin{aligned} \|\mathbf{x}_{\text{supp}(\mathcal{A}(\Phi\mathbf{x}+\mathbf{e}))^c}\|_2^2 &\leq \|\mathbf{x}_{\text{supp}(\mathcal{A}(\Phi\mathbf{x}+\mathbf{e}))}\|_2^2 + \left\| (\mathbf{x} - \mathcal{A}(\Phi\mathbf{x} + \mathbf{e}))_{\text{supp}(\mathcal{A}(\Phi\mathbf{x}+\mathbf{e}))} \right\|_2^2 \\ &= \|\mathbf{x} - \mathcal{A}(\Phi\mathbf{x} + \mathbf{e})\|_2^2 \leq (\Gamma'(\mathbf{e}))^2. \end{aligned}$$

Thus, if $\|\mathbf{x}\|_2 > (1/\beta')\Gamma'(\mathbf{e})$ then $\|\mathbf{x}_{\text{supp}(\mathcal{A}(\Phi\mathbf{x}+\mathbf{e}))^c}\|_2 \leq \beta'((1/\beta')\Gamma'(\mathbf{e})) \leq \beta'\|\mathbf{x}\|_2$. \square

Lemma 1 demonstrates that many nontrivial SIP triples of the type we are interested in exist. Of course, using a compressive sensing method in order to create a SIP triple seems slightly nonsensical given that one would generally want to create a SIP triple in order to develop a new compressive sensing method in the first place. This immediately raises the question of whether nontrivial SIP triples exist which do not in themselves already effectively serve as a compressive sensing method. The answer to that question is “yes”, and the easiest example is the SIP triple which the original CoSaMP method is itself already implicitly utilizes. Given $\mathbf{x} \in \mathbb{C}^N$ and $s \in [N]$ let $\mathbf{x}|_s \in \mathbb{C}^N$ be the vector obtained from \mathbf{x} by setting all but its s -largest magnitude entries to 0. The following lemma explicitly demonstrates the SIP triple on which the original CoSaMP algorithm [33] is implicitly based.

Lemma 2 (The CoSaMP SIP Triple). *Let $\Phi \in \mathbb{C}^{m \times N}$ have the RIP of order $(2s, 0.1)$ (so that its Restricted Isometry Constants (RIC)s satisfy $\delta_s \leq \delta_{2s} \leq 0.1$), $s \in [N]$, $\beta \in (0.2223, 1)$, and define $\mathcal{A} : \mathbb{C}^m \rightarrow \mathcal{P}([N])$ by $\mathcal{A}(\mathbf{y}) := \text{supp}((\Phi^*\mathbf{y})|_s)$, and $\Gamma : \mathbb{C}^m \rightarrow [0, \infty)$ by $\Gamma(\mathbf{e}) := \left(\frac{2.34}{\beta-0.2223}\right)\|\mathbf{e}\|_2$. Then, the triple $(\Phi, \mathcal{A}, \Gamma)$ has the SIP of order (s, β) .*

Proof. Lemma 4.2 of [33] implies that $\|\mathbf{x}_{\mathcal{A}(\Phi\mathbf{x}+\mathbf{e})^c}\|_2 \leq 0.2223\|\mathbf{x}\|_2 + 2.34\|\mathbf{e}\|_2$ holds for all $\mathbf{e} \in \mathbb{C}^m$ and $\mathbf{x} \in \mathbb{C}^N$ with $\|\mathbf{x}\|_0 \leq s$. Suppose, furthermore, that $\|\mathbf{x}\|_2 > \Gamma(\mathbf{e}) = \left(\frac{2.34}{\beta-0.2223}\right)\|\mathbf{e}\|_2$. Then,

$$\begin{aligned} \|\mathbf{x}_{\mathcal{A}(\Phi\mathbf{x}+\mathbf{e})^c}\|_2 &\leq 0.2223\|\mathbf{x}\|_2 + 2.34\|\mathbf{e}\|_2 < 0.2223\|\mathbf{x}\|_2 + 2.34\left(\frac{\beta - 0.2223}{2.34}\right)\|\mathbf{x}\|_2 \\ &= \beta\|\mathbf{x}\|_2 \end{aligned}$$

holds. Also, $\Gamma(\mathbf{0}) = 0$. \square

In Section 3 we will demonstrate that the original SIP triple implicitly used by the CoSaMP algorithm can be replaced with any other SIP triple of similar quality without substantively changing the performance of the resulting CoSaMP variant as a compressive sensing algorithm. Before we can do this, however, we will require some additional notation and preliminary infrastructural results.

2 Notation and Preliminaries

Recall that $D \in \mathbb{N}$ is the number of variables in the function of interest $f : \times_{j \in [D]} \mathcal{D}_j \rightarrow \mathbb{C}$ (where $\mathcal{D}_j \subset \mathbb{R}$ for all $j \in [D]$, and $\mathcal{D} := \times_{j \in [D]} \mathcal{D}_j$). Vectors $\mathbf{n} \in [N]^D$ with $\|\mathbf{n}\|_0 \leq d \leq D$ will always index a basis function in

$$\mathcal{B} := \mathcal{B}_{N,d} = \{T_{\mathbf{n}} : \mathcal{D} \rightarrow \mathbb{C} \mid \mathbf{n} \in [N]^D \text{ with } \|\mathbf{n}\|_0 \leq d\}, \quad (2.1)$$

where we have suppressed the basis subscripts for ease of discussion. In addition, we further assume that the BOS \mathcal{B} is a product basis so that $T_{\mathbf{n}}(\boldsymbol{\xi})$ satisfies (1.3) as above.

2.1 Restrictions and Partial Evaluations

The following notation will be utilized heavily during the analysis of the proposed support identification procedure. Let $\mathcal{S} \subset [D]$, $\mathbf{w} \in \times_{j \in \mathcal{S}} \mathcal{D}_j$ with $w_j \in \mathcal{D}_j$, and $\mathbf{n} \in [N]^D$. The function $T_{\mathcal{S};\mathbf{n}} : \times_{j \in \mathcal{S}} \mathcal{D}_j \rightarrow \mathbb{C}$ is defined to be

$$T_{\mathcal{S};\mathbf{n}}(\mathbf{w}) := \prod_{j \in \mathcal{S}} T_{j;n_j}(w_j). \quad (2.2)$$

Then, the set

$$\mathcal{B}_{\mathcal{S}} := \{T_{\mathcal{S};\mathbf{n}} \mid \mathbf{n} \in [N]^D \text{ with } \|\mathbf{n}\|_0 \leq d\} \quad (2.3)$$

is a BOS with respect to the probability measure $\mu_{\mathcal{S}} := \otimes_{j \in \mathcal{S}} \mu_j$ over $\mathcal{D}_{\mathcal{S}} := \times_{j \in \mathcal{S}} \mathcal{D}_j \subset \mathbb{R}^{|\mathcal{S}|}$ with BOS constant $K_{\mathcal{S}} \leq \min \left\{ \prod_{j \in \mathcal{S}} K_j, K_{\infty}^d K_0^{\max\{|\mathcal{S}|-d, 0\}} \right\}$. For any set E , let $\mathcal{P}(E)$ denote the power set of E containing all possible subsets of E . Given any vector $\mathbf{v} \in \mathbb{C}^p$ and $\mathcal{T} \subset [p]$ we will let $\mathbf{v}_{\mathcal{T}} \in \mathbb{C}^p$ have entries

$$(v_{\mathcal{T}})_j = \begin{cases} v_j & \text{if } j \in \mathcal{T} \\ 0 & \text{if } j \notin \mathcal{T} \end{cases}.$$

For $t \in [p]$, we let $\mathbf{v}_t \in \mathbb{C}^p$ a vector restricting \mathbf{v} to its t largest-magnitude entries. Let $\mathcal{S}^c := [D] \setminus \mathcal{S}$ for all $\mathcal{S} \subset [D]$. We will then construct $f_{\mathcal{S};\mathbf{w}} : \mathcal{D}_{\mathcal{S}^c} \rightarrow \mathbb{C}$ from $f : \mathcal{D} \rightarrow \mathbb{C}$ by defining

$$f_{\mathcal{S};\mathbf{w}}(\mathbf{z}) := f(\boldsymbol{\xi})$$

where $\boldsymbol{\xi} \in \mathcal{D} \subset \mathbb{R}^D$ is the unique vector with $\boldsymbol{\xi}_{\mathcal{S}} = \mathbf{w}$ and $\boldsymbol{\xi}_{\mathcal{S}^c} = \mathbf{z}$. In this context, we define the permutation function $\varrho_{\mathcal{S}} : \mathcal{D}_{\mathcal{S}} \times \mathcal{D}_{\mathcal{S}^c} \rightarrow \mathcal{D}$ given by

$$\varrho_{\mathcal{S}}(\mathbf{w}, \mathbf{z}) = \boldsymbol{\xi} \text{ such that } \boldsymbol{\xi}_{\mathcal{S}} = \mathbf{w} \text{ and } \boldsymbol{\xi}_{\mathcal{S}^c} = \mathbf{z}. \quad (2.4)$$

This yields the alternative characterization $f_{\mathcal{S};\mathbf{w}}(\mathbf{z}) = f(\varrho_{\mathcal{S}}(\mathbf{w}, \mathbf{z}))$.

The restricted vectors of the input vector $\boldsymbol{\xi}$ such as $\boldsymbol{\xi}_{\mathcal{S}}$ and $\boldsymbol{\xi}_{\mathcal{S}^c}$ have the reduced dimensions. However, the coefficient vectors such as $\tilde{\mathbf{c}}$ and $\tilde{\mathbf{r}}$ will maintain the full dimension even though they are restricted to some subset of indices.

If $\mathbf{n}, \mathbf{m} \in \mathcal{I}_{N,d}$ and $\mathcal{S} \subset [D]$ then we will define $(\mathbf{n}, \mathbf{m})_{\mathcal{S}} \in \mathcal{I}_{N,d}$ to be the vector $\mathbf{n}_{\mathcal{S}} + \mathbf{m}_{\mathcal{S}^c}$. Furthermore, for a given $\mathbf{v} \in \mathbb{C}^{|\mathcal{I}_{N,d}|}$, $\mathbf{n} \in \mathcal{I}_{N,d}$, and $\mathcal{S} \subset [D]$, we will let the vector $\mathbf{v}_{\mathcal{S};\mathbf{n}} \in \mathbb{C}^{|\mathcal{I}_{N,d}|}$ indexed by $\mathbf{k} \in \mathcal{I}_{N,d}$ have entries given by

$$(v_{\mathcal{S};\mathbf{n}})_{\mathbf{k}} = \begin{cases} v_{\mathbf{k}}, & \text{if } \mathbf{k}_{\mathcal{S}} = \mathbf{n}_{\mathcal{S}} \\ 0 & \text{otherwise} \end{cases}. \quad (2.5)$$

Note that $\mathbf{v}_{\mathcal{S};\mathbf{n}} \in \mathbb{C}^{|\mathcal{I}_{N,d}|}$ will only have at most $N^{|\mathcal{S}^c|}$ nonzero entries corresponding to the entries of $\mathbf{v} \in \mathbb{C}^{|\mathcal{I}_{N,d}|}$, $v_{\mathbf{m}} \in \mathbb{C}$, whose indices \mathbf{m} match those of \mathbf{n} on \mathcal{S} (i.e., so that $\mathbf{m}_{\mathcal{S}} = \mathbf{n}_{\mathcal{S}}$).

The following calculation will be repeated sufficiently often that it merits being referred to as a lemma. It concerns the partial sum approximation to f from (1.7) in \mathcal{B} given by (1.8). Recall that $\tilde{\mathbf{c}}$ contains only the entries of the sequence \mathbf{c} corresponding to the indices in $\mathcal{I}_{N,d}$.

Moreover, the next lemma also demonstrates the usage of the newly introduced notation. Its statement will be used later in the proofs of Lemmas 10 and 11.

Lemma 3. *Let $\mathcal{S} \subset [D]$, $\mathbf{w} \in \mathcal{D}_{\mathcal{S}} = \times_{j \in \mathcal{S}} \mathcal{D}_j$ with $w_j \in \mathcal{D}_j$, and $\mathbf{n} \in \mathcal{I}_{N,d}$. Then*

$$\left\langle \tilde{f}_{\mathcal{S};\mathbf{w}}, T_{\mathcal{S}^c;\mathbf{n}} \right\rangle_{(\mathcal{D}_{\mathcal{S}^c}, \mu_{\mathcal{S}^c})} = \left\langle \tilde{\mathbf{c}}_{\mathcal{S}^c;\mathbf{n}}, \overline{\Phi_{\mathcal{S};\mathbf{n};\mathbf{w}}} \right\rangle$$

where \tilde{f} is as in (1.8), and $\Phi_{\mathcal{S};\mathbf{n};\mathbf{w}} \in \mathbb{C}^{|\mathcal{I}_{N,d}|}$ is a vector indexed by $\mathbf{k} \in \mathcal{I}_{N,d}$ with entries

$$(\Phi_{\mathcal{S};\mathbf{n};\mathbf{w}})_{\mathbf{k}} := \begin{cases} T_{\mathcal{S};\mathbf{k}}(\mathbf{w}) & \text{if } \mathbf{k}_{\mathcal{S}^c} = \mathbf{n}_{\mathcal{S}^c} \\ 0 & \text{otherwise} \end{cases}. \quad (2.6)$$

Proof. Computing the inner product one quickly sees that

$$\begin{aligned} \left\langle \tilde{f}_{\mathcal{S};\mathbf{w}}, T_{\mathcal{S}^c;\mathbf{n}} \right\rangle_{(\mathcal{D}_{\mathcal{S}^c}, \mu_{\mathcal{S}^c})} &= \int_{\mathcal{D}_{\mathcal{S}^c}} \tilde{f}_{\mathcal{S};\mathbf{w}}(\mathbf{z}) \overline{T_{\mathcal{S}^c;\mathbf{n}}(\mathbf{z})} d\mu_{\mathcal{S}^c}(\mathbf{z}) \\ &= \int_{\mathcal{D}_{\mathcal{S}^c}} \left(\sum_{\mathbf{k} \in \mathcal{I}_{N,d}} \tilde{c}_{\mathbf{k}} T_{\mathcal{S};\mathbf{k}}(\mathbf{w}) T_{\mathcal{S}^c;\mathbf{k}}(\mathbf{z}) \right) \overline{T_{\mathcal{S}^c;\mathbf{n}}(\mathbf{z})} d\mu_{\mathcal{S}^c}(\mathbf{z}) \end{aligned}$$

$$= \sum_{\mathbf{k} \in \mathcal{I}_{N,d}} (\tilde{c}_{\mathcal{S}^c; \mathbf{n}})_{\mathbf{k}} T_{\mathcal{S}; \mathbf{k}}(\mathbf{w}).$$

The stated result follows. \square

Let $m \in \mathbb{N}$ and $\mathbf{n} \in \mathcal{I}_{N,d}$. For any matrix $A \in \mathbb{C}^{m \times |\mathcal{I}_{N,d}|}$, we define $(A)_{\mathbf{n}}$ be the column of A corresponding to the index \mathbf{n} . Also, we can choose multiple columns, e.g., for $\mathbf{n}_1, \mathbf{n}_2 \in \mathcal{I}_{N,d}$, $(A)_{\{\mathbf{n}_1, \mathbf{n}_2\}}$ refers the columns of A corresponding to the indices \mathbf{n}_1 and \mathbf{n}_2 . More generally, for any $\mathcal{S} \subset \mathcal{I}_{N,d}$ the matrix $(A)_{\mathcal{S}} = A_{\mathcal{S}} \in \mathbb{C}^{m \times |\mathcal{S}|}$ will consist of the columns of A indexed by \mathcal{S} .

2.2 Sampling Matrices associated to a BOS and Restricted Isometry Constants

Given a BOS as in (2.1), let $\{\boldsymbol{\xi}_{\ell}\}_{\ell \in [m]} \subset \mathcal{D}$ be sampling points drawn independently at random according to μ with corresponding samples $\{\tilde{\mathbf{y}}_{\ell} := \tilde{f}(\boldsymbol{\xi}_{\ell})\}_{\ell \in [m]}$ from \tilde{f} in (1.8). The *random sampling matrix* $\Phi \in \mathbb{C}^{m \times |\mathcal{I}_{N,d}|}$ associated with the points $\{\boldsymbol{\xi}_{\ell}\}_{\ell \in [m]}$ and the BOS has entries given by

$$\Phi_{\ell, \mathbf{n}} = T_{\mathbf{n}}(\boldsymbol{\xi}_{\ell}) \quad (2.7)$$

with indices $\ell \in [m]$ and $\mathbf{n} \in \mathcal{I}_{N,d}$. One can see that, e.g., $\tilde{\mathbf{y}} = \Phi \tilde{\mathbf{c}}$ will hold in this case. Furthermore, results from the compressive sensing literature guarantee that $\frac{1}{\sqrt{m}}\Phi$ will also have well-behaved *restricted isometry constants* as soon as m is sufficiently large.

Definition 2 (See Definition 6.1 in [17]). *The s -th restricted isometry constant δ_s of a matrix $A \in \mathbb{C}^{m \times N}$ is the smallest $\delta \geq 0$ such that*

$$(1 - \delta)\|\mathbf{x}\|_2^2 \leq \|A\mathbf{x}\|_2^2 \leq (1 + \delta)\|\mathbf{x}\|_2^2$$

holds for all s -sparse vectors $\mathbf{x} \in \mathbb{C}^N$. The matrix A is said to satisfy the restricted isometry property (RIP) of order (s, δ) if $1 > \delta \geq \delta_s \geq 0$.

Theorem 2 (See Theorem 12.32 and Remark 12.33 in [17]). *Let $\Phi \in \mathbb{C}^{m \times |\mathcal{I}_{N,d}|}$ be the random sampling matrix associated to a BOS with constant $K \geq 1$ for $m, |\mathcal{I}_{N,d}|, s \in \mathbb{Z}^+ \setminus \{1\}$. If, for $\delta, p \in (0, 1)$,*

$$m \geq aK^2\delta^{-2}s \cdot \max \left\{ \ln^2(s) \ln(|\mathcal{I}_{N,d}|) \ln(m), \ln(p^{-1}) \right\},$$

then with probability at least $1 - p$ the restricted isometry constant δ_s of $\frac{1}{\sqrt{m}}\Phi$ satisfies $\delta_s \leq \delta$ so that Φ has the RIP of order (s, δ) . Here the constant $a \in \mathbb{R}^+$ is universal.

In addition, one can also, e.g., bound the ℓ_2 operator norm of the random sampling matrix Φ in the event that it has the RIP. We have the following consequence of Theorem 2.

Lemma 4 (See Proposition 3.5 in [33]). *Suppose $A \in \mathbb{C}^{m \times |\mathcal{I}_{N,d}|}$ has the restricted isometry property (RIP) of order (s, δ) . Then,*

$$\|A\mathbf{x}\|_2 \leq \sqrt{1 + \delta} \left(\frac{\|\mathbf{x}\|_1}{\sqrt{s}} + \|\mathbf{x}\|_2 \right)$$

holds for all $\mathbf{x} \in \mathbb{C}^{|\mathcal{I}_{N,d}|}$.

We are now prepared to develop the new CoSaMP variants on which our new sublinear-time algorithms will be based.

3 Robust Sublinear-Time Sparse Approximation via CoSaMP with Fast Support Identification

Algorithm 1 CoSaMP with the new support identification

```

1: procedure CoSaMPnewSupportID
2: Input:  $\mathbf{y}_{\text{SID}} = \Phi_{\text{SID}}\mathbf{x}_s + \mathbf{e}_{\text{SID}}$ ,  $\Phi_{\text{SID}}$ ,  $\mathcal{A}$ ,  $\mathbf{y}_{\text{CE}} = \Phi_{\text{CE}}\mathbf{x}_s + \mathbf{e}_{\text{CE}}$ ,  $\Phi_{\text{CE}}$ ,  $\kappa$ ,  $s$ ,  $d$ ,  $\tilde{d}$ 
3: Output:  $s$ -sparse approximation  $\mathbf{a}$  of  $\mathbf{x}$ 
4:    $\mathbf{a}^0 = \mathbf{0}$  {Initial approximation}
5:    $\mathbf{v}_{\text{SID}} \leftarrow \mathbf{y}_{\text{SID}}$ 
6:    $k \leftarrow 0$ 
7:   repeat
8:      $k \leftarrow k + 1$ 
9:      $\tilde{\Omega} \leftarrow \mathcal{A}(\mathbf{v}_{\text{SID}})$  { $|\tilde{\Omega}| \leq 2s$ , New support identification step (e.g., Algorithm 2)}
10:     $\Omega \leftarrow \tilde{\Omega} \cup \text{supp}(\mathbf{a}^{k-1})$  {Merge supports}
11:     $\Phi' \leftarrow \frac{1}{\sqrt{m_{\text{CE}}}} \Phi_{\text{CE}}|_{\Omega}$ 
12:     $\mathbf{b}|_{\Omega} \leftarrow (\Phi')^\dagger \frac{\mathbf{y}_{\text{CE}}}{\sqrt{m_{\text{CE}}}}$  {approximated using 3 LS iterations (Richardson's or CG)}
13:     $\mathbf{a}^k \leftarrow (\mathbf{b}|_{\Omega})_s$  {Prune to obtain next approximation}
14:     $\mathbf{v}_{\text{SID}} \leftarrow \mathbf{y}_{\text{SID}} - \Phi_{\text{SID}}\mathbf{a}^k$  {Update current samples I}
15:     $\mathbf{v}_{\text{CEold}} \leftarrow \mathbf{v}_{\text{CE}}$ ,  $\mathbf{v}_{\text{CE}} \leftarrow \mathbf{y}_{\text{CE}} - \Phi_{\text{CE}}\mathbf{a}^k$  {Update current samples II}
16:    until  $\|\mathbf{v}_{\text{CE}}\|_2^2 > \|\mathbf{v}_{\text{CEold}}\|_2^2$ , or  $k \geq \kappa$  {Halting criteria}
17:    If  $\|\mathbf{v}_{\text{CE}}\|_2^2 > \|\mathbf{v}_{\text{CEold}}\|_2^2$  then  $\mathbf{a} \leftarrow \mathbf{a}^{k-1}$  else  $\mathbf{a} \leftarrow \mathbf{a}^k$ 
18: end procedure

```

In this section we analyze a generalized CoSaMP variant which uses any support identification method satisfying the SIP introduced in Definition 1 above (see Algorithm 1).

In Theorem 3 we provide error guarantees as well as the general sampling and runtime complexities that one can obtain for such CoSaMP variants with a particular choice of halting criteria. Later, in Section 4, we then propose a new admissible support identification method which runs in sublinear time for BOPBs with sufficiently small BOS constants (see Algorithm 2). This method is proven to satisfy the SIP as stated in Theorem 4 of this section. Finally, combining Algorithms 1 and 2, we obtain Theorem 5 which combines the error guarantees from Theorem 3 due to the SIP with the specific sampling and runtime complexities of the support identification algorithm presented in Section 4 for BOPBs. We hasten to point out that the modularity of this proof approach makes it easier to improve upon than prior works have been. If a better (e.g., faster) support identification method satisfying the SIP is developed for a particular basis in the future it can immediately replace the one from Section 4 and produce an improved CoSaMP type algorithm with a better performance for that particular basis.

We assume herein that the function f in (1.7) can be written as

$$f := \tilde{f} + e' \quad (3.1)$$

where $\tilde{f} : \mathcal{D} \rightarrow \mathbb{C}$ is as per (1.8) with the coefficient vector $\tilde{\mathbf{c}} \in \mathbb{C}^{\mathcal{I}_{N,d}}$ in \mathcal{B} , and where $e' : \mathcal{D} \rightarrow \mathbb{C}$ is bounded so that $\|e'\|_\infty \leq \gamma$. Now, we rewrite f as

$$f = \tilde{f} + e' = \tilde{f}_s^{\text{opt}} + \underbrace{\left(\tilde{f} - \tilde{f}_s^{\text{opt}} + e' \right)}_{=: e}. \quad (3.2)$$

Our goal is to approximate the best s -term approximation $\tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}}$ of $\tilde{\mathbf{c}}$, which is the coefficient vector of \tilde{f}_s^{opt} . Since CoSaMP from [33] approximates the best s -term of a given vector efficiently while allowing mild noise on the samples, we modify the CoSaMP algorithm in order to make it handle our high-dimensional problem more efficiently. Since the analysis of our CoSaMP type algorithm will be based on [33], it is helpful to introduce the connection between our notation and the notation from [33]. Toward that end, going forward we will set $\mathbf{x}_s := \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}}$, $\mathbf{x} := \tilde{\mathbf{c}}$, and $\Phi := \frac{1}{\sqrt{m_{\text{CE}}}} \Phi_{\text{CE}}$ in the notation of [33]. The samples $\mathbf{u} = \Phi \mathbf{x} + \mathbf{e}$ in [33] can then be viewed as containing renormalized function evaluations of f , and accordingly, \mathbf{e} contains renormalized function evaluations of the e defined in (3.2). In particular, $\mathbf{e} = \frac{\mathbf{e}_{\text{CE}}}{\sqrt{m_{\text{CE}}}}$.

In the pseudocode of Algorithm 1, most of the steps are identical to the original CoSaMP except the “New support identification step”, “Update current samples I & II”, and “Halting criteria” lines. The inputs $\mathbf{y}_{\text{SID}} \in \mathbb{C}^{m_{\text{SID}}}$ and $\mathbf{y}_{\text{CE}} \in \mathbb{C}^{m_{\text{CE}}}$ of Algorithm 1 contain function evaluations of f which will be used for support identification and coefficient estimation (i.e., via least squares), respectively. Accordingly, $\mathbf{e}_{\text{SID}} \in \mathbb{C}^{m_{\text{SID}}}$ and $\mathbf{e}_{\text{CE}} \in \mathbb{C}^{m_{\text{CE}}}$ appearing in Theorem 3 contain the corresponding function evaluations of

e from (3.2), and $\Phi_{\text{SID}} \in \mathbb{C}^{m_{\text{SID}} \times |\mathcal{I}_{N,d}|}$ and $\Phi_{\text{CE}} \in \mathbb{C}^{m_{\text{CE}} \times |\mathcal{I}_{N,d}|}$ have the function evaluations of $T_{\mathbf{n}}$ for $\mathbf{n} \in \mathcal{I}_{N,d}$ at the corresponding evaluation points. Note that \mathbf{e}_{SID} and \mathbf{e}_{CE} do not change over the iterations of Algorithm 1. We define $\mathbf{e}'_{\text{SID}} \in \mathbb{C}^{m_{\text{SID}}}$ and $\mathbf{e}'_{\text{CE}} \in \mathbb{C}^{m_{\text{CE}}}$ as the vectors whose entries are the function evaluations of e' from (3.1). Each row number (m_{SID} and m_{CE}) is, therefore, the total number of function evaluations used for support identification and the coefficient estimation, respectively. In the k -th iteration, Algorithm 1 starts with an s -sparse approximation \mathbf{a}^{k-1} of \mathbf{x}_s and then tries to approximate the at most $2s$ -sparse residual vector $\mathbf{r}^{k-1} := \mathbf{x}_s - \mathbf{a}^{k-1}$. The support identification procedure \mathcal{A} in the ‘‘New support identification’’ step begins approximating \mathbf{r}^{k-1} by finding a support set $\tilde{\Omega} \subset \mathcal{I}_{N,d}$ of cardinality at most $2s$ which contains the indices of the entries where most of the energy of \mathbf{r}^{k-1} is located. As noted above, any support identification method satisfying the SIP can substitute the ‘‘New support identification step’’ in Algorithm 1 in order to accomplish this task – the algorithm developed and analyzed in Section 4 is a specific instance.

After the support identification, in the ‘‘Merge supports’’ step, a new support set Ω of cardinality at most $3s$ is then formed from the union of $\tilde{\Omega}$ with the support of the current approximation \mathbf{a}^{k-1} . At this stage Ω should contain the overwhelming majority of the important (i.e., energetic) index vectors for \mathbf{r}^{k-1} . As a result, restricting the columns of the sampling matrix Φ_{CE} to those in Ω (or constructing them on the fly in a low memory setting) in order to solve for $\mathbf{b}_{\Omega} := \operatorname{argmin}_{\mathbf{u} \in \mathbb{C}^{|\Omega|}} \frac{1}{\sqrt{m_{\text{CE}}}} \|(\Phi_{\text{CE}})_{\Omega} \mathbf{u} - \mathbf{y}_{\text{CE}}\|_2$ should yield accurate estimates for the true coefficients of $\mathbf{x}_s = \tilde{\mathbf{c}}_{\Omega_{f,s}}^{\text{opt}}$ indexed by the elements of Ω ,

$\left(\tilde{\mathbf{c}}_{\Omega_{f,s}}^{\text{opt}} \right)_{\Omega}$.⁵ The vector $(\mathbf{b}_{\Omega})_s$ then becomes the next approximation \mathbf{a}^k of $\mathbf{x}_s = \tilde{\mathbf{c}}_{\Omega_{f,s}}^{\text{opt}}$.

Theorem 3 provides the error guarantees for $\left\| \tilde{\mathbf{c}}_{\Omega_{f,s}}^{\text{opt}} - \mathbf{a} \right\|_2$, as well as the runtime complexity of Algorithm 1 in terms of the provided support identification algorithm’s runtime.

Theorem 3. *Let $\bar{\Gamma} \geq 0$, $\beta \in (0, 0.2228]$, $\kappa \in \mathbb{N}$ and $\delta \in (0, 0.025]$ be fixed, and let K be the BOS constant of (1.6). Suppose that $\mathbf{x}_s := \tilde{\mathbf{c}}_{\Omega_{f,s}}^{\text{opt}}$ is s -sparse with $\mathbf{y}_{\text{SID}} = \Phi_{\text{SID}} \mathbf{x}_s + \mathbf{e}_{\text{SID}}$ and $\mathbf{y}_{\text{CE}} = \Phi_{\text{CE}} \mathbf{x}_s + \mathbf{e}_{\text{CE}}$ where the triple $(\Phi_{\text{SID}} \in \mathbb{C}^{m_{\text{SID}} \times |\mathcal{I}_{N,d}|}, \mathcal{A} : \mathbb{C}^{m_{\text{SID}}} \rightarrow \mathcal{P}(|\mathcal{I}_{N,d}|), \Gamma : \mathbb{C}^{m_{\text{SID}}} \rightarrow [0, \infty))$ has the SIP of order $(2s, \beta)$ and $\Gamma(\mathbf{e}_{\text{SID}}) \leq \bar{\Gamma}$, and $\frac{1}{\sqrt{m_{\text{CE}}}} \Phi_{\text{CE}} \in \mathbb{C}^{m_{\text{CE}} \times |\mathcal{I}_{N,d}|}$ has RIP constant $\delta_{2s} \leq \delta$ and $m_{\text{CE}} = \mathcal{O}(sK^2 \log^4 |\mathcal{I}_{N,d}|)$. Suppose that $\mathbf{r}^k := \mathbf{x}_s - \mathbf{a}^k$ is a $2s$ -sparse vector such that $\mathbf{v}_{\text{SID}} = \Phi_{\text{SID}}(\mathbf{x}_s - \mathbf{a}^k) + \mathbf{e}_{\text{SID}} = \Phi_{\text{SID}} \mathbf{r}^k + \mathbf{e}_{\text{SID}}$. Furthermore, suppose that the support identification procedure \mathcal{A} ’s output always has car-*

⁵In practice, it suffices to approximate the least-squares solution \mathbf{b}_{Ω} by an iterative least-squares approach such as Richardson’s iteration or conjugate gradient [6, 14] since computing the exact least squares solution can be expensive when s is large. The argument of [33] shows that it is enough to take three iterations for Richardson’s iteration or conjugate gradient if the initial condition is set to \mathbf{a}^{k-1} , and if Φ_{CE} has an RIP constant $\delta_{2s} < 0.025$. In fact, both of these methods have similar runtime performance.

dinality at most $2s$, that it runs in $\mathcal{O}(\mathcal{L}_{\mathcal{A}})$ -time, and that it uses $m_{\text{SID}} = \mathcal{O}(\mathcal{L}'_{\mathcal{A}})$ function evaluations. Then, for all $k \geq 0$, the signal approximation \mathbf{a}^k in Algorithm 1 is s -sparse and satisfies

$$\left\| \mathbf{x}_s - \mathbf{a}^{k+1} \right\|_2 \leq 0.5 \left\| \mathbf{x}_s - \mathbf{a}^k \right\|_2 + \frac{2.124}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2,$$

as long as $\|\mathbf{r}^k\|_2 = \|\mathbf{x}_s - \mathbf{a}^k\|_2 > \bar{\Gamma}$. In particular, if $\min_{j \in [k]} \|\mathbf{r}^j\|_2 > \bar{\Gamma}$ then

$$\left\| \mathbf{x}_s - \mathbf{a}^k \right\|_2 \leq 2^{-k} \|\mathbf{x}_s\|_2 + \frac{4.248}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2. \quad (3.3)$$

As a consequence, CoSaMP with any such support identification method \mathcal{A} will produce an s -sparse approximation \mathbf{a} that satisfies

$$\|\mathbf{x}_s - \mathbf{a}\|_2 \leq \max \left\{ 1.03\bar{\Gamma} + 2.03 \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}}, 2^{-\kappa} \|\mathbf{x}_s\|_2 + 4.3 \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}}, 8.625 \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}} \right\}. \quad (3.4)$$

The sampling complexity of Algorithm 1 will be $m_{\text{SID}} + m_{\text{CE}} = \mathcal{O}(\mathcal{L}'_{\mathcal{A}} + sK^2 \log^4 |\mathcal{I}_{N,d}|)$. The runtime complexity of Algorithm 1 will be $\mathcal{O}(\mathcal{L}_{\mathcal{A}} + s^2 K^2 \mathcal{L}_{\Phi} \log^4 |\mathcal{I}_{N,d}| + s \mathcal{L}_{\Phi} m_{\text{SID}}) \cdot \kappa$, where $\mathcal{O}(\mathcal{L}_{\Phi})$ is the runtime complexity of computing any desired matrix entry $(\Phi_{\text{CE}})_{j,\ell}$, or $(\Phi_{\text{SID}})_{j,\ell}$, for any valid choice of j, ℓ .

Proof. When $k < \kappa$ and $\|\mathbf{r}^k\|_2 > \bar{\Gamma} \geq \Gamma(\mathbf{e}_{\text{SID}})$, we obtain

$$\begin{aligned} \left\| \mathbf{x}_s - \mathbf{a}^{k+1} \right\|_2 &\leq 2 \|\mathbf{x}_s - \mathbf{b}\|_2 && \text{(Lemma 4.5 in [33], } \mathbf{a}^{k+1} = (\mathbf{b}|_{\Omega})_s, \mathbf{b} = \mathbf{b}|_{\Omega}) \\ &\leq 2.224 \left\| (\mathbf{x}_s)_{\Omega^c} \right\|_2 + 0.0044 \|\mathbf{r}^k\|_2 + \frac{2.124}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2 && \text{(Corollary 5.3 in [33])} \\ &\leq 2.224 \left\| \mathbf{r}_{\Omega^c}^k \right\|_2 + 0.0044 \|\mathbf{r}^k\|_2 + \frac{2.124}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2 && \text{(Lemma 4.3 in [33])} \\ &\leq 2.224 \cdot \beta \|\mathbf{r}^k\|_2 + 0.0044 \|\mathbf{r}^k\|_2 + \frac{2.124}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2 && \text{(The SIP assumption)} \\ &\leq 0.5 \|\mathbf{r}^k\|_2 + \frac{2.124}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2 \\ &= 0.5 \left\| \mathbf{x}_s - \mathbf{a}^k \right\|_2 + \frac{2.124}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2. \end{aligned}$$

In order to obtain the bound in (3.3) we may now simply solve the recursion for the final error after noting that

$$(1 + 0.5 + 0.25 + \dots) \cdot \frac{2.124}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2 = \frac{4.248}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2.$$

If the last $k \geq \kappa$ in Algorithm 1, and $\|\mathbf{r}^k\|_2 > \bar{\Gamma} \geq \Gamma(e_{\text{SID}})$ for all $k < \kappa$, then

$$\|\mathbf{x}_s - \mathbf{a}\|_2 \leq 2^{-\kappa} \|\mathbf{x}_s\|_2 + \frac{4.248}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2.$$

On the other hand, if the last $k \geq \kappa$ in Algorithm 1, $\|\mathbf{v}_{\text{CE}}\|_2 \leq \|\mathbf{v}_{\text{CEold}}\|_2$ for all $k < \kappa$, and $\|\mathbf{r}^k\|_2 \leq \bar{\Gamma}$ for some $k < \kappa$, then

$$\begin{aligned} \frac{1}{\sqrt{m_{\text{CE}}}} \|\mathbf{y}_{\text{CE}} - \Phi_{\text{CE}} \mathbf{a}^\kappa\|_2 &\leq \frac{1}{\sqrt{m_{\text{CE}}}} \left\| \mathbf{y}_{\text{CE}} - \Phi_{\text{CE}} \mathbf{a}^k \right\|_2 \\ &\leq \left\| \frac{1}{\sqrt{m_{\text{CE}}}} \Phi_{\text{CE}} (\mathbf{x}_s - \mathbf{a}^k) + \frac{\mathbf{e}_{\text{CE}}}{\sqrt{m_{\text{CE}}}} \right\|_2 \\ &\leq \left\| \frac{1}{\sqrt{m_{\text{CE}}}} \Phi_{\text{CE}} (\mathbf{x}_s - \mathbf{a}^k) \right\|_2 + \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}} \\ &\leq \sqrt{1 + \delta} \|\mathbf{x}_s - \mathbf{a}^k\|_2 + \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}} \\ &\leq \bar{\Gamma} \sqrt{1 + \delta} + \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}}, \end{aligned}$$

and

$$\begin{aligned} \frac{1}{\sqrt{m_{\text{CE}}}} \|\mathbf{y}_{\text{CE}} - \Phi_{\text{CE}} \mathbf{a}^\kappa\|_2 &\geq \frac{1}{\sqrt{m_{\text{CE}}}} (\|\Phi_{\text{CE}} (\mathbf{x}_s - \mathbf{a}^\kappa)\|_2 - \|\mathbf{e}_{\text{CE}}\|_2) \\ &\geq \sqrt{1 - \delta} \|\mathbf{x}_s - \mathbf{a}^\kappa\|_2 - \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}}. \end{aligned}$$

By combining the upper and lower bounds, we obtain

$$\begin{aligned} \|\mathbf{x}_s - \mathbf{a}\|_2 \leq \|\mathbf{x}_s - \mathbf{a}^\kappa\|_2 &\leq \frac{\sqrt{1 + \delta} \bar{\Gamma}}{\sqrt{1 - \delta}} + \frac{2\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}\sqrt{1 - \delta}} \\ &\leq 1.0254\bar{\Gamma} + \frac{2.0255}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2. \end{aligned}$$

Now assume that the first condition $\|\mathbf{v}_{\text{CE}}\|_2^2 > \|\mathbf{v}_{\text{CEold}}\|_2^2$ of the halting criteria in line 16 of Algorithm 1 holds. There are two possible cases : (i) $\|\mathbf{r}^{k-1}\|_2 \leq \bar{\Gamma}$ and (ii) $\|\mathbf{r}^{k-1}\|_2 > \bar{\Gamma}$. The case (i) implies that $\|\mathbf{x}_s - \mathbf{a}\|_2 \leq \bar{\Gamma}$. For the case (ii), note first that $\|\mathbf{x}_s - \mathbf{a}^k\|_2 \leq 0.5\|\mathbf{x}_s - \mathbf{a}^{k-1}\|_2 + \frac{2.124}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2$. Also, from the halting criterion,

$$\begin{aligned} \frac{1}{\sqrt{m_{\text{CE}}}} \left\| \Phi_{\text{CE}} (\mathbf{x}_s - \mathbf{a}^k) + \mathbf{e}_{\text{CE}} \right\|_2 &\geq \frac{1}{\sqrt{m_{\text{CE}}}} \left\| \Phi_{\text{CE}} (\mathbf{x}_s - \mathbf{a}^{k-1}) + \mathbf{e}_{\text{CE}} \right\|_2, \\ \left\| \frac{1}{\sqrt{m_{\text{CE}}}} \Phi_{\text{CE}} (\mathbf{x}_s - \mathbf{a}^k) \right\|_2 + \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}} &\geq \left\| \frac{1}{\sqrt{m_{\text{CE}}}} \Phi_{\text{CE}} (\mathbf{x}_s - \mathbf{a}^{k-1}) \right\|_2 - \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}}, \end{aligned}$$

$$\begin{aligned}\sqrt{1+\delta}\|\mathbf{x}_s - \mathbf{a}^k\|_2 + \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}} &\geq \sqrt{1-\delta}\|\mathbf{x}_s - \mathbf{a}^{k-1}\|_2 - \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}}, \\ \|\mathbf{x}_s - \mathbf{a}^k\|_2 &\geq \sqrt{\frac{1-\delta}{1+\delta}}\|\mathbf{x}_s - \mathbf{a}^{k-1}\|_2 - \frac{2}{\sqrt{1+\delta}}\frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}}.\end{aligned}$$

By combining the upper and lower bounds of $\|\mathbf{x}_s - \mathbf{a}^k\|_2$, we obtain

$$\begin{aligned}\|\mathbf{x}_s - \mathbf{a}\|_2 = \|\mathbf{x}_s - \mathbf{a}^{k-1}\|_2 &\leq \frac{\left(2.124 + \frac{2}{\sqrt{1+\delta}}\right)\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}\left(\sqrt{\frac{1-\delta}{1+\delta}} - 0.5\right)} \\ &\leq 8.625\frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}}.\end{aligned}\tag{3.5}$$

The support identification algorithm \mathcal{A} is assumed to have $\mathcal{O}(\mathcal{L}_{\mathcal{A}})$ runtime complexity in line 9. A conjugate gradient least square solver can approximate line 12 with $\mathcal{O}(s^2K^2\log^4|\mathcal{I}_{N,d}|)$ runtime complexity per iteration (see, e.g., Chapter 7 of [6], and Section 3 of [23]). Furthermore, a constant number of iterations (e.g. three in [33]) suffices. Lines 11, 14, and 15 require the generation of an $m_{\text{CE}} \times \mathcal{O}(s)$ or $m_{\text{SID}} \times \mathcal{O}(s)$ submatrix of either Φ_{CE} or Φ_{SID} , respectively. This will take $\mathcal{O}(s\mathcal{L}_{\Phi}m_{\text{CE}} + s\mathcal{L}_{\Phi}m_{\text{SID}})$ -time. Finally, the iteration number of the entire CoSaMP loop is bounded by κ , so that the overall runtime complexity is $\mathcal{O}\left((\mathcal{L}_{\mathcal{A}} + s^2K^2\mathcal{L}_{\Phi}\log^4|\mathcal{I}_{N,d}| + s\mathcal{L}_{\Phi}m_{\text{SID}}) \cdot \kappa\right)$. With respect to the sampling complexity, the support identification requires $m_{\text{SID}} = \mathcal{O}(\mathcal{L}'_{\mathcal{A}})$ function evaluations and the conjugate gradient method requires $m_{\text{CE}} = \mathcal{O}(sK^2\log^4|\mathcal{I}_{N,d}|)$ function evaluations [6, 23], and thus the overall sampling complexity is $m_{\text{SID}} + m_{\text{CE}} = \mathcal{O}(\mathcal{L}'_{\mathcal{A}} + sK^2\log^4|\mathcal{I}_{N,d}|)$. \square

Results concerning randomized constructions of RIP matrices $\frac{1}{\sqrt{m_{\text{CE}}}}\Phi_{\text{CE}} \in \mathbb{C}^{m_{\text{CE}} \times |\mathcal{I}_{N,d}|}$ for BOBPs with $\delta_{2s} \leq \delta$ and $m_{\text{CE}} = \mathcal{O}(sK^2\log^4|\mathcal{I}_{N,d}|)$ are well known (see, e.g., Theorem 2 and Chapter 12 of [17]). Our next result gives a qualitatively similar construction of a triple $(\Phi_{\text{SID}}, \mathcal{A}, \bar{\Gamma})$ with what is essentially the SIP for BOPBs (see Proposition 1 for an explicit SIP statement regarding this triple). More specifically, Theorem 4 constructs a support identification procedure with the properties required by Theorem 3, and also bounds its computational and sampling requirements. We remind the reader that the error vector $\mathbf{e}_{\text{SID}} \in \mathbb{C}^{m_{\text{SID}}}$ appearing in both Theorems 3 and 4 does not change from iteration to iteration in the analysis of Algorithm 1.

Theorem 4. (*Sublinear-Time Support Identification for BOPBs*). *There exists an algorithm $\mathcal{A} : \mathbb{C}^{m_{\text{SID}}} \rightarrow \mathcal{P}(\mathcal{I}_{N,d})$ that always outputs a set of at most $2s$ index vectors $\in \mathcal{I}_{N,d}$, and a sampling strategy for randomly selecting a set of m_{SID} grid points $\{\xi_{\ell}\}_{\ell \in [m_{\text{SID}}]} \subset \mathcal{D}$,*

such that the random sampling matrix $\Phi_{\text{SID}} \in \mathbb{C}^{m_{\text{SID}} \times |\mathcal{I}_{N,d}|}$ associated with $\{\xi_\ell\}_{\ell \in [m_{\text{SID}}]}$ as per (2.7) will have the following property with probability ≥ 0.99 :

$\mathcal{A}(\Phi_{\text{SID}} \mathbf{r}^k + \mathbf{e}_{\text{SID}}) =^6 \mathcal{A}(\Phi_{\text{SID}}(\tilde{\mathbf{c}} - \mathbf{a}^k) + \mathbf{e}'_{\text{SID}})$ will output a set $\tilde{\Omega} \subset \mathcal{I}_{N,d}$ such that

$$\left\| \mathbf{r}_{\tilde{\Omega}^c}^k \right\|_2 \leq 0.2086 \left\| \mathbf{r}^k \right\|_2 + 2.4172 \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 \leq 0.2203 \left\| \mathbf{r}^k \right\|_2 \quad (3.6)$$

holds for any $\mathbf{r}^k = \mathbf{x}_s - \mathbf{a}^k = \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} - \mathbf{a}^k$ satisfying $\left\| \mathbf{r}^k \right\|_2 > \bar{\Gamma}$, where

$$\bar{\Gamma} := \left(25\sqrt{23s} + 1 \right) \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + 18\sqrt{23} \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_1 + 22\gamma\sqrt{23s}. \quad (3.7)$$

In order to achieve this property with probability ≥ 0.99 it suffices that

$$m_{\text{SID}} = \mathcal{O}(\mathcal{L}'_{\mathcal{A}}) = \mathcal{O}\left(DK_{\infty}^{4\tilde{d}} s^3 d^4 \cdot \log^4\left(\frac{DN}{d}\right) \log^2(s) \log^2(D) \right)$$

if the BOS constants K_j are 1 for all but at most $\tilde{d} \in \mathbb{Z} \cap [0, D]$ BOS basis sets \mathcal{B}_j , and that

$$m_{\text{SID}} = \mathcal{O}(\mathcal{L}'_{\mathcal{A}}) = \mathcal{O}\left(DK_{\infty}^{4d} s^3 d^4 \cdot \log^4\left(\frac{DN}{d}\right) \log^2(s) \log^2(D) \right)$$

if $K_0 = 1$. In the first case the runtime complexity of \mathcal{A} will be

$$\mathcal{O}(\mathcal{L}_{\mathcal{A}}) = \mathcal{O}\left((s^5 + s^3 N) DK_{\infty}^{4\tilde{d}} d^4 \cdot \log^4\left(\frac{DN}{d}\right) \log^2(s) \log^2(D) \right),$$

and in the second case its runtime complexity will be

$$\mathcal{O}(\mathcal{L}_{\mathcal{A}}) = \mathcal{O}\left((s^5 + s^3 N) DK_{\infty}^{4d} d^4 \cdot \log^4\left(\frac{DN}{d}\right) \log^2(s) \log^2(D) \right)$$

when $K_0 = 1$.

Proof. See Section 4. This is a slight restatement of Theorem 7. The algorithm \mathcal{A} is Algorithm 2, and the m_{SID} points $\{\xi_\ell\}_{\ell \in [m_{\text{SID}}]} \subset \mathcal{D}$ should be randomly selected as per the first paragraph of Theorem 6. The runtime and sampling complexities then also follow from Theorem 6.

⁶Note that $\mathcal{A}(\Phi_{\text{SID}} \mathbf{r}^k + \mathbf{e}_{\text{SID}}) = \mathcal{A}(\Phi_{\text{SID}}(\mathbf{x}_s - \mathbf{a}^k) + \mathbf{e}_{\text{SID}}) = \mathcal{A}\left(\Phi_{\text{SID}}\left(\tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} - \mathbf{a}^k\right) + \mathbf{e}_{\text{SID}}\right) = \mathcal{A}\left(\Phi_{\text{SID}}(\tilde{\mathbf{c}} - \mathbf{a}^k) + \mathbf{e}'_{\text{SID}}\right)$ where $\mathbf{e}'_{\text{SID}} := \mathbf{e}_{\text{SID}} - \Phi_{\text{SID}}\left(\tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}}\right)$.

The first inequality in (3.6) follows directly from Theorem 7. In order to see that the second inequality

$$0.2086 \left\| \mathbf{r}^k \right\|_2 + 2.4172 \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 \leq 0.2203 \left\| \mathbf{r}^k \right\|_2$$

holds whenever $\left\| \mathbf{r}^k \right\|_2 > \bar{\Gamma}$, we note that

$$\begin{aligned} \left\| \mathbf{r}^k \right\|_2 > \bar{\Gamma} &= \left(25\sqrt{23s} + 1 \right) \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + 18\sqrt{23} \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_1 + 22\gamma\sqrt{23s} \\ &\geq \left(25\sqrt{23s} + 1 \right) \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + 18\sqrt{23} \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 \\ &\geq \left(43\sqrt{23} + 1 \right) \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2. \end{aligned}$$

Thus, one can see that

$$0.2086 \left\| \mathbf{r}^k \right\|_2 + 2.4172 \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 < \left(0.2086 + \frac{2.4172}{43\sqrt{23} + 1} \right) \left\| \mathbf{r}^k \right\|_2$$

which yields the desired effective SIP constant $\beta = 0.2203 \in (0, 0.2228]$. \square

The following proposition is a variant of Theorem 4 that more formally establishes exactly the type of SIP triple

$$\left(\Phi_{\text{SID}} \in \mathbb{C}^{m_{\text{SID}} \times |\mathcal{I}_{N,d}|}, \mathcal{A} : \mathbb{C}^{m_{\text{SID}}} \rightarrow \mathcal{P}([\mathcal{I}_{N,d}]), \Gamma : \mathbb{C}^{m_{\text{SID}}} \rightarrow [0, \infty) \right)$$

defined in Section 1. Its main contribution is to explicitly define a function $\Gamma : \mathbb{C}^{m_{\text{SID}}} \rightarrow [0, \infty)$ as per Definition 1 which produces a SIP triple when combined with the algorithm \mathcal{A} and matrix Φ_{SID} from Theorem 4. We hasten to point out, however, that only a valid upper bound $\bar{\Gamma}$ of $\Gamma(\mathbf{e}_{\text{SID}})$ as per (3.7) is actually necessary in order to apply Theorem 3 – one doesn't actually have to know the exact form of the best achievable function Γ . Nonetheless, for completeness we provide a function Γ in the next proposition which allows us to formally satisfy Definition 1 as stated in Section 1.

Proposition 1. *Let $\mathcal{A} : \mathbb{C}^{m_{\text{SID}}} \rightarrow \mathcal{P}(\mathcal{I}_{N,d})$ and $\Phi_{\text{SID}} \in \mathbb{C}^{m_{\text{SID}} \times |\mathcal{I}_{N,d}|}$ be the algorithm and random sampling matrix referred to by Theorem 4, where $\bigcup_{j \in [2D-1]} \left\{ \varrho_{S_j} \left(\mathbf{w}_\ell^j, \mathbf{z}_k^j \right) \right\}_{\ell \in [m_1], k \in [m_2]}$ denotes the $m_{\text{SID}} = m_1 m_2 (2D - 1)$ random sampling*

points⁷ used to create Φ_{SID} . In addition, define

$$\Gamma(\mathbf{e}_{\text{SID}}) := C \frac{\sqrt{s}}{\sqrt{m_1 m_2}} \max_{j \in [2D-1]} \|\mathbf{e}_{\text{SID}}^j\|_2,$$

where $C \in \mathbb{R}^+$ is an absolute constant ≤ 29 fixed below⁸, and where $\mathbf{e}_{\text{SID}}^j \in \mathbb{C}^{m_1 m_2}$ corresponds to the portion of $\mathbf{e}_{\text{SID}} \in \mathbb{C}^{m_{\text{SID}}}$ formed by evaluating e in (3.2) at the evaluation points $\{\varrho_{\mathcal{S}_j}(\mathbf{w}_\ell^j, \mathbf{z}_k^j)\}_{\ell \in [m_1], k \in [m_2]}$ for each $j \in [2D-1]$. Then, with probability ≥ 0.99 the triple $(\Phi_{\text{SID}}, \mathcal{A}, \Gamma)$ formed using the random evaluation points $\bigcup_{j \in [2D-1]} \{\varrho_{\mathcal{S}_j}(\mathbf{w}_\ell^j, \mathbf{z}_k^j)\}_{\ell \in [m_1], k \in [m_2]}$ will have both of the following properties:

(i) $(\Phi_{\text{SID}}, \mathcal{A}, \Gamma)$ will have the SIP of order $(2s, \beta = 0.2203)$, and

(ii) $\bar{\Gamma} := \bar{\Gamma}(\bar{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{f,s}^{\text{opt}}}, \gamma)$ in (3.7) will satisfy $\bar{\Gamma} \geq \Gamma(\mathbf{e}_{\text{SID}})$ for all inputs $\Phi_{\text{SID}} \mathbf{r}^k + \mathbf{e}_{\text{SID}} = \Phi_{\text{SID}}(\bar{\mathbf{c}} - \mathbf{a}^k) + \mathbf{e}'_{\text{SID}}$ with $\|\mathbf{e}'_{\text{SID}}\|_\infty \leq \gamma$.

Proof. The fact that $\bar{\Gamma}$ in (3.7) satisfies $\bar{\Gamma} \geq \Gamma(\mathbf{e}_{\text{SID}})$ is ultimately a consequence of Lemmas 7 and 14. The SIP holding for $(\Phi_{\text{SID}}, \mathcal{A}, \Gamma)$ follows from the fact that Theorem 4 still holds if the condition $\|\mathbf{r}^k\|_2 > \bar{\Gamma}$ is replaced by the condition $\|\mathbf{r}^k\|_2 > \Gamma(\mathbf{e}_{\text{SID}})$. This can be seen by tracing through Theorem 4's proof beginning with the proof of Theorem 8 where one need not apply Lemma 14, and from which an alternate version of Theorem 9 with (4.9) involving $\Gamma(\mathbf{e}_{\text{SID}})$ instead of its current right-hand side trivially follows. With such an alternate form of Theorem 9 in hand one can then immediately recover a similar variant of Theorem 6 involving $\Gamma(\mathbf{e}_{\text{SID}})$ which, in turn, can then provide an alternate (though less easily stated and interpretable) version of Theorem 4 involving the condition $\|\mathbf{r}^k\|_2 > \Gamma(\mathbf{e}_{\text{SID}})$. \square

Finally, in Theorem 5, it is shown that Algorithm 2 can be utilized as the support identification algorithm \mathcal{A} in a SIP triple $(\Phi_{\text{SID}}, \mathcal{A}, \Gamma : \mathbb{C}^{m_{\text{SID}}} \rightarrow [0, \infty))$ for use in Theorem 3. The sublinear runtime and sampling complexities of Algorithm 2 listed in Theorem 4 then result in a new sublinear-time and memory efficient compressive sensing approach

⁷See the input of Algorithm 2 for a description of the sampling points and note that the $2D-1$ blocks have been reindexed for ease of discussion, and that the index sets \mathcal{S}_j must therefore correspond to either $\{j\}$ or $[j+1]$ accordingly. For a description of how to generate the component points $\mathbf{w}_\ell^j, \mathbf{z}_k^j$ we refer the reader to Theorem 6.

⁸See (4.20) in Theorem 10 for a definition of Γ with explicit constants, where we further point out that α is fixed to be $\sqrt{23}$ in Theorem 6. When looking at Theorem 10 one should keep in mind that the matrix $\mathcal{E}_S^h \in \mathbb{C}^{m_1 \times m_2}$ therein is nothing other than a matricized version of $\mathbf{e}_{\text{SID}}^j$ with $S = \mathcal{S}_j$ for any desired choice of $j \in [2D-1]$.

for BOPB-compressible functions $f : \mathcal{D} \rightarrow \mathbb{C}$. We would like to remind the reader before stating this main result that $\mathbf{x}_s = \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}}$ is s -sparse with $\mathbf{y}_{\text{SID}} = \Phi_{\text{SID}}\mathbf{x}_s + \mathbf{e}_{\text{SID}}$ and $\mathbf{y}_{\text{CE}} = \Phi_{\text{CE}}\mathbf{x}_s + \mathbf{e}_{\text{CE}}$. Furthermore, the triple $(\Phi_{\text{SID}}, \mathcal{A}, \Gamma)$ constructed from the support identification procedure in Algorithm 2 with $\alpha := \sqrt{23}$ satisfies the SIP of order $(2s, \beta)$ with $\beta \in [0.2203, 0.2228]$ with high probability (see Theorem 4 and Proposition 1), and the matrix $\frac{1}{\sqrt{m_{\text{CE}}}}\Phi_{\text{CE}}$ has a RIP constant $\delta_{2s} \leq \delta$ for $\delta \in (0, 0.025]$ with high probability (see Theorem 2). Finally, \mathcal{A} always outputs a set of cardinality at most $2s$ as noted in Algorithm 1.

Theorem 5. (*Sublinear-Time Compressive Sensing for BOPB-compressible Signals*). *Let $N, d \in \mathbb{N} \setminus \{1\}$, $s < |\mathcal{I}_{N,d}|/2$, $\delta \in (0, 0.025]$, $\eta \in (0, \infty)$, $\kappa = \left\lceil \log_2 \left(\left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 / \eta \right) \right\rceil$, K the BOS constant of (1.6), $\mathbf{y}_{\text{SID}} = \Phi_{\text{SID}}\tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} + \mathbf{e}_{\text{SID}}$, and $\mathbf{y}_{\text{CE}} = \Phi_{\text{CE}}\tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} + \mathbf{e}_{\text{CE}}$ where we assume that both $\mathbf{e}'_{\text{SID}} := \mathbf{e}_{\text{SID}} - \Phi_{\text{SID}}\left(\tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}}\right)$ and $\mathbf{e}'_{\text{CE}} := \mathbf{e}_{\text{CE}} - \Phi_{\text{CE}}\left(\tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}}\right)$ have $\|\mathbf{e}'_{\text{SID}}\|_\infty \leq \gamma$ and $\|\mathbf{e}'_{\text{CE}}\|_\infty \leq \gamma$, respectively. Suppose further that the triple $(\Phi_{\text{SID}}, \mathcal{A}, \Gamma)$ with $\bar{\Gamma} \in \mathbb{R}^+$ such that $\bar{\Gamma} \geq \Gamma(\mathbf{e}_{\text{SID}})$ satisfies the SIP of order $(2s, \beta)$ with $\beta \in [0.2203, 0.2228]$ as per Theorem 4, and that $\frac{1}{\sqrt{m_{\text{CE}}}}\Phi_{\text{CE}} \in \mathbb{C}^{m_{\text{CE}} \times |\mathcal{I}_{N,d}|}$ has a RIP constant $\delta_{2s} \leq \delta$ and $m_{\text{CE}} = \mathcal{O}(sK^2 \log^4 |\mathcal{I}_{N,d}|) = \mathcal{O}(sK^2 d^4 \cdot \log^4(\frac{DN}{d}))$ (see (4.23)). Then, for each $k \geq 0$ the signal approximation \mathbf{a}^k in Algorithm 1 is s -sparse and satisfies*

$$\left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} - \mathbf{a}^{k+1} \right\|_2 \leq 0.5 \left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} - \mathbf{a}^k \right\|_2 + \frac{2.124}{\sqrt{m_{\text{CE}}}} \|\mathbf{e}_{\text{CE}}\|_2, \quad (3.8)$$

as long as

$$\left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} - \mathbf{a}^k \right\|_2 > \bar{\Gamma} := (25\sqrt{23}s + 1) \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + 18\sqrt{23} \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_1 + 22\gamma\sqrt{23s}.$$

As a consequence, Algorithm 1 produces an s -sparse approximation \mathbf{a} that satisfies

$$\left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} - \mathbf{a} \right\|_2 \leq \max \left\{ 1.03\bar{\Gamma} + 2.03 \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}}, 2^{-\kappa} \left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + 5 \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}}, 9 \frac{\|\mathbf{e}_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}} \right\} \quad (3.9)$$

$$\leq C \left(\sqrt{s} \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_1 + \gamma\sqrt{s} \right) + \eta, \quad (3.10)$$

where $C \in \mathbb{R}^+$ is an absolute universal constant.

In order to achieve (3.10) for all such possible inputs \mathbf{y}_{SID} and \mathbf{y}_{CE} with probability ≥ 0.99 it suffices that

$$m := m_{\text{SID}} + m_{\text{CE}} = \mathcal{O} \left(DK_\infty^{4\tilde{d}} s^3 d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right)$$

if the BOS constants K_j are 1 for all but at most $\tilde{d} \in \mathbb{Z} \cap [0, D]$ BOS basis sets \mathcal{B}_j (BOPB of type I), and that

$$m = \mathcal{O} \left(DK_\infty^{4d} s^3 d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right)$$

if $K_0 = 1$ (BOPB of type II).

In the BOPB of type I, the runtime complexity of the entire algorithm will be

$$\mathcal{O} \left((s^5 + s^3 N) D^2 K_\infty^{4\tilde{d}} d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \log \left(\left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 / \eta \right) \right),$$

and in the BOPB of type II, the runtime complexity will be

$$\mathcal{O} \left((s^5 + s^3 N) D^2 K_\infty^{4d} d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \log \left(\left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 / \eta \right) \right).$$

Here we have assumed that the runtime complexity of computing any desired matrix entry $(\Phi_{\text{CE}})_{j,\ell}$, or $(\Phi_{\text{SID}})_{j,\ell}$, for any valid choice of j, ℓ is $\mathcal{O}(\mathcal{L}_\Phi) = \mathcal{O}(ND)$ -time.

Proof. The result follows by combining Theorems 3 and 4 which immediately yields (3.9), as well as the stated runtime and sampling complexities. Note that Theorem 4 assumes that we are sampling from a BOPB-sparse function with arbitrary additive noise $e' : \mathcal{D} \rightarrow \mathbb{C}$ that has $\|e'\|_\infty \leq \gamma$, which leads to the restriction on $\|e'_{\text{SID}}\|_\infty$ and $\|e'_{\text{CE}}\|_\infty$. To obtain (3.10) one can simply substitute our choice of κ into (3.9) and use Lemma 4 to see that

$$\begin{aligned} \frac{\|e_{\text{CE}}\|_2}{\sqrt{m_{\text{CE}}}} &= \left\| \frac{e'_{\text{CE}}}{\sqrt{m_{\text{CE}}}} + \frac{1}{\sqrt{m_{\text{CE}}}} \Phi_{\text{CE}} \left(\tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right) \right\|_2 \\ &\leq \gamma + \sqrt{1 + \delta} \left(\frac{\left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_1}{\sqrt{s}} + \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 \right). \end{aligned}$$

Finally, we note that the runtime and sampling complexity bounds have been simplified by collecting and removing dominated terms along with the fact that $K \leq K_\infty^{\tilde{d}}$ (BOPB of type I) or $K \leq K_\infty^d$ (BOPB of type II) as discussed in Section 1.1. \square

With Theorem 5 in hand we may now prove our main result concerning function approximation in a Hilbert space $L^2(\mathcal{D}, \mu)$ spanned by a countable orthonormal product basis $\{T_{\mathbf{n}} \mid \mathbf{n} \in \mathbb{N}^D\} \supset \mathcal{B}$.

Corollary 1. (Main Result). Let $\eta \in (0, \infty)$ and $s, d, N \in \mathbb{N} \setminus \{1\}$ with $d \leq D$ and $s < |\mathcal{I}_{N,d}|/2$. There exists a finite set of grid points $\mathcal{G} \subset \mathcal{D}$, an algorithm $\mathcal{H} : \mathbb{C}^{|\mathcal{G}|} \rightarrow (\mathcal{I}_{N,d} \times \mathbb{C})^s$, and an absolute universal constant $C' \in \mathbb{R}^+$ such that the function $a : \mathcal{D} \rightarrow \mathbb{C}$ defined by $a(\boldsymbol{\xi}) := \sum_{(\mathbf{n}, a_{\mathbf{n}}) \in \mathcal{H}(f(\mathcal{G}))} a_{\mathbf{n}} T_{\mathbf{n}}(\boldsymbol{\xi})$ satisfies

$$\|f - a\|_{L^2(\mathcal{D}, \mu)} \leq \|f - \tilde{f}\|_{L^2(\mathcal{D}, \mu)} + C' \left(\sqrt{s} \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_2 + \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_1 + \gamma \sqrt{s} \right) + \eta$$

for all $f = \sum_{\mathbf{n} \in \mathbb{N}^D} c_{\mathbf{n}} T_{\mathbf{n}} \in L^2(\mathcal{D}, \mu)$ with $\gamma := \|f - \tilde{f}\|_{\infty} = \sup_{\boldsymbol{\xi} \in \mathcal{D}} |(f - \tilde{f})(\boldsymbol{\xi})| < \infty$, where $\tilde{f} : \mathcal{D} \rightarrow \mathbb{C}$ is the finite dimensional approximation to f defined as per (1.8).

If the BOS constants K_j are 1 for all but at most $\tilde{d} \in \mathbb{Z} \cap [0, D]$ BOS basis sets \mathcal{B}_j then

$$|\mathcal{G}| = \mathcal{O} \left(DK_{\infty}^{4\tilde{d}} s^3 d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right),$$

and the algorithm \mathcal{H} will have runtime complexity

$$\mathcal{O} \left((s^5 + s^3 N) D^2 K_{\infty}^{4\tilde{d}} d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \log \left(\left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_2 / \eta \right) \right).$$

If $K_0 = 1$ then

$$|\mathcal{G}| = \mathcal{O} \left(DK_{\infty}^{4d} s^3 d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right),$$

and the algorithm \mathcal{H} will have runtime complexity

$$\mathcal{O} \left((s^5 + s^3 N) D^2 K_{\infty}^{4d} d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \log \left(\left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_2 / \eta \right) \right).$$

Here we have assumed that any desired basis function $T_{\mathbf{n}} \in \mathcal{B}$ can be evaluated at any desired point in \mathcal{D} in $\mathcal{O}(ND)$ -time.

Proof. This follows from Theorem 5. The algorithm $\mathcal{H} : \mathbb{C}^{|\mathcal{G}|} \rightarrow (\mathcal{I}_{N,d} \times \mathbb{C})^s$ is Algorithm 1 using Algorithm 2 for line 9. The set of grid points $\mathcal{G} \subset \mathcal{D}$ is the union of the evaluation points used to create the random sampling matrices Φ_{SID} and Φ_{CE} from Theorem 5 so that $f(\mathcal{G}) = (\mathbf{y}_{\text{SID}}, \mathbf{y}_{\text{CE}}) \in \mathbb{C}^{m_{\text{SID}} + m_{\text{CE}}}$. And, the error bound follows from (3.10) and the triangle inequality since

$$\begin{aligned} \|f - a\|_{L^2(\mathcal{D}, \mu)} &\leq \|f - \tilde{f}\|_{L^2(\mathcal{D}, \mu)} + \|\tilde{f} - a\|_{L^2(\mathcal{D}, \mu)} \\ &= \|f - \tilde{f}\|_{L^2(\mathcal{D}, \mu)} + \|\tilde{\mathbf{c}} - \mathbf{a}\|_2 \end{aligned}$$

$$\begin{aligned}
&\leq \|f - \tilde{f}\|_{L^2(\mathcal{D}, \mu)} + \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_2 + \left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} - \mathbf{a} \right\|_2 \\
&\leq \|f - \tilde{f}\|_{L^2(\mathcal{D}, \mu)} + (C + 1) \left(\sqrt{s} \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_2 + \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_1 + \gamma \sqrt{s} \right) + \eta
\end{aligned}$$

where the absolute constant C is from Theorem 5. \square

Next, in Section 4, we will focus on developing Algorithm 2 and demonstrating that it performs as desired. We hasten to note before beginning, however, that the development of another support identification method satisfying the SIP with lower runtime or sampling complexity could be used to create a new and potentially superior version of Theorem 5 in the future. We leave the development of such improved methods in the hands of the sufficiently interested and clever reader.

4 Sublinear-Time Support Identification

We assume herein that the function $h : \mathcal{D} \rightarrow \mathbb{C}$ of D variables,

$$h := \tilde{h} + e', \quad (4.1)$$

where $\tilde{h} : \mathcal{D} \rightarrow \mathbb{C}$ is as per (1.8) with coefficient vector $\tilde{\mathbf{r}} \in \mathbb{C}^{\mathcal{I}_{N,d}}$ in the BOS \mathcal{B} as per (2.1),

$$\tilde{h}(\boldsymbol{\xi}) := \sum_{n \in \mathcal{I}_{N,d}} \tilde{r}_n T_n(\boldsymbol{\xi}), \quad (4.2)$$

and where $e' : \mathcal{D} \rightarrow \mathbb{C}$ is bounded so that $\sup_{\boldsymbol{\xi} \in \mathcal{D}} |e'(\boldsymbol{\xi})| \leq \gamma$. In terms of our problem setting about f , the function \tilde{h} is each residual function $\tilde{f} - a$ where a is the function constructed from the approximation \mathbf{a}^k that Algorithm 1 produces in each iteration. In order to escape exponential sampling dependence on the dimension D we will further assume below the BOPB of type I or II (see Section 1.1). In addition, motivated by Section 3, we will be most interested in the case where $\left\| h - \tilde{h}_{2s}^{\text{opt}} \right\|_{L^2(\mathcal{D}, \mu)} \lesssim \left\| \tilde{h}_{2s}^{\text{opt}} \right\|_{L^2(\mathcal{D}, \mu)}$.

In particular, we will almost exclusively represent h as $h = \tilde{h}_{2s}^{\text{opt}} + \left(\tilde{h} - \tilde{h}_{2s}^{\text{opt}} + e' \right)$ below where we hope that $e_h := \tilde{h} - \tilde{h}_{2s}^{\text{opt}} + e'$ has a relatively small L^2 -norm compared to that of $\tilde{h}_{2s}^{\text{opt}}$.

In order to approximate h we seek to find a near-optimal set of basis functions from \mathcal{B} on which to approximately project h . In particular, we would be quite pleased to identify all of $\Omega_{\tilde{h}, 2s}^{\text{opt}}$ – that is, all the basis functions which compose $\tilde{h}_{2s}^{\text{opt}}$ – if possible given that $h \approx \tilde{h}_{2s}^{\text{opt}}$. This appears a bit too ambitious goal in general, however. Instead, we will

focus on the easier goal of identifying all the entries of $\Omega_{\tilde{h}, 2s}^{\text{opt}}$ which individually contribute a nontrivial amount of energy to the total L^2 -norm of $\tilde{h}_{2s}^{\text{opt}}$. We will represent (portions of) these basis element indices via the following sets of (partial) energetic indices.

Let $\mathcal{S} \subseteq [D]$, $s' \in \mathbb{N}$, $\alpha \in (1, \infty)$ be a fixed constant to be determined later. We define the set of *energetic partial index vectors of $\tilde{h}_{s'}^{\text{opt}}$ in $N^{\mathcal{S}}$* to be

$$\Omega_{\mathcal{S}}^{\alpha, s'} := \left\{ \mathbf{n}_{\mathcal{S}} \mid \mathbf{n} \in \mathcal{I}_{N, d} \ \& \ \left\| \left(\tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}} \right)_{\mathcal{S}; \mathbf{n}} \right\|_2 \geq \frac{\left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}} \right\|_2}{\alpha \sqrt{s'}} \right\} \subseteq N^{\mathcal{S}}, \quad (4.3)$$

where $N^{\mathcal{S}} := \left\{ \mathbf{n}_{\mathcal{S}} \mid \mathbf{n} \in \mathcal{I}_{N, d} \right\} \subseteq \mathcal{I}_{N, d} \subseteq [N]^D$. Note in particular that $N^{[D]} = \mathcal{I}_{N, d}$ so

that $\Omega_{[D]}^{\alpha, s'}$ contains all $\mathbf{n} \in \Omega_{\tilde{h}, s'}^{\text{opt}}$ whose associated entry has $|\tilde{r}_{\mathbf{n}}| \geq \frac{\left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}} \right\|_2}{\alpha \sqrt{s'}}$. Furthermore, it is also important to note that $\Omega_{[D]}^{\alpha, s'} \subseteq \Omega_{\tilde{h}, s'}^{\text{opt}}$ holds for all $s' \in [|\mathcal{I}_{N, d}|] \setminus \{0\}$. More generally, $\Omega_{\mathcal{S}}^{\alpha, s'} \subseteq \Omega_{s', \mathcal{S}}^{\text{opt}} := \left\{ \mathbf{q}_{\mathcal{S}} \mid \mathbf{q} \in \Omega_{\tilde{h}, s'}^{\text{opt}} \right\}$ holds for all $\mathcal{S} \subseteq [D]$ and $s' \in [|\mathcal{I}_{N, d}|] \setminus \{0\}$.

Our next lemma shows that identifying a superset of $\Omega_{[D]}^{\alpha, 2s}$ is enough to ensure that we will find a set of basis elements that can approximate $\tilde{h}_{2s}^{\text{opt}}$ (and therefore h) well. In particular, we will find the support of the majority of the energy of $\tilde{h}_{2s}^{\text{opt}}$, $\left\| \tilde{h}_{2s}^{\text{opt}} \right\|_{L^2(\mathcal{D}, \mu)} = \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_2$. With respect to Section 3, the next lemma shows that any support set we discover which contains $\Omega_{[D]}^{\alpha, 2s}$ will be sufficiently informative to guarantee that CoSaMP will make progress during its current iteration.

Lemma 5. *Let $\alpha \geq \sqrt{23}$. If $\Omega_{[D]}^{\alpha, 2s} \subseteq \tilde{\Omega} \subseteq \mathcal{I}_{N, d}$ then*

$$\left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}} \cap \tilde{\Omega}^c} \right\|_2 \leq 0.2086 \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_2.$$

Proof. Setting $\mathbf{r}' := \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}}$, one can see that

$$\left\| \mathbf{r}'_{\tilde{\Omega}^c} \right\|_2^2 = \sum_{\mathbf{n} \in \Omega_{\tilde{h}, 2s}^{\text{opt}} \cap \tilde{\Omega}^c} |r'_{\mathbf{n}}|^2 < 2s \cdot \frac{\left\| \mathbf{r}' \right\|_2^2}{\alpha^2 2s} \leq \frac{\left\| \mathbf{r}' \right\|_2^2}{23}$$

since $\tilde{\Omega}^c \cap \Omega_{[D]}^{\alpha, 2s} = \emptyset$. □

In light of Lemma 5 above we will now turn our attention to identifying $\Omega_{[D]}^{\alpha, 2s}$ in a computationally and sample efficient fashion. In particular, we seek to identify $\Omega_{[D]}^{\alpha, 2s}$ as quickly as possible while simultaneously using as few fixed and nonadaptive samples from $h = \tilde{h}_{2s}^{\text{opt}} + e_h$ as possible. This is accomplished via Algorithm 2 below. Theorem 6 then proves that it works as intended.

Theorem 6. *Let $\{\mathbf{w}_\ell^j\}_{\ell \in [m_1]} \subset \mathcal{D}_j$ be m_1 points drawn independently at random according to μ_j , and $\{\mathbf{z}_k^j\}_{k \in [m_2]} \subset \mathcal{D}_{[D] \setminus \{j\}}$ be m_2 points drawn independently at random according to $\mu_{[D] \setminus \{j\}}$, for all $j \in [D]$. Furthermore, let $\{\mathbf{w}_\ell^{D-1+j}\}_{\ell \in [m_1]} \subset \mathcal{D}_{[j+1]}$ be m_1 points drawn independently at random according to $\mu_{[j+1]}$, and $\{\mathbf{z}_k^{D-1+j}\}_{k \in [m_2]} \subset \mathcal{D}_{[D] \setminus [j+1]}$ be m_2 points drawn independently at random according to $\mu_{[D] \setminus [j+1]}$, for all $j \in [D] \setminus \{0\}$. If m_1 and m_2 are chosen to be sufficiently large for all $j \in [2D-1]$ then the following property will hold with probability ≥ 0.99 :*

Algorithm 2 will output a set $\tilde{\Omega} \supset \Omega_{[D]}^{\alpha, 2s}$ for all $h = \tilde{h}_{2s}^{\text{opt}} + e_h$ as per (4.1) with coefficient vector $\tilde{\mathbf{r}} \in \mathbb{C}^{\mathcal{I}^{N,d}}$ in the BOS \mathcal{B} satisfying

$$\left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_2 > 25\sqrt{23}s \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_2 + 18\sqrt{23} \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_1 + 22\gamma\sqrt{23s}. \quad (4.4)$$

In order to achieve this property with probability ≥ 0.99 it suffices for Algorithm 2 to utilize a total number of function evaluations from h that is of size

$$m_{\text{SID}} = m_1 m_2 (2D-1) = \mathcal{O} \left(DK_\infty^{4\tilde{d}} s^3 d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right)$$

if the BOS constants K_j are 1 for all but at most $\tilde{d} \in \mathbb{Z} \cap [0, D]$ BOS basis sets \mathcal{B}_j (BOPB of type I), and that is of size

$$m'_{\text{SID}} = m_1 m_2 (2D-1) = \mathcal{O} \left(DK_\infty^{4\tilde{d}} s^3 d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right)$$

if $K_0 = 1$ (BOPB of type II).

In the BOPB of type I, the runtime complexity of Algorithm 2 will be

$$\mathcal{O} \left((s^5 + s^3 N) DK_\infty^{4\tilde{d}} d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right),$$

Algorithm 2 Implemented Support Identification (Special Case of Algorithm 3)

1: **procedure** SupportID

2: **Parameters:** $N \in \mathbb{N}$, $D \in \mathbb{N}$, $\alpha \geq \sqrt{23}$, sparsity $s \in \mathbb{N}$.

3: **Input:** $\mathbf{v}_{\text{SID}} \in \mathbb{C}^{m_1 m_2 (2D-1)}$ split into $2D-1$ blocks. The first D blocks $\mathbf{v}_{\text{SID},j} := \{h(\varrho_{\{j\}}(\mathbf{w}_\ell^j, \mathbf{z}_k^j))\}_{\ell \in [m_1], k \in [m_2]}$, $j \in [D]$, belong to entry identification where $\mathbf{w}_\ell^j \in \mathcal{D}_{\{j\}}$, $\mathbf{z}_k^j \in \mathcal{D}_{[D] \setminus \{j\}}$, and $\varrho_{\{j\}}$ as per (2.4). The last $D-1$ blocks $\mathbf{v}_{\text{SID},D-1+j} := \{h(\varrho_{[j+1]}(\mathbf{w}_\ell^{D-1+j}, \mathbf{z}_k^{D-1+j}))\}_{\ell \in [m_1], k \in [m_2]}$, $j \in [D] \setminus \{0\}$, belong to the pairing where $\mathbf{w}_\ell^{D-1+j} \in \mathcal{D}_{[j+1]}$ and $\mathbf{z}_k^{D-1+j} \in \mathcal{D}_{[D] \setminus [j+1]}$.

4: **Output:** A set $\tilde{\Omega} \supset \Omega_{[D]}^{\alpha, 2s}$ with $|\tilde{\Omega}| \leq 2s$.

5: **for** $j = 0$ **up to** $D-1$ **do**

6: $\mathbf{E}_{j,n}^{\text{EI}} \leftarrow \frac{1}{m_2} \sum_{k \in [m_2]} \left| \frac{1}{m_1} \sum_{\ell \in [m_1]} (\mathbf{v}_{\text{SID},j})_{\ell,k} \overline{T_{j;n}(\mathbf{w}_\ell^j)} \right|^2$ for each $n \in [N]$, see

also (4.16), with $(\mathbf{v}_{\text{SID},j})_{\ell,k} = h(\varrho_{\{j\}}(\mathbf{w}_\ell^j, \mathbf{z}_k^j))$.

7: $\mathcal{N}_j \leftarrow \{n \in [N] \mid \min(2s, N)\text{-largest values } \mathbf{E}_{j,n}^{\text{EI}}\}$.

8: **end for**

9: $\mathcal{T}_D \leftarrow \mathcal{N}_0$.

10: **for** $j = 1$ **up to** $D-1$ **do**

11: $\mathcal{T}'_{D+j} \leftarrow \{\mathbf{n} + \mathbf{m} \mid \mathbf{n} \in \mathcal{T}_{D+j-1}, \mathbf{m} \in \mathcal{N}_j\} \cap \mathcal{I}_{N,d} \subseteq N^{[j+1]}$.

12: $\mathbf{E}_{j,n}^{\text{P}} \leftarrow \frac{1}{m_2} \sum_{k \in [m_2]} \left| \frac{1}{m_1} \sum_{\ell \in [m_1]} (\mathbf{v}_{\text{SID},D-1+j})_{\ell,k} \overline{T_{[j+1];\mathbf{n}}(\mathbf{w}_\ell^{D-1+j})} \right|^2$ for each $\mathbf{n} \in \mathcal{T}'_{D+j}$, see also (4.16).

13: $\mathcal{T}_{D+j} \leftarrow \{\mathbf{n} \in \mathcal{T}'_{D+j} \mid \min(2s, |\mathcal{T}'_{D+j}|)\text{-largest values } \mathbf{E}_{j,\mathbf{n}}^{\text{P}}\}$.

14: **end for**

15: Return $\tilde{\Omega} \leftarrow \mathcal{T}_{2D-1}$ (Note that it will always be true that $|\tilde{\Omega}| \leq 2s$.)

16: **end procedure**

and in the BOPB of type II, the runtime complexity will be

$$\mathcal{O} \left((s^5 + s^3 N) DK_\infty^{4d} d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right).$$

Proof. See Section 4.1. The desired result follows from a simplified version of Theorem 9 with $\alpha = \sqrt{23}$. \square

The index j in Theorem 6 belongs to three different sets, $[D]$, $[D] \setminus \{0\}$ and $[2D - 1]$. To explain, the set $[2D - 1]$ comprehends all j 's belonging to the first two sets, $[D]$ and $[D] \setminus \{0\}$.

Theorem 6 combined with Lemma 5 is enough to guarantee that Algorithm 2 can identify a support set $\tilde{\Omega}$ that contains the majority of the energy of the $2s$ -sparse vector $\tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}}$. However, Theorem 3 in Section 3 requires that $\|\mathbf{r}_{\tilde{\Omega}^c}\|_2$ should be relatively small, where $\mathbf{r} \in \mathbb{C}^{\mathcal{I}_{N,d}}$ is the $2s$ -sparse vector $\mathbf{r} := \mathbf{x}_s - \mathbf{a}^k = \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} - \mathbf{a}^k$ ($\mathbf{r} = \mathbf{r}^k$ in Section 3).⁹ As a result we must now relate this \mathbf{r} to the coefficients $\tilde{\mathbf{r}} := \tilde{\mathbf{c}} - \mathbf{a}^k = \mathbf{x} - \mathbf{a}^k$ of the function $\tilde{h} := \tilde{f} - \sum_{\mathbf{n} \in \mathcal{I}_{N,d}} \tilde{a}_n^k T_n$ whose noisy samples we are passing into Algorithm 2 in line 9 of Algorithm 1. The following lemma can be used to relate $\|\mathbf{r}\|_2$ to $\|\tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}}\|_2$.

Lemma 6. *Let $s \in \lceil |\mathcal{I}_{N,d}| / 2 \rceil$, $\tilde{\mathbf{c}}, \mathbf{a}^k \in \mathbb{C}^{\mathcal{I}_{N,d}}$ where $\|\mathbf{a}^k\|_0 \leq s$, and recall that $\tilde{\mathbf{r}} := \tilde{\mathbf{c}} - \mathbf{a}^k$, $\tilde{h}(\boldsymbol{\xi}) := \sum_{\mathbf{n} \in \mathcal{I}_{N,d}} \tilde{r}_n T_n(\boldsymbol{\xi})$, and $\tilde{f}(\boldsymbol{\xi}) := \sum_{\mathbf{n} \in \mathcal{I}_{N,d}} \tilde{c}_n T_n(\boldsymbol{\xi})$. One can see that*

$$\|\mathbf{r}\|_2 = \left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} - \mathbf{a}^k \right\|_2 \leq \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_2 + \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_2.$$

Proof. Let $\mathcal{Q} := \left(\text{supp}(\mathbf{a}^k) \cup \Omega_{\tilde{f}, s}^{\text{opt}} \right) \cap \text{supp}(\tilde{\mathbf{r}})$, and note that $|\Omega_{\tilde{h}, 2s}^{\text{opt}}| = \min \{2s, |\text{supp}(\tilde{\mathbf{r}})|\} \geq |\mathcal{Q}|$. As a result one can see that

$$\begin{aligned} \left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} - \mathbf{a}^k \right\|_2 &= \left\| \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} - \mathbf{a}^k + \tilde{\mathbf{c}}_{\text{supp}(\mathbf{a}^k) \setminus \Omega_{\tilde{f}, s}^{\text{opt}}} - \tilde{\mathbf{c}}_{\text{supp}(\mathbf{a}^k) \setminus \Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_2 \\ &= \left\| \left(\tilde{\mathbf{c}} - \mathbf{a}^k \right)_{\Omega_{\tilde{f}, s}^{\text{opt}} \cup \text{supp}(\mathbf{a}^k)} - \tilde{\mathbf{c}}_{\text{supp}(\mathbf{a}^k) \setminus \Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_2 \\ &\leq \left\| \left(\tilde{\mathbf{c}} - \mathbf{a}^k \right)_{\Omega_{\tilde{f}, s}^{\text{opt}} \cup \text{supp}(\mathbf{a}^k)} \right\|_2 + \left\| \tilde{\mathbf{c}}_{\text{supp}(\mathbf{a}^k) \setminus \Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_2 \\ &= \|\tilde{\mathbf{r}}_{\mathcal{Q}}\|_2 + \left\| \tilde{\mathbf{c}}_{\text{supp}(\mathbf{a}^k) \setminus \Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_2 \\ &\leq \|\tilde{\mathbf{r}}_{\mathcal{Q}}\|_2 + \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_2 \\ &\leq \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_2 + \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f}, s}^{\text{opt}}} \right\|_2. \end{aligned}$$

⁹Recall that $\tilde{\mathbf{c}} \in \mathbb{C}^{\mathcal{I}_{N,d}}$ is the coefficient vector of \tilde{f} as per (3.1), and that $\mathbf{a}^k \in \mathbb{C}^{\mathcal{I}_{N,d}}$ is CoSaMP's s -sparse approximation to $\mathbf{x} = \tilde{\mathbf{c}} \in \mathbb{C}^{\mathcal{I}_{N,d}}$ in its k^{th} -iteration.

as we wished to show. \square

The next lemma upper bounds the best $2s$ -term approximation error of $\tilde{\mathbf{r}}$ by the best s -term approximation error of $\tilde{\mathbf{c}}$. It will allow us to relate the condition (4.4) under which Algorithm 2 succeeds to $\tilde{\mathbf{c}}$.

Lemma 7. *Let $s \in [|\mathcal{I}_{N,d}|/2]$, $\tilde{\mathbf{c}}, \mathbf{a}^k \in \mathbb{C}^{\mathcal{I}_{N,d}}$ where $\|\mathbf{a}^k\|_0 \leq s$, and recall that $\tilde{\mathbf{r}} := \tilde{\mathbf{c}} - \mathbf{a}^k$, $\tilde{h}(\boldsymbol{\xi}) := \sum_{\mathbf{n} \in \mathcal{I}_{N,d}} \tilde{r}_{\mathbf{n}} T_{\mathbf{n}}(\boldsymbol{\xi})$, and $\tilde{f}(\boldsymbol{\xi}) := \sum_{\mathbf{n} \in \mathcal{I}_{N,d}} \tilde{c}_{\mathbf{n}} T_{\mathbf{n}}(\boldsymbol{\xi})$. One can see that $\left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}} \right\|_p \leq \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_p$ holds for all $p \geq 1$. As a consequence, it will always be the case that*

$$\begin{aligned} 25\sqrt{23}s \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}} \right\|_2 + 18\sqrt{23} \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}} \right\|_1 + 22\gamma\sqrt{23}s \\ \leq 25\sqrt{23}s \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + 18\sqrt{23} \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_1 + 22\gamma\sqrt{23}s =: \Gamma. \end{aligned} \quad (4.5)$$

Proof. A quick calculation reveals that

$$\begin{aligned} \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}} \right\|_p^p &\leq \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{f},s}^{\text{opt}} \cup \text{supp}(\mathbf{a}^k)} \right\|_p^p = \sum_{\mathbf{n} \in \mathcal{I}_{N,d} \setminus (\Omega_{\tilde{f},s}^{\text{opt}} \cup \text{supp}(\mathbf{a}^k))} |\tilde{c}_{\mathbf{n}} - a_{\mathbf{n}}^k|^p \\ &= \sum_{\mathbf{n} \in \mathcal{I}_{N,d} \setminus (\Omega_{\tilde{f},s}^{\text{opt}} \cup \text{supp}(\mathbf{a}^k))} |\tilde{c}_{\mathbf{n}}|^p \\ &\leq \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_p^p, \end{aligned}$$

as we wished to show. \square

We are now able to assert that our support identification algorithm will work for all $2s$ -sparse vectors $\tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}}$ whose norms are sufficiently large with respect to the best s -term approximation error Γ defined above in (4.5).

Lemma 8. *Let the \mathbf{w}_ℓ^j and \mathbf{z}_k^j for $j \in [2D-1]$ in Algorithm 2 be chosen independently at random as per Theorem 6 above. Then, the following property will hold with probability ≥ 0.99 :*

Algorithm 2 will output a set $\tilde{\Omega} \subset \mathcal{I}_{N,d}$ with $|\tilde{\Omega}| \leq 2s$ that will also have

$$\left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}} \cap \tilde{\Omega}^c} \right\|_2 \leq 0.2086 \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}} \right\|_2 \quad (4.6)$$

for all $h = \tilde{h}_{2s}^{\text{opt}} + e_h$ as per (4.1) with coefficient vector $\tilde{\mathbf{r}} \in \mathbb{C}^{\mathcal{I}_{N,d}}$ in the BOS \mathcal{B} satisfying $\left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}} \right\|_2 > \Gamma$, where Γ is defined in (4.5).

The runtime and sampling complexities of Algorithm 2 will remain as in Theorem 6 above.

Proof. By Lemma 7, $\left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}} \right\|_2 > \Gamma$ implies that (4.4) holds. Thus, the result follows from Theorem 6 combined with Lemma 5. \square

The following theorem is the main theorem of this section. It proves that the support set $\tilde{\Omega}$ found by Algorithm 2 will also contain the majority of the energy of the $2s$ -sparse vector $\mathbf{r} := \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} - \mathbf{a}^k \in \mathbb{C}^{\mathcal{I}_{N,d}}$, as needed in Section 3.

Theorem 7. (*Support Identification*). Let $s \in \lceil |\mathcal{I}_{N,d}|/2 \rceil$, $\tilde{\mathbf{c}}, \mathbf{a}^k \in \mathbb{C}^{\mathcal{I}_{N,d}}$ where $\|\mathbf{a}^k\|_0 \leq s$, and recall that $\tilde{\mathbf{r}} := \tilde{\mathbf{c}} - \mathbf{a}^k$, $\tilde{h}(\boldsymbol{\xi}) := \sum_{\mathbf{n} \in \mathcal{I}_{N,d}} \tilde{r}_{\mathbf{n}} T_{\mathbf{n}}(\boldsymbol{\xi})$, and $\tilde{f}(\boldsymbol{\xi}) := \sum_{\mathbf{n} \in \mathcal{I}_{N,d}} \tilde{c}_{\mathbf{n}} T_{\mathbf{n}}(\boldsymbol{\xi})$. Suppose that the \mathbf{w}_{ℓ}^j and \mathbf{z}_k^j in Algorithm 2 are chosen independently at random as per Theorem 6 above. Then the following property will hold with probability ≥ 0.99 :

Algorithm 2 will output a set $\tilde{\Omega} \subset \mathcal{I}_{N,d}$ with

$$\left\| \mathbf{r}_{\tilde{\Omega}^c} \right\|_2 \leq 0.2086 \|\mathbf{r}\|_2 + 2.4172 \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2$$

for any $\mathbf{r} = \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} - \mathbf{a}^k$ satisfying $\|\mathbf{r}\|_2 > \bar{\Gamma}$, where

$$\bar{\Gamma} := \left(25\sqrt{23s} + 1 \right) \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + 18\sqrt{23} \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_1 + 22\gamma\sqrt{23s}. \quad (4.7)$$

The runtime and sampling complexities of Algorithm 2 will remain as in Theorem 6 above.

Proof. Let $\mathbf{r}' := \tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}}$. Note that $\tilde{\mathbf{r}} - \mathbf{r} = \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}}$, and so

$$\|\mathbf{r}' - \mathbf{r}\|_2 = \|\mathbf{r}' - \tilde{\mathbf{r}} + \tilde{\mathbf{r}} - \mathbf{r}\|_2$$

$$\begin{aligned}
&\leq \|\mathbf{r}' - \tilde{\mathbf{r}}\|_2 + \|\tilde{\mathbf{r}} - \mathbf{r}\|_2 \\
&= \|\tilde{\mathbf{r}} - \mathbf{r}'\|_2 + \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 \\
&\leq 2 \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2, \tag{4.8}
\end{aligned}$$

where the last inequality holds by Lemma 7. Thus, we have that the following holds whenever (4.6) does:

$$\begin{aligned}
\|\mathbf{r}_{\tilde{\Omega}^c}\|_2 &= \|\mathbf{r}_{\tilde{\Omega}^c} - \mathbf{r}'_{\tilde{\Omega}^c} + \mathbf{r}'_{\tilde{\Omega}^c}\|_2 \\
&\leq \|(\mathbf{r} - \mathbf{r}')_{\tilde{\Omega}^c}\|_2 + \|\mathbf{r}'_{\tilde{\Omega}^c}\|_2 \\
&\leq \|(\mathbf{r} - \mathbf{r}')_{\tilde{\Omega}^c}\|_2 + 0.2086 \|\mathbf{r}'\|_2 \\
&\leq 2 \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + 0.2086 \|\mathbf{r}' - \mathbf{r} + \mathbf{r}\|_2 \\
&\leq 2 \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + 0.2086 (\|\mathbf{r}' - \mathbf{r}\|_2 + \|\mathbf{r}\|_2) \\
&\leq 2 \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + 0.2086 \left(2 \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + \|\mathbf{r}\|_2 \right) \\
&= 2.4172 \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2 + 0.2086 \|\mathbf{r}\|_2,
\end{aligned}$$

where the second inequality holds if (4.6) does, and the third and fifth inequalities hold by (4.8).

To finish we note that (4.6) will indeed hold by Lemma 8 as long as $\left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}} \right\|_2 > \Gamma = \bar{\Gamma} - \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2$. And, $\left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}} \right\|_2 > \bar{\Gamma} - \left\| \tilde{\mathbf{c}} - \tilde{\mathbf{c}}_{\Omega_{\tilde{f},s}^{\text{opt}}} \right\|_2$ will hold whenever $\|\mathbf{r}\|_2 > \bar{\Gamma}$ holds by Lemma 6. \square

We will now focus on proving Theorem 6.

4.1 Proof of Theorem 6: Identifying $\Omega_{[D]}^{\alpha,2s}$ for $\tilde{h}_{2s}^{\text{opt}}$ Using Samples from $h = \tilde{h}_{2s}^{\text{opt}} + e_h$

Our strategy for finding $\Omega_{[D]}^{\alpha,2s}$ will involve building it up from a sequence of energetic partial index vectors of $\tilde{h}_{2s}^{\text{opt}}$ that correspond to, e.g., the disjoint subsets of indices

$$\mathcal{S}_j^{\text{EI}} = \{j\} \text{ for all } j \in [D].^{10}$$

Note that the energetic partial index vectors in this case will contain the entries of the index vectors which have large associated values in $\tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}}$. That is,

$$\Omega_{\mathcal{S}_j^{\text{EI}}}^{\alpha,2s} = \Omega_{\{j\}}^{\alpha,2s} \supseteq \left\{ ne_j \mid \exists \mathbf{n} \in \Omega_{\tilde{h},2s}^{\text{opt}} \text{ with } |\tilde{r}_{\mathbf{n}}| \geq \frac{\left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},2s}^{\text{opt}}} \right\|_2}{\alpha\sqrt{2s}} \text{ whose } j^{\text{th}} \text{ entry is } n \in [N] \right\},$$

where $e_j \in \mathcal{I}_{N,d}$ is the j^{th} standard basis vector. As a result, the set $\Omega_{\mathcal{S}_j^{\text{EI}}}^{\alpha,2s}$ effectively contains all the j^{th} -entries of the largest-magnitude coefficient vector indices in $\Omega_{\tilde{h},2s}^{\text{opt}}$. Furthermore, it is trivial to find a reasonably small superset of $\Omega_{\mathcal{S}_j^{\text{EI}}}^{\alpha,2s}$ when, e.g., N is not too large – one can simply use the set $N^{\{j\}} = \{ne_j \mid n \in [N]\} \supset \Omega_{\mathcal{S}_j^{\text{EI}}}^{\alpha,2s}$.

Of course, the sets $\Omega_{\mathcal{S}_0^{\text{EI}}}^{\alpha,2s}, \dots, \Omega_{\mathcal{S}_{D-1}^{\text{EI}}}^{\alpha,2s}$ are of limited utility in their own right when it comes to finding $\Omega_{[D]}^{\alpha,2s}$. Our strategy will therefore be to use these sets to build up a sequence of new energetic partial index sets $\Omega_{\mathcal{S}_1^{\text{P}}}^{\alpha,2s}, \dots, \Omega_{\mathcal{S}_{D-1}^{\text{P}}}^{\alpha,2s}$ each of which corresponds to an increasingly large subset of indices $\mathcal{S}_j^{\text{P}} \subseteq [D]$. In particular, if we define $\mathcal{S}_j^{\text{P}} := \cup_{\ell=0}^j \mathcal{S}_\ell^{\text{EI}} = \cup_{\ell=0}^j \{\ell\} = [j+1]$ for all $j \in [D] \setminus \{0\}$ we will eventually obtain a superset of $\Omega_{\mathcal{S}_{D-1}^{\text{P}}}^{\alpha,2s} = \Omega_{[D]}^{\alpha,2s}$ (as desired) in a process that is analogous to the ‘‘Pairing’’ method utilized in [12]. The following lemma is the basis for building up $\Omega_{[D]}^{\alpha,2s}$ by combining energetic partial index vectors of $\tilde{h}_{2s}^{\text{opt}}$ that correspond to smaller index sets $\mathcal{S}_1, \mathcal{S}_2 \subset [D]$ in this fashion. Recall that $\mathcal{P}(N^{\mathcal{S}})$ denotes the power set of $N^{\mathcal{S}}$ for any given $\mathcal{S} \subseteq [D]$.

Lemma 9. *Let $s' \in \mathbb{N}$, $\alpha \in (1, \infty)$, and $\mathcal{S}_1, \mathcal{S}_2 \subset [D]$ be disjoint. If $\mathcal{T}_1 \in \mathcal{P}(N^{\mathcal{S}_1})$ and $\mathcal{T}_2 \in \mathcal{P}(N^{\mathcal{S}_2})$ are such that $\Omega_{\mathcal{S}_1}^{\alpha,s'} \subseteq \mathcal{T}_1$ and $\Omega_{\mathcal{S}_2}^{\alpha,s'} \subseteq \mathcal{T}_2$, then*

$$\Omega_{\mathcal{S}_1 \cup \mathcal{S}_2}^{\alpha,s'} \subseteq \mathcal{T}_{1,2} := \{\mathbf{n} + \mathbf{m} \mid \mathbf{n} \in \mathcal{T}_1, \mathbf{m} \in \mathcal{T}_2\} \cap \mathcal{I}_{N,d} \subseteq N^{\mathcal{S}_1 \cup \mathcal{S}_2}.$$

¹⁰Here the ‘‘EI’’ in the superscript of $\mathcal{S}_j^{\text{EI}}$ stands for ‘‘Entry Identification’’ in the terminology of [12]. In fact many other valid choices for these sets also exist – please see Algorithm 3 for the general criteria they must satisfy.

Proof. Let $\mathbf{r}' := \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}}$, and note that for all $\mathbf{n} \in \mathcal{I}_{N,d}$ it is the case that

$$\{\mathbf{m} \in \mathcal{I}_{N,d} \mid \mathbf{m}_{\mathcal{S}_1 \cup \mathcal{S}_2} = \mathbf{n}_{\mathcal{S}_1 \cup \mathcal{S}_2}\} \subseteq \{\mathbf{m} \in \mathcal{I}_{N,d} \mid \mathbf{m}_{\mathcal{S}_1} = \mathbf{n}_{\mathcal{S}_1}\} \cap \{\mathbf{m} \in \mathcal{I}_{N,d} \mid \mathbf{m}_{\mathcal{S}_2} = \mathbf{n}_{\mathcal{S}_2}\}$$

holds. As a consequence, for any $\mathbf{n} \in \mathcal{I}_{N,d}$ it will be the case that both

$$\|\mathbf{r}'_{\mathcal{S}_1; \mathbf{n}}\|_2^2 = \sum_{\substack{\mathbf{m} \in \mathcal{I}_{N,d} \text{ s.t.} \\ \mathbf{m}_{\mathcal{S}_1} = \mathbf{n}_{\mathcal{S}_1}}} |r'_m|^2 \geq \sum_{\substack{\mathbf{m} \in \mathcal{I}_{N,d} \text{ s.t.} \\ \mathbf{m}_{\mathcal{S}_1 \cup \mathcal{S}_2} = \mathbf{n}_{\mathcal{S}_1 \cup \mathcal{S}_2}}} |r'_m|^2 = \|\mathbf{r}'_{\mathcal{S}_1 \cup \mathcal{S}_2; \mathbf{n}}\|_2^2$$

and

$$\|\mathbf{r}'_{\mathcal{S}_2; \mathbf{n}}\|_2^2 = \sum_{\substack{\mathbf{m} \in \mathcal{I}_{N,d} \text{ s.t.} \\ \mathbf{m}_{\mathcal{S}_2} = \mathbf{n}_{\mathcal{S}_2}}} |r'_m|^2 \geq \sum_{\substack{\mathbf{m} \in \mathcal{I}_{N,d} \text{ s.t.} \\ \mathbf{m}_{\mathcal{S}_1 \cup \mathcal{S}_2} = \mathbf{n}_{\mathcal{S}_1 \cup \mathcal{S}_2}}} |r'_m|^2 = \|\mathbf{r}'_{\mathcal{S}_1 \cup \mathcal{S}_2; \mathbf{n}}\|_2^2$$

hold. These inequalities in turn imply that $\Omega'_1 := \{\mathbf{n}_{\mathcal{S}_1} \mid \mathbf{n} \in \Omega_{\mathcal{S}_1 \cup \mathcal{S}_2}^{\alpha, s'}\} \subseteq \Omega_{\mathcal{S}_1}^{\alpha, s'} \subseteq \mathcal{T}_1$ and $\Omega'_2 := \{\mathbf{n}_{\mathcal{S}_2} \mid \mathbf{n} \in \Omega_{\mathcal{S}_1 \cup \mathcal{S}_2}^{\alpha, s'}\} \subseteq \Omega_{\mathcal{S}_2}^{\alpha, s'} \subseteq \mathcal{T}_2$. Finally, the fact that \mathcal{S}_1 and \mathcal{S}_2 are disjoint now implies that

$$\Omega_{\mathcal{S}_1 \cup \mathcal{S}_2}^{\alpha, s'} \subseteq \{\mathbf{n} + \mathbf{m} \mid \mathbf{n} \in \Omega'_1, \mathbf{m} \in \Omega'_2\} \subseteq \{\mathbf{n} + \mathbf{m} \mid \mathbf{n} \in \mathcal{T}_1, \mathbf{m} \in \mathcal{T}_2\}$$

is true as desired. \square

Let $\mathcal{S}_0^{\text{P}} := \mathcal{S}_0^{\text{EI}}$. Note that applying Lemma 9 repeatedly with, e.g., $\mathcal{S}_1 = \mathcal{S}_j^{\text{P}}, \mathcal{S}_2 = \mathcal{S}_{j+1}^{\text{EI}}, \mathcal{T}_1 = \mathcal{T}_{1,2}$ from the $(j-1)^{\text{st}}$ application of Lemma 9,¹¹ and $\mathcal{T}_2 = N^{\mathcal{S}_{j+1}^{\text{EI}}}$ for $j = 0, 1, \dots, D-2$ will yield a superset of $\Omega_{[D]}^{\alpha, 2s}$ on its $(D-1)^{\text{st}}$ application. However, the cardinality of the resulting superset $\mathcal{T}_{1,2}$ of $\Omega_{\mathcal{S}_1 \cup \mathcal{S}_2}^{\alpha, 2s}$ will also balloon to $|\mathcal{T}_1| \cdot |\mathcal{T}_2|$ at each step, eventually becoming exponentially large in D on the $(D-1)^{\text{st}}$ application of Lemma 9 in the worst case. In order to prevent this worst case exponential growth in the size of the resulting sets $\mathcal{T}_{1,2}$ we will interleave the applications of Lemma 9 with the use of an *energetic-index sieve function* $\mathcal{F}_{\mathcal{S}_1 \cup \mathcal{S}_2}^{2s} : \mathcal{P}(N^{\mathcal{S}_1 \cup \mathcal{S}_2}) \rightarrow \mathcal{P}(N^{\mathcal{S}_1 \cup \mathcal{S}_2})$ as in (4.19) which reduces the cardinality of any $\mathcal{T}_{1,2} \supseteq \Omega_{\mathcal{S}_1 \cup \mathcal{S}_2}^{\alpha, 2s}$ to $2s$ without losing any of $\Omega_{\mathcal{S}_1 \cup \mathcal{S}_2}^{\alpha, 2s}$. These sieve functions will allow Lemma 9 to be applied repeatedly as above while maintaining output sets of small cardinality at all stages, which we can see how they work in lines 12 and 13 of Algorithm 2.

The next theorem proves the existence of a set-valued function $\mathcal{F}_{\mathcal{S}}^{2s} : \mathcal{P}(N^{\mathcal{S}}) \rightarrow \mathcal{P}(N^{\mathcal{S}})$ for any given $\mathcal{S} \subseteq [D]$ which, when given any subset $\mathcal{T} \subset N^{\mathcal{S}}$ containing $\Omega_{\mathcal{S}}^{\alpha, 2s}$ as per (4.3)

¹¹with, e.g., $\mathcal{T}_1 = N^{\mathcal{S}_0^{\text{EI}}}$ when $j = 0$

as input, will output a smaller subset $\mathcal{T}' \subset \mathcal{T}$ of cardinality at most $2s$ which still contains $\Omega_S^{\alpha, 2s}$. Note that these set valued functions necessarily depend on the function $\tilde{h}_{2s}^{\text{opt}}$ in question via the definition of $\Omega_S^{\alpha, 2s}$. However, it is crucial to note that all the \mathcal{F}_S^{2s} considered herein only utilize a few point samples from $h = \tilde{h}_{2s}^{\text{opt}} + e_h$ (i.e., noisy point samples from $\tilde{h}_{2s}^{\text{opt}}$) on a fixed and nonadaptive grid. More specifically, the grid on which each \mathcal{F}_S^{2s} samples h depends only on $2s, \mathcal{S}$, and the BOPB \mathcal{B} with respect to which h is presumed to be approximately sparse, and not at all on the particular function h in question.

Theorem 8 (Existence of Low-Complexity Energetic-Index Sieve Functions). *Choose $t' \in [2D]$ and any desired $\mathcal{S}_0, \dots, \mathcal{S}_{t'} \subseteq [D]$. For all $j \in [t']$ there exists an associated energetic-index sieve function $\mathcal{F}_{\mathcal{S}_j}^{2s} : \mathcal{P}(N^{\mathcal{S}_j}) \rightarrow \mathcal{P}(N^{\mathcal{S}_j})$ for which both*

1. $\Omega_{\mathcal{S}_j}^{\alpha, 2s} \cap \mathcal{T} \subseteq \mathcal{F}_{\mathcal{S}_j}^{2s}(\mathcal{T})$ holds for all $\mathcal{T} \in \mathcal{P}(N^{\mathcal{S}_j})$, and
2. $|\mathcal{F}_{\mathcal{S}_j}^{2s}(\mathcal{T})| \leq 2s$ holds for all $\mathcal{T} \in \mathcal{P}(N^{\mathcal{S}_j})$,

are true for all $h : \mathcal{D} \rightarrow \mathbb{C}$ as above (4.1) that satisfy

$$\left\| \tilde{h}_{2s}^{\text{opt}} \right\|_{L^2(\mathcal{D}, \mu)} = \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_2 > 25\alpha\sqrt{s} \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_2 + 18\alpha \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_1 + 22\alpha\gamma\sqrt{s}.^{12}$$

Furthermore, each $\mathcal{F}_{\mathcal{S}_j}^{2s} : \mathcal{P}(N^{\mathcal{S}_j}) \rightarrow \mathcal{P}(N^{\mathcal{S}_j})$ is computed using evaluations of any given $h = \tilde{h}_{2s}^{\text{opt}} + e_h$ at $m_1^j m_2^j$ fixed and nonadaptive grid points $\left\{ \left(\mathbf{w}_\ell^j, \mathbf{z}_k^j \right) \right\}_{\ell \in [m_1^j], k \in [m_2^j]} \subset \mathcal{D}$, $m_1^j, m_2^j \in \mathbb{N}$, where $\mathbf{w}_\ell^j \in \mathcal{D}_{\mathcal{S}_j}$ and $\mathbf{z}_k^j \in \mathcal{D}_{\mathcal{S}_j^c}$ for all $j \in [t']$, $\ell \in [m_1^j]$, and $k \in [m_2^j]$.

If the BOS constants K_j are 1 for all but at most $\tilde{d} \in \mathbb{Z} \cap [0, D]$ BOS basis sets \mathcal{B}_j , then each such $\mathcal{F}_{\mathcal{S}_j}^{2s} : \mathcal{P}(N^{\mathcal{S}_j}) \rightarrow \mathcal{P}(N^{\mathcal{S}_j})$ above requires only

$$m_j = m_1^j m_2^j = \mathcal{O} \left(K_\infty^{4\tilde{d}} s^3 d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right)$$

evaluations of any given $h = \tilde{h}_{2s}^{\text{opt}} + e_h$ at m_j fixed and nonadaptive grid points $\subset \mathcal{D}$.¹³ As a result, $\mathcal{F}_{\mathcal{S}_j}^{2s}(\mathcal{T})$ can be computed in just $\mathcal{O}(m_j |\mathcal{T}|)$ -time for any given $\mathcal{T} \in \mathcal{P}(N^{\mathcal{S}_j})$ and h in this case.¹⁴ If, on the other hand, $K_0 = 1$ then each such $\mathcal{F}_{\mathcal{S}_j}^{2s} : \mathcal{P}(N^{\mathcal{S}_j}) \rightarrow \mathcal{P}(N^{\mathcal{S}_j})$

¹²The constants here have been rounded up to the nearest integer from those implied by Theorem 10 and Lemma 14 after substituting $s' = 2s$.

¹³It is important to emphasize here that the grid on which we must evaluate each function f is a fixed grid which does not change depending on h .

¹⁴Herein we assume that h has been evaluated in advance on our non-adaptive grid so that its values

requires only

$$m'_j = m_1^j m_2^j = \mathcal{O} \left(K_\infty^{4d} s^3 d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right)$$

evaluations of any given h at m'_j fixed and nonadaptive grid points in \mathcal{D} . As a result, $\mathcal{F}_{\mathcal{S}_j}^{2s}(\mathcal{T})$ can be computed in just $\mathcal{O}(m'_j |\mathcal{T}|)$ -time for any given $\mathcal{T} \in \mathcal{P}(N^{\mathcal{S}_j})$ and h in this case.

Proof. See Section 4.2 below. The proof follows by applying Theorem 10 with $e_h = \tilde{h} - \tilde{h}_{2s}^{\text{opt}} + (h - \tilde{h}) = h - \tilde{h}_{2s}^{\text{opt}}$ for each set $\mathcal{S}_0, \dots, \mathcal{S}_{t'}$. After recalling that $\sup_{\xi \in \mathcal{D}} \left| (h - \tilde{h})(\xi) \right| \leq \gamma$ we can see that Lemma 14 will also apply in each case. Finally, the runtime and sampling complexity bounds follow from Lemma 15 and Remark 3. \square

Remark 1. It is important to note that Section 4.2 proves more than mere existence of the collection of low-complexity energetic-index sieve functions promised in Theorem 8. In fact it proves their existence by proving that one can generate such a collection with high probability \geq , e.g., 0.99 by letting $\{\mathbf{w}_\ell^j\}_{\ell \in [m_1^j]} \subset \mathcal{D}_{\mathcal{S}_j}$ be m_1^j sampling points drawn independently at random according to $\mu_{\mathcal{S}_j}$, and by letting $\{\mathbf{z}_k^j\}_{k \in [m_2^j]} \subset \mathcal{D}_{\mathcal{S}_j^c}$ be m_2^j sampling points drawn independently at random according to $\mu_{\mathcal{S}_j^c}$, for all $j \in [t']$. This is done by showing that randomly selecting the nonadaptive grid points in this fashion ultimately guarantees that their related random sampling matrices in (4.10) and (4.13) have well behaved restricted isometry constants. See Remark 3 for additional details and related discussion.

With Lemma 9 and Theorem 8 in hand one can now see that Algorithm 3 will be guaranteed to return a superset $\tilde{\Omega}$ of $\Omega_{[D]}^{\alpha, 2s}$ whose cardinality is at most $2s$.

Theorem 9. Let $\mathcal{S}_0^{\text{EI}}, \dots, \mathcal{S}_t^{\text{EI}} \subset [D]$ form a partition of $[D]$ for $t \in [D] \setminus \{0\}$ and set $\mathcal{S}_j^{\text{P}} := \cup_{\ell=0}^j \mathcal{S}_\ell^{\text{EI}} \subseteq [D]$ for all $j \in [t+1] \setminus \{0\}$ as per Algorithm 3. Let $\mathcal{F}_{\mathcal{S}_j^{\text{EI}}}^{2s} : \mathcal{P}(N^{\mathcal{S}_j^{\text{EI}}}) \rightarrow \mathcal{P}(N^{\mathcal{S}_j^{\text{EI}}})$ and $\mathcal{F}_{\mathcal{S}_j^{\text{P}}}^{2s} : \mathcal{P}(N^{\mathcal{S}_j^{\text{P}}}) \rightarrow \mathcal{P}(N^{\mathcal{S}_j^{\text{P}}})$ be their associated energetic-index sieve functions. When executed using these energetic-index sieve functions Algorithm 3 will output a set $\tilde{\Omega}$ with $|\tilde{\Omega}| \leq 2s$ that will also have $\Omega_{[D]}^{\alpha, 2s} \subset \tilde{\Omega}$ provided that $h = \tilde{h}_{2s}^{\text{opt}} + e_h$ has

$$\left\| \tilde{h}_{2s}^{\text{opt}} \right\|_{L^2(\mathcal{D}, \mu)} = \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_2 > 25\alpha\sqrt{s} \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_2 + 18\alpha \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, 2s}^{\text{opt}}} \right\|_1 + 22\alpha\gamma\sqrt{s}. \quad (4.9)$$

at each grid point can be retrieved in $\mathcal{O}(1)$ -time. In addition, note that setting $d = D$ above still leads to sampling and runtime complexities for each sieve function that scale only polynomially in D . This is due to \tilde{d} being independent of d .

Algorithm 3 Support Identification

- 1: **procedure Generalized SupportID**
 - 2: **Parameters:** $s \in \mathbb{N}$, $t \in [D] \setminus \{0\}$, A Partition of $[D]$ into $\mathcal{S}_0^{\text{EI}}, \dots, \mathcal{S}_t^{\text{EI}} \subset [D]$, and the Associated Pairing Index Sets $\mathcal{S}_j^{\text{P}} := \cup_{\ell=0}^j \mathcal{S}_\ell^{\text{EI}} \subseteq [D]$ for all $j \in [t+1] \setminus \{0\}$.
 - 3: **Input:** $\mathbf{v}_{\text{SID}} \in \mathbb{C}^{\sum_{j=0}^{2t} m_1^j m_2^j}$ split into $2t + 1$ blocks $\left\{ h \left(\varrho_{\mathcal{S}_j^{\text{EI}}} \left(\mathbf{w}_\ell^j, \mathbf{z}_k^j \right) \right) \right\}_{\ell \in [m_1^j], k \in [m_2^j]}^{j \in [t+1]} \cup \left\{ h \left(\varrho_{\mathcal{S}_{j-t}^{\text{P}}} \left(\mathbf{w}_\ell^j, \mathbf{z}_k^j \right) \right) \right\}_{\ell \in [m_1^j], k \in [m_2^j]}^{j \in [2t+1] \setminus [t+1]}$ indexed by j w/ $\mathbf{w}_\ell^j \in \mathcal{D}_{\mathcal{S}_j^{\text{EI}}} \forall j \in [t+1]$ & $\ell \in [m_1^j]$, $\mathbf{w}_\ell^j \in \mathcal{D}_{\mathcal{S}_{j-t}^{\text{P}}} \forall j \in [2t+1] \setminus [t+1]$ & $\ell \in [m_1^j]$, $\mathbf{z}_k^j \in \mathcal{D}_{(\mathcal{S}_j^{\text{EI}})^c} \forall j \in [t+1]$ & $k \in [m_2^j]$, & $\mathbf{z}_k^j \in \mathcal{D}_{(\mathcal{S}_{j-t}^{\text{P}})^c} \forall j \in [2t+1] \setminus [t+1]$ & $k \in [m_2^j]$.
 - 4: **Output:** A set $\tilde{\Omega} \supset \Omega_{[D]}^{\alpha, 2s}$
 - 5: Compute $\mathcal{N}_j \leftarrow \mathcal{F}_{\mathcal{S}_j^{\text{EI}}}^{2s} \left(N^{\mathcal{S}_j^{\text{EI}}} \right)$ using $\left\{ h \left(\varrho_{\mathcal{S}_j^{\text{EI}}} \left(\mathbf{w}_\ell^j, \mathbf{z}_k^j \right) \right) \right\}_{\ell \in [m_1^j], k \in [m_2^j]}$ for each $j \in [t+1]$
 - 6: $\mathcal{T}_t \leftarrow \mathcal{N}_0$
 - 7: **for** $j = t + 1$ **up to** $2t$ **do**
 - 8: $\mathcal{T}'_j \leftarrow \left\{ \mathbf{n} + \mathbf{m} \mid \mathbf{n} \in \mathcal{T}_{j-1}, \mathbf{m} \in \mathcal{N}_{j-t} \right\} \cap \mathcal{I}_{N,d} \subseteq N^{\mathcal{S}_{j-t}^{\text{P}}}$
 - 9: $\mathcal{T}_j \leftarrow \mathcal{F}_{\mathcal{S}_{j-t}^{\text{P}}}^{2s} \left(\mathcal{T}'_j \right)$ using $\left\{ h \left(\varrho_{\mathcal{S}_{j-t}^{\text{P}}} \left(\mathbf{w}_\ell^j, \mathbf{z}_k^j \right) \right) \right\}_{\ell \in [m_1^j], k \in [m_2^j]}$
 - 10: **end for**
 - 11: Return $\tilde{\Omega} \leftarrow \mathcal{T}_{2t}$
 - 12: **end procedure**
-

The total number of function evaluations required¹⁵ by Algorithm 3 is

$$m_{\text{SID}} = \mathcal{O} \left(t K_\infty^{4\tilde{d}} s^3 d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right)$$

if the BOS constants K_j are 1 for all but at most $\tilde{d} \in \mathbb{Z} \cap [0, D]$ BOS basis sets \mathcal{B}_j , and is

$$m'_{\text{SID}} = \mathcal{O} \left(t K_\infty^{4d} s^3 d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right)$$

if $K_0 = 1$.

The runtime complexity of Algorithm 3 will be

$$\mathcal{O} \left(\left(s^5 + s^3 \max_{j \in [t+1]} |N^{\mathcal{S}_j^{\text{EI}}}| \right) t K_\infty^{4\tilde{d}} d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right)$$

¹⁵In the bounds below t may be upper bounded by D .

if the BOS constants K_j are 1 for all but at most $\tilde{d} \in \mathbb{Z} \cap [0, D]$ BOS basis sets \mathcal{B}_j , and

$$\mathcal{O} \left(\left(s^5 + s^3 \max_{j \in [t+1]} |N^{\mathcal{S}_j^{\text{EI}}}| \right) t K_\infty^{4d} d^4 \cdot \log^4 \left(\frac{DN}{d} \right) \log^2(s) \log^2(D) \right)$$

if $K_0 = 1$.

Proof. The proof follows directly from Lemma 9 and Theorem 8. \square

Remark 2. Note that Algorithm 2 is a special case of Algorithm 3 with $t = D - 1$, $\mathcal{S}_j^{\text{EI}} := \{j\}$ for all $j \in [D]$, $\mathcal{S}_j^{\text{P}} := \cup_{\ell=0}^j \mathcal{S}_\ell^{\text{EI}} = \cup_{\ell=0}^j \{\ell\} \subseteq [D]$ for all $j \in [D] \setminus \{0\}$, and where the sieve functions $\mathcal{F}_{\mathcal{S}_j^{\text{EI}}}^{2s}$, $\mathcal{F}_{\mathcal{S}_{j-t}^{\text{P}}}^{2s}$ have been written down explicitly using (4.16), (4.18), and (4.19). Therein the $\mathcal{F}_{\mathcal{S}_j^{\text{EI}}}^{2s} \left(N^{\mathcal{S}_j^{\text{EI}}} \right)$ are computed for all $j \in [t+1]$ by lines 5 – 8 of Algorithm 2, and each $\mathcal{F}_{\mathcal{S}_{j-t}^{\text{P}}}^{2s} \left(\mathcal{T}'_j \right)$ in Algorithm 3 is computed by lines 12 – 13 of Algorithm 2.

Though dedicated to proving Theorem 8, this next subsection will be initially focussed on learning $\Omega_{\mathcal{S}}^{\alpha, s'}$ for arbitrary BOPB-sparse functions with $h = \tilde{h} = \tilde{h}_{\tilde{h}, s'}^{\text{opt}}$ for which $\mathbf{r}' = \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}}$. It will then be generalized to cover more general functions h of the type discussed above (4.1) toward its end as an extension of the noisy sparse case. A proof of Theorem 8 may then be obtained by setting $s' = 2s$.

4.2 Proof of Theorem 8: Generalized Entry Identification & Pairing

In the vast majority of this subsection we will be considering an arbitrary function $\tilde{h}: \mathcal{D} \rightarrow \mathbb{C}$ of D variables as per (4.2) whose coefficient vector $\tilde{\mathbf{r}} \in \mathbb{C}^{\mathcal{I}_{N,d}}$ is only nonzero for entries indexed by index vectors $\mathbf{q} \in \mathcal{I}_{N,d}$. In particular, we will be focussing almost exclusively on the development of efficient strategies for learning about the support of the coefficient vector $\tilde{\mathbf{r}}$ of such \tilde{h} in the special case where $\tilde{\mathbf{r}}$ is s' -sparse so that $\tilde{h} = \tilde{h}_{\tilde{h}, s'}^{\text{opt}}$ and $\tilde{\mathbf{r}} = \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}}$. Our first lemma does this by telling us how to estimate the ℓ^2 -norm of

any $\mathbf{r}'' = \left(\tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}} \right)_{\mathcal{S}; \mathbf{n}} = \tilde{\mathbf{r}}_{\{\mathbf{q} \in \Omega_{\tilde{h}, s'}^{\text{opt}} \mid \mathbf{q}_{\mathcal{S}} = \mathbf{n}_{\mathcal{S}}\}} \in \mathbb{C}^{\mathcal{I}_{N,d}}$ in that case (i.e., how to estimate

of the energy of all the coefficients of $\tilde{h}_{\tilde{h}, s'}^{\text{opt}}$ whose index vectors $\mathbf{q} \in \mathcal{I}_{N,d}$ match another fixed index vector $\mathbf{n} \in \mathcal{I}_{N,d}$ in all index positions $\mathcal{S} \subset [D]$) by using just a few inner products with “simpler” functions of only $|\mathcal{S}| < D$ variables. The idea is that these inner products will be easy to approximate numerically for $|\mathcal{S}|$ small. As a result, one can hope to learn about the index vectors of the nonzero entries of any such \mathbf{r}'' by approximately computing just a few inner products involving functions of just a few variables in order

to, e.g., discover values of \mathbf{n} for which $\|\mathbf{r}''\|_2$ is large.

Lemma 10. *Let $\delta \in (0, 3/4]$, $\mathcal{S} \subset [D]$, and $\{\mathbf{z}_k\}_{k \in [m_2]} \subset \mathcal{D}_{\mathcal{S}^c}$ be m_2 sampling points drawn independently at random according to $\mu_{\mathcal{S}^c}$ in order to form a zero-padded random sampling matrix $\Phi_{\mathcal{S}^c; \mathbf{0}} \in \mathbb{C}^{m_2 \times |\mathcal{I}_{N,d}|}$ for the BOS $\mathcal{B}_{\mathcal{S}^c}$ as in (2.3) with entries*

$$(\Phi_{\mathcal{S}^c; \mathbf{0}})_{k, \mathbf{q}} := \begin{cases} T_{\mathcal{S}^c; \mathbf{q}}(\mathbf{z}_k) & \text{if } \mathbf{q}_{\mathcal{S}} = \mathbf{0} \text{ \& } \mathbf{q} \in \mathcal{I}_{N,d} \\ 0 & \text{otherwise} \end{cases} \quad (4.10)$$

indexed by $k \in [m_2]$ and $\mathbf{q} \in \mathcal{I}_{N,d}$. Suppose the nonzero columns of $\frac{1}{\sqrt{m_2}} \Phi_{\mathcal{S}^c; \mathbf{0}}$ have the restricted isometry property (RIP) of order (s, δ) . Then, for all $\mathbf{n} \in \mathcal{I}_{N,d} \subseteq [N]^D$, vectors of additive evaluation errors $\mathbf{e}^h \in \mathbb{C}^{m_2}$, and functions \tilde{h} as per (4.2) one will have

$$\left| \sqrt{\sum_{k \in [m_2]} \frac{1}{m_2} \left| \left\langle (\tilde{h}_{s'}^{\text{opt}})_{\mathcal{S}^c; \mathbf{z}_k}, T_{\mathcal{S}; \mathbf{n}} \right\rangle_{(\mathcal{D}_{\mathcal{S}}, \mu_{\mathcal{S}})} + e_k^h \right|^2} - \|\mathbf{r}''\|_2 \right| \leq \frac{2}{3} \delta \|\mathbf{r}''\|_2 + \frac{\|\mathbf{e}^h\|_2}{\sqrt{m_2}} \quad (4.11)$$

where $\mathbf{r}'' := \tilde{\mathbf{r}}_{\{\mathbf{q} \in \Omega_{\tilde{h}, s'}^{\text{opt}} \mid \mathbf{q}_{\mathcal{S}} = \mathbf{n}_{\mathcal{S}}\}}$.

Proof. Consider the zero-padded random sampling matrix $\Phi_{\mathcal{S}^c; \mathbf{n}} \in \mathbb{C}^{m_2 \times |\mathcal{I}_{N,d}|}$ for the BOS $\mathcal{B}_{\mathcal{S}^c}$ as in (2.3) with entries

$$(\Phi_{\mathcal{S}^c; \mathbf{n}})_{k, \mathbf{q}} := \begin{cases} T_{\mathcal{S}^c; \mathbf{q}}(\mathbf{z}_k) & \text{if } \mathbf{q}_{\mathcal{S}} = \mathbf{n}_{\mathcal{S}} \text{ \& } \mathbf{q} \in \mathcal{I}_{N,d} \\ 0 & \text{otherwise} \end{cases} \quad (4.12)$$

indexed by $k \in [m_2]$ and $\mathbf{q} \in \mathcal{I}_{N,d}$. Note that $\mathbf{q}_{\mathcal{S}} = \mathbf{n}_{\mathcal{S}} \text{ \& } \mathbf{q} \in \mathcal{I}_{N,d} \implies (\mathbf{q}, \mathbf{0})_{\mathcal{S}^c} \in \mathcal{I}_{N,d}$ for all $d \in [D+1] \setminus \{0\}$, $\mathcal{S}^c \subset [D]$, and $\mathbf{n}, \mathbf{q} \in \mathcal{I}_{N,d}$. As a result, the nonzero columns of $\Phi_{\mathcal{S}^c; \mathbf{0}}$ will contain the nonzero columns of $\Phi_{\mathcal{S}^c; \mathbf{n}}$ as a subset.¹⁶ This further implies that the matrix consisting of the nonzero columns of $\frac{1}{\sqrt{m_2}} \Phi_{\mathcal{S}^c; \mathbf{n}}$ will also have the restricted isometry property (RIP) of order (s, δ) .

Applying Lemma 3 together with the definition of $\tilde{h}_{s'}^{\text{opt}}$ we now have that

$$\sum_{k \in [m_2]} \frac{1}{m_2} \left| \left\langle (\tilde{h}_{s'}^{\text{opt}})_{\mathcal{S}^c; \mathbf{z}_k}, T_{\mathcal{S}; \mathbf{n}} \right\rangle_{(\mathcal{D}_{\mathcal{S}}, \mu_{\mathcal{S}})} + e_k^h \right|^2 = \sum_{k \in [m_2]} \left| \frac{1}{\sqrt{m_2}} \left\langle \left(\tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}} \right)_{\mathcal{S}; \mathbf{n}}, \overline{\Phi_{\mathcal{S}^c; \mathbf{n}; \mathbf{z}_k}} \right\rangle + \frac{e_k^h}{\sqrt{m_2}} \right|^2.$$

¹⁶Note that the nonzero columns of $\Phi_{\mathcal{S}^c; \mathbf{n}}$ will be indexed by different \mathbf{q} in $\Phi_{\mathcal{S}^c; \mathbf{0}}$. However, this reindexing will ultimately just represent a permutation of the nonzero columns of $\Phi_{\mathcal{S}^c; \mathbf{n}}$ as a submatrix of $\Phi_{\mathcal{S}^c; \mathbf{0}}$. And, permuting the columns of a matrix does not change its restricted isometry constants.

Noting now that each vector $\Phi_{\mathcal{S}^c; \mathbf{n}; \mathbf{z}_k}$ as per (2.6) can be replaced by an equivalent row of $\Phi_{\mathcal{S}^c; \mathbf{n}}$ in (4.12) we can further see that

$$\begin{aligned} \sqrt{\sum_{k \in [m_2]} \frac{1}{m_2} \left| \left\langle (\tilde{h}_{s'}^{\text{opt}})_{\mathcal{S}^c; \mathbf{z}_k}, T_{\mathcal{S}; \mathbf{n}} \right\rangle_{(\mathcal{D}_{\mathcal{S}}, \mu_{\mathcal{S}})} + e_k^h \right|^2} &= \sqrt{\sum_{k \in [m_2]} \left| \frac{1}{\sqrt{m_2}} \left(\Phi_{\mathcal{S}^c; \mathbf{n}}(\tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}}^{\text{opt}})_{\mathcal{S}; \mathbf{n}} \right)_k + \frac{e_k^h}{\sqrt{m_2}} \right|^2} \\ &= \left\| \frac{1}{\sqrt{m_2}} \Phi_{\mathcal{S}^c; \mathbf{n}}(\tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}}^{\text{opt}})_{\mathcal{S}; \mathbf{n}} + \frac{\mathbf{e}^h}{\sqrt{m_2}} \right\|_2 \\ &= \left\| \frac{1}{\sqrt{m_2}} \Phi_{\mathcal{S}^c; \mathbf{n}} \mathbf{r}'' + \frac{\mathbf{e}^h}{\sqrt{m_2}} \right\|_2. \end{aligned}$$

Using that $\Phi_{\mathcal{S}^c; \mathbf{n}}$ has the restricted isometry property (RIP) of order (s, δ) together with the (reverse) triangle inequality on this last line we now can see that

$$\begin{aligned} \sqrt{1 - \delta} \|\mathbf{r}''\|_2 - \frac{\|\mathbf{e}^h\|_2}{\sqrt{m_2}} &\leq \sqrt{\sum_{k \in [m_2]} \frac{1}{m_2} \left| \left\langle (\tilde{h}_{s'}^{\text{opt}})_{\mathcal{S}^c; \mathbf{z}_k}, T_{\mathcal{S}; \mathbf{n}} \right\rangle_{(\mathcal{D}_{\mathcal{S}}, \mu_{\mathcal{S}})} + e_k^h \right|^2} \\ &\leq \sqrt{1 + \delta} \|\mathbf{r}''\|_2 + \frac{\|\mathbf{e}^h\|_2}{\sqrt{m_2}}. \end{aligned}$$

After subtracting $\|\mathbf{r}''\|_2$ from the quantities in the inequality above, we finally use the bounds $(\sqrt{1 + \delta} - 1) \frac{\sqrt{1 + \delta} + 1}{\sqrt{1 + \delta} + 1} = \frac{\delta}{\sqrt{1 + \delta} + 1} < \frac{\delta}{2}$ and $(\sqrt{1 - \delta} - 1) \frac{\sqrt{1 - \delta} + 1}{\sqrt{1 - \delta} + 1} = \frac{-\delta}{\sqrt{1 - \delta} + 1} \geq -\frac{2}{3}\delta$ to finish the proof. \square

Lemma 10 yields an alternate entry identification technique to that provided in Section 4.1 of [12]. In particular, if $\mathcal{S}^c = [D] \setminus \{j\}$ for some $j \in [D]$ the inner products $\left\langle (\tilde{h}_{s'}^{\text{opt}})_{\mathcal{S}^c; \mathbf{z}_k}, T_{\mathcal{S}; \mathbf{n}} \right\rangle_{(\mathcal{D}_{\mathcal{S}}, \mu_{\mathcal{S}})}$ are just one-dimensional integrals that can be computed to high accuracy for any desired $\mathbf{n} \in \mathcal{I}_{N, d}$ using only $\mathcal{O}(N)$ function evaluations of $(\tilde{h}_{s'}^{\text{opt}})_{\mathcal{S}^c; \mathbf{z}_k} : \mathcal{D}_{\mathcal{S}} \rightarrow \mathbb{C}$ via, e.g., a quadrature rule whenever the the basis functions in the j^{th} -dimension, $\mathcal{B}_j = \{T_{\mathcal{S}; \mathbf{n}} \mid \mathbf{n} \in \mathcal{I}_{N, d}\}$, are polynomials of degree at most N . If \mathcal{B}_j is either the Fourier or Chebyshev basis and N is very large then these one-dimensional integrals can also be computed for all $\mathbf{n} \in \mathcal{I}_{N, d}$ in sublinear-in- N time since $(\tilde{h}_{s'}^{\text{opt}})_{\mathcal{S}^c; \mathbf{z}_k}$ will be \mathcal{B}_j -sparse (see, e.g., [20, 21, 19, 24, 25, 3, 22, 26, 39, 31, 23, 5]).

When $|\mathcal{S}^c| \ll D - 1$ the situation becomes more difficult. However, to efficiently evaluate the higher-dimensional inner products $\left\langle (\tilde{h}_{s'}^{\text{opt}})_{\mathcal{S}^c; \mathbf{z}_k}, T_{\mathcal{S}; \mathbf{n}} \right\rangle_{(\mathcal{D}_{\mathcal{S}}, \mu_{\mathcal{S}})}$ that arise in these settings one can instead utilize non-adaptive random sampling techniques motivated by compressive sensing theory. The following lemma does this by quantifying how well the estimator

$$\frac{1}{m_1} \sum_{\ell \in [m_1]} (\tilde{h}_{s'}^{\text{opt}})_{\mathcal{S}^c; \mathbf{z}_k}(\mathbf{w}_\ell) \overline{T_{\mathcal{S}; \mathbf{n}}(\mathbf{w}_\ell)}$$

based on the m_1 randomly chosen grid points $\{\mathbf{w}_\ell\}_{\ell \in [m_1]} \subset \mathcal{D}_S$ approximates all such $\left\langle (\tilde{h}_{s'}^{\text{opt}})_{S^c; \mathbf{z}_k}, T_{S; \mathbf{n}} \right\rangle_{(\mathcal{D}_S, \mu_S)}$.

Lemma 11. *Let $\tilde{\delta} \in (0, 1)$, $S \subset [D]$, and $\{\mathbf{w}_\ell\}_{\ell \in [m_1]} \subset \mathcal{D}_S$ be m_1 sampling points drawn independently at random according to μ_S in order to form a zero-padded random sampling matrix $\Phi_{S; \mathbf{0}} \in \mathbb{C}^{m_1 \times |\mathcal{I}_{N,d}|}$ for the BOS \mathcal{B}_S in (2.3) with entries*

$$(\Phi_{S; \mathbf{0}})_{\ell, \mathbf{q}} := \begin{cases} T_{S; \mathbf{q}}(\mathbf{w}_\ell) & \text{if } \mathbf{q}_{S^c} = \mathbf{0} \text{ \& } \mathbf{q} \in \mathcal{I}_{N,d} \\ 0 & \text{otherwise} \end{cases} \quad (4.13)$$

indexed by $\ell \in [m_1]$ and $\mathbf{q} \in \mathcal{I}_{N,d}$. Suppose the nonzero columns of $\frac{1}{\sqrt{m_1}} \Phi_{S; \mathbf{0}}$ have the restricted isometry property (RIP) of order $(2, \tilde{\delta})$, and let

$$e_k^h := \frac{1}{m_1} \sum_{\ell \in [m_1]} (\tilde{h}_{s'}^{\text{opt}})_{S^c; \mathbf{z}_k}(\mathbf{w}_\ell) \overline{T_{S; \mathbf{n}}(\mathbf{w}_\ell)} - \left\langle (\tilde{h}_{s'}^{\text{opt}})_{S^c; \mathbf{z}_k}, T_{S; \mathbf{n}} \right\rangle_{(\mathcal{D}_S, \mu_S)}$$

for any desired $\mathbf{n} \in \mathcal{I}_{N,d} \subseteq [N]^D$, function \tilde{h} as per (4.2), and point $\mathbf{z}_k \in \mathcal{D}_{S^c}$. Then,

$$|e_k^h| \leq \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}} \right\|_2 \tilde{\delta} \left\| (\Phi_{S^c; \mathbf{0}; \mathbf{z}_k})_{\Omega_{\tilde{h}, s'}^{\text{opt}}} \right\|_2 \quad (4.14)$$

where $\Phi_{S^c; \mathbf{0}; \mathbf{z}_k}$ is defined as in (2.6).

Proof. We begin by noting that

$$\frac{1}{m_1} \sum_{\ell \in [m_1]} (\tilde{h}_{s'}^{\text{opt}})_{S^c; \mathbf{z}_k}(\mathbf{w}_\ell) \overline{T_{S; \mathbf{n}}(\mathbf{w}_\ell)} = \sum_{\mathbf{q} \in \Omega_{\tilde{h}, s'}^{\text{opt}}} \tilde{r}_{\mathbf{q}} T_{S^c; \mathbf{q}}(\mathbf{z}_k) \nu_S(\mathbf{q}, \mathbf{n}) \quad (4.15)$$

where

$$\begin{aligned} \nu_S(\mathbf{q}, \mathbf{n}) &:= \sum_{\ell \in [m_1]} \frac{1}{\sqrt{m_1}} T_{S; \mathbf{q}}(\mathbf{w}_\ell) \frac{1}{\sqrt{m_1}} \overline{T_{S; \mathbf{n}}(\mathbf{w}_\ell)} \\ &= \left\langle \frac{1}{\sqrt{m_1}} (\Phi_{S; \mathbf{0}})_{(\mathbf{q}, \mathbf{0})_S}, \frac{1}{\sqrt{m_1}} (\Phi_{S; \mathbf{0}})_{(\mathbf{n}, \mathbf{0})_S} \right\rangle. \end{aligned}$$

Appealing to standard results concerning coherence in, e.g., Chapter 6 of [17] one can see that $|\nu_S(\mathbf{q}, \mathbf{n}) - 1| \leq \tilde{\delta}$ holds if $\mathbf{n}_S = \mathbf{q}_S$, and that $|\nu_S(\mathbf{q}, \mathbf{n})| \leq \tilde{\delta}$ holds if $\mathbf{n}_S \neq \mathbf{q}_S$.

Let $\Omega' := \{\mathbf{q} \in \Omega_{\tilde{h}, s'}^{\text{opt}} \mid \mathbf{q}_S = \mathbf{n}_S\} \subset \Omega_{\tilde{h}, s'}^{\text{opt}}$, and $\Omega'' := \Omega_{\tilde{h}, s'}^{\text{opt}} \setminus \Omega'$. Using (4.15) one has that

$$\left| \frac{1}{m_1} \sum_{\ell \in [m_1]} (\tilde{h}_{s'}^{\text{opt}})_{S^c; \mathbf{z}_k}(\mathbf{w}_\ell) \overline{T_{S; \mathbf{n}}(\mathbf{w}_\ell)} - \sum_{\mathbf{q} \in \Omega'} \tilde{r}_{\mathbf{q}} T_{S^c; \mathbf{q}}(\mathbf{z}_k) \right|$$

$$\begin{aligned}
&= \left| \sum_{\mathbf{q} \in \Omega'} \tilde{r}_{\mathbf{q}} T_{\mathcal{S}^c; \mathbf{q}}(\mathbf{z}_k) (\nu_{\mathcal{S}}(\mathbf{q}, \mathbf{n}) - 1) + \sum_{\mathbf{q} \in \Omega''} \tilde{r}_{\mathbf{q}} T_{\mathcal{S}^c; \mathbf{q}}(\mathbf{z}_k) \nu_{\mathcal{S}}(\mathbf{q}, \mathbf{n}) \right| \\
&\leq \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, \tilde{h}', \tilde{h}''}}^{\text{opt}} \right\|_2 \sqrt{\sum_{\mathbf{q} \in \Omega'} |T_{\mathcal{S}^c; \mathbf{q}}(\mathbf{z}_k) (\nu_{\mathcal{S}}(\mathbf{q}, \mathbf{n}) - 1)|^2 + \sum_{\mathbf{q} \in \Omega''} |T_{\mathcal{S}^c; \mathbf{q}}(\mathbf{z}_k) \nu_{\mathcal{S}}(\mathbf{q}, \mathbf{n})|^2}
\end{aligned}$$

where the last inequality follows from Cauchy-Schwarz. Continuing from this the last line we can further see that

$$\begin{aligned}
&\left| \frac{1}{m_1} \sum_{\ell \in [m_1]} (\tilde{h}_{s'}^{\text{opt}})_{\mathcal{S}^c; \mathbf{z}_k}(\mathbf{w}_{\ell}) \overline{T_{\mathcal{S}; \mathbf{n}}(\mathbf{w}_{\ell})} - \sum_{\mathbf{q} \in \Omega'} \tilde{r}_{\mathbf{q}} T_{\mathcal{S}^c; \mathbf{q}}(\mathbf{z}_k) \right| \\
&\leq \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, \tilde{h}', \tilde{h}''}}^{\text{opt}} \right\|_2 \tilde{\delta} \sqrt{\sum_{\mathbf{q} \in \Omega'} |T_{\mathcal{S}^c; \mathbf{q}}(\mathbf{z}_k)|^2 + \sum_{\mathbf{q} \in \Omega''} |T_{\mathcal{S}^c; \mathbf{q}}(\mathbf{z}_k)|^2} \\
&= \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, \tilde{h}', \tilde{h}''}}^{\text{opt}} \right\|_2 \tilde{\delta} \left\| (\Phi_{\mathcal{S}^c; \mathbf{0}; \mathbf{z}_k})_{\Omega_{\tilde{h}, \tilde{h}', \tilde{h}''}}^{\text{opt}} \right\|_2.
\end{aligned}$$

To finish we note that

$$\sum_{\mathbf{q} \in \Omega'} \tilde{r}_{\mathbf{q}} T_{\mathcal{S}^c; \mathbf{q}}(\mathbf{z}_k) = \left\langle \left(\tilde{\mathbf{r}}_{\Omega_{\tilde{h}, \tilde{h}', \tilde{h}''}}^{\text{opt}} \right)_{\mathcal{S}; \mathbf{n}}, \overline{\Phi_{\mathcal{S}^c; \mathbf{n}; \mathbf{z}_k}} \right\rangle = \left\langle \left(\tilde{h}_{s'}^{\text{opt}} \right)_{\mathcal{S}^c; \mathbf{z}_k}, T_{\mathcal{S}; \mathbf{n}} \right\rangle_{(\mathcal{D}_{\mathcal{S}}, \mu_{\mathcal{S}})}$$

by Lemma 3. The desired result follows. \square

Choose any $\mathbf{n} \in \mathcal{I}_{N, d}$ and $\mathcal{S} \subset [D]$ you like. Using Lemma 11 to approximate the inner product appearing in Lemma 10's (4.11) then yields the following estimator for accurately approximating the ℓ^2 -norm of $\mathbf{r}'' := \tilde{\mathbf{r}}_{\{\mathbf{q} \in \Omega_{\tilde{h}, \tilde{h}', \tilde{h}'}}^{\text{opt}} \mid \mathbf{q}_{\mathcal{S}} = \mathbf{n}_{\mathcal{S}}\}} \in \mathbb{C}^{\mathcal{I}_{N, d}}$ for the coefficient vector $\tilde{\mathbf{r}}$ of any function \tilde{h} as in (4.2). The estimator is defined for any function $u : \mathcal{D} \rightarrow \mathbb{C}$, $\mathcal{S} \subset [D]$, and $\mathbf{n} \in \mathcal{I}_{N, d}$ to be

$$E_{\mathcal{S}; \mathbf{n}}^u := \frac{1}{m_2} \sum_{k \in [m_2]} \left| \frac{1}{m_1} \sum_{\ell \in [m_1]} u_{\mathcal{S}^c; \mathbf{z}_k}(\mathbf{w}_{\ell}) \overline{T_{\mathcal{S}; \mathbf{n}}(\mathbf{w}_{\ell})} \right|^2 \quad (4.16)$$

for fixed nodes $\{\mathbf{w}_{\ell}\}_{\ell \in [m_1]} \subset \mathcal{D}_{\mathcal{S}}$, and $\{\mathbf{z}_k\}_{k \in [m_2]} \subset \mathcal{D}_{\mathcal{S}^c}$. Note that (4.16) is essentially identical to the pairing energy estimator defined in Section 4.2 of [12]. The following lemma provides an error guarantee for this estimator that matches the quality of those in [12] despite having a simpler proof (see Lemma 7 in [12]).

Lemma 12. Let $\mathcal{S} \subset [D]$, $\delta \in (0, 3/4]$, $\tilde{\delta} \in (0, 1/s']$, and $m_1, m_2 \in \mathbb{N}$. Furthermore, suppose that $\{\mathbf{w}_\ell\}_{\ell \in [m_1]} \subset \mathcal{D}_{\mathcal{S}}$, and $\{\mathbf{z}_k\}_{k \in [m_2]} \subset \mathcal{D}_{\mathcal{S}^c}$ satisfy the RIP assumptions concerning (4.13) and (4.10) in Lemmas 11 and 10, respectively. Then, for all $\mathbf{n} \in \mathcal{I}_{N,d} \subseteq [N]^D$ and functions \tilde{h} as per (4.2) one will have

$$\left| \sqrt{E_{\mathcal{S};\mathbf{n}}^{\tilde{h}^{\text{opt}}}} - \|\mathbf{r}''\|_2 \right| \leq \frac{2}{3}\delta\|\mathbf{r}''\|_2 + \sqrt{\frac{7}{4}}\sqrt{\tilde{\delta}} \cdot \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},s'}^{\text{opt}}} \right\|_2,$$

where $\mathbf{r}'' := \tilde{\mathbf{r}}_{\{\mathbf{q} \in \Omega_{\tilde{h},s'}^{\text{opt}} \mid \mathbf{q}_{\mathcal{S}} = \mathbf{n}_{\mathcal{S}}\}}$.

Proof. Applying Lemma 10 we can immediately see that

$$\left| \sqrt{E_{\mathcal{S};\mathbf{n}}^{\tilde{h}^{\text{opt}}}} - \|\mathbf{r}''\|_2 \right| \leq \frac{2}{3}\delta\|\mathbf{r}''\|_2 + \frac{\|\mathbf{e}^h\|_2}{\sqrt{m_2}} \quad (4.17)$$

where $\mathbf{e}^h \in \mathbb{C}^{m_2}$ has its entries given by

$$e_k^h := \frac{1}{m_1} \sum_{\ell \in [m_1]} (\tilde{h}_{s'}^{\text{opt}})_{\mathcal{S}^c; \mathbf{z}_k}(\mathbf{w}_\ell) \overline{T_{\mathcal{S};\mathbf{n}}(\mathbf{w}_\ell)} - \left\langle (\tilde{h}_{s'}^{\text{opt}})_{\mathcal{S}^c; \mathbf{z}_k}, T_{\mathcal{S};\mathbf{n}} \right\rangle_{(\mathcal{D}_{\mathcal{S}}, \mu_{\mathcal{S}})}.$$

Thus, it suffices to bound $\|\mathbf{e}^h\|_2$ in order to obtain our final result.

Applying Lemma 11 we can see that

$$\begin{aligned} \|\mathbf{e}^h\|_2^2 &\leq \sum_{k \in [m_2]} \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},s'}^{\text{opt}}} \right\|_2^2 \tilde{\delta}^2 \left\| (\Phi_{\mathcal{S}^c; \mathbf{0}; \mathbf{z}_k})_{\Omega_{\tilde{h},s'}^{\text{opt}}} \right\|_2^2 \\ &= \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},s'}^{\text{opt}}} \right\|_2^2 \tilde{\delta}^2 \left\| (\Phi_{\mathcal{S}^c; \mathbf{0}})_{\Omega_{\tilde{h},s'}^{\text{opt}}} \right\|_{\mathbb{F}}^2, \end{aligned}$$

where we have used that $(\Phi_{\mathcal{S}^c; \mathbf{0}; \mathbf{z}_k})_{\Omega_{\tilde{h},s'}^{\text{opt}}}$ are the rows of the submatrix $(\Phi_{\mathcal{S}^c; \mathbf{0}})_{\Omega_{\tilde{h},s'}^{\text{opt}}} \in \mathbb{C}^{m_2 \times s'}$ of $\Phi_{\mathcal{S}^c; \mathbf{0}}$ in (4.10). Using the RIP property of the nonzero columns of $\frac{1}{\sqrt{m_2}}\Phi_{\mathcal{S}^c; \mathbf{0}}$ we can now finish bounding $\|\mathbf{e}^h\|_2^2$ by noting that

$$\|\mathbf{e}^h\|_2^2 \leq \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},s'}^{\text{opt}}} \right\|_2^2 \tilde{\delta}^2 \cdot s' \cdot m_2(1 + \delta) \leq \frac{7}{4}m_2 \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},s'}^{\text{opt}}} \right\|_2^2 \tilde{\delta}$$

where we have used that $\delta \in (0, 3/4]$ and that $\tilde{\delta} \in (0, 1/s']$. Substituting this last bound into (4.17) now finishes the proof. \square

Though useful, Lemma 12 presupposes that one has access to noiseless samples from $\tilde{h}_{s'}^{\text{opt}}$. This will rarely be the case in practice. The next lemma bounds the error of the estimator (4.16) in the setting where one instead has noisy samples from $\tilde{h}_{s'}^{\text{opt}}$. Such noisy samples will be represented with the help of an arbitrary additive noise/error function, $e_h : \mathcal{D} \rightarrow \mathbb{C}$, below.

Lemma 13. *Let $\mathcal{S} \subset [D]$, $\delta \in (0, 3/4]$, $\tilde{\delta} \in (0, 1/s']$, and $m_1, m_2 \in \mathbb{N}$. Furthermore, suppose that $\{\mathbf{w}_\ell\}_{\ell \in [m_1]} \subset \mathcal{D}_\mathcal{S}$, and $\{\mathbf{z}_k\}_{k \in [m_2]} \subset \mathcal{D}_{\mathcal{S}^c}$ satisfy the RIP assumptions concerning (4.13) and (4.10) in Lemmas 11 and 10, respectively. Then, for all $\mathbf{n} \in \mathcal{I}_{N,d} \subseteq [N]^D$, \tilde{h} as per (4.2), and additive error functions $e_h : \mathcal{D} \rightarrow \mathbb{C}$ one will have*

$$\left| \sqrt{E_{\mathcal{S};\mathbf{n}}^{\tilde{h}_{s'}^{\text{opt}} + e_h}} - \|\mathbf{r}''\|_2 \right| \leq \frac{2}{3}\delta\|\mathbf{r}''\|_2 + \sqrt{\frac{7}{4}}\tilde{\delta} \cdot \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},s'}^{\text{opt}}} \right\|_2 + \frac{\sqrt{2}\|\mathcal{E}_\mathcal{S}^h\|_{\mathbb{F}}}{\sqrt{m_1 m_2}},$$

where $\mathbf{r}'' := \tilde{\mathbf{r}}_{\{\mathbf{q} \in \Omega_{\tilde{h},s'}^{\text{opt}} \mid \mathbf{q}_\mathcal{S} = \mathbf{n}_\mathcal{S}\}}$ and $\mathcal{E}_\mathcal{S}^h \in \mathbb{C}^{m_1 \times m_2}$ has entries $(\mathcal{E}_\mathcal{S}^h)_{\ell,k} = e_h(\varrho_\mathcal{S}(\mathbf{w}_\ell, \mathbf{z}_k))$ with the permutation function $\varrho_\mathcal{S} : \mathcal{D}_\mathcal{S} \times \mathcal{D}_{\mathcal{S}^c} \rightarrow \mathcal{D}$ defined in (2.4).

Proof. Note that

$$\sqrt{E_{\mathcal{S};\mathbf{n}}^{\tilde{h}_{s'}^{\text{opt}}}} - \sqrt{E_{\mathcal{S};\mathbf{n}}^{e_h}} \leq \sqrt{E_{\mathcal{S};\mathbf{n}}^{\tilde{h}_{s'}^{\text{opt}} + e_h}} \leq \sqrt{E_{\mathcal{S};\mathbf{n}}^{\tilde{h}_{s'}^{\text{opt}}}} + \sqrt{E_{\mathcal{S};\mathbf{n}}^{e_h}}$$

by the (reverse) triangle inequality. As a result one can immediately see that

$$\begin{aligned} \left| \sqrt{E_{\mathcal{S};\mathbf{n}}^{\tilde{h}_{s'}^{\text{opt}} + e_h}} - \|\mathbf{r}''\|_2 \right| &\leq \left| \sqrt{E_{\mathcal{S};\mathbf{n}}^{\tilde{h}_{s'}^{\text{opt}} + e_h}} - \sqrt{E_{\mathcal{S};\mathbf{n}}^{\tilde{h}_{s'}^{\text{opt}}}} \right| + \left| \sqrt{E_{\mathcal{S};\mathbf{n}}^{\tilde{h}_{s'}^{\text{opt}}}} - \|\mathbf{r}''\|_2 \right| \\ &\leq \sqrt{E_{\mathcal{S};\mathbf{n}}^{e_h}} + \frac{2}{3}\delta\|\mathbf{r}''\|_2 + \sqrt{\frac{7}{4}}\tilde{\delta} \cdot \left\| \tilde{\mathbf{r}}_{\Omega_{\tilde{h},s'}^{\text{opt}}} \right\|_2 \end{aligned}$$

where the bound on the second term above follows from Lemma 12. It remains to show that $\sqrt{E_{\mathcal{S};\mathbf{n}}^{e_h}} \leq \frac{\sqrt{2}\|\mathcal{E}_\mathcal{S}^h\|_{\mathbb{F}}}{\sqrt{m_1 m_2}}$.

Define $\mathbf{v} \in \mathbb{C}^{m_2}$ by $v_k := \frac{1}{\sqrt{m_1}} \left\langle (\mathcal{E}_\mathcal{S}^h)_k, \frac{1}{\sqrt{m_1}} (\Phi_{\mathcal{S};\mathbf{0}})_{(\mathbf{n},\mathbf{0})_\mathcal{S}} \right\rangle$ where $\Phi_{\mathcal{S};\mathbf{0}}$ is defined in (4.13), and note that $|v_k| \leq \frac{1}{\sqrt{m_1}} \|(\mathcal{E}_\mathcal{S}^h)_k\|_2 \left\| \frac{1}{\sqrt{m_1}} (\Phi_{\mathcal{S};\mathbf{0}})_{(\mathbf{n},\mathbf{0})_\mathcal{S}} \right\|_2$. Furthermore,

$$\left\| \frac{1}{\sqrt{m_1}} (\Phi_{\mathcal{S};\mathbf{0}})_{(\mathbf{n},\mathbf{0})_\mathcal{S}} \right\|_2 = \left\| \frac{1}{\sqrt{m_1}} (\Phi_{\mathcal{S};\mathbf{0}})_{\{(\mathbf{n},\mathbf{0})_\mathcal{S}, (\tilde{\mathbf{n}},\mathbf{0})_\mathcal{S}\}} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right\|_2 \leq \sqrt{1 + \tilde{\delta}} \leq \sqrt{1 + \frac{1}{s'}} \leq \sqrt{2}$$

for any $\tilde{\mathbf{n}} \neq \mathbf{n}$. Thus, $|v_k| \leq \sqrt{\frac{2}{m_1}} \|(\mathcal{E}_\mathcal{S}^h)_k\|_2$. As a result, $\sqrt{E_{\mathcal{S};\mathbf{n}}^{e_h}} = \frac{1}{\sqrt{m_2}} \|\mathbf{v}\|_2 \leq \frac{\sqrt{2}\|\mathcal{E}_\mathcal{S}^h\|_{\mathbb{F}}}{\sqrt{m_1 m_2}}$. \square

For any given $\mathcal{S} \subseteq [D]$ we denote the power set of $N^{\mathcal{S}}$ by $\mathcal{P}(N^{\mathcal{S}})$. In the final theorem of this subsection we will prove that the energy estimator in (4.16) can be used for an arbitrary s' -sparse function $\tilde{h}_{s'}^{\text{opt}}$ to define a new set-valued function $\mathcal{F}_{\mathcal{S}}^{s'} : \mathcal{P}(N^{\mathcal{S}}) \rightarrow \mathcal{P}(N^{\mathcal{S}})$ for each \mathcal{S} which, when given any subset $\mathcal{T} \subset N^{\mathcal{S}}$ containing the heavy set $\Omega_{\mathcal{S}}^{\alpha, s'}$ as per (4.3) as input, will output a smaller subset $\mathcal{T}' \subset \mathcal{T}$ which still contains $\Omega_{\mathcal{S}}^{\alpha, s'}$. These set-valued functions were also called “energetic-index sieve function” in Section 4.1 and will then be used to iteratively build up subsets \mathcal{T} of controlled cardinality for larger and larger sets of indices \mathcal{S} until we eventually have a set of full index vectors $\mathcal{T}'' \subset \mathcal{I}_{N, d}$ which contains all of $\Omega_{[D]}^{\alpha, s'}$. This set of full index vectors \mathcal{T}'' will then be able to be used as an accurate estimate of $\Omega_{\tilde{h}, s'}^{\text{opt}}$, the support of $\tilde{h}_{s'}^{\text{opt}}$.

Before we can state our final theorem we must define the set-valued functions $\mathcal{F}_{\mathcal{S}}^{s'} : \mathcal{P}(N^{\mathcal{S}}) \rightarrow \mathcal{P}(N^{\mathcal{S}})$ in question. For a given $\mathcal{S} \subseteq [D]$, $\mathcal{T} \subseteq N^{\mathcal{S}}$, \tilde{h} as per (4.2), and additive error function $e_h : \mathcal{D} \rightarrow \mathbb{C}$, let an ordering of the elements of \mathcal{T} , $\mathbf{n}_1, \mathbf{n}_2, \dots, \mathbf{n}_{|\mathcal{T}|} \in \mathcal{T}$, be defined by

$$E_{\mathcal{S}; \mathbf{n}_1}^{\tilde{h}_{s'}^{\text{opt}} + e_h} \geq E_{\mathcal{S}; \mathbf{n}_2}^{\tilde{h}_{s'}^{\text{opt}} + e_h} \geq E_{\mathcal{S}; \mathbf{n}_3}^{\tilde{h}_{s'}^{\text{opt}} + e_h} \geq \dots \geq E_{\mathcal{S}; \mathbf{n}_{|\mathcal{T}|}}^{\tilde{h}_{s'}^{\text{opt}} + e_h} \quad (4.18)$$

with ties broken lexicographically. We define $\mathcal{F}_{\mathcal{S}}^{s'}$ based on this ordering by

$$\mathcal{F}_{\mathcal{S}}^{s'}(\mathcal{T}) := \{\mathbf{n}_1, \mathbf{n}_2, \dots, \mathbf{n}_{\min(s', |\mathcal{T}|)}\} \subseteq \mathcal{T}. \quad (4.19)$$

The following theorem proves that $\Omega_{\mathcal{S}}^{\alpha, s'} \cap \mathcal{T} \subseteq \mathcal{F}_{\mathcal{S}}^{s'}(\mathcal{T})$ provided that the additive error e_h is sufficiently mild.

Theorem 10 (Entry Identification and Pairing). *Let $\mathcal{S} \subseteq [D]$ with $|\mathcal{S}| > 0$, $\delta \in (0, 1/2]$, $\tilde{\delta} \in (0, \frac{1}{256\alpha^2 s'}]$, and $m_1, m_2 \in \mathbb{N}$. Furthermore, suppose that $\{\mathbf{w}_\ell\}_{\ell \in [m_1]} \subset \mathcal{D}_{\mathcal{S}}$, and $\{\mathbf{z}_k\}_{k \in [m_2]} \subset \mathcal{D}_{\mathcal{S}^c}$ satisfy the RIP assumptions concerning (4.13) and (4.10) in Lemmas 11 and 10, respectively. Then, $\Omega_{\mathcal{S}}^{\alpha, s'} \cap \mathcal{T} \subseteq \mathcal{F}_{\mathcal{S}}^{s'}(\mathcal{T})$ for all s' -sparse $\tilde{h} = \tilde{h}_{s'}^{\text{opt}}$, $\mathcal{T} \subseteq N^{\mathcal{S}}$, and additive error functions $e_h : \mathcal{D} \rightarrow \mathbb{C}$ provided that*

$$\|\tilde{\mathbf{r}}\|_2 = \|\tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}}\|_2 > \frac{6\alpha\sqrt{s'}}{\sqrt{m_1 m_2}} \cdot \|\mathcal{E}_{\mathcal{S}}^h\|_{\text{F}} \quad (4.20)$$

holds, where $\mathcal{E}_{\mathcal{S}}^h \in \mathbb{C}^{m_1 \times m_2}$ has entries $(\mathcal{E}_{\mathcal{S}}^h)_{\ell, k} = e_h(\varrho_{\mathcal{S}}(\mathbf{w}_\ell, \mathbf{z}_k))$ with the permutation function $\varrho_{\mathcal{S}} : \mathcal{D}_{\mathcal{S}} \times \mathcal{D}_{\mathcal{S}^c} \rightarrow \mathcal{D}$ defined in (2.4).

Proof. We will focus on the case where $|\mathcal{T}| > s'$ since the result holds trivially when $|\mathcal{T}| \leq s'$. Suppose for the sake of contradiction that $\mathbf{m} \in \Omega_{\mathcal{S}}^{\alpha, s'} \cap \mathcal{T}$, but that $\mathbf{m} \notin \mathcal{F}_{\mathcal{S}}^{s'}(\mathcal{T})$. It must then be the case that $E_{\mathcal{S}; \mathbf{k}}^{\tilde{h}_{s'}^{\text{opt}} + e_h} \geq E_{\mathcal{S}; \mathbf{m}}^{\tilde{h}_{s'}^{\text{opt}} + e_h}$ for some $\mathbf{k} \in \mathcal{T}$ with $\mathbf{k}_{\mathcal{S}} \notin \Omega_{s', \mathcal{S}}^{\text{opt}} :=$

$\{\mathbf{q}_{\mathcal{S}} \mid \mathbf{q} \in \Omega_{\tilde{h}, s'}^{\text{opt}}\}$ since $\Omega_{\mathcal{S}}^{\alpha, s'} \subset \Omega_{s', \mathcal{S}}^{\text{opt}}$ and $|\mathcal{F}_{\mathcal{S}}^{s'}(\mathcal{T})| = s' \geq |\Omega_{s', \mathcal{S}}^{\text{opt}}|$. Thus, $\|\tilde{\mathbf{r}}_{\mathcal{S}; \mathbf{k}}\|_2 = 0$. As a result, Lemma 13 implies that

$$\sqrt{E_{\mathcal{S}; \mathbf{k}}^{\tilde{h}_{s'}^{\text{opt}} + e_h}} \leq \sqrt{\frac{7}{4}\tilde{\delta}} \cdot \|\tilde{\mathbf{r}}\|_2 + \frac{\sqrt{2}\|\mathcal{E}_{\mathcal{S}}^h\|_{\text{F}}}{\sqrt{m_1 m_2}} \leq \frac{\sqrt{7}}{32} \frac{\|\tilde{\mathbf{r}}\|_2}{\alpha\sqrt{s'}} + \frac{\sqrt{2}\|\mathcal{E}_{\mathcal{S}}^h\|_{\text{F}}}{\sqrt{m_1 m_2}} =: \epsilon. \quad (4.21)$$

On the other hand, Lemma 13 also implies that $\sqrt{E_{\mathcal{S}; \mathbf{m}}^{\tilde{h}_{s'}^{\text{opt}} + e_h}} \geq \|\tilde{\mathbf{r}}_{\mathcal{S}; \mathbf{m}}\|_2(1 - \frac{2}{3}\delta) - \epsilon$. Combining this with (4.21), we have $\frac{2\epsilon}{1 - \frac{2}{3}\delta} \geq \|\tilde{\mathbf{r}}_{\mathcal{S}; \mathbf{m}}\|_2$. Since $\delta \leq \frac{1}{2}$, it must also be the case that

$$3\epsilon \geq \frac{2\epsilon}{1 - \frac{2}{3}\delta} \geq \|\tilde{\mathbf{r}}_{\mathcal{S}; \mathbf{m}}\|_2 \geq \frac{\|\tilde{\mathbf{r}}\|_2}{\alpha\sqrt{s'}}.$$

However, it is impossible that $4\alpha\sqrt{s'}\epsilon \geq \|\tilde{\mathbf{r}}\|_2$ since by assumption

$$3\alpha\sqrt{s'}\epsilon = 3\alpha\sqrt{s'} \left(\frac{\sqrt{7}}{32} \frac{\|\tilde{\mathbf{r}}\|_2}{\alpha\sqrt{s'}} + \frac{\sqrt{2}\|\mathcal{E}_{\mathcal{S}}^h\|_{\text{F}}}{\sqrt{m_1 m_2}} \right) < 3\alpha\sqrt{s'} \left(\frac{\sqrt{7}}{32} \frac{\|\tilde{\mathbf{r}}\|_2}{\alpha\sqrt{s'}} + \frac{\sqrt{2}\|\tilde{\mathbf{r}}\|_2}{6\alpha\sqrt{s'}} \right) < \|\tilde{\mathbf{r}}\|_2.$$

Hence, $\mathbf{m} \in \Omega_{\mathcal{S}}^{\alpha, s'} \cap \mathcal{T} \implies \mathbf{m} \in \mathcal{F}_{\mathcal{S}}^{s'}(\mathcal{T})$. \square

Theorem 10 forms the basis of our support identification strategy. As such, it behooves us to investigate its associate resource demands and error performance more closely. We do this in the next subsection.

4.2.1 Associated Runtime, Sampling, and Error Bounds

The following lemmas provide evaluation complexity, sampling, and error bounds for the set valued functions $\mathcal{F}_{\mathcal{S}}^{s'} : \mathcal{P}(N^{\mathcal{S}}) \rightarrow \mathcal{P}(N^{\mathcal{S}})$ defined in (4.18) – (4.19). We will begin by providing more meaningful error bounds for the case where the function \tilde{h} in question is not exactly BOPB-sparse.

Lemma 14. *Let $\mathcal{S} \subseteq [D]$ with $|\mathcal{S}| > 0$, $\delta \in (0, 1/2]$, $\tilde{\delta} \in (0, \frac{1}{256\alpha^2 s'})$, $\gamma \in \mathbb{R}^+$, and $m_1, m_2 \in \mathbb{N}$. Furthermore, suppose that $\{\mathbf{w}_{\ell}\}_{\ell \in [m_1]} \subset \mathcal{D}_{\mathcal{S}}$, and $\{\mathbf{z}_k\}_{k \in [m_2]} \subset \mathcal{D}_{\mathcal{S}^c}$ satisfy the RIP assumptions concerning (4.13) and (4.10) in Lemmas 11 and 10, respectively. Finally, suppose also that $e_h = \tilde{h} - \tilde{h}_{s'}^{\text{opt}} + e'$ for an arbitrary function $e' : \mathcal{D} \rightarrow \mathbb{C}$ with $\sup_{\boldsymbol{\xi} \in \mathcal{D}} |e'(\boldsymbol{\xi})| \leq \gamma$. Then, the additive sampling error $\mathcal{E}_{\mathcal{S}}^h \in \mathbb{C}^{m_1 \times m_2}$ satisfies*

$$\frac{\|\mathcal{E}_{\mathcal{S}}^h\|_{\text{F}}}{\sqrt{m_1 m_2}} \leq \sqrt{\frac{771}{512}} \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}} \right\|_2 + \sqrt{\frac{771}{512 s'}} \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}} \right\|_1 + \gamma$$

where $\mathcal{E}_{\mathcal{S}}^h$ has entries $(\mathcal{E}_{\mathcal{S}}^h)_{\ell, k} = e_h(\varrho_{\mathcal{S}}(\mathbf{w}_{\ell}, \mathbf{z}_k))$ as in Theorem 10.

Proof. Note that

$$\begin{aligned}
\frac{\|\mathcal{E}_S^h\|_F}{\sqrt{m_1 m_2}} &\leq \gamma + \frac{1}{\sqrt{m_1 m_2}} \sqrt{\sum_{\ell, k} \left| \left(\tilde{h} - \tilde{h}_{s'}^{\text{opt}} \right) \left(\varrho_S(\mathbf{w}_\ell, \mathbf{z}_k) \right) \right|^2} \\
&= \gamma + \frac{1}{\sqrt{m_1 m_2}} \sqrt{\sum_{\ell, k} \left| \left\langle \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}}, \Phi_{(\ell, k)}^* \right\rangle \right|^2} \\
&= \left\| \frac{1}{\sqrt{m_1 m_2}} \Phi \left(\tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}} \right) \right\|_2 + \gamma, \tag{4.22}
\end{aligned}$$

where $\Phi \in \mathbb{C}^{m_1 m_2 \times |\mathcal{I}_{N, d}|}$ has entries given by $\Phi_{(\ell, k), \mathbf{n}} = T_{\mathbf{n}}(\varrho_S(\mathbf{w}_\ell, \mathbf{z}_k))$. Note also that $\frac{1}{\sqrt{m_1 m_2}} \Phi$ consists of a subset of the columns of the Kronecker product $\left(\frac{1}{\sqrt{m_1}} \Phi_{S; \mathbf{0}} \right) \otimes \left(\frac{1}{\sqrt{m_2}} \Phi_{S^c; \mathbf{0}} \right)$ where $\Phi_{S; \mathbf{0}} \in \mathbb{C}^{m_1 \times |\mathcal{I}_{N, d}|}$ is defined in (4.13), and $\Phi_{S^c; \mathbf{0}} \in \mathbb{C}^{m_2 \times |\mathcal{I}_{N, d}|}$ is defined in (4.10). Furthermore, Proposition 6.6 of [17] implies that the nonzero columns of $\frac{1}{\sqrt{m_1}} \Phi_{S; \mathbf{0}}$ also has the RIP of order $(s', \frac{1}{256\alpha^2})$ since it has the RIP of order $(2, \frac{1}{256\alpha^2 s'})$. Hence, $\frac{1}{\sqrt{m_1 m_2}} \Phi$ has the RIP of order $(s', (1 + \frac{1}{256\alpha^2})(1 + \frac{1}{2}) - 1)$ by Lemma 2 of [15], and consequently of order $(s', \frac{259}{512})$ for $\alpha \geq 1$. Returning to (4.22), we can now use Lemma 4 to see that

$$\frac{\|\mathcal{E}_S^h\|_F}{\sqrt{m_1 m_2}} \leq \sqrt{\frac{771}{512}} \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}} \right\|_2 + \sqrt{\frac{771}{512 s'}} \left\| \tilde{\mathbf{r}} - \tilde{\mathbf{r}}_{\Omega_{\tilde{h}, s'}^{\text{opt}}} \right\|_1 + \gamma$$

as desired. \square

The next lemma tells us how many evaluation points we need to randomly generate in Lemmas 11 and 10 before we can be sure to have the RIP properties required by both Theorem 10 and Lemma 14 above hold with high probability.

Lemma 15. *Let $\mathcal{S} \subseteq [D]$, $\delta \in (0, 1/2]$, and $\tilde{\delta} = (0, \frac{1}{256\alpha^2 s'}]$. Furthermore, suppose that $m_1, m_2, s', N, D \in \mathbb{Z}^+ \setminus \{1\}$, $d \in \mathbb{Z} \cap [1, D]$, and $p \in (0, 1)$ satisfy*

$$m_1 \geq a_1 \alpha^4 K_S^2 (s')^2 \cdot \max \left\{ d \ln \left(\frac{DN}{d} \right) \ln(m_1), \ln(p^{-1}) \right\},$$

and

$$m_2 \geq a_2 K_{S^c}^2 \delta^{-2} s' \cdot \max \left\{ d \ln^2(s') \ln \left(\frac{DN}{d} \right) \ln(m_2), \ln(p^{-1}) \right\},$$

where $a_1, a_2 \in \mathbb{R}^+$ are universal constants. Then, the samples $\{\mathbf{w}_\ell\}_{\ell \in [m_1]} \subset \mathcal{D}_S$ and $\{\mathbf{z}_k\}_{k \in [m_2]} \subset \mathcal{D}_{S^c}$ will both simultaneously satisfy their respective RIP assumptions concerning (4.13) and (4.10) in Lemmas 11 and 10 above with probability at least $1 - p$.

Proof. The bounds on both m_1 and m_2 follow from applications of Theorem 2. To bound m_1 we note that the normalized nonzero columns of (4.13) need to have the RIP of order $(2, \tilde{\delta})$, and have an associated BOS constant of $K_{\mathcal{S}}$. Furthermore, there will never be more than

$$|\mathcal{I}_{N,d}| = \binom{D}{d} N^d \leq \left(\frac{eD}{d}\right)^d N^d = \left(\frac{eDN}{d}\right)^d \quad (4.23)$$

nonzero columns of (4.13) for any choice of $\mathcal{S} \subseteq [D]$. As a consequence we can see that it suffices to have

$$\begin{aligned} m_1 &\geq a'_1 K_{\mathcal{S}}^2 \tilde{\delta}^{-2} \cdot \max \left\{ d \ln \left(\frac{DN}{d} \right) \ln(m_1), \ln(p^{-1}) \right\} \\ &\geq a_1 \alpha^4 K_{\mathcal{S}}^2 (s')^2 \cdot \max \left\{ d \ln \left(\frac{DN}{d} \right) \ln(m_1), \ln(p^{-1}) \right\} \end{aligned}$$

in order to satisfy the required RIP conditions for (4.13) with probability at least $1 - p/2$.

To bound m_2 we note that the normalized nonzero columns of (4.10) need to have the RIP of order (s', δ) , and have an associated BOS constant of $K_{\mathcal{S}^c}$. As a result, (4.23) together with Theorem 2 implies that it suffices to have

$$m_2 \geq a_2 K_{\mathcal{S}^c}^2 \delta^{-2} s' \cdot \max \left\{ d \ln^2(s') \ln \left(\frac{DN}{d} \right) \ln(m_2), \ln(p^{-1}) \right\}$$

in order to satisfy the required RIP conditions for (4.10) with probability at least $1 - p/2$. The final desired probability of success now results from the union bound. \square

Remark 3. *To simplify the appearance of our bounds from Lemma 15 we will make use of the following additional facts and mild assumptions. First, we will assume hereafter that both m_1 and m_2 are less than $|\mathcal{I}_{N,d}|$. We consider this a reasonable assumption given that the techniques presented herein should only be used in situations where this is the case. Furthermore, we will use $\delta = 1/2$ above as this is its largest valid parameter setting, and will also consider α to be a universal constant given that it is ultimately set to a fixed value. Finally, we will also replace our probability of failure parameter p by $c/2D$ for some small constant $c < 0.01$ (for example) in anticipation of wanting to survive a union bound involving $2D - 1$ applications of Lemma 15 for $2D - 1$ different sets of random samples. This will allow us to assert that any at most $2D - 1$ different set valued functions $\mathcal{F}_{\mathcal{S}}^{s'}$ will all simultaneously satisfy both Theorem 10 and Lemma 14 with a “high probability” of at least 0.99. Utilizing these simplifications we obtain the simplified sufficient sampling conditions*

$$m_1 \geq c'_1 K_{\mathcal{S}}^2 (s')^2 \cdot d^2 \ln^2 \left(\frac{DN}{d} \right) \ln(D),$$

$$m_2 \geq c'_2 K_{\mathcal{S}^c}^2 s' \cdot d^2 \ln^2(s') \ln^2\left(\frac{DN}{d}\right) \ln(D)$$

for new absolute constants $c'_1, c'_2 \in \mathbb{R}^+$.

Finally, and perhaps most controversially, we will make the additional assumption above that either (i) the BOS constants K_j are 1 for all but at most $\tilde{d} \in \mathbb{Z} \cap [0, D]$ BOS basis sets \mathcal{B}_j (note that \tilde{d} can be independent of d), or else that (ii) $K_0 = 1$. In either case we will have that both $K_{\mathcal{S}}$ and $K_{\mathcal{S}^c}$ will be bounded above by a constant that depends on only \tilde{d} or d , respectively, for all $\mathcal{S} \subseteq [D]$. In particular, in case (i) we will have that $K_{\mathcal{S}}$ and $K_{\mathcal{S}^c}$ are both at most $K_{\infty}^{\tilde{d}}$, and in case (ii) that $K_{\mathcal{S}}$ and $K_{\mathcal{S}^c}$ are both at most K_{∞}^d . Utilizing this final assumption now allows us to bound the total number of samples we need in order to compute any $2D - 1$ informative $\mathcal{F}_{\mathcal{S}}^{s'}$ sets for any given $2D - 1$ sets \mathcal{S} with high probability (w.h.p.) by either

$$m_1 m_2 \geq c_1 K_{\infty}^{4\tilde{d}} (s')^3 d^4 \cdot \ln^4\left(\frac{DN}{d}\right) \ln^2(s') \ln^2(D)$$

in case (i) (note here that letting $d = D$ still avoids exponential dependence on D in this setting), and by

$$m_1 m_2 \geq c_2 K_{\infty}^{4d} (s')^3 d^4 \cdot \ln^4\left(\frac{DN}{d}\right) \ln^2(s') \ln^2(D)$$

in case (ii), where $c_1, c_2 \in \mathbb{R}^+$ are absolute constants.

We are now ready to demonstrate the numerical performance of our proposed method.

5 Empirical Evaluation

In this section, Algorithm 1 in combination with Algorithm 2 is evaluated numerically for the exactly sparse case with noisy measurements as well as the approximately sparse case. The algorithms were implemented in MATLAB and are publicly available.¹⁷ For the entry identification, we use the pairing approach. In addition to the stopping criterion “ $\|\mathbf{v}_{\text{CE}}\|_2^2 > \|\mathbf{v}_{\text{CEold}}\|_2^2$ or $k \geq \kappa$ ” in line 16 in Algorithm 1, we also stop Algorithm 1 if $\text{supp}(\mathbf{a}^k) = \text{supp}(\mathbf{a}^{k-1}) = \text{supp}(\mathbf{a}^{k-2})$, i.e., the identified index vectors are the same for three consecutive iterations. All time measurements were performed on a computer with 2 x 6-core Intel Xeon CPU E5-2620 v3 (2.40GHz), 64 GB RAM, using 12 threads.

¹⁷See “SHT II: Best s -Term Approximation Guarantees for Bounded Orthonormal Product Bases in Sublinear-Time” on Mark Iwen’s code page <https://www.math.msu.edu/~markiwen/Code.html>.

5.1 Exactly sparse case and noisy measurements

We start with the exactly sparse case, and we consider tensor product basis functions with different bases in $d = D$ spatial dimensions, where we choose d up to 100. We set $N = 200$ and use $\mathcal{I}_{N,d} = \mathcal{I}_{200,d}$ as search space of possible basis indices, where e.g. $|\mathcal{I}_{200,50}| \approx 10^{115}$ and $|\mathcal{I}_{200,100}| \approx 10^{230}$. We set the maximum number of iterations $\kappa := 20$, and we always use $m_{\text{CE}} := 50s$ samples for the coefficient estimation where $s = |\mathcal{S}|$. For every data point in every plot below, we use 100 different randomly generated trial signals

$$f(\boldsymbol{\xi}) = \sum_{\mathbf{n} \in \mathcal{S}} c_{\mathbf{n}} T_{\mathbf{n}}(\boldsymbol{\xi}), \quad (5.1)$$

where we draw the function's support set $\mathcal{S} \subset \mathcal{I}_{N,d}$ uniformly at random without repetition and the coefficients $c_{\mathbf{n}} \in \{-1, 1\}$ uniformly at random.

Below, a trial will always refer to the execution of Algorithm 1 on a particular randomly generated trial function f as defined in (5.1). A failed trial will refer to any trial where Algorithm 1 failed to recover the correct support set \mathcal{S} for f .

We assume that the function evaluations of f are contaminated with (white) Gaussian noise, i.e., we provide Algorithm 1 with noisy samples

$$\mathbf{y}' = \mathbf{y} + \mathbf{g}' = \mathbf{y} + \sigma \frac{\|\mathbf{y}\|_2}{\|\mathbf{g}\|_2} \mathbf{g},$$

where \mathbf{y} contains noiseless samples from f , $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, I)$, and $\sigma \in \mathbb{R}^+$ is used to control the Signal to Noise Ratio (SNR) defined herein by

$$\text{SNR}_{\text{db}} := 10 \log_{10} \left(\frac{\|\mathbf{y}\|_2^2}{\|\mathbf{g}'\|_2^2} \right) = -10 \log_{10}(\sigma^2).$$

In the following subsections, we consider different types of basis functions. First, in Section 5.1.1, we use mixed bases in up to 100 spatial dimensions, which consist of Fourier, Chebyshev, and Legendre bases. Afterwards, we use bases which only consist of Fourier type in Section 5.1.2, Chebyshev type in Section 5.1.3, and Legendre type in Section 5.1.4.

5.1.1 Mixed bases

First, we consider basis functions T_{n_j} of mixed type: T_{n_0} , $T_{n_{\lfloor d/2 \rfloor}}$, and $T_{n_{d-2}}$ are of Chebyshev type; T_{n_1} , $T_{n_{\lfloor d/2 \rfloor - 1}}$, and $T_{n_{d-1}}$ are of (preconditioned) Legendre type; and the remaining $d - 6$ basis functions T_{n_j} , $j \in [d] \setminus \{0, 1, \lfloor d/2 \rfloor - 1, \lfloor d/2 \rfloor, d - 2, d - 1\}$ are of

Fourier type. Preconditioned Legendre type means that instead of using standard Legendre polynomials $L_n(x) = (2n-1)/n x L_{n-1}(x) - (n-1)/n L_{n-2}(x)$, $L_1 := x$, $L_0 := 1$, with BOS constant $K = \sqrt{2n+1}$, we apply the preconditioning method from [37], i.e., we use the preconditioned Legendre polynomials $Q_n(x) := \sqrt{\pi/2} (1-x^2)^{1/4} L_n(x)$ with BOS constant $K = \sqrt{3}$ and choose the sampling nodes randomly with respect to the Chebyshev measure for the basis functions T_{n_1} , $T_{n_{\lfloor d/2 \rfloor - 1}}$, and $T_{n_{d-1}}$. Consequently, the overall BOS constant is $\sqrt{2}^3 \cdot \sqrt{3}^3 \cdot 1^{d-6} = \sqrt{6}^3$ independent of the spatial dimension $d = D$. For the entry identification and pairing steps, we set the parameter $m_2 = \#\mathbf{z}_{j,k} := 4s$ for different sparsities $s = |\mathcal{S}|$. The parameter $m_1 = \#\mathbf{w}_{j,\ell}$ is chosen as cs , where the constant $c \geq 1$ does not depend on the sparsity s , which is distinctly smaller than the theoretical results of $m_1 \sim s^2$ in Theorem 6 and Lemma 15.

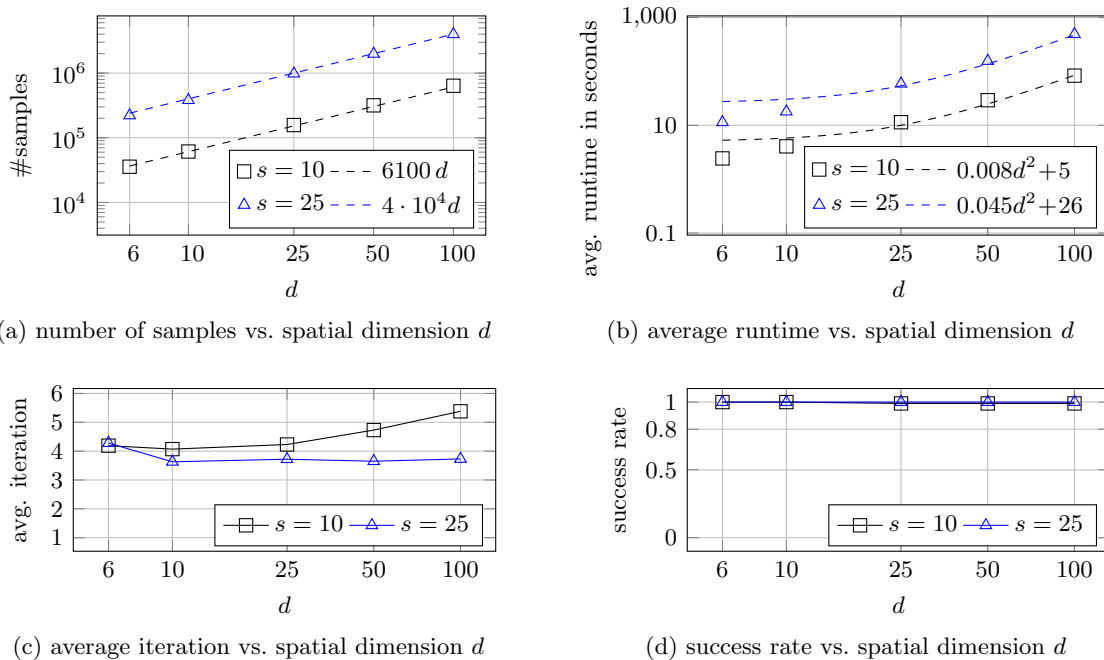


Figure 1: Number of samples, runtime, number of iterations, success rates vs. spatial dimension $d = D \in \{6, 10, 25, 50, 100\}$ for mixed bases (3 Chebyshev, 3 Legendre, $d - 6$ Fourier), $N = 200$, sparsity $s \in \{10, 25\}$, $\text{SNR}_{\text{db}} = 10$, $m_1 = 8s$, $m_2 = 4s$.

In Figure 1, we visualize the obtained results in dependence of the spatial dimensions $d = D \in \{6, 10, 25, 50, 100\}$ for sparsity $s \in \{10, 25\}$ and signal to noise ratio $\text{SNR}_{\text{db}} = 10$. In Figure 1a, we plot the number of samples with respect to the spatial dimension d . We observe that the number of samples grows nearly linearly in d . Additionally, we plot

the average runtime of the 100 test runs with respect to d in Figure 1b, and we observe that it grows approximately like $\sim d^2$. When having a look at the average number of iterations in Figure 1c, we observe that 4.1 to 5.4 iterations were required for sparsity $s = 10$ and around 4 iterations for sparsity $s = 25$. For the considered test setting, the observed success rate was 100% for sparsity $s = 25$ and at least 99% for sparsity $s = 10$, cf. Figure 1d.

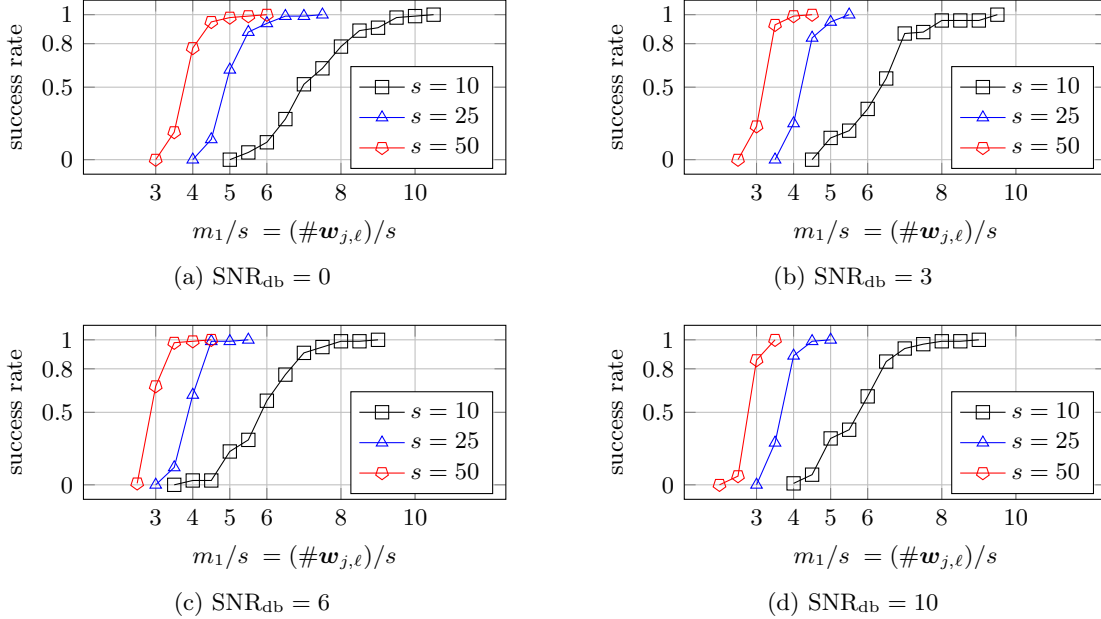


Figure 2: Success rate vs. m_1/s for mixed bases, $d = D = 50$, $N = 200$, sparsity $s \in \{10, 25, 50\}$, $m_2 = \#\mathbf{z}_{j,k} = 4s$, and $\text{SNR}_{\text{db}} \in \{0, 3, 6, 10\}$.

For different choices of the parameter $m_1 = \#\mathbf{w}_{j,\ell} \in \{2.5s, 3s, 3.5s, \dots, 10s\}$, we investigate the success rate for spatial dimension $d = D = 50$ and sparsities $s \in \{10, 25, 50\}$ in Figure 2, where we set the signal to noise ratio SNR_{db} to 0, 3, 6, and 10 in Figure 2a, 2b, 2c, and 2d, respectively. We observe that the success rates increase for growing parameter m_1 . Moreover, the transition between 0% success rate and 99%–100% success rate occurs relatively fast. Additionally, the value m_1/s where the success rate reaches 99% seems to decrease for increasing sparsity s and for increasing signal to noise ratio.

In Figure 3, we plot the used number of samples and average runtime as a function of the sparsity $s \in \{10, 25, 50\}$ for spatial dimension $d = D = 50$ and for each signal to noise ratio $\text{SNR}_{\text{db}} \in \{0, 3, 6, 10\}$. We observe that the plots only differ slightly for the different signal to noise ratios $\text{SNR}_{\text{db}} \in \{3, 6, 10\}$, i.e. the numbers of samples and runtimes seem




Figure 3: Number of samples and average runtime vs. sparsity s for mixed bases, $\geq 99\%$ success rate, $d = D = 50$, $N = 200$.

to depend only mildly on the signal to noise ratios for $\geq 99\%$ success rate. In the case SNR_{db} = 0, i.e., when the energy of the signal and of the noise match, the runtimes for sparsities $s \in \{10, 25\}$ are similar to the ones of SNR_{db} $\in \{3, 6, 10\}$ and approximately double for $s = 50$.

Additionally, we repeat the tests of Figure 2 for spatial dimension $d = D = 100$ and sparsities $s \in \{10, 25\}$, and we visualize the corresponding results in Figure 4. We obtain results analogously to the previous ones.

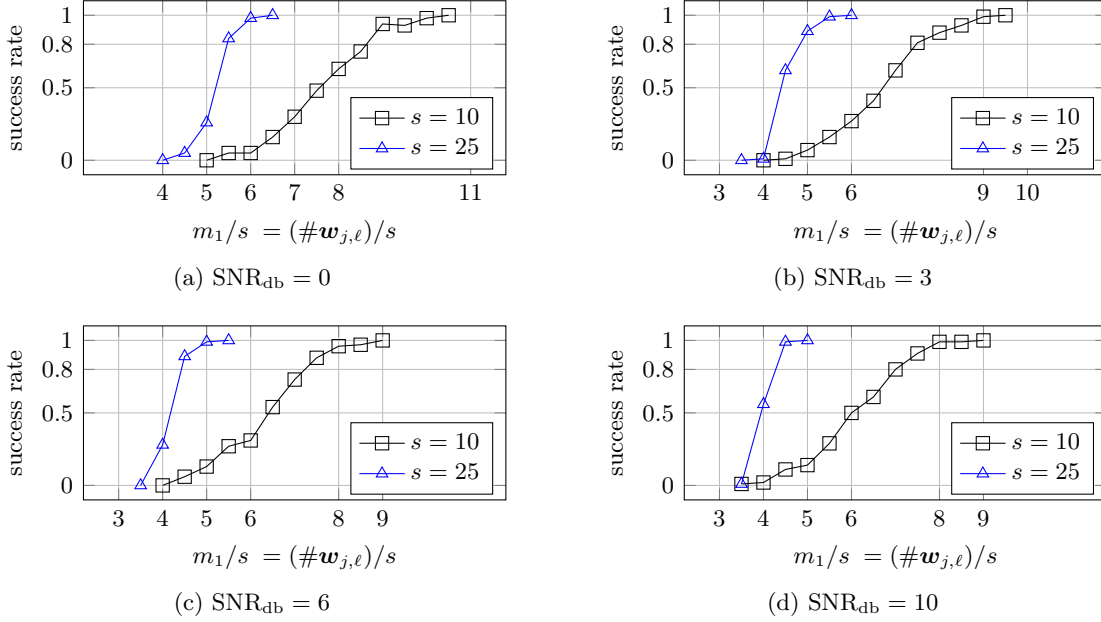



Figure 4: Success rate vs. m_1/s for mixed bases, $d = D = 100$, $N = 200$, sparsity $s \in \{10, 25\}$, $m_2 = \#\mathbf{z}_{j,k} = 4s$, and $\text{SNR}_{\text{db}} \in \{0, 3, 6, 10\}$.


5.1.2 Fourier bases

As for the case of mixed bases in Section 5.1.1, we now perform the numerical tests for tensor products of Fourier bases and show the results in Figure 5. Here, the overall BOS constant K is 1 independent of the spatial dimension d . Due to the smaller BOS constant, we can reduce the parameters m_1 to $5s$ and m_2 to s while still obtaining a success rate of 100%, cf. Figure 5d. As in Section 5.1.1, the numbers of samples in Figure 5a grow nearly linearly in d and the average runtimes in Figure 5b approximately like $\sim d^2$. The average number of iterations in Figure 5c is between 3 and 4.


In Figure 6, we depict the success rate as a function of $m_1/s \in \{1, 1.5, 2, 2.5, 3\}$ for sparsities $s \in \{10, 25, 50, 100\}$ and signal to noise ratio $\text{SNR}_{\text{db}} = 10$ in spatial dimensions $d \in \{6, 8, 10, 12\}$. We observe a very small dependence on the spatial dimension d . For $m_1 = 3s$, the success rate is 100% in each considered case. Furthermore, there is a rapid transition between full and zero success rate, i.e., the success rate is 0% for $m_1 = s$ each time.




(a) number of samples vs. spatial dimension d



(b) average runtime vs. spatial dimension d



(c) average iteration vs. spatial dimension d



(d) success rate vs. spatial dimension d

Figure 5: Number of samples, runtime, number of iterations, success rate vs. spatial dimension $d = D \in \{6, 10, 25, 50, 100\}$ for Fourier bases, $N = 200$, sparsity $s \in \{10, 25\}$, $\text{SNR}_{\text{db}} = 10$, $m_1 = 5s$, $m_2 = s$.


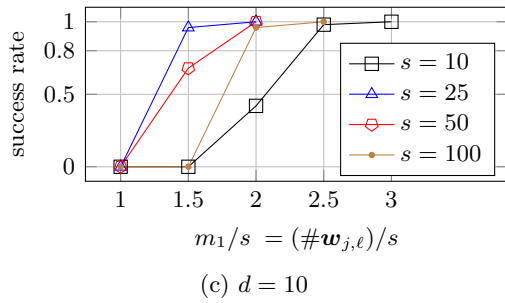




Figure 6: Success rate vs. m_1/s for Fourier bases, spatial dimension $d = D \in \{6, 8, 10, 12\}$, $N = 200$, sparsity $s \in \{10, 25, 50, 100\}$, $\text{SNR}_{\text{db}} = 10$, $m_2 = s$.

5.1.3 Chebyshev bases

Next, we consider the tensor products of Chebyshev basis functions. Here we expect larger numbers of samples and runtimes compared to the Fourier case in Section 5.1.2 due to the BOS constant $K = \sqrt{2}^d$ for Chebyshev and $K = 1$ for Fourier. In particular, for fixed sparsity s and fixed success rate, the numbers of samples and runtimes might grow for increasing spatial dimension d .

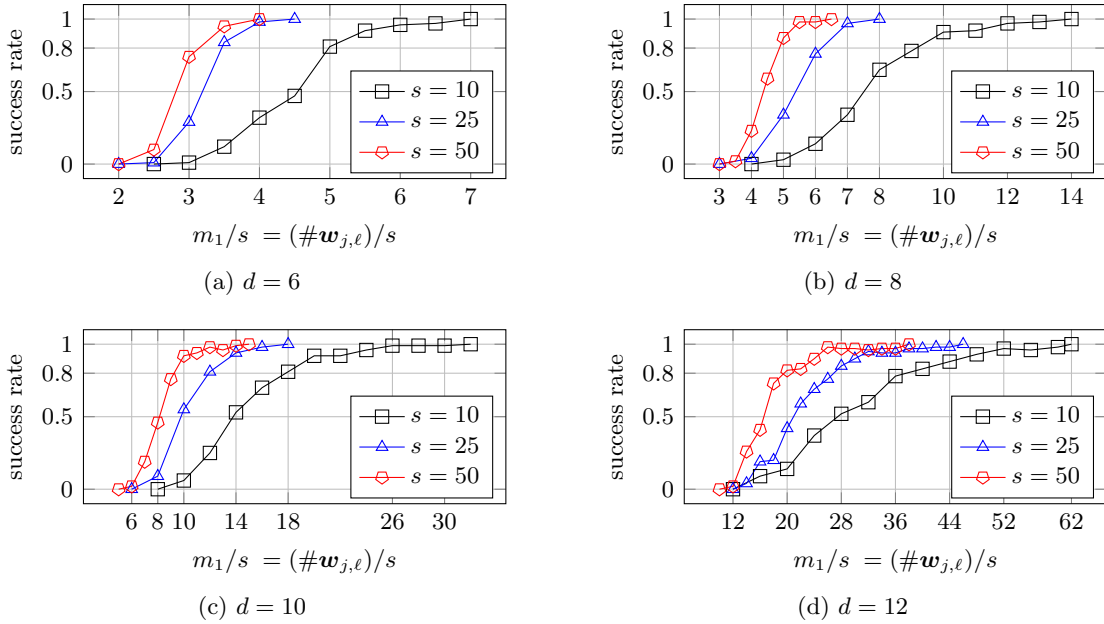


Figure 7: Success rate vs. m_1/s for Chebyshev bases, spatial dimension $d = D \in \{6, 8, 10, 12\}$, $N = 200$, sparsity $s \in \{10, 25, 50\}$, $\text{SNR}_{\text{db}} = 10$, $m_2 = 4s$.

In Figure 7, we depict the success rate as a function of m_1/s for sparsities $s \in \{10, 25, 50\}$ and signal to noise ratio $\text{SNR}_{\text{db}} = 10$ in spatial dimensions $d \in \{6, 8, 10, 12\}$. As predicted, we observe that we have to increase $m_1 = \#\mathbf{w}_{j,\ell}$ distinctly for growing spatial dimension d and fixed sparsity s . For instance, for $s = 25$, we observe a 98% success rate for $m_1 = 4s$ and $d = 6$, but obtain a success rate of only 4% for $d = 8$. For $d = 12$, we had to choose $m_1 = 42s$ to achieve a success rate of 98%.

In Figure 8, we investigate the dependence of the spatial dimension d on the number of samples and average runtime for the case of $\geq 99\%$ success rate in more detail. For our test cases, we observe that the numbers of samples grow approximately like between $\sim 2^{0.69d}$ and $\sim 2^{0.75d}$ as well as the runtimes approximately like between $\sim 2^{0.58d}$ and

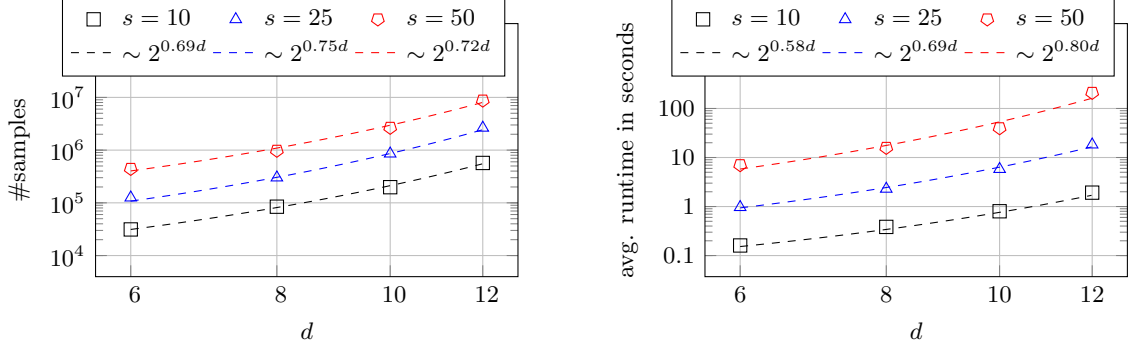


Figure 8: Number of samples and average runtime vs. spatial dimension $d \in \{6, 8, 10, 12\}$ for Chebyshev bases, $N = 200$, sparsity $s \in \{10, 25, 50\}$, $\text{SNR}_{\text{db}} = 10$, $m_2 = 4s$, success rate $\geq 99\%$.

$\sim 2^{0.80d}$. In each case, this is distinctly less than the worst case upper bounds in Theorem 5 suggest.

5.1.4 Preconditioned Legendre bases

Here, we consider the tensor products of preconditioned Legendre basis functions Q_n with BOS constant $K = \sqrt{3}^d$, cf. Section 5.1.1. In Figure 9, we show the success rates as a function of m_1/s for sparsities $s \in \{10, 25, 50\}$ and signal to noise ratio $\text{SNR}_{\text{db}} = 10$ in spatial dimensions $d \in \{6, 8, 10, 12\}$. As in the case of Chebyshev bases, we observe that we have to increase m_1 distinctly for growing spatial dimension d and fixed sparsity s . For instance, for $s = 25$, we observe a 97% success rate for $m_1 = 4s$ and $d = 6$ as well as 100% for $m_1 = 4.5s$ and $d = 6$, but obtained a success rate of only 2% for $m_1 = 4s$ and $d = 8$. Moreover, we had to choose $m_1 = 18s$ to have a success rate of 100% for $d = 10$ and $m_1 = 56s$ for $d = 12$.

When comparing the obtained results with the ones for the Chebyshev bases, we do not numerically observe the higher BOS constant $K = \sqrt{3}^d$ here for $d = 6, 8, 10$. The plots in Figure 9 look very similar to the ones in Figure 7. For $d = 12$, the values of m_1 where a success rate of $\geq 99\%$ is reached are slightly larger than the ones in the Chebyshev case.

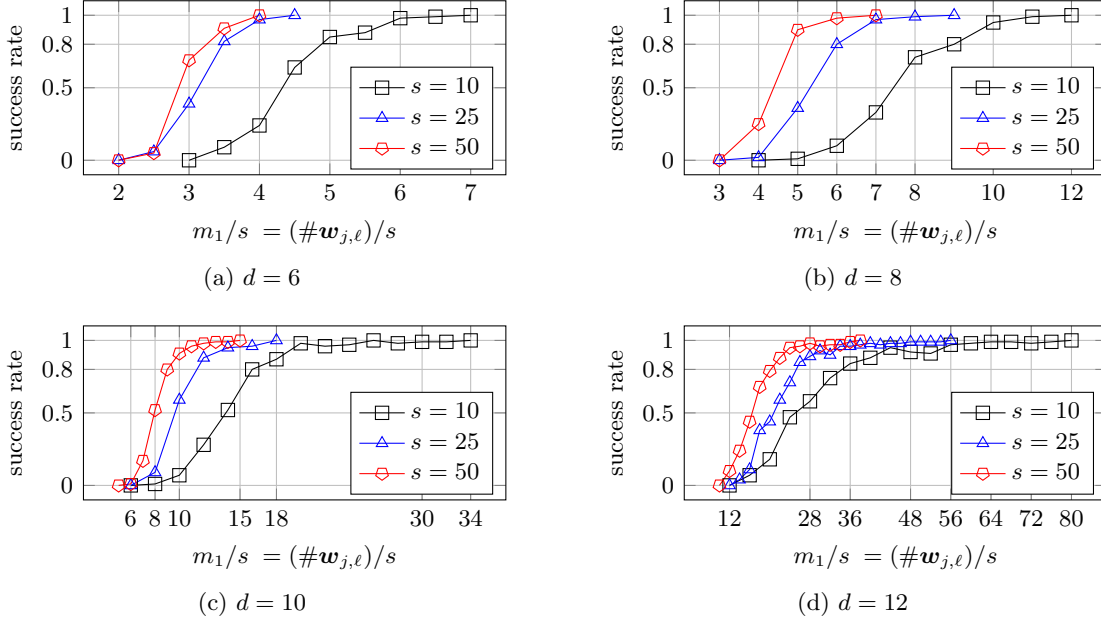


Figure 9: Success rate vs. m_1/s for preconditioned Legendre bases, spatial dimension $d = D \in \{6, 8, 10, 12\}$, $N = 200$, sparsity $s \in \{10, 25, 50\}$, $\text{SNR}_{\text{db}} = 10$, $m_2 = 4s$.

5.2 Approximately sparse case

After considering exactly sparse test functions f in Section 5.1, we continue with examples for the approximately sparse case, i.e., our test functions under consideration will have infinitely many non-zero basis coefficients c_n .

5.2.1 Fourier type with $D = 10$

We use the 10-variate periodic test function $f: \mathbb{T}^{10} \rightarrow \mathbb{R}$,

$$f(\boldsymbol{\xi}) := \prod_{t \in \{0,2,7\}} N_2(\xi_t) + \prod_{t \in \{1,4,5,9\}} N_4(\xi_t) + \prod_{t \in \{3,6,8\}} N_6(\xi_t), \quad (5.2)$$

from [34, Section 3.3] and [27, Section 5.3] with infinitely many non-zero Fourier coefficients c_n , where $\mathbb{T} \simeq [0, 1)$ is the torus and $N_m: \mathbb{T} \rightarrow \mathbb{R}$ is the B-Spline of order $m \in \mathbb{N}$,

$$N_m(x) := C_m \sum_{n \in \mathbb{Z}} \text{sinc}\left(\frac{\pi}{m}n\right)^m (-1)^n e^{2\pi i n x},$$

with a constant $C_m > 0$ such that $\|N_m\|_{L^2(\mathbb{T})} = 1$. We remark that each B-Spline N_m of order $m \in \mathbb{N}$ is a piece-wise polynomial of degree $m - 1$. We approximate the function f by multivariate trigonometric polynomials a using Algorithm 1. The obtained basis index sets $\tilde{\Omega}$ should “consist of” the union of three lower dimensional manifolds, a three-dimensional hyperbolic cross in the dimensions 1, 3, 8; a four-dimensional hyperbolic cross in the dimensions 2, 5, 6, 10; and a three-dimensional hyperbolic cross in the dimensions 4, 7, 9. All tests are performed 10 times and the relative $L^2(\mathbb{T}^{10})$ approximation error

$$\frac{\|f - a\|_{L^2(\mathbb{T}^{10})}}{\|f\|_{L^2(\mathbb{T}^{10})}} = \frac{\sqrt{\|f\|_{L^2(\mathbb{T}^{10})}^2 - \sum_{\mathbf{n} \in \tilde{\Omega}} |\hat{f}_{\mathbf{n}}|^2 + \sum_{\mathbf{n} \in \tilde{\Omega}} |a_{\mathbf{n}} - \hat{f}_{\mathbf{n}}|^2}}{\|f\|_{L^2(\mathbb{T}^{10})}}$$

is computed each time, where the approximant $a := \sum_{\mathbf{n} \in \tilde{\Omega}} a_{\mathbf{n}} e^{2\pi i \mathbf{n} \cdot \circ}$.

We set the parameters $N = 64$, $d = D = 10$, $m_2 = s$, and we always use $m_{\text{CE}} := 50s$ samples for the coefficient estimation where $s = |\mathcal{S}|$. For our tests, we consider two different parameter combinations: $m_1 = 3s$ and $\kappa = 20$, as well as $m_1 = 8s$ and $\kappa = 10$. The obtained results, i.e., the numbers of samples, average runtimes, average iterations, and relative $L^2(\mathbb{T}^{10})$ errors are plotted as a function of the sparsity $s \in \{100, 200, 500, 1000\}$ in Figure 10. Due to the parameter choices for m_1 and m_2 , we observe that the numbers of samples grow quadratically for increasing sparsity s . The average runtimes grow approximately like $\sim s^3 \min\{s, N\}$ and this means $\sim s^3$ for fixed N . Moreover, the average numbers of iterations are much smaller than its imposed maximum κ in most cases. The relative $L^2(\mathbb{T}^{10})$ errors decrease for increasing sparsity s having a value of approximately 10^{-2} for sparsity $s = 1000$. Again, we emphasize the extremely high power of Algorithm 1, which is able to determine the $s = 1000$ approximately largest bases coefficients and the corresponding indices for our test function out of $|\mathcal{I}_{N,d}| = N^d = 64^{10} \approx 10^{18}$ allowed indices.

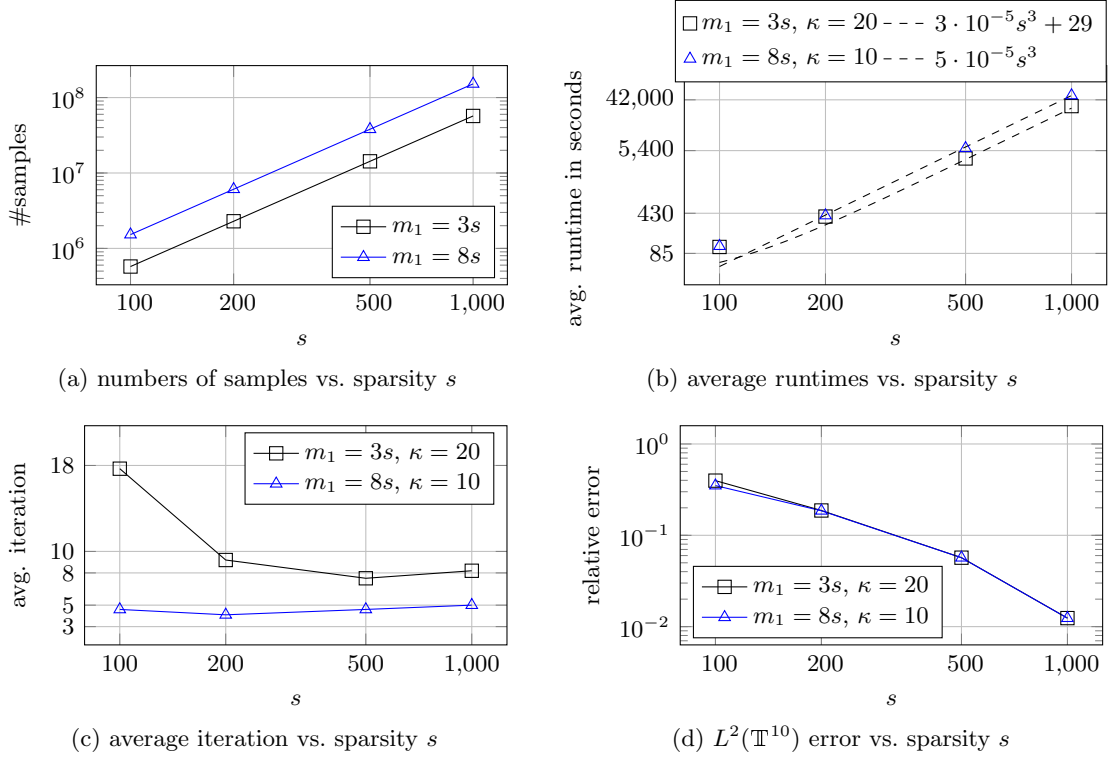


Figure 10: Number of samples, runtime, number of iterations, $L^2(\mathbb{T}^{10})$ error vs. sparsity $s \in \{100, 200, 500, 1000\}$ for Fourier basis and test function (5.2).

5.2.2 Chebyshev and Legendre type with $D = 7$




Figure 11: B-splines B_3 and B_5 considered in interval $[-1, 1]$.

Next, we apply Algorithm 1 on the 7-variate test function $f: [-1, 1]^7 \rightarrow \mathbb{R}$,


$$f(\boldsymbol{\xi}) := \prod_{t \in \{0,2,5\}} B_3(\xi_t) + \prod_{t \in \{1,3,4,6\}} B_5(\xi_t) \quad (5.3)$$

similar as in [35], where $B_3: \mathbb{R} \rightarrow \mathbb{R}$ is a shifted, scaled and dilated B-spline of order 3 and $B_5: \mathbb{R} \rightarrow \mathbb{R}$ is a shifted, scaled and dilated B-spline of order 5, see Figure 11 for


illustration. We remark that the absolute values of the Chebyshev coefficients c_n , $n \in \mathbb{N}_0$, of B_3 and B_5 decay like $\sim n^{-3}$ and $\sim n^{-5}$, respectively. The obtained basis index sets $\tilde{\Omega}$ should “consist of” the union of two lower dimensional manifolds, a three-dimensional hyperbolic cross in the dimensions 0, 2, 5; and a four-dimensional hyperbolic cross in the dimensions 1, 3, 4, 6. All tests are performed 10 times and the relative $L^2([-1, 1]^7, \mu_C)$ approximation error $\|f - a\|_{L^2([-1, 1]^7, \mu_C)} / \|f\|_{L^2([-1, 1]^7, \mu_C)}$ is computed each time, where the approximant $a := \sum_{\mathbf{n} \in \tilde{\Omega}} a_{\mathbf{n}} T_{\mathbf{n}}$, $T_{\mathbf{n}}$ is the Chebyshev product basis, and $\mu_C(\boldsymbol{\xi}) := \pi^{-D} \prod_{j \in [D]} (1 - \xi_j^2)^{-1/2}$ is the Chebyshev product measure.

We set the parameters $N = 64$, $d = D = 7$, $m_2 = 4s$, and we always use $m_{CE} := 50s$ samples for the coefficient estimation where $s = |\mathcal{S}|$. We consider two different parameter combinations: $m_1 = 4s$ and $\kappa = 20$, as well as $m_1 = 8s$ and $\kappa = 10$. The obtained results, i.e., the numbers of samples, average runtimes, average iterations, and relative $L^2([-1, 1]^7, \mu_C)$ errors are plotted as a function of the sparsity $s \in \{25, 50, 100, 200, 500\}$ in Figure 12. Due to the parameter choices for m_1 and m_2 , we observe that the numbers of samples grow quadratically for increasing sparsity s . The average runtimes grow approximately like $\sim s^3 \min\{s, N\}$ and this means $\sim s^3$ for fixed N . Moreover, the average numbers of iterations are well below its imposed maximum κ for $m_1 = 4s$ and $\kappa = 20$ as well as close to κ for $m_1 = 8s$ and $\kappa = 10$. The relative $L^2([-1, 1]^7, \mu_C)$ errors decrease for increasing sparsity s having a value of approximately $2.3 \cdot 10^{-4}$ for sparsity $s = 500$. We emphasize that Algorithm 1 is able to easily determine the $s = 500$ approximately largest basis coefficients and the corresponding basis indices for our test function out of $|\mathcal{I}_{N,d}| = N^d = 64^7 \approx 4.4 \cdot 10^{12}$ allowed indices.


In addition, we use the preconditioned Legendre polynomials Q_n from Section 5.1.1 as basis functions, i.e. $T_{\mathbf{n}}$ is now the Legendre product basis in the approximant $a := \sum_{\mathbf{n} \in \tilde{\Omega}} a_{\mathbf{n}} T_{\mathbf{n}}$. Besides that, we keep all parameters identical but determine the relative $L^2([-1, 1]^7, \mu_L)$ approximation error $\|f - a\|_{L^2([-1, 1]^7, \mu_L)} / \|f\|_{L^2([-1, 1]^7, \mu_L)}$, which corresponds to the Legendre basis and uses the probability measure $\mu_L \equiv 2^{-D}$. The results are shown in Figure 13. Here, we observe that the numbers of iterations are higher by up to $\approx 50\%$ compared to the Chebyshev case in Figure 12, and that they reach the imposed maximum of $\kappa := 20$ for $m_1 = 4s$ and $\kappa := 10$ for $m_2 = 8s$ in several cases. Correspondingly, the runtimes are also higher by up to $\approx 50\%$. The obtained relative $L^2([-1, 1]^7, \mu_L)$ errors are similar, but we also remark that we cannot compare these errors directly to the relative $L^2([-1, 1]^7, \mu_C)$ errors of the Chebyshev basis since they are measured with respect to different probability measures, $\mu_C(\boldsymbol{\xi}) := \pi^{-D} \prod_{j \in [D]} (1 - \xi_j^2)^{-1/2}$ for Chebyshev and $\mu_L \equiv 2^{-D}$ for Legendre.




(a) number of samples vs. sparsity s



(b) average runtime vs. sparsity s




(c) average iteration vs. sparsity s




(d) relative $L^2([-1, 1]^7, \mu_C)$ error vs. sparsity s


Figure 12: Number of samples, runtime, number of iterations, relative $L^2([-1, 1]^7, \mu_C)$ error vs. sparsity $s \in \{25, 50, 100, 200, 500\}$ for Chebyshev basis and test function (5.3).




(a) number of samples vs. sparsity s



(b) average runtime vs. sparsity s



(c) average iteration vs. sparsity s



(d) relative $L^2([-1, 1]^7, \mu_L)$ error vs. sparsity s

Figure 13: Number of samples, runtime, number of iterations, relative $L^2([-1, 1]^7, \mu_L)$ error vs. sparsity $s \in \{25, 50, 100, 200, 500\}$ for Legendre basis and test function (5.3).

5.2.3 Mixed type with $D = 10$

Finally, we combine parts of the test functions from the previous two subsections. We consider the 10-variate test function $f: \tilde{\mathcal{D}} \rightarrow \mathbb{R}$, $\tilde{\mathcal{D}} := [-1, 1] \times \mathbb{T} \times [-1, 1]^3 \times \mathbb{T}^2 \times [-1, 1] \times \mathbb{T}^2$,

$$f(\boldsymbol{\xi}) := B_3(\xi_0)B_3(\xi_2)N_4(\xi_8) + B_5(\xi_3)B_5(\xi_4)N_2(\xi_1)N_2(\xi_6) + B_3(\xi_7)N_2(\xi_5)N_2(\xi_9) \quad (5.4)$$


In spatial dimensions $j = 0, 2, 3, 4, 7$, we use Chebyshev basis functions as well as Fourier basis functions in the remaining spatial dimensions $j = 1, 5, 6, 8, 9$. All tests are performed 10 times and the relative $L^2(\tilde{\mathcal{D}}, \mu_{F,C})$ approximation error $\|f - a\|_{L^2(\tilde{\mathcal{D}}, \mu_{F,C})} / \|f\|_{L^2(\tilde{\mathcal{D}}, \mu_{F,C})}$ is computed each time, where the approximant $a := \sum_{\mathbf{n} \in \tilde{\Omega}} a_{\mathbf{n}} T_{\mathbf{n}}$,

$$T_{\mathbf{n}}(\boldsymbol{\xi}) := \left(\prod_{j \in \{0,2,3,4,7\}} \cos(n_j \arccos \xi_j) \right) \left(\prod_{j \in \{1,5,6,8,9\}} e^{2\pi i n_j \xi_j} \right)$$


is the mixed product basis, and $\mu_{F,C}(\boldsymbol{\xi}) := \left(\frac{2}{\pi}\right)^5 \prod_{j \in \{0,2,3,4,7\}} (1 - \xi_j^2)^{-1/2}$ is the corresponding probability measure.

Here we set the parameters $N = 64$, $d = D = 10$, $m_2 = 4s$, and we always use $m_{CE} := 50s$ samples for the coefficient estimation where $s = |\mathcal{S}|$. We consider two different parameter combinations from the previous subsection: $m_1 = 4s$ and $\kappa = 20$, as well as $m_1 = 8s$ and $\kappa = 10$. The obtained results, i.e., the numbers of samples, average runtimes, average iterations, and relative $L^2(\tilde{\mathcal{D}}, \mu_{F,C})$ errors are plotted as a function of the sparsity $s \in \{25, 50, 100, 200\}$ in Figure 14. As before, the numbers of samples grow quadratically for increasing sparsity s . The average runtimes grow approximately like $\sim s^3 \min\{s, N\}$ and this means $\sim s^3$ for fixed N . Moreover, the average numbers of iterations are well below its imposed maximum κ . The relative $L^2(\tilde{\mathcal{D}}, \mu_{F,C})$ errors decrease for increasing sparsity s having a value of approximately $4.9 \cdot 10^{-3}$ for sparsity $s = 500$. We emphasize that Algorithm 1 is able to easily determine the $s = 500$ approximately largest basis coefficients and the corresponding indices for our test function out of $|\mathcal{I}_{N,d}| = N^d = 64^{10} \approx 10^{18}$ possible indices.


Additionally, we use the preconditioned Legendre polynomials Q_n from Section 5.1.1 as basis functions in the spatial dimensions $j = 0, 2, 3, 4, 7$ instead of the Chebyshev polynomials. Besides that, we keep all parameters identical but determine now the relative $L^2(\tilde{\mathcal{D}}, \mu_{F,L})$ approximation error with respect to the probability measure $\mu_{F,L}(\boldsymbol{\xi}) \equiv 2^{-5}$ which corresponds to the current choice of bases. The results are presented in Figure 15. As before, we observe that the numbers of iterations are higher, now by up to $\approx 100\%$ compared to using Chebyshev polynomials in Figure 15. Correspondingly, the runtimes also double in some cases. The obtained relative errors are similar, but we again remark




(a) number of samples vs. sparsity s



(b) average runtime vs. sparsity s



(c) average iteration vs. sparsity s



(d) relative $L^2(\tilde{\mathcal{D}}, \mu_{F,C})$ error vs. sparsity s

Figure 14: Number of samples, runtime, number of iterations, relative $L^2(\tilde{\mathcal{D}}, \mu_{F,C})$ error vs. sparsity $s \in \{25, 50, 100, 200\}$ for mixed Fourier+Chebyshev basis and test function (5.4).

that we cannot compare these errors directly since they are measured with respect to different probability measures.

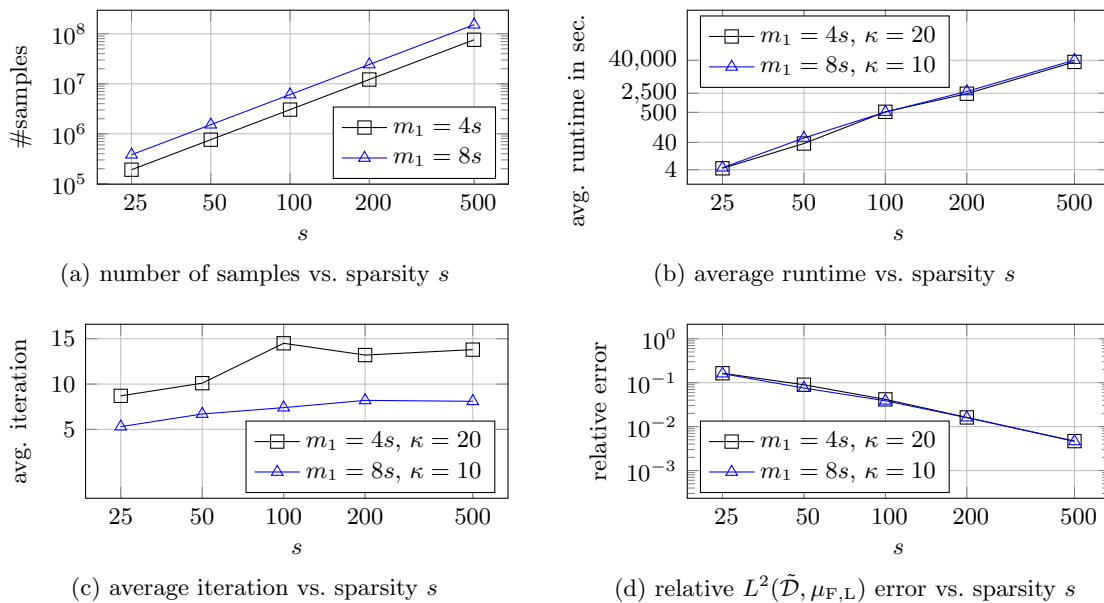


Figure 15: Number of samples, runtime, number of iterations, relative $L^2(\tilde{\mathcal{D}}, \mu_{F,L})$ error vs. sparsity $s \in \{25, 50, 100, 200\}$ for mixed Fourier+Legendre basis and test function (5.4).

Acknowledgements

Mark Iwen was supported in part by NSF DMS-1912706, and would like to dedicate this paper to his ever bright, hard working, and spirited wife Tsveta, and to the prosperity of their newborn daughter Evgenia. Evgenia – I am anxious to know you are healthy, eager to see you are happy, and already sad at the distant prospect your moving out. May you be more like your mother than like me for your own sake!

References

- [1] B. Adcock. Infinite-dimensional ℓ^1 minimization and function approximation from pointwise data. *Constructive Approximation*, 45(3):345–390, 2017.
- [2] B. Adcock, S. Brugiapaglia, and C. G. Webster. Compressed sensing approaches for polynomial approximation of high-dimensional functions. In *Compressed Sensing and its Applications*, pages 93–124. Springer International Publishing, 2017.
- [3] J. Bailey, M. A. Iwen, and C. V. Spencer. On the design of deterministic matrices for

- fast recovery of Fourier compressible functions. *SIAM Journal on Matrix Analysis and Applications*, 33(1):263–289, 2012.
- [4] S. Bittens and G. Plonka. Sparse fast DCT for vectors with one-block support. *Numerical Algorithms*, pages 1–35, 2018.
- [5] S. Bittens, R. Zhang, and M. A. Iwen. A deterministic sparse FFT for functions with structured Fourier sparsity. *Advances in Computational Mathematics*, 45:519–561, 2019.
- [6] A. Björck. *Numerical Methods for Least Squares Problems*. Society for Industrial and Applied Mathematics (SIAM), 1996.
- [7] J.-L. Bouchot, H. Rauhut, and C. Schwab. Multi-level Compressed Sensing Petrov-Galerkin discretization of high-dimensional parametric PDEs. *ArXiv e-prints*, 2017.
- [8] H.-J. Bungartz and M. Griebel. Sparse grids. *Acta Numerica*, 13:147–269, 2004.
- [9] A. Chkifa, N. Dexter, H. Tran, and C. Webster. Polynomial approximation via compressed sensing of high-dimensional functions on lower sets. *Mathematics of Computation*, 87(311):1415–1450, 2018.
- [10] B. Choi, A. Christlieb, and Y. Wang. Multi-dimensional sublinear sparse Fourier algorithm. *arXiv preprint arXiv:1606.07407*, 2016.
- [11] B. Choi, A. Christlieb, and Y. Wang. Multiscale High-Dimensional Sparse Fourier Algorithms for Noisy Data. *arXiv e-prints*, page arXiv:1907.03692, 2019.
- [12] B. Choi, M. Iwen, and F. Krahmer. Sparse harmonic transforms: A new class of sublinear-time algorithms for learning functions of many variables. *arXiv preprint arXiv:1808.04932*, 2018.
- [13] A. Cohen, W. Dahmen, and R. DeVore. Compressed sensing and best k -term approximation. *Journal of the American Mathematical Society*, 22(1):211–231, 2009.
- [14] G. Dahlquist and A. Björck. *Numerical Methods in Scientific Computing, Volume I*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, USA, 2008.
- [15] M. F. Duarte and R. G. Baraniuk. Kronecker compressive sensing. *IEEE Transactions on Image Processing*, 21(2):494–504, 2012.
- [16] D. Dũng, V. N. Temlyakov, and T. Ullrich. Hyperbolic cross approximation. *arXiv preprint arXiv:1601.03978*, 2016.
- [17] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Springer New York, 2013.

- [18] A. Gilbert, A. Gu, C. Re, A. Rudra, and M. Wootters. Sparse Recovery for Orthogonal Polynomial Transforms. *arXiv preprint arXiv:1907.08362*, 2019.
- [19] A. Gilbert, M. Iwen, and M. Strauss. Empirical evaluation of a sub-linear time sparse DFT algorithm. *Communications in Mathematical Sciences*, 5(4):981–998, 2007.
- [20] A. C. Gilbert, P. Indyk, M. A. Iwen, and L. Schmidt. Recent developments in the sparse Fourier transform: A compressed Fourier transform for big data. *IEEE Signal Processing Magazine*, 31(5):91–100, 2014.
- [21] A. C. Gilbert, S. Muthukrishnan, and M. Strauss. Improved time bounds for near-optimal sparse Fourier representations. In *Proceedings of SPIE*, volume 5914, page 59141A, 2005.
- [22] H. Hassanieh, P. Indyk, D. Katabi, and E. Price. Simple and practical algorithm for sparse Fourier transform. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1183–1194. Society for Industrial and Applied Mathematics (SIAM), 2012.
- [23] X. Hu, M. Iwen, and H. Kim. Rapidly computing sparse Legendre expansions via sparse Fourier transforms. *Numerical Algorithms*, 74(4):1029–1059, 2017.
- [24] M. A. Iwen. A deterministic sub-linear time sparse Fourier algorithm via non-adaptive compressed sensing methods. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 20–29. Society for Industrial and Applied Mathematics (SIAM), 2008.
- [25] M. A. Iwen. Combinatorial sublinear-time Fourier algorithms. *Foundations of Computational Mathematics*, 10(3):303–338, 2010.
- [26] M. A. Iwen. Improved approximation guarantees for sublinear-time Fourier algorithms. *Applied and Computational Harmonic Analysis*, 34(1):57–82, 2013.
- [27] L. Kämmerer, D. Potts, and T. Volkmer. High-dimensional sparse FFT based on sampling along multiple rank-1 lattices. *arXiv preprint arXiv:1711.05152*, 2017.
- [28] M. Kapralov. Sparse Fourier transform in any constant dimension with nearly-optimal sample complexity in sublinear time. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 264–277. ACM Press, 2016.
- [29] M. Kapralov, A. Velingker, and A. Zandieh. Dimension-independent Sparse Fourier Transform. *arXiv e-prints*, page arXiv:1902.10633, 2019.
- [30] Y. Mansour. Randomized interpolation and approximation of sparse polynomials. In *Proceedings of the 19th International Colloquium on Automata, Languages and Programming*, ICALP '92, pages 261–272, London, UK, 1992. Springer-Verlag.

- [31] S. Merhi, R. Zhang, M. A. Iwen, and A. Christlieb. A new class of fully discrete sparse Fourier transforms: Faster stable implementations with guarantees. *Journal of Fourier Analysis and Applications*, 25(3):751–784, 2019.
- [32] L. Morotti. Explicit universal sampling sets in finite vector spaces. *Applied and Computational Harmonic Analysis*, 43(2):354–369, 2017.
- [33] D. Needell and J. A. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, 26(3):301–321, 2009.
- [34] D. Potts and T. Volkmer. Sparse high-dimensional FFT based on rank-1 lattice sampling. *Applied and Computational Harmonic Analysis*, 41(3):713–748, 2016.
- [35] D. Potts and T. Volkmer. Multivariate sparse FFT based on rank-1 Chebyshev lattice sampling. In *2017 International Conference on Sampling Theory and Applications (SampTA)*, pages 504–508. IEEE, 2017.
- [36] H. Rauhut. Random sampling of sparse trigonometric polynomials. *Applied and Computational Harmonic Analysis*, 22(1):16–42, 2007.
- [37] H. Rauhut and R. Ward. Sparse Legendre expansions via ℓ_1 -minimization. *Journal of Approximation Theory*, 164(5):517–533, 2012.
- [38] C. Schwab and R. A. Todor. Karhunen–Loève approximation of random fields by generalized fast multipole methods. *Journal of Computational Physics*, 217(1):100–122, 2006.
- [39] B. Segal and M. A. Iwen. Improved sparse Fourier approximation results: Faster implementations and stronger guarantees. *Numerical Algorithms*, 63(2):239–263, 2013.
- [40] J. Shen and L.-L. Wang. Sparse spectral approximations of high-dimensional problems based on hyperbolic cross. *SIAM Journal on Numerical Analysis*, 48(3):1087–1109, 2010.