

Encoder Blind Combinatorial Compressed Sensing

Michael Murray, Jared Tanner

Abstract—In its most elementary form, compressed sensing studies the design of decoding algorithms to recover a sufficiently sparse vector or code from a lower dimensional linear measurement vector. Typically it is assumed that the decoder has access to the encoder matrix, which in the combinatorial case is sparse and binary. In this paper we consider the problem of designing a decoder to recover a set of sparse codes from their linear measurements alone, that is without access to encoder matrix. To this end we study the matrix factorisation task of recovering both the encoder and sparse coding matrices from the associated linear measurement matrix. The contribution of this paper is a computationally efficient decoding algorithm, Decoder-Expander Based Factorisation, with strong performance guarantees. Under mild assumptions on the sparse coding matrix and by deploying a novel random encoder matrix, we prove that Decoder-Expander Based Factorisation recovers both the encoder and sparse coding matrix at the optimal measurement rate with high probability and from a near optimal number of measurement vectors. In addition, our experiments demonstrate the efficacy and computational efficiency of our algorithm in practice. Beyond compressed sensing, our results may be of interest for researchers working in areas as diverse as linear sketching, coding theory, matrix compression and dictionary learning.

Index Terms—Combinatorial Compressed Sensing, Matrix Factorisation, Expander Graphs, Linear Sketching, Dictionary Recovery

I. INTRODUCTION

COMPRESSED sensing (CS) [1], [2], [3], [4], [5], [6], in its simplest setting, studies the design of sensing or encoder matrices $\mathbf{A} \in \mathbb{R}^{m \times n}$ and decoder algorithms in order to recover a sparse vector $\mathbf{x} \in \mathbb{R}^n$ from $m \ll n$ linear measurements, $\mathbf{y} = \mathbf{A}\mathbf{x} \in \mathbb{R}^m$. A primary motivation for compressed sensing is the recovery and reconstruction of sparse signals using a number of measurements well below that suggested by the Nyquist-Shannon sampling theorem. In particular, instead of sampling a signal at a high rate and then performing compression, which can be wasteful, one can directly sample or sense the data in a compressed form, thereby reducing the number of measurements required for reconstruction. We refer the reader to [7], [8] for an in-depth overview and survey of the topic.

Typically, when analysing and designing compressed sensing algorithms, it is assumed that the decoder has access to both the measurement vector and encoder matrix. Here we

instead study the design of an encoder matrix and decoder algorithm to perform compressed sensing at the optimal measurement rate under the assumption that the decoder does not have access to the encoder matrix. In this context, we refer to the decoder algorithm as being ‘blind’ to the specific encoder matrix used to generate the linear measurements available to it. We therefore refer to this problem as encoder blind compressed sensing. At first glance, it does not seem at all likely that this problem can be efficiently solved. However, under a few mild additional assumptions, and by granting the decoder access to multiple measurement vectors, we will show that both the encoder and sparse coding matrices can be recovered with high probability exactly and efficiently from remarkably few samples.

The surprising nature of this result may be of interest not only in compressed sensing but also in the context of matrix factorisation problems involving a sparse factor matrix. We note that encoder blind CS is reminiscent of both blind compressed sensing [9] and dictionary learning [7], [10], [11], [12], [13]. However, there are important distinctions, notably in terms of the restrictions placed on the matrix factors. First, in [9] it is assumed that the decoder has knowledge of the encoder matrix, but does not have access to the sparsifying basis of the data in question, the sole objective being to recover the data directly from a set of linear measurements. By contrast, our goal here is to simultaneously recover the encoder matrix and sparse representations of data from a set of linear measurements. Furthermore, while in [9] the encoder matrix is drawn from a dense, Gaussian distribution in order to ensure uniqueness, here we assume the encoder is drawn, as in combinatorial compressed sensing, from a sparse, binary distribution. Insight into the reasons behind this choice will be given shortly. Finally, compared with the problem we study, the goal of dictionary learning is to compute sparse representations, or codes, of data, rather than recovering sparse codes from a set of linear measurements. To be clear, while both amount to performing a factorisation of the form $\mathbf{Y} \approx \mathbf{A}\mathbf{X}$, where $\mathbf{Y} \in \mathbb{R}^{m \times N}$ is the data or compressed data matrix, $\mathbf{A} \in \mathbb{R}^{m \times n}$ is a dictionary or encoder matrix, and $\mathbf{X} \in \mathbb{R}^{n \times N}$ is a sparse matrix, in our case we know by construction that such a factorisation exists, furthermore, we also have freedom to design \mathbf{A} in order to achieve better compression and facilitate recovery. While both problems are challenging, dictionary learning is significantly more so as there are no guarantees that such a factorisation exists and it is typically not clear in advance what properties the dictionary needs in order to sparsely represent a given data sample.

Our insight into how the problem of encoder blind CS might be solved arises from observations on combinatorial compressed sensing [14]. Consider an encoder matrix which is sparse, binary, has a fixed number $d \ll m$ ones per column

Manuscript received July 2nd, 2021; revised XXX. This work is supported by the Alan Turing Institute under the EPSRC grant EP/N510129/1 and the Ana Leaf Foundation. This paper was presented in part at SPARS 2019 and SIAM LA21.

Michael Murray (email: mmurray@math.ucla.edu) and Jared Tanner (email: tanner@maths.ox.ac.uk) conducted this work while both were at the Mathematical Institute, University of Oxford (Address: Radcliffe Observatory, Andrew Wiles Building, Woodstock Rd, Oxford OX2 6GG) and The Alan Turing Institute (Address: 2QR, 96 Euston Rd, London NW1 2DB).

and is the adjacency matrix of (k, ϵ, d) -expander graph: such graphs satisfy the unique neighbour property, defined and discussed in detail in Section II-A, which bounds the overlap in support between any subset of k columns of \mathbf{A} . With $\mathbf{y} = \mathbf{A}\mathbf{x}$, this property implies that there exists a threshold on the number of times a sum of two or more entries of \mathbf{x} can appear in \mathbf{y} . As a result, and if sums over different sets of nonzeros in \mathbf{x} result in different values, then nonzero entries in \mathbf{x} can trivially be recovered by identifying the entries in \mathbf{y} whose frequency of appearance is above this threshold. Furthermore, the location of a nonzero of \mathbf{x} can then be recovered by finding the column of \mathbf{A} whose support maximally overlaps the set row locations of \mathbf{y} in which the nonzero appears. Once a nonzero value of \mathbf{x} and its location have been recovered, then its contribution can be removed from the measurement vector, thereby allowing for the identification of further nonzeros and their locations. As long as the overlap in support in each subset of k columns of \mathbf{A} is sufficiently small, then this iterative, peeling approach will fully decode \mathbf{y} , we refer the reader to [15] for further details.

Consider now the same setting but assume that the decoder is not granted access to \mathbf{A} : can the same principles and techniques be applied? Observe that although it may be possible to recover certain nonzeros of \mathbf{x} in the same manner as before, without access to \mathbf{A} it is not possible to recover the locations of these nonzeros, or remove their contribution from \mathbf{y} . Without the ability to peel away the contributions of the recovered nonzeros from \mathbf{y} , it may not be possible to recover all the nonzeros in \mathbf{x} . To apply the same iterative, peeling technique as before, it therefore seems necessary to also recover the encoder matrix. Ignoring for now the recovery of the locations, then identifying a nonzero value of \mathbf{x} in \mathbf{y} , which does not require access to \mathbf{A} , facilitates the recovery of at least part of the support of a column of \mathbf{A} . If we have access to $N \geq 1$ measurements vectors, stored in the columns of a matrix $\mathbf{Y} \in \mathbb{R}^{m \times N}$, and can identify certain nonzeros in each whose locations correspond to the support of the same column of \mathbf{A} , then by unioning these locations it seems plausible that a column of \mathbf{A} could be recovered in its entirety. We therefore study the matrix factorisation problem of recovering both the encoder $\mathbf{A} \in \mathbb{R}^{m \times n}$ and the N k -sparse column codes $\mathbf{X} \in \mathbb{R}^{n \times N}$ from $\mathbf{Y} = \mathbf{A}\mathbf{X}$. We remark that to recover all of the columns of \mathbf{A} it is necessary that $N \geq n/k$.

The problem we study here is daunting, indeed, before even considering computational feasibility and efficiency, the conditions ensuring uniqueness of such a factorisation are far from clear. In particular, for any permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ ¹, as $\mathbf{Y} = \mathbf{A}\mathbf{X} = \mathbf{X}\mathbf{P}^T\mathbf{P}\mathbf{A}$ then, regardless of the choice of \mathbf{A} , without any information in addition to \mathbf{Y} the decoder algorithm can at best only hope to recover the sparse codes up to row permutation. We therefore allow the decoder and encoder to agree on a unique ordering of the columns of \mathbf{A} , discussed in detail in Section II-C, in advance. Note that under this assumption recovery up to permutation is sufficient for exact recovery. Our key contribution in this paper

is the design of an encoder matrix and an efficient decoding algorithm, Decoder-Expander Based Factorisation (D-EBF), which, under certain assumptions on the sparse coding matrix, can solve the aforementioned matrix factorisation problem at the optimal measurement rate $m = \mathcal{O}(k \log(n/k))$ with high probability from just $N = \Omega(\frac{n}{k} \log^2(n))$ measurement vectors. Note this sample complexity differs from the tight lower bound $N \geq n/k$ only in terms of additional multiplicative logarithmic factors, which arise due to solving a variant of the coupon-collector and dixie-cup problems [16], [17], [18]. The application of such tools in the context of compressed sensing, dictionary learning and related fields is to the best of our knowledge novel. In addition to our theoretical analysis, in Section V we demonstrate the computational efficiency and efficacy of D-EBF in practice on synthetic, mid-sized problems.

II. PROBLEM DEFINITION AND SUMMARY OF RESULTS

The purpose of this section is fourfold: first to provide the necessary background material, second to define our data model and assumptions, third to summarise the key results and findings of this paper under these assumptions. To this end, in Section II-A we provide a background on combinatorial compressed sensing and expander graphs. Then, and building on prior work [19], in Section II-B we introduce and motivate the Permuted Striped Block (PSB) model. This model consists of a set of assumptions placed on the encoder sparse coding matrices. In Section II-C, we present and discuss the key theoretical result of this paper, Theorem II.3. This theorem, which we prove in Section IV, provides performance guarantees for the Decoder-Expander Based Factorisation (D-EBF) algorithm, defined and described in detail in Section III. Finally, in Section II-D we discuss related work on dictionary recovery.

A. Compressed sensing and expander graphs

In compressed sensing many of the theorems guaranteeing the recovery of the sparse code rely on the encoder matrix satisfying one or more properties, in particular the nullspace property [20], [21], a condition on the mutual coherence [21], [22] or the restricted isometry property (RIP) [23], [24], [3]. Although constructing matrices which satisfy a condition on the mutual coherence can be accomplished in a computationally efficient manner, it is NP-hard to compute both the nullspace constant (NSC) and restricted isometry constant (RIC) of a matrix [25]. On the other hand, it is recognised that stronger conditions can be proved via the nullspace property and RIP [8]. A perhaps surprising result is that a number of random matrix constructions, notably Gaussian, Bernoulli and partial Fourier, satisfy RIP with probability approaching one exponentially fast in the problem size [8]. Furthermore, not only are these random constructions popular from the perspective of deriving improved recovery guarantees, but also play a key role in applications: for instance, Gaussian matrices achieve the optimal measurement rate $m = \mathcal{O}(k \log(n/k))$ with high probability while partial Fourier matrices are the natural sensing mechanism in Tomography [26].

¹Note here that $\mathcal{P}^{n \times n}$ is used to denote the set of $n \times n$ permutation matrices.

Combinatorial compressed sensing (CCS) [14] analyses the compressed sensing problem in the setting where the encoder matrix is sparse and binary. Not only are these types of encoder matrix the relevant sensing mechanism in certain applications, e.g., the single pixel camera [27], but they also have benefits in terms of reduced computation and memory overheads compared with dense alternatives. Encoder matrices of this form are often interpreted as the adjacency matrices of an unbalanced bipartite graph. Furthermore, in [14] it was shown that when the encoder matrix is the adjacency matrix of a (k, ϵ, d) -expander graph, then the optimal measurement rate can be achieved.

Definition II.1 ((k, ϵ, d) -expander graph). Consider a left d -regular bipartite graph $G = ([n], [m], E)$, for any $\mathcal{S} \subseteq [n]$ let $\mathcal{N}(\mathcal{S}) := \{j \in [m] : \exists l \in \mathcal{S} \text{ s.t. } (l, j) \in E\}$ be the subset of nodes in $[m]$ connected to a node in \mathcal{S} . With $\epsilon \in (0, 1)$, then G is a (k, ϵ, d) -expander graph iff

$$|\mathcal{N}(\mathcal{S})| > (1 - \epsilon)d|\mathcal{S}| \quad \forall \quad \mathcal{S} \in [n]^{(\leq k)}. \quad (1)$$

Here $[n]^{(\leq k)}$ denotes the set of subsets of $[n]$ with cardinality at most k and ϵ is the expansion parameter, which plays a role analogous to that of the RIC of a dense matrix in proving recovery guarantees [14]. In this paper we define the *adjacency matrix of a (k, ϵ, d) -expander graph* $G = ([n], [m], E)$ as an $m \times n$ binary matrix \mathbf{A} , where $a_{j,l} = 1$ iff there is an edge between node $j \in [m]$ and $l \in [n]$ and is 0 otherwise². In all that follows we will use $\mathcal{E}_{k,\epsilon,d}^{m \times n} \subset \{0, 1\}^{m \times n}$ to denote the set of (k, ϵ, d) -expander graph adjacency matrices of dimension $m \times n$. The existence of optimal expanders, in the sense of achieving the optimal measurement rate $m = \mathcal{O}(k \log(n/k))$, was proved in [28], [29]. With regard to constructing such encoder matrices however, the current best deterministic constructions only achieve a measurement rate of $m = \mathcal{O}(k^{1+\gamma})$ for some constant $\gamma > 0$ [30]. As a result it is common to resort to random constructions, sampling A from a distribution so that with high probability $A \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$. A popular and natural construction in this regard is to sample A so that its columns are mutually independent and identically distributed, with the support of each column being chosen uniform at random from all possible supports of size d . Current state of the art bounds on the probability that this particular random matrix is a (k, ϵ, d) -expander satisfying the optimal measurement rate can be found in [19]. In particular, we highlight the following result.

Lemma II.1 ([19, Lemma 3.1]). Let A be an $m \times n$ random, binary matrix with mutually independent, identically distributed columns, whose supports are drawn uniform at random across all possible supports of cardinality d . For any $\epsilon \in (0, 1/2)$, suppose as $(k, m, n) \rightarrow \infty$ that $k/m \rightarrow \rho$ and $m/n \rightarrow \alpha_2$, where $\rho, \alpha_2 \in (0, 1)$ are constants. As long as $\rho < (1 - \gamma)\rho_{BT}(\alpha_2, \epsilon, d)$ for some constant $\gamma \in (0, 1)$, then the probability that $A \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ approaches one exponentially fast in n .

²We note that the adjacency matrices of graphs are often defined to describe the connectivity between all nodes in the graph, however, as we only consider bipartite graphs, in the definition adopted here the edges between nodes in the same group are not set to zero but rather are omitted entirely.

Here $\rho_{BT}(\alpha_2, \epsilon, d)$ denotes the phase transition curve, which is computed numerically, see [19, Equation 29]. Following [14], a series of iterative, greedy CCS algorithms were proposed: Sparse Matching Pursuit (SMP) [31], Sequential Sparse Matching Pursuit (SSMP) [32], Left Degree Dependent Signal Recovery (LDDSR) [33], Expander Recovery (ER) [15] and ℓ_0 -decode [34, Algorithms 1 & 2]. These algorithms are specifically designed for the CCS setting and utilise certain key properties of expander graphs, namely the unique neighbour property.

Theorem II.2 (Unique neighbour property [29, Lemma 1.1]). Suppose that $G = ([n], [m], E)$ is an unbalanced, left d -regular bipartite graph. Let $\mathcal{S} \in [n]^{(\leq k)}$ and define

$$\mathcal{N}_1(\mathcal{S}) = \{j \in \mathcal{N}(\mathcal{S}) : |\mathcal{N}(j) \cap \mathcal{S}| = 1\},$$

where $\mathcal{N}(j)$ is the subset of nodes in $[n]$ connected to the node $j \in [m]$. If G is a (k, ϵ, d) -expander then

$$|\mathcal{N}_1(\mathcal{S})| > (1 - 2\epsilon)d|\mathcal{S}| \quad \forall \quad \mathcal{S} \in [n]^{(\leq k)}. \quad (2)$$

A proof of Theorem II.2 in the notation used here is available in [34, Appendix A]. The unique neighbour property is used to prove that expander encoders satisfy RIP in the ℓ_1 norm [14]. Moreover, and critical to our purpose, if $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ then the unique neighbour property ensures that certain entries in \mathbf{x} appear repeatedly in the measurement vector $\mathbf{y} = \mathbf{A}\mathbf{x}$. In Section III-B we study sufficient conditions for identifying such entries, which, being the sum of a single nonzero entry in \mathbf{x} , we refer to as singleton values. The unique neighbour property is exploited by CCS algorithms such as LDDSR, ER and ℓ_0 -decode to guarantee that at each iteration a contraction in $\|\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}\|_0$ occurs, where here $\hat{\mathbf{x}}$ denotes the reconstruction of \mathbf{x} . We likewise leverage this fact in the design of D-EBF: in particular, the locations in which an entry of \mathbf{x} appears in \mathbf{y} also provides access to part of the support of a column of \mathbf{A} . The key idea behind the D-EBF algorithm is to combine and cluster partial supports, extracted from different columns of \mathbf{Y} , so as to recover the columns of \mathbf{A} . We define and discuss in detail the D-EBF algorithm in Section III.

B. The Permuted Striped Block (PSB) model

In this section we introduce and motivate a particular distribution over a set of measurement matrices, which we call the Permuted Striped Block (PSB) model. This distribution arises from conditions placed on both the encoder matrix, which we have freedom to design, as well as modelling assumptions placed on the sparse coding matrix.

Based on the discussion at the end of Section II-A, a strong candidate encoder is one sampled from the distribution of a random binary matrix whose column supports are mutually independent and have a fixed cardinality of size d . With k, m and n such that Lemma II-A is satisfied, then for sufficiently large problems such a matrix will with high probability be the adjacency matrix of a (k, ϵ, d) -expander graph and satisfy the unique neighbour property. However, the downside of this construction is that it does not rule out the possibility of there existing a very dense, in terms of the number of ones, row. The approach we adopt in the design of D-EBF relies on eventually

observing each one in a column of \mathbf{A} through the extraction of partial supports. As a result, if a row is sufficiently dense then the one entries in said row are unlikely to ever be identified. In fact, the proof of Lemma IV.4, see Appendix E, which plays a key role in providing the guarantees for D-EBF summarised in Theorem II.3, relies on the number of ones per row of \mathbf{A} being bounded. To this end we consider a variant on the random construction already discussed which maintains the same key properties while in addition bounding the number of ones per row. How to sample from the distribution of this random matrix is described below in Algorithm 1 below. A detailed definition and description of this algorithm is provided in Appendix A-B.

Algorithm 1 GENERATE-ENCODER(m, n, d)

- 1: Initialise A as the $m \times n$ zero matrix.
 - 2: **while** there are still zero columns in A **do**
 - 3: Sample a permutation of $[m]$ uniform at random.
 - 4: Split permutation into groups of d consecutive elements.
 - 5: Assign each group to be the support of one of the zero columns of A .
 - 6: **end while**
-

Algorithm 1 takes as inputs the dimensions of the encoder m and n , the column sparsity d , and returns a binary matrix. Note also that draws of different permutations are considered to be mutually independent. Encoder matrices generated using Algorithm 1 have a fixed number d ones per column and a maximum of $\lceil n/\lfloor m/d \rfloor \rceil$ ones per row. We emphasise that this upper bound on the number of ones per row of A will be crucial for proving our recovery results, in particular Lemma IV.4. Note also that the support of each column of A is dependant on the supports of at most $\lfloor m/d \rfloor - 1$ other columns, and that the intersection in support of these sets of dependent columns is empty by construction.

We claim here that Lemma II-A applies without modification to encoder matrices generated using Algorithm 1. This result follows in exactly the same manner as proved in [19] by observing that any pair of columns of A generated using Algorithm 1 either have independently drawn supports, as considered in the original case in [19], or, if they are dependent, then they are disjoint by construction. For brevity we do not replicate the original proof in detail here, instead referring the reader to [19] for further details.

Turning our attention now to modelling assumptions on the sparse coding matrix, we remark that multimeasurement vector compressed sensing [35], [36], [37] also studies the compressed sensing problem in the context of a set of $N > 1$ measurement vectors. Typically in this context \mathbf{X} is structured such that sparse recovery of one column aids the recovery of others. For example, in the joint sparse setting the columns of \mathbf{X} share a common support. As our goal here is to use the different measurement vectors of \mathbf{Y} to recover different parts of the encoder matrix we adopt a very different model: in particular we assume that the supports of each column are chosen uniformly at random over all possible supports of size k , and that the columns are mutually independent of one another. This distribution feels neither adversarial or

overly sympathetic to our objective, and is adopted in order to capture a sense of how well an algorithm might typically perform at the desired factorisation task. In addition to the distribution placed on the support, we also assume the the nonzero coefficients in each column of the sparse coding matrix are dissociated.

Definition II.2 (Dissociated vector, see Definition 4.32 of [38]). A vector $\mathbf{x} \in \mathbb{R}^N$ is said to be dissociated, which we will denote as $\mathbf{x} \in \chi_k^n$, iff for any pair of subsets $\mathcal{T}_1, \mathcal{T}_2 \subseteq \text{supp}(\mathbf{x})$, $\mathcal{T}_1 \neq \mathcal{T}_2$, it holds that $\sum_{j \in \mathcal{T}_1} x_j \neq \sum_{i \in \mathcal{T}_2} x_i$. For $N > 1$ $\mathbf{X} \in \chi_k^{n \times N}$ iff $\mathbf{x}_i \in \chi_k^n$ for all $i \in [N]$.

Although at first glance this condition may appear restrictive, it is fulfilled almost surely for isotropic vectors and more generally for any random vector whose nonzeros are drawn independently from a continuous distribution. This property plays a key role in the design of D-EBF as discussed in Section III-B, and is also adopted in [34] where it plays a key role in enabling sparse recovery in the context of massive problems, i.e., $n = 2^{26}$, $m/n = 10^{-3}$ and $k/n = 0.3$, in a matter of only seconds using standard computing infrastructure. We are now ready to introduce the PSB model.

Definition II.3 (PSB model). Let $d, k, m, n, N \in \mathbb{N}$ such that $k < m < n$, $k/n = \alpha_1$ and $m/n = \alpha_2 := \theta \alpha_1 \log(1/\alpha_1)$ where $\alpha_1, \alpha_2 \in (0, 1)$ and $\theta \in \mathbb{R}_{>0}$ are constants, and suppose there exists a constant $\gamma \in (0, 1)$ such that $k/m < (1 - \gamma)\rho_{BT}(\alpha_2, 1/6, d)$. Consider now the following random matrices.

- $A := [A_1 \ A_2 \dots \ A_n]$ is a random, binary matrix of size $m \times n$ whose distribution can be sampled from using Algorithm 8. The columns of A each have exactly d ones, while the rows have at most $\lceil n/\lfloor d/m \rfloor \rceil$ ones by construction.
- $X := [X_1 \ X_2 \dots \ X_N]$ is a random, real matrix of size $n \times N$, whose columns are mutually independent, have exactly k nonzero entries, are dissociated and whose supports are chosen uniformly at random across all possible supports with cardinality k .

A random, real $m \times N$ matrix Y is sampled from the PSB model, which we will denote as $Y \sim \text{PSB}(d, k, m, n, N)$, iff $Y := AX$.

With regard to the choice of parameters, observe that, under the assumptions listed, if an algorithm recovers the sparse codes from a measurement matrix drawn from the PSB model, then it does so at the optimal measurement rate $m = \mathcal{O}(k \log(n/k))$. Second, as $k/n = \alpha_1$ then $m/n = -\theta \alpha_1 \log(\alpha_1) =: \alpha_2$ and $k/m = -1/\theta \log(\alpha_1)$ are constants. Therefore, by construction, it follows that the conditions stated in Lemma II-A are satisfied, and as a result the probability that $A \in \mathcal{E}_{k, \epsilon, d}^{m \times n}$ with $\epsilon \leq 1/6$ approaches one exponentially fast as $n \rightarrow \infty$.

In order to explain the origin of its name, observe that matrices drawn from the PSB model can be expressed as the sum of n rank one matrices. The supports of these matrices, under appropriate row and column permutations, can be organised into a single dense block. These permuted




Fig. 1. Visualisation of two examples of PSB matrices. In both cases the matrix consists of a sum of four rank 1 matrices, each with a support cardinality of 16. The left hand plot corresponds to the case where the support of each rank one matrix is arranged into a block. The right hand plot takes the same rank one matrices as in the left hand plot, but before summing them applies a random row and column permutation to each. The white squares indicate a zero entry and nonzero entries in the same column have the same coefficient value. As a result, aside from where there are overlaps, each column of a block is a single stripe of colour.

blocks are also striped in the sense that the entries in any given column of a block have the same value. We provide a visualisation in Figure 1.

C. Summary of contributions

The contributions of this paper are twofold: first, we provide a novel algorithm, Decoder-Expander Based Factorisation (D-EBF), detailed in Algorithm 6 in Section III-D, which is designed to recover $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ and $\mathbf{X} \in \chi_k^{n \times N}$ up to permutation from the measurement matrix $\mathbf{Y} = \mathbf{A}\mathbf{X}$. Second, we analyse the performance of D-EBF in the context of measurement matrices sampled from the PSB model, proving that it recovers the matrix factors up to permutation with high probability. The theoretical guarantees for D-EBF are summarised in Theorem II.3.

Theorem II.3. *Let $Y \sim \text{PSB}(d, k, m, n, N)$ as per Definition II.3. Under the assumption that $A \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$, consider the reconstructions of the matrix factors of Y returned by D-EBF,*

$$\hat{A}, \hat{X} \leftarrow \text{D-EBF}(Y, m, n, N, \epsilon, d).$$

Then the following statements are true.

- 1) **The reconstructions are accurate up to permutation:** there exists a random permutation $P \in \mathcal{P}^{n \times n}$ such that $\text{supp}(\hat{A}P^T) \subseteq \text{supp}(A)$, $\text{supp}(P\hat{X}) \subseteq \text{supp}(X)$ and $\hat{x}_{j,i} = x_{j,i}$ for all $(j, i) \in \text{supp}(\hat{X})$.
- 2) **On the uniqueness of the factorisation:** if $Y = \hat{A}\hat{X}$ then this factorisation is unique up to permutation, i.e., there exists a random permutation $P \in \mathcal{P}^{n \times n}$ such that $\hat{A}P^T = A$ and $P\hat{X} = X$.
- 3) **D-EBF is successful with high probability in n :** Suppose in addition to the assumptions of the PSB model that $N \geq \nu \frac{8d}{3\tau(n)} \frac{n}{k} (\ln^2(n) + \ln(n))$, where $\tau(n) = \mathcal{O}(1)$ ³ and $\nu > 1$ is a constant. Then the probability that there exists a permutation $P \in \mathcal{P}^{n \times n}$ such that $\hat{A}P^T = A$ and $P\hat{X} = X$ is greater than $1 - \mathcal{O}\left(n^{-\sqrt{\nu}+1} \log^2(n)\right)$.

A few remarks regarding the D-EBF algorithm and Theorem II.3 are in order.

³A definition of $\tau(n)$ is provided in Lemma IV.4

- **Probability that A is an expander:** recall from the discussion in Section II-B that, with A defined as in the PSB model, then the probability that $A \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ goes to one exponentially as $n \rightarrow \infty$. Furthermore, the probability of this event can be lower bounded using the bounds derived in [19]. As a result, from the definition of conditional probability, upper bounds on the probability of all three statements of Theorem II.3 can easily be derived. It therefore also follows that the three statements of Theorem II.3 still hold with high probability even without conditioning on the event $A \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$. We have opted not to present Theorem II.3 in this form for two reasons: first in order to more clearly highlight the contributions of this paper and second because the lower bound on the probability that $A \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ derived in [19] is loose and overly pessimistic.
- **Sample complexity of D-EBF:** perhaps the most remarkable aspect of Theorem II.3 is the required sample complexity. Indeed, observing that ν and d are constants and that $\tau(n) = \mathcal{O}(1)$, $n/k = \mathcal{O}(1)$, then as long as $N = \Omega\left(\frac{n}{k} \log^2(n)\right)$ D-EBF succeeds in recovering A and X up to permutation with high probability. This is equal to the lower bound $N \geq \frac{n}{k}$ up to logarithmic factors. For measurement matrices drawn from the PSB model we hypothesise that this is likely to be optimal as one $\log(n)$ factor arises from a coupon collector argument, inherent to the way X is sampled, and the second $\log(n)$ factor is needed to achieve the stated rate of convergence in probability.
- **Computational complexity of D-EBF:** in Section III-D we will ascertain that the per while loop iteration complexity cost of D-EBF is $\mathcal{O}(k^2(n + N)N)$. A trivial asymptotic upper bound of $\mathcal{O}(nk^2(n + N)N)$ follows from our method of proof, however this is highly pessimistic and analysing and bounding the number of iterations of the while loop of D-EBF in a meaningful manner is beyond the scope of Theorem II.3. Our preliminary experimental investigations on moderately sized problems indicate that the number of iterations required in practice is not onerous and, as expected, reduces as N increases. We leave a proper analysis to future work.
- **Input arguments required by D-EBF:** observe that D-EBF requires not only the product matrix Y and its dimensions, but also certain parameters of the encoder matrix: namely the number of columns n , the expansion parameter ϵ and the column sparsity d . As will be discussed in Section V, in practice n is not required. Indeed, ND-EBF can be run in an online fashion, processing input vectors as and when they are received and dynamically updating a list of column vectors representing \hat{A} . Although the the column sparsity d is known in advance by the encoder, and can therefore be reasonably agreed upon in advance with the decoder, the same cannot be said of the expansion parameter ϵ , as even with access to \mathbf{A} computing ϵ is an NP-complete problem [39]. However, as discussed and demonstrated empirically in section V, in practice ND-EBF can still function effectively given

access only to an upper bound on ϵ instead. In our experiments we use $1/6$ for simplicity but tighter upper bounds on the expansion parameter can be computed with relative ease, see e.g., [40].

As shall be discussed in more detail in Section V, as long as $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and N is sufficiently large, then in practice D-EBF requires only knowledge of the column sparsity d in advance to compute the factors \mathbf{A} and \mathbf{X} of \mathbf{Y} up to permutation. In the context of compressed sensing, the fact that the D-EBF algorithm is only able to recover the sparse codes up to permutation may be problematic. However, if the encoder and decoder agree in advance on a protocol with which to identify the original encoder matrix then this issue can be overcome. One such protocol to this end, based on a particular ordering of the columns, is as follows.

- 1) The encoder generates an encoder matrix A as per the PSB model and then, interpreting each column as a rational number written in binary, reorders the columns by the size of their corresponding rational number, for example, from largest to smallest. The specific ordering used is agreed in advance with the decoder.
- 2) The encoder then generates Y by multiplying X by the ordered encoder matrix and then sends Y and d to the decoder.
- 3) The decoder applies the D-EBF algorithm to recover the encoder matrix and sparse codes up to permutation. The decoder then interprets each of the columns of the reconstructed encoder matrix as a rational number written in binary, and reorders the columns of \hat{A} and the corresponding rows of \hat{X} by the size of their corresponding rational number.

If $A \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ then each column of A maps to a unique rational number and therefore the columns of A have a unique ordering. As a result, if the decoder knows the column ordering protocol, e.g., smallest to largest or largest to smallest, and sparsity d in advance, then recovery up to permutation implies full recovery of the location information also. We note that this condition can be viewed as a property of the distribution from which the encoder matrix is sampled. In order to decode a set of linear measurements, D-EBF does not need access to the specific encoder matrix used, but does need information as to the distribution from which this encoder was drawn.

D. Connections and related work

The PSB model and its associated factorisation task are most closely related to other matrix factorisations in which sparsity also plays a prominent role, such as in dictionary learning and subspace clustering. Dictionary learning [7], [10], [11], [12], [13] is a prominent matrix factorisation technique in data science, in which the columns of the observed data matrix are assumed to lie, at least approximately, on a union of low dimensional subspaces. This structure can be expressed as the product of an overcomplete matrix, known as a dictionary, and a sparse matrix. To be clear, given a data matrix \mathbf{Y} , dictionary learning methods seek to compute a dictionary \mathbf{A} and a sparse coding matrix \mathbf{X} such that $\mathbf{Y} \approx \mathbf{A}\mathbf{X}$. Compared

with dictionary learning, subspace clustering [41], [42], [43] adopts the additional structural assumption that the data lies on a union of low dimensional subspaces which are independent of one another.

More specifically, the literature most relevant to this work is that on dictionary recovery, a field aiming to provide recovery guarantees for dictionary learning. Compared with dictionary learning, dictionary recovery [44], [45], [46], [47], [48], [49], [50] presupposes that $\mathbf{Y} := \mathbf{A}\mathbf{X}$, the goal being to recover \mathbf{A} and \mathbf{X} up to permutation. It is common in dictionary recovery to place further structural assumptions on the factor matrices so as to facilitate the development of strong guarantees, albeit at the expense of model expressiveness. Two popular structural assumptions are that the dictionary is complete [48], [49], [50] or incoherent [44], [46]. Similar to the PSB model, the sparse coding matrix is often assumed to have been drawn from a sparse distribution, e.g., Bernoulli-Gaussian [44], [46] or Bernoulli-uniform [48], [49], [50]. This work diverges from the majority of the literature on dictionary recovery in regard to the construction of the encoder matrix vs the dictionary, as our goal is data compression rather than data expression. Furthermore, in dictionary recovery incoherence or completeness properties are typically adopted in order to design algorithms and derive guarantees. We instead operate under the assumption that the encoder matrix is the adjacency matrix of a (k, ϵ, d) -expander graph. As a result we use a very different set of ideas and tools in our algorithmic approach and proof technique.

In the context of matrix factorisation and dictionary recovery, we are aware of only one other work [51] in which A is constructed so as to leverage the properties of expander graphs. In this work the authors, guided by the notion of reversible neural networks, considered the learning of a deep network as a layerwise nonlinear dictionary learning problem. This work differs substantially from [51] in a number of respects. First, in terms of the problem setup, the primary difference is that in [51] the problem of recovering A and X from $Y = \sigma(\mathbf{A}\mathbf{X})$ is studied, where σ denotes the elementwise unit step function and both X and Y are binary. We emphasise that the problem we study here and that studied in [51] are not really comparable. Second, in regard to algorithmic approach, in this prior work A is recovered up to permutation using a non-iterative approach, involving first the computation of all row wise correlations of Y and then applying a clustering technique adapted from the graph square root problem. By contrast, and as described in detail in Section III-D, our method iteratively recovers parts of A and X directly from the columns of the residual by simultaneously leveraging both the unique neighbour property of A and the fact that the columns of X are dissociated. This iterative approach enables us to prove much stronger and in fact near optimal sample complexity guarantees.

III. DECODER-EXPANDER BASED FACTORISATION (D-EBF)

In this section we present the Decoder-Expander Based Factorisation (D-EBF) algorithm as a method for factorising

$\mathbf{Y} = \mathbf{A}\mathbf{X}$, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ and $\mathbf{X} \in \chi_k^{n \times N}$. In particular, in Section III-A we provide a high level overview of the D-EBF algorithm, before describing in detail the key steps in Sections III-B and III-C. Finally, we provide a full definition of D-EBF in Section III-D. In the body of this section we focus on providing high level summaries of each algorithm in order to better highlight the key ideas. Detailed versions and discussions of each algorithm are provided in Appendix A.

A. Overview of the D-EBF algorithm

The D-EBF algorithm iteratively computes a reconstruction of \mathbf{A} and \mathbf{X} by recovering parts of them from the observed matrix \mathbf{Y} , removing the contributions of these parts to form a residual, and then repeating these steps on said residual. The D-EBF algorithm starts each iteration by searching, independently, the columns of the residual for singleton values - nonzeros which are the sum of one nonzero in the corresponding column of \mathbf{X} . In Section III-B we prove, under certain assumptions, that singleton values can be identified by the number of times they appear in a given column of the residual. In addition, as the columns of \mathbf{X} are dissociated, then the locations in which a singleton value appears indicate part of the support of a column of \mathbf{A} . Furthermore, if a singleton value appears in sufficiently many locations then the associated partial support is unique to one column of \mathbf{A} . After identifying these frequently occurring singleton values from the columns of the residual, the D-EBF algorithm clusters the associated partial supports by their unique column of \mathbf{A} . The union of the partial supports in each cluster is then used to reconstruct, at least partially, a column of \mathbf{A} . The completely recovered columns, i.e., those whose supports have cardinality d , are then used to recover further entries of \mathbf{X} using a sparse decoding algorithm from the combinatorial compressed sensing literature [31], [32], [34], [33], [15]. Finally, the contributions from the recovered parts of \mathbf{A} and \mathbf{X} are removed from \mathbf{Y} to form a new residual. This process is then repeated on the subsequent residuals until either no additional entries from \mathbf{A} or \mathbf{X} are recovered, or \mathbf{Y} is fully decoded. We summarise this high level approach in Algorithm 2.

In the following sections we proceed to define each of the high level steps of Algorithm 2 in detail. In Section III-B we study steps 3, 4 and 5, which concern the extraction and clustering of singleton values and partial supports. In Section III-C we consider step 6, the decoding step, and review sparse decoding algorithms from the combinatorial compressed sensing literature. Finally, in Section III-D we present a detailed version of Algorithm 2.

B. Singleton values and partial supports

As referenced to in Section II-A, expander graphs play major role both in the design of D-EBF and the proof of Theorem II.3. In Lemma III.1 we summarise key properties of the adjacency matrices of (k, ϵ, d) -expander graphs.

Lemma III.1 (Properties of the adjacency matrix of a (k, ϵ, d) -expander graph). *If $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ then any submatrix \mathbf{A}_S of \mathbf{A} , where $S \in [n]^{\leq k}$, satisfies the following.*

Algorithm 2 Summary of D-EBF($\mathbf{Y}, m, n, N, \epsilon, d$)

- 1: Initialise residual as observed data \mathbf{Y}
 - 2: **while** not converged **do**
 - 3: Extract singleton values and associated partial supports from the residual
 - 4: Cluster the partial supports by their column of \mathbf{A}
 - 5: Update the reconstructions of \mathbf{A} and \mathbf{X} using the clustered partial supports and the singleton values respectively
 - 6: Run a decode algorithm using the fully recovered columns of \mathbf{A} to recover more entries of \mathbf{X} .
 - 7: Update residual by removing contributions from recovered parts of \mathbf{A} and \mathbf{X}
 - 8: **end while**
-

- 1) *There are more than $(1 - \epsilon)d|S|$ rows in \mathbf{A}_S that have at least one non-zero.*
- 2) *There are more than $(1 - 2\epsilon)d|S|$ rows in \mathbf{A}_S that have only one non-zero.*
- 3) *The overlap in support of the columns of \mathbf{A}_S is upper bounded as follows,*

$$\left| \bigcap_{l \in S} \text{supp}(\mathbf{a}_l) \right| < 2\epsilon d.$$

A proof of Lemma III.1 is provided in Appendix C. The inspiration for studying singleton values arises from property 2 of Lemma III.1. Indeed, by conditioning on \mathbf{A} being the adjacency matrix of a (k, ϵ, d) -expander graph, then the submatrix of \mathbf{A} associated with the support of a column of \mathbf{X} has potentially many rows with a single nonzero entry. This implies that there are potentially many entries in each column of \mathbf{Y} which correspond directly to an entry in \mathbf{X} . Denoting the j th row of \mathbf{A} as $\tilde{\mathbf{a}}_j$, we are now ready to provide the definition of a singleton value.

Definition III.1 (Singleton value). *Consider a vector $\mathbf{r} = \mathbf{A}\mathbf{z}$ where $\mathbf{A} \in \{0, 1\}^{m \times n}$ and $\mathbf{z} \in \mathbb{R}^n$. A singleton value of \mathbf{r} is an entry $r_j \neq 0$ such that $|\text{supp}(\tilde{\mathbf{a}}_j) \cap \text{supp}(\mathbf{z})| = 1$, hence $r_j = x_l$ for some $l \in \text{supp}(\mathbf{z})$.*

We call a binary vector whose nonzeros coincide with some subset of the nonzeros of a column of \mathbf{A} a partial support of that column.

Definition III.2 (Partial support). *A partial support $\mathbf{w} \in \{0, 1\}^m$ of a column $\mathbf{a}_l \in \{0, 1\}^m$ of $\mathbf{A} \in \{0, 1\}^{m \times n}$ is a binary vector satisfying $\text{supp}(\mathbf{w}) \subseteq \text{supp}(\mathbf{a}_l)$. Furthermore, \mathbf{w} is said to originate from \mathbf{a}_l iff $\text{supp}(\mathbf{w}) \subseteq \text{supp}(\mathbf{a}_l)$ and $\text{supp}(\mathbf{w}) \not\subseteq \text{supp}(\mathbf{a}_h)$ for all $h \in [n] \setminus \{l\}$.*

We now introduce some useful notation for what follows.

- Let the function $f : \mathbb{R} \times \mathbb{R}^m \rightarrow [m]$ count the number of times a real number $\alpha \in \mathbb{R}$ appears in some vector $\mathbf{r} \in \mathbb{R}^m$, i.e., $f(\alpha, \mathbf{r}) := |\{j \in [m] : r_j = \alpha\}|$.
- Let the function $g : \mathbb{R} \times \mathbb{R}^m \rightarrow \{0, 1\}^m$ return a binary vector whose nonzeros correspond to the locations in which $\alpha \in \mathbb{R}$ appears in a vector $\mathbf{r} \in \mathbb{R}^m$, i.e., with $\mathbf{w} = g(\alpha, \mathbf{r})$ then $w_j = 1$ iff $r_j = \alpha$ and is 0 otherwise.

Considering Definition II-B, observe that if \mathbf{z} is dissociated then the locations in which a singleton value appears in \mathbf{r} defines a partial support of a column of \mathbf{A} . Furthermore, under the assumption that \mathbf{A} is the adjacency matrix of a (k, ϵ, d) -expander graph and that \mathbf{z} is dissociated, it is possible to derive the following sufficient condition with which to identify singleton values.

Corollary III.1.1 (Sufficient condition for identifying singleton values (i)). *Consider a vector $\mathbf{r} = \mathbf{A}\mathbf{z}$ where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ and $\mathbf{z} \in \chi_k^n$. For $\alpha \in \mathbb{R} \setminus \{0\}$, if $f(\alpha, \mathbf{r}) \geq 2\epsilon d$ then there exists an $l \in \text{supp}(\mathbf{z})$ such that $\alpha = z_l$.*

A proof of Corollary III.1.1 is provided in Appendix C. Under the assumptions stated this corollary implies that any value which appears in \mathbf{r} at least $2\epsilon d$ times must be a singleton value. However, this corollary does not provide insight into whether or not there exist singleton values in \mathbf{r} appearing in at least $2\epsilon d$ locations. To resolve this matter, in Lemma III.2, adapted from [34, Theorem 4.6], we show, under certain assumptions, that there always exist a positive number of singleton values which appear more than $(1 - 2\epsilon)d$ times in \mathbf{r} . This result is needed for the proof of Theorem II.3.

Lemma III.2 (Existence of frequently occurring singleton values, adapted from [34, Theorem 4.6]). *Consider a vector $\mathbf{r} = \mathbf{A}\mathbf{z}$ where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ and $\mathbf{z} \in \chi_k^n$. For $l \in \text{supp}(\mathbf{z})$, let $\Omega_l := \{j \in [m] : r_j = z_l\}$ be the set of row indices $j \in [m]$ of \mathbf{r} such that $r_j = z_l$. Defining $\mathcal{T} := \{l \in \text{supp}(\mathbf{z}) : |\Omega_l| > (1 - 2\epsilon)d\}$ as the set of singleton values which each appear more than $(1 - 2\epsilon)d$ times in \mathbf{r} , then*

$$|\mathcal{T}| \geq \frac{|\text{supp}(\mathbf{z})|}{(1 + 2\epsilon)d}.$$

A proof of Lemma III.2 is provided in Appendix C. If $\epsilon \leq 1/4$ then $(1 - 2\epsilon)d \geq 2\epsilon d$, combining this with Corollary III.1.1 and Lemma III.2 guarantees, under the necessary assumptions, that it is always possible to identify at least one singleton value which appears more than $2\epsilon d$ times in \mathbf{r} . The results presented so far concerning the extraction of singleton values and partial supports assume that the residual analysed is of the form $\mathbf{r} = \mathbf{A}\mathbf{z}$, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ and $\mathbf{z} \in \chi_k^n$. Due to the iterative nature of D-EBF, typically the residual under consideration is instead of the form $\mathbf{r} = \mathbf{y} - \hat{\mathbf{A}}\hat{\mathbf{x}}$, where $\hat{\mathbf{A}}$ and $\hat{\mathbf{x}}$ are the estimates or reconstructions of \mathbf{A} and \mathbf{x} respectively. Under certain assumptions, the following result maintains that the sufficient condition given in Corollary III.1.1 also holds for residuals of this form.

Lemma III.3 (Sufficient condition for identifying singleton values (ii)). *Consider a vector $\mathbf{y} = \mathbf{A}\mathbf{x}$ where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ and $\mathbf{x} \in \chi_k^n$. Let $\hat{\mathbf{A}} \in \{0, 1\}^{m \times n}$ and $\hat{\mathbf{x}} \in \chi_k^n$ be such that a column $\hat{\mathbf{a}}_l$ of $\hat{\mathbf{A}}$ is nonzero iff $\hat{x}_l \neq 0$, and there exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\text{supp}(\hat{\mathbf{A}}\mathbf{P}^T) \subseteq \text{supp}(\mathbf{A})$, $\text{supp}(\mathbf{P}\hat{\mathbf{x}}) \subseteq \text{supp}(\mathbf{x})$ and $\hat{x}_{P(l)} = x_l$ for all $l \in \text{supp}(\hat{\mathbf{x}})$, where $P : [n] \rightarrow [n]$ denotes the row permutation caused by pre-multiplication with \mathbf{P} . Consider the residual $\mathbf{r} = \mathbf{y} - \hat{\mathbf{A}}\hat{\mathbf{x}}$, if for some $\alpha \in \mathbb{R} \setminus \{0\}$ it holds that $f(\alpha, \mathbf{r}) \geq 2\epsilon d$ is satisfied, then there exists an $l \in [n]$ such that $\alpha = x_l$.*

Furthermore, with $\mathbf{w} = g(\alpha, \mathbf{r})$ then $\text{supp}(\mathbf{w}) \subseteq \text{supp}(\mathbf{a}_l)$ is a partial support of \mathbf{a}_l .

A proof of Lemma III.2 is provided in Appendix C. To perform step 4 of Algorithm 2, it is necessary that the partial supports extracted across all columns of \mathbf{Y} , or the residual of \mathbf{Y} , can be clustered accurately and efficiently. To this end we present Lemma III.4, which provides a necessary and sufficient condition with which to cluster sufficiently large partial supports.

Lemma III.4 (Clustering partial supports). *Consider a pair of partial supports \mathbf{w}_1 and \mathbf{w}_2 , extracted from $\mathbf{r}_1 = \mathbf{A}\mathbf{z}_1$ and $\mathbf{r}_2 = \mathbf{A}\mathbf{z}_2$ respectively, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ and $\mathbf{z}_1, \mathbf{z}_2 \in \chi_k^n$. If $\epsilon \leq 1/6$, $|\text{supp}(\mathbf{w}_1)| > (1 - 2\epsilon)d$ and $|\text{supp}(\mathbf{w}_2)| > (1 - 2\epsilon)d$, then \mathbf{w}_1 and \mathbf{w}_2 originate from the same column of \mathbf{A} iff $\mathbf{w}_1^T \mathbf{w}_2 \geq 2\epsilon d$.*

A proof of Lemma III.4 is provided in Appendix C. As highlighted in the proof, $\mathbf{w}_1^T \mathbf{w}_2 \geq 2\epsilon d$ is sufficient to conclude that any two partial supports originate from the same column. However, without adding the additional assumptions on ϵ and on the size of the partial supports, we cannot conclude that $\mathbf{w}_1^T \mathbf{w}_2 < 2\epsilon d$ implies \mathbf{w}_1 and \mathbf{w}_2 do not originate from the same column. This limits us, for now, to clustering only fairly large partial supports, which have at least 2/3rds of the total nonzero entries of their respective columns of \mathbf{A} . Fortunately, under the same conditions as Lemma III.4, Lemma III.2 guarantees the existence of partial supports of this size. Note however, that if we have access to a complete column \mathbf{a}_l of \mathbf{A} , then $\mathbf{w}_1^T \mathbf{a}_l \geq 2\epsilon d$ is sufficient to conclude that \mathbf{w}_1 originates from \mathbf{a}_l . This implies that it is possible to assign partial supports consisting of only around 1/3rd of the total nonzeros to a column of \mathbf{A} . For this reason we will later introduce the decode step, discussed in section III-C, to try and match smaller partial supports to the columns of $\hat{\mathbf{A}}$ which have d nonzeros.

We are now in a position to define and discuss the subroutines corresponding to steps 3, 4 and 5 of Algorithm 2. We emphasise before proceeding that the subroutines we will present are designed around the assumption that $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$. Furthermore our presentation here is focused on communicating the key, and often fairly simple ideas involved. For readers interested in implementing these algorithms themselves, we refer them to Appendix A, where more detailed definitions and in-depth discussions can be found.

Consider the column vector $\mathbf{y} = \mathbf{A}\mathbf{x}$ and suppose that $\hat{\mathbf{A}}$ and $\hat{\mathbf{x}}$ are the current reconstructions of \mathbf{A} and \mathbf{x} . Algorithm 3 processes the residual $\mathbf{r} = \mathbf{y} - \hat{\mathbf{A}}\hat{\mathbf{x}}$, taking as input arguments \mathbf{r} , the dimension m of \mathbf{r} , the expansion parameter ϵ , the column sparsity d , as well as the current reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{x}}$. The algorithm identifies singleton values and partial supports, either matching them with a nonzero column of $\hat{\mathbf{A}}$ or identifying them as belonging to a column of \mathbf{A} not yet observed. To this end, Algorithm 3 returns the unmatched singleton values and partial supports, stored as entries of the vector \mathbf{q} and columns of the matrix \mathbf{W} respectively, the number of unmatched partial supports p , the updated reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{x}}$ and a Boolean variable UPDATED, which

indicates whether or not the reconstructions have in fact been updated. A detailed definition and description of Algorithm 3 is given in Appendix A, Algorithm 9. The computational complexity of Algorithm 3 is $\mathcal{O}(k^2n)$ if the sparseness of \mathbf{r} , $\hat{\mathbf{A}}$ and $\hat{\mathbf{x}}$ and is taken advantage of, and is $\mathcal{O}(m^2n^2)$ if not. Again, we refer the reader to Appendix A-B for a more detailed discussion on the computational complexity of this algorithm.

Algorithm 3 Summary EXTRACT&MATCH($\mathbf{r}, m, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{x}}$)

- 1: Initialise \mathbf{q} as an empty vector, \mathbf{W} as an empty matrix and p as 0.
 - 2: Compute the number of times each unique entry of \mathbf{r} appears in \mathbf{r} .
 - 3: **for** each unique entry r' in \mathbf{r} **do**
 - 4: Construct the partial support \mathbf{w} associated with r .
 - 5: **if** r' appears more than $(1 - 2\epsilon)d$ **then**
 - 6: Find the column $\hat{\mathbf{a}}_l$ of $\hat{\mathbf{A}}$ whose support overlaps most with $\text{supp}(\mathbf{w})$.
 - 7: **if** the overlap is at least $2\epsilon d$ **then**
 - 8: Set $\text{supp}(\mathbf{a}_l) \leftarrow \text{supp}(\mathbf{a}_l) \cup \text{supp}(\mathbf{w})$.
 - 9: Assign r' to the corresponding entry in $\hat{\mathbf{x}}$.
 - 10: **else**
 - 11: Add \mathbf{w} as a column of \mathbf{W} .
 - 12: Add r' as an entry of \mathbf{q}
 - 13: $p \leftarrow p + 1$
 - 14: **end if**
 - 15: **else**
 - 16: **if** there exists $l \in [n]$ such that $r' = \hat{x}_l$ **then**
 - 17: Set $\text{supp}(\mathbf{a}_l) \leftarrow \text{supp}(\mathbf{a}_l) \cup \text{supp}(\mathbf{w})$.
 - 18: **end if**
 - 19: **end if**
 - 20: **end for**
 - 21: Return $\mathbf{q}, \mathbf{W}, p, \hat{\mathbf{A}}, \hat{\mathbf{x}}, \text{UPDATED}$
-

We now turn our attention to clustering the unmatched partial supports and singleton values extracted across all columns of the residual $\mathbf{R} = \mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}$. This step of D-EBF is performed by Algorithm 4, which takes as inputs the unmatched singleton values, stored in the entries of the vector $\mathbf{q} \in \mathbb{R}^p$, the unmatched partial supports stored in the columns of the matrix $\mathbf{W} \in \{0, 1\}^{m \times p}$, ordered so that \mathbf{w}_i is the partial support associated with the singleton q_i , the number of unmatched singleton values p , the smallest zero column index $\eta := \min_{l \in [n]} \{\hat{\mathbf{a}}_l = \mathbf{0}_m\}$, the vector $\mathbf{h} \in [N]^p$ which stores the indices of the columns of \mathbf{R} from which each singular value was extracted, i.e., h_i is the column index of \mathbf{R} from which q_i was extracted, the expansion parameter ϵ and column sparsity d of \mathbf{A} , and the current reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$. A detailed definition and description of Algorithm 4 is provided in Appendix A-C, Algorithm 10.

The main expense of Algorithm 10 is the computation of the inner products between different partial supports, equivalent to a matrix product between $\hat{\mathbf{A}}$ and \mathbf{W} . As $p = \mathcal{O}(kN)$ then if the sparsity of \mathbf{W} is not leveraged this matrix product costs $\mathcal{O}(k^2mN^2)$. If \mathbf{W} is stored using a sparse format, as would be natural, then, as each column of \mathbf{W} is d sparse with d fixed, the computational cost can be reduced to $\mathcal{O}(k^2N^2)$.

Algorithm 4 Summary CLUSTER&ADD($\mathbf{q}, \mathbf{W}, p, \eta, \mathbf{h}, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{X}}$)

- 1: Compute all pairwise inner products between unmatched partial supports, i.e., columns of \mathbf{W} .
 - 2: Cluster partial supports: two partial supports are put in the same group iff their pairwise inner product is at least $2\epsilon d$.
 - 3: **for** each group \mathcal{C} of partial supports **do**
 - 4: Set $\text{supp}(\hat{\mathbf{a}}_\eta) \leftarrow \bigcup_{j \in \mathcal{C}} \text{supp}(\mathbf{w}_j)$
 - 5: Update $\hat{x}_{\eta, h_j} \leftarrow q_j$
 - 6: Increment $\eta \leftarrow \eta + 1$
 - 7: **end for**
 - 8: Return $\hat{\mathbf{A}}, \hat{\mathbf{X}}, \eta$
-

C. Using a decoder algorithm to improve recovery

While the nonzeros of \mathbf{A} can be recovered from potentially many of the columns of \mathbf{R} , each nonzero of \mathbf{X} can be recovered only from one column. This makes the recovery, up to row permutation, of \mathbf{X} typically more challenging than than the recovery, up to column permutation, of \mathbf{A} . Due to Lemma III.4, in Algorithm 3 partial supports which are at most $(1 - 2\epsilon)d$ are not clustered or matched with a column of $\hat{\mathbf{A}}$. This is because the updates of prior iterations only guarantee that each nonzero column of the reconstruction has more than $(1 - 2\epsilon)d$ nonzeros. However, if we know that $\hat{\mathbf{a}}_l$ is complete, i.e. there exists a $h \in [n]$ such that $\hat{\mathbf{a}}_l = \mathbf{a}_h$, then $\hat{\mathbf{a}}_l^T \mathbf{w} \geq 2\epsilon d$ is sufficient to conclude that a partial support \mathbf{w} originates from $\hat{\mathbf{a}}_l$. Therefore, it is possible to match smaller partial supports, i.e., in practice those of size approximately $d/3$ rather than $2d/3$, to columns of the reconstruction which are complete. To this end, while Algorithm 9 focuses on clustering partial supports with cardinality larger than $(1 - 2\epsilon)d$ and matching them to a nonzero column of $\hat{\mathbf{A}}$, the decode algorithm, summarised in Algorithm 5, focuses on matching any partial support whose cardinality is at least $2\epsilon d$ to a complete column of $\hat{\mathbf{A}}$. More generally, the decode step should be considered a placeholder for one of the many algorithms from the combinatorial compressed sensing literature, aiming to at least partially decode $\hat{\mathbf{X}}$ given \mathbf{R} and the subset of complete columns of the encoder reconstruction $\hat{\mathbf{A}}_S$. In particular, [34, Algorithms 1 & 2] are particularly appropriate for deployment in this context as they are based on a model in which the columns of \mathbf{X} also satisfy the dissociated property, and, as can be observed in [34], are able to recover \mathbf{X} from \mathbf{Y} given \mathbf{A} even when k is large, i.e., $k \approx m/3$. For clarity, and to allow for the reuse of certain tools and results, in our analysis we will consider the decoder subroutine summarised in Algorithm 5, which is inspired by Expander Recovery (ER) [15]. For a more detailed description and discussion of this algorithm we refer the reader to Appendix A-D, Algorithm 11.

Without leveraging sparsity, the overall the cost per iteration of the while loop is $\mathcal{O}(m^2n)$. As $\hat{\mathbf{x}}$ has at most k nonzeros then the total computational cost of Algorithm 11 is $\mathcal{O}(km^2n)$. However, if \mathbf{r} , $\hat{\mathbf{A}}_S$ and $\hat{\mathbf{x}}$ are stored using a sparse format, then the per iteration cost of the while loop is $\mathcal{O}(kn)$, resulting in a reduced total cost $\mathcal{O}(k^2n)$. Again, we refer the reader to Appendix A-D for a more detailed description of the algorithm and discussion of its complexity.

Algorithm 5 Summary DECODE($\mathbf{y}, m, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{x}}$)

```

1:  $\mathcal{S} \leftarrow \{l \in [n] : |\text{supp}(\mathbf{a}_l)| = d\}$ 
2: while  $\hat{\mathbf{x}}$  is updated do
3:   Compute the residual  $\mathbf{r} = \mathbf{y} - \hat{\mathbf{A}}\hat{\mathbf{x}}$ 
4:   Compute the number of times each unique entry of  $\mathbf{r}$ 
     appears in  $\mathbf{r}$ .
5:   for each unique entry  $r'$  in  $\mathbf{r}$  do
6:     Construct the partial support  $\mathbf{w}$  associated with  $r$ .
7:     if  $r'$  appears at least  $2\epsilon d$  times then
8:       Find the complete column  $\hat{\mathbf{a}}_l$  of  $\hat{\mathbf{A}}_{\mathcal{S}}$  whose support
         has the largest intersection with  $\text{supp}(\mathbf{w})$ .
9:       if the overlap is at least  $2\epsilon d$  then
10:        Assign  $r'$  to the corresponding entry in  $\hat{\mathbf{x}}$ .
11:       end if
12:     end if
13:   end for
14: end while
15: Return  $\hat{\mathbf{x}}, \text{UPDATED}$ 

```

D. Detailed summary of the D-EBF algorithm

Algorithm 6 provides a detailed version of Algorithm 2. The algorithm takes as inputs the product matrix $\mathbf{Y} \in \mathbb{R}^{m \times N}$, the dimensions m, n and N of \mathbf{A} and \mathbf{X} , and the expansion parameter ϵ and column sparsity d of \mathbf{A} . The algorithm returns the estimates $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ of \mathbf{A} and \mathbf{X} respectively, aiming to recover them up to permutation. We emphasise again that the D-EBF algorithm implicitly assumes that $\mathbf{A} \in \mathcal{E}_{k, \epsilon, d}^{m \times n}$ with $\epsilon \leq 1/6$.

The while loop starting on line 6 of Algorithm 6 runs until either $\hat{\mathbf{A}}\hat{\mathbf{X}} = \mathbf{Y}$, or no changes to the residual \mathbf{R} of \mathbf{Y} occur. The for loop starting on line 11 extracts partial supports and singleton values from each column of \mathbf{R} as per Algorithm 9, storing the unmatched singleton values and partial supports in \mathbf{q} and \mathbf{W} respectively. On line 15 the vector \mathbf{h} is used to keep track of each singleton value's column of origin, while on line 16 p is used to count the number of singleton values extracted across all columns of \mathbf{R} at the current iteration. So long as at least one unmatched singleton value is found, then on line 19 the associated unmatched partial supports are clustered and used to construct new nonzero columns in $\hat{\mathbf{A}}$. Note that the nonzero columns of $\hat{\mathbf{A}}$ are added from left to right, with the variable η tracking the zero column of $\hat{\mathbf{A}}$ with the smallest index. To further grow the support of $\hat{\mathbf{X}}$, the for loop starting on line 22 applies Algorithm 11 to each of the columns of \mathbf{Y} . Alternatively, and as discussed in Section III-C, some other decoding algorithm from the combinatorial compressed sensing literature could be used instead. Finally, given the updates to $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$, the residual is then recomputed on line 25 in preparation for singleton value and partial support extraction and clustering at the next iteration.

In terms of computational cost the key contributors are lines 12, 19, 23 and 25. We now present the computational cost first without and then with taking advantage of sparsity. As per the discussion in Section III-B, the for loop running from line 11 to 17 has a cost of $\mathcal{O}(m^2 n^2 N)$ or $\mathcal{O}(k^2 n N)$ respectively. Observe also that by transferring the matching

Algorithm 6 D-EBF($\mathbf{Y}, m, n, N, \epsilon, d$)

```

1:  $\hat{\mathbf{A}} \leftarrow \text{zeros}(m, n)$ 
2:  $\hat{\mathbf{X}} \leftarrow \text{zeros}(n, N)$ 
3:  $\mathbf{R} \leftarrow \mathbf{Y}$ 
4: UPDATED  $\leftarrow$  TRUE
5:  $\eta \leftarrow 1$ 
6: while  $\mathbf{R} \neq \mathbf{0}_{m \times N}$  and UPDATED=TRUE do
7:    $\mathbf{q} \leftarrow []$ 
8:    $\mathbf{W} \leftarrow []$ 
9:    $\mathbf{h} \leftarrow []$ 
10:   $p \leftarrow 0$ 
11:  for  $i = 1 : N$  do
12:     $\mathbf{u}, \mathbf{V}, c, \hat{\mathbf{A}}, \hat{\mathbf{x}}_i, \text{UPDATED}$ 
       $\leftarrow \text{EXTRACT\&MATCH}(\mathbf{r}_i, m, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{x}}_i)$ 
13:     $\mathbf{q} \leftarrow [\mathbf{q}^T; \mathbf{u}^T]$ 
14:     $\mathbf{W} \leftarrow [\mathbf{W}; \mathbf{V}]$ 
15:     $\mathbf{h} \leftarrow [\mathbf{h}, i \times \text{zeros}(c)]$ 
16:     $p \leftarrow p + c$ 
17:  end for
18:  if  $p > 0$  then
19:     $\hat{\mathbf{A}}, \hat{\mathbf{X}}, \eta \leftarrow \text{CLUSTER\&ADD}(\mathbf{q}, \mathbf{W}, p, \eta, \mathbf{h}, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{X}})$ 
20:    UPDATED  $\leftarrow$  TRUE
21:  end if
22:  for  $i=1:N$  do
23:     $\hat{\mathbf{x}}_i, \text{UPDATED} \leftarrow \text{DECODE}(\mathbf{y}_i, m, N, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{x}}_i)$ 
24:  end for
25:   $\mathbf{R} \leftarrow \mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}$ 
26: end while
27: Return  $\hat{\mathbf{A}}, \hat{\mathbf{X}}$ 

```

process, lines 17-24 of Algorithm 9, to Algorithm 10, then singleton values and partial supports can be extracted from each column of \mathbf{R} in parallel, using N processors with a computational cost of $\mathcal{O}(k^2 n)$ per processor. Indeed, this step is trivially parallelizable across the columns of \mathbf{R} . Likewise, from the discussion in Section III-B, line 19 has a cost of $\mathcal{O}(k^2 m N^2)$ or $\mathcal{O}(k^2 N^2)$ respectively. In comparison this step cannot be parallelized. Turning our attention to the decoding step of the algorithm, then, as per the discussion in Section III-C, the for loop running from line 22 to 24 has a cost of $\mathcal{O}(k m^2 n N)$ or $\mathcal{O}(k^2 n N)$ respectively. This step is again trivially parallelizable across the columns of \mathbf{Y} . Finally, the matrix product on line 25 has a cost of $\mathcal{O}(m n N)$ or $\mathcal{O}(k n N)$ respectively. Therefore, the while loop running from lines 6 to 26 of Algorithm 6 has a per iteration cost of $\mathcal{O}(m(m n^2 + k^2 N) N)$ if sparsity is not leveraged, or $\mathcal{O}(k^2(n + N) N)$ if it is. Under certain assumptions we are also able to provide accuracy guarantees for Algorithm 6, as detailed in Lemma III.5. Finally, as per the discussion following Lemma III.5 below, we remark here that D-EBF terminates in at most kN iterations. We emphasise that this upper bound on the number of iterations is highly pessimistic and very loose in practice.

Lemma III.5 (Accuracy of D-EBF). *Let $\mathbf{Y} = \mathbf{A}\mathbf{X}$, where $\mathbf{A} \in \mathcal{E}_{k, \epsilon, d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{X} \in \chi_k^{n \times N}$. Consider Algorithm 6, D-EBF: in regard to the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ computed*

at any point during the run-time of D-EBF, we will say \mathcal{Q} is true iff the following hold.

- (a) A column $\hat{\mathbf{a}}_l$ of $\hat{\mathbf{A}}$ is nonzero iff the l th row of $\hat{\mathbf{X}}$ is nonzero.
- (b) For all $l \in [n]$ if $\hat{\mathbf{a}}_l$ is nonzero then $\text{supp}(\hat{\mathbf{a}}_l) > (1-2\epsilon)d$.
- (c) There exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\text{supp}(\hat{\mathbf{A}}\mathbf{P}^T) \subseteq \text{supp}(\mathbf{A})$, $\text{supp}(\mathbf{P}\hat{\mathbf{X}}) \subseteq \text{supp}(\mathbf{X})$ and $\hat{x}_{P(l),h} = x_{P(l),h}$ for all $h \in [N]$ and $l \in \text{supp}(\hat{\mathbf{x}}_l)$, where $P: [n] \rightarrow [n]$ denotes the row permutation caused by pre-multiplication with \mathbf{P} .

Suppose that D-EBF exits after the completion of iteration $t_f \in \mathbb{N}$ of the while loop, then the following statements are true.

- 1) At any point during the run time of D-EBF \mathcal{Q} is true. Therefore the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ returned by D-EBF also satisfy the conditions for \mathcal{Q} to be true.
- 2) For any iteration of the while loop $t < t_f$, let $\hat{\mathbf{A}}^{(1)}$ and $\hat{\mathbf{X}}^{(1)}$ denote the reconstructions at the start of the while loop, line 6, and $\hat{\mathbf{A}}^{(2)}$ and $\hat{\mathbf{X}}^{(2)}$ denote the reconstructions at the end of the while loop, line 26. Then either $\text{supp}(\hat{\mathbf{A}}^{(1)}) \subset \text{supp}(\hat{\mathbf{A}}^{(2)})$ and $\text{supp}(\hat{\mathbf{X}}^{(1)}) \subseteq \text{supp}(\hat{\mathbf{X}}^{(2)})$, or $\text{supp}(\hat{\mathbf{A}}^{(1)}) \subseteq \text{supp}(\hat{\mathbf{A}}^{(2)})$ and $\text{supp}(\hat{\mathbf{X}}^{(1)}) \subset \text{supp}(\hat{\mathbf{X}}^{(2)})$.
- 3) Consider the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ returned by D-EBF and the associated residual $\mathbf{R} = \mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}$: there exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$ and $\mathbf{P}\hat{\mathbf{X}} = \mathbf{X}$ iff $\mathbf{R} = \mathbf{0}_{m \times N}$. As a result, $\mathbf{R} = \mathbf{0}_{m \times N}$ is a sufficient condition to ensure that \mathbf{Y} has a unique, up to permutation, factorisation of the form $\mathbf{Y} = \mathbf{A}\mathbf{X}$, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{X} \in \chi_k^{n \times N}$.

A proof of Lemma III.5 is provided in Appendix D. Observe that the exit iteration t_f satisfies $t_f \leq kN$. This follows from the fact that at every iteration a new, as of yet unseen singleton value must be identified and extracted from the residual in order for an update to the reconstructions to occur, and there are kN nonzero entries in \mathbf{X} . We emphasise that under the PSB model this bound is highly pessimistic and in practice we typically observe far fewer iterations as per Section V. Furthermore, under the conditions of Theorem II.3, and by studying the Naive Decoder-Expander Based Factorisation (ND-EBF) algorithm discussed in Section IV-A, it can be shown that $t_f \leq N$ with high probability. We leave analysing and deriving tighter bounds on the number of iterations of D-EBF for future work.

The implication of Lemma III.5 is that as long as $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$, then the reconstructions returned by D-EBF are accurate up to permutation, and if they are in addition complete up to permutation then this factorisation is unique. Without further assumptions on \mathbf{X} it is not possible to derive guarantees in respect to fully recovering the matrix factors up to permutation. For instance, if a row of \mathbf{X} is a zero row vector then the corresponding column of \mathbf{A} contributes to no columns of \mathbf{Y} and therefore cannot be hoped to be recovered. In Section IV we will derive such guarantees by assuming that Y is drawn from the PSB model, defined in Definition II.3.

IV. THEORETICAL GUARANTEES FOR D-EBF UNDER THE PSB MODEL

In this section we derive Theorem II.3. Analysing D-EBF directly is challenging, therefore our approach is instead to study a simpler, surrogate algorithm, which we use to lower bound the performance of D-EBF. To this end we introduce the Naive Decoder-Expander Based Factorisation Algorithm (ND-EBF). Although highly suboptimal from a computational perspective, and therefore clearly not recommended for deployment in practice, this algorithm allows us to analyse the parameter regime in which D-EBF is successful with high probability. The structure of this section is as follows: in Section IV-A we present and describe the ND-EBF algorithm, provide accuracy guarantees analogous to Lemma III.5 and connect ND-EBF with D-EBF. Then, in Section IV-B, we use these results to prove Theorem II.3.

A. Naive Decoder-Expander Based Factorisation (ND-EBF)

We provide a high-level summary of the ND-EBF algorithm below in Algorithm 7. The key distinction between ND-EBF and D-EBF is that at each iteration one and only column of \mathbf{A} is attempted to be reconstructed. If the new, reconstructed column of \mathbf{A} has its full compliment of ones, i.e., d of them, then ND-EBF proceeds to the next iteration, otherwise it terminates. By design, this reduces the recovery problem to that of recovering a column at every iteration, which we shall observe is easier to analyse. For a detailed description and discussion of the ND-EBF algorithm we refer the reader to Appendix A-F and Algorithm 13.

Algorithm 7 Summary of ND-EBF($\mathbf{Y}, m, n, N, \epsilon, d$)

- 1: Initialise $\mathbf{R} \leftarrow \mathbf{Y}$, $\text{RUN} \leftarrow \text{TRUE}$ and $\eta \leftarrow 0$
 - 2: **while** $\text{RUN} = \text{TRUE}$ **do**
 - 3: Extract singleton values and associated partial supports from the residual
 - 4: Identify the largest cluster \mathcal{C}^* of partial supports.
 - 5: Set $\text{supp}(\hat{\mathbf{a}}_\eta) = \bigcup_{i \in \mathcal{C}^*} \text{supp}(\mathbf{w}_i)$
 - 6: **if** $|\text{supp}(\hat{\mathbf{a}}_\eta)| = d$ **then**
 - 7: Run a decode algorithm using the fully recovered columns of \mathbf{A} to recover more entries of \mathbf{X} .
 - 8: Update $\mathbf{R} \leftarrow \mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}$ and $\eta \rightarrow \eta + 1$.
 - 9: **else**
 - 10: $\text{RUN} \leftarrow \text{FALSE}$
 - 11: **end if**
 - 12: **end while**
-

Lemma IV.1 provides accuracy guarantees for ND-EBF, similar to those afforded to D-EBF by Lemma III.5.

Lemma IV.1 (Accuracy of ND-EBF). *Let $\mathbf{Y} = \mathbf{A}\mathbf{X}$, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{X} \in \chi_k^{n \times N}$. Consider Algorithm 13, ND-EBF: in regard to the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ computed at any point during the run-time of ND-EBF, we will say, for $\eta \in [n]$, that $\mathcal{Q}(\eta)$ is true iff the following hold.*

- (a) *The column vector $\hat{\mathbf{a}}_l$ and row vector $\hat{\mathbf{x}}_l$ are nonzero iff $l < \eta$.*
- (b) *If $l < \eta$ then $|\text{supp}(\hat{\mathbf{a}}_l)| = d$.*

(c) *There exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\text{supp}(\hat{\mathbf{A}}\mathbf{P}^T) \subseteq \text{supp}(\mathbf{A})$, $\text{supp}(\mathbf{P}\hat{\mathbf{X}}) \subseteq \text{supp}(\mathbf{X})$ and $\hat{x}_{P(l),i} = x_{P(l),i}$ for all $i \in [N]$ and $l \in \text{supp}(\hat{\mathbf{x}}_l)$, where $P : [n] \rightarrow [n]$ denotes the row permutation caused by pre-multiplication with \mathbf{P} .*

Suppose that ND-EBF exits after the completion of iteration $\eta_f \in [n]$ of the while loop, then the following statements are true.

- 1) *For all $\eta \in [\eta_f]$, then at the start of the η th iteration of the while loop $\mathcal{Q}(\eta)$ is true.*
- 2) *Consider the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ returned by ND-EBF: $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ always satisfy (c) and there exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$ and $\mathbf{P}\hat{\mathbf{X}} = \mathbf{X}$ iff $\eta_f = n$ and $\mathcal{Q}(n+1)$ is true.*
- 3) *Consider the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ returned by ND-EBF and the associated residual $\mathbf{R} = \mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}$: there exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$ and $\mathbf{P}\hat{\mathbf{X}} = \mathbf{X}$ iff $\mathbf{R} = \mathbf{0}_{m \times N}$. As a result, $\mathbf{R} = \mathbf{0}_{m \times N}$ is a sufficient condition to ensure that \mathbf{Y} has a unique, up to permutation, factorisation of the form $\mathbf{Y} = \mathbf{A}\mathbf{X}$, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{X} \in \mathcal{X}_k^{n \times N}$.*

A proof of Lemma IV.1 is provided in Appendix D. In short, the implication of Lemma IV.1 is that as long as $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$, then the reconstructions returned by ND-EBF are accurate up to permutation, and, if they are in addition complete up to permutation, then the factorisation of this form is unique. The following corollary will allow us to focus our analysis on the recovery of \mathbf{A} .

Corollary IV.1.1 (Recovery of \mathbf{A} up to permutation is sufficient to recover both factors up to permutation). *Let $\mathbf{Y} = \mathbf{A}\mathbf{X}$, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{X} \in \mathcal{X}_k^{n \times N}$. Suppose that $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ are the reconstructions of \mathbf{A} and \mathbf{X} returned by either ND-EBF or D-EBF. If there exists a $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$ then $\mathbf{P}\hat{\mathbf{X}} = \mathbf{X}$.*

A proof of Corollary IV.1.1 is provided in Appendix E. To complete this section we provide Lemma IV.2: the key take-away of this lemma is that if ND-EBF successfully recovers the matrix factors of \mathbf{Y} up to permutation, then so to will D-EBF.

Lemma IV.2 (If ND-EBF successfully computes the matrix factors up to permutation then so will D-EBF). *Let $\mathbf{Y} = \mathbf{A}\mathbf{X}$ where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{X} \in \mathcal{X}_k^{n \times N}$. Let $\hat{\mathbf{A}}^{(1)}$ and $\hat{\mathbf{X}}^{(1)}$ denote the reconstructions returned by ND-EBF and $\hat{\mathbf{A}}^{(2)}$ and $\hat{\mathbf{X}}^{(2)}$ the reconstructions returned by D-EBF. If there exists a $\mathbf{P}_1 \in \mathcal{P}^{n \times n}$ such that $\hat{\mathbf{A}}^{(1)}\mathbf{P}_1^T = \mathbf{A}$ and $\mathbf{P}_1\hat{\mathbf{X}}^{(1)} = \mathbf{X}$, then there exists a $\mathbf{P}_2 \in \mathcal{P}^{n \times n}$ such that $\hat{\mathbf{A}}^{(2)}\mathbf{P}_2^T = \mathbf{A}$ and $\mathbf{P}_2\hat{\mathbf{X}}^{(2)} = \mathbf{X}$.*

A proof of Lemma IV.2 is provided in Appendix E.

B. Proof of Theorem II.3

Together, Corollary IV.1.1 and Lemma IV.2 allow us to lower bound the performance D-EBF by analysing the conditions under which ND-EBF recovers a column of \mathbf{A} at each iteration of the while loop, lines 6-31 of Algorithm 13. To

this end, we first lower bound N in order to lower bound with high probability the number of nonzeros per row of X .

Lemma IV.3 (Nonzeros per row in X). *For some $\beta \in \mathbb{N}$ and $\mu > 1$, if $N \geq \beta (\mu \frac{n}{k} \ln(n) + 1)$ then the probability that the random matrix X , as defined in the PSB model Definition II.3, has at least β nonzeros per row is more than $(1 - n^{-(\mu-1)})^\beta$.*

A proof of Lemma IV.3 is given in Appendix E. The proof of this lemma involves analysing a generalisation of the coupon-collector and dixie-cup problems [16], [17], [18]. In the classic coupon collector problem objects are drawn one at a time from a set of size n , typically uniformly at random with replacement until each object has been seen once. The generalisation studied here is equivalent to analysing the number of draws of k objects from $[n]$, uniform at random with replacement, until all objects have been observed β times with high probability. Assuming a lower bound on the number of nonzeros per row of X and that Y is drawn from the PSB model, Lemma IV.4 provides a lower bound on the probability that a column of A is recovered at iteration $\eta \in [n]$ of the while loop, lines 6-31, of ND-EBF.

Lemma IV.4 (Probability that ND-EBF recovers a column at iteration η). *Let $Y = AX$ as per the PSB model, detailed in Definition II.3. In addition, assume that $A \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$, and that each row of X has at least $\beta(n) = (1+2\epsilon d)L(n)$ nonzeros, where $L : \mathbb{N} \rightarrow \mathbb{N}$. Suppose that ND-EBF, Algorithm 13, is deployed to try and compute the factors of Y and that the algorithm reaches iteration $\eta \in [n]$ of the while loop. Then there is a unique column $A_{\ell(\eta)}$ of A , satisfying $A_{\ell(\eta)} \neq \hat{A}_l$ for all $l \in [\eta - 1]$, such that*

$$\mathbb{P}(\hat{A}_\eta = A_{\ell(\eta)}) > 1 - de^{-\tau(n)L(n)},$$

where $\tau(n) := -\ln\left(1 - \left(1 - \frac{b}{\alpha_1 n}\right)^{\alpha_1 n}\right)$ is $\mathcal{O}(1)$.

A proof of Lemma IV.4 is given in Appendix E. With Lemmas IV.3 and IV.4 in place we are ready to prove Theorem II.3. Statements 1 and 2 of Theorem II.3 are immediate consequences of Lemma III.5, therefore all that is left to prove is Statement 3. To provide a brief recap, our objective is to recover up to permutation the random factor matrices A and X , as defined in the PSB model in Definition II.3, from the random matrix $Y = AX$. Our strategy at a high level is as follows: using Lemma IV.2 and Corollary IV.1.1 we lower bound the probability that D-EBF successfully factorises Y up to permutation by lower bounding the probability that ND-EBF recovers A up to permutation. ND-EBF recovers A up to permutation iff at each iteration of the while loop, lines 6-31 of Algorithm 13, a new column of A is recovered. We lower bound the probability of this event in turn using Lemma IV.4. For the proof Theorem II.3 we adopt the following notation: let $\Lambda_{\text{D-EBF}}^*$ and $\Lambda_{\text{ND-EBF}}^*$ denote the events that D-EBF and ND-EBF recover A and X up to permutation respectively, i.e., there exists a random permutation $P \in \mathcal{P}^{n \times n}$ such that $\hat{A}P^T = A$ and $P\hat{X} = X$. Let $\Lambda_0 := \{A \in \mathcal{E}_{k,\epsilon,d}^{m \times n}\} \cap \{\epsilon \leq 1/6\}$ be the event that the random matrix A is the adjacency matrix of a (k, ϵ, d) -expander graph with expansion parameter $\epsilon \leq 1/6$. For $\eta \in [n]$, let Λ_η denote the event that at the end

of η th iteration of the while loop of ND-EBF, there exists a unique column $A_{\ell(\eta)}$ of A satisfying $A_{\ell(\eta)} = \hat{A}_\eta$ and $A_{\ell(\eta)} \neq \hat{A}_l$ for all $l \in [n] \setminus \{\eta\}$. Finally, let Λ_{n+1} denotes the event that each row of X has at least $\beta(n) := (1 + 2\epsilon)dL(n)$ nonzeros per row where $L(n) := \lceil \frac{\gamma \ln(n)}{\tau(n)} \rceil$ and $\gamma > 1$ is a constant.

Proof. Suppose that ND-EBF is deployed instead of D-EBF to recover the matrix factors of Y . From Corollary IV.1.1, for ND-EBF to recover both factors up to permutation it suffices to recover A up to permutation. Clearly then

$$\mathbb{P}(\Lambda_{ND-EBF}^* \mid \Lambda_0) = \mathbb{P}\left(\bigcap_{h=1}^n \Lambda_h \mid \Lambda_0\right).$$

We now apply Bayes' Theorem and condition on Λ_{n+1} ,

$$\begin{aligned} & \mathbb{P}(\Lambda_{ND-EBF}^* \mid \Lambda_0) \\ &= \mathbb{P}\left(\bigcap_{\eta=1}^n \Lambda_\eta \mid \Lambda_0\right) \\ &= \frac{\mathbb{P}\left(\bigcap_{\eta=1}^n \Lambda_\eta, \Lambda_0\right)}{\mathbb{P}(\Lambda_0)} \\ &\geq \frac{\mathbb{P}\left(\bigcap_{\eta=1}^n \Lambda_\eta, \Lambda_0 \mid \Lambda_{n+1}\right) \mathbb{P}(\Lambda_{n+1})}{\mathbb{P}(\Lambda_0)} \\ &= \frac{\mathbb{P}\left(\bigcap_{\eta=1}^n \Lambda_\eta \mid \Lambda_0, \Lambda_{n+1}\right) \mathbb{P}(\Lambda_0 \mid \Lambda_{n+1}) \mathbb{P}(\Lambda_{n+1})}{\mathbb{P}(\Lambda_0)} \\ &= \mathbb{P}\left(\bigcap_{\eta=1}^n \Lambda_\eta \mid \Lambda_0, \Lambda_{n+1}\right) \mathbb{P}(\Lambda_{n+1}) \\ &= \mathbb{P}(\Lambda_{n+1}) \prod_{\eta=1}^n \mathbb{P}\left(\Lambda_\eta \mid \bigcap_{l=0}^{\eta-1} \Lambda_l, \Lambda_{n+1}\right). \end{aligned}$$

In the above, line 2 follows as a result of Bayes Theorem, line 3 is an application of the law of total probability and line 4 is derived using the probability chain rule. The equality on line 5 follows as A and X are drawn independently of one another, therefore, given that Λ_0 is a property of A and Λ_{n+1} a property of X , $\mathbb{P}(\Lambda_0 \mid \Lambda_{n+1}) = \mathbb{P}(\Lambda_0)$. Finally, line 6 is once again an application of the probability chain rule. Assume for now that $N \geq \beta(n) \left(\mu \frac{n}{k} \ln(n) + 1\right)$, where $\mu > 1$ is some constant. Then as a consequence of Lemma IV.3 it follows that

$$\mathbb{P}(\Lambda_{n+1}) > \left(1 - n^{-(\mu-1)}\right)^{\beta(n)}.$$

Applying Lemma IV.4 it also holds for any $\eta \in [n]$ that

$$\mathbb{P}\left(\Lambda_\eta \mid \bigcap_{l=0}^{\eta-1} \Lambda_l, \Lambda_{n+1}\right) > 1 - de^{-\tau(n)L(n)}.$$

where $\tau(n)$ is $\mathcal{O}(1)$. As a result

$$\begin{aligned} & \mathbb{P}(\Lambda_{ND-EBF}^* \mid \Lambda_0) \\ &> \left(1 - n^{-(\mu-1)}\right)^{\beta(n)} \left(1 - de^{-\tau(n)L(n)}\right)^n. \end{aligned}$$

Recalling that $L(n) := \lceil \frac{\gamma \ln(n)}{\tau(n)} \rceil$, where $\gamma > 1$ is an arbitrary constant, then analysing the right-hand factor it follows that

$$\begin{aligned} \left(1 - de^{-\tau(n)L(n)}\right)^n &\geq \left(1 - de^{\ln(1/n^\gamma)}\right)^n \\ &= \left(1 - \frac{d}{n^\gamma}\right)^n \\ &= \sum_{i=0}^{\infty} \binom{n}{i} (-1)^i \left(\frac{d}{n^\gamma}\right)^i \\ &= 1 - \frac{d}{n^{\gamma-1}} + \sum_{i=2}^{\infty} \binom{n}{i} (-1)^i \left(\frac{d}{n^\gamma}\right)^i \\ &= 1 - \mathcal{O}(n^{-\gamma+1}). \end{aligned}$$

Above, the inequality on the first line follows from the fact that $\lceil \frac{\gamma \ln(n)}{\tau(n)} \rceil \geq \frac{\gamma \ln(n)}{\tau(n)}$. The equality on the third line follows by applying the binomial series expansion. Analysing now the left-hand factor derived from Lemma IV.3, note that as $\beta(n) = (1 + 2\epsilon)d \lceil \frac{\gamma \ln(n)}{\tau(n)} \rceil$ and $\tau(n) = \mathcal{O}(1)$ then $\beta(n) = \mathcal{O}(\log(n))$. As a result

$$\begin{aligned} \left(1 - n^{-(\mu-1)}\right)^{\beta(n)} &= \sum_{i=0}^{\infty} \binom{\beta(n)}{i} (-1)^i (n^{-\mu+1})^i \\ &= 1 - \frac{\beta(n)}{n^{\mu-1}} + \sum_{i=2}^{\infty} \binom{\beta(n)}{i} (-1)^i (n^{-\mu+1})^i \\ &= 1 - \mathcal{O}(n^{-\mu+1} \log(n)). \end{aligned}$$

Therefore we arrive at the asymptotic lower bound

$$\begin{aligned} & \mathbb{P}(\Lambda_{ND-EBF}^* \mid \Lambda_0) \\ &> \left(1 - \mathcal{O}(n^{-\mu+1} \log(n))\right) \left(1 - \mathcal{O}(n^{-\gamma+1})\right). \end{aligned}$$

Given that $\epsilon \leq 1/6$ implies $(1 + 2\epsilon) \leq \frac{4}{3}$, $\lceil \frac{\gamma \ln(n)}{\tau(n)} \rceil \leq \frac{2\gamma \ln(n)}{\tau(n)}$, and $\mu, \frac{n}{k} > 1$,

$$\begin{aligned} \beta(n) \left(\mu \frac{n}{k} \ln(n) + 1\right) &= (1 + 2\epsilon)d \lceil \frac{\gamma \ln(n)}{\tau(n)} \rceil \left(\mu \frac{n}{k} \ln(n) + 1\right) \\ &\leq \gamma \mu \frac{8d}{3\tau(n)} \frac{n}{k} (\ln^2(n) + \ln(n)) \\ &\leq N \end{aligned}$$

by construction. Given that $\gamma, \mu > 1$ are arbitrary constants, let $\gamma = \mu$ and define $\nu := \gamma^2$, clearly it must also hold that $\nu > 1$. It follows from Lemma IV.2 that if $N \geq \nu \frac{8d}{3\tau(n)} \frac{n}{k} (\ln^2(n) + \ln(n))$ then

$$\begin{aligned} \mathbb{P}(\Lambda_{D-EBF}^* \mid \Lambda_0) &\geq \mathbb{P}(\Lambda_{ND-EBF}^* \mid \Lambda_0) \\ &> 1 - \mathcal{O}(n^{-\sqrt{\nu}+1} \log(n)) \end{aligned}$$

as claimed. This concludes the proof of Theorem II.3. \square

V. EXPERIMENTS

A key takeaway of Theorem II.3 is that there are parameter regimes in which, at least asymptotically, D-EBF will successfully factorise a matrix Y drawn from the PSB model. In this section we investigate this matter empirically. First, in Section V-A, we highlight certain adjustments to D-EBF which enable it to be deployed in practice: in particular, we discuss how to circumvent the issue of not having access to the expansion parameter ϵ and highlight potential extensions

in which d is unknown and or not fixed. Second, in Section V-B, we conduct experiments demonstrating the efficacy of D-EBF in factorising even mid-sized problems, i.e., $n = 10^3$. We emphasise that in these experiments we sample encoder matrices as per Algorithm 8, Appendix A-A. In particular, we have no guarantees, for a given k , that the encoder matrices deployed are in fact the adjacency matrices of expander graphs with $\epsilon \leq 1/6$. In fact, in most cases it is very likely only some weaker expansion property holds. Despite this D-EBF remains remarkably effective.

A. Deploying D-EBF in practice

As highlighted in Section II-C, in practice the decoder may not have access to the expansion parameter of the encoder matrix. One reason for this, already highlighted in Section II-C, is that computing the expansion parameter is an NP-complete problem. D-EBF is designed for factorisation problems in which $\epsilon \leq 1/6$, we therefore restrict our attention to this setting and consider how the decoder might still be able to succeed in factorising \mathbf{Y} without knowledge of ϵ in advance. To this end, suppose that $1/6$ is used by D-EBF instead of ϵ : as $2\epsilon d \leq d/3$ using this value places a stricter requirement on the frequency with which an entry of the residual must occur in order to be accepted as a candidate singleton value. In addition, as $(1-2\epsilon)d \geq 2d/3$, all singleton values occurring more than $(1-2\epsilon)d$ times in a given column are still extracted. Therefore, from the perspective of extracting singleton values and partial supports, the only implication of using $1/6$ instead of ϵ is that some singleton values and their associated partial supports may be missed.

In regard to clustering however, as $(1-4\epsilon)d \geq d/3$, then using $1/6$ instead of ϵ loses the guarantees afforded by Lemma III.4. To be clear, consider two partial supports \mathbf{w}_1 and \mathbf{w}_2 satisfying $|\text{supp}(\mathbf{w}_1)| > 2d/3$ and $|\text{supp}(\mathbf{w}_2)| > 2d/3$: if $\mathbf{w}_1^T \mathbf{w}_2 \geq d/3 \geq 2\epsilon d$ then we can conclude that these two partial supports originate from the same column, however, $\mathbf{w}_1^T \mathbf{w}_2 < d/3$ does not guarantee these partial supports originate from different columns. If D-EBF is deployed as described in Algorithm 6, this may result in duplicate reconstructions of the same column of \mathbf{A} in $\hat{\mathbf{A}}$. This issue can be overcome by adding in an additional column merge subroutine at the end of each iteration. An example merge subroutine is provided in Algorithm 14, Appendix A-G. In summary, as long as $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$, then in practice, even without access to the expansion parameter ϵ , D-EBF can still be deployed by using the upper bound $1/6$ instead. In fact, the same reasoning applies to any bound on ϵ made available to the decoder in advance. We demonstrate the efficacy of D-EBF with the program expansion parameter set to $1/6$ in Section V-B. Furthermore, at least for certain parameter settings, the performance of the D-EBF seems fairly robust to the underlying choice of ϵ as illustrated in Appendix B-B.

As defined in Algorithm 6, D-EBF requires knowledge of the column sparsity d and number of columns n of \mathbf{A} . We speculate that it may be possible to estimate or use upper and lower bounds on d and therefore still deploy D-EBF without access to d in advance. We also hypothesise that

similar results and algorithms may be derivable when \mathbf{A} is not the adjacency matrix of a fixed degree expander, with instead the degree of each node being bounded in some interval with high probability. We leave the study of such questions to future work and proceed under the assumption that the decoder has access to d in advance. With regard to requiring access to n , we claim this is an artefact of our problem setup rather than a necessity in practice. Indeed, instead of initialising $\hat{\mathbf{A}}$ as an $m \times n$ array of zeros, the reconstruction of \mathbf{A} could instead be kept dynamic, with partial or complete columns of \mathbf{A} being added to $\hat{\mathbf{A}}$ as and when they are recovered. Furthermore, D-EBF can be run in advance of receiving all of the columns of \mathbf{Y} . Consider the situation in which the columns of \mathbf{Y} are received in sequence: as each new column arrives D-EBF can be ran only on this new column, in a so called turnstile model, or in combination with the residuals of previous columns, using the partial reconstruction of \mathbf{A} already acquired as an initialisation point. In either case, after the decoder has seen sufficiently many columns to enable it to recover the encoder matrix up to permutation, then future columns can be decoded efficiently as and when they arrive, using a decoding algorithm from the CCS literature.

To summarise, in addition to \mathbf{Y} , then in practice D-EBF requires access only to the column sparsity d of the encoder in advance. Furthermore, D-EBF can easily be adapted to process subsets of columns of \mathbf{Y} , allowing it to be deployed for example in data streaming contexts.

B. Performance of D-EBF on mid-sized problems

As a result of Lemma II-A and Theorem II.3, then for sufficiently large, sparse problems the assumptions upon which D-EBF is based hold with high probability. However, understanding the regimes in which D-EBF will succeed or fail at a practical level is not clear. Indeed, the probability bounds derived in [19] are loose, making it hard to specify exactly how large n needs to be in order for the encoder \mathbf{A} , sampled from the PSB model, to satisfy $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ with high probability. Additionally, even if this assumption holds the probability bound provided in Theorem II.3 is asymptotic and also likely to be loose. The purpose of this section is to provide a preliminary investigation into the empirical performance of D-EBF, testing the parameter settings for which the algorithm is likely to be successful. In particular, as already highlighted our theory suggests that D-EBF will likely be successful on large, sparse problems and unsuccessful on small, dense ones. We therefore focus our attention on mid-sized problems, varying N and the ratio k/n , in order to better understand the transition point between success and failure.

The parameter settings for our experiments are as follows: $n = 1000$, $m = 800$ and $d = 10$ are kept fixed, with k varied between 1% and 10% of n and $N \in \{100, 200, 300\}$. For each of these parameter settings at least 20 and typically more trials were conducted. For each trial an \mathbf{A} and an \mathbf{X} were generated as per the PSB model, Definition II.3, with the coefficients of \mathbf{X} being sampled from the uniform distribution over an interval bounded away from 0 (in this case $[1, 10]$). For each trial D-EBF was deployed to recover \mathbf{A} and \mathbf{X} up

to permutation from $Y = AX$ with the ϵ program input parameter set to $1/6$. We note here that in our experiments D-EBF, Algorithm 6, was not amended as per the discussion in Section V-A to include an additional merge subroutine as detailed in Algorithm A-G, Appendix A-G. This choice was taken for reasons of clarity and consistency with our theoretical discussion of D-EBF. With \hat{A} and \hat{X} denoting the reconstructions returned by D-EBF, then for each trial the Frobenius norm of the residual $\|Y - \hat{A}\hat{X}\|_F$ was computed as a percentage of $\|Y\|_F$. In addition, for trial the number of iterations of the while loop of D-EBF, lines 6-25 of Algorithm 6, was counted and likewise averaged over the 10 trials. The outcomes of these experiments are shown in Figures 2 and 3.




Fig. 2. Plot shows the percentage $\|Y - \hat{A}\hat{X}\|_F / \|Y\|_F$ error for each trial generated as per the PSB model, Definition II.3. Here \hat{A} and \hat{X} denote the reconstructions returned by D-EBF. The line plots represent the average performance while the opaque coloured regions indicate the 95% confidence interval. The parameters k and N are varied while $n = 1000$, $m = 800$, $d = 10$ are fixed. In all trials the expansion parameter ϵ of A is not known or computed in advance and D-EBF is deployed using $1/6$ in its stead.

In Figure 2 we observe that Y is most likely to be successfully factorized when N is large and k/n is neither small but not too small. Indeed, a larger N will typically imply more partial supports extracted at each iteration of the while loop and therefore a greater chance of recovering A up to permutation. This in turn allows for the recovery of X up to permutation. Similar to the situation in which N is small, when k/n is small fewer partial supports are extracted per column of Y at each iteration of the while loop, making the recovery of A harder. This results in the incomplete factorisation of Y for $k/n < 3\%$. Note also for $k/n < 3\%$, that as N increases Y on average becomes closer to being fully factorized: this is because the partial supports extracted from

the additional columns compensate for the reduced number of partial supports extracted per column. We would expect even for small k/n that as N increases Y will converge towards being fully factorized. When k/n is large a similar issue occurs in terms of few or even no partial supports being extracted per column of Y . However, this is due to the fact that with m and n fixed then as k increases the probability that $A \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ converges towards 0. When this assumption is not satisfied it is likely no longer possible to identify singleton values using frequency of occurrence, let alone cluster their associated partial supports. This can be observed in the failure to make any progress in factorising Y when $k/n = 10\%$. Furthermore, unlike the former situation, increasing N in this context is unlikely to yield any improvement. The middle ground, $3\% \leq k/n \leq 8\%$, appears ripe for factorisation, striking the right balance between k being large enough so that each column of A contributes to many of the columns of Y , while being sufficiently small so that singleton values can still be identified via the frequency with which they appear. We note also that the distribution of reconstruction errors is tightly concentrated around its mean outside of the transition region $7\% < k/n < 9\%$. Furthermore, outside this transition region it appears that D-EBF will either consistently succeed or consistently fail respectively. Inside this transition region there is less consistency, with D-EBF making progress in some trials but little in others. One may also observe that trials in which $N = 300$ counterintuitively seem to perform less well than those in which $N = 200$ in this region: this is explained by the fact that as soon as duplicate columns are identified D-EBF terminates. Such duplicates appear sooner for problems with large N and large k/n . If one were to include the aforementioned merge subroutine then this issue would be avoided.

Turning our attention to the Figure 3, we observe that as N increases then the number of iterations of the while loop of D-EBF generally decreases. This is to be expected as a larger N means more partial supports extracted at each iteration, and therefore we would expect A to be recovered in fewer iterations. Observe for $N = 100$ and $k/n = 7\%$ that the iterative approach of D-EBF is invaluable in being able to factorise Y . To summarise, these experiments demonstrate that it is possible to practically decode a set of linear measurements without access to the full encoder matrix. While Theorem II.3 provides asymptotic guarantees, these experiments illustrate that D-EBF is successful even on relatively small to medium sized problems. We expect for large problems the performance of D-EBF to improve, allowing for problems with larger k/n ratios to be factorized successfully.

In Appendix B we provide two more experiments: in Figure 4, Appendix B-A, we show over 100 trials, in which X is fixed and A is resampled each time, that, at least under the parameter settings considered, reconstruction errors are still tightly concentrated around 0. This implies that the ability to compress the sparse codes and accurately reconstruct them using D-EBF is fairly agnostic to the specific choice of encoder matrix used. Additionally, in Figure 5, Appendix B-B, we show over 100 trials the spread of reconstruction errors for different ϵ program parameters. Under the parameter setting




Fig. 3. Plot shows the number of iterations the while loop of D-EBF, lines 6-25, required for convergence over at least 20 trials each generated as per the PSB model, Definition II.3. Here $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ denote the reconstructions returned by D-EBF. Lines denote the average number of iterations over the trials, the opaque coloured regions indicate the 95% confidence interval. The parameters k and N are varied while $n = 1000$, $m = 800$, $d = 10$ are fixed. In all trials the expansion parameter ϵ of A is not known or computed in advance, D-EBF is deployed using $1/6$ in its stead.

considered, D-EBF appears fairly robust to the choice of ϵ . We leave a more comprehensive empirical study of the D-EBF algorithm to future work.

VI. CONCLUDING REMARKS AND POTENTIAL EXTENSIONS

In this paper we studied the feasibility of performing multi-measurement combinatorial compressed sensing in the setting where the decoder does not have access to the encoder matrix. To this end we presented a particular random encoder matrix, see Algorithm 8, and the Decoder-Expander Based Factorisation algorithm, summarised in Algorithm 6. We proved that the combination of these two allows for the successful factorisation of a measurement matrix Y drawn from PSB model, thereby recovering both the encoder and sparse coding matrices up to permutation, with high probability and with near optimal sample complexity. In particular, the decoder only needs knowledge of the column sparsity d of the encoder matrix in advance. We further illustrated how this combination also enables recovery of the location information by applying an additional column ordering procedure on the encoder matrix, agreed on in advance, both prior to encoding and after decoding. This condition can be viewed as a property of the distribution from which the encoder matrix is sampled. We then demonstrated that tD-EBF performs well in practice even on mid-sized problems. In addition, the per while loop iteration

computational complexity of D-EBF is $\mathcal{O}(k^2(n+N)N)$: this, combined with experimental evidence indicating only a modest number of iterations are required for convergence, as well as ample opportunities for parallelization, suggest that the computational cost of D-EBF at scale is very manageable. Some potential extensions of this work are as follows.

- **Relaxing the $\epsilon \leq 1/6$ condition:** Lemma III.4 guarantees that partial supports can be trivially clustered if the binary factor matrix A is a (k, ϵ, d) -expander graph with expansion parameter $\epsilon \leq 1/6$. This expansion parameter is guaranteed over all $\binom{n}{k}$ sets of k columns in A . As Y contains at most N such sets of k columns, and as N would typically be exponentially smaller than this binomial coefficient, it is expected that similar bounds could be derived under a model for which it holds only with high probability that a subset of k columns of A satisfies the expansion bound.
- **Relaxing the dissociated condition condition:** similar to the the first point, it may suffice if most if not all possible subsets of the nonzeros of a sparse code are dissociated. However, adversarial choices of nonzero values in X may require combinatorial searches on certain entries of Y and could result in Y no longer having a unique factorisation.
- **Stability to noise and projections of arbitrary sparse matrices onto the set of PSB matrices:** more developed matrix factorizations, such as PCA, subspace clustering, and dictionary learning, are all known to be effective for matrices which only approximately achieve the modelling assumptions in context. An extension of the PSB model would be to show that the factors are stable to additive noise and remain efficient and reliable to compute. Stability to additive noise one might expect to be achieved via related work in combinatorial compressed sensing, such as [52]. This could be further augmented by having repeated measurement vectors contained in Y . In order for the PSB model to be more generally applicable, approximation rates are required. In particular, a derivation of the distance between an arbitrary sparse matrix \mathbf{Y} and its projection onto the set of PSB matrices, which would presumably show better approximation as the number of columns n of the encoder matrix \mathbf{A} is increased. To be clear, this would entail bounds of the form $\|\mathbf{Y} - P_{PSB(n)}(\mathbf{Y})\| \leq f(n)$, where $P_{PSB(n)}(\cdot)$ projects to the nearest PSB matrix with parameters k, m, n, d, N under some matrix norm, and $f(n)$ is a rapidly decreasing function of n and or potentially d and k . One would also expect expressivity benefits from relaxing the fixed degree condition, instead allowing the number of nonzeros per column of A to vary between an upper and lower bound.

ACKNOWLEDGEMENTS

The authors would like to thank Robert Calderbank for his helpful comments and advice.

APPENDIX A

DETAILED DEFINITIONS OF ALGORITHMS

A. GENERATE-ENCODER

Algorithm 8 GENERATE-ENCODER(m, n, d)

```

1:  $A \leftarrow \text{zeros}(m, n)$ 
2:  $\alpha \leftarrow \lfloor m/d \rfloor$ 
3:  $\beta \leftarrow \lceil n/\alpha \rceil$ 
4:  $l \leftarrow 1$ 
5: for  $i = 1 : \beta$  do
6:    $Q \sim U(\pi([m]))$ 
7:    $\omega \leftarrow \min\{n - l + 1, \alpha\}$ 
8:   for  $j = 1 : \omega$  do
9:      $c \leftarrow (j - 1)d$ 
10:    for  $r = 1 : d$  do
11:       $A_{Q_{c+r}, l} \leftarrow 1$ 
12:    end for
13:     $l \leftarrow l + 1$ 
14:  end for
15: end for
16: Return  $A$ 

```

Algorithm 8 takes as inputs the dimensions of the encoder m and n , the column sparsity d , and returns a binary matrix. The variable α stores the number of columns whose supports can be assigned using a single permutation of $[m]$, β is the number of permutations that are needed to be drawn in order to assign a support of cardinality d to each column of A , ω is the number of columns of A to assign a support to using the current permutation of $[m]$, c denotes the entry of the permutation vector from which to start assigning a support, and finally l is the index of the column next in line to be assigned a support. On line 6 a permutation of the set $[n]$ is drawn uniformly at random, note here that $Q \in [n]^n$ is a random vector holding this permutation and Q_i is the i th element of this vector. Note also that draws of different permutations are considered to be mutually independent. The for loop on lines 8-14 then uses the permutation Q to assign a support to columns l through to $l + \omega - 1$ of A . This is then repeated by drawing another permutation until all columns of A have been assigned a support.

B. EXTRACT&MATCH

Consider the column vector $\mathbf{y} = \mathbf{A}\mathbf{x}$ and suppose that $\hat{\mathbf{A}}$ and $\hat{\mathbf{x}}$ are the current reconstructions of \mathbf{A} and \mathbf{x} . Algorithm 9 processes the residual $\mathbf{r} = \mathbf{y} - \hat{\mathbf{A}}\hat{\mathbf{x}}$, taking as input arguments \mathbf{r} , the dimension m of \mathbf{r} , the expansion parameter ϵ , the column sparsity d , as well as the current reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{x}}$. The algorithm identifies singleton values and partial supports, either matching them with a nonzero column of $\hat{\mathbf{A}}$ or identifying them as belonging to a column of \mathbf{A} not yet observed. To this end Algorithm 9 returns the unmatched singleton values and partial supports, stored as entries of the vector \mathbf{q} and columns of the matrix \mathbf{W} respectively, the number of unmatched partial supports p , the updated reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{x}}$ and a Boolean variable UPDATED, which indicates whether or not the reconstructions have in fact been updated.

The outer for loop and subsequent if statement, lines 5 and 6 of Algorithm 9 respectively, iterate through each nonzero

Algorithm 9 EXTRACT&MATCH($\mathbf{r}, m, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{x}}$)

```

1:  $\mathbf{q} \leftarrow []$ 
2:  $\mathbf{W} \leftarrow []$ 
3:  $p \leftarrow 0$ 
4: UPDATED  $\leftarrow$  FALSE
5: for  $i = 1 : m$  do
6:   if  $r_i \neq 0$  then
7:      $c \leftarrow 1$ 
8:      $\mathbf{w} \leftarrow \text{zeros}(m)$ 
9:      $w_i \leftarrow 1$ 
10:    for  $j = (i + 1) : m$  do
11:      if  $r_i = r_j$  then
12:         $w_j \leftarrow 1$ 
13:         $r_j \leftarrow 0$ 
14:         $c \leftarrow c + 1$ 
15:      end if
16:    end for
17:    if  $c > (1 - 2\epsilon)d$  then
18:       $\mathbf{t} \leftarrow \hat{\mathbf{A}}^T \mathbf{w}$ 
19:       $\kappa \leftarrow \text{argmax}_{k \in [n]} \{t_k\}$ 
20:      if  $t_\kappa \geq 2\epsilon d$  then
21:         $\hat{\mathbf{x}}_\kappa \leftarrow r_i$ 
22:         $\hat{\mathbf{a}}_\kappa \leftarrow \sigma(\hat{\mathbf{a}}_\kappa + \mathbf{w})$ 
23:        UPDATED  $\leftarrow$  TRUE
24:      else
25:         $\mathbf{q} \leftarrow [\mathbf{q}^T; r_i]$ 
26:         $\mathbf{W} \leftarrow [\mathbf{W}; \mathbf{w}]$ 
27:         $p \leftarrow p + 1$ 
28:      end if
29:    else
30:      for  $l \in \text{supp}(\hat{\mathbf{x}})$  do
31:        if  $r_i = \hat{x}_l$  then
32:           $\hat{\mathbf{a}}_l \leftarrow \sigma(\hat{\mathbf{a}}_l + \mathbf{w})$ 
33:          UPDATED  $\leftarrow$  TRUE
34:        end if
35:      end for
36:    end if
37:  end if
38: end for
39: Return  $\mathbf{q}, \mathbf{W}, p, \hat{\mathbf{A}}, \hat{\mathbf{x}}, \text{UPDATED}$ 

```

entry of the input vector \mathbf{r} to check for singleton values: this is performed by the inner for loop, starting on line 10, which identifies the locations in \mathbf{r} where the current entry being checked appears. The variable c counts the number of appearances of the current entry; if $c > (1 - 2\epsilon)d$ then it is presumed that the current entry is a singleton value. Subsequently, on lines 18 and 19 the column $\hat{\mathbf{a}}_\kappa$ of $\hat{\mathbf{A}}$ whose support overlaps most with the associated candidate partial support is identified. Considering line 20, then if the overlap is sufficiently large, i.e., at least $2\epsilon d$ as per Lemma III.4, then the partial support is assumed to originate from the identified column $\hat{\mathbf{a}}_\kappa$ and the reconstructions are updated accordingly as per lines 21 and 22. Note here that $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ denotes the unit threshold function, meaning $\sigma(x) = 1$ for all $x \geq 1$ and is 0 otherwise. Note also, guided by Lemma III.4 and given that

$\hat{\mathbf{A}}$ may contain incomplete columns⁴, that Algorithm 9 only attempts to match partial supports with cardinality larger than $(1 - 2\epsilon)d$ to a column of $\hat{\mathbf{A}}$. If the overlap is not sufficient, then, again guided by Lemma III.4, we may conclude that the partial support under consideration does not match any existing nonzero column of the reconstruction. As a result, on lines 25 and 26 the current entry of the residual and its associated candidate partial support are added to \mathbf{q} and \mathbf{W} respectively. On line 27 the variable p , used to count the number of unmatched singleton values extracted from \mathbf{r} , is updated accordingly. If $c \leq (1 - 2\epsilon)d$ then the for loop starting on line 30 checks if the entry under consideration corresponds to a previously identified singleton value, now stored in $\hat{\mathbf{x}}$. If so, then on line 32 the partial support associated with the current entry is used to update the appropriate column of $\hat{\mathbf{A}}$, regardless of its size.

In regard to the computational complexity of Algorithm 9, if the sparsity of \mathbf{r} , $\hat{\mathbf{A}}$ and $\hat{\mathbf{x}}$ is not taken advantage of then the for loops starting on lines 5 and 10 iterate through $\mathcal{O}(m)$ entries, and the for loop starting on line 30 through $\mathcal{O}(n)$ entries. As the latter two for loops are nested inside the first, and $m = \mathcal{O}(n)$, then this corresponds to $\mathcal{O}(mn)$ for loop iterations. Based on the matrix vector product on line 18, each iteration has a cost of $\mathcal{O}(mn)$, therefore the total cost of Algorithm 9 is $\mathcal{O}(m^2n^2)$. If the sparsity of \mathbf{r} , $\hat{\mathbf{A}}$ and $\hat{\mathbf{x}}$ is taken advantage of, then as \mathbf{r} has at most kd nonzero entries and d is constant, the for loops starting on lines 5, 10 and 30 each iterate through $\mathcal{O}(k)$ nonzero entries. Again, as the latter two for loops are nested inside the first this corresponds to $\mathcal{O}(k^2)$ for loop iterations. Based on the sparse matrix vector product on line 18, each iteration has a cost of $\mathcal{O}(n)$, therefore when sparsity is leveraged the total cost of Algorithm 9 is $\mathcal{O}(k^2n)$.

C. CLUSTER&ADD

Algorithm 10 CLUSTER&ADD($\mathbf{q}, \mathbf{W}, p, \eta, \mathbf{h}, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{X}}$)

```

1:  $\mathbf{b} \leftarrow \text{zeros}(p)$ 
2:  $\mathbf{G} \leftarrow \mathbf{W}^T \mathbf{W}$ 
3: for  $i = 1 : p$  do
4:   if  $b_i = 0$  then
5:      $\hat{\mathbf{a}}_\eta \leftarrow \mathbf{w}_i$ 
6:      $\hat{x}_{\eta, h_i} \leftarrow q_i$ 
7:      $b_i \leftarrow 1$ 
8:     for  $j = (i + 1) : p$  do
9:       if  $g_{i,j} \geq 2\epsilon d$  and  $b_j = 0$  then
10:         $\hat{\mathbf{a}}_\eta \leftarrow \sigma(\hat{\mathbf{a}}_\eta + \mathbf{w}_j)$ 
11:         $\hat{x}_{\eta, h_j} \leftarrow q_j$ 
12:         $b_j \leftarrow 1$ 
13:       end if
14:     end for
15:      $\eta \leftarrow \eta + 1$ 
16:   end if
17: end for
18: Return  $\hat{\mathbf{A}}, \hat{\mathbf{X}}, \eta$ 

```

⁴Observe we can only guarantee that a nonzero column of $\hat{\mathbf{A}}$ has more than $(1 - 2\epsilon)d$ nonzeros.

The p dimensional binary vector \mathbf{b} , initialised as the zero vector on line 1 of Algorithm 10, is used to indicate whether or not a partial support has been assigned to a column of $\hat{\mathbf{A}}$, one indicating true and zero false. The entries of the gram matrix \mathbf{G} , calculated on line 2, are used to cluster the unmatched partial supports. The outer for loop, starting on line 3, first checks on line 4 if the current partial support under consideration has been used to update a column of the reconstruction already. If it hasn't, then on line 5 it is used as the basis for a new nonzero column $\hat{\mathbf{a}}_\eta$ of $\hat{\mathbf{A}}$. On line 6 the corresponding singleton value is then used to update the appropriate entry in $\hat{\mathbf{X}}$. The inner for loop, starting on line 8, checks the inner product of the partial support in question with all other partial supports that have not already been processed, as per the if statement on line 9. On line 10, any partial support that overlaps the new nonzero column $\hat{\mathbf{a}}_\eta$ sufficiently, as per Lemma III.4, and which has not yet been assigned to a column, is then used to update $\hat{\mathbf{a}}_\eta$. On line 11 the corresponding singleton values are used to update the appropriate entries in $\hat{\mathbf{X}}$. Once this process is complete, then on line 15 η is updated accordingly, with new nonzero columns being introduced to $\hat{\mathbf{A}}$ from left to right.

D. DECODE

Algorithm 11 attempts to decode a column vector $\mathbf{y} = \mathbf{A}\mathbf{x}$, taking as inputs \mathbf{y} and its dimension m , the expansion parameter ϵ and column sparsity d of \mathbf{A} , and the corresponding reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{x}}$ of \mathbf{A} and \mathbf{x} . The algorithm returns an updated reconstruction of the sparse code $\hat{\mathbf{x}}$ and a Boolean variable UPDATED, which indicates whether or not an update to $\hat{\mathbf{x}}$ has actually occurred. Algorithm 11 operates in a manner similar to Algorithm 9, the key difference being that the partial supports are compared with and matched to only the complete columns of $\hat{\mathbf{A}}$, which are identified on line 1. As in the case of Algorithm 9, the role of the nested for loops starting on lines 8 and 13 is to identify and extract singleton values and partial supports. However, as the columns of $\hat{\mathbf{A}}_S$ are complete, then the required overlap in support as per line 20 is only $2\epsilon d$: this follows from the fact that any partial support \mathbf{w} , irrespective of its size, originates from a column \mathbf{a}_l of \mathbf{A} if $\mathbf{w}^T \mathbf{a}_l \geq 2\epsilon d$. The Boolean variable UPDATED is used to track whether or not at any point an update to $\hat{\mathbf{x}}$ has occurred. The Boolean variable RUN is used to determine whether the while loop running from lines 4 to 32 should iterate: this occurs as long as a new nonzero is added to $\hat{\mathbf{x}}$ on line 24. As a result, at each iterate prior to the termination iterate, then on line 6 the contribution of at least one column of $\hat{\mathbf{A}}_S$ is removed from \mathbf{y} . This update to the residual potentially reveals new singleton values, which can be matched with another column of $\hat{\mathbf{A}}_S$, allowing for further updates to $\hat{\mathbf{x}}$. If $\hat{\mathbf{x}}$ is not updated during the previous iterate, then the residual at the current and previous iterates will be the same. Therefore no progress can be made and so the algorithm terminates.

Without leveraging sparsity, the outer for loop starting on line 8 iterates through $\mathcal{O}(m)$ elements. Each iteration involves an inner for loop, starting on line 13, which also iterates through $\mathcal{O}(m)$ elements, and a matrix vector product on

Algorithm 11 DECODE($\mathbf{y}, m, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{x}}$)

```

1:  $\mathcal{S} \leftarrow \{l \in [n] : |\text{supp}(\mathbf{a}_l)| = d\}$ 
2: UPDATED $\leftarrow$ FALSE
3: RUN $\leftarrow$ TRUE
4: while RUN = TRUE do
5:   RUN $\leftarrow$ FALSE
6:    $\mathbf{r} \leftarrow \mathbf{y} - \hat{\mathbf{A}}\hat{\mathbf{x}}$ 
7:   if  $\mathbf{r} \neq \mathbf{0}_m$  then
8:     for  $i = 1 : m$  do
9:       if  $r_i \neq 0$  then
10:         $c \leftarrow 1$ 
11:         $\mathbf{w} \leftarrow \text{zeros}(m)$ 
12:         $w_i \leftarrow 1$ 
13:        for  $j = (i + 1) : m$  do
14:          if  $r_i = r_j$  then
15:             $w_j \leftarrow 1$ 
16:             $r_j \leftarrow 0$ 
17:             $c \leftarrow c + 1$ 
18:          end if
19:        end for
20:        if  $c \geq 2\epsilon d$  then
21:           $\mathbf{t} \leftarrow \hat{\mathbf{A}}_{\mathcal{S}}^T \mathbf{w}$ 
22:           $\kappa \leftarrow \text{argmax}_{k \in \mathcal{S}} \{t_k\}$ 
23:          if  $t_\kappa \geq 2\epsilon d$  then
24:             $\hat{x}_\kappa \leftarrow r_i$ 
25:            UPDATED $\leftarrow$ TRUE
26:            RUN $\leftarrow$ TRUE
27:          end if
28:        end if
29:      end if
30:    end for
31:  end if
32: end while
33: Return  $\hat{\mathbf{x}}, \text{UPDATED}$ 

```

line 21. The cost of each of these is $\mathcal{O}(m)$ and $\mathcal{O}(mn)$ respectively, therefore overall the cost per iteration of the while loop is $\mathcal{O}(m^2n)$. As $\hat{\mathbf{x}}$ has $\mathcal{O}(k)$ nonzeros then the while loop starting on line 4 iterates $\mathcal{O}(k)$ times. Therefore, without taking advantage of sparsity, the total computational cost of Algorithm 11 is $\mathcal{O}(km^2n)$. However, if \mathbf{r} , $\hat{\mathbf{A}}_{\mathcal{S}}$ and $\hat{\mathbf{x}}$ are stored using a sparse format, then the per iteration cost of the while loop is $\mathcal{O}(kn)$, resulting in a reduced total cost of $\mathcal{O}(k^2n)$.

E. MAXCLUSTER&ADD

In order to define the ND-EBF algorithm we must introduce the subroutine detailed in Algorithm 12. This subroutine plays a role analogous to that of Algorithm 10 in D-EBF. The key difference between these two subroutines is that while Algorithm 10 attempts to cluster all partial supports in \mathbf{W} , and then use each of these clusters to reconstruct a column of \mathbf{A} , Algorithm 12 attempts to identify only the largest cluster of partial supports in \mathbf{W} , and then use this one cluster to reconstruct a single column of \mathbf{A} . Algorithm 12 takes as inputs the following: the unmatched singleton values, stored in the entries of the vector $\mathbf{q} \in \mathbb{R}^p$, the unmatched partial supports

stored in the columns of the matrix $\mathbf{W} \in \{0, 1\}^{m \times p}$ and ordered so that \mathbf{w}_i is the partial support associated with the singleton q_i , the number of unmatched singleton values p , the index $\eta := \min\{l \in [n] : \hat{\mathbf{a}}_l = \mathbf{0}_m\}$, the vector $\mathbf{h} \in [N]^p$ which stores the indices of the columns of \mathbf{R} from which each singleton value was extracted, i.e., h_i is the column index of \mathbf{R} from which q_i was extracted, the expansion parameter ϵ and column sparsity d of \mathbf{A} , and the current reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$. The algorithm returns updated reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$: as can be observed at most one new nonzero column of $\hat{\mathbf{A}}$ and row of $\hat{\mathbf{X}}$ are updated.

Algorithm 12 MAXCLUSTER&ADD($\mathbf{q}, \mathbf{W}, p, \eta, \mathbf{h}, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{X}}$)

```

1:  $\mathbf{b} \leftarrow \text{zeros}(p)$ 
2:  $\mathbf{G} \leftarrow \mathbf{W}^T \mathbf{W}$ 
3:  $\mathcal{C}^* \leftarrow \emptyset$ 
4: for  $i = 1 : p$  do
5:   if  $b_i = 0$  then
6:      $\mathcal{C} \leftarrow \{i\}$ 
7:     for  $j = (i + 1) : p$  do
8:       if  $g_{i,j} \geq 2\epsilon d$  and  $b_j = 0$  then
9:          $\mathcal{C} \leftarrow \mathcal{C} \cup \{j\}$ 
10:      end if
11:    end for
12:    if  $|\mathcal{C}| > |\mathcal{C}^*|$  then
13:       $\mathcal{C}^* \leftarrow \mathcal{C}$ 
14:    end if
15:  end if
16: end for
17: for  $i \in \mathcal{C}^*$  do
18:    $\hat{\mathbf{a}}_\eta \leftarrow \sigma(\hat{\mathbf{a}}_\eta + \mathbf{w}_i)$ 
19:    $\hat{x}_{\eta, h_i} \leftarrow q_i$ 
20: end for
21: Return  $\hat{\mathbf{A}}, \hat{\mathbf{X}}$ 

```

The p dimensional binary vector \mathbf{b} , initialised as the zero vector on line 1, is used to indicate whether or not a partial support has been assigned to a column of $\hat{\mathbf{A}}$, one indicating true and zero false. The entries of the gram matrix \mathbf{G} , calculated on line 2, are used to cluster the unmatched partial supports. On line 3 the index set $\mathcal{C}^* \subseteq [p]$ is initialised as empty and is used to denote the cluster of partial supports in \mathbf{W} with the largest cardinality identified so far. The outer for loop, starting on line 4, first checks, as per the if statement on line 5, if the partial support \mathbf{w}_i currently under consideration has been used to update a column of the reconstruction already. If this is not the case, then on line 6 the index set $\mathcal{C} \subseteq [p]$, which is used to denote the set of partial supports in \mathbf{W} which originate from the same column of \mathbf{A} as \mathbf{w}_i , is initialised. The inner for loop, lines 7-11, then identifies and finds the indices of any such partial supports. On lines 12 and 13, if $|\mathcal{C}| > |\mathcal{C}^*|$ then a larger cluster than \mathcal{C}^* has been identified and so \mathcal{C}^* is updated accordingly. Finally, after the completion of the outer for loop then, on lines 18 and 19, the partial supports and singleton values belonging to the largest cluster \mathcal{C}^* are used to update just the η th column of $\hat{\mathbf{A}}$ and row of $\hat{\mathbf{X}}$ respectively.

F. ND-EBF

Algor

Algorithm 13 ND-EBF($\mathbf{Y}, m, n, N, d, \epsilon$)

```

1:  $\hat{\mathbf{A}} \leftarrow \text{zeros}(m, n)$ 
2:  $\hat{\mathbf{X}} \leftarrow \text{zeros}(n, N)$ 
3:  $\mathbf{R} \leftarrow \mathbf{Y}$ 
4:  $\eta \leftarrow 1$ 
5: RUN←TRUE
6: while  $\eta \leq n$  and RUN=TRUE do
7:    $\mathbf{q} \leftarrow \emptyset$ 
8:    $\mathbf{W} \leftarrow \emptyset$ 
9:    $\mathbf{h} \leftarrow \emptyset$ 
10:   $p \leftarrow 0$ 
11:  for  $i = 1 : N$  do
12:     $\mathbf{u}, \mathbf{V}, c, \_, \hat{\mathbf{x}}_{i, \_}$ 
       $\leftarrow \text{EXTRACT\&MATCH}(\mathbf{r}_i, m, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{x}}_i)$ 
13:     $\mathbf{q} \leftarrow [\mathbf{q}^T; \mathbf{u}^T]$ 
14:     $\mathbf{W} \leftarrow [\mathbf{W}; \mathbf{V}]$ 
15:     $\mathbf{h} \leftarrow [\mathbf{h}, i \times \text{zeros}(c)]$ 
16:     $p \leftarrow p + c$ 
17:  end for
18:  if  $p > 0$  then
19:     $\hat{\mathbf{A}}, \hat{\mathbf{X}}$ 
       $\leftarrow \text{MAXCLUSTER\&ADD}(\mathbf{q}, \mathbf{W}, p, \eta, \mathbf{h}, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{X}})$ 
20:    if  $|\text{supp}(\hat{\mathbf{a}}_\eta)| = d$  then
21:      for  $i = 1 : N$  do
22:         $\hat{\mathbf{x}}_{i, \_} \leftarrow \text{DECODE}(\mathbf{y}_i, m, N, \epsilon, d, \hat{\mathbf{A}}_S, \hat{\mathbf{x}}_i)$ 
23:      end for
24:       $\mathbf{R} \leftarrow \mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}$ 
25:       $\eta \leftarrow \eta + 1$ 
26:    else
27:      RUN←FALSE
28:    end if
29:  else
30:    RUN←FALSE
31:  end if
32: end while
33: Return  $\hat{\mathbf{A}}, \hat{\mathbf{X}}$ 

```

As referenced to earlier, the way Algorithm 6 operates is very similar to that of D-EBF, the key difference being that while D-EBF maximally utilises all partial supports and singleton values extracted at each iteration, ND-EBF seeks only to utilise those associated with the largest cluster. The extraction of singleton values and partial supports, lines 11-17, is done in exactly the same fashion as in D-EBF. If at least one singleton value and partial support are extracted, i.e., $p > 0$, then on line 19 Algorithm 12 is deployed in order to recover a single column of \mathbf{A} from the largest cluster of partial supports in \mathbf{W} . If this reconstruction is not complete, meaning the condition on line 20 is not satisfied, then the algorithm terminates. In order to recover further entries of \mathbf{X} and grow the support of $\hat{\mathbf{X}}$, then on line 22 Algorithm 11, or some other combinatorial decoding algorithm, is deployed. Given the updates to both $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$, the residual is then recomputed on line 24 in preparation for singleton value and partial support

extraction and clustering at the next iteration. Analogous to Lemma III.5 we provide the following accuracy guarantees for ND-EBF.

G. MERGE

As discussed in Section V, overestimating the expansion parameter ϵ by using $1/6$ instead can cause issues in practice during the clustering stage of D-EBF. Recalling the example provided in Section V, consider two partial supports \mathbf{w}_1 and \mathbf{w}_2 satisfying $|\text{supp}(\mathbf{w}_1)| > 2d/3$ and $|\text{supp}(\mathbf{w}_2)| > 2d/3$. If $\mathbf{w}_1^T \mathbf{w}_2 \geq d/3 \geq 2\epsilon d$ then we can conclude that these two partial supports originate from the same column, however $\mathbf{w}_1^T \mathbf{w}_2 < d/3$ does not guarantee that these partial supports originate from different columns. If D-EBF is deployed as described in Algorithm 6, this may result in duplicate reconstructions of the same column of \mathbf{A} in $\hat{\mathbf{A}}$.

The purpose of the Algorithm 14 is to merge duplicate columns in $\hat{\mathbf{A}}$, as well as the corresponding rows of $\hat{\mathbf{X}}$, by checking the inner product between columns of $\hat{\mathbf{A}}$. Assuming that $\epsilon \leq 1/6$, then, if any pair of columns has an inner product larger than $d/3$ and given that $d/3 \geq 2\epsilon d$, we may conclude that these two columns and their respective rows in $\hat{\mathbf{X}}$ should be combined.

Algorithm 14 MERGE($d, \hat{\mathbf{A}}, \hat{\mathbf{X}}$)

```

1:  $\mathbf{G} \leftarrow \hat{\mathbf{A}}^T \hat{\mathbf{A}}$ 
2: for  $i = 1 : n$  do
3:   if  $|\text{supp}(\hat{\mathbf{a}}_i)| \geq d/3$  then
4:     for  $j = (i+1) : n$  do
5:       if  $g_{i,j} \geq d/3$  then
6:          $\hat{\mathbf{a}}_i \leftarrow \sigma(\hat{\mathbf{a}}_i + \hat{\mathbf{a}}_j)$ 
7:          $\hat{\mathbf{a}}_j \leftarrow \mathbf{0}_m$ 
8:          $\hat{\mathbf{x}}_i \leftarrow \hat{\mathbf{x}}_i + \tilde{\mathbf{x}}_j$ 
9:          $\hat{\mathbf{x}}_j \leftarrow \mathbf{0}_N^T$ 
10:       end if
11:     end for
12:   end if
13: end for
14: Return  $\hat{\mathbf{A}}, \hat{\mathbf{X}}$ 

```

APPENDIX B
OTHER EXPERIMENTSA. Assessing reconstruction accuracy of D-EBF: fixed \mathbf{X}

Figure 4 illustrates that, at least in the parameter regime tested, the distribution of reconstruction errors for reconstructions returned by D-EBF across 100 trials, where in all trials the sparse coding matrix \mathbf{X} is fixed but at each trial the encoder matrix is resampled, is tightly concentrated and close to 0. We note that these results are in close agreement with those displayed in Figure 2.

B. Assessing reconstruction accuracy: varying ϵ

Figure 5 suggests that, at least in the parameter regime tested, there is happily a welcome robustness to the choice of

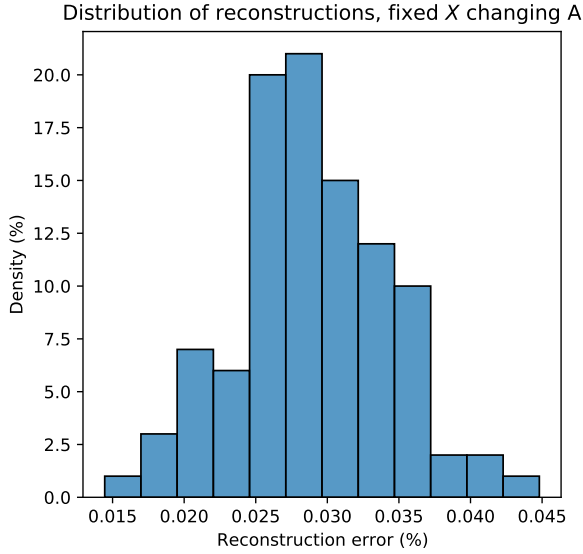


Fig. 4. Distribution of reconstruction errors $\|\mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}\|_F^2 / \|\mathbf{Y}\|_F^2$ of reconstructions $\hat{\mathbf{A}}$, $\hat{\mathbf{X}}$ returned by D-EBF with \mathbf{X} fixed and \mathbf{A} resampled in each trial. Parameter settings for experiment: number of trials $T=100$, $m = 800$, $n = 1000$, $d = 10$, $N = 100$, $k = 50$. In all trials D-EBF was deployed using value $1/6$.

ϵ used by D-EBF. Despite our concerns that choosing $\epsilon > 1/6$ might cause issue when it comes to clustering partial supports and identifying singleton values, these are not realised at least in the problem considered here. In fact, choosing a larger value for ϵ seems to help in this case as it facilitates the extraction of more partial supports which is particularly important for this problem as $N = 100$ is relatively small.

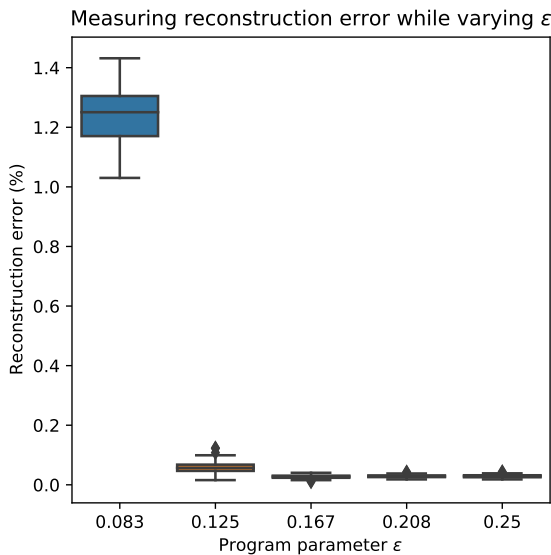


Fig. 5. Distribution of reconstruction errors $\|\mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}\|_F^2 / \|\mathbf{Y}\|_F^2$ of reconstructions $\hat{\mathbf{A}}$, $\hat{\mathbf{X}}$ returned by D-EBF varying the program proxy expansion parameter ϵ . Parameter settings for experiment: number of trials $T=100$, $m = 800$, $n = 1000$, $d = 10$, $N = 100$, $k = 50$.

APPENDIX C EXTRACTION OF SINGLETON VALUES AND PARTIAL SUPPORTS

In this appendix we provide proofs of the results given in Section III concerning the extraction of singleton values and the clustering of partial support. To quickly recap some key pieces of notation we recall the following function definitions.

- Let the function $f : \mathbb{R} \times \mathbb{R}^m \rightarrow [m]$ count the number of times a real number $\alpha \in \mathbb{R}$ appears in some vector $\mathbf{r} \in \mathbb{R}^m$, $f(\alpha, \mathbf{r}) := |\{j \in [m] : r_j = \alpha\}|$.
- Let the function $g : \mathbb{R} \times \mathbb{R}^m \rightarrow \{0, 1\}^m$ return a binary vector whose nonzeros correspond to the locations in which $\alpha \in \mathbb{R}$ appears in a vector $\mathbf{r} \in \mathbb{R}^m$, i.e., with $\mathbf{w} = g(\alpha, \mathbf{r})$ then $w_j = 1$ iff $r_j = \alpha$ and is 0 otherwise.

We first prove Lemma III.1, the three properties of expander graph adjacency matrices identified in this lemma will provide the basis for much of what follows.

Lemma III.1 (Properties of the adjacency matrix of a (k, ϵ, d) -expander graph) *If $\mathbf{A} \in \mathcal{E}_{k, \epsilon, d}^{m \times n}$ then any submatrix \mathbf{A}_S of \mathbf{A} , where $S \in [n]^{(\leq k)}$, satisfies the following.*

- 1) *There are more than $(1 - \epsilon)d|S|$ rows in \mathbf{A}_S that have **at least one non-zero**.*
- 2) *There are more than $(1 - 2\epsilon)d|S|$ rows in \mathbf{A}_S that have **only one non-zero**.*
- 3) *The overlap in support of the columns of \mathbf{A}_S is upper bounded as follows,*

$$\left| \bigcap_{l \in S} \text{supp}(\mathbf{a}_l) \right| < 2\epsilon d.$$

Proof. Property 1 follows directly from Definition II.1 and property 2 from Theorem II.2. To prove property 3, consider the pairwise overlap in support between any two distinct columns \mathbf{a}_l and \mathbf{a}_h of $\mathbf{A} \in \mathcal{E}_{k, \epsilon, d}^{m \times n}$. From Definition II.1 it follows that

$$|\text{supp}(\mathbf{a}_l) \cup \text{supp}(\mathbf{a}_h)| > (1 - \epsilon)2d.$$

Using the inclusion-exclusion principle then

$$|\text{supp}(\mathbf{a}_l)| + |\text{supp}(\mathbf{a}_h)| - |\text{supp}(\mathbf{a}_l) \cap \text{supp}(\mathbf{a}_h)| > (1 - \epsilon)2d.$$

Given that $|\text{supp}(\mathbf{a}_l)| = |\text{supp}(\mathbf{a}_h)| = d$, then a simple rearrangement shows that

$$|\text{supp}(\mathbf{a}_l) \cap \text{supp}(\mathbf{a}_h)| < 2d - (1 - \epsilon)2d = 2\epsilon d.$$

Therefore, for any $S \subseteq [n]$ with $|S| \geq 2$ and for any $l, h \in S$ such that $l \neq h$, then

$$\left| \bigcap_{l \in S} \text{supp}(\mathbf{a}_l) \right| \leq |\text{supp}(\mathbf{a}_l) \cap \text{supp}(\mathbf{a}_h)| < 2\epsilon d.$$

□

Using Lemma III.1 it is possible to derive the following sufficient condition for identifying singleton values, based on the frequency with which they appear.

Corollary III.1.1 (Sufficient condition for identifying singleton values) (i) *Consider a vector $\mathbf{r} = \mathbf{A}\mathbf{z}$ where $\mathbf{A} \in \mathcal{E}_{k, \epsilon, d}^{m \times n}$ and $\mathbf{z} \in \chi_k^n$. With $\alpha \in \mathbb{R} \setminus \{0\}$ if $f(\alpha, \mathbf{r}) \geq 2\epsilon d$ then there exists an $l \in \text{supp}(\mathbf{x})$ such that $\alpha = x_l$.*

Proof. By construction if $\alpha \in \mathbb{R} \setminus \{0\}$ and $f(\alpha, \mathbf{r}) > 0$ then there exists a $\mathcal{S} \subseteq \text{supp}(\mathbf{z})$ such that $\alpha = \sum_{l \in \mathcal{S}} x_l$. Suppose that $|\mathcal{S}| \geq 2$. Then

$$\begin{aligned} f(\alpha, \mathbf{r}) &= |\{j \in [m] : \mathcal{S} = \text{supp}(\tilde{\mathbf{a}}_j) \cap \text{supp}(\mathbf{z})\}| \\ &\leq |\{j \in [m] : \mathcal{S} \subseteq \text{supp}(\tilde{\mathbf{a}}_j) \cap \text{supp}(\mathbf{z})\}| \\ &\leq |\{j \in [m] : \mathcal{S} \subseteq \text{supp}(\tilde{\mathbf{a}}_j)\}| \\ &= \left| \bigcap_{l \in \mathcal{S}} \text{supp}(\mathbf{a}_l) \right| \\ &< 2\epsilon d, \end{aligned}$$

where the final inequality follows from Lemma III.1. Therefore if $\alpha \in \mathbb{R} \setminus \{0\}$ satisfies $f(\alpha, \mathbf{r}) \geq 2\epsilon d$ it must follow that α is a singleton value, i.e., $|\mathcal{S}| = 1$, and therefore there exists an $l \in \text{supp}(\mathbf{x})$ such that $\alpha = x_l$. \square

Corollary III.1.1 provides a sufficient condition for identifying singleton values, however it does not provide any guarantees that there will be any singleton values appearing at least $2\epsilon d$ in any such \mathbf{r} . This existence of high frequency singleton values is resolved by the following lemma.

Lemma III.2 (Existence of frequently occurring singleton values, adapted from [34, Theorem 4.6]) Consider a vector $\mathbf{r} = \mathbf{A}\mathbf{z}$ where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ and $\mathbf{z} \in \chi_k^n$. For $l \in \text{supp}(\mathbf{z})$, let $\Omega_l := \{j \in [m] : r_j = z_l\}$ be the set of row indices $j \in [m]$ of \mathbf{r} such that $r_j = z_l$. Defining $\mathcal{T} := \{l \in \text{supp}(\mathbf{z}) : |\Omega_l| > (1 - 2\epsilon)d\}$ as the set of singleton values which each appear more than $(1 - 2\epsilon)d$ times in \mathbf{r} , then

$$|\mathcal{T}| \geq \frac{|\text{supp}(\mathbf{z})|}{(1 + 2\epsilon)d}.$$

Proof. For typographical ease let $\gamma := \lceil (1 - 2\epsilon)d \rceil$ and define $\mathcal{U} := \text{supp}(\mathbf{z}) \setminus \mathcal{T}$ as the set of singleton values which appear at most $(1 - 2\epsilon)d$ times. Additionally, for this proof we adopt the following graph inspired notation.

- $\mathcal{N}_1(\mathbf{r}) := \bigcup_{l \in \text{supp}(\mathbf{z})} \Omega_l$ is the set of row indices of \mathbf{r} corresponding to singleton values, by Theorem II.2 it holds that $|\mathcal{N}_1(\mathbf{r})| \geq \gamma |\text{supp}(\mathbf{z})|$.
- $\mathcal{N}_1^{\mathcal{T}}(\mathbf{r}) := \bigcup_{l \in \mathcal{T}} \Omega_l$ is the set of row indices of \mathbf{r} corresponding to singleton values that appear more than $(1 - 2\epsilon)d$ times. Therefore $\gamma |\mathcal{T}| \leq |\mathcal{N}_1^{\mathcal{T}}(\mathbf{r})| \leq d |\mathcal{T}|$.
- $\mathcal{N}_1^{\mathcal{U}}(\mathbf{r}) := \bigcup_{l \in \mathcal{U}} \Omega_l$ is the set of row indices of \mathbf{r} corresponding to singleton values that appear at most $(1 - 2\epsilon)d$ times. Note that $|\mathcal{N}_1^{\mathcal{U}}(\mathbf{r})| \leq (\gamma - 1) |\mathcal{U}|$.

Clearly $|\mathcal{N}_1(\mathbf{r})| = |\mathcal{N}_1^{\mathcal{T}}(\mathbf{r})| + |\mathcal{N}_1^{\mathcal{U}}(\mathbf{r})|$, it therefore follows that

$$\begin{aligned} |\mathcal{N}_1(\mathbf{r})| &\geq \gamma |\text{supp}(\mathbf{z})| \\ &= \gamma (|\mathcal{T}| + |\mathcal{U}|) \\ &= \gamma |\mathcal{T}| + |\mathcal{U}| + (\gamma - 1) |\mathcal{U}| \\ &\geq \gamma |\mathcal{T}| + |\mathcal{U}| + |\mathcal{N}_1^{\mathcal{U}}(\mathbf{r})|. \end{aligned}$$

Substituting $|\mathcal{N}_1(\mathbf{r})| = |\mathcal{N}_1^{\mathcal{T}}(\mathbf{r})| + |\mathcal{N}_1^{\mathcal{U}}(\mathbf{r})|$ then

$$\begin{aligned} |\mathcal{N}_1^{\mathcal{T}}(\mathbf{r})| &\geq \gamma |\mathcal{T}| + |\mathcal{U}| \\ &= |\mathcal{T}| + |\mathcal{U}| + (\gamma - 1) |\mathcal{T}| \\ &= |\text{supp}(\mathbf{z})| + (\gamma - 1) |\mathcal{T}|. \end{aligned}$$

As $|\mathcal{N}_1^{\mathcal{T}}(\mathbf{r})| \leq d |\mathcal{T}|$ then $d |\mathcal{T}| \geq |\text{supp}(\mathbf{z})| + (\gamma - 1) |\mathcal{T}|$, rearranging gives

$$|\mathcal{T}| \geq \frac{|\text{supp}(\mathbf{z})|}{d - \gamma + 1}.$$

Finally, as $\gamma \geq d - 2\epsilon d$ then

$$|\mathcal{T}| \geq \frac{|\text{supp}(\mathbf{z})|}{1 + 2\epsilon d}$$

as claimed. Note in addition that so long as \mathbf{r} is nonzero, meaning $|\text{supp}(\mathbf{z})| \geq 1$, then it is always possible to extract at least one partial support with cardinality greater than $(1 - 2\epsilon)d$. \square

In addition, the sufficient condition concerning the extraction of singleton values and partial supports, Corollary III.1.1, can be extended to the following case.

Lemma III.3 [Sufficient condition for identifying singleton values (ii)] Consider a vector $\mathbf{y} = \mathbf{A}\mathbf{x}$ where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ and $\mathbf{x} \in \chi_k^n$. Let $\hat{\mathbf{A}} \in \{0, 1\}^{m \times n}$ and $\hat{\mathbf{x}} \in \chi_k^n$ be such that a column $\hat{\mathbf{a}}_l$ of $\hat{\mathbf{A}}$ is nonzero iff $\hat{x}_l \neq 0$, and there exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\text{supp}(\hat{\mathbf{A}}\mathbf{P}^T) \subseteq \text{supp}(\mathbf{A})$, $\text{supp}(\mathbf{P}\hat{\mathbf{x}}) \subseteq \text{supp}(\mathbf{x})$ and $\hat{x}_{P(l)} = x_{P(l)}$ for all $l \in \text{supp}(\hat{\mathbf{x}})$, where $P : [n] \rightarrow [n]$ denotes the row permutation caused by pre-multiplication with \mathbf{P} . Consider the residual $\mathbf{r} = \mathbf{y} - \hat{\mathbf{A}}\hat{\mathbf{x}}$, if for some $\alpha \in \mathbb{R} \setminus \{0\}$ it holds that $f(\alpha, \mathbf{r}) \geq 2\epsilon d$ is satisfied, then there exists an $l \in [n]$ such that $\alpha = x_l$. Furthermore, with $\mathbf{w} = g(\alpha, \mathbf{r})$ then $\text{supp}(\mathbf{w}) \subseteq \text{supp}(\mathbf{a}_l)$ is a partial support of \mathbf{a}_l .

Proof. As in Corollary III.1.1, then by construction if $\alpha \in \mathbb{R} \setminus \{0\}$ and $f(\alpha, \mathbf{r}) > 0$ then there exists a $\mathcal{S} \subseteq \text{supp}(\mathbf{z})$ such that $\alpha = \sum_{l \in \mathcal{S}} x_l$. In order to simplify our notation for what follows, let $\mathbf{V} := \hat{\mathbf{A}}\mathbf{P}^T$ and $\mathbf{u} := \mathbf{P}\hat{\mathbf{x}}$. By construction $\text{supp}(\mathbf{V}) \subseteq \text{supp}(\mathbf{A})$ and $\text{supp}(\mathbf{u}) \subseteq \text{supp}(\mathbf{x})$. In addition, with $\Omega_j := \text{supp}(\tilde{\mathbf{a}}_j) \cap \text{supp}(\mathbf{x})$ and $\Gamma_j := \text{supp}(\tilde{\mathbf{v}}_j) \cap \text{supp}(\mathbf{u})$, then $\Gamma_j \subseteq \Omega_j$ for all $j \in [m]$. As $u_l = x_l$ for all $l \in \text{supp}(\mathbf{u})$ and $\mathbf{r} = \mathbf{y} - \mathbf{V}\mathbf{u}$ then

$$r_j = \sum_{l \in \Omega_j} x_l - \sum_{h \in \Gamma_j} x_h = \sum_{l \in (\Omega_j \setminus \Gamma_j)} x_l.$$

It therefore follows that each entry r_j is a sum over a subset of the nonzero values of \mathbf{x} used in the sum of y_j . As a result, we may write $\mathbf{r} := \mathbf{B}\mathbf{x}$ where $\mathbf{B} \in \{0, 1\}^{m \times n}$ and $\text{supp}(\mathbf{B}) \subseteq \text{supp}(\mathbf{A})$. Therefore, for any $\mathcal{S} \subseteq \text{supp}(\mathbf{x})$ such that $|\mathcal{S}| \geq 2$, then from Lemma III.1 it follows that

$$\left| \bigcap_{l \in \mathcal{S}} \text{supp}(\mathbf{b}_l) \right| \leq \left| \bigcap_{l \in \mathcal{S}} \text{supp}(\mathbf{a}_l) \right| < 2\epsilon d.$$

Therefore, if $\alpha = \sum_{l \in \mathcal{S}} x_l$ with $|\mathcal{S}| \geq 2$ then $f(\alpha, \mathbf{r}) < 2\epsilon d$. As a result, for some $\alpha \in \mathbb{R}$ if $f(\alpha, \mathbf{r}) \geq 2\epsilon d$ then α is a singleton and there exists an $l \in [n]$ such that $\alpha = x_l$. With $\mathbf{w} = g(\alpha, \mathbf{r})$, then as \mathbf{x} is dissociated it follows that $\text{supp}(\mathbf{w}) \subseteq \text{supp}(\mathbf{b}_l) \subseteq \text{supp}(\mathbf{a}_l)$ and so \mathbf{w} is a partial support of \mathbf{a}_l . \square

The next result provides guarantees with regards to clustering partial supports.

Lemma III.4 (Clustering partial supports) Consider a pair of partial supports \mathbf{w}_1 and \mathbf{w}_2 , extracted from $\mathbf{r}_1 = \mathbf{A}\mathbf{z}_1$

and $\mathbf{r}_2 = \mathbf{A}\mathbf{z}_2$ respectively, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ and $\mathbf{z}_1, \mathbf{z}_2 \in \chi_k^n$. If $\epsilon \leq 1/6$, $|\text{supp}(\mathbf{w}_1)| > (1 - 2\epsilon)d$ and $|\text{supp}(\mathbf{w}_2)| > (1 - 2\epsilon)d$, then \mathbf{w}_1 and \mathbf{w}_2 originate from the same column of \mathbf{A} iff $\mathbf{w}_1^T \mathbf{w}_2 \geq 2\epsilon d$.

Proof. Suppose that \mathbf{w}_1 and \mathbf{w}_2 originate from i th and j th column of \mathbf{A} respectively. From Lemma III.1, and observing that $|\text{supp}(\mathbf{a}_i) \cap \text{supp}(\mathbf{a}_j)| = \mathbf{a}_i^T \mathbf{a}_j$ for binary vectors \mathbf{a}_i and \mathbf{a}_j , then $i \neq j$ iff $\mathbf{a}_i^T \mathbf{a}_j < 2\epsilon d$. Observe also that as $\text{supp}(\mathbf{w}_1) \subseteq \text{supp}(\mathbf{a}_i)$ and $\text{supp}(\mathbf{w}_2) \subseteq \text{supp}(\mathbf{a}_j)$ then $\mathbf{w}_1^T \mathbf{w}_2 \leq \mathbf{a}_i^T \mathbf{a}_j$. Suppose that $\mathbf{w}_1^T \mathbf{w}_2 \geq 2\epsilon d$ and assume that $i \neq j$. This is a contradiction however as it must hold that $\mathbf{a}_i^T \mathbf{a}_j < 2\epsilon d$, but

$$\mathbf{a}_i^T \mathbf{a}_j \geq \mathbf{w}_1^T \mathbf{w}_2 \geq 2\epsilon d.$$

Therefore $\mathbf{w}_1^T \mathbf{w}_2 \geq 2\epsilon d$ implies that $i = j$, i.e., \mathbf{w}_1 and \mathbf{w}_2 originate from the same column of \mathbf{A} . Now assume as per the statement of the lemma that $|\text{supp}(\mathbf{w}_1)| > (1 - 2\epsilon)d$, $|\text{supp}(\mathbf{w}_2)| > (1 - 2\epsilon)d$ and $\epsilon \leq 1/6$, which implies that $(1 - 2\epsilon)d \geq 2\epsilon d$. Suppose $\mathbf{w}_1^T \mathbf{w}_2 < 2\epsilon d$, if $i = j$ then \mathbf{w}_1 and \mathbf{w}_2 must differ by less than $4\epsilon d$ nonzeros. This implies that

$$\mathbf{w}_1^T \mathbf{w}_2 > (1 - 4\epsilon)d \geq 2\epsilon d$$

which is a contradiction. Therefore, under the assumptions given it follows that $\mathbf{w}_1^T \mathbf{w}_2 < 2\epsilon d$ implies $i \neq j$. Combining these results then \mathbf{w}_1 and \mathbf{w}_2 originate from the same column of \mathbf{A} iff $\mathbf{w}_1^T \mathbf{w}_2 \geq 2\epsilon d$ as claimed. \square

APPENDIX D

ACCURACY OF D-EBF AND ND-EBF

In this appendix we prove results concerning the operation of D-EBF and ND-EBF. Before proceeding we emphasise a point regarding notation: many of the variables discussed are updated iteratively throughout the run-time of an algorithm. Therefore, in order to avoid a proliferation of distinct variables representing the same program variable at different stages of run-time, unless otherwise stated we will treat each variable as a program variables. We start by proving accuracy guarantees for the subroutines common to both D-EBF and ND-EBF, starting with EXTRACT&MATCH, detailed in Algorithm 9.

Lemma D.1 (Accuracy of Algorithm 9). *Consider a vector $\mathbf{y} = \mathbf{A}\mathbf{x}$, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{x} \in \chi_k^n$. Let $\hat{\mathbf{A}}^{(1)} \in \{0, 1\}^{m \times n}$ and $\hat{\mathbf{x}}^{(1)} \in \chi_k^n$ be such that a column $\hat{\mathbf{a}}_l^{(1)}$ of $\hat{\mathbf{A}}^{(1)}$ is nonzero iff $\hat{x}_l^{(1)} \neq 0$, and if it is nonzero then $|\text{supp}(\hat{\mathbf{a}}_l^{(1)})| > (1 - 2\epsilon)d$. In addition, assume that there exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\text{supp}(\hat{\mathbf{A}}^{(1)}\mathbf{P}^T) \subseteq \text{supp}(\mathbf{A})$, $\text{supp}(\mathbf{P}\hat{\mathbf{x}}^{(1)}) \subseteq \text{supp}(\mathbf{x})$ and $\hat{x}_{P(l)}^{(1)} = x_{P(l)}$ for all $l \in \text{supp}(\hat{\mathbf{x}}^{(1)})$, where $P : [n] \rightarrow [n]$ denotes the row permutation caused by pre-multiplication with \mathbf{P} . Define the residual $\mathbf{r} := \mathbf{y} - \hat{\mathbf{A}}^{(1)}\hat{\mathbf{x}}^{(1)}$ and consider the variables returned by Algorithm 9,*

$$\mathbf{q}, \mathbf{W}, p, \hat{\mathbf{A}}^{(2)}, \hat{\mathbf{x}}^{(2)} \leftarrow \text{EXTRACT\&MATCH}(\mathbf{r}, m, d, \epsilon, \hat{\mathbf{A}}^{(1)}, \hat{\mathbf{x}}^{(1)})$$

Then the following hold.

- 1) If $p \geq 1$ then for each $i \in [p]$ there exists a unique $l \in \text{supp}(\mathbf{x})$ such that $q_i = x_l$ and $\text{supp}(\mathbf{w}_i) \subseteq \text{supp}(\mathbf{a}_l)$

with $|\text{supp}(\mathbf{w}_i)| > (1 - 2\epsilon)d$. In addition, $\mathbf{a}_i^T \hat{\mathbf{a}}_h^{(1)} < 2\epsilon d$ for all $h \in [n]$.

- 2) The updated reconstructions satisfy the following: a column $\hat{\mathbf{a}}_i^{(2)}$ of $\hat{\mathbf{A}}^{(2)}$ is nonzero iff $\hat{x}_i^{(2)} \neq 0$, $\text{supp}(\hat{\mathbf{A}}^{(1)}\mathbf{P}^T) \subseteq \text{supp}(\hat{\mathbf{A}}^{(2)}\mathbf{P}^T) \subseteq \text{supp}(\mathbf{A})$, $\text{supp}(\mathbf{P}\hat{\mathbf{x}}^{(1)}) \subseteq \text{supp}(\mathbf{P}\hat{\mathbf{x}}^{(2)}) \subseteq \text{supp}(\mathbf{x})$ and $\hat{x}_{P(l)}^{(2)} = x_{P(l)}$ for all $l \in \text{supp}(\hat{\mathbf{x}}^{(2)})$.

Proof. We prove both statements by analysing the updates to the reconstructions, which take place on lines 21, 22, 25, 26 and 32. To this end, we first prove that the variables r_i and \mathbf{w} in Algorithm 9 are a singleton value and partial support respectively if $c > (1 - 2\epsilon)d$. We then show that when \mathbf{w} and a nonzero column of $\hat{\mathbf{A}}^{(1)}$ are partial supports originating from the same column in \mathbf{A} , then \mathbf{w} is correctly used to update this column of the reconstruction and r_i the corresponding row of $\hat{\mathbf{x}}^{(1)}$. In the case where \mathbf{w} originates from a different column of \mathbf{A} to any of the nonzero columns of $\hat{\mathbf{A}}^{(1)}$, we show that \mathbf{w} and r_i are correctly added to \mathbf{W} and \mathbf{q} respectively.

By the assumptions of the lemma the residual \mathbf{r} satisfies the conditions of Lemma III.3. Therefore, and as $\epsilon \leq 1/6$, if on line 17 $c = f(r_i, \mathbf{r}) > (1 - 2\epsilon)d \geq 2\epsilon d$ then there exists an $l \in \text{supp}(\mathbf{x})$ such that $r_i = x_l$. Furthermore, with $\mathbf{w} = g(r_i, \mathbf{r})$ then $\text{supp}(\mathbf{w}) \subseteq \text{supp}(\mathbf{a}_l)$ and $|\text{supp}(\mathbf{w})| = f(r_i, \mathbf{r}) > (1 - 2\epsilon)d$. We conclude then that if $c > (1 - 2\epsilon)d$ then r_i and \mathbf{w} are a singleton and partial support respectively. All that remains to be shown is that the partial support \mathbf{w} is either correctly matched to an existing nonzero column of $\hat{\mathbf{A}}^{(1)}$, or it is correctly identified as originating from a column of \mathbf{A} for which there is not currently a partial or full reconstruction. We now assume that \mathbf{w} is a partial support of \mathbf{a}_l and examine each of these two cases in turn.

Suppose there exists a nonzero column $\hat{\mathbf{a}}_h^{(1)}$ of $\hat{\mathbf{A}}^{(1)}$ which is a partial support originating from \mathbf{a}_l , i.e., $\text{supp}(\hat{\mathbf{a}}_h^{(1)}) \subseteq \text{supp}(\mathbf{a}_l)$. As $|\text{supp}(\hat{\mathbf{a}}_h^{(1)})| > (1 - 2\epsilon)d \geq 2\epsilon d$ then by Lemma III.1 it follows that $\text{supp}(\hat{\mathbf{a}}_h^{(1)}) \not\subseteq \text{supp}(\mathbf{a}_k)$ for all $k \neq h$. From Lemma III.4 then $\mathbf{w}^T \hat{\mathbf{a}}_h^{(1)} \geq 2\epsilon d$ and $\mathbf{w}^T \hat{\mathbf{a}}_k^{(1)} < 2\epsilon d$ for all $k \neq h$. As a result, on line 19 it must follow that $\kappa = h$ and on line 20 that $t_\kappa \geq 2\epsilon d$. Therefore, the partial support \mathbf{w} is identified as matching with the reconstruction of \mathbf{a}_l in $\hat{\mathbf{A}}^{(1)}$, and \hat{x}_l and $\hat{\mathbf{a}}_h$ are then correctly updated on lines 21 and 22 accordingly. Consider now the other case, in which there does not exist a nonzero column $\hat{\mathbf{a}}_h^{(1)}$ of $\hat{\mathbf{A}}^{(1)}$ which is a partial support originating from \mathbf{a}_l . Therefore, by Lemma III.4, for any $h \in [n]$ then $\mathbf{w}^T \hat{\mathbf{a}}_h^{(1)} < 2\epsilon d$ and so the condition on line 20 is not satisfied. As a result, the singleton value r_i and partial support \mathbf{w} are correctly used to update \mathbf{q} and \mathbf{W} on lines 25 and 26, and as a consequence statement 1 of the lemma immediately follows. For Statement 2, in addition observe that line 32 introduces no erroneous nonzeros as \mathbf{x} is dissociated. Note also that as updates only occur to the existing nonzero columns of $\hat{\mathbf{A}}^{(1)}$, the sets of column indices corresponding to the nonzero columns of $\hat{\mathbf{A}}^{(1)}$ and $\hat{\mathbf{A}}^{(2)}$ are equal. Given these observations Statement 2 must also be true. \square

Under the similar assumptions we are also able to provide the following accuracy guarantees for the DECODE subroutine, detailed in Algorithm 11.

Lemma D.2 (Accuracy of Algorithm 11). Consider a vector $\mathbf{y} = \mathbf{A}\mathbf{x}$, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{x} \in \chi_k^n$. Let $\hat{\mathbf{A}} \in \{0,1\}^{m \times n}$ and $\hat{\mathbf{x}}^{(1)} \in \chi_k^n$ be such that a column $\hat{\mathbf{a}}_l$ of $\hat{\mathbf{A}}$ is nonzero iff $\hat{x}_l^{(1)} \neq 0$, and if it is nonzero then $|\text{supp}(\hat{\mathbf{a}}_l)| > (1 - 2\epsilon)d$. In addition, assume there exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\text{supp}(\hat{\mathbf{A}}\mathbf{P}^T) \subseteq \text{supp}(\mathbf{A})$, $\text{supp}(\mathbf{P}\hat{\mathbf{x}}^{(1)}) \subseteq \text{supp}(\mathbf{x})$ and $\hat{x}_{P(l)}^{(1)} = x_{P(l)}$ for all $l \in \text{supp}(\hat{\mathbf{x}}^{(1)})$, where $P : [n] \rightarrow [n]$ denotes the row permutation caused by pre-multiplication with \mathbf{P} . Consider the variables computed during and returned by DECODE, Algorithm 11,

$$\hat{\mathbf{x}}^{(2)}, \text{UPDATED} \leftarrow \text{DECODE}(\mathbf{y}, m, \epsilon, d, \hat{\mathbf{A}}, \hat{\mathbf{x}}^{(1)}).$$

Let \mathcal{Q} denote the event that $\text{supp}(\mathbf{P}\hat{\mathbf{x}}) \subseteq \text{supp}(\mathbf{x})$ and $\hat{x}_{P(l)} = x_{P(l)}$ for all $l \in \text{supp}(\hat{\mathbf{x}})$. Then the following hold.

- 1) At any point during the run-time of DECODE \mathcal{Q} is true.
- 2) The sparse code reconstruction returned by DECODE satisfies $\text{supp}(\mathbf{P}\hat{\mathbf{x}}^{(1)}) \subseteq \text{supp}(\mathbf{P}\hat{\mathbf{x}}^{(2)})$. Furthermore, recalling that $\mathcal{S} = \{l \in [n] : \hat{\mathbf{a}}_l = d\}$ and defining $P(\mathcal{S}) := \{h \in [n] : \exists l \in \mathcal{S} \text{ s.t. } P(l) = h\}$, if $\text{supp}(\mathbf{x}) \subseteq P(\mathcal{S})$ then $\mathbf{P}\hat{\mathbf{x}}^{(2)} = \mathbf{x}$.
- 3) The while loop, lines 5-31, exits after at most k iterations.

Proof. First we prove Statement 1 by induction. Clearly, by the assumptions on $\hat{\mathbf{x}}^{(1)}$, then \mathcal{Q} is true at the start of iteration 1. In addition, if \mathcal{Q} is true at the end of an iteration it must be true at the beginning of the next iteration. To prove Statement 1 it therefore suffices to prove that if \mathcal{Q} is true at the start, i.e., line 4, of iteration t , then it is also true at the end, i.e., line 31, of iteration t . Suppose then that \mathcal{Q} is true at the start of iteration t , then the residual at the start of iteration t satisfies the conditions required for Lemma III.3. As a result, if on line 20 it holds that $c = f(r_i, \mathbf{r}) \geq 2\epsilon d$, then there exists an $l \in \text{supp}(\mathbf{x})$ such that $r_i = x_l$, meaning r_i is a singleton value. Furthermore, with $\mathbf{w} = g(r_i, \mathbf{r})$ then $|\text{supp}(\mathbf{w})| \geq 2\epsilon d$ and there exists a unique $l \in [n]$ such that $\text{supp}(\mathbf{w}) \subseteq \text{supp}(\mathbf{a}_l)$. Similar to Lemma D.1, we now explore the two possible cases: either there exists a unique $h \in \mathcal{S}$ such that $\hat{\mathbf{a}}_h = \mathbf{a}_l$ or there does not.

Suppose then that there exists a column $\hat{\mathbf{a}}_h$ such that $\hat{\mathbf{a}}_h = \mathbf{a}_l$. Therefore $\mathbf{w}^T \hat{\mathbf{a}}_h \geq 2\epsilon d$ and from Lemma III.4 it follows that $\mathbf{w}^T \hat{\mathbf{a}}_k < 2\epsilon d$ for all $k \neq h$. Therefore on line 22 it must be true that $\kappa = h$. As a result the condition on line 23 is satisfied and so the partial support \mathbf{w} is identified as matching with the complete reconstruction of \mathbf{a}_l in $\hat{\mathbf{A}}$. Subsequently the corresponding entry in $\hat{\mathbf{x}}$ is correctly updated as per line 24. Suppose now for all $h \in \mathcal{S}$ it holds that $\hat{\mathbf{a}}_h \neq \mathbf{a}_l$. From Lemma III.1 it therefore follows that $t_h = \mathbf{w}^T \hat{\mathbf{a}}_h \leq \mathbf{a}_l^T \mathbf{a}_{P(h)} < 2\epsilon d$ for all $h \in \mathcal{S}$. Therefore $t_\kappa < 2\epsilon d$ and hence the condition on line 23 is not satisfied. As a result, under this assumption \mathbf{w} correctly does not match with any complete column of the reconstruction and so no update occurs to $\hat{\mathbf{x}}$. As $\hat{\mathbf{A}}$ is unchanged and only the row entries of $\hat{\mathbf{x}}$ with index in \mathcal{S} are updated, then this implies that \mathcal{Q} is true at the end of iteration t . This concludes the proof by induction of Statement 1.

The first part of Statement 2 follows immediately from Statement 1. For the second part of Statement 2, let $\mathbf{V} := \hat{\mathbf{A}}\mathbf{P}^T$ and $\mathbf{u} := \mathbf{P}\hat{\mathbf{x}}^{(1)}$. By assumption $\text{supp}(\mathbf{V}) \subseteq \text{supp}(\mathbf{A})$

and $\text{supp}(\mathbf{u}) \subseteq \text{supp}(\mathbf{x})$, therefore, at any iteration of the while loop, the residual calculated on line 6 can be expressed as

$$\mathbf{r} = \mathbf{y} - \hat{\mathbf{A}}\hat{\mathbf{x}} = \mathbf{A}\mathbf{x} - \mathbf{V}\mathbf{u}.$$

For typographical ease let $\mathcal{C} := \text{supp}(\mathbf{x})$ and suppose that $\mathcal{C} \subseteq P(\mathcal{S})$, we proceed by contradiction and assume that Statement 2 is false. If Statement 2 is false then this implies that the while loop of DECODE exits at some iteration with $\mathbf{r} \neq \mathbf{0}_m$. Therefore, and as Statement 1 is true, it must follow that $f(x_l, \mathbf{r}) < 2\epsilon d$ for all $l \in \mathcal{C}$. Due to our assumptions and the proof of Statement 1 then $\mathbf{V}_{P(\mathcal{S})} = \mathbf{A}_{P(\mathcal{S})}$, $\text{supp}(\mathbf{u}) \subseteq \mathcal{C}$ and $u_l = x_l$ for all $l \in \text{supp}(\mathbf{u})$. Therefore

$$\mathbf{r} = \mathbf{A}_\mathcal{C}\mathbf{x}_\mathcal{C} - \mathbf{V}_\mathcal{C}\mathbf{u}_\mathcal{C} = \mathbf{A}_\mathcal{C}\mathbf{z}$$

where $\mathbf{A}_\mathcal{C} \in \mathcal{E}_{k,\epsilon,d}^{m \times k}$ with $\epsilon \leq 1/6$ and $\mathbf{z} := \mathbf{x}_\mathcal{C} - \mathbf{u}_\mathcal{C} \in \chi_k^k$. As $\mathbf{A}_\mathcal{C} \in \mathcal{E}_{k,\epsilon,d}^{m \times k}$ with $\epsilon \leq 1/6$ and $\mathbf{z} \in \chi_k^k$, $\mathbf{z} \neq \mathbf{0}_k$, then Lemma III.2 ensures the existence an $l \in \text{supp}(\mathbf{x}) \setminus \text{supp}(\mathbf{u})$ such that $f(x_l, \mathbf{r}) > (1 - 2\epsilon)d \geq 2\epsilon d$, which is a contradiction. Therefore Statement 2 must also be true.

Finally, to prove Statement 3, observe that at each iteration prior to the exit iteration the contribution of at least one column of $\hat{\mathbf{A}}_\mathcal{S}$ is removed from the residual \mathbf{r} . As \mathbf{x} is k sparse then the maximum number of possible iterates is k . \square

With Lemmas D.1 and D.2 in place then we are able to prove accuracy guarantees for Algorithm 6, D-EBF.

Lemma III.5 (Accuracy of D-EBF) Let $\mathbf{Y} = \mathbf{A}\mathbf{X}$, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{X} \in \chi_k^{n \times N}$. Consider Algorithm 6, D-EBF: in regard to the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ computed at any point during the run-time of D-EBF, we will say \mathcal{Q} is true iff the following hold.

- (a) A column $\hat{\mathbf{a}}_l$ of $\hat{\mathbf{A}}$ is nonzero iff the l th row of $\hat{\mathbf{X}}$ is nonzero.
- (b) For all $l \in [n]$ if $\hat{\mathbf{a}}_l$ is nonzero then $\text{supp}(\hat{\mathbf{a}}_l) > (1 - 2\epsilon)d$.
- (c) There exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\text{supp}(\hat{\mathbf{A}}\mathbf{P}^T) \subseteq \text{supp}(\mathbf{A})$, $\text{supp}(\mathbf{P}\hat{\mathbf{X}}) \subseteq \text{supp}(\mathbf{X})$ and $\hat{x}_{P(l),h} = x_{P(l),h}$ for all $h \in [N]$ and $l \in \text{supp}(\hat{\mathbf{x}}_l)$, where $P : [n] \rightarrow [n]$ denotes the row permutation caused by pre-multiplication with \mathbf{P} .

Suppose that D-EBF exits after the completion of iteration $t_f \in \mathbb{N}$ of the while loop, then the following statements are true.

- 1) At any point during the run time of D-EBF \mathcal{Q} is true. Therefore the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ returned by D-EBF also satisfy the conditions for \mathcal{Q} to be true.
- 2) For any iteration of the while loop $t < t_f$, let $\hat{\mathbf{A}}^{(1)}$ and $\hat{\mathbf{X}}^{(1)}$ denote the reconstructions at the start of the while loop, line 6, and $\hat{\mathbf{A}}^{(2)}$ and $\hat{\mathbf{X}}^{(2)}$ denote the reconstructions at the end of the while loop, line 26. Then either $\text{supp}(\hat{\mathbf{A}}^{(1)}) \subset \text{supp}(\hat{\mathbf{A}}^{(2)})$ and $\text{supp}(\hat{\mathbf{X}}^{(1)}) \subseteq \text{supp}(\hat{\mathbf{X}}^{(2)})$ or $\text{supp}(\hat{\mathbf{A}}^{(1)}) \subseteq \text{supp}(\hat{\mathbf{A}}^{(2)})$ and $\text{supp}(\hat{\mathbf{X}}^{(1)}) \subset \text{supp}(\hat{\mathbf{X}}^{(2)})$.
- 3) Consider the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ returned by D-EBF and the associated residual $\mathbf{R} = \mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}$: there exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$ and $\mathbf{P}\hat{\mathbf{X}} = \mathbf{X}$ iff $\mathbf{R} = \mathbf{0}_{m \times N}$. As a result, $\mathbf{R} = \mathbf{0}_{m \times N}$ is a sufficient condition to ensure that \mathbf{Y} has a unique,

up to permutation, factorisation of the form $\mathbf{Y} = \mathbf{A}\mathbf{X}$ where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{X} \in \chi_k^{n \times N}$.

Proof. We first prove statement 1 by induction, showing that if at the start of an iteration of the while loop \mathcal{Q} is true, then \mathcal{Q} will also be true throughout and up to the end of this iteration. This implies that \mathcal{Q} will also be true at the start of the next iteration, and therefore by induction we will be able to conclude that Statement 1 is true. At the start of the first iterate \mathcal{Q} is trivially true as this corresponds to the initialisation $\hat{\mathbf{A}} = \mathbf{0}_{m \times n}$, $\hat{\mathbf{X}} = \mathbf{0}_{n \times N}$. Assume then at any iterate $t \in [t_f]$ that \mathcal{Q} is true on line 6, this implies \mathcal{Q} is also true at the start of EXTRACT&MATCH. This ensures that the conditions on $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ required by Lemma D.1 are satisfied and therefore the reconstructions returned by EXTRACT&MATCH, which are passed to CLUSTER&ADD on line 20, are also such that \mathcal{Q} is true. In addition, by Statement 1 of Lemma D.1, then for all $i \in [p]$ it follows that q_i and \mathbf{w}_i are a singleton value and partial support respectively, with $|\text{supp}(\mathbf{w}_i)| > (1 - 2\epsilon)d$. Furthermore, these partial supports originate from columns of \mathbf{A} for which there is not yet a corresponding nonzero partial or complete column in $\hat{\mathbf{A}}$. As a result, the only manner in which the reconstructions returned by CLUSTER&ADD on line 20 can result in \mathcal{Q} being false is if there is a clustering error, i.e., there exists a pair of partial supports whose inner product is larger than $2\epsilon d$ despite originating from different columns. This however would contradict Lemma III.4, therefore the reconstructions returned by CLUSTER&ADD, which are inputted into DECODE on line 23, ensure that \mathcal{Q} is true. This in turn implies that the conditions of Lemma D.2 are satisfied, and therefore, from Statement 1 of Lemma D.2, it follows that the reconstructions returned after the completion of the for loop, running on lines 22-24, are also such that \mathcal{Q} is true. Therefore, if \mathcal{Q} is true at the beginning of an iteration it is also true throughout said iteration and at its end. As this implies that \mathcal{Q} is true at the beginning of the next iteration, and given that \mathcal{Q} is true at the start of the first iteration, then by induction Statement 1 must be true.

Given that Statement 1 is true, and as by inspection at no point during the run-time of D-EBF are nonzeros removed from the reconstructions, then clearly $\text{supp}(\hat{\mathbf{A}}^{(1)}) \subseteq \text{supp}(\hat{\mathbf{A}}^{(2)})$ and $\text{supp}(\hat{\mathbf{X}}^{(1)}) \subseteq \text{supp}(\hat{\mathbf{X}}^{(2)})$. In addition, as it is assumed that $t < t_f$, then by the construction of D-EBF an update must occur to either $\hat{\mathbf{A}}$ or $\hat{\mathbf{X}}$ during the iteration, therefore Statement 2 must also be true.

Finally, to prove Statement 3, clearly if there exists a $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\mathbf{A} = \hat{\mathbf{A}}\mathbf{P}^T$ and $\mathbf{X} = \mathbf{P}\hat{\mathbf{X}}$ then $\mathbf{Y} = \mathbf{A}\mathbf{X} = \hat{\mathbf{A}}\mathbf{P}^T\mathbf{P}\hat{\mathbf{X}}$ and so $\mathbf{R} = \mathbf{0}_{m \times N}$. For the converse, as Statement 1 is true, then there exists a $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\mathbf{V} := \hat{\mathbf{A}}\mathbf{P}^T$ satisfies $\text{supp}(\mathbf{V}) \subseteq \text{supp}(\mathbf{A})$, $\mathbf{U} := \mathbf{P}\hat{\mathbf{X}}$ satisfies $\text{supp}(\mathbf{U}) \subseteq \text{supp}(\mathbf{X})$ and $u_{i,j} = x_{i,j}$ for all $(i,j) \in \text{supp}(\mathbf{U})$. If $\mathbf{R} = \mathbf{0}_{m \times N}$ then $\mathbf{Y} = \hat{\mathbf{A}}\hat{\mathbf{X}} = \hat{\mathbf{A}}\mathbf{P}^T\mathbf{P}\hat{\mathbf{X}} = \mathbf{V}\mathbf{U}$. As $\mathbf{Y} = \mathbf{A}\mathbf{X} = \mathbf{V}\mathbf{U}$ and $\text{supp}(\mathbf{V}) \subseteq \text{supp}(\mathbf{A})$, $\text{supp}(\mathbf{U}) \subseteq \text{supp}(\mathbf{X})$ and $u_{i,j} = x_{i,j}$ for all $(i,j) \in \text{supp}(\mathbf{U})$, then $\mathbf{A} = \mathbf{V}$ and $\mathbf{X} = \mathbf{U}$. \square

Analogous accuracy guarantees can be proved for ND-EBF.

Lemma IV.1 (Accuracy of ND-EBF) Let $\mathbf{Y} = \mathbf{A}\mathbf{X}$, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{X} \in \chi_k^{n \times N}$. Consider

Algorithm 13, ND-EBF: in regard to the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ computed at any point during the run-time of ND-EBF, we will say, for $\eta \in [n]$, that $\mathcal{Q}(\eta)$ is true iff the following hold.

- (a) The column vector $\hat{\mathbf{a}}_l$ and row vector $\hat{\mathbf{x}}_l$ are nonzero iff $l < \eta$.
- (b) If $l < \eta$ then $|\text{supp}(\hat{\mathbf{a}}_l)| = d$.
- (c) There exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\text{supp}(\hat{\mathbf{A}}\mathbf{P}^T) \subseteq \text{supp}(\mathbf{A})$, $\text{supp}(\mathbf{P}\hat{\mathbf{X}}) \subseteq \text{supp}(\mathbf{X})$ and $\hat{x}_{P(l),i} = x_{P(l),i}$ for all $i \in [N]$ and $l \in \text{supp}(\hat{\mathbf{x}}_l)$, where $P : [n] \rightarrow [n]$ denotes the row permutation caused by pre-multiplication with \mathbf{P} .

Suppose that ND-EBF exits after the completion of iteration $\eta_f \in [n]$ of the while loop, then the following statements are true.

- 1) For all $\eta \in [\eta_f]$, then at the start of the η th iteration of the while loop $\mathcal{Q}(\eta)$ is true.
- 2) Consider the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ returned by ND-EBF: $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ always satisfy (c) and there exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$ and $\mathbf{P}\hat{\mathbf{X}} = \mathbf{X}$ iff $\eta_f = n$ and $\mathcal{Q}(n+1)$ is true.
- 3) Consider the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ returned by ND-EBF and the associated residual $\mathbf{R} = \mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}}$: there exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$ and $\mathbf{P}\hat{\mathbf{X}} = \mathbf{X}$ iff $\mathbf{R} = \mathbf{0}_{m \times N}$. As a result, $\mathbf{R} = \mathbf{0}_{m \times N}$ is a sufficient condition to ensure that \mathbf{Y} has a unique, up to permutation, factorisation of the form $\mathbf{Y} = \mathbf{A}\mathbf{X}$, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{X} \in \chi_k^{n \times N}$.

Proof. We will prove Statement 1 by induction: the base case $\mathcal{Q}(1)$ is trivially true at the start of the first iteration as this corresponds to the initialisation $\hat{\mathbf{A}}^{(0)} = \mathbf{0}_{m \times n}$ and $\hat{\mathbf{X}}^{(0)} = \mathbf{0}_{n \times N}$. It therefore suffices to prove for any $\eta \in [n_f - 1]$, that if $\mathcal{Q}(\eta)$ is true at the start of iteration η , then $\mathcal{Q}(\eta + 1)$ is true at the end of iteration η , as this in turn implies $\mathcal{Q}(\eta + 1)$ is true at the start of iteration $\eta + 1$. Assume then for arbitrary $\eta \in [n_f - 1]$ that $\mathcal{Q}(\eta)$ is true at the start of iteration η , it follows that the reconstructions $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ inputted into EXTRACT&MATCH satisfy the conditions required for Lemma D.1. Therefore, by Lemma D.1, for all $i \in [p]$ it follows that q_i and \mathbf{w}_i are a singleton value and partial support respectively, $|\text{supp}(\mathbf{w}_i)| > (1 - 2\epsilon)d$ and \mathbf{w}_i originates from a column $\mathbf{a}_l \neq \hat{\mathbf{a}}_h$ for all $h \in [n]$. Note also that as $\text{supp}(\hat{\mathbf{a}}_l) = d$ for any $l < \eta$ and $\hat{\mathbf{a}}_l = \mathbf{0}_m$ otherwise, updates to $\hat{\mathbf{X}}$ during EXTRACT&MATCH occur only on rows with a row index less than η . Furthermore, as per line 12 ND-EBF ignores updates to $\hat{\mathbf{A}}$ performed by EXTRACT&MATCH. Therefore $\mathcal{Q}(\eta)$ must still be true after the completion of EXTRACT&MATCH. Consider now the updates to the reconstructions performed by the MAXCLUSTER&ADD subroutine on line 19 of ND-EBF. This subroutine updates only the η th column and row of $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ respectively. As $\eta < \eta_f$ then $p > 0$ and $|\text{supp}(\hat{\mathbf{a}}_\eta)| = d$. Furthermore, the corresponding row of $\hat{\mathbf{X}}$ is also updated by MAXCLUSTER&ADD with $\hat{x}_{\eta,i} = q_i$ for all $i \in \mathcal{C}^*$. Therefore conditions (a) and (b) of $\mathcal{Q}(\eta + 1)$ must be true at the completion of MAXCLUSTER&ADD. Assume now that condition (c) is not satisfied by the reconstructions returned

by MAXCLUSTER&ADD: then $\hat{\mathbf{a}}_\eta \neq \mathbf{a}_h$ for all $h \in [n]$ and as a result there must have occurred a clustering error. This implies in turn that there exists a pair of columns of \mathbf{W} whose inner product is at least $2\epsilon d$ despite both columns being partial supports originating from different columns of \mathbf{A} . This however would contradict Lemma III.4, therefore by contradiction condition (c) is satisfied and as a result we may conclude that at the exit of MAXCLUSTER&ADD $\mathcal{Q}(\eta + 1)$ is true. As a result, the reconstructions inputted to DECODE on line 22 are such that the conditions required for Lemma D.2 are satisfied, hence by Lemma D.2 it follows that $\text{supp}(\mathbf{P}\hat{\mathbf{X}}) \subseteq \text{supp}(\mathbf{X})$ and $\hat{x}_{P(l),i} = x_{P(l),i}$ for all $(l, i) \in \text{supp}(\hat{\mathbf{X}})$. As DECODE only updates the rows of $\hat{\mathbf{X}}$ corresponding to the complete columns of $\hat{\mathbf{A}}$, then at the end of iteration η of the while loop of ND-EBF it follows that $\mathcal{Q}(\eta + 1)$ is true. This concludes the proof by induction of Statement 1.

Turning our attention to Statement 2 then to prove the reconstructions returned by ND-EBF always satisfy condition (c), observe from Statement 1 that $\mathcal{Q}(\eta_f)$ is true at the start of iteration η_f . Therefore, by the same arguments used in the proof of Statement 1, the reconstructions returned by MAXCLUSTER&ADD at the η_f th iteration satisfy condition (c). As η_f is the final iteration of the while loop of ND-EBF, then by construction there are two possibilities: (i) $|\text{supp}(\mathbf{a}_\eta)| < d$ or (ii) $|\text{supp}(\mathbf{a}_\eta)| = d$ and $\eta_f = n$. In the case of (i), ND-EBF exits with no further updates to the reconstructions, therefore the reconstructions returned by ND-EBF must satisfy (c). In the case of (ii), as the reconstructions satisfy condition (c) and $|\text{supp}(\hat{\mathbf{a}}_h)| = d$ for all $h \in [n]$, then there must also exist a permutation $\mathbf{P} \in \mathcal{P}^{n \times n}$ such $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$. It follows that the inputs to DECODE during the n th iteration of the while loop satisfy the conditions required for Lemma D.2. As $\mathcal{S} = \{l \in [N] : |\text{supp}(\hat{\mathbf{a}}_l)| = d\} = [n]$, then $\text{supp}(\mathbf{x}_i) \subseteq P(\mathcal{S})$ for all $i \in [N]$ and, due to Statement 2 of Lemma D.2, it therefore follows that $\mathbf{P}\hat{\mathbf{X}} = \mathbf{X}$. In summary, in either case the reconstructions returned by ND-EBF satisfy condition (c).

Concerning the second part of Statement 2, observe that if $\mathcal{Q}(n+1)$ is true then $\eta_f = n$. Otherwise, if $\eta_f < n$ then ND-EBF exits with $|\text{supp}(\hat{\mathbf{a}}_h)| < d$ for all $\eta_f < h \leq n$, and so $\mathcal{Q}(n+1)$ cannot be true. Suppose there exists a permutation matrix $\mathbf{P} \in \mathcal{P}^{n \times n}$ such $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$ and $\mathbf{P}\hat{\mathbf{X}} = \mathbf{X}$. By inspection it is clear that $\mathcal{Q}(n+1)$ must be true which in turn implies $\eta_f = n$. The converse follows by observing that $\eta_f = n$ and $\mathcal{Q}(n+1)$ being true imply that case (ii) must occur, which we have already shown implies the existence of a permutation $\mathbf{P} \in \mathcal{P}^{n \times n}$ such $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$ and $\mathbf{P}\hat{\mathbf{X}} = \mathbf{X}$.

The proof of Statement 3 follows in the same manner as the proof of Statement 3 in Lemma III.5. Clearly, if there exists a $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\mathbf{A} = \hat{\mathbf{A}}\mathbf{P}^T$ and $\mathbf{X} = \mathbf{P}\hat{\mathbf{X}}$ then $\mathbf{Y} = \hat{\mathbf{A}}\mathbf{P}^T\mathbf{P}\hat{\mathbf{X}} = \hat{\mathbf{A}}\hat{\mathbf{X}}$ and therefore $\mathbf{R} = \mathbf{Y} - \hat{\mathbf{A}}\hat{\mathbf{X}} = \mathbf{0}_{m \times N}$. For the converse, as Statement 2 is true then there exists a $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\mathbf{V} := \hat{\mathbf{A}}\mathbf{P}^T$ satisfies $\text{supp}(\mathbf{V}) \subseteq \text{supp}(\mathbf{A})$, $\mathbf{U} := \mathbf{P}\hat{\mathbf{X}}$ satisfies $\text{supp}(\mathbf{U}) \subseteq \text{supp}(\mathbf{X})$ and $u_{i,j} = x_{i,j}$ for all $(i, j) \in \text{supp}(\mathbf{U})$. As $\mathbf{A}\mathbf{X} = \hat{\mathbf{A}}\mathbf{P}^T\mathbf{P}\hat{\mathbf{X}} = \mathbf{V}\mathbf{U}$, $\text{supp}(\mathbf{V}) \subseteq \text{supp}(\mathbf{A})$, $\text{supp}(\mathbf{U}) \subseteq \text{supp}(\mathbf{X})$ and $u_{i,j} = x_{i,j}$

for all $(i, j) \in \text{supp}(\mathbf{U})$, then $\mathbf{A} = \mathbf{V}$ and $\mathbf{X} = \mathbf{U}$. \square

APPENDIX E

SUPPORTING LEMMAS FOR THEOREM II.3

The following corollary highlights that it suffices to recover the encoder up to permutation.

Corollary IV.1.1 (Recovery of \mathbf{A} up to permutation is sufficient to recover both factors up to permutation) *Let $\mathbf{Y} = \mathbf{A}\mathbf{X}$, where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{X} \in \chi_k^{n \times N}$. Suppose that $\hat{\mathbf{A}}$ and $\hat{\mathbf{X}}$ are the reconstructions of \mathbf{A} and \mathbf{X} returned by either ND-EBF or D-EBF. If there there exists a $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$ then $\mathbf{P}\hat{\mathbf{X}} = \mathbf{X}$.*

Proof. In the case of ND-EBF this result follows directly from Statement 2 of Lemma IV.1. For D-EBF, observe that the last update to $\hat{\mathbf{A}}$ occurs on line 19 during the final iteration of the while loop prior to the algorithm terminating. Therefore, if there exists a $\mathbf{P} \in \mathcal{P}^{n \times n}$ such that the returned reconstruction of \mathbf{A} satisfies $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$, then from line 20 and onward it must likewise hold that $\hat{\mathbf{A}}\mathbf{P}^T = \mathbf{A}$. Therefore, at the input to DECODE on line 23 it follows that $\mathcal{S} = \{l \in [n] : \hat{\mathbf{a}}_l = d\} = [n]$ and as a result $\text{supp}(\mathbf{x}_i) \subset P(\mathcal{S})$ for all $i \in [N]$. In addition, due to Statement 1 of Lemma III.5, it follows that the input reconstructions to DECODE on line 23 satisfy the conditions of Lemma D.2. As a result, by Statement 2 of Lemma D.2, the sparse code $\hat{\mathbf{X}}$ returned by DECODE satisfies $\mathbf{P}\hat{\mathbf{X}} = \mathbf{X}$. This concludes the proof of the corollary. \square

We now proceed to the derivation of Lemma IV.2, to prove this result however we require the following result.

Lemma E.1. *Let $\mathbf{y} = \mathbf{A}\mathbf{x}$ where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{x} \in \chi_k^n$. For $b \in \{1, 2\}$, let $\hat{\mathbf{A}}^{(b)}$ and $\hat{\mathbf{x}}^{(b)}$ be such that there exists a $\mathbf{P}_b \in \mathcal{P}^{n \times n}$ $\text{supp}(\hat{\mathbf{A}}^{(b)}\mathbf{P}_b^T) \subseteq \text{supp}(\mathbf{A})$ and $\hat{x}_{P_b(l)}^{(b)} = x_{P_b(l)}$ for all $l \in \text{supp}(\hat{\mathbf{x}}^{(b)})$, where $P_b : [n] \rightarrow [n]$ is the row permutation corresponding to pre-multiplication by \mathbf{P}_b . In addition, assume that $\text{supp}(\hat{\mathbf{A}}^{(1)}\mathbf{P}_1^T) \subseteq \text{supp}(\hat{\mathbf{A}}^{(2)}\mathbf{P}_2^T)$ and $\text{supp}(\mathbf{P}_1\hat{\mathbf{x}}^{(1)}) \subseteq \text{supp}(\mathbf{P}_2\hat{\mathbf{x}}^{(2)})$. Considering the sparse codes returned by DECODE,*

$$\hat{\mathbf{x}}^{(b)}, _ \leftarrow \text{DECODE}(\mathbf{y}, m, \hat{\mathbf{A}}^{(b)}, \hat{\mathbf{x}}^{(b)}, \epsilon, d)$$

then $\text{supp}(\hat{\mathbf{x}}^{(1)}) \subseteq \text{supp}(\hat{\mathbf{x}}^{(2)})$.

Proof. For typographical ease let $\mathbf{V}^{(1)} = \hat{\mathbf{A}}^{(1)}\mathbf{P}_1^T$, $\mathbf{u}^{(1)} = \mathbf{P}_1\hat{\mathbf{x}}^{(1)}$, $\mathbf{V}^{(2)} = \hat{\mathbf{A}}^{(2)}\mathbf{P}_2^T$ and $\mathbf{u}^{(2)} = \mathbf{P}_2\hat{\mathbf{x}}^{(2)}$. Observe for $b \in \{1, 2\}$ that if

$$\mathbf{u}^{(b)}, _ \leftarrow \text{DECODE}(\mathbf{y}, m, \mathbf{V}^{(b)}, \mathbf{u}^{(b)}, \epsilon, d) \quad (3)$$

then $\mathbf{u}^{(b)} = \mathbf{P}_b\hat{\mathbf{x}}^{(b)}$. Given equation 3 it therefore suffices to prove $\text{supp}(\mathbf{u}^{(1)}) \subseteq \text{supp}(\mathbf{u}^{(2)})$. Let $\mathcal{Q}(t)$ be true iff at the end of iteration t of the while loop, starting on line 4 of DECODE, it holds that $\text{supp}(\mathbf{u}^{(1)}) \subseteq \text{supp}(\mathbf{u}^{(2)})$. We proceed by induction: observe that $\mathcal{Q}(0)$, which we may interpret as the event that at the start of the first iteration $\text{supp}(\mathbf{u}^{(1)}) \subseteq \text{supp}(\mathbf{u}^{(2)})$, is trivially true by the assumptions of the lemma. Therefore it suffices to prove that if $\mathcal{Q}(t-1)$ is true then $\mathcal{Q}(t)$ is true.

Assume $\mathcal{Q}(t-1)$ is true, analysing the residual computed on line 6 of DECODE it follows that

$$\begin{aligned}\mathbf{r}^{(1)} &= \mathbf{A}\mathbf{x} - \mathbf{V}^{(1)}\mathbf{u}^{(1)}, \\ \mathbf{r}^{(2)} &= \mathbf{A}\mathbf{x} - \mathbf{V}^{(2)}\mathbf{u}^{(2)}.\end{aligned}$$

Define, as in the proof of Lemma III.1, for $j \in [m]$ the sets $\Omega_j := \text{supp}(\tilde{\mathbf{a}}_j) \cap \text{supp}(\mathbf{x})$, $\Gamma_j^{(1)} := \text{supp}(\tilde{\mathbf{v}}_j^{(1)}) \cap \text{supp}(\mathbf{u}^{(1)})$ and $\Gamma_j^{(2)} := \text{supp}(\tilde{\mathbf{v}}_j^{(2)}) \cap \text{supp}(\mathbf{u}^{(2)})$. It follows for any entry $j \in [m]$ that

$$\begin{aligned}r_j^{(1)} &= \sum_{l \in \Omega_j} x_l - \sum_{h \in \Gamma_j^{(1)}} x_h = \sum_{l \in \Omega_j \setminus \Gamma_j^{(1)}} x_l \\ r_j^{(2)} &= \sum_{l \in \Omega_j} x_l - \sum_{h \in \Gamma_j^{(2)}} x_h = \sum_{l \in \Omega_j \setminus \Gamma_j^{(2)}} x_l\end{aligned}$$

As $\Gamma_j^{(1)} \subseteq \Gamma_j^{(2)}$, then $\Omega_j \setminus \Gamma_j^{(2)} \subseteq \Omega_j \setminus \Gamma_j^{(1)}$. Therefore any entry $r_j^{(2)}$ is a sum over a subset of the nonzeros in \mathbf{x}_i used in the sum of $r_{j,i}^{(1)}$. Consider any row index $\kappa \in [n]$ such that $u_\kappa^{(1)} = 0$ at the start of iteration t and $u_\kappa^{(1)} = x_\kappa$ at the end of iteration t . This implies that $f(x_\kappa, \mathbf{r}^{(1)}) \geq 2\epsilon d$ and therefore $\mathbf{v}_\kappa^{(1)} = \mathbf{v}_\kappa^{(2)} = \mathbf{a}_\kappa$. If $u_\kappa^{(2)} = x_\kappa$ at the start of iteration t then clearly $u_\kappa^{(2)} = x_\kappa$ at the end of iteration t . If $u_\kappa^{(2)} = 0$ at the start of iteration t , then $\kappa \in \Omega_j \setminus \Gamma_j^{(2)}$ for all $j \in \text{supp}(\mathbf{a}_\kappa)$. In addition, as $\Omega_j \setminus \Gamma_j^{(2)} \subseteq \Omega_j \setminus \Gamma_j^{(1)}$ then $f(x_\kappa, \mathbf{r}^{(2)}) \geq f(x_\kappa, \mathbf{r}^{(1)}) \geq 2\epsilon d$, and therefore line 24 of DECODE ensures $u_\kappa^{(2)} = x_\kappa$ at the end of iteration t . Therefore we have proved, assuming $\mathcal{Q}(t-1)$ is true, that if $u_\kappa^{(1)} = 0$ at the start of iteration t and $u_\kappa^{(1)} = x_\kappa$ at the end of iteration t , then $u_\kappa^{(2)} = x_\kappa$ at the end of iteration t . Therefore $\text{supp}(\mathbf{u}^{(1)}) \subseteq \text{supp}(\mathbf{u}^{(2)})$ at the end of iteration t and so $\mathcal{Q}(t)$ is true. Given that we have proved both the base case and the inductive step the proof is complete. \square

With this result in place we may now proceed to the proof of Lemma IV.2.

Lemma IV.2 (If ND-EBF successfully computes the matrix factors up to permutation then so will D-EBF) *Let $\mathbf{Y} = \mathbf{A}\mathbf{X}$ where $\mathbf{A} \in \mathcal{E}_{k,\epsilon,d}^{m \times n}$ with $\epsilon \leq 1/6$ and $\mathbf{X} \in \chi_k^{n \times N}$. Let $\hat{\mathbf{A}}^{(1)}$ and $\hat{\mathbf{X}}^{(1)}$ denote the reconstructions returned by ND-EBF and $\hat{\mathbf{A}}^{(2)}$ and $\hat{\mathbf{X}}^{(2)}$ the reconstructions returned by D-EBF. If there exists a $\mathbf{P}_1 \in \mathcal{P}^{n \times n}$ such that $\hat{\mathbf{A}}^{(1)}\mathbf{P}_1^T = \mathbf{A}$ and $\mathbf{P}_1\hat{\mathbf{X}}^{(1)} = \mathbf{X}$, then there exists a $\mathbf{P}_2 \in \mathcal{P}^{n \times n}$ such that $\hat{\mathbf{A}}^{(2)}\mathbf{P}_2^T = \mathbf{A}$ and $\mathbf{P}_2\hat{\mathbf{X}}^{(2)} = \mathbf{X}$.*

Proof. Before proceeding we emphasise certain points concerning the notation used in this proof. In particular, in what follows we use $\hat{\mathbf{A}}^{(1)}$ and $\hat{\mathbf{X}}^{(1)}$ to denote the reconstructions computed any point during the run-time of ND-EBF. Likewise, we use $\hat{\mathbf{A}}^{(2)}$ and $\hat{\mathbf{X}}^{(2)}$ to denote the reconstructions computed at any point during the run-time of D-EBF. In addition, let $\ell(\eta)$, where $\ell: [n] \rightarrow [n]$, denote the column of \mathbf{A} recovered by ND-EBF during iteration η , i.e., $\hat{\mathbf{a}}_\eta^{(1)} = \mathbf{a}_{\ell(\eta)}$. Furthermore, let $\tau(\eta)$, where $\tau: [n] \rightarrow [n]$, denote the index of the column reconstruction of $\mathbf{a}_{\ell(\eta)}$ in $\hat{\mathbf{A}}^{(2)}$. Note that as a result of Lemma III.5, at any point during the run-time of D-EBF either $\hat{\mathbf{a}}_{\tau(\eta)}^{(2)} = \mathbf{0}_m$ or $\mathbf{a}_{\ell(\eta)}^T \hat{\mathbf{a}}_{\tau(\eta)}^{(2)} > (1-2\epsilon)d$. Finally, let $\mathcal{Q}(\eta)$ be true iff at the end of iteration η of the while loop of D-EBF

$\hat{\mathbf{a}}_{\tau(\eta)}^{(2)} = \hat{\mathbf{a}}_{\ell}^{(1)}$ for all $\ell \in [\eta]$, and for any $\ell \in [\eta]$, $\kappa \in [N]$ such that $\hat{x}_{\ell,\kappa}^{(1)} \neq 0$, then $\hat{x}_{\tau(\eta),\kappa}^{(2)} = \hat{x}_{\ell,\kappa}^{(1)}$.

It suffices to prove that if ND-EBF recovers $\mathbf{a}_{\ell(\eta)}$ by the end of η th iteration of its while loop, then D-EBF also recovers $\mathbf{a}_{\ell(\eta)}$ by the end of η th iteration of its while loop. Indeed, if this is true then if ND-EBF recovers \mathbf{A} up to permutation then so will D-EBF. The result claimed then follows from Corollary IV.1.1. Our task therefore is to prove for all $\eta \in [n]$ that $\mathcal{Q}(\eta)$ is true. We proceed by induction. Observe that $\mathcal{Q}(0)$ is trivially true as both algorithms are initialised with $\hat{\mathbf{A}}^{(1)} = \hat{\mathbf{A}}^{(2)} = \mathbf{0}_{m \times n}$ and $\hat{\mathbf{X}}^{(1)} = \hat{\mathbf{X}}^{(2)} = \mathbf{0}_{n \times N}$. It suffices then to show for $\eta \in [n]$ that $\mathcal{Q}(\eta-1)$ true implies $\mathcal{Q}(\eta)$ true. Under the assumption that $\mathcal{Q}(\eta-1)$ is true, we now analyse the residuals computed at the beginning of iteration η of the respective while loops of ND-EBF and D-EBF. From Statement 1 of Lemma IV.1, then there exists a $\mathbf{P}_1 \in \mathcal{P}^{n \times n}$ such that $\mathbf{V}^{(1)} := \hat{\mathbf{A}}^{(1)}\mathbf{P}_1^T$ satisfies $\text{supp}(\mathbf{V}^{(1)}) \subseteq \text{supp}(\mathbf{A})$, $\mathbf{U}^{(1)} := \mathbf{P}_1\hat{\mathbf{X}}^{(1)}$ satisfies $\text{supp}(\mathbf{U}^{(1)}) \subseteq \text{supp}(\mathbf{X})$, and $u_{i,j}^{(1)} = x_{i,j}$ for all $(i,j) \in \text{supp}(\mathbf{U}^{(1)})$. Likewise, from Statement 1 of Lemma III.5, there exists a $\mathbf{P}_2 \in \mathcal{P}^{n \times n}$ such that $\mathbf{V}^{(2)} := \hat{\mathbf{A}}^{(2)}\mathbf{P}_2^T$ satisfies $\text{supp}(\mathbf{V}^{(2)}) \subseteq \text{supp}(\mathbf{A})$, $\mathbf{U}^{(2)} := \mathbf{P}_2\hat{\mathbf{X}}^{(2)}$ satisfies $\text{supp}(\mathbf{U}^{(2)}) \subseteq \text{supp}(\mathbf{X})$, and $u_{i,j}^{(2)} = x_{i,j}$ for all $(i,j) \in \text{supp}(\mathbf{U}^{(2)})$. From Lemma D.2, for any $l \geq \eta$ then $\mathbf{v}_{P_1(l)}^1$ and $\tilde{\mathbf{u}}_{P_1(l)}^1$ are zero vectors, therefore as $\mathcal{Q}(\eta-1)$ is true $\text{supp}(\mathbf{V}^{(1)}) \subseteq \text{supp}(\mathbf{V}^{(2)})$ and $\text{supp}(\mathbf{U}^{(1)}) \subseteq \text{supp}(\mathbf{U}^{(2)})$. The respective residuals inputted to EXTRACT&MATCH for each algorithm can be written as

$$\begin{aligned}\mathbf{R}^{(1)} &= \mathbf{A}\mathbf{X} - \mathbf{V}^{(1)}\mathbf{U}^{(1)}, \\ \mathbf{R}^{(2)} &= \mathbf{A}\mathbf{X} - \mathbf{V}^{(2)}\mathbf{U}^{(2)}.\end{aligned}$$

Considering an arbitrary column index $i \in [N]$ of the residuals, define, as in the proof of Lemma III.1, $\Omega_j := \text{supp}(\tilde{\mathbf{a}}_j) \cap \text{supp}(\mathbf{x}_i)$, $\Gamma_j^{(1)} := \text{supp}(\tilde{\mathbf{v}}_j^{(1)}) \cap \text{supp}(\mathbf{u}_i^{(1)})$ and $\Gamma_j^{(2)} := \text{supp}(\tilde{\mathbf{v}}_j^{(2)}) \cap \text{supp}(\mathbf{u}_i^{(2)})$, observing that $\Gamma_j^{(1)} \subseteq \Gamma_j^{(2)}$. The entry $(j,i) \in [m] \times [N]$ in each of these residuals can therefore be written as

$$\begin{aligned}r_{j,i}^{(1)} &= \sum_{l \in \Omega_j} x_l - \sum_{h \in \Gamma_j^{(1)}} x_h = \sum_{l \in \Omega_j \setminus \Gamma_j^{(1)}} x_l, \\ r_{j,i}^{(2)} &= \sum_{l \in \Omega_j} x_l - \sum_{h \in \Gamma_j^{(2)}} x_h = \sum_{l \in \Omega_j \setminus \Gamma_j^{(2)}} x_l\end{aligned}$$

As $\Gamma_j^{(1)} \subseteq \Gamma_j^{(2)}$, then $\Omega_j \setminus \Gamma_j^{(2)} \subseteq \Omega_j \setminus \Gamma_j^{(1)}$. Therefore, any entry $r_{j,i}^{(2)}$ is a sum over a subset of the nonzeros in \mathbf{x}_i used in the sum of $r_{j,i}^{(1)}$.

Suppose that at the end of iteration η of the while loop of D-EBF it holds that $\hat{\mathbf{a}}_{\tau(\eta)}^{(2)} = \mathbf{a}_{\ell(\eta)}$. Letting $\mathcal{S}_1 := \{\ell \in [n] : \text{supp}(\hat{\mathbf{a}}_\ell^{(1)}) = d\}$ and $\mathcal{S}_2 := \{\ell \in [n] : \text{supp}(\hat{\mathbf{a}}_\ell^{(2)}) = d\}$, then, as we are also assuming that $\mathcal{Q}(\eta-1)$ is true, it must follow that $\mathcal{S}_1 \subseteq \mathcal{S}_2$. As a result, by Lemma E.1, if for some column index $\kappa \in [N]$ it holds that $\hat{x}_{\eta,\kappa}^{(1)} = x_{\ell(\eta),\kappa}$, then $\hat{x}_{\tau(\eta),\kappa}^{(2)} = x_{\ell(\eta),\kappa}$ and as a result $\mathcal{Q}(\eta)$ must be true. Therefore, to prove the inductive step it suffices to show that $\hat{\mathbf{a}}_{\tau(\eta)}^{(2)} = \mathbf{a}_{\ell(\eta)}$ at the end of iteration η .

To this end, consider then an arbitrary $\gamma \in \text{supp}(\mathbf{a}_{\ell(\eta)})$. By the construction of ND-EBF and the assumption that ND-EBF

recovers \mathbf{A} up to permutation, then $\hat{a}_{\gamma,\eta}^{(1)} = 0$ at the start of iteration η and $\hat{a}_{\gamma,\eta}^{(1)} = 1$ at the end of iteration η . Therefore, considering the residual at the start of iteration η , there must exist a $\kappa \in [N]$ such that $f(x_{\ell(\eta),\kappa}, \mathbf{r}_\kappa^{(1)}) > (1 - 2\epsilon)d$ and $r_{\gamma,\kappa}^{(1)} = x_{\ell(\eta),\kappa}$. Recall now that $r_{j,\kappa}^{(2)}$ is a sum over a subset of the nonzeros in \mathbf{x}_κ used in the sum of $r_{j,\kappa}^{(1)}$, therefore either $r_{j,\kappa}^{(2)} = 0$ or $r_{j,\kappa}^{(2)} = x_{\ell(\eta),\kappa}$. If $r_{j,\kappa}^{(2)} = 0$ then it must already be true at the beginning of iteration η that $\hat{\mathbf{a}}_{\gamma,\tau(\eta)}^{(2)} = 1$. If at the beginning of iteration η it holds that $\hat{\mathbf{x}}_{\tau(\eta),\kappa}^{(2)} = x_{\ell(\eta),\kappa}$ and $r_{j,\kappa}^{(2)} = x_{\ell(\eta),\kappa}$, then the for loop on lines 30-35 of EXTRACT&MATCH ensures by the end of iteration η that $\hat{\mathbf{a}}_{\gamma,\tau(\eta)}^{(2)} = 1$. If at the beginning of iteration η it holds instead that $\hat{\mathbf{x}}_{\tau(\eta),\kappa}^{(2)} = 0$ and $r_{j,\kappa}^{(2)} = x_{\ell(\eta),\kappa}$, then $\ell(\eta) \in \Omega_p \setminus \Gamma_p$ for all $p \in \text{supp}(\mathbf{a}_{\ell(\eta)})$. Therefore, as $f(x_{\ell(\eta),\kappa}, \mathbf{r}_\kappa^{(1)}) > (1 - 2\epsilon)d$ then $f(x_{\ell(\eta),\kappa}, \mathbf{r}_\kappa^{(2)}) > (1 - 2\epsilon)d$. Therefore during the subroutine EXTRACT&MATCH on line 13 of D-EBF, the partial support $\mathbf{w} = g(x_{\ell(\eta),\kappa}, \mathbf{r}_\kappa^{(2)})$, which satisfies $\gamma \in \text{supp}(\mathbf{w})$ and $|\text{supp}(\mathbf{w})| > (1 - 2\epsilon)d$, is extracted. During the CLUSTER&ADD subroutine on line 20 then by Lemma III.5 it follows that $\hat{\mathbf{a}}_{\tau(\eta)}^{(2)} \leftarrow \sigma(\hat{\mathbf{a}}_{\tau(\eta)}^{(2)} + \mathbf{w})$. Therefore by the end of iteration η $\hat{a}_{\gamma,\tau(\eta)}^{(2)} = 1$. As $\gamma \in \text{supp}(\mathbf{a}_{\ell(\kappa)})$ was arbitrary, then it follows that $\hat{\mathbf{a}}_{\tau(\eta)}^{(2)} = \mathbf{a}_{\ell(\kappa)}$ at the end of iteration η . Therefore, if $\mathcal{Q}(\eta - 1)$ is true then $\mathcal{Q}(\eta)$ is true. This concludes the proof of the lemma. \square

The third lemma we prove in this appendix provides a lower bound on the probability that the random matrix X , as defined in Definition II.3, has at least $\beta \in \mathbb{Z}_{>0}$ nonzeros per row. This result plays a key role in the proof of Theorem II.3.

Lemma IV.3 (Nonzeros per row in X) For some $\beta \in \mathbb{N}$ and $\mu > 1$, if $N \geq \beta(\mu \frac{n}{k} \ln(n) + 1)$ then the probability that the random matrix X , as defined in the PSB model Definition II.3, has at least β nonzeros per row is more than $(1 - n^{-(\mu-1)})^\beta$.

Proof. Let $X(k, \beta, r^c)$ denote a random binary matrix constructed by concatenating i.i.d. random, binary column vectors whose supports are drawn uniformly at random across all possible supports with a fixed cardinality k , until there are at least β non-zeros per row. Here r^c indicates that the elements of the supports of each column are chosen without replacement. Let then $X(k, \beta, r)$ denote a random binary matrix constructed by concatenating i.i.d. random, binary column vectors with replacement, i.e., the support of each column is generated by taking k i.i.d samples uniform from $[n]$ with replacement, until there are at least β non-zeros per row. Let $\eta[X(k, \beta, r^c)]$ be the random variable which counts the number of columns of $X(k, \beta, r^c)$, note here that we adopt square brackets purely for typographical clarity. Also note that the supports of the columns of the random matrix X as defined in Definition II.3 are drawn in the same manner as those of the random matrix $X(k, \beta, r^c)$. The difference between these random matrices is that the number of columns N of X is fixed in advance instead of, as is the case for $X(k, \beta, r^c)$, drawing enough columns

until there are at least β nonzeros per row. Therefore

$$\begin{aligned} & \mathbb{P}(\text{"X has at least } \beta \text{ non-zeros per row"}) \\ &= \mathbb{P}(\eta[X(k, \beta, r^c)] \leq N). \end{aligned}$$

Our goal is to lower bound $\mathbb{P}(\eta[X(k, \beta, r^c)] \leq N)$. We first note this problem can be interpreted as a generalisation of the coupon collector and dixie cup problems [16], [17], [18], which correspond to analysing the expectations of $\eta[X(1, 1, r^c)]$ and $\eta[X(\beta, 1, r^c)]$ respectively. For context, in the classic coupon collector problem, objects are drawn one at a time from a set of size n , typically uniformly at random with replacement until each object has been seen once. We again emphasise that in the notation adopted here r^c refers to sampling the entries in the support of a given column without replacement, not the supports of the columns themselves which are mutually independent by construction. The generalisation studied here is equivalent to analysing the number of draws of k objects from $[n]$, uniform at random with replacement⁵, until each object has been picked at least β times with high probability. As the distribution of $\eta[X(k, \beta, r^c)]$ is challenging to analyse directly, we instead bound the quantity of interest using the classic coupon collector distribution, for which Chernoff style bounds can more easily be derived.

Using the notation described above, let $(X^{(i)}(k, 1, r^c))_{i=1}^\beta$ be mutually independent and identically distributed random matrices and define

$$\tilde{X}(k, \beta, r^c) := [X^{(1)}(k, 1, r^c), X^{(2)}(k, 1, r^c) \dots X^{(\beta)}(k, 1, r^c)].$$

As each of the submatrices $X^{(i)}(k, 1, r^c)$ may contain rows with more than 1 nonzero in them then clearly it holds that $\eta[X(k, \beta, r^c)] \leq \eta[\tilde{X}(k, \beta, r^c)]$, and therefore

$$\mathbb{P}(\eta[X(k, \beta, r^c)] \leq N) \geq \mathbb{P}(\eta[\tilde{X}(k, \beta, r^c)] \leq N).$$

Observe that

$$\bigcap_{i=1}^\beta \{\eta[X^{(i)}(k, 1, r^c)] \leq \frac{N}{\beta}\} \implies \{\eta[\tilde{X}(k, \beta, r^c)] \leq N\},$$

$$\{\eta[\tilde{X}(k, \beta, r^c)] \leq N\} \not\implies \bigcap_{i=1}^\beta \{\eta[X^{(i)}(k, 1, r^c)] \leq \frac{N}{\beta}\},$$

therefore $\bigcap_{i=1}^\beta \{\eta[X^{(i)}(k, 1, r^c)] \leq \frac{N}{\beta}\} \subseteq \{\eta[\tilde{X}(k, \beta, r^c)] \leq N\}$. As a result

$$\begin{aligned} \mathbb{P}(\eta[X(k, \beta, r^c)] \leq N) &\geq \mathbb{P}(\eta[\tilde{X}(k, \beta, r^c)] \leq N) \\ &\geq \mathbb{P}\left(\bigcap_{i=1}^\beta \{\eta[X^{(i)}(k, 1, r^c)] \leq \frac{N}{\beta}\}\right) \\ &= \prod_{i=1}^\beta \mathbb{P}\left(\eta[X^{(i)}(k, 1, r^c)] \leq \frac{N}{\beta}\right) \\ &= \left[\mathbb{P}\left(\eta[X(k, 1, r^c)] \leq \frac{N}{\beta}\right)\right]^\beta \end{aligned}$$

where the equalities on the third and fourth lines follow from the assumptions of mutual independence and identical distribution respectively. If the support of a column is sampled with

⁵Recall that r^c refers to sampling without replacement the entries in a given column, the supports of different columns are sampled with replacement.

replacement then this will potentially decrease its cardinality relative to if it were sampled without replacement. It follows that

$$\mathbb{P}\left(\eta[X(k, 1, r^c)] \leq \frac{N}{\beta}\right) \geq \mathbb{P}\left(\eta[X(k, 1, r)] \leq \frac{N}{\beta}\right).$$

Observe that sampling k elements of the support of a column with replacement is equivalent to unioning the supports of k mutually independent, identically distributed column vectors, each with only one element in their support, drawn uniformly at random across all n possible locations. To this end, consider the random matrix $X(1, 1)$, where the r argument has been dropped as sampling with or without replacement when there is only one nonzero per column are equivalent. The columns of this matrix are i.i.d. random vectors of dimension n with a single nonzero, whose location is drawn uniformly at random. Furthermore, sufficiently many of these columns are drawn and concatenated so as to ensure that $X(1, 1)$ has at least one nonzero per row. Denote the l th column of $X(1, 1)$ as Z_l . Let X' be a random binary matrix which is a deterministic function of $X(1, 1)$ with $\text{supp}(X'_l) := \bigcup_{i=(l-1)k+1}^{lk} \text{supp}(Z_i)$ for all $l \in \lceil \lceil \eta[X(1, 1)]/k \rceil \rceil$ ⁶. Therefore, given that for any $x \in \mathbb{R}$ it holds that $\lceil x \rceil \leq x + 1$, then

$$\begin{aligned} \mathbb{P}\left(\eta[X(k, 1, r^c)] \leq \frac{N}{\beta}\right) &\geq \mathbb{P}\left(\eta[X(k, 1, r)] \leq \frac{N}{\beta}\right) \\ &= \mathbb{P}\left(\eta[X'] \leq \frac{N}{\beta}\right) \\ &= \mathbb{P}\left(\lceil \eta[X(1, 1)]/k \rceil \leq \frac{N}{\beta}\right) \\ &\geq \mathbb{P}\left(\eta[X(1, 1)] \leq k\left(\frac{N}{\beta} - 1\right)\right). \end{aligned}$$

By inspection $\eta[X(1, 1)]$ is equivalent to the number of draws required to see each coupon at least once in the classic coupon collector problem, for which the following bound with $t > 0$ is well known (see e.g., [53]),

$$\mathbb{P}(\eta[X(1, 1)] > n \ln(n) + tn) < e^{-t}.$$

Letting $k\left(\frac{N}{\beta} - 1\right) = n \ln(n) + tn$, then rearranging gives $t = \frac{k(N-\beta) - \beta n \ln(n)}{\beta n}$. Assuming $n > 1$, then for some $\mu > 1$ if $N = \beta\left(\mu \frac{n}{k} \ln(n) + 1\right)$, then $t = (\mu - 1) \ln(n)$ and therefore

$$\begin{aligned} &\mathbb{P}\left(\eta[X(1, 1)] \leq k\left(\frac{N}{\beta} - 1\right)\right) \\ &= 1 - \mathbb{P}\left(\eta[X(1, 1)] > k\left(\frac{N}{\beta} - 1\right)\right) \\ &> 1 - e^{-(\mu-1)\ln(n)} \\ &= 1 - n^{-(\mu-1)}. \end{aligned}$$

Therefore, if X has $N \geq \beta\left(\mu \frac{n}{k} \ln(n) + 1\right)$ columns whose supports are sampled as described in Definition II.3, then

$$\mathbb{P}(\text{“X has at least } \beta \text{ non-zeros per row”}) > \left(1 - n^{-(\mu-1)}\right)^\beta$$

⁶Note that if $\eta[X(1, 1)]$ is not a multiple of k , then additional columns can be drawn so that the support of the last column of X' is still the union of the supports of k columns.

as claimed. \square

The final lemma we provide in this appendix lower bounds the probability that at each iteration of its while loop ND-EBF is able to recover a column of \mathbf{A} . This lemma also plays a key role in the proof of Theorem II.3.

Lemma IV.4 (Probability that ND-EBF recovers a column at iteration η) *Let $Y = AX$ as per the PSB model, detailed in Definition II.3. In addition, assume that $A \in \mathcal{E}_{k, \epsilon, d}^{m \times n}$ with $\epsilon \leq 1/6$, and that each row of X has at least $\beta(n) = (1 + 2\epsilon d)L(n)$ nonzeros, where $L : \mathbb{N} \rightarrow \mathbb{N}$. Suppose ND-EBF, Algorithm 13, is deployed to try and compute the factors of Y and that the algorithm reaches iteration $\eta \in [n]$ of the while loop. Then there is a unique column $A_{\ell(\eta)}$ of A , satisfying $A_{\ell(\eta)} \neq \hat{A}_l$ for all $l \in [\eta - 1]$, such that*

$$\mathbb{P}(\hat{A}_\eta = A_{\ell(\eta)}) > 1 - de^{-\tau(n)L(n)},$$

where $\tau(n) := -\ln\left(1 - \left(1 - \frac{b}{\alpha_1 n}\right)^{\alpha_1 n}\right)$ is $\mathcal{O}(1)$.

Proof. Consider the start of iteration η of the while loop of ND-EBF, then from Lemma IV.1 the following must hold: \hat{A}_l and \tilde{X}_l are nonzero iff $l < \eta$, if $l < \eta$ then $|\text{supp}(\hat{A}_l)| = d$, there exists a random permutation matrix $P \in \mathcal{P}^{n \times n}$ such that $\text{supp}(\hat{A}P^T) \subseteq \text{supp}(A)$, $\text{supp}(P\tilde{X}) \subseteq \text{supp}(X)$ and $\tilde{X}_{P(l), i} = X_{P(l), i}$ for all $i \in [N]$ and $l \in \text{supp}(\tilde{X}_l)$, where $P : [n] \rightarrow [n]$ denotes the row permutation caused by pre-multiplication with the random matrix P . With $V := \hat{A}P^T$ and $U := P\tilde{X}$ then by the construction of ND-EBF the residual can be expressed as

$$R = Y - \hat{A}\hat{X} = AX - VU = A(X - U).$$

Defining $Z := X - U$, as $X \in \chi_k^{n \times N}$ and $U_{l, i} = X_{l, i}$ for all $(l, i) \in \text{supp}(U)$, then $Z \in \chi_k^{n \times N}$, $\text{supp}(Z) \subseteq \text{supp}(X)$ and $Z_{l, i} = X_{l, i}$ for all $(l, i) \in \text{supp}(Z)$. Consider now a column $i \in [N]$ of the residual, inputted to EXTRACT&MATCH on line 12 of ND-EBF. By the construction of DECODE as well as Lemma D.2, any partial support extracted from R_i must originate from a column A_h such that $A_h \neq \hat{A}_l$ for all $l \in [\eta - 1]$, otherwise this partial support would have been extracted by DECODE during iteration $\eta - 1$. Let

$$\mathcal{T}_i := \{\alpha \in \mathbb{R} : f(\alpha, R_i) > (1 - 2\epsilon)d\}$$

denote the set of singleton values appearing more than $(1 - 2\epsilon)d$ extracted from R_i during EXTRACT&MATCH. Then using Lemma III.2 we can lower bound the number of partial supports extracted across all columns of the residual at

iteration η as follows.

$$\begin{aligned}
 \sum_{i \in [N]} |\mathcal{T}_i| &> \frac{1}{(1+2\epsilon)d} \sum_{i=1}^N |\text{supp}(Z_i)| \\
 &= \frac{1}{(1+2\epsilon)d} \sum_{j=1}^n |\text{supp}(\tilde{Z}_j)| \\
 &\geq \frac{1}{(1+2\epsilon)d} \sum_{j=\eta}^n |\text{supp}(\tilde{X}_j)| \\
 &\geq \frac{(n-\eta+1)\beta(n)}{(1+2\epsilon)d} \\
 &= (n-\eta+1)L(n).
 \end{aligned}$$

The inequality on the first line follows directly from Lemma III.2. The equality on the the second line is clearly true as $\sum_{i=1}^N |\text{supp}(Z_i)| = |\text{supp}(Z)| = \sum_{j=1}^n |\text{supp}(\tilde{Z}_j)|$. The inequality on the third line follows from Lemma IV.1. The fourth inequality and final equality on line 5 arise as a result of the fact that we are conditioning on there being at least $\beta(n) = (1+2\epsilon)dL(n)$ nonzeros per row of X . As there are more than $(n-\eta+1)L(n)$ partial supports, all of which belong to one of the $n-\eta+1$ columns of A not yet recovered, then by the pigeon hole principle there exists at least one cluster of partial supports with cardinality at least $L(n)$. Therefore, on line 17 of MAXCLUSTER&ADD it must follow that $|\mathcal{C}^*| \geq L(n)$.

For clarity, let the index of each partial support in \mathcal{C}^* correspond to the column of Y from which it originates, for example, W_i is extracted from R_i and the corresponding column of X is therefore X_i . Note that as $\epsilon \leq 1/6$ then each column of Y can only provide a maximum of one partial support with cardinality at least $2\epsilon d$ per column of A . Therefore this choice of indexing does not result in any ambiguity. Line 18 of MAXCLUSTER&ADD can be written equivalently as

$$\hat{A}_\eta = \sigma \left(\sum_{i \in \mathcal{C}^*} W_i \right).$$

By Lemmas III.4 and IV.1 there exists a unique column $A_{\ell(\eta)}$ such that $\text{supp}(\hat{A}_\eta) \subseteq \text{supp}(A_{\ell(\eta)})$. We now proceed to upper bound the probability that $\hat{A}_\eta \neq A_{\ell(\eta)}$. Given that $|\mathcal{C}^*| \geq L(n)$ then

$$\begin{aligned}
 &\mathbb{P}(\hat{A}_\eta \neq A_{\ell(\eta)}) \\
 &= \mathbb{P} \left(\left(\bigcup_{i=1}^{|\mathcal{C}^*|} \text{supp}(W_i) \right) \neq \text{supp}(A_{\ell(\eta)}) \right) \\
 &\leq \mathbb{P} \left(\left(\bigcup_{i=1}^{L(n)} \text{supp}(W_i) \right) \neq \text{supp}(A_{\ell(\eta)}) \right) \\
 &= \mathbb{P} \left(\bigcup_{j \in \text{supp}(A_{\ell(\eta)})} \left(j \notin \left(\bigcup_{i=1}^{L(n)} \text{supp}(W_i) \right) \right) \right)
 \end{aligned}$$

Given that $|\text{supp}(A_\ell)| = d$, then applying the union bound it follows that

$$\begin{aligned}
 &\mathbb{P}(\hat{A}_\eta \neq A_{\ell(\eta)}) \\
 &\leq d \mathbb{P} \left(j \notin \left(\bigcup_{i=1}^{L(n)} \text{supp}(W_i) \right) \mid j \in \text{supp}(A_{\ell(\eta)}) \right) \\
 &= d \mathbb{P} \left(\bigcap_{i=1}^{L(n)} (j \notin \text{supp}(W_i)) \mid j \in \text{supp}(A_{\ell(\eta)}) \right).
 \end{aligned}$$

If $j \notin \text{supp}(W_i)$ then there exists a $h \in \text{supp}(Z_i) \setminus \{\ell(\eta)\}$ such that $j \in \text{supp}(A_h)$. To be clear, if $j \in \text{supp}(A_{\ell(\eta)})$ but $j \notin \text{supp}(W_i)$, then there must exist a column A_h in the submatrix $A_{\text{supp}(Z_i) \setminus \{\ell(\eta)\}}$ which has a 1 in its j th entry or row, thereby resulting in $R_{j,i}$ not being a singleton value. Letting $\zeta := \lceil n/\lfloor m/d \rfloor \rceil$ then

$$\begin{aligned}
 &\mathbb{P} \left(\bigcap_{i=1}^{L(n)} (j \notin \text{supp}(W_i)) \mid j \in \text{supp}(A_{\ell(\eta)}) \right) \\
 &\leq \mathbb{P} \left(\bigcap_{i=1}^{L(n)} \left(j \in \left(\bigcup_{h \in \text{supp}(Z_i) \setminus \{\ell(\eta)\}} \text{supp}(A_h) \right) \right) \right) \\
 &\leq \mathbb{P} \left(\bigcap_{i=1}^{L(n)} \left(j \in \left(\bigcup_{h \in \text{supp}(Z_i) \setminus \{\ell(\eta)\}} \text{supp}(A_h) \right) \mid |\text{supp}(\tilde{A}_j)| = \zeta \right) \right) \\
 &\leq \mathbb{P} \left(\bigcap_{i=1}^{L(n)} \left(j \in \left(\bigcup_{h \in \text{supp}(X_i) \setminus \{\ell(\eta)\}} \text{supp}(A_h) \right) \mid |\text{supp}(\tilde{A}_j)| = \zeta \right) \right) \\
 &= \prod_{i=1}^{L(n)} \mathbb{P} \left(j \in \left(\bigcup_{h \in \text{supp}(X_i) \setminus \{\ell(\eta)\}} \text{supp}(A_h) \right) \mid |\text{supp}(\tilde{A}_j)| = \zeta \right).
 \end{aligned}$$

The inequality on the second line follows from the number of nonzeros per row of A being bounded by construction, therefore the events $j \in \text{supp}(A_{\ell(\eta)})$ and $j \in \text{supp}(A_h)$ for $h \neq \ell(\eta)$ are negatively correlated. The inequality on the third line of the above follows by considering the j th row to achieve the upper bound ζ of nonzeros per row. The inequality on the fourth line follows from the fact that $\text{supp}(Z_i) \subseteq \text{supp}(X_i)$ for all $i \in [N]$. Lastly, the equality on the fifth line follows from the mutual independence of the supports of the columns of X . Observe that

$$\begin{aligned}
 &j \in \left(\bigcup_{h \in \text{supp}(X_i) \setminus \{\ell(\eta)\}} \text{supp}(A_h) \right) \\
 &\iff (\text{supp}(X_i) \cap \text{supp}(\tilde{A}_j)) \setminus \{\ell(\eta)\} \neq \emptyset.
 \end{aligned}$$

Furthermore, note that by the construction of X in the PSB model, Definition II.3, that the probability that none of the supports of the columns of A with index in $\text{supp}(X_i) \setminus \{\ell(\eta)\}$ contain j is simply a count of the number of draws of $\text{supp}(X_i) \setminus \{\ell(\eta)\}$ that contain none of the column indices

in $\text{supp}(\tilde{A}_j) \setminus \{\ell(\eta)\}$, divided by the total number of possible draws of $\text{supp}(X_i) \setminus \{\ell(\eta)\}$. As a result

$$\begin{aligned} \mathbb{P}\left(j \in \left(\bigcup_{r \in \text{supp}(X_i) \setminus \{\ell(\eta)\}} \text{supp}(A_r)\right) \mid |\text{supp}(\tilde{A}_j)| = \zeta\right) \\ = 1 - \binom{n-\zeta}{k-1} \binom{n-1}{k-1}^{-1}, \end{aligned}$$

from which it follows that

$$\mathbb{P}(\hat{A}_l \neq A_l) \leq d \left(1 - \binom{n-\zeta}{k-1} \binom{n-1}{k-1}^{-1}\right)^{L(n)}. \quad (4)$$

To simplify equation 4, we bound the binomial coefficients as follows:

$$\begin{aligned} \binom{n-\zeta}{k-1} \binom{n-1}{k-1}^{-1} &= \frac{(n-\zeta)!(n-k)!}{(n-\zeta-k+1)!(n-1)!} \\ &= \frac{(n-\zeta)(n-\zeta-1)\dots(n-\zeta-k+2)}{(n-1)(n-2)\dots(n-k+1)} \\ &\geq \left(\frac{n-\zeta-k+2}{n-k+1}\right)^{k-1}. \end{aligned}$$

Therefore

$$\mathbb{P}(\hat{A}_l \neq A_l) \leq d \left(1 - \left(\frac{n-\zeta-k+2}{n-k+1}\right)^{k-1}\right)^{L(n)}. \quad (5)$$

Recalling from Definition II.3 that $k = \alpha_1 n$, $m = \alpha_2 n$ and $\zeta = \lceil n/\lfloor m/d \rfloor \rceil \leq \frac{nd}{m+d} + 1$, then

$$\begin{aligned} \left(\frac{n-\zeta-k+2}{n-k+1}\right)^{k-1} &\geq \left(1 - \frac{nd}{(m+d)(n-k+1)}\right)^{\alpha_1 n - 1} \\ &> \left(1 - \frac{nd}{m(n-k)}\right)^{\alpha_1 n} \\ &= \left(1 - \frac{d}{\alpha_2(1-\alpha_1)n}\right)^{\alpha_1 n} \\ &= \left(1 - \frac{b}{\alpha_1 n}\right)^{\alpha_1 n} \end{aligned}$$

where $b := \frac{d\alpha_1}{\alpha_2(1-\alpha_1)}$. As a result

$$\begin{aligned} \mathbb{P}(\hat{A}_\eta \neq A_{\ell(\eta)}) &\leq d \left(1 - \left(\frac{n-\zeta-k+2}{n-k+1}\right)^{k-1}\right)^{L(n)} \\ &\leq d \left(1 - \left(1 - \frac{b}{\alpha_1 n}\right)^{\alpha_1 n}\right)^{L(n)} \\ &= de^{-\tau(n)L(n)} \end{aligned}$$

where the final equality is derived by taking the exponential of the logarithm, and $\tau(n) := -\ln\left(1 - \left(1 - \frac{b}{\alpha_1 n}\right)^{\alpha_1 n}\right)$. Note that $\tau(x)$ for $x \in \mathbb{R}_{>0}$ is continuous and $\lim_{n \rightarrow \infty} \tau(n) = -\ln(1 - e^{-b})$, therefore $\tau(n) = \mathcal{O}(1)$ as claimed. \square

REFERENCES

- [1] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [2] E. J. Candes and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [3] E. Candes and T. Tao, "Decoding by linear programming," *Information Theory, IEEE Transactions on*, vol. 51, pp. 4203–4215, 01 2006.
- [4] E. J. Candes and J. Romberg, "Quantitative robust uncertainty principles and optimally sparse decompositions," *Foundations of Computational Mathematics*, vol. 6, no. 2, pp. 227–254, 2006. [Online]. Available: <https://doi.org/10.1007/s10208-004-0162-x>
- [5] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.20124>
- [6] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, pp. 1289–1306, 2004.
- [7] M. Elad, *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*, 1st ed. Springer Publishing Company, Incorporated, 2010.
- [8] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. Birkhäuser Basel, 2013.
- [9] S. Gleichman and Y. C. Eldar, "Blind compressed sensing," *IEEE Transactions on Information Theory*, vol. 57, no. 10, pp. 6958–6975, 2011.
- [10] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4311–4322, 2006.
- [11] H. Lee, A. Battle, R. Raina, and A. Ng, "Efficient sparse coding algorithms," in *Advances in Neural Information Processing Systems*.
- [12] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Supervised dictionary learning," in *Proceedings of the 21st International Conference on Neural Information Processing Systems*, ser. NIPS'08, 2008, p. 1033–1040.
- [13] B. A. Olshausen and D. J. Fieldt, "Sparse coding with an overcomplete basis set: a strategy employed by v1," *Vision Research*, vol. 37, pp. 3311–3325, 1997.
- [14] R. Berinde, A. C. Gilbert, P. Indyk, H. Karloff, and M. J. Strauss, "Combining geometry and combinatorics: A unified approach to sparse signal recovery," in *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, Sep. 2008, pp. 798–805.
- [15] S. Jafarpour, W. Xu, B. Hassibi, and A. R. Calderbank, "Efficient and robust compressed sensing using optimized expander graphs," *IEEE Transactions on Information Theory*, vol. 55, pp. 4299–4308, 2009.
- [16] M. Ferrante and A. Tagliavini, "On the coupon-collector's problem with several parallel collections," *arXiv e-prints*, Sep. 2016.
- [17] A. Doumas and V. Papanicolaou, "The coupon collector's problem revisited: generalizing the double dixie cup problem of Newman and Shepp," *ESAIM: Probability and Statistics*, vol. 20, June 2016.
- [18] W. Stadje, "The collector's problem with group drawings," *Advances in Applied Probability*, vol. 22, 12 1990.
- [19] B. Bah and J. Tanner, "On the construction of sparse matrices from expander graphs," *Frontiers in Applied Mathematics and Statistics*, vol. 4, p. 39, 2018. [Online]. Available: <https://www.frontiersin.org/article/10.3389/fams.2018.00039>
- [20] A. Cohen, W. Dahmen, and R. DeVore, "Compressed sensing and best k-term approximation," *Journal of the American Mathematical Society*, vol. 22, no. 1, pp. 211–231, Jan. 2009.
- [21] D. L. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Transactions on Information Theory*, vol. 47, no. 7, pp. 2845–2862, 2001.
- [22] D. L. Donoho and M. Elad, "Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization," *Proceedings of the National Academy of Sciences*, vol. 100, no. 5, pp. 2197–2202, 2003. [Online]. Available: <https://www.pnas.org/content/100/5/2197>
- [23] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, pp. 253–263, 12 2008.
- [24] J. Blanchard, C. Cartis, and J. Tanner, "Compressed sensing: how sharp is the restricted isometry property?" *SIAM Review*, vol. 53, 04 2010.
- [25] A. M. Tillmann and M. E. Pfetsch, "The computational complexity of the restricted isometry property, the nullspace property, and related

- concepts in compressed sensing," *IEEE Transactions on Information Theory*, vol. 60, no. 2, pp. 1248–1259, 2014.
- [26] J. C. Ye, "Compressed sensing mri: a review from signal processing perspective," *BMC Biomedical Engineering*, vol. 1, no. 1, p. 8, 2019. [Online]. Available: <https://doi.org/10.1186/s42490-019-0006-z>
- [27] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk, "Single-pixel imaging via compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 83–91, 2008.
- [28] L. A. Bassalygo and M. S. Pinsker, "On the complexity of an optimal non-blocking commutation scheme without reorganization," *Problems Information Transmission*, vol. 9, pp. 64–66, 1973.
- [29] M. Capalbo, O. Reingold, S. Vadhan, and A. Wigderson, "Randomness conductors and constant-degree lossless expanders," in *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*. ACM, 2002, pp. 659–668.
- [30] V. Guruswami, C. Umans, and S. Vadhan, "Unbalanced expanders and randomness extractors from parvaresh–vardy codes," *J. ACM*, vol. 56, no. 4, Jul. 2009. [Online]. Available: <https://doi.org/10.1145/1538902.1538904>
- [31] R. Berinde, P. Indyk, and M. Ruzic, "Practical near-optimal sparse recovery in the ℓ_1 norm," in *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, 2008, pp. 198–205.
- [32] R. Berinde and P. Indyk, "Sequential sparse matching pursuit," in *Proceedings of the 47th Annual Allerton Conference on Communication, Control, and Computing*, ser. Allerton '09. IEEE Press, 2009, p. 36–43.
- [33] W. Xu and B. Hassibi, "Efficient compressive sensing with deterministic guarantees using expander graphs," in *2007 IEEE Information Theory Workshop*, 2007, pp. 414–419.
- [34] R. Mendoza-Smith and J. Tanner, "Expander ℓ_0 -decoding," *Applied and Computational Harmonic Analysis*, vol. 45, no. 3, pp. 642 – 667, 2018. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1063520317300210>
- [35] J. Chen and X. Huo, "Theoretical results on sparse representations of multiple-measurement vectors," *IEEE Transactions on Signal Processing*, vol. 54, no. 12, pp. 4634–4643, 2006.
- [36] S. F. Cotter, B. D. Rao, Kjersti Engan, and K. Kreutz-Delgado, "Sparse solutions to linear inverse problems with multiple measurement vectors," *IEEE Transactions on Signal Processing*, vol. 53, no. 7, pp. 2477–2488, 2005.
- [37] M. Mishali and Y. C. Eldar, "Reduce and boost: Recovering arbitrary sets of jointly sparse vectors," *IEEE Transactions on Signal Processing*, vol. 56, no. 10, pp. 4692–4702, 2008.
- [38] T. Tao and V. H. Vu, *Additive Combinatorics*, ser. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2006.
- [39] N. Alon, "Eigenvalues and expanders," *Combinatorica*, vol. 6, no. 2, pp. 83–96, Jun. 1986.
- [40] T. Høholdt and H. Janwa, "Eigenvalues and expansion of bipartite graphs," *Designs, Codes and Cryptography*, vol. 65, no. 3, pp. 259–273, 2012. [Online]. Available: <https://doi.org/10.1007/s10623-011-9598-6>
- [41] E. Elhamifar and R. Vidal, "Sparse subspace clustering," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 2790–2797.
- [42] L. Parsons, E. Haque, and H. Liu, "Subspace clustering for high dimensional data: a review," *SIGKDD Explorations*, vol. 6, pp. 90–105, 2004.
- [43] R. Vidal, "A tutorial on subspace clustering."
- [44] A. Agarwal, A. Anandkumar, and P. Netrapalli, "A clustering approach to learning sparsely used overcomplete dictionaries," *IEEE Transactions on Information Theory*, vol. 63, no. 1, pp. 575–592, 2017.
- [45] S. Arora, A. Bhaskara, R. Ge, and T. Ma, "More algorithms for provable dictionary learning," *CoRR*, vol. abs/1401.0579, 2014. [Online]. Available: <http://arxiv.org/abs/1401.0579>
- [46] S. Arora, R. Ge, and A. Moitra, "New algorithms for learning incoherent and overcomplete dictionaries," in *Proceedings of The 27th Conference on Learning Theory*, ser. Proceedings of Machine Learning Research, vol. 35. PMLR, Jun 2014, pp. 779–806. [Online]. Available: <http://proceedings.mlr.press/v35/arora14.html>
- [47] B. Barak, J. A. Kelner, and D. Steurer, "Dictionary learning and tensor decomposition via the sum-of-squares method," in *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, ser. STOC '15, 2015, p. 143–151. [Online]. Available: <https://doi.org/10.1145/2746539.2746605>
- [48] D. Spielman, H. Wang, and J. Wright, "Exact recovery of sparsely-used dictionaries," *IJCAI International Joint Conference on Artificial Intelligence*, vol. 23, June 2012.
- [49] J. Sun, Q. Qu, and J. Wright, "Complete dictionary recovery over the sphere I: Overview and the geometric picture," *IEEE Transactions on Information Theory*, vol. 63, no. 2, pp. 853–884, 2017.
- [50] —, "Complete dictionary recovery over the sphere II: Recovery by riemannian trust-region method," *IEEE Transactions on Information Theory*, vol. 63, no. 2, pp. 885–914, 2017.
- [51] S. Arora, A. Bhaskara, R. Ge, and T. Ma, "Provable bounds for learning some deep representations," in *Proceedings of the 31st International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 32, no. 1. PMLR, June 2014, pp. 584–592. [Online]. Available: <http://proceedings.mlr.press/v32/arora14.html>
- [52] R. Mendoza-Smith, J. W. Tanner, and F. Wechsung, "A robust parallel algorithm for combinatorial compressed sensing," *IEEE Transactions on Signal Processing*, vol. 66, no. 8, pp. 2167–2177, 2018.
- [53] B. Doerr, *Probabilistic Tools for the Analysis of Randomized Optimization Heuristics*. Springer International Publishing, 2020, pp. 1–87.