There and Back Again:
A General Approach to Learning Sparse Models

Vatsal Sharan
Stanford University
vsharan@stanford.edu

Kai Sheng Tai
Stanford University
kst@cs.stanford.edu

Peter Bailis
Stanford University
pbailis@cs.stanford.edu

Gregory Valiant
Stanford University
valiant@stanford.edu

Abstract

We propose a simple and efficient approach to learning sparse models. Our approach consists of (1) projecting the data into a lower dimensional space, (2) learning a dense model in the lower dimensional space, and then (3) recovering the sparse model in the original space via compressive sensing. We apply this approach to Non-negative Matrix Factorization (NMF), tensor decomposition and linear classification—showing that it obtains $10\times$ compression with negligible loss in accuracy on real data, and obtains up to $5\times$ speedups. Our main theoretical contribution is to show the following result for NMF: if the original factors are sparse, then their projections are the sparsest solutions to the projected NMF problem. This explains why our method works for NMF and shows an interesting new property of random projections: they can preserve the solutions of non-convex optimization problems such as NMF.

1 Introduction

In what settings can sparse models be trained more efficiently than their dense analogs? Sparse models commonly arise in feature selection (e.g., diagnostic tasks [1, 2, 3, 4]), in problems with few, highly indicative features (e.g., time-series prediction [5, 6]), and in resource-constrained environments (i.e., a sparse model is needed for more efficient inference [7, 8]). In these tasks, sparsity can reduce the data requirements for a desired predictive power, and improve interpretability and inference speed; and a range of popular techniques provide means of learning these sparse models from dense and/or high-dimensional data [9, 10]. In this work, we study an additional dimension of sparse statistical models: the extent to which a model's sparsity can be leverage for a *computational* advantage.

We present a general and efficient approach to learning sparse models that explicitly leverages the target model sparsity. We adapt now-classic techniques from compressive sensing [11, 12] in a simple framework for learning sparse models from high-dimensional data: (1) project data to a lower-dimensional subspace using a random projection, (2) learn a dense model in this lower-dimensional subspace, then (3) recover a sparse model in the original dimensionality using compressive sensing. This approach has several benefits. First, it allows us to perform the bulk of model training in an often much lower-dimensional dense subspace, offering in significant efficiency gains when data is high-dimensional and the target model is sparse. Second, it is useful in settings where the dataset is very large and cannot be fit in memory; our approach allows us to work with a much smaller, low dimensional projection but still recover the original model. Third, it provides a conceptually simple approach to learning sparse models, combining powerful ideas from dimensionality reduction and compressive sensing.

Although low-dimensional projections are extremely popular tools for speeding up a variety of optimization and learning problems [13, 14, 15], our framework is conceptually different, as we

are interested in being able to recover the sparse solution in the original space from the (possibly dense) solution in the lower-dimensional space. Note that this is not even information theoretically possible for many tasks even if the original model or solution is sparse: if the solution in the lower-dimensional space is not a projection of the original solution, then performing a compressive sensing based recovery has no hope of recovering the original solution (even if the original solution is sparse).

Our contributions. We develop a simple framework for learning sparse models and investigate it in the context of Non-negative Matrix Factorization (NMF), tensor decomposition, and linear classification. We show that our approach recovers the true solution with a negligible error across several synthetic and real datasets. Empirically, sparse learning allows $10\times$ compression with negligible loss in accuracy on real data, and obtains up to $5\times$ speedups in terms of wall clock time to obtain the same accuracy.

Theoretically establishing *why* such speed-ups and high fidelity reconstructions are possible and justified requires care: we establish novel uniqueness guarantees for NMF which demonstrate that if the original factors are sparse, then the projection of the factors of the original, high-dimensional NMF problem is the *sparsest* possible factorization for the low-dimensional projected NMF problem. As NMF algorithms implicitly [16] or explicitly [17] learn sparse factorizations, our results explain why the factorization is preserved by projection, explaining the empirical success of our procedure. Our results do not assume separability [18] of the NMF factors, a commonly used—but very strong—assumption for showing uniqueness of NMF. We also derive simple conditions for uniqueness and provable recovery for tensor decomposition in the projected space. These theoretical insights, backed by our empirical findings, illustrate the benefits of random projections to learn sparse models—learning is faster, without compromising model quality.

Our results on the uniqueness of NMF after random projection—and our empirical findings—open a new perspective on the properties of random projections. It is well known that random projections can preserve distances and other geometric properties (e.g., [19]), but our results show that random projections can also preserve the solutions to non-convex optimization problems such as NMF. Our approach successfully leverages this property to efficiently learn sparse models, and we believe this aproach has several promising applications beyond our target settings.

We briefly outline the three main classes of sparse models we consider in this paper:

Non-negative Matrix Factorization (NMF). The goal of NMF is to approximate a non-negative matrix $M \in \mathbb{R}^{m \times n}$ with a product of two non-negative matrices $W \in \mathbb{R}^{m \times r}$ and $H \in \mathbb{R}^{r \times n}$. The matrix W is regarded as the dictionary, and H stores the combination weights of each column of M in terms of the dictionary W. Applications of NMF include document clustering [20] and image analysis [16, 21]. A crucial property of NMF that enables it to learn good representation of the data is that the coefficients H are typically sparse [16, 17], so that each document or image is expressed as a combination of few basis vectors in the dictionary. In settings such as document clustering or image classification, the columns of M correspond to samples from the data set (e.g., documents or images). In many scenarios, $n \gg m$, as the number of samples is large for big datasets. We are interested in this setting, where n is very large but the factor H is sparse.

Tensor Decomposition. The goal of tensor decomposition is to decompose a tensor $T \in \mathbb{R}^{m_1 \times m_2 \times n}$ with rank r in terms of its factors $T = \sum_{i=1}^{r} A_i \otimes B_i \otimes C_i$. Here, A_i denotes the ith column of the factor $A \in \mathbb{R}^{m_1 \times r}$ and $\otimes$ denotes the outer product: $a, b, c \in \mathbb{R}^n$ then $a \otimes b \otimes c \in \mathbb{R}^{n \times n \times n}$ and $(a \otimes b \otimes b)_{ijk} = a_i b_j c_k$. As in NMF, we consider settings where the third mode corresponding to the factor C represents data points and is very large—but sparse. The first two modes could correspond to features for each data point; for example in a topic modeling setup [22], the first two modes could represent word co-occurrences, with the third mode representing different documents.

Linear Classification. We also consider the problem of finding the linear classifier to classify m data points in n dimensions. In modern datasets, the number of features n is often very large, containing numerous irrelevant features [1, 2, 3, 4]. In these settings, the best (and/or most interpretable) classifier is sparse, as it only relies on the few features which are predictive/indicative.

2   Related Work

Sparse Recovery. In the compressive sensing or sparse recovery framework, there is a sparse signal $x \in \mathbb{R}^n$, and we are given measurements $Ax, A \in \mathbb{R}^{d \times n}$, where $d \ll n$. The goal is to recover x using the measurements Ax, and the property that x is sparse. Compressive sensing has been

extensively studied by several different communities, and many algorithms have been proposed which have a high compression rate and allow fast decoding [23, 24, 11]. Particularly relevant to our work is the work on sparse recovery using sparse projection matrices [25, 26, 27, 12], and we use these results to efficiently recover back the solution in the original space from the solution in the projected space. As is evident, the key difference of our approach from compressive sensing is that we do not have direct access to the measurements Aw^* of the solution w^* in the projected space, and instead hope to recover an approximation to Aw^* by solving the optimization problem in the projected space.

Non-negative Matrix Factorization. Many algorithms have been proposed for NMF, the most popular ones are multiplicative updates [16] and Alternating Least Squares (ALS). As mentioned in the introduction, for many modern datasets $n \gg m$ as the number of samples is very large. Approaches such as ALS and multiplicative weights do not scale to this online setting where the data set is very large and does not fit in memory. For the online setting, many algorithms [28, 29, 30] have been proposed for the *dictionary learning problem*, where the goal is to only learn the matrix W, but these approaches are unable to learn the coefficient matrix H. In the context of using random projection for NMF, Wang and Li [31] obtained speedups compared to ALS by projecting down the matrix via a random projection. However, they do not attempt to directly recover the factors in the original space by using the factors in the projected space and exploiting properties of the original solution such as its sparsity.

Tensor Decomposition. The most popular algorithm for tensor decomposition is Alternating Least Squares (ALS), which proceeds by fixing two of the factor matrices and solving for the third one via a least squares problem, and continuing this till convergence. Recently, there has been a burst of exciting work on speeding up ALS via sampling and sketching approaches [32, 33, 34, 35]. These approaches obtain significant speedups over the naive ALS algorithm, but they do not leverage the fact that the factors of the tensor are sparse, which is what we aim to exploit in our paper.

Random Projections for Classification. Say we have data points $x \in \mathbb{R}^n$ with labels $y \in \{0, 1\}$ which are separated by a linear classifier $w^* \in \mathbb{R}^n$, therefore $y = \text{sign}(x^T w^*)$. If we perform a random projection of the points to d dimensions, are they still separable? A number of papers [36, 37, 38] have examined this basic question, and a one line summary of the results is that if the original data points are separable with a large margin, then a random projection preserves the margin. In our framework however, the goal is not to just classify accurately in the projected space, but to recover the optimal classifier in the projected space. Zhang et al. [39] also consider the question of obtaining the optimal classifier in the original space from the optimal classifier in the projected space. They show that a novel approach based on the dual of the optimization problem in the projected space can recover the original classifier if the data matrix is approximately low rank.[1] In our framework, we instead aim to employ the sparsity of the original classifier w^* in order to be able to recover it.

3 Our Approach

We instantiate our framework in the context of NMF, tensor decomposition and linear classification. In each case, we obtain a lower-dimensional problem via a random projection (described in detail in Sec. 3.1). Denote the random projection matrix as $P \in \mathbb{R}^{d \times n}$, where n is the original dimension and $d < n$. After solving the projected problem, we obtain an approximation to the solution in the original space by solving a sparse recovery problem via the following linear program (LP):

$$\hat{w}^* = \underset{w}{\text{minimize}} \quad \|w\|_1, \quad \text{subject to} \quad Pw = \tilde{w}^* \tag{1}$$

where $w \in \mathbb{R}^n$ and $\tilde{w}^* \in \mathbb{R}^d$ is the solution of the projected problem[2]. In the compressive sensing literature, this technique is also known as *basis pursuit* [40].

We now describe the specifics of each application of our method.

Non-negative Matrix Factorization. Consider the problem of factorizing a non-negative matrix $M \in \mathbb{R}^{m \times n}$ into a product of non-negative matrices $M = WH$. We are interested in the setting where matrix H is sparse and $n \gg m$. We first project M to $\tilde{M} = MP^T$. This corresponds to projecting each row of H from n dimensions to d dimensions. We then factorize $\tilde{M} = \tilde{W}\tilde{H}$ by

---

[1]Zhang et al. [39] also obtain guarantees when the true classifier w^* is sparse, but require that the data points have most of their mass on the non-zero entries of w^* in this case, and hence are actually close to low rank.

[2]In settings like NMF where the desired solution is non-negative, we can add the additional constraint $x \geq 0$ to Eq. 1. In practice, we find that projecting the solutions of Eq. 1 post-hoc onto the non-negative orthant is often an effective heuristic. Omitting the additional constraint typically yields faster solution times.

solving the NMF problem for $\tilde{M}$. We set our estimate $\hat{W}$ of W to be $\tilde{W}$. To recover H from $\tilde{H}$, we solve the compressive sensing problem in Eq. 1 for each row of H_i of H. Therefore, for a rank-r approximation, we solve r instances of Eq. 1. The outputs of the compressive sensing step are our estimates of the rows of H.

Tensor decomposition. The goal here is to decompose a tensor $T \in \mathbb{R}^{m_1 \times m_2 \times n}$ in terms of r factors: $T = \sum_{i=1}^{r} A_i \otimes B_i \otimes C_i$ where $n \gg m_1, m_2$ and C is sparse. We project the tensor T along the 3rd mode to a projected tensor $\tilde{T} \in \mathbb{R}^{m_1 \times m_2 \times d}$, which has factorization $\tilde{T} = \sum_{i=1}^{r} A_i \otimes B_i \otimes \tilde{C}_i$, where $\tilde{C}_i = PC_i$. To understand this projection, denote $T_{(n)}$ as the mode n matricization of the tensor T, which is the flattening of the tensor T along the nth mode obtained by stacking the matrix slices together as columns of $T_{(n)}$. To project T, we project its mode 3 matricization $T_{(3)} \in \mathbb{R}^{m_1 m_2 \times n}$ to obtain the projection of the mode 3 matricization $\tilde{T}_{(3)} \in \mathbb{R}^{m_1 m_2 \times d}$ of $\tilde{T}$ by computing $\tilde{T}_{(3)} = T_{(3)} P^T$. We decompose $\tilde{T}$ in the projected space to obtain the factors $\tilde{A}, \tilde{B}, \tilde{C}$. Our estimates $\hat{A}, \hat{B}, \hat{C}$ for the original factors are $\hat{A} = \tilde{A}, \hat{B} = \tilde{B}$ and $\hat{C}$ is obtained by solving the compressive sensing problem in Eq. 1 for each column C_i. As above, this yields r instances of Eq. 1.

Linear classification. For the problem of classifying data points $x^{(i)} \in \mathbb{R}^n$, $1 \leq i \leq m$, we first perform a random projection to project the points into d dimensions using the projection matrix P. Hence the projected points $\tilde{x}^{(i)} \in \mathbb{R}^d$ are given by $\tilde{x}^{(i)} = Px^{(i)}$. We then solve the classification problem in the projected space to learn a classifier $\tilde{w}^*$. Finally, we use compressive sensing to recover an estimate of the true classifier $\hat{w}^*$ from $\tilde{w}^*$.

3.1 Choosing the projection matrix P

There are several requirements which our projection matrix P must satisfy. The first, more obvious, requirement is that we should be able to solve the compressive sensing problem to recover the sparse solution in the original space from the solution in the projected space. A slightly more subtle requirement is that P should be chosen such that if w^* is the optimal solution in the original space, then the solution in the projected space is $\tilde{w}^* = Pw^*$, so that compressive sensing based recovery can then be carried out. This requires uniqueness of the optimization problem in the projected space, the requirement for which could depend on the original optimization problem. In Section 4 we analyze these requirements for NMF and tensor decomposition. Finally, for computational efficiency, we want projections which are efficient to compute, and which have the property that the sparse recovery problem can be solved efficiently using them.

Consider a sparse projection matrix $P \in \{0,1\}^{d \times n}$ such that every column of P has exactly p non-zero entries, we refer to p as the number of buckets, because each coordinate in the original n dimensional space contributes to p locations in the projected d dimensional space. Let the support of the non-zero entries of each column of P be chosen independently and uniformly at random. The key property of P which enables compressive sensing is that P is the adjacency matrix of a *bipartite expander* [41]. Consider a bipartite graph G with n nodes on the left part and d nodes on the right part such that every node in the left part has degree p. We call G a $(\gamma n, \alpha)$ expander if every subset of at most $t \leq \gamma n$ nodes in U has at least αtp neighbors in V. It is not difficult to show that a randomly chosen matrix P with p non-zero entries per column is the adjacency matrix of a $(\gamma n, 4/5)$ expander for $\gamma n = d/(pe^5)$ with high probability if $p = O(\log n)$. We formally prove this as Lemma 5 in Appendix D. Note that as P is sparse, the projection step can be carried out efficiently.

For the basis pursuit problem described in (1), seminal results in compressive sensing [23, 24, 11] show that using dense projection P guarantees recovery with $d = O(k \log n)$ if the original solution is k-sparse. This property was subsequently extended to sparse projection matrices, and we will leverage this fact. Theorem 3 of Berinde et al. [12] shows that if x is k-sparse, then with P as defined above x can be recovered from Px using the LP if P is the adjacency matrix of a $(2k, 4/5)$ expander. A randomly chosen P with $d = O(k \log n)$ will be a $(2k, 4/5)$ expander with high probability (Lemma 5), and hence can be used for sparse recovery of k-sparse signals. Efficient recovery algorithms [26, 27] which are much faster than solving the LP are also known for sparse projection matrices. Hence using P defined as above ensures that the sparse recovery step is computationally efficient and provably works if the original solution w^* is k-sparse and the solution $\tilde{w}^*$ in the projected space is $\tilde{w}^* = Pw^*$.

However, we still need to ensure that the solution $\tilde{w}^* = Pw^*$. We analyze when this is the case for NMF and tensor decomposition next in Section 4. We summarize these conditions in Table 1 for

Table 1: Parameters for the projection matrix $P \in \{0, 1\}^{d \times n}$ suggested by our theoretical analysis for NMF and tensor decomposition (TD). p is the number of non-zero entries per column of P.

Task	Parameter	Number of buckets p	Projection dimension d
NMF	r: rank of matrix	$O(\log n)$	$O((r + k) \log n)$
TD	r: rank of tensor	$O(\log n)$	$\max\{r, O(k \log n)\}$

NMF and tensor decomposition. Note that the bounds suggested by our theoretical results are tight up to logarithmic factors, as projecting the matrix or the tensor to dimension below the rank will not preserve uniqueness of the decomposition.

4 Guarantees for our approach

In this section, we will show uniqueness guarantees for performing non-negative matrix factorization and tensor decomposition on the projected matrix or tensor. This will establish that the factors in the projected space correspond to a projection of the factors in the original space, a prerequisite for being able to use compressive sensing to recover the factors in the original space.

4.1 Guarantees for non-negative matrix factorization

Donoho and Stodden [18] showed that non-negative matrix factorization is unique under the separability condition. However, separability is a very strong assumption. Instead of using separability, we will show uniqueness results for NMF using sparsity of the weight matrix H. Our uniqueness results for NMF in the original space (Part (a) of Theorem 1) are similar in principle to uniqueness results obtained for sparse coding by Spielman et al. [42]. We show that if a matrix M has a true rank r factorization $M = WH$ and H is sparse, then for any other rank r factorization $M = W'H'$, H' has strictly more non-zero entries than H. As NMF has been found to recover sparse solutions [16, 17], and many algorithms explicitly enforce sparsity [17], it is reasonable to hope that the recovery procedure will be successful if the true solution is the sparsest possible solution. The main contribution of Theorem 1 is Part (b), which shows uniqueness results in the projected space. This is significantly more challenging than showing uniqueness in the original space, because in the original space we could use the fact the entries of H are drawn independently at random to simplify our arguments. However we lose access to this independence after the random projection step as the entries of the projected matrix are no longer independent. Instead, the proof in the projected space uses the expansion properties of the projection matrix P to control the dependencies, enabling us to prove uniqueness results in the projected space.

Theorem 1 only assumes that W is full column rank, and that the non-zero entries of H are independent Gaussian random variables. The Gaussian assumption is not necessary and is stated for simplicity—any continuous probability distribution for the non-zero entries of H suffices. Theorem 1 shows that if the rows of H are k-sparse, then projecting into $\Omega(p(r + k))$ dimensions preserves uniqueness, with failure probability at most $re^{-\beta k}/n$ for some constant $\beta > 0$.

Theorem 1. *Consider a rank r matrix $M \in \mathbb{R}^{m \times n}$ which has factorization $M = WH$ where $W \in \mathbb{R}^{m \times r}$ is full column rank, $H \in \mathbb{R}^{r \times n}$ and $H = B \odot Y$ where each row of B has exactly k non-zero entries chosen uniformly at random, each entry of Y is drawn independently from $N(0, 1)$ and $\odot$ denotes the element wise product. Assume $k > C$, where C is a fixed constant. Consider the projection matrix $P \in \{0, 1\}^{d \times n}$ where each column of P has exactly p non-zero entries, $p = O(\log n)$ and P is a $(\gamma n, 4/5)$ expander with $\gamma = d/(npe^5)$. Assume $d = \Omega(p(r + k))$. Let $\tilde{M} = MP^T$. Note that $\tilde{M}$ has a factorization $\tilde{M} = W\tilde{H}$ where $\tilde{H} = HP^T$. Then,*

- *(a) For any other rank r factorization $M = W'H'$, $\|H\|_0 < \|H'\|_0$ with failure probability at most $re^{-\beta k}/n$ for some $\beta > 0$.*
- *(b) For any other rank r factorization $\tilde{M} = W'H'$, $\|\tilde{H}\|_0 < \|H'\|_0$ with failure probability at most $re^{-\beta k}/n$ for some $\beta > 0$.*

Proof. We prove Part (a) of Theorem 1 in Appendix B. Here, we provide the outline of the proof for the more challenging Part (b), moving proof details to Appendix B. We will argue about the row space of the matrix $\tilde{H}$. We first claim that the row space of $\tilde{M}$ equals the row space of $\tilde{H}$. To verify, note that the row space of $\tilde{M}$ is a subspace of the row space of $\tilde{H}$. As W is full column rank, the rank of the row space of $\tilde{M}$ equals the rank of the row space of $\tilde{H}$. Therefore, the row space of $\tilde{M}$ equals

the row space of $\tilde{H}$. By the same argument, for any alternative factorization $\tilde{M} = W'H'$, the row space of H' must equal the row space of $\tilde{M}$ – which equals the row space of $\tilde{H}$. As the row space of H' equals the row space of $\tilde{H}$, therefore H' must lie in the row space of $\tilde{H}$. Therefore, our goal will be to prove that the rows of $\tilde{H}$ are the sparsest vectors in the row space of $\tilde{H}$, which implies that for any other alternative factorization $\tilde{M} = W'H'$, $\|\tilde{H}\|_0 < \|H'\|_0$.

The outline of the proof is as follows. First, we argue that the if we take any subset S of the rows of $\tilde{H}$, then the number of columns which have non-zero entries in at least one of the rows in S is large. This implies that taking a linear combination of all the S rows will result in a vector with a large number of non-zero entries—unless the non-zero entries cancel in many of the columns. But using the properties of the projection matrix P and the fact that the non-zero entries of the original matrix H are drawn from a continuous distribution, we show this happens with zero probability.

Lemma 1 shows that the number of columns which are have at least one zero entry in a subset S of the rows of H grows proportionately with the size of S. The proof proceeds by showing that choosing B such that each row has k randomly chosen non-zero entries ensures expansion for B with high probability, and the fact that P is given to be an expander.

Lemma 1. *For any subset S of the rows of $\tilde{H}$, define $N(S)$ to be the subset of the columns of $\tilde{H}$ which have a non-zero entry in at least one of the rows in S. Then for every subset S of rows of $\tilde{H}$, $|N(S)| \geq \min\{16|S|kp/25, d/200\}$ with failure probability $re^{-\beta k}/n$.*

We now prove the second part of the argument–that any linear combination of rows in S cannot have much fewer non-zero entries than $N(S)$, as the probability that many of the non-zero entries get canceled is zero. Lemma 2 is the key to showing this. Define a vector x as *fully dense* if all its entries are non-zero.

Lemma 2. *For any subset S of the rows of $\tilde{H}$, let U be the submatrix of $\tilde{H}$ corresponding to the S rows and $N(S)$ columns. Then with probability one, every subset of the columns of U of size at least $|S|p$ does not have any fully dense vector in its left null space.*

Proof. Without loss of generality, assume that S corresponds to the first $|S|$ rows of $\tilde{H}$, and $N(S)$ corresponds to the first $|N(S)|$ columns of $\tilde{H}$. We will partition the columns of U into $|S|$ groups $\{\mathcal{G}_1, \ldots, \mathcal{G}_{|S|}\}$. Each group will have size at most p. To select the first group, we choose any entry y_1 of Y which appears in the first column of U. For example, if the first row of H has a one in its first column, and $P(1,1) = 1$, then the random variable $Y_{1,1}$ appears in the first column of U. Say we choose $y_1 = Y_{1,1}$. We then choose $\mathcal{G}_1$ to be the set of all columns where y_1 appears. We then remove the set of columns $\mathcal{G}_1$ from U. To select the second group, we pick any one of the remaining columns, and choose any entry y_2 of Y which appears in that column of U. $\mathcal{G}_2$ is the set of all columns where y_2 appears. We repeat this procedure to obtain $|S|$ groups, each of which will have size at most p as every variable appears in p columns.

Let $\mathcal{N}_j$ be the left null space of the first j groups of columns. We define $\mathcal{N}_0 = \mathbb{R}^{|S|}$. We will now show that either $\text{rank}(\mathcal{N}_i) = |S| - i$ or $\mathcal{N}_i$ does not contain a fully dense vector. We prove this by induction. Consider the jth step, at which we have j groups $\{\mathcal{G}_1, \ldots, \mathcal{G}_j\}$. By the induction hypothesis, either $\mathcal{N}_j$ does not contain any fully dense vector, or $\text{rank}(\mathcal{N}_j) = |S| - j$. If $\mathcal{N}_j$ does not contain any fully dense vector, then we are done as this implies that $\mathcal{N}_{j+1}$ also does not contain any fully dense vector. Assume that $\mathcal{N}_j$ contains a fully dense vector x. Choose any column U_c which has not been already been assigned to one of the sets. By the following elementary proposition, the probability that x is orthogonal to U_c is zero. We provide a proof in the appendix.

Proposition 1. *Let $v = (v_1, \ldots, v_n) \in \mathbb{R}^n$ be a vector of n independent $N(0,1)$ random variables. For any subset $S \subseteq \{1, \ldots, n\}$, let $v(S) \in \mathbb{R}^{|S|}$ refer to the subset of v corresponding to the indices in S. Consider t such subsets $S_1, \ldots, S_t$. Let each set S_i defines some linear relation $\alpha_{S_i}^T v(S_i) = 0$, for some $\alpha_{S_i} \in \mathbb{R}^{|S_i|}$. Assume that the variable v_i appear in the set S_i. Then the probability distribution of the set of variables $\{v_{t+1}, \ldots, v_n\}$ conditioned on the linear relations defined by $S_1, \ldots, S_t$ is still continuous and its density function has full support. In particular, any linear combination of the set of variables $\{v_{t+1}, \ldots, v_n\}$ has zero probability of being zero.*

If $\mathcal{N}_j$ contains a fully dense vector, then with probability one, $\text{rank}(\mathcal{N}_{j+1}) = \text{rank}(\mathcal{N}_j) - 1 = n - j - 1$. This proves the induction argument. Therefore, with probability one, either $\text{rank}(\mathcal{N}_{|S|}) = 0$ or $\mathcal{N}_{|S|}$ does not contain a fully dense vector and Lemma 2 follows. $\qquad\square$
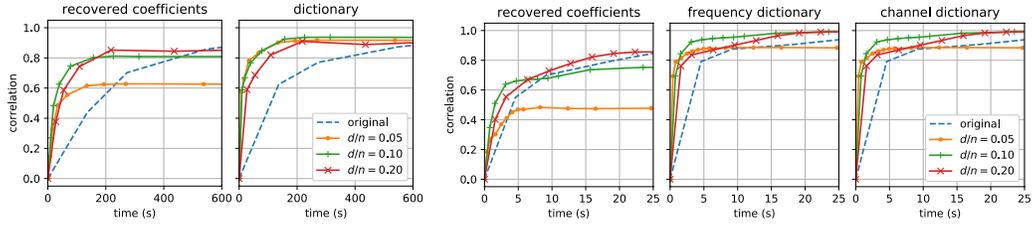
Figure 1: Pearson correlation of recovered factors for the `Enron` dataset with NMF *(left)* and EEG data with non-negative tensor factorization *(right)*. 'Recovered coefficients' here refers to the factors obtained along the larger mode of the matrix/tensor which was projected down. The 'dictionary' refers to the factors along the modes which were not projected down.

Table 2: Frobenius errors and runtimes (mins) for NMF on original and compressed data.

Dataset	$n \times d$	orig. err.	orig. time	comp. err.	comp. time
MNIST [46]	$60,000 \times 786$	1.42E3	3.1	1.78E3 (+25.3%)	1.5 (48%)
RCV1 [47]	$21,531 \times 18,758$	6.41E3	14.8	6.50E3 (+ 1.4%)	2.0 (14%)
Enron [48]	$39,861 \times 28,102$	4.51E3	18.8	4.97E3 (+10.0%)	5.3 (28%)
NYTimes [48]	$131,688 \times 102,659$	1.04E4	300.2	1.07E4 (+ 3.3%)	72.1 (24%)

Rank $r = 10$; $m = \lfloor 0.1n \rfloor$; compressed times include projection and LP recovery

We now complete the proof for the second part of Theorem 1. Note that the rows of $\tilde{H}$ have at most kp non-zero entries. Consider any set S of rows of $\tilde{H}$. Consider any linear combination $v \in \mathbb{R}^d$ of the set S rows, such that all the combination weights $x \in \mathbb{R}^{|S|}$ are non-zero. By Lemma 1, $|N(S)| \geq \min\{16|S|kp/25, d/200\}$ with failure probability $re^{-\beta k}/n$. We claim that v has more than $|N(S)| - |S|p$ non zero entries. We prove by contradiction. Assume that v has $|N(S)| - |S|$ or fewer non zero entries. Consider the submatrix U of $\tilde{H}$ corresponding to the S rows and $N(S)$ columns. If v has $|N(S)| - |S|p$ or fewer non zero entries, then there must be a subset S' of the $|N(S)|$ columns of U with $|S'| = |S|p$, such that each of the columns in S' has at least one non-zero entry, and the fully dense vector x lies in the left null space of S'. But by Lemma 2, the probability of this happening is zero. Hence v has more than $|N(S)| - |S|p$ non zero entries. Lemma 3 obtains a lower bound on $|N(S)| - |S|p$ using simple algebra.

Lemma 3. $|N(S)| - |S|p \geq 6kp/5$ *for* $|S| > 1$ *for* $d \geq 400p(r + k)$.

Hence any linear combination of more than one row of $\tilde{H}$ has at least $6kp/5$ non-zero entries with failure probability $re^{-\beta k}/n$. Hence the rows of $\tilde{H}$ are the sparsest vectors in the row space of $\tilde{H}$ with failure probability $re^{-\beta k}/n$. □

4.2 Guarantees for tensor decomposition

Showing uniqueness for tensor decomposition after random projection is not difficult, as tensor decompsition is unique for every general condition on the factor matrices [43, 44]. For example, tensors have a unique decomposition if all the factor matrices are full rank. The following elementary fact follows from this property whenever the projection matrix P is full column rank, which is true with high probability over the randomness in choosing P.

Fact 1. *Consider a n-dimensional rank r tensor $T = \sum_{i=1}^{r} w_i A_i \otimes B_i \otimes C_i$. If A, B and C are full column rank, then the tensor has a unique decomposition after projecting into any dimension $d \geq r$.*

Hence tensors have a unique decomposition on projecting to dimension at least the rank r. For tensor decomposition, we can go beyond uniqueness and prove that efficient recovery of the factors is also possible in the projected space. We show this using known guarantees for the tensor power method [45], a popular tensor decomposition algorithm. Our results for provable recovery require projection dimension $d = O(r^4 \log r)$. This is higher than the uniqueness requirement, possibly due to gaps in the analysis of the tensor power method. We state these results in Appendix C.
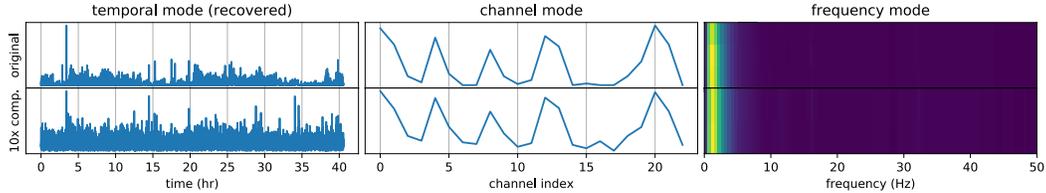
5 Experiments

Figure 2: Visualization of a sample factor from the decomposition of the EEG signal tensor. The top row is from the decomposition in the original space. The bottom row is from the decomposition of a compressed version of the tensor where the temporal dimension is reduced by 10x. Notice that we are able to recover the major features of all 3 factors, including the one compressed by 10x.

We evaluate the practical feasibility of sparse model recovery through a series of experiments on real and synthetic data. Our experiments were run on a server equipped with a Xeon E7-4870 CPU at 2.4GHz with 1TB of RAM. All timing measurements were taken under single-threaded execution.

Non-negative Matrix Factorization. We evaluate the performance of compressed NMF on a number of standard non-negative datasets. In our experiments, we used MATLAB's built-in `nnmf` with Alternating Least Squares. Unless otherwise stated, the optimization was run to convergence with a tolerance of 10^{-5}. We set the rank to 10. In Table 2 we give normalized Frobenius errors along with wall clock runtimes for NMF in the original and compressed space. We use LP-based reconstruction for the recovery step. Table 2 shows that the compressed NMF runs achieve achieve up to $5\times$ speedup at the cost of slightly increased error.

As we are usually interested in learning interesting factors via NMF rather than just minimizing squared error, the linear correlation between the recovered factors and ground-truth factors gives a finer-grained evaluation of recovery quality than Frobenius error alone. We use the
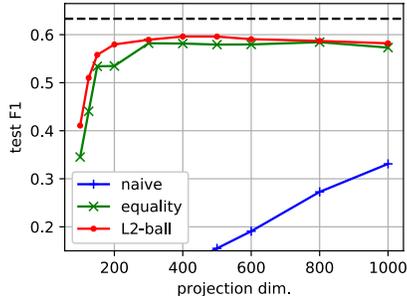

Figure 3: Sparse recovery of a linear classifier on the CENSUSINCOME dataset (original $d = 123$) with 2500 additional distractor features. The dashed line shows the F1 score of an L1-regularized logistic regressor trained in the uncompressed space. Plotted scores are medians of 5 runs. Observe that we can obtain a $6.5\times$ compression with a 6% loss in accuracy.

factors obtained in the original space at convergence as a proxy for ground-truth in the absence of cluster labels for evaluation. Since the two sets of factors may not be in the same order, we compute correlation scores via a maximum bipartite matching. In Fig. 1, we plot the correlation of the factors from the Enron word-document count matrix for various compression factors. At a fixed quality level of 0.5 (as measured by correlation), we are able to achieve $4\times$ speedups by operating in the compressed space, and are able to recover back the solution in the original space to small error using compressive sensing, even for $10\times$ compression.

Tensor Decomposition. The CHB-MIT Scalp EEG Database [49] is a collection of scalp electroencephalogram (EEG) recordings of children susceptible to epileptic seizures. Each EEG recording consists of 23 real-valued time series; each channel corresponds to a measurement of electrical activity in the different region of the brain. We preprocess one EEG trace from the dataset (`CHB01`) by computing the spectrogram of each channel. This yields a non-negative tensor of size $27804 \times 303 \times 23$ (time×frequency×channel) corresponding to approximately 40 hours of data. We expect that this data should exhibit the desired sparsity property since EEG signals typically exhibit well-localized time-frequency structure. In our experiments, we set the rank to 10. We compute a non-negative decomposition of this tensor using projected CP-ALS [50]. The tensor was projected along the first (temporal) mode using the random projection described in Section 3.1. As with NMF, we treat the factors obtained in the original space at convergence as ground truth and evaluate the mean correlation of the factors via a maximum bipartite matching.

In Fig. 1 (right), we plot the correlation of each factor over time for three compression ratios. As in our NMF experiments, we observe that for the uncompressed dictionary factors (second and third figures of the EEG data), our approach converges to a good solution with 0.5 correlation $5\times$ faster

8

than running the decomposition along the original dimension. For the compressed factors (first figure of the EEG data) we get a more modest $2\times$ speedup for $d/n \geq 0.1$. Fig 2 shows that we are able to accurately recover the original factors even with a $10\times$ compression. We provide additional details regarding this experiment and a qualitative analysis of the resulting decomposition in the Appendix.

Linear Classification. We present empirical results on recovering sparse linear classifiers trained on a compressed feature space. We use the 2-class CENSUSINCOME dataset [48] augmented with 2,500 random binary features sampled iid from $\mathrm{Bernoulli}(0.5)$. Including the distractor features, the total dimensionality is $n = 2623$. We use the standard train-test split of 32,561/16,281. For a range of projection dimensions d, we trained logistic regressors on the compressed data and recovered linear classifiers in the original space using sparse recovery. The sparsity of the random projections was fixed at $s = \lceil m/50 \rceil$. In addition to performing sparse recovery by solving the equality-constrained ℓ_1-minimization problem (Eq. 1), we also experimented with less-stringent ℓ_2-ball constraints [51]: $\|Pw - \tilde{w}^*\|_2 \leq \varepsilon$. In our experiments, we used the heuristic setting $\varepsilon = \frac{1}{2}\|P^T\tilde{w}^*\|_\infty$. We also compare the test performance of the recovered classifiers against the test performance obtained by the naive approach of projecting the test data using the random projection P and classifying directly with the weights $\tilde{w}^*$ learnt in the compressed space. Our results for this experiment are summarized in Fig. 3. The ℓ_2-ball constraint slightly outperforms the stricter equality constraint across all values of d. At $d = 400$, we achieve $\sim 6.5\times$ compression of the feature space at a loss of 3.7 F1 points versus a classifier trained in the uncompressed space (i.e. 94% of the baseline score). Both recovery techniques significantly outperform the naive approach, highlighting the importance of the denoising effect realized by sparse recovery.

6 Conclusion

We proposed a simple framework for learning sparse models by combining ideas from dimension reduction and compressive sensing, and theoretical and empirical results illustrate the viability of our approach for NMF and tensor decomposition as well as preliminary results for linear classifiers. A natural direction of future work is to extend these theoretical and experimental results to more models and applications that require learning sparse models.

References

[1] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

[2] Isabelle Guyon and André Elisseeff. An introduction to variable and feature selection. *Journal of machine learning research*, 3(Mar):1157–1182, 2003.

[3] Myunghwan Kim, Roshan Sumbaly, and Sam Shah. Root cause detection in a service-oriented architecture. In *ACM SIGMETRICS Performance Evaluation Review*, volume 41, pages 93–104. ACM, 2013.

[4] Peter Bailis, Edward Gan, Samuel Madden, Deepak Narayanan, Kexin Rong, and Sahaana Suri. MacroBase: Prioritizing attention in fast data. In *SIGMOD*. ACM, 2017.

[5] Jushan Bai and Serena Ng. Forecasting economic time series using targeted predictors. *Journal of Econometrics*, 146(2):304–317, 2008.

[6] Ling Huang, Jinzhu Jia, Bin Yu, Byung-Gon Chun, Petros Maniatis, and Mayur Naik. Predicting execution time of computer programs using sparse polynomial regression. In *Advances in neural information processing systems*, pages 883–891, 2010.

[7] Jeremy Fowers, Kalin Ovtcharov, Karin Strauss, Eric S Chung, and Greg Stitt. A high memory bandwidth FPGA accelerator for sparse matrix-vector multiplication. In *FCCM*, 2014.

[8] Song Han, Junlong Kang, Huizi Mao, Yiming Hu, Xin Li, Yubin Li, Dongliang Xie, Hong Luo, Song Yao, Yu Wang, et al. Ese: Efficient speech recognition engine with sparse lstm on fpga. In *FPGA*, 2017.

[9] Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical Learning with Sparsity: The Lasso and Generalizations*. CRC Press, 2016.

[10] Irina Rish and Genady Grabarnik. *Sparse Modeling: Theory, Algorithms, and Applications*. CRC Press, 2014.

[11] Emmanuel J Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE transactions on information theory*, 52(12):5406–5425, 2006.

[12] Radu Berinde, Anna C Gilbert, Piotr Indyk, Howard Karloff, and Martin J Strauss. Combining geometry and combinatorics: A unified approach to sparse signal recovery. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, pages 798–805. IEEE, 2008.

[13] David P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1–2):1–157, 2014. ISSN 1551-305X. doi: 10.1561/0400000060. URL http://dx.doi.org/10.1561/0400000060.

[14] Petros Drineas and Michael W. Mahoney. Randnla: Randomized numerical linear algebra. *CACM*, 59(6):80–90, May 2016. ISSN 0001-0782. doi: 10.1145/2842602. URL http://doi.acm.org/10.1145/2842602.

[15] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.

[16] Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

[17] Patrik O Hoyer. Non-negative matrix factorization with sparseness constraints. *Journal of machine learning research*, 5(Nov):1457–1469, 2004.

[18] David Donoho and Victoria Stodden. When does non-negative matrix factorization give a correct decomposition into parts? In *Advances in neural information processing systems*, pages 1141–1148, 2004.

[19] Daniel M Kane and Jelani Nelson. Sparser johnson-lindenstrauss transforms. *Journal of the ACM (JACM)*, 61(1):4, 2014.

[20] Wei Xu, Xin Liu, and Yihong Gong. Document clustering based on non-negative matrix factorization. In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, pages 267–273. ACM, 2003.

[21] Stan Z Li, Xin Wen Hou, Hong Jiang Zhang, and Qian Sheng Cheng. Learning spatially localized, parts-based representation. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 1, pages I–I. IEEE, 2001.

[22] Animashree Anandkumar, Rong Ge, Daniel J Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of Machine Learning Research*, 15(1):2773–2832, 2014.

[23] David L Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.

[24] Emmanuel J Candes. The restricted isometry property and its implications for compressed sensing. *Comptes Rendus Mathematique*, 346(9-10):589–592, 2008.

[25] Anna Gilbert and Piotr Indyk. Sparse recovery using sparse matrices. *Proceedings of the IEEE*, 98(6):937–947, 2010.

[26] Radu Berinde, Piotr Indyk, and Milan Ruzic. Practical near-optimal sparse recovery in the l1 norm. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, pages 198–205. IEEE, 2008.

[27] Piotr Indyk and Milan Ruzic. Near-optimal sparse recovery in the l1 norm. In *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*, pages 199–207. IEEE, 2008.

[28] Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro. Online learning for matrix factorization and sparse coding. *Journal of Machine Learning Research*, 11(Jan):19–60, 2010.

[29] Naiyang Guan, Dacheng Tao, Zhigang Luo, and Bo Yuan. Online nonnegative matrix factorization with robust stochastic approximation. *IEEE Transactions on Neural Networks and Learning Systems*, 23(7):1087–1099, 2012.

[30] Fei Wang, Chenhao Tan, Ping Li, and Arnd Christian König. Efficient document clustering via online nonnegative matrix factorizations. In *Proceedings of the 2011 SIAM International Conference on Data Mining*, pages 908–919. SIAM, 2011.

[31] Fei Wang and Ping Li. Efficient nonnegative matrix factorization with random projections. In *Proceedings of the 2010 SIAM International Conference on Data Mining*, pages 281–292. SIAM, 2010.

[32] Yining Wang, Hsiao-Yu Tung, Alexander J Smola, and Anima Anandkumar. Fast and guaranteed tensor decomposition via sketching. In *Advances in Neural Information Processing Systems*, pages 991–999, 2015.

[33] Zhao Song, David Woodruff, and Huan Zhang. Sublinear time orthogonal tensor decomposition. In *Advances in Neural Information Processing Systems*, pages 793–801, 2016.

[34] Casey Battaglino, Grey Ballard, and Tamara G Kolda. A practical randomized cp tensor decomposition. *arXiv preprint arXiv:1701.06600*, 2017.

[35] Evangelos Papalexakis, Christos Faloutsos, and Nicholas Sidiropoulos. Parcube: Sparse parallelizable tensor decompositions. *Machine Learning and Knowledge Discovery in Databases*, pages 521–536, 2012.

[36] Rosa I Arriaga and Santosh Vempala. An algorithmic theory of learning: Robust concepts and random projection. In *Foundations of Computer Science, 1999. 40th Annual Symposium on*, pages 616–623. IEEE, 1999.

[37] Maria-Florina Balcan, Avrim Blum, and Santosh Vempala. Kernels as features: On kernels, margins, and low-dimensional mappings. *Machine Learning*, 65(1):79–94, 2006.

[38] Qinfeng Shi, Chunhua Shen, Rhys Hill, and Anton Hengel. Is margin preserved after random projection? In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 591–598, 2012.

[39] Lijun Zhang, Mehrdad Mahdavi, Rong Jin, Tianbao Yang, and Shenghuo Zhu. Random projections for classification: A recovery approach. *IEEE Transactions on Information Theory*, 60(11):7300–7316, 2014.

[40] Scott Shaobing Chen, David L Donoho, and Michael A Saunders. Atomic decomposition by basis pursuit. *SIAM review*, 43(1):129–159, 2001.

[41] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.

[42] Daniel A Spielman, Huan Wang, and John Wright. Exact recovery of sparsely-used dictionaries. In *COLT*, pages 37–1, 2012.

[43] Joseph B Kruskal. Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear algebra and its applications*, 18(2): 95–138, 1977.

[44] Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51 (3):455–500, 2009.

[45] Anima Anandkumar, Prateek Jain, Yang Shi, and Uma Naresh Niranjan. Tensor vs. matrix methods: Robust tensor decomposition under block sparse perturbations. In *Artificial Intelligence and Statistics*, pages 268–276, 2016.

[46] Yann LeCun. The mnist database of handwritten digits. *http://yann. lecun. com/exdb/mnist/*, 1998.

[47] Massih Amini, Nicolas Usunier, and Cyril Goutte. Learning from multiple partially observed views-an application to multilingual text categorization. In *Advances in neural information processing systems*, pages 28–36, 2009.

[48] M. Lichman. UCI machine learning repository, 2013. URL http://archive.ics.uci.edu/ml.

[49] Ali H Shoeb and John V Guttag. Application of machine learning to epileptic seizure detection. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 975–982, 2010.

[50] Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51 (3):455–500, 2009.
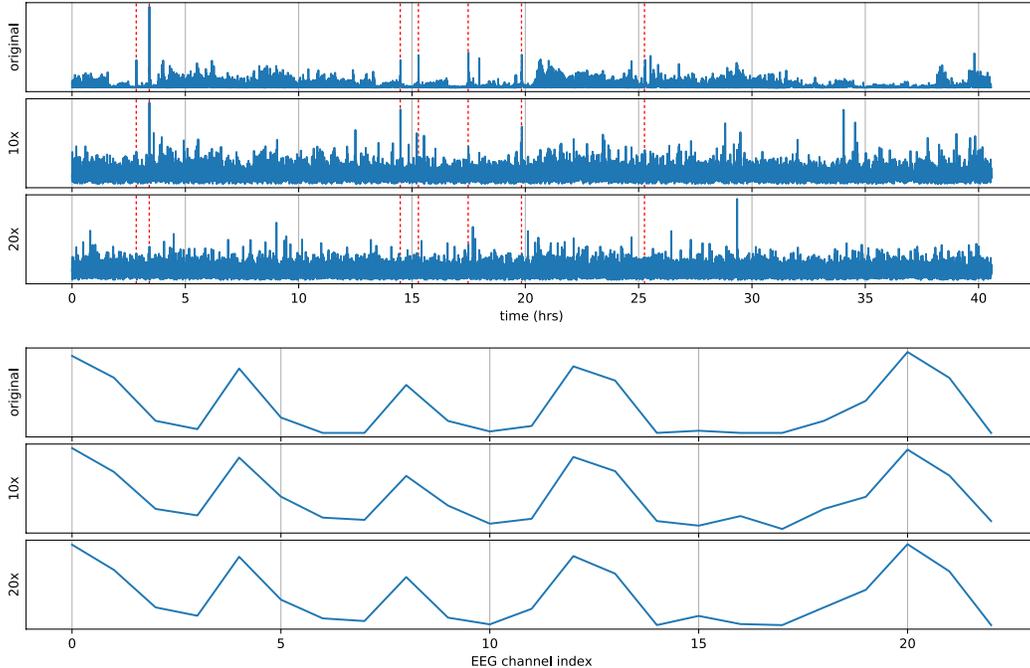
Figure 4: Visualization of a factor that correlates with the onset of seizures in the patient. The top set of plots shows the temporal mode while the bottom set shows the channel mode. $10\times$ and $20\times$ labels refer to the level of compression. Red dotted lines indicate the actual onset times of seizures. The $10\times$ compression level preserves most of the peaks, but at $20\times$ compression these peaks are lost. The channel dictionary is well-preserved across these levels of compression.

[51] Emmanuel J Candes, Justin K Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on pure and applied mathematics*, 59(8): 1207–1223, 2006.

[52] Vatsal Sharan and Gregory Valiant. Orthogonalized ALS: A theoretically principled tensor decomposition algorithm for practical use. *arXiv preprint arXiv:1703.01804*, 2017.

[53] Animashree Anandkumar, Rong Ge, and Majid Janzamin. Guaranteed non-orthogonal tensor decomposition via alternating Rank-1 updates. *arXiv preprint arXiv:1402.5180*, 2014.

A Supplementary Experimental Results

A.1 Tensor decomposition of EEG data

Preprocessing. Each channel is individually whitened with a mean and standard deviation estimated from segments of data known to not contain any periods of seizure. The spectrogram is computed with a Hann window of size 512 (corresponding to two seconds of data). The window overlap is set to 64. In order to capture characteristic sequences across time windows, we transform the spectrogram by concatenating groups of sequential windows, following Shoeb and Guttag [49]. We concatenate groups of size three.

Detection of seizures. The original goal of the work by [49] was to train patient-specific classifiers to detect the onset of epileptic seizures. We find that a tensor decomposition of the time series yields a factor that correlates with the onset of seizures, as shown in Fig. 4. This example illustrates the tradeoff between the compression factor and the fidelity of the recovered modes.

A.2 NMF experiments on synthetic data

We perform experiments on NMF using synthetic data where we can control the sparsity of the true factors to verify how our approach performs if the factors are very sparse. We take $n = m$ and $r = 20$. We vary n from 1000 to 10000. For $M = WH$, we set W to be the absolute value of a

(a) Time taken for solving NMF in original space and compressed space

(b) Quality of solution for solving NMF in original space and compressed space
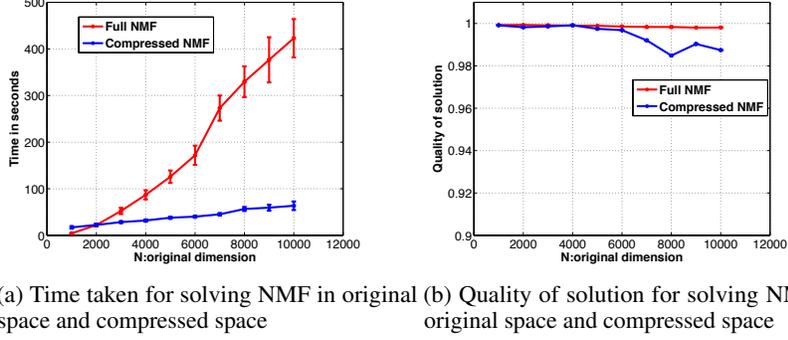
Figure 5: Synthetic experiments on NMF. The total number of non zero entries per row of H is fixed to be $k = 20$ and rank $r = 20$. We project down to $k \log n$ dimensions. We obtain significant asymptotic speedups with negligible change in accuracy in this case.

matrix $G \in \mathbb{R}^{n \times r}$ such that the entries of G are independent Gaussian random variables. For each row of H, we fix the number of non-zero entries to be $k = 20$. We project down to the theoretical limit of $d = O(k \log n)$. We can see from Fig. 5a and 5b that we can get significant speedups in this setting, at negligible degradation in accuracy. Of course, this setting is highly artificial, but serves as a sanity check that compressed NMF works.

B Additional proofs for Section 4.1: Uniqueness for NMF

B.1 Proof of Part (a) of Theorem 1

Proof. As in Part (b), we will argue about the row space of the matrix H. By the same argument as in Part (b), for any alternative factorization $M = W'H'$, the row space of H' must equal the row space of M, which equals the row space of H. As the row space of H' equals the row space of H, therefore H' must lie in the row space of H. Hence our goal we will prove that the rows of H are the sparsest vectors in the row space of H, which implies that for any other factorization $M = W'H'$, $\|H\|_0 < \|H'\|_0$. Lemma 4 shows that the number of columns which are have at least one zero entry in a subset S of the rows of H grows proportionately with the size of S.

Lemma 4. *For any subset S of the rows of H, define $N(S)$ to be the subset of the columns of H which have a non-zero entry in at least one of the rows in S. Then for every subset S of rows of H, $|N(S)| \geq \min\{4|S|k/5, n/200\}$ with high probability.*

Proof. Consider a bipartite graph G with r nodes on the left part U corresponding to the r rows of H, and n nodes on the right part V corresponding to the n columns or indices of each factor. The ith node in U has an edge to k nodes in V corresponding to the non-zero indices of the ith row of H. Note that $|N(S)|$ is the neighborhood of the set of nodes S in G. From Part 1 of Lemma 5, the graph G is a $(\gamma_1 r, 4/5)$ expander with failure probability $re^{-\beta k}/n$ for $\gamma_1 = n/(rke^5)$ and a fixed constant $\beta > 0$.

Lemma 5. *Randomly choose a bipartite graph G with n_1 vertices on the left part U and n_2 vertices on the right part V such that every vertex in U has degree D. Then,*

1. *For every $n_1, n_2, n_1 < n_2$, G is a $(\gamma n_1, 4/5)$ expander for $D \geq c$ for some fixed constant c and $\gamma n_1 = \frac{n_2}{De^5}$ except with probability $n_1 e^{-\beta D}/n_2$ for a fixed constant $\beta > 0$.*

2. *For every $n_1, n_2, n_2 < n_1$, G is a $(\gamma n_1, 4/5)$ expander for $D \geq c \log n_1$ for some fixed constant c and $\gamma n_1 = \frac{n_2}{De^5}$ except with probability $1/n_2$.*

As G is a $(\gamma_1 r, 4/5)$ expander, every set of $|S| \leq \gamma_1 r$ nodes has at least $4|S|k/5$ neighbors. A set of size $|S| > \gamma_1 r$ nodes, must include a subset of size $\gamma_1 r$ which has $4n/(5e^5) \geq n/200$ neighbours, and hence every set of size $|S| > \gamma_1 r$ has at least $n/200$ neighbors. Therefore, for every subset of S rows, $|N(S)| \geq \min\{4|S|k/5, n/200\}$ with failure probability $re^{-\beta k}/n$. □

We now show the second part of the proof—that for any subset S of the rows, every linear combination of all S rows must have $|N(S)| - |S|$ non-zero entries. The following simple Lemma is a restatement of a result from from Spielman et al. [42]. We provide a proof in the appendix for completeness. Define a vector x as *fully dense* if all its entries are non-zero.

Lemma 6. *Let $\Omega \in \{0,1\}^{n \times n}$ be any binary matrix with at least one nonzero in each column. Let $U = \Omega \odot Y$ where each entry of Y is drawn independently from $N(0,1)$. Then with probability one in the random matrix V, the left nullspace of U does not contain a fully dense vector.*

Proof. Let $\mathcal{N}_i$ be the null space of the first i columns of U, we define $\mathcal{N}_0 = \mathbb{R}^n$. We will show that with probability one, either N_i does not contain a fully dense vector, or $\text{rank}(\mathcal{N}_i) = n - i$. The proof proceeds by induction. Consider $i = 1$. Consider any full dense vector $x \in \mathbb{R}^n$. As every column of U has at least one non-zero entry, the probability that x is orthogonal to the first column U_1 is exactly zero. This is because the non-zero entries of U_1 are independent Gaussian random variables, hence the probability that any fixed linear combination is zero is exactly zero. Hence $\text{rank}(\mathcal{N}_1) = \text{rank}(\mathcal{N}_0) - 1 = n - 1$ with probability 1. Consider the jth step. By the induction hypothesis, either $\mathcal{N}_j$ does not contain any fully dense vector, or $\text{rank}(\mathcal{N}_j) = n - j$. If $\mathcal{N}_j$ does not contain any fully dense vector, then we are done as this implies that $\mathcal{N}_{j+1}$ also does not contain a fully dense vector. Let $\mathcal{N}_j$ contain a fully dense vector x. The $(j+1)$th column U_{j+1} of U has at least one non-zero entry, and the entries are independent of the previous columns. The probability that x is orthogonal to U_{j+1} is zero, by the previous argument. Hence with probability one, either $\mathcal{N}_{j+1}$ does not contain a fully dense vector, or $\text{rank}(\mathcal{N}_{j+1}) = \text{rank}(\mathcal{N}_{j+1}) - 1 = n - j - 1$. This proves the induction argument. Therefore, with probability 1, either $\text{rank}(\mathcal{N}_n) = 0$ or $\mathcal{N}_n$ does not contain a fully dense vector. Hence the left null space of U does not contain a fully dense vector. $\square$

We now complete the proof for the first part of Theorem 1. Consider any set S of the rows of H. Consider any linear combination $v \in \mathbb{R}^n$ of the set S rows, such that all the combination weights $x \in \mathbb{R}^{|S|}$ are non-zero. By Lemma 4, $|N(S)| \geq \min\{4|S|k/5, n/200\}$ with failure probability $re^{-\beta k}/n$. We claim that v has more than $|N(S)| - |S|$ non zero entries. We prove by contradiction. Assume that v has at most $|N(S)| - |S|$ non zero entries. Consider the submatrix U of H corresponding to the S rows and $N(S)$ columns. If v has $|N(S)| - |S|$ or fewer non zero entries, then there must be a subset S' of the $|N(S)|$ columns of U with $|S'| = |S|$, each of which has at least one non-zero entry, and such that the fully dense vector x lies in the left null space of S'. But by Lemma 6, the probability of this happening is zero. Hence v has more than $|N(S)| - |S|$ non zero entries. Note that $|N(S)| - |S| \geq 7k/5$. Hence any linear combination of more than one row of H has more than $7k/5$ non-zero entries with failure probability $re^{-\beta k}/n$. Hence the rows of H are the sparsest vectors in the row space of H with failure probability $re^{-\beta k}/n$. $\square$

B.2 Additional proofs for Part (b) of Theorem 1

Lemma 1. *For any subset S of the rows of $\tilde{H}$, define $N(S)$ to be the subset of the columns of $\tilde{H}$ which have a non-zero entry in at least one of the rows in S. Then for every subset S of rows of $\tilde{H}$, $|N(S)| \geq \min\{16|S|kp/25, d/200\}$ with failure probability $re^{-\beta k}/n$.*

Proof. Recall that in the proof of Lemma 4, we considered a bipartite graph corresponding to the r rows and the n indices. After the projection step, the n indices are projected to d dimensions, and the projection matrix is a $(\gamma_2 n, 4/5)$ expander with $\gamma_2 = d/(nke^5)$. We can now consider a tripartite graph, by adding a third set W with d nodes. We add an edge from a node i in V to node j in W if $P(j, i) = 1$. For any subset S of rows of $\tilde{H}$, $N(S)$ are the set of nodes in W which are reachable from the nodes S in U. By Lemma 4, the number of neighbors in V of any set S of nodes in U is of size at least $\min\{4|S|k/5, n/200\}$ with failure probability $re^{-\beta k}/n$. As the projection matrix P is a $(\gamma_2 n, 4/5)$ expander with $\gamma_2 = d/(npe^5)$, every subset of size t in V has at least $\min\{4tp/5, d/200\}$ neighbors in W. By combining this argument with Lemma 4, it follows that for every subset of S rows of H, $|N(S)| \geq \min\{16|S|kp/25, d/200\}$ with failure probability $re^{-\beta k}/n$. $\square$

Proposition 1. *Let $v = (v_1, \ldots, v_n) \in \mathbb{R}^n$ be a vector of n independent $N(0,1)$ random variables. For any subset $S \subseteq \{1, \ldots, n\}$, let $v(S) \in \mathbb{R}^{|S|}$ refer to the subset of v corresponding to the indices in S. Consider t such subsets $S_1, \ldots, S_t$. Let each set S_i defines some linear relation $\alpha_{S_i}^T v(S_i) = 0$, for*

some $\alpha_{S_i} \in \mathbb{R}^{|S_i|}$. Assume that the variable v_i appear in the set S_i. Then the probability distribution of the set of variables $\{v_{t+1}, \ldots, v_n\}$ conditioned on the linear relations defined by $S_1, \ldots, S_t$ is still continuous and its density function has full support. In particular, any linear combination of the set of variables $\{v_{t+1}, \ldots, v_n\}$ has zero probability of being zero.

Proof. We prove by induction. For the base case, note that without any linear constraints, the set of n random variables $\{v_1, \cdots, v_n\}$ is continuous and has full support as the random variables v_i are independent Gaussian. Consider the jth step, when linear constraints defined by the sets $S_1, \cdots, S_j$ have been imposed on the variables. We claim that the distribution of the set of random variables $\{v_{j+1}, \cdots, v_n\}$ is continuous and has full support after imposition of the constraints $S_1, \cdots, S_j$. By the induction hypothesis, the distribution of the set of random variables $\{v_j, \cdots, v_n\}$ is continuous and has full support after imposition of the constraints $S_1, \cdots, S_{j-1}$. Note that the linear constraint S_j can be satisfied for any assignment to the subset of variables $\{v_{j+1}, \cdots, v_n\}$ which appear in the constraint S_j, as v_j can be chosen appropriately because by the induction hypothesis it has full support conditioned on the previous constraints $S_1, \cdots, S_{j-1}$. Hence the probability distribution of the set of variables $\{v_{j+1}, \cdots, v_n\}$ is still continuous and has full support after adding the constraint S_j. $\qquad\square$

Lemma 3. $|N(S)| - |S|p \geq 6kp/5$ for $|S| > 1$ for $d \geq 400p(r+k)$.

Proof. For $2 \leq |S| \leq d/(128kp)$,

$$|N(S)| - |S|p \geq (16kp/25)|S| - p|S| = 30kp/25 + kp(16|S| - 30)/25 - p|S|$$
$$\geq 6kp/5 + p\Big(k(16|S| - 30) - |S|\Big)$$

For $|S| \geq 2$ and $k \geq 2$, $k(16|S| - 30) - |S| \geq 0$, hence $|N(S)| - |S|p \geq 6kp/5$ for $2 \leq |S| \leq d/(128kp)$. For $|S| > d/(128kp)$, $|N(S)| \geq d/200$. Therefore, $|N(S)| - |S|p \geq d/200 - rp \geq 2kp$ for $d \geq 400p(r+k)$. $\qquad\square$

C   Guaranteed recovery for tensors in the projected space

We can prove a stronger result for symmetric, incoherent tensors and guarantee accurate recovery in the compressed space using the tensor power method. The tensor power method is the tensor analog of the matrix power method for finding eigenvectors. It is equivalent to finding a rank 1 factorization using the Alternating Least Squares (ALS) algorithm. Incoherent tensors are tensors for which the factors have small inner products with other. We define the incoherence $\mu = \max_{i \neq j}\{|A_i^T A_j|\}$. Our guarantees for tensor decomposition follow from the analysis of the tensor power method by Sharan and Valiant [52]. Proposition 2 shows guarantees for recovering one of the true factors, multiple random initializations can then be used for the tensor power method to recover back all the factors (see Anandkumar et al. [53]).

Proposition 2. *Consider a n-dimensional rank r tensor $T = \sum_{i=1}^{r} w_i A_i \otimes A_i \otimes A_i$. Let $c_{\max} = \max_{i \neq j} |A_i^T A_j|$ be the incoherence between the true factors and $\gamma = \frac{w_{\max}}{w_{\min}}$ be the ratio of the largest and smallest weight. Assume γ is a constant and $\mu \leq o(r^{-2})$. Consider a projection matrix $P \in \{0, \pm 1\}^{n \times d}$ where every row has exactly p non-zero entries, chosen uniformly and independently at random and the non-zero entries have uniformly and independently distributed signs. We take $d = O(r^4 \log r)$ and $p = O(r^2 \log r)$. Let $\tilde{A} = AP$ and $\tilde{T}$ be the d dimensional projection of T, hence $\tilde{T} = \sum_{i=1}^{k} w_i \tilde{A}_i \otimes \tilde{A}_i \otimes \tilde{A}_i$. Then,*

1. *For the original tensor decomposition problem, if the initialization $x_0 \in \mathbb{R}^n$ is chosen uniformly at random from the unit sphere, then with high probability the tensor power method converges to one of the true factors of A (say the first factor A_1) in $O(r(\log r + \log \log n))$ steps, and the estimate A_1' satisfies $\| A_1 - A_1' \|_2^2 \leq O(r \max\{\mu^2, 1/n^2\})$.*

2. *For the projected tensor decomposition problem, if the initialization $x_0 \in \mathbb{R}^d$ is chosen uniformly at random from the unit sphere, then with high probability the tensor power method converges to one of the true factors of $\tilde{T}$ (say the first factor $\tilde{A}_1$) in $O(r(\log r + \log \log d))$ steps, and the estimate $\tilde{A}'$ satisfies $\| \tilde{A}_1 - \tilde{A}_1' \|_2^2 \leq O(r \max\{\mu^2, 1/d^2\})$.*

Proof. Our proof relies on Theorem 3 of Sharan and Valiant [52] and sparse Johnson Lindenstrauss transforms due to Kane and Nelson [19]. Claim 1 of Proposition 2 is Theorem 3 of Sharan and Valiant [52]. To show Claim 2 we need to ensure that the incoherence parameter in the projected space is small. We use the Johnson Lindenstrauss property of our projection matrix to ensure this. A matrix M is regarded as a Johnson Lindenstrauss matrix if it preserves the norm of a randomly chosen unit vector x up to a factor of $(1 \pm \epsilon)$, with failure probabilty δ.

$$\mathbb{P}_x[(1 - \epsilon) < \parallel Mx \parallel_2^2 < (1 + \epsilon)] > 1 - \delta$$

We use the results of Kane and Nelson [19] who show that with high probability a matrix $P \in \{0, \pm 1\}^{n \times d}$ where every row has p non-zero entries, chosen uniformly and independently at random and the non-zero entries have uniformly and independently distributed signs, preserves pairwise distances to within a factor ϵ for $d = O(\epsilon^{-2} \log(1/\delta))$ and $p = \Theta(\epsilon^{-1} \log(1/\delta))$.

It is easy to verify that inner-products are preserved to within an additive error ϵ if the pairwise distances are preserved to within a factors of $(1 \pm \epsilon)$. By choosing $\delta = 1/r^3$ and doing a union bound over all the r^2 pairs of factors, the factors are $(\mu \pm \epsilon)$ incoherent in the projected space with high probability if they were μ incoherent in the original space. Setting $\epsilon = r^{-2} \log^{-1} r$ ensures that $\mu + \epsilon = o(r^{-2})$. Claim 2 now again follows from Theorem 3 of Sharan and Valiant [52]. $\qquad \square$

D Proof of expander property: Lemma 5

In this section, we provide a proof of our claim that randomly chosen projection matrix is an expander with the desired expansion. This is Part 2 of Lemma 5. Part 1 is used in the proof of uniqueness for NMF (Lemma 4).

Lemma 5. *Randomly choose a bipartite graph G with n_1 vertices on the left part U and n_2 vertices on the right part V such that every vertex in U has degree D. Then,*

1. *For every $n_1, n_2, n_1 < n_2$, G is a $(\gamma n_1, 4/5)$ expander for $D \geq c$ for some fixed constant c and $\gamma n_1 = \frac{n_2}{De^5}$ except with probability $n_1 e^{-\beta D}/n_2$ for a fixed constant $\beta > 0$.*

2. *For every $n_1, n_2, n_2 < n_1$, G is a $(\gamma n_1, 4/5)$ expander for $D \geq c \log n_1$ for some fixed constant c and $\gamma n_1 = \frac{n_2}{De^5}$ except with probability $1/n_2$.*

Proof. Consider any subset $S \subset U$ with $|S| \leq \gamma n_1$. Let $\mathbb{P}(N(S) \subseteq M)$ denote the probability of the event that the neighborhood of S is entirely contained in $M \subset V$. $\mathbb{P}(N(S) \subseteq M) \leq \left(\frac{|M|}{n_2}\right)^{D|S|}$. We will upper bound the probability of G not being an expander by upper-bounding the probability of each subset $S \subset U$ with $|S| \leq \gamma n_1$ not expanding. Let $\mathbb{P}(\bar{S})$ denote the probability of the neighborhood of S being entirely contained in a subset $M \subset V$ with $M < \alpha|S|D$. By a union bound,

$$\mathbb{P}(G \text{ is not a } (\gamma n_1, \alpha) \text{ expander}) \leq \sum_{\substack{S \subset U \\ |S| \leq \gamma n_1}} \mathbb{P}(\bar{S})$$

$$\leq \sum_{\substack{S \subset U \\ |S| \leq \gamma n_1}} \sum_{\substack{M \subset V \\ M = \alpha|S|D}} \mathbb{P}(N(S) \subseteq M)$$

$$= \sum_{s=1}^{\gamma n_1} \sum_{\substack{S \subset U \\ |S| = s}} \sum_{\substack{M \subset V \\ M = \alpha|S|D}} \left(\frac{\alpha|S|D}{n_2}\right)^{D|S|}$$

$$\leq \sum_{s=1}^{\gamma n_1} \binom{n_1}{s} \binom{n_2}{\alpha Ds} \left(\frac{\alpha Ds}{n_2}\right)^{Ds}$$

16

By using the relation $\binom{n}{k} \leq \left(\frac{ne}{k}\right)^k$, we get:

$$\mathbb{P}(G \text{ is not a } (\gamma n_1, \alpha) \text{ expander}) \leq \sum_{s=1}^{\gamma n_1} \left(\frac{n_1 e}{s}\right)^s \left(\frac{n_2 e}{\alpha D s}\right)^{\alpha D s} \left(\frac{\alpha D s}{n_2}\right)^{D s}$$

$$= \sum_{s=1}^{\gamma n_1} \left[\left(\frac{n_1 e}{s}\right)\left(\frac{n_2 e}{\alpha D s}\right)^{\alpha D}\left(\frac{\alpha D s}{n_2}\right)^{D}\right]^s$$

$$\leq \sum_{s=1}^{\gamma n_1} x_s^s$$

where $x_s = \left(\frac{n_1 e}{s}\right)\left(\frac{n_2 e}{\alpha D s}\right)^{\alpha D}\left(\frac{\alpha D s}{n_2}\right)^{D}$. x_s can be bounded as follows-

$$x_s = \left(\frac{n_1 e}{s}\right)\left(\frac{\alpha D s e^{1/(1-\alpha)}}{n_2}\right)^{(1-\alpha)D}$$

$$\leq \left(\frac{e}{\gamma}\right)\left(\frac{\alpha D \gamma n_1 e^{1/(1-\alpha)}}{n_2}\right)^{(1-\alpha)D}$$

$$\leq \left(\frac{n_1 e^{1+1/(1-\alpha)}}{n_2}\right) D \alpha^{(1-\alpha)D}$$

$$\leq \left(\frac{n_1 e^6}{n_2}\right) D e^{-D/25} = x$$

where in the last step we set $\alpha = 4/5$. Hence we can upper bound the probability of G not being an expander as follows—

$$\mathbb{P}(G \text{ is not a } (\gamma n_1, \alpha) \text{ expander}) \leq \sum_{s=1}^{\infty} x^s \leq \frac{x}{1-x}$$

The two parts of Lemma 5 follow by plugging in the respective values for n_1, n_2 and D. $\qquad \square$