

# Approximate Guarantees for Dictionary Learning

Aditya Bhaskara

Wai Ming Tai\*

University of Utah, Salt Lake City, Utah, USA

BHASKARA@CS.UTAH.EDU

WMTAI@CS.UTAH.EDU

**Editors:** Alina Beygelzimer and Daniel Hsu

## Abstract

In the dictionary learning (or sparse coding) problem, we are given a collection of signals (vectors in  $\mathbb{R}^d$ ), and the goal is to find a “basis” in which the signals have a sparse (approximate) representation. The problem has received a lot of attention in signal processing, learning, and theoretical computer science. The problem is formalized as factorizing a matrix  $X(d \times n)$  (whose columns are the signals) as  $X = AY$ , where  $A$  has a prescribed number  $m$  of columns (typically  $m \ll n$ ), and  $Y$  has columns that are  $k$ -sparse (typically  $k \ll d$ ). Most of the known theoretical results involve assuming that the columns of the unknown  $A$  have certain *incoherence* properties, and that the coefficient matrix  $Y$  has random (or partly random) structure.

The goal of our work is to understand what can be said in the absence of such assumptions. Can we still find  $A$  and  $Y$  such that  $X \approx AY$ ? We show that this is possible, if we allow violating the bounds on  $m$  and  $k$  by appropriate factors that depend on  $k$  and the desired approximation. Our results rely on an algorithm for what we call the *threshold correlation* problem, which turns out to be related to hypercontractive norms of matrices. We also show that our algorithmic ideas apply to a setting in which some of the columns of  $X$  are outliers, thus giving similar guarantees even in this challenging setting.

**Keywords:** Dictionary learning, sparse coding, approximation algorithms, pursuit algorithms

## 1. Introduction

The problem of dictionary learning or *sparse coding* was introduced in the neuroscience literature by (Olshausen and Field, 1997) as a way of understanding how the visual cortex perceives images. Their hypothesis was that it learns to recognize “basic patterns” that allow for a sparse representation of all natural images. This idea was subsequently used extensively in the machine learning and signal processing literature, with applications including compression, feature extraction, denoising and more. Mathematically, the sparse coding problem is the following: given a collection of vectors  $x_1, x_2, \dots, x_n \in \mathbb{R}^d$ , does there exist a small ‘dictionary’ (small set of vectors in  $\mathbb{R}^d$ ) such that every  $x_i$  can be written as a *sparse combination* of the elements of the dictionary? Formally, if  $X \in \mathbb{R}^{d \times n}$  is the matrix whose columns are the vectors  $x_i$  and  $m, k$  are parameters, then the goal is to find  $A \in \mathbb{R}^{d \times m}$  such that  $X \approx AY$ , for some matrix  $Y$ , each of whose columns is  $k$ -sparse.

The early algorithms for the problem, including ones by (Olshausen and Field, 1997) and (Aharon et al., 2006) are iterative in nature, and find a decomposition that minimizes  $\|X - AY\|_F$ , subject to sparsity constraints. These algorithms lacked theoretical convergence guarantees. More recently, there has been significant progress in terms of designing algorithms that are efficient (polynomial time and sample complexity), and are guaranteed to recover an unknown dictionary  $A$ . The work of

---

\* Thanks to support by NSF 1514520

Spielman (Spielman et al., 2013) studied the setting in which  $A$  is full column rank (which implies  $m \leq d$ ), and  $Y$  has entries that are independently drawn from a sparse random distribution, and  $k < \sqrt{d}$ . Many later works including (Arora et al., 2014; Agarwal et al., 2014; Barak et al., 2015; Gribonval et al., 2015), have developed powerful algorithmic tools that handle the case  $d > m$ , as well as more general distributions of  $Y$ . Recently, (Awasthi and Vijayaraghavan, 2018) considered semi-random models in which a small portion of  $Y$  is random and the rest can be adversarial (in support). All of these results have the following high level structure: assuming that  $A$  satisfies certain properties (full-rank, incoherent columns, etc.) and that  $Y$  has certain random (or partly random) structure, one can efficiently recover  $A$  up to a desired accuracy, in polynomial time. The recoverability of the dictionary  $A$  is known as *identifiability*, and is an important requirement in all the algorithms discussed above.

The question we consider is the following: suppose we only wish to find  $A$  and  $Y$  such that  $\|X - AY\|_F$  is small. Can we solve the problem even in the cases where identifiability fails to hold? Note that while identifiability is important for some applications where the dictionary elements are believed to have semantic information, it is not crucial for applications such as compression, denoising, etc. Here, the recoverability of  $X$  up to a small error is more important. Our main result is showing that such a decomposition can be efficiently obtained, with small losses in the parameters  $m$  and  $k$  that depend on the quality of the approximation desired. We show that such a dependence is necessary in order to avoid known intractability results (Tillmann, 2015). Our approximation guarantees are similar in spirit to bi-criteria approximations known for problems such as  $k$ -means clustering (see (Ostrovsky et al., 2006; Jaiswal et al., 2012)). We will now describe in detail the problem settings we consider along with more motivation.

### 1.1. Problem formulations and motivation

**Approximate dictionary learning** Suppose we are given a collection of signals  $x_1, x_2, \dots, x_n \in \mathbb{R}^d$ , with the promise that there exists a dictionary  $A^*$  of  $m$  vectors, such that each of the  $x_i$  can be approximated by a  $k$ -sparse linear combination of the dictionary elements, so that the total error in the approximation is at most  $\gamma^* \|X\|_F^2$ . (Here  $\gamma^*$  can be sub-constant.) The goal is to find a dictionary  $A'$  with  $m'$  vectors, such that each  $x_i$  can be approximated by a  $k'$ -sparse linear combination of the dictionary elements, and the total error in the approximation is  $\leq (\gamma^* + \varepsilon) \|X\|_F^2$ , for some prescribed  $\varepsilon > 0$ . We highlight that there are *no additional assumptions* on the vectors or the combination (except a norm bound  $\Lambda$ , as we will see).

This way of formalizing the question is similar to the one done in the work of (Olshausen and Field, 1997) which introduced sparse coding. As long as  $m'$  and  $k'$  are close enough to  $m, k$ , the obtained representation  $A'Y'$  would still achieve a non-trivial “compression” of the original matrix. Intuitively, we are promised that  $X$  has a “representation complexity” (measured by the number of parameters used to approximate it well) of  $dm + kn$  (which is  $\ll nd$ , the naïve representation), and we are aiming to construct a good approximation with roughly  $dm' + k'n$  parameters.

Our results (Theorem 4) show how to obtain  $m'$  and  $k'$  that are small, when  $k \ll d$ . One way of looking at our result is that if the number of observed vectors ( $n$ ) is large, then assuming there exists a basis in which each vector can be “compressed” to  $k$  real numbers, we can efficiently find a basis in which each vector can be expressed with  $k'$  ( $\approx k^2/\text{poly}(\varepsilon)$ ) (Corollary 5) real numbers, up to an  $\varepsilon$  error. In the setting where  $k$  is small ( $\ll \sqrt{d}$ ), this can be a significantly smaller representation. One key thing to consider is the dependence on  $\varepsilon$ . The bounds we give on  $m', k'$  have  $\text{poly}(1/\varepsilon)$  factors

in them. It is natural to ask if a sub-polynomial, say  $\log(1/\varepsilon)$  dependence is possible. However, we note that existing hardness results (see (Tillmann, 2015)) rule out this possibility. Indeed, even for the “easier” problem of sparse recovery, (Foster et al., 2015) showed that such a guarantee is impossible.

Our algorithm proceeds by finding  $A'$  one column at a time. To this end, the following problem turns out to be of importance. We note that such a question implicitly arises in recent works such as (Awasthi and Vijayaraghavan, 2018).

**The threshold correlation problem** We now define the problem that plays a central role in our algorithm for dictionary learning. In what follows,  $\mathcal{B}_2^d$  denotes the unit ball in  $\mathbb{R}^d$ , i.e.,  $\mathcal{B}_2^d = \{x \in \mathbb{R}^d : \|x\| \leq 1\}$ .

Suppose we are given  $n$  unit vectors  $v_1, v_2, \dots, v_n \in \mathcal{B}_2^d$ . We can define the vector that is *most correlated* with these vectors as the  $x \in \mathbb{R}^d$  with  $\|x\| = 1$  that maximizes  $\sum_i \langle x, v_i \rangle^2$ . Such a vector can clearly be found as the top left singular vector of the matrix whose columns are  $v_i$ . Now, consider a variant in which we are also given a threshold  $\tau$ , and we wish to maximize  $\sum_i \langle x, v_i \rangle^2$ , but *only over indices*  $i$  that satisfy  $\langle x, v_i \rangle^2 \geq \tau$ . Intuitively, we only “get credit” if the squared inner product with  $v_i$  is above a threshold. We call this the threshold correlation ( $\tau$ -TC) problem, and formally define it below (with weights).

**Definition 1 ( $\tau$ -TC problem)** Let  $v_1, v_2, \dots, v_n \in \mathcal{B}_2^d$ , and let  $w_i \in \mathbb{R}_{\geq 0}$  be non-negative weights. Let  $\tau \in [0, 1]$  be given. The goal in the threshold correlation ( $\tau$ -TC) problem is to find a unit vector  $x \in \mathbb{R}^d$  that maximizes

$$\sum_{i \in [n]} [\mathbf{1}_{\langle x, v_i \rangle^2 \geq \tau}] w_i \langle x, v_i \rangle^2,$$

where  $[\mathbf{1}_P]$  for a predicate  $P$  is the indicator that is 1 if  $P$  holds and 0 otherwise.

The two extremes  $\tau = 0$  and  $\tau = 1$  can be solved easily. The former is simply the problem of computing the largest singular vector. The latter is solved by returning the most frequent vector (taking into account the weights  $w_i$ ) of length exactly 1 among the  $v_i$ . Developing algorithms for arbitrary  $\tau \in [0, 1]$  is a natural question. In our setting, it turns out that bi-criteria approximations for the  $\tau$ -TC problem give interesting guarantees for dictionary learning.

**Definition 2 ( $(\alpha, \beta)$ -approximation)** Let  $v_1, v_2, \dots, v_n, w_1, \dots, w_n$  and  $\tau$  be the inputs for an instance of the  $\tau$ -TC problem. Suppose the optimum objective value is  $\text{OPT}$ . For  $0 \leq \alpha, \beta \leq 1$ , a  $(\alpha, \beta)$ -approximation algorithm is one that returns a unit vector  $x \in \mathbb{R}^d$  such that

$$\sum_{i \in [n]} [\mathbf{1}_{\langle x, v_i \rangle^2 \geq \alpha \cdot \tau}] w_i \langle x, v_i \rangle^2 \geq \beta \cdot \text{OPT}.$$

In other words, the threshold is reduced by a factor  $\alpha$ , and the objective value by a factor  $\beta$ .

While natural by itself, the  $\tau$ -TC problem turns out to be related to the well-studied question of approximating hypercontractive norms of matrices (see e.g., (Barak et al., 2012; Bhaskara and Vijayaraghavan, 2011; Bhattiprolu et al., 2018)), as we will see.

**Approximate dictionary learning with outliers** Finally, we consider the dictionary learning problem where some of the columns (upper bounded in number) are allowed to be *outliers*, i.e., they do not need to have a sparse approximation in terms of the dictionary. This is quite common to assume in practice.

**The objective.** The first question is to define the right objective for capturing this version of the problem. The natural choice is to use the same objective as before, i.e., find  $A$  and  $Y$  so as to minimize  $\|X - AY\|_F^2$  (subject to the constraints as before). However, consider a scenario in which all the columns of  $X$  are unit vectors, and say 10% of the columns are outliers. Suppose the optimum objective value is  $0.05 \|X\|_F^2$ , with *all the error* coming from the outliers. (I.e., the inliers have a perfect representation in terms of a sparse dictionary.) Now, if an algorithm achieves the optimum objective value of  $0.05 \|X\|_F^2$ , it could potentially be because of *every column* having an error of 0.05 in the obtained representation. This is not desirable, as the dictionary is not doing well enough on the inliers.

To remedy this, we propose the following objective, which is closely related to the objectives studied in various robust estimation problems.

**Definition 3 (Dictionary learning with outliers)** *Given parameters  $m, k$ , and a collection of observations  $x_1, x_2, \dots, x_n \in \mathbb{R}^d$  that form the columns of the matrix  $X$ , the goal is to decompose  $X$  as  $AY + N$ , where  $A \in \mathbb{R}^{d \times m}$ , every column of  $Y$  has at most  $k$  non-zero entries, and  $N$  is a matrix with at most  $\rho n$  nonzero columns. The objective is to minimize the error in the decomposition, i.e., minimize  $\|X - AY - N\|_F^2$ .*

While we technically think of the fraction  $\rho$  as small, our algorithm makes no such assumption. This is not too surprising, given that the guarantee we aim for is additive in the norm of the entire matrix. The problem of allowing outliers in dictionary learning was studied in (Gribonval et al., 2015), but the focus there is quite different from ours.

## 1.2. Our results

Our main result gives a connection between an approximation for the  $\tau$ -TC problem and dictionary learning.

**Theorem 4** *Let  $X \in \mathbb{R}^{d \times n}$  be an instance of dictionary learning for which we know that there exist matrices  $A^* \in \mathbb{R}^{d \times m}$ ,  $Y^* \in \mathbb{R}^{m \times n}$ , where*

- (a) *the columns of  $A^*$  are all unit vectors,*
- (b) *each column of  $Y^*$  is  $k$ -sparse, and satisfies the norm bound  $\|Y_i^*\|^2 \leq \Lambda \|x_i\|^2$ , and*
- (c) *we have  $\|X - A^*Y^*\|_F^2 \leq \gamma^* \|X\|_F^2$ .*

*Let  $\varepsilon > 0$  be an accuracy parameter, and suppose that there exists an efficient  $(\alpha, \beta)$ -approximation algorithm for the  $\tau$ -TC problem, when  $\tau = \varepsilon^2/k\Lambda$  (for parameters  $\alpha, \beta$  depending on  $\tau$ ). Then there is an efficient algorithm that outputs matrices  $A', Y'$  such that (a)  $A'$  has at most  $O(m\Lambda/\beta\varepsilon)$  columns, (b) every column of  $Y'$  has at most  $O(k\Lambda/\alpha\varepsilon^2)$  non-zero entries, and (c)  $\|X - A'Y'\|_F^2 \leq (\gamma^* + \varepsilon) \|X\|_F^2$ .*

Using our approximation algorithm for the  $\tau$ -TC problem, we obtain the following corollary.

**Corollary 5** *Under the assumptions of Theorem 4, the number of columns of the output  $A'$  is at most  $O\left(\frac{mk^2\Lambda^3}{\varepsilon^5}\right)$ . Also, every column of the output  $Y'$  has at most  $O\left(\frac{k^2\Lambda^2}{\varepsilon^4}\right)$ .*

Note that these bounds are off from  $m, k$  by factors that only depend on  $k, \varepsilon$  and  $\Lambda$  (and not  $d, m$ ). For instance, if  $k$  is small and  $\varepsilon, \Lambda$  are constants, we obtain a non-trivial *compression* of the input signals. As noted before, a  $\text{poly}(1/\varepsilon)$  dependence on the parameter  $\varepsilon$  is essential due to hardness results.

**Aside: the parameter  $\Lambda$ .** Our bounds also depend on the parameter  $\Lambda$ , which is the norm of the coefficient vector used to represent  $x_i / \|x_i\|$  in the promised solution. For intuition on  $\Lambda$ , suppose we have a unit vector  $v$  (in our case, an  $x_i$ ) that we express as a linear combination of unit vectors  $u_1, u_2, \dots, u_k$  (in our case a subset of the columns of  $A^*$ ),  $v = \xi_1 u_1 + \dots + \xi_k u_k$ . If  $u_i$  are (near) orthogonal, we can *guarantee* that  $\|\xi\| = O(1)$ . However, near orthogonality (or large least-singular-value) is not a necessity, depending on  $v$ . The factor  $\Lambda$  arises in our analysis as it is related to how well  $v$  is correlated with the  $u_i$ . As an extreme example, consider  $d = 2$ , and  $v = (0, 1)$ ,  $u_1 = (1, 0)$  and  $u_2 = (\sqrt{1 - \gamma^2}, \gamma)$ , where  $\gamma$  is tiny. In this case,  $v$  can be represented as a linear combination of  $u_1$  and  $u_2$ , but neither of the vectors is “well-correlated” with  $v$ . This is, in essence, because the coefficients used in the linear combination are large ( $\Lambda = \Theta(1/\gamma^2)$ ). Having small  $\Lambda$  allows us to argue that one of the  $u_i$  has high correlation with  $v$ . (See Lemma 10.) Thus a dependence on  $\Lambda$  is essential for any iterative algorithm of the type we consider.

The next result we show is a bi-criteria approximation for the  $\tau$ -TC problem, in the sense of Definition 2.

**Theorem 6** *For every  $\tau \in (0, 1)$ , there is a polynomial time  $\left(\frac{\tau}{4}, \frac{\tau^2}{32}\right)$ -approximation algorithm for the  $\tau$ -TC problem.*

It is a very interesting open problem to improve at least one of the parameters in the bi-criteria approximation. We can show that we cannot expect constant factor approximations. To this end, we prove the following result, connecting the complexity of  $\tau$ -TC to the question of approximating hypercontractive norms of matrices. Informally, we show that constant factor bi-criteria approximations to the  $\tau$ -TC problem imply a logarithmic approximation to the so-called  $2 \mapsto p$  norms (which is unlikely, due to the work of Barak et al. (2012)). We refer to Section 4.2 for the details.

Finally, for the problem of dictionary learning with outliers, we show that an algorithm almost identical to the one in Theorem 4 (with a slightly modified analysis) yields the following.

**Theorem 7** *Let  $X \in \mathbb{R}^{d \times n}$  be a matrix of observations for which we know that there exist matrices  $A^* \in \mathbb{R}^{d \times m}$ ,  $Y^* \in \mathbb{R}^{m \times n}$ ,  $N^* \in \mathbb{R}^{d \times n}$ , where*

- (a) *the columns of  $A^*$  are all unit vectors,*
- (b) *each column of  $Y$  is  $k$ -sparse, and satisfies the norm bound  $\|Y_i\|^2 \leq \Lambda \|x_i\|^2$ ,*
- (c) *the matrix  $N^*$  has at most  $pn$  non-zero columns, and*
- (d) *we have  $\|X - A^*Y^*\|_F^2 \leq \gamma^* \|X\|_F^2$ .*

*Let  $\varepsilon > 0$  be an accuracy parameter, and suppose that there exists an efficient  $(\alpha, \beta)$ -approximation algorithm for the  $\tau$ -TC problem, when  $\tau = \varepsilon^2/k\Lambda$  (for parameters  $\alpha, \beta$  depending on  $\tau$ ). Then there is an efficient algorithm that outputs matrices  $A', Y', N'$  such that (a)  $A'$  has at most  $O(m\Lambda/\beta\varepsilon^3)$  columns, (b) every column of  $Y'$  has at most  $O(k\Lambda/\alpha\varepsilon^2)$  non-zero entries, (c)  $N'$  has at most  $pn$  non-zero columns, and (d)  $\|X - A'Y'\|_F^2 \leq (\gamma^* + \varepsilon) \|X\|_F^2$ .*

Note that the only difference in the parameters between Theorems 4 and 7 is a mildly worse dependence on  $\varepsilon$  in the bound on the number of columns.

### 1.3. Related work

There has been a large body of work in the signal processing literature on the problem of dictionary learning and also the related question of sparse recovery. Sparse recovery is the problem of reconstructing a sparse vector  $x$  given  $Ax$ , where  $A$  is now a *known* dictionary. Iterative *pursuit* algorithms of the kind we consider have been extensively studied in this context (see, e.g., (Davis et al., 1997) and (Tropp, 2004)). Our iterative algorithms are also reminiscent of Frank-Wolfe iteration, which is a powerful technique for sparse approximation. We refer to (Clarkson, 2010; Shalev-Shwartz et al., 2010) and references therein. In this context, our results may be viewed as showing that greedy pursuit with an appropriate subroutine—an algorithm for the  $\tau$ -TC problem—can give approximate guarantees for dictionary learning.

Another related work is that of (Blum et al., 2016). They consider a problem in which we are given  $x_1, x_2, \dots, x_n \in \mathbb{R}^d$ , and the goal is to come up with a small set of points  $\mathcal{P}$  such that all the  $x_i$  lie in  $\text{conv}(\mathcal{P})$ , the convex hull of  $\mathcal{P}$ . This is equivalent to requiring a decomposition in which the columns of the  $Y$  matrix are convex combinations. However, it turns out that the results do not imply any bound for dictionary learning (in the traditional setting that we study).

Our work shows an interesting connection between the  $\tau$ -TC problem and hypercontractive norms. Indeed, our algorithm implies an approximation algorithm for the  $2 \mapsto p$  norm problem, in the case when the norm is *large enough*. Now, there are other more sophisticated sub-exponential time algorithms for hypercontractive norms Barak et al. (2012). It is an interesting question to investigate if these techniques imply better approximations for the  $\tau$ -TC problem. From our results, this would translate to better guarantees for approximate dictionary learning.

## 2. Preliminaries and overview

### 2.1. Notation

We will refer to the input vectors  $x_i$  to the dictionary learning problem as observations, signals, or simply as input vectors. We will also use standard notation for norms of vectors and matrices.  $\|v\|_p$  refers to the  $\ell_p$  norm. We will drop the subscript when we are referring to the  $\ell_2$  norm. For a matrix  $A$ , we use  $A_i$  to refer to the  $i$ 'th column vector of  $A$ .  $\|A\|_F$  is the Frobenius norm of the matrix, defined as  $\sqrt{\sum_{i,j} A_{ij}^2}$ .

### 2.2. Paper outline

The rest of the paper is organized as follows. In Section 3, we give our main algorithm for approximate dictionary learning, assuming an approximation algorithm for the  $\tau$ -TC problem. This will establish Theorem 4. The outline of the argument is presented at the opening of the section.

Next in Section 4, we give a bi-criteria approximation for the threshold correlation problem, establishing 6. The argument involves a simple clustering procedure, and shows that one of the given vectors is itself a good enough solution for the  $\tau$ -TC problem. We then establish a connection between the  $\tau$ -TC problem and the question of approximating  $2 \mapsto p$  norms of matrices, thereby showing a hardness result for  $\tau$ -TC.

Finally in Section D (whose brief outline is in Section 5), we consider the problem of dictionary learning with outliers defined above, and show that the iterative framework of Section 3 extends to a

setting in which some of the columns have no sparse approximation using the chosen dictionary. The argument relies heavily on the methods in Section 3, even though the structure is slightly different.

### 3. Algorithm and analysis

In this section, we will describe the algorithm in detail, along with its analysis. The goal will be to establish Theorem 4. Let us start with a rough outline of the algorithm, and introduce some notation. Recall that  $X \in \mathbb{R}^{d \times n}$  is the matrix whose columns are the observations  $x_i \in \mathbb{R}^d$ . The goal is to return an approximation decomposition  $A'Y'$ , as described in Theorem 4.

**Algorithm outline.** The algorithm proceeds by building the matrix  $A'$  iteratively, one column at a time, and simultaneously the matrix  $Y'$  one row at a time. In each iteration, we maintain a “current approximation” to each of the  $x_i$  (which will be the  $i$ th column of the current  $A'Y'$ ). This will be the projection of  $x_i$  onto a *subset* of the columns of  $A'$ . We denote the error in the approximation in iteration  $t$  by  $z_i^{(t)}$ . I.e., if  $A'$  and  $Y'$  are the matrices at time  $t$ , then  $z_i^{(t)} = x_i - (A'Y')_i$ . We then find a vector  $v$  that has a high correlation with many of the  $z_i^{(t)}$  by solving an appropriate instance of the  $\tau$ -TC problem. The  $v$  is then added as a new column to  $A'$ . A new row is then added to  $Y'$ , and the  $i$ th entry is non-zero only if  $|\langle v, z_i^{(t)} \rangle|$  is above an appropriately chosen threshold. This procedure is formally described in Algorithm 1.

**Analysis.** There are two main steps in the analysis. The first is showing that the *residual* problem always has a good solution for  $v$ . It turns out that this is a non-trivial step in the analysis because a natural set-cover style argument, which involves considering the optimal dictionary  $A^*$  and showing that one of the columns allows us to make sufficient progress, faces two difficulties. Firstly, the columns of  $A^*$  are arbitrary. There turns out to be a significant difference in the guarantee one can obtain when the columns of  $A^*$  are (near) orthogonal, and when they are arbitrary. Another (more serious) difficulty arises due to the fact that the residual vector  $z_i^{(t)}$  is a projection of  $x_i$  onto some (potentially arbitrary) subspace  $T$  (formed by a subset of the columns in the current  $A'$ ). While it is not difficult to argue that one of the columns of  $A^*$  has a high inner product with  $x_i$ , it is trickier to say the same about  $z_i^{(t)}$ .

The second step in the analysis is to use the guarantees of the bi-criteria approximation to the  $\tau$ -TC problem to bound the number of iterations and the sparsity. Both of these bounds are relatively straightforward. The relation between the number of iterations and the error is shown via an analysis similar to that of Frank-Wolfe iteration (see Clarkson (2010)). As for the sparsity, the manner in which we update  $Y'$  (non-zero only if  $\langle z_i^{(t)}, v \rangle$  is above a threshold) implies that every column in  $Y'$  is sparse.

Let us now proceed with the details of the analysis. Apart from the notation  $X, A^*, Y^*, \varepsilon$  from the statement of Theorem 4, we also define the following quantities:

- $z_i^{(t)}$  is the “residual” vector, as described in the algorithm (and the outline above).
- $\gamma_i := \|(X - A^*Y^*)_i\|^2 / \|x_i\|^2$ , i.e., the optimum error in column  $i$ , as a fraction of its mass. Clearly,

$$\gamma^* = \frac{1}{\|X\|_F^2} \sum_i \|x_i\|^2 \gamma_i.$$

---

**Algorithm 1** DICTAPPROX( $A \in \mathbb{R}^{d \times n}$ ,  $k, m, \Lambda, \varepsilon$ )
 

---

- 1: Initialize  $z_i^{(0)} = x_i$ . Set  $\tau = \frac{\varepsilon^2}{k\Lambda}$ , and  $M = mk$ .
- 2: Initialize  $A', Y'$  to empty matrices
- 3: **for**  $t = 1 : M$  **do**
- 4: Let  $v$  be the output of an  $(\alpha, \beta)$ -approximation algorithm for  $\tau$ -TC, with parameters being:

$$\tau = \frac{\varepsilon^2}{k\Lambda}, \text{ the vectors } \frac{z_i^{(t)}}{\|x_i\|}, \text{ and weights } w_i = \|x_i\|^2, \text{ for } 1 \leq i \leq n.$$

- 5: Add  $v$  as a column to  $A'$
  - 6: Add an empty row to  $Y'$
  - 7: **for**  $i \in [n]$  satisfying  $\langle z_i^{(t)}, v \rangle^2 \geq \alpha\tau$  **do**
  - 8: set  $z_i^{(t+1)} = z_i^{(t)} - \langle z_i^{(t)}, v \rangle v$
  - 9: set  $Y'[t, i] = \langle z_i^{(t)}, v \rangle$
  - 10: **end for**
  - 11: **end for**
  - 12: Return  $A', Y'$
- 

If we view  $\|x_i\|^2 / \|X\|_F^2$  as defining a probability distribution over indices, the above is equivalent to  $\gamma^* = \mathbb{E}_i[\gamma_i]$ .

- $S_i^*$  is the support of  $Y_i^*$ , i.e., the subset of the dictionary used to represent  $x_i$  in the optimum solution.
- $\theta_i^{(t)} = \left\| z_i^{(t)} \right\|^2 / \|x_i\|^2$ , i.e., the error in column  $i$  at time  $t$ , as a fraction of its mass.
- Let  $\psi^{(t)}$  (which will be the potential) be  $\mathbb{E}_i[\theta_i^{(t)}]$ , i.e.,

$$\psi^{(t)} := \frac{1}{\|X\|_F^2} \sum_i \|x_i\|^2 \theta_i^{(t)} = \frac{1}{\|X\|_F^2} \sum_i \left\| z_i^{(t)} \right\|^2.$$

- For any real number  $\xi$ , we denote  $(\xi)_+ = \max\{0, \xi\}$ .

The quantity  $\psi^{(t)}$  defined above will be the potential function we use in the analysis. Let us first start with a simple observation:

**Observation 8** At any step,  $z_i^{(t)}$  is the projection of  $x_i$  onto the subspace orthogonal to the space spanned by a subset of the columns of the current  $A'$ .

**Proof** The proof follows immediately from the fact that  $z_i^{(0)} = x_i$ , and the update step (step 8). ■

The observation also makes sure that the instance of  $\tau$ -TC we are solving is indeed valid. I.e., the vectors  $z_i^{(t)} / \|x_i\|$  are in the unit ball.

Our goal now is to show the following bound on the convergence of the algorithm:



**Theorem 9** After  $t$  iterations of the main loop in the algorithm, the error  $\psi^{(t)}$  is bounded by  $\gamma^* + \frac{16m\Lambda}{\beta t}$ . Further, the number of non-zero entries in each column of  $Y'$  at the end of the algorithm is at most  $1/\alpha\tau$ , where  $\tau = \varepsilon^2/k\Lambda$  as in the algorithm.

This clearly implies Theorem 4 by setting  $t$  appropriately. The goal is thus to show Theorem 9.

### 3.1. Increment lemma

As described in the outline, a key step of our argument is showing that in each step, there exists a vector with a large inner product with a large number of the residual vectors.

**Lemma 10** Let  $u, s_1, s_2, \dots, s_k$  be unit vectors in  $\mathbb{R}^d$ , and let  $u = \sum_i \alpha_i s_i + z$ . Let  $T$  be any subspace of  $\mathbb{R}^d$ , and define  $\theta = \|u - \Pi_T u\|$ . Then, we have

$$\sum_i \langle u - \Pi_T u, s_i \rangle^2 \geq \frac{(\theta^2 - \|z\|^2)_+^2}{4(\sum_i \alpha_i^2)},$$

where  $(\beta)_+$  denotes  $\max(0, \beta)$ , for any real number  $\beta$ .

**Proof** We may assume that  $\theta^2 \geq \|z\|^2$ , else there is nothing to prove. Now define  $x = \sum_i \alpha_i s_i$  and  $u^\perp := u - \Pi_T u$ , for convenience. We will show that

$$\langle u^\perp, x \rangle \geq \frac{\theta^2 - \|z\|^2}{2}, \quad (1)$$

and the desired result then follows by an application of the Cauchy-Schwartz inequality, which implies that  $(\sum_i \alpha_i \langle u^\perp, s_i \rangle) \leq (\sum_i \alpha_i^2)(\sum_i \langle u^\perp, s_i \rangle^2)$ . We can now show (1) as follows.

$$\begin{aligned} \langle u^\perp, x \rangle &= \langle u, x \rangle - \langle \Pi_T u, x \rangle \\ &= \frac{\|u\|^2 + \|x\|^2 - \|z\|^2}{2} - \langle \Pi_T u, x \rangle \\ &\geq \frac{\|u\|^2 + \|x\|^2 - \|z\|^2}{2} - \|\Pi_T u\| \|x\| \\ &\geq \frac{1 + \|x\|^2 - \|z\|^2}{2} - \frac{\|\Pi_T u\|^2 + \|x\|^2}{2} \\ &= \frac{\|u^\perp\|^2 - \|z\|^2}{2}. \end{aligned}$$

In the last step, we used the Pythagoras theorem, concluding that  $1 - \|\Pi_T u\|^2 = \|u^\perp\|^2$ . This concludes the proof of the lemma.  $\blacksquare$

As an immediate consequence we have the following (recall that  $S_i^*$  is the subset of the dictionary vectors used to represent  $x_i$  in the optimum solution).

**Lemma 11** Let  $S_i^* = \{s_{i,1}^*, s_{i,2}^*, \dots, s_{i,k}^*\}$ . For every  $i \in [n]$  and iteration  $t$ , we have the following:

$$\sum_{j \in [k]} \langle z_i^{(t)}, s_{i,j}^* \rangle^2 \geq \frac{(\theta_i^{(t)} - \gamma_i)_+^2}{4\Lambda} \|x_i\|^2.$$

This follows directly from Lemma 10. The next lemma is a convexity argument, relating the sum of the “per-column” improvements to the overall progress. We defer the proof to Appendix A.

**Lemma 12** *In every iteration  $t$ , we have*

$$\sum_{i \in [n]} (\theta_i^{(t)} - \gamma_i)_+^2 \|x_i\|^2 \geq (\psi^{(t)} - \gamma^*)_+^2 \|X\|_F^2.$$

The next lemma is crucial to the argument. It shows the existence of a good solution to the instance of the  $\tau$ -TC problem, unless the current error is already close to the optimum.

**Lemma 13** *Consider any iteration  $t$  for which  $\psi^{(t)} \geq \gamma^* + \varepsilon$ . There exists a vector  $a \in A^*$ , and a subset  $R$  of  $[n]$  such that*

$$\langle z_i^{(t)}, a \rangle^2 \geq \frac{\varepsilon^2}{16k\Lambda} \|x_i\|^2 \text{ for all } i \in R,$$

and additionally,

$$\sum_{i \in R} \langle z_i^{(t)}, a \rangle^2 \geq \frac{(\psi^{(t)} - \gamma^*)^2}{16m\Lambda} \|X\|_F^2.$$

**Proof** The rough idea of the proof is to use Lemma 11 along with Lemma 12 to conclude that there exists a column  $a$  of the optimal dictionary  $A^*$  such that  $\sum_i \langle z_i^{(t)}, a \rangle^2$  is large. However, we also need every term to be large enough. This requires additional pruning steps. We defer the details to Appendix B. ■

As a corollary, we show that the vector found by the algorithm satisfies appropriate guarantees.

**Corollary 14** *Let  $v$  be the vector output by an  $(\alpha, \beta)$ -approximation to the  $\tau$ -TC problem, with*

$$\tau = \frac{\varepsilon^2}{16k\Lambda}, \text{ the vectors } \frac{z_i^{(t)}}{\|x_i\|}, \text{ and weights } w_i = \|x_i\|^2, \text{ for } 1 \leq i \leq n.$$

*Then there exists a subset  $R' \subseteq [n]$  such that  $\langle z_i^{(t)}, v \rangle^2 \geq \alpha\tau \|x_i\|^2$  for all  $i \in R'$ , and*

$$\sum_{i \in R'} \langle z_i^{(t)}, v \rangle^2 \geq \frac{\beta(\psi^{(t)} - \gamma^*)^2}{16m\Lambda} \|X\|_F^2.$$

The corollary follows from Lemma 13 and the definition of an  $(\alpha, \beta)$ -approximation (Definition 2). Next, we are ready to prove Theorem 9.

**Proof** [of Theorem 9] From Corollary 14, we have that

$$\psi^{(t)} - \psi^{(t+1)} \geq C(\psi^{(t)} - \gamma^*)^2, \text{ where } C = \frac{\beta}{16m\Lambda}.$$

Writing  $a_t = (\psi^{(t)} - \gamma^*)$ , the above may be written as  $a_t - a_{t+1} \geq C \cdot a_t^2 \geq C \cdot a_t a_{t+1}$ . Rearranging,

$$\frac{1}{a_{t+1}} - \frac{1}{a_t} \geq C.$$

Using  $a_0 \leq 1$ , we get that for all  $t$ ,  $\frac{1}{a_t} \geq Ct$ , or

$$\psi^{(t)} - \gamma^* = a_t \leq \frac{1}{Ct} = \frac{16m\Lambda}{\beta t}.$$

This completes the proof of the error bound. Next, to bound the sparsity per column, note that every time we use one of the vectors  $v$  to update  $z_i^{(t)}$ , the quantity  $\|z_i^{(t)}\|^2$  drops by at least  $\alpha\tau \|x_i\|^2$ . Thus the number of such updates must be at most  $1/(\alpha\tau)$ . This completes the proof of the theorem. ■

#### 4. Approximability of the threshold correlation problem

We first present our main algorithmic result for the  $\tau$ -TC problem (Theorem 6). It is a bi-criteria approximation, in the sense of Definition 2.

##### 4.1. Bi-criteria approximation

The key lemma behind our algorithm is the following.

**Lemma 15** *Let  $v_1, v_2, \dots, v_q \in \mathcal{B}_2^d$ , and suppose that there exists a unit vector  $x$  such that  $\langle x, v_i \rangle^2 \geq \tau$  for all  $i \in [q]$ . Suppose also that we are given weights  $w_i \geq 0$  for each  $i$ . Then there exists an index  $\ell$  such that*

$$\sum_i w_i \langle v_\ell, v_i \rangle^2 [\mathbf{1}_{\langle v_\ell, v_i \rangle^2 \geq \tau^2/4}] \geq \frac{\tau^2}{32} \sum_i w_i \langle x, v_i \rangle^2.$$

First, let us see how Theorem 6 follows from Lemma 15.

**Proof** [of Theorem 6] Consider any instance of the  $\tau$ -TC problem, and let  $x$  be the optimum vector. Let  $S$  be the set of indices  $\{i \in [n] : \langle x, v_i \rangle^2 \geq \tau\}$ . Applying Lemma 15 to this set of vectors gives us a  $v_\ell$  satisfying the conclusion of the lemma. As this is a vector of length  $\leq 1$ , scaling it to a unit vector only improves the objective. Thus the algorithm can simply iterate over all the vectors  $v_i$  and pick the one that satisfies the conclusion of the lemma. This completes the proof. ■

Let us thus prove Lemma 15.

**Proof** [of Lemma 15] For every  $i \in [q]$ , let us write  $v_i = \alpha_i x + u_i$ , with  $\langle x, u_i \rangle = 0$ . By assumption,  $\alpha_i^2 \geq \tau$  for all  $i \in [q]$ . Let  $I \subset [q]$  be the set of indices with  $\alpha_i > 0$ . By replacing  $x$  by  $-x$  if necessary, we may assume that

$$\sum_{i \in I} w_i \alpha_i^2 \geq \frac{1}{2} \sum_{i \in [q]} w_i \alpha_i^2. \quad (2)$$

For any two  $i, j \in I$ , we have

$$\langle v_i, v_j \rangle = \alpha_i \alpha_j + \langle u_i, u_j \rangle. \quad (3)$$

Thus we have that  $\langle v_i, v_j \rangle < \alpha_i \alpha_j / 2 \implies \langle u_i, u_j \rangle < -\alpha_i \alpha_j / 2 \leq -\tau/2$ . Now, consider the following clustering procedure for the vectors  $v_i$ ,  $i \in I$ . Start with any unclustered  $v_i$ , and place all  $j$  such that  $\langle v_i, v_j \rangle \geq \alpha_i \alpha_j / 2$  in cluster  $C_i$  (in this case, we say that  $i$  ‘‘captures’’  $j$ ). We repeat this

procedure until no unclustered  $v_i$  remains. We claim that in this process, at most  $2\lceil\frac{1}{\tau}\rceil + 1$  clusters can be formed. To see this, consider the indices  $i$  that occur in the procedure above, and call them  $i_1, i_2, \dots, i_t$ . Because  $i_{r+1}$  is not captured by  $i_1, \dots, i_r$ , we have  $\langle v_{i_a}, v_{i_b} \rangle < (\alpha_{i_a} \alpha_{i_b})/2$  for all  $a, b \in [t]$ . From (3), this means that  $\langle u_{i_a}, u_{i_b} \rangle < -\tau/2$  for all  $a, b \in [t]$ .

Now, we appeal to the simple observation that for any integer  $k \geq 1$ , we cannot have  $k + 1$  unit vectors whose pairwise inner products are all  $< -\frac{1}{k}$ . (This contradicts the fact that the Gram matrix is positive semidefinite.)

Thus, at most  $(2\lceil\frac{1}{\tau}\rceil + 1) < 4/\tau$  clusters are formed. This means that there exists one cluster  $C$ , say the one formed using the vector  $v_\ell$ , such that

$$\sum_{i \in C} w_i \alpha_i^2 \geq \frac{\tau}{4} \cdot \sum_{i \in I} w_i \alpha_i^2. \quad (4)$$

(I.e., if we view  $w_i \alpha_i^2$  as the *mass* of index  $i$ , one of the clusters has at least  $\tau/4$  of the mass, due to the bound on the number of clusters.)

Next, for all  $i \in C$ , we have (by the way clustering was performed) that

$$\langle v_\ell, v_i \rangle^2 \geq \frac{\alpha_\ell^2 \alpha_i^2}{4} \geq \frac{\tau \alpha_i^2}{4}. \quad (5)$$

The RHS is always  $\geq \tau^2/4$ . This implies that

$$\sum_{i \in [q]} w_i \langle v_\ell, v_i \rangle^2 [\mathbf{1}_{\langle v_\ell, v_i \rangle^2 \geq \tau^2/4}] \geq \sum_{i \in C} w_i \langle v_\ell, v_i \rangle^2 \quad (6)$$

Now using (5) followed by (4) and finally (2), we obtain

$$\sum_{i \in [q]} w_i \langle v_\ell, v_i \rangle^2 [\mathbf{1}_{\langle v_\ell, v_i \rangle^2 \geq \tau^2/4}] \geq \frac{\tau}{4} \sum_{i \in C} w_i \alpha_i^2 \geq \frac{\tau^2}{16} \sum_{i \in I} w_i \alpha_i^2 \geq \frac{\tau^2}{32} \sum_{i \in [q]} w_i \alpha_i^2.$$

This completes the proof. ■

## 4.2. Inapproximability – connection to matrix norms

The main result of this section is to prove that the approximability (even for bi-criteria approximations) of the  $\tau$ -TC problem is related to the well-studied question of approximating matrix norms.

Recall the problem of computing  $q \mapsto p$  norm of a matrix.

**Definition 16** Let  $p, q \geq 1$ , and let  $A \in \mathbb{R}^{n \times d}$  be a matrix. We define the  $q \mapsto p$  norm of  $A$  as

$$\|A\|_{q \mapsto p} = \max_{x \in \mathbb{R}^d, x \neq 0} \frac{\|Ax\|_p}{\|x\|_q}.$$

We will specifically consider the case of  $q = 2$  and  $p > 2$ . This is a special case of the so-called *hypercontractive* norms, which refers to the case  $p > q$ . Note that a  $\theta > 0$  approximation to the  $q \mapsto p$  norm problem is a vector  $x$  such that

$$\frac{\|Ax\|_p}{\|x\|_q} \geq \theta \|A\|_{q \mapsto p}.$$

**Theorem 17** *Suppose we have an  $(\alpha, \beta)$ -approximation to the  $\tau$ -TC problem, where  $\alpha, \beta$  are independent of  $\tau$  (e.g., constants). Then there exists an  $\Omega\left(\frac{\alpha^{1/p-1/2}\beta^{1/p}}{\log^{1/p}n}\right)$  approximation to the  $2 \mapsto p$  norm problem.*

We defer the proof to Appendix C.

**Known hardness results for the  $2 \mapsto 4$  norm.** (Barak et al., 2012) proved that assuming the exponential time hypothesis (ETH), it is impossible to obtain a  $2^{(\log n)^{1/4}}$  approximation to the  $2 \mapsto 4$  norm problem for  $n \times n$  matrices in time  $< n^{(\log n)^{1/2}}$  (which is super-polynomial). Together with Theorem 17, this implies that achieving constant  $\alpha, \beta$  via polynomial time algorithms is impossible assuming ETH.

## 5. Dictionary learning with outliers

It turns out that the technique of iteratively *peeling off* mass is useful in obtaining approximations even when we have outliers among the signals  $x_i$ . We will indeed be able to use ideas very similar to those in Section 3 to obtain efficient algorithms even when a  $\rho$  fraction of the columns of  $X$  are arbitrary outliers.

The full details will be presented in Appendix D. We now mention a few key differences. First, we can longer just use the fact that if a lot of mass is uncovered, then we must be able to make progress (because even in the optimal solution, a lot of the *overall* mass is uncovered). In essence, we need to be able to perform the earlier algorithm and the analysis only on the inliers, without knowing the inliers! This turns out to be possible, by maintaining proxies for the appropriate parameters that arise in the analysis.

We defer the details to Appendix D.

## References

- Alekh Agarwal, Animashree Anandkumar, Prateek Jain, Praneeth Netrapalli, and Rashish Tandon. Learning sparsely used overcomplete dictionaries. In *Proceedings of The 27th Conference on Learning Theory, COLT 2014, Barcelona, Spain, June 13-15, 2014*, pages 123–137, 2014. URL <http://jmlr.org/proceedings/papers/v35/agarwal14a.html>.
- M. Aharon, M. Elad, and A. Bruckstein. *rmk*-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322, Nov 2006. ISSN 1053-587X. doi: 10.1109/TSP.2006.881199.
- Sanjeev Arora, Rong Ge, and Ankur Moitra. New algorithms for learning incoherent and overcomplete dictionaries. In *Proceedings of The 27th Conference on Learning Theory, COLT 2014, Barcelona, Spain, June 13-15, 2014*, pages 779–806, 2014. URL <http://jmlr.org/proceedings/papers/v35/arora14.html>.
- Pranjal Awasthi and Aravindan Vijayaraghavan. Towards learning sparsely used dictionaries with arbitrary supports. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 283–296, 2018. doi: 10.1109/FOCS.2018.00035. URL <https://doi.org/10.1109/FOCS.2018.00035>.

- Boaz Barak, Fernando G. S. L. Brandão, Aram Wettroth Harrow, Jonathan A. Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 307–326, 2012. doi: 10.1145/2213977.2214006. URL <https://doi.org/10.1145/2213977.2214006>.
- Boaz Barak, Jonathan A. Kelner, and David Steurer. Dictionary learning and tensor decomposition via the sum-of-squares method. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC '15*, pages 143–151, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3536-2. doi: 10.1145/2746539.2746605. URL <http://doi.acm.org/10.1145/2746539.2746605>.
- Aditya Bhaskara and Aravindan Vijayaraghavan. Approximating matrix p-norms. In *ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA*, pages 497–511, 2011.
- Vijay Bhattiprolu, Mrinalkanti Ghosh, Venkatesan Guruswami, Euiwoong Lee, and Madhur Tulsiani. Inapproximability of matrix  $p \rightarrow q$  norms. *CoRR*, abs/1802.07425, 2018. URL <http://arxiv.org/abs/1802.07425>.
- Avrim Blum, Sariel Har-Peled, and Benjamin Raichel. Sparse approximation via generating point sets. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 548–557, 2016. doi: 10.1137/1.9781611974331.ch40. URL <http://dx.doi.org/10.1137/1.9781611974331.ch40>.
- Kenneth L. Clarkson. Coresets, sparse greedy approximation, and the frank-wolfe algorithm. *ACM Trans. Algorithms*, 6(4):63:1–63:30, September 2010. ISSN 1549-6325. doi: 10.1145/1824777.1824783. URL <http://doi.acm.org/10.1145/1824777.1824783>.
- G. Davis, S. Mallat, and M. Avellaneda. Adaptive greedy approximations. *Constructive Approximation*, 13(1):57–98, Mar 1997. ISSN 1432-0940. doi: 10.1007/BF02678430. URL <https://doi.org/10.1007/BF02678430>.
- Dean P. Foster, Howard J. Karloff, and Justin Thaler. Variable selection is hard. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, pages 696–709, 2015. URL <http://jmlr.org/proceedings/papers/v40/Foster15.html>.
- Rémi Gribonval, Rodolphe Jenatton, and Francis Bach. Sparse and spurious: dictionary learning with noise and outliers. *IEEE Transactions on Information Theory*, page 22, 2015. doi: 10.1109/TIT.2015.2472522. URL <https://hal.archives-ouvertes.fr/hal-01025503>. This is a substantially revised version of a first draft that appeared as a preprint titled "Local stability and robustness of sparse dictionary learning in the presence of noise", <http://hal.inria.fr/hal-00737152>.
- Ragesh Jaiswal, Amit Kumar, and Sandeep Sen. A simple  $d^2$ -sampling based PTAS for k-means and other clustering problems. *CoRR*, abs/1201.4206, 2012. URL <http://arxiv.org/abs/1201.4206>.

- Bruno A. Olshausen and David J. Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision Research*, 37(23):3311 – 3325, 1997. ISSN 0042-6989. doi: [http://dx.doi.org/10.1016/S0042-6989\(97\)00169-7](http://dx.doi.org/10.1016/S0042-6989(97)00169-7). URL <http://www.sciencedirect.com/science/article/pii/S0042698997001697>.
- Rafail Ostrovsky, Yuval Rabani, Leonard J. Schulman, and Chaitanya Swamy. The effectiveness of lloyd-type methods for the k-means problem. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21-24 October 2006, Berkeley, California, USA, Proceedings*, pages 165–176, 2006. doi: 10.1109/FOCS.2006.75. URL <https://doi.org/10.1109/FOCS.2006.75>.
- Shai Shalev-Shwartz, Nathan Srebro, and Tong Zhang. Trading accuracy for sparsity in optimization problems with sparsity constraints. *SIAM J. on Optimization*, 20(6):2807–2832, August 2010. ISSN 1052-6234. doi: 10.1137/090759574. URL <http://dx.doi.org/10.1137/090759574>.
- Daniel A. Spielman, Huan Wang, and John Wright. Exact recovery of sparsely-used dictionaries. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*, pages 3087–3090. AAAI Press, 2013. ISBN 978-1-57735-633-2. URL <http://dl.acm.org/citation.cfm?id=2540128.2540583>.
- A. M. Tillmann. On the computational intractability of exact and approximate dictionary learning. *IEEE Signal Processing Letters*, 22(1):45–49, Jan 2015. ISSN 1070-9908. doi: 10.1109/LSP.2014.2345761.
- J. A. Tropp. Greed is good: algorithmic results for sparse approximation. *IEEE Transactions on Information Theory*, 50(10):2231–2242, Oct 2004. ISSN 0018-9448. doi: 10.1109/TIT.2004.834793.

## Appendix A. Proof of Lemma 12

**Proof** We may assume that  $\psi^{(t)} \geq \gamma^*$ , as there is nothing to prove otherwise.

Now, using the fact that  $\mathbb{E}[Z^2] \geq (\mathbb{E}[Z])^2$  for any random variable  $Z$ , we have that

$$\sum_i \frac{\|x_i\|^2}{\|X\|_F^2} (\theta_i^{(t)} - \gamma_i)_+^2 \geq \left( \sum_i \frac{\|x_i\|^2}{\|X\|_F^2} (\theta_i^{(t)} - \gamma_i)_+ \right)^2. \quad (7)$$

Next, using the fact that  $(\beta)_+ \geq \beta$ , we have

$$\sum_i \frac{\|x_i\|^2}{\|X\|_F^2} (\theta_i^{(t)} - \gamma_i)_+ \geq \sum_i \frac{\|x_i\|^2}{\|X\|_F^2} (\theta_i^{(t)} - \gamma_i) = \psi^{(t)} - \gamma^*.$$

We have noted that  $\psi^{(t)} \geq \gamma^*$ , and thus the RHS is  $\geq 0$ . Thus we can square the inequality above, and together with (7), we have

$$\sum_i \frac{\|x_i\|^2}{\|X\|_F^2} (\theta_i^{(t)} - \gamma_i)_+^2 \geq (\psi^{(t)} - \gamma^*)^2.$$

This completes the proof. ■

## Appendix B. Proof of Lemma 13

The rough idea of the proof is to use Lemma 11 along with Lemma 12 to conclude that there exists a column  $a$  of the optimal dictionary  $A^*$  such that  $\sum_i \langle z_i^{(t)}, a \rangle^2$  is large. However, we also need every term to be large enough. This requires additional pruning steps.

First, we perform a pruning on the indices  $i \in [n]$ . Let

$$I := \left\{ i \in [n] : (\theta_i^{(t)} - \gamma_i)_+^2 \geq \frac{(\psi^{(t)} - \gamma^*)^2}{2} \right\}.$$

By definition, the following holds (the LHS is a sum over indices *not in*  $I$ ):

$$\sum_{i \in [n] \setminus I} (\theta_i^{(t)} - \gamma_i)_+^2 \|x_i\|^2 < \frac{(\psi^{(t)} - \gamma^*)^2}{2} \|X\|_F^2.$$

Thus from Lemma 12, we have that

$$\sum_{i \in I} (\theta_i^{(t)} - \gamma_i)_+^2 \|x_i\|^2 \geq \frac{1}{2} \sum_{i \in [n]} (\theta_i^{(t)} - \gamma_i)_+^2 \|x_i\|^2. \quad (8)$$

Next, for each  $i \in I$ , we consider the bound from Lemma 11, and prune out terms that contribute too little to the sum. For every  $i \in I$ , define  $J_i$  as

$$J_i = \left\{ j \in [k] : \langle z_i^{(t)}, s_{i,j}^* \rangle^2 \geq \frac{(\theta_i^{(t)} - \gamma_i)_+^2}{8k\Lambda} \|x_i\|^2 \right\}.$$

The same argument as above (because  $|S_i^*| = k$ ) implies that

$$\sum_{j \in J_i} \langle z_i^{(t)}, s_{i,j}^* \rangle^2 \geq \frac{1}{2} \sum_{j \in [k]} \langle z_i^{(t)}, s_{i,j}^* \rangle^2. \quad (9)$$

Now, if we imagine each  $i \in I$  contributing a mass of  $\langle z_i^{(t)}, s_{i,j}^* \rangle^2$  to the bin  $s_{i,j}^*$  for each  $j \in J_i$ , we can conclude that there exists a bin (one of the  $[m]$  columns in the dictionary) with a total mass that is at least

$$\frac{1}{m} \sum_{i \in I} \sum_{j \in J_i} \langle z_i^{(t)}, s_{i,j}^* \rangle^2.$$

Combining equations (8) and (9) (and Lemmas 11 and 12), this is at least

$$\frac{1}{16m\Lambda} \sum_{i \in [n]} (\theta_i^{(t)} - \gamma_i)_+^2 \|x_i\|^2 \geq \frac{(\psi^{(t)} - \gamma^*)^2}{16m\Lambda} \|X\|_F^2.$$

Furthermore, for each  $i \in I$  and  $j \in J_i$ , by our definitions, we have

$$\langle z_i^{(t)}, s_{i,j}^* \rangle^2 \geq \frac{(\psi^{(t)} - \gamma^*)^2}{16k\Lambda} \|x_i\|^2 \geq \frac{\varepsilon^2}{16k\Lambda} \|x_i\|^2.$$

This completes the proof.



### Appendix C. Proof of Theorem 17

**Proof** Consider an instance  $A$  of the  $2 \mapsto p$  norm problem. Let  $v_i$  denote the rows of  $A$ . Without loss of generality, we assume that  $\|v_i\| \leq 1$  by dividing  $A$  by  $\max_i \|v_i\|$ . We wish to find  $x$  with  $\|x\| = 1$  such that  $\sum_{i \in [n]} \langle v_i, x \rangle^p$  is maximized.

Let  $Q$  denote the optimum value of  $(\sum_i \langle v_i, x \rangle^p)^{1/p}$ . Namely,  $\sum_i \langle v_i, x_0 \rangle^p = Q^p$  for some unit vector  $x_0$ . Denote  $I = \left\{ i \in [n] \mid \langle v_i, x_0 \rangle^p \geq \frac{Q^p}{2n} \right\}$ . It is easy to see that  $\sum_{i \in I} \langle v_i, x_0 \rangle^p \geq Q^p/2$ .

By dividing the terms  $\langle v_i, x \rangle^2$  into levels (by powers of 2), we have that there exists a level  $(\tau_j, \tau_{j+1})$  such that the sum

$$\sum_{i \in I, \tau_j \leq \langle v_i, x_0 \rangle^2 \leq \tau_{j+1}} \langle v_i, x_0 \rangle^p \geq \frac{Q^p}{2(\log n + 1)}$$

where  $\tau_j = \frac{Q^p}{2n} \cdot 2^j$  for  $j = 0, 1, \dots, (\log n) + 1$ .

Now, suppose we consider the instance of the  $\tau$ -TC problem with this value of  $\tau = \tau_j$ , and the vectors  $v_i$  (with its corresponding weight 1). Let  $\text{OPT}_{TC}$  denote the optimum value of this problem. We clearly have

$$\text{OPT}_{TC} = \max_{\|x\|=1} \sum_{i \in [n], \langle v_i, x \rangle^2 \geq \tau} \langle v_i, x \rangle^2 \geq \frac{Q^p \tau}{2(\log n + 1)(2\tau)^{p/2}}.$$

Thus, consider an  $(\alpha, \beta)$ -approximation to the  $\tau$ -TC problem. This will end up finding a vector  $x$  such that

$$\sum_{i \in [n], \langle v_i, x \rangle^2 \geq \alpha\tau} \langle v_i, x \rangle^2 \geq \beta \cdot \frac{Q^p \tau}{2(\log n + 1)(2\tau)^{p/2}}.$$

For any such vector, we have

$$\sum_{i \in [n]} \langle v_i, x \rangle^p \geq \sum_{i \in [n], \langle v_i, x \rangle^2 \geq \alpha\tau} (\alpha\tau)^{(p-2)/2} \langle v_i, x \rangle^2 \geq (\alpha\tau)^{(p-2)/2} \cdot \frac{\beta Q^p \tau}{2(\log n + 1)(2\tau)^{p/2}} = \frac{\alpha^{(p-2)/2} \beta Q^p}{2^{p/2+1} (\log n + 1)}.$$

■

### Appendix D. Dictionary learning with outliers

As we discussed earlier, the technique of iteratively *peeling off* mass is useful in obtaining approximations even when we have outliers among the signals  $x_i$ . We will now prove this formally, thus establishing Theorem 7. We begin by recalling the definitions of  $A^*$ ,  $Y^*$ , and  $N^*$  (the latter is the matrix that only consists of the outlier columns, and is a  $d \times n$  matrix that is non-zero in only  $\rho n$  columns). Additionally, we define the following:

1.  $\mathcal{I}$ : the indices of the inliers (of course, unknown to the algorithm).
2.  $X_{\mathcal{I}}$ : the submatrix of  $X$  restricted to the set of inliers. It has dimensions  $d \times (1 - \rho)n$ .
3. As before,  $\gamma^* = \frac{1}{\|X_{\mathcal{I}}\|_F^2} \sum_{i \in \mathcal{I}} \|(X - A^*Y^*)_i\|^2$ .

4. As before,  $\psi^{(t)} = \frac{1}{\|X_{\mathcal{I}}\|_F^2} \sum_{i \in \mathcal{I}} \|z_i^{(t)}\|^2$ , where  $z_i^{(t)}$  is the residual (as maintained by the algorithm) in column  $i$  at iteration  $t$ .
5.  $\Phi^{(t)} = \sum_{i \in [n]} \|z_i^{(t)}\|^2$ .
6.  $\widehat{\psi}^{(t)} = \min_{T \subset [n], |T|=(1-\rho)n} \sum_{i \in T} \|z_i^{(t)}\|^2$ .

The two new parameters are  $\Phi^{(t)}$  and  $\widehat{\psi}^{(t)}$ . These are quantities the algorithm can compute at any time step.  $\widehat{\psi}^{(t)}$  is a quantity that we use as a proxy for  $\|X_{\mathcal{I}}\|_F^2 \cdot \psi^{(t)}$ . For any  $t$ , we have

$$\widehat{\psi}^{(t)} \leq \|X_{\mathcal{I}}\|_F^2 \cdot \psi^{(t)}.$$

**Algorithm.** The algorithm is almost identical to the one from Section 3, with one main difference: instead of running for a prescribed number of iterations, we execute the main loop as long as

$$\frac{1}{\|X\|_F^2} (\Phi^{(t)} - \Phi^{(t+1)}) \geq \frac{\beta \varepsilon^3}{16m\Lambda}.$$

As another minor difference, in the end, the algorithm declares the  $\rho n$  columns with the largest values of  $\|z_i^{(t)}\|$  as the outliers (which is the natural choice). The procedure is described in detail in Algorithm 2.

---

**Algorithm 2** OUTLIERDICTAPPROX( $A \in \mathbb{R}^{d \times n}$ ,  $k$ ,  $m$ ,  $\Lambda$ ,  $\rho$ ,  $\varepsilon$ )

---

- 1: Initialize  $z_i^{(0)} = x_i$ . Set  $\tau = \frac{\varepsilon^2}{k\Lambda}$ .
- 2: Initialize  $A'$ ,  $Y'$  to empty matrices
- 3: **while**  $\Phi^{(t+1)} < \Phi^{(t)} - \frac{\beta \varepsilon^3}{16m\Lambda}$  **do**
- 4:   Let  $v$  be the output of an  $(\alpha, \beta)$ -approximation algorithm for  $\tau$ -TC, with parameters being:

$$\tau = \frac{\varepsilon^2}{k\Lambda}, \text{ the vectors } \frac{z_i^{(t)}}{\|x_i\|}, \text{ and weights } w_i = \|x_i\|^2, \text{ for } 1 \leq i \leq n.$$

- 5:   Add  $v$  as a column to  $A'$
  - 6:   Add an empty row to  $Y'$
  - 7:   **for**  $i \in [n]$  satisfying  $\langle z_i^{(t)}, v \rangle^2 \geq \alpha \tau$  **do**
  - 8:     set  $z_i^{(t+1)} = z_i^{(t)} - \langle z_i^{(t)}, v \rangle v$
  - 9:     set  $Y'[t, i] = \langle z_i^{(t)}, v \rangle$
  - 10:   **end for**
  - 11: **end while**
  - 12: Return  $A'$ ,  $Y'$ , and declare the  $\rho n$  columns with the largest  $\|z_i^{(t)}\|$  as the outliers.
- 

**Observation 18** *At any iteration  $t$ , the current decomposition (using the current  $A'$ ,  $Y'$  and declaring the columns with largest  $z_i^{(t)}$  norm as outliers) achieves error  $\widehat{\psi}^{(t)}$ .*

The proof is immediate from the definition of  $\widehat{\psi}^{(t)}$ . One consequence of this is that we may assume that  $\widehat{\psi}^{(0)} \geq \varepsilon \|X\|_F^2$  (i.e., the Frobenius norm of the entire matrix  $X$ ). Otherwise, we return  $A'$ ,  $X'$  being 0 matrices (with say one column and row respectively) and declare the  $\rho n$  vectors of largest norm as outliers  $N'$  which clearly satisfies the guarantee in Theorem 7.

We can now directly use our arguments from Section 3.

**Proof** [of Theorem 7] Consider any iteration  $t$ . Clearly, we have that  $\Phi^{(t)} - \Phi^{(t+1)} \geq \|X_{\mathcal{I}}\|^2 (\psi^{(t)} - \psi^{(t+1)})$ . This is simply because for the columns not in  $\mathcal{I}$ ,  $\|z_i^{(t+1)}\| \leq \|z_i^{(t)}\|$ . Thus the (known) quantity  $\Phi^{(t)}$  is used as a proxy for the (unknown) quantity  $\psi^{(t)}$ .

Now, from Corollary 14, we have that if  $\psi^{(t)} \geq \gamma^* + \varepsilon$ ,

$$\Phi^{(t)} - \Phi^{(t+1)} \geq \|X_{\mathcal{I}}\|_F^2 (\psi^{(t)} - \psi^{(t+1)}) \geq \frac{\beta \|X_{\mathcal{I}}\|_F^2}{16m\Lambda} \cdot (\psi^{(t)} - \gamma^*)^2 \geq \frac{\beta \varepsilon^2 \|X_{\mathcal{I}}\|_F^2}{16m\Lambda} \geq \frac{\beta \varepsilon^3 \|X\|_F^2}{16m\Lambda}. \quad (10)$$

In the very last inequality we used the observation from above that  $\widehat{\psi}^{(0)} \geq \varepsilon \|X\|_F^2$ , and clearly,  $\|X_{\mathcal{I}}\|_F^2 \geq \widehat{\psi}^{(0)}$ , because the definition of  $\widehat{\psi}$  ignores the  $\rho n$  columns of largest length (while  $\|X_{\mathcal{I}}\|$  ignores the outliers).

This means that when the algorithm terminates, we have  $\psi^{(t)} < \gamma^* + \varepsilon$ , and therefore also

$$\widehat{\psi}^{(t)} < \|X_{\mathcal{I}}\|_F^2 (\gamma^* + \varepsilon) = \|X - A^*Y^* - N^*\|_F^2 + \varepsilon \|X_{\mathcal{I}}\|_F^2.$$

This establishes that when the algorithm terminates, we have a good bound on the error in our approximation. Next, we need to bound the number of iterations. For this, we simply use the crude bound obtained by the drop in  $\Phi^{(t)}$ .  $\Phi^{(0)} \leq \|X\|_F^2$ , and  $\Phi^{(t)} \geq 0$  for any  $t$ . Thus the number of iterations is at most  $\frac{16m\Lambda}{\varepsilon^3\beta}$ .

The proof of the sparsity per column of the output  $Y'$  is precisely the same as in Section 3, as the update to  $z$  as well as the value of  $\tau$  are the same.  $\blacksquare$