Recovery Guarantee of Non-negative Matrix Factorization via Alternating Updates: Supplementary Material

Yuanzhi Li, Yingyu Liang, Andrej Risteski
Computer Science Department at Princeton University
35 Olden St, Princeton, NJ 08540
`{yuanzhil, yingyul, risteski}@cs.princeton.edu`

Abstract

Non-negative matrix factorization is a popular tool for decomposing data into feature and weight matrices under non-negativity constraints. It enjoys practical success but is poorly understood theoretically. This paper proposes an algorithm that alternates between decoding the weights and updating the features, and shows that assuming a generative model of the data, it provably recovers the ground-truth under fairly mild conditions. In particular, its only essential requirement on features is linear independence. Furthermore, the algorithm uses ReLU to exploit the non-negativity for decoding the weights, and thus can tolerate adversarial noise that can potentially be as large as the signal, and can tolerate unbiased noise much larger than the signal. The analysis relies on a carefully designed coupling between two potential functions, which we believe is of independent interest.

1 Introduction

In this paper, we study the problem of non-negative matrix factorization (NMF), where given a matrix $\mathbf{Y} \in \mathbb{R}^{m \times N}$, the goal to find a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and a *non-negative* matrix $\mathbf{X} \in \mathbb{R}^{n \times N}$ such that $\mathbf{Y} \approx \mathbf{A}\mathbf{X}$.[1] $\mathbf{A}$ is often referred to as *feature matrix* and $\mathbf{X}$ referred as *weights*. NMF has been extensively used in extracting a parts representation of the data (e.g., [LS97, LS99, LS01]). It has been shown that the non-negativity constraint on the coefficients forcing features to combine, but not cancel out, can lead to much more interpretable features and improved downstream performance of the learned features.

Despite all the practical success, however, this problem is poorly understood theoretically, with only few provable guarantees known. Moreover, many of the theoretical algorithms are based on heavy tools from algebraic geometry (e.g., [AGKM12]) or tensors (e.g. [AKF+12]), which are still not as widely used in practice primarily because of computational feasibility issues or sensitivity to assumptions on $\mathbf{A}$ and $\mathbf{X}$. Some others depend on specific structure of the feature matrix, such as separability [AGKM12] or similar properties [BGKP16].

A natural family of algorithms for NMF alternate between decoding the weights and updating the features. More precisely, in the decoding step, the algorithm represents the data as a non-negative combination of the current set of features; in the updating step, it updates the features using the decoded representations. This meta-algorithm is popular in practice due to ease of implementation, computational efficiency, and empirical quality of the recovered features. However, even less theoretical analysis exists for such algorithms.

---

[1]In the usual formulation of the problem, $\mathbf{A}$ is also assumed to be non-negative, which we will not require in this paper.

30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

This paper proposes an algorithm in the above framework with provable recovery guarantees. To be specific, the data is assumed to come from a generative model $y = \mathbf{A}^* x^* + \nu$. Here, $\mathbf{A}^*$ is the ground-truth feature matrix, x^* are the non-negative ground-truth weights generated from an unknown distribution, and ν is the noise. Our algorithm can provably recover $\mathbf{A}^*$ under mild conditions, even in the presence of large adversarial noise.

Overview of main results. The existing theoretical results on NMF can be roughly split into two categories. In the first category, they make heavy structural assumptions on the feature matrix $\mathbf{A}^*$ such as separability ([AGM12]) or allowing running time exponential in n ([AGKM12]). In the second one, they impose strict distributional assumptions on x^* ([AKF$^+$12]), where the methods are usually based on the method of moments and tensor decompositions and have poor tolerance to noise, which is very important in practice.

In this paper, we present a very simple and natural alternating update algorithm that achieves the best of both worlds. First, we have minimal assumptions on the feature matrix $\mathbf{A}^*$: the only essential condition is linear independence of the features. Second, it is robust to adversarial noise ν which in some parameter regimes be potentially be on the same order as the signal $\mathbf{A}^* x^*$, and is robust to unbiased noise potentially even higher than the signal by a factor of $O(\sqrt{n})$. The algorithm does not require knowing the distribution of x^*, and allows a fairly wide family of interesting distributions. We get this at a rather small cost of a mild "warm start". Namely, we initialize each of the features to be "correlated" with the ground truth features. This type of initialization is often used in practice as well, for example in LDA-c, the most popular software for topic modeling ([lda16]).

A major feature of our algorithm is the significant robustness to noise. In the presence of adversarial noise on each entry of y up to level C_ν, the noise level $\|\nu\|_1$ can be in the same order as the signal $\mathbf{A}^* x^*$. Still, our algorithm is able to output a matrix $\mathbf{A}$ such that the *final* $\|\mathbf{A}^* - \mathbf{A}\|_1 \leq O(\|\nu\|_1)$ in the order of the noise in *one* data point. If the noise is unbiased (i.e., $\mathbb{E}[\nu|x^*] = 0$), the noise level $\|\nu\|_1$ can be $\Omega(\sqrt{n})$ times larger than the signal $\mathbf{A}^* x^*$, while we can still guarantee $\|\mathbf{A}^* - \mathbf{A}\|_1 \leq O(\|\nu\|_1 \sqrt{n})$ – so our algorithm is not only tolerant to noise, but also has very strong denoising effect. Note that even for the unbiased case the noise can potentially be correlated with the ground-truth in very complicated manner, and also, all our results are obtained only requiring the columns of $\mathbf{A}^*$ are independent.

Technical contribution. The success of our algorithm crucially relies on exploiting the non-negativity of x^* by a ReLU thresholding step during the decoding procedure. Similar techniques have been considered in prior works on matrix factorization, however to the best of our knowledge, the analysis (e.g., [AGMM15]) requires that the decodings are correct in all the intermediate iterations, in the sense that the supports of x^* are recovered with no error. Indeed, we cannot hope for a similar guarantee in our setting, since we consider adversarial noise that could potentially be the same order as the signal. Our major technical contribution is a way to deal with the erroneous decoding through out all the intermediate iterations. We achieve this by a coupling between two potential functions that capture different aspects of the working matrix $\mathbf{A}$. While analyzing iterative algorithms like alternating minimization or gradient descent in non-convex settings is a popular topic in recent years, the proof usually proceeds by showing that the updates are approximately performing gradient descent on an objective with some local or hidden convex structure. Our technique diverges from the common proof strategy, and we believe is interesting in its own right.

Organization. After reviewing related work, we define the problem in Section 3 and describe our main algorithm in Section 4. To emphasize the key ideas, we first present the results and the proof sketch for a simplified yet still interesting case in Section 5, and then present the results under much more general assumptions in Section 6. The complete proof is provided in the appendix.

2 Related work

Non-negative matrix factorization relates to several different topics in machine learning.

Non-negative matrix factorization. The area of non-negative matrix factorization (NMF) has a rich empirical history, starting with the practical algorithm of [LS97].On the theoretical side, [AGKM12] provides a fixed-parameter tractable algorithm for NMF, which solves algebraic equations and thus has poor noise tolerance. [AGKM12] also studies NMF under separability assumptions about the

2

features. [BGKP16] studies NMF under heavy noise, but also needs assumptions related to separability, such as the existence of dominant features. Also, their noise model is different from ours.

Topic modeling. A closely related problem to NMF is topic modeling, a common generative model for textual data [BNJ03, Ble12]. Usually, $\|x^*\|_1 = 1$ while there also exist work that assume $x_i^* \in [0, 1]$ and are independent [ZX12]. A popular heuristic in practice for learning $\mathbf{A}^*$ is *variational inference*, which can be interpreted as alternating minimization in KL divergence norm. On the theory front, there is a sequence of works by based on either spectral or combinatorial approaches, which need certain "non-overlapping" assumptions on the topics. For example, [AGH$^+$13] assume the topic-word matrix contains "anchor words": words which appear in a single topic. Most related is the work of [AR15] who analyze a version of the variational inference updates when documents are long. However, they require strong assumptions on both the warm start, and the amount of "non-overlapping" of the topics in the topic-word matrix.

ICA. Our generative model for x^* will assume the coordinates are independent, therefore our problem can be viewed as a non-negative variant of ICA with high levels of noise. Results here typically are not robust to noise, with the exception of [AGMS12] that tolerates Gaussian noise. However, to best of our knowledge, no result in this setting is provably robust to adversarial noise.

Non-convex optimization. The framework of having a "decoding" for the samples, along with performing an update for the model parameters has proven successful for dictionary learning as well. The original empirical work proposing such an algorithm (in fact, it suggested that the V1 layer processes visual signals in the same manner) was due to [OF97]. Even more, similar families of algorithms based on "decoding" and gradient-descent are believed to be neurally plausible as mechanisms for a variety of tasks like clustering, dimension-reduction, NMF, etc ([PC15a, PC14]). A theoretical analysis came latter for dictionary learning due to [AGMM15] under the assumption that the columns of $\mathbf{A}^*$ are incoherent. The technique is not directly applicable to our case, as we don't wish to have any assumptions on the matrix $\mathbf{A}^*$. For instance, if $\mathbf{A}^*$ is non-negative and columns with l_1 norm 1, incoherence effectively means the the columns of $\mathbf{A}^*$ have very small overlap.

3 Problem definition and assumptions

Given a matrix $\mathbf{Y} \in \mathbb{R}^{m \times N}$, the goal of non-negative matrix factorization (NMF) is to find a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and a non-negative matrix $\mathbf{X} \in \mathbb{R}^{n \times N}$, so that $\mathbf{Y} \approx \mathbf{AX}$. The columns of $\mathbf{Y}$ are called data points, those of $\mathbf{A}$ are features, and those of $\mathbf{X}$ are weights. We note that in the original NMF, $\mathbf{A}$ is also assumed to be non-negative, which is not required here. We also note that typically $m \gg n$, i.e., the features are a few representative components in the data space. This is different from dictionary learning where overcompleteness is often assumed.

The problem in the worst case is NP-hard [AGKM12], so some assumptions are needed to design provable efficient algorithms. In this paper, we consider a generative model for the data point

$$y = \mathbf{A}^* x^* + \nu \tag{1}$$

where $\mathbf{A}^*$ is the ground-truth feature matrix, x^* is the ground-truth non-negative weight from some unknown distribution, and ν is the noise. Our focus is to recover $\mathbf{A}^*$ given access to the data distribution, assuming some properties of $\mathbf{A}^*$, x^*, and ν. To describe our assumptions, we let $[\mathbf{M}]^i$ denote the i-th row of a matrix $\mathbf{M}$, $[\mathbf{M}]_j$ its i-th column, $\mathbf{M}_{i,j}$ its (i, j)-th entry. Denote its column norm, row norm, and symmetrized norm as $\|\mathbf{M}\|_1 = \max_j \sum_i |\mathbf{M}_{i,j}|$, $\|\mathbf{M}\|_\infty = \max_i \sum_j |\mathbf{M}_{i,j}|$, and $\|\mathbf{M}\|_s = \max \{\|\mathbf{M}\|_1, \|\mathbf{M}\|_\infty\}$, respectively.

We assume the following hold for parameters $C_1, c_2, C_2, \ell, C_\nu$ to be determined in our theorems.

(**A1**) The columns of $\mathbf{A}^*$ are linearly independent.

(**A2**) For all $i \in [n]$, $x_i^* \in [0, 1]$, $\mathbb{E}[x_i^*] \leq \frac{C_1}{n}$ and $\frac{c_2}{n} \leq \mathbb{E}[(x_i^*)^2] \leq \frac{C_2}{n}$, and x_i^*'s are independent.

(**A3**) The initialization $\mathbf{A}^{(0)} = \mathbf{A}^*(\mathbf{\Sigma}^{(0)} + \mathbf{E}^{(0)}) + \mathbf{N}^{(0)}$, where $\mathbf{\Sigma}^{(0)}$ is diagonal, $\mathbf{E}^{(0)}$ is off-diagonal, and

$$\mathbf{\Sigma}^{(0)} \succeq (1 - \ell)\mathbf{I}, \quad \left\|\mathbf{E}^{(0)}\right\|_s \leq \ell.$$

We consider two noise models.

3

(N1) Adversarial noise: only assume that $\max_i |\nu_i| \leq C_\nu$ almost surely.

(N2) Unbiased noise: $\max_i |\nu_i| \leq C_\nu$ almost surely, and $\mathbb{E}[\nu|x^*] = 0$.

Remarks We make several remarks about each of the assumptions.
(A1) is the assumption about $\mathbf{A}^*$. It only requires the columns of $\mathbf{A}^*$ to be linear independent, which is very mild and needed to ensure identifiability. Otherwise, for instance, if $(\mathbf{A}^*)_3 = \lambda_1(\mathbf{A}^*)_1 + \lambda_2(\mathbf{A}^*)_2$, it is impossible to distinguish between the case when $x_3^* = 1$ and the case when $x_2^* = \lambda_1$ and $x_1^* = \lambda_2$. In particular, we do not restrict the feature matrix to be non-negative, which is more general than the traditional NMF and is potentially useful for many applications. We also do not make incoherence or anchor word assumptions that are typical in related work.

(A2) is the assumption on x^*. First, the coordinates are non-negative and bounded by 1; this is simply a matter of scaling. Second, the assumption on the moments requires that, roughly speaking, each feature should appear with reasonable probability. This is expected: if the occurrences of the features are extremely unbalanced, then it will be difficult to recover the rare ones. The third requirement on independence is motivated by that the features should be different so that their occurrences are not correlated. Here we do not stick to a specific distribution, since the moment conditions are more general, and highlight the essential properties our algorithm needs. Example distributions satisfying our assumptions will be discussed later.

The warm start required by **(A3)** means that each feature $A_i^{(0)}$ has a large fraction of the ground-truth feature $\mathbf{A}_i^*$ and a small fraction of the other features, plus some noise outside the span of the ground-truth features. We emphasize that $\mathbf{N}^{(0)}$ is the component of $\mathbf{A}^{(0)}$ outside the column space of $\mathbf{A}^*$, and is not the difference between $\mathbf{A}^{(0)}$ and $\mathbf{A}^*$. This requirement is typically achieved in practice by setting the columns of $\mathbf{A}^{(0)}$ to reasonable "pure" data points that contains one major feature and a small fraction of some other features (e.g. [lda16, AR15]); in this initialization, it is generally believed that $\mathbf{N}^{(0)} = 0$. But we state our theorems to allow some noise $\mathbf{N}^{(0)}$ for robustness in the initialization.

The adversarial noise model **(N1)** is very general, only imposing an upper bound on the entry-wise noise level. Thus, ν can be correlated with x^* in some complicated unknown way. **(N2)** additionally requires it to be zero mean, which is commonly assumed and will be exploited by our algorithm to tolerate larger noise.

4   Main algorithm

Algorithm 1 Purification

Input: initialization $\mathbf{A}^{(0)}$, threshold α, step size η, scaling factor r, sample size N, iterations T
  1: **for** $t = 0, 1, 2, ..., T-1$ **do**
  2:   Draw examples $y_1, \ldots, y_N$.
  3:   (Decode) Compute $\mathbf{A}^\dagger$, the pseudo-inverse of $\mathbf{A}^{(t)}$ with minimum $\|(\mathbf{A})^\dagger\|_\infty$.
       Set $x = \phi_\alpha(\mathbf{A}^\dagger y)$ for each example y.   // ϕ_α is ReLU activation; see (2) for the definition
  4:   (Update) Update the feature matrix
$$\mathbf{A}^{(t+1)} = (1-\eta)\,\mathbf{A}^{(t)} + r\eta\hat{\mathbb{E}}\left[(y-y')(x-x')^\top\right]$$
       where $\hat{\mathbb{E}}$ is over independent uniform y, y' from $\{y_1, \ldots, y_N\}$, and x, x' are their decodings.
Output: $\mathbf{A} = \mathbf{A}^{(T)}$

Our main algorithm is presented in Algorithm 1. It keeps a working feature matrix and operates in iterations. In each iteration, it first compute the weights for a batch of N examples (*decoding*), and then uses the computed weights to update the feature matrix (*updating*).

The decoding is simply multiplying the example by the pseudo-inverse of the current feature matrix and then passing it through the rectified linear unit (ReLU) ϕ_α with offset α. The pseudo-inverse with minimum infinity norm is used so as to maximize the robustness to noise (see the theorems). The ReLU function ϕ_α operates element-wisely on the input vector v, and for an element v_i, it is

4

defined as

$$\phi_\alpha(v_i) = \max\{v_i - \alpha, 0\}. \tag{2}$$

To get an intuition why the decoding makes sense, suppose the current feature matrix is the ground-truth. Then $\mathbf{A}^\dagger y = \mathbf{A}^\dagger \mathbf{A}^* x^* + \mathbf{A}^\dagger \nu = x^* + \mathbf{A}^\dagger \nu$. So we would like to use a small $\mathbf{A}^\dagger$ and use threshold to remove the noise term.

In the encoding step, the algorithm move the feature matrix along the direction $\mathbb{E}\left[(y - y')(x - x')^\top\right]$. To see intuitively why this is a good direction, note that when the decoding is perfect and there is no noise, $\mathbb{E}\left[(y - y')(x - x')^\top\right] = \mathbf{A}^*$, and thus it is moving towards the ground-truth. Without those ideal conditions, we need to choose a proper step size, which is tuned by the parameters η and r.

5   Results for a simplified case

Our intuitions can be demonstrated in a simplified setting with (**A1**), (**A2'**), (**A3**), and (**N1**), where

   (**A2'**) x_i^*'s are independent, and $x_i^* = 1$ with probability s/n and 0 otherwise for a constant $s > 0$.

Furthermore, let $\mathbf{N}^{(0)} = 0$. This is a special case of our general assumptions, with $C_1 = c_2 = C_2 = s$ where s is the parameter in (**A2'**). It is still an interesting setting; as far as we know, there is no existing guarantee of alternating type algorithms for it.

To present our results, we let $(\mathbf{A}^*)^\dagger$ denote the matrix satisfying $(\mathbf{A}^*)^\dagger \mathbf{A}^* = \mathbf{I}$; if there are multiple such matrices we let it denote the one with minimum $\|(\mathbf{A}^*)^\dagger\|_\infty$.

Theorem 1 (Simplified case, adversarial noise). *There exists a* absolute constant $\mathcal{G}$ *such that when Assumption (A1)(A2')(A3) and (N1) are satisfied with* $l = 1/10$, $C_\nu \leq \frac{\mathcal{G}c}{\max\{m,n\|(\mathbf{A}^*)^\dagger\|_\infty\}}$ *for some* $0 \leq c \leq 1$, *and* $\mathbf{N}^{(0)} = 0$, *then there exist* α, η, r *such that for every* $0 < \epsilon, \delta < 1$ *and* $N = \mathrm{poly}(n, m, 1/\epsilon, 1/\delta)$ *the following holds with probability at least* $1 - \delta$.

After $T = O\left(\ln\frac{1}{\epsilon}\right)$ *iterations, Algorithm 1 outputs a solution* $\mathbf{A} = \mathbf{A}^*(\mathbf{\Sigma} + \mathbf{E}) + \mathbf{N}$ *where* $\mathbf{\Sigma} \succeq (1 - \ell)\mathbf{I}$ *is diagonal,* $\|\mathbf{E}\|_1 \leq \epsilon + c$ *is off-diagonal, and* $\|\mathbf{N}\|_1 \leq c$.

Remarks. Consequently, when $\|\mathbf{A}^*\|_1 = 1$, we can do normalization $\hat{\mathbf{A}}_i = \mathbf{A}_i/\|\mathbf{A}_i\|_1$, and the normalized output $\hat{\mathbf{A}}$ satisfies

$$\|\hat{\mathbf{A}} - \mathbf{A}^*\|_1 \leq \epsilon + 2c.$$

So under mild conditions and with proper parameters, our algorithm recovers the ground-truth in a geometric rate. It can achieve arbitrary small recovery error in the noiseless setting, and achieve error up to the noise limit even with adversarial noise whose level is comparable to the signal.

The condition on ℓ means that a constant warm start is sufficient for our algorithm to converge, which is much better than previous work such as [AR15]. Indeed, in that work, the ℓ needs to even depend on the dynamic range of the entries of $\mathbf{A}^*$ which is problematic in practice.

It is shown that with large adversarial noise, the algorithm can still recover the features up to the noise limit. When $m \geq n\|(\mathbf{A}^*)^\dagger\|_\infty$, each data point has adversarial noise with ℓ_1 norm as large as $\|\nu\|_1 = C_\nu m = \Omega(c)$, which is in the same order as the signal $\|\mathbf{A}^* x^*\|_1 = O(1)$. Our algorithm still works in this regime. Furthermore, the *final* error $\|\mathbf{A} - \mathbf{A}^*\|_1$ is $O(c)$, in the same order as the *adversarial* noise in *one* data point.

Note the appearance of $\|(\mathbf{A}^*)^\dagger\|_\infty$ is not surprising. The case when the columns are the canonical unit vectors for instance, which corresponds to $\|(\mathbf{A}^*)^\dagger\|_\infty = 1$, is expected to be easier than the case when the columns are nearly the same, which corresponds to large $\|(\mathbf{A}^*)^\dagger\|_\infty$.

A similar theorem holds for the unbiased noise model.

Theorem 2 (Simplified case, unbiased noise). *If Assumption (A1)(A2')(A3) and (N2) are satisfied with* $C_\nu = \frac{\mathcal{G}c\sqrt{n}}{\max\{m,n\|(\mathbf{A}^*)^\dagger\|_\infty\}}$ *and the other parameters set as in Theorem 1, then the same guarantee in holds.*

5

Remarks. With unbiased noise which is commonly assumed in many applications, the algorithm can tolerate noise level $\sqrt{n}$ larger than the adversarial case. When $m \geq n\|(\mathbf{A}^*)^\dagger\|_\infty$, each data point has adversarial noise with ℓ_1 norm as large as $\|\nu\|_1 = C_\nu m = \Omega(c\sqrt{n})$, which can be $\Omega(\sqrt{n})$ times larger than the signal $\|\mathbf{A}^*x^*\|_1 = O(1)$. The algorithm can recover the ground-truth in this heavy noise regime. Furthermore, the *final* error $\|\mathbf{A} - \mathbf{A}^*\|_1$ is $O\left(\|\nu\|_1/\sqrt{n}\right)$, which is only $O(1/\sqrt{n})$ fraction of the noise in *one* data point. This is very strong denoising effect and a bit counter-intuitive. It is possible since we exploit the average of the noise for cancellation, and also use thresholding to remove noise spread out in the coordinates.

5.1 Analysis: intuition

A natural approach typically employed to analyze algorithms for non-convex problems is to define a function on the intermediate solution $\mathbf{A}$ and the ground-truth $\mathbf{A}^*$ measuring their distance and then show that the function decreases at each step. However, a single potential function will not be enough in our case, as we argue below, so we introduce a novel framework of maintaining two potential functions which capture different aspects of the intermediate solutions.

Let us denote the intermediate solution and the update as (omitting the superscript (t))

$$\mathbf{A} = \mathbf{A}^*(\boldsymbol{\Sigma} + \mathbf{E}) + \mathbf{N}, \quad \hat{\mathbb{E}}[(y - y')(x - x')^\top] = \mathbf{A}^*(\widetilde{\boldsymbol{\Sigma}} + \widetilde{\mathbf{E}}) + \widetilde{\mathbf{N}},$$

where $\boldsymbol{\Sigma}$ and $\widetilde{\boldsymbol{\Sigma}}$ are diagonal, $\mathbf{E}$ and $\widetilde{\mathbf{E}}$ are off-diagonal, and $\mathbf{N}$ and $\widetilde{\mathbf{N}}$ are the terms outside the span of $\mathbf{A}^*$ which is caused by the noise. To cleanly illustrate the intuition behind ReLU and the coupled potential functions, we focus on the noiseless case and assume that we have infinite samples.

Since $\mathbf{A}_i = \boldsymbol{\Sigma}_{i,i}\mathbf{A}_i^* + \sum_{j\neq i}\mathbf{E}_{j,i}\mathbf{A}_j^*$, if the ratio between $\|\mathbf{E}_i\|_1 = \sum_{j\neq i}|\mathbf{E}_{j,i}|$ and $\boldsymbol{\Sigma}_{i,i}$ gets smaller, then the algorithm is making progress; if the ratio is large at the end, a normalization of $\mathbf{A}_i$ gives a good approximation of $\mathbf{A}_i^*$. So it suffices to show that $\boldsymbol{\Sigma}_{i,i}$ is always about a constant while $\|\mathbf{E}_i\|_1$ decreases at each iteration. We will focus on $\mathbf{E}$ and consider the update rule in more detail to argue this. After some calculation, we have

$$\mathbf{E} \leftarrow (1 - \eta)\mathbf{E} + r\eta\widetilde{\mathbf{E}}, \qquad\qquad \widetilde{\mathbf{E}} = \mathbb{E}[(x^* - (x')^*)(x - x')^\top], \qquad (3)$$

where x, x' are the decoding for $x^*, (x')^*$ respectively:

$$x = \phi_\alpha\left((\boldsymbol{\Sigma} + \mathbf{E})^{-1}x^*\right), \qquad\qquad x' = \phi_\alpha\left((\boldsymbol{\Sigma} + \mathbf{E})^{-1}(x')^*\right). \qquad (4)$$

To see why the ReLU function matters, consider the case when we do not use it.

$$\widetilde{\mathbf{E}} = \mathbb{E}(x^* - (x')^*)\left[\mathbf{A}^\dagger\mathbf{A}^*(x^* - (x')^*)\right]^\top = \mathbb{E}\left[(x^* - (x')^*)(x^* - (x')^*)^\top\right]\left[(\boldsymbol{\Sigma} + \mathbf{E})^{-1}\right]^\top$$
$$\propto \left[(\boldsymbol{\Sigma} + \mathbf{E})^{-1}\right]^\top \approx \boldsymbol{\Sigma}^{-1} - \boldsymbol{\Sigma}^{-1}\mathbf{E}\boldsymbol{\Sigma}^{-1}.$$

where we used Taylor expansion and the fact that $\mathbb{E}\left[(x^* - (x')^*)(x^* - (x')^*)^\top\right]$ is a scaling of identity. Hence, if we think of $\boldsymbol{\Sigma}$ as approximately $\mathbf{I}$ and take an appropriate r, the update to the matrix $\mathbf{E}$ is approximately $\mathbf{E} \leftarrow \mathbf{E} - \eta\mathbf{E}^\top$. Since we do not have control over the signs of $\mathbf{E}$ throughout the iterations, the problematic case is when the entries of $\mathbf{E}^\top$ and $\mathbf{E}$ roughly match in signs, which would lead to the entries of $\mathbf{E}$ increasing.

Now we consider the decoding to see why ReLU is important. Ignoring the higher order terms and regarding $\boldsymbol{\Sigma} = \mathbf{I}$, we have

$$x = \phi_\alpha\left((\boldsymbol{\Sigma} + \mathbf{E})^{-1}x^*\right) \approx \phi_\alpha\left(\boldsymbol{\Sigma}^{-1}x^* - \boldsymbol{\Sigma}^{-1}\mathbf{E}\boldsymbol{\Sigma}^{-1}x^*\right) \approx \phi_\alpha\left(x^* - \mathbf{E}x^*\right). \qquad (5)$$

The problematic term is $\mathbf{E}x^*$. These errors when summed up will be comparable or even larger than the signals, and the algorithm will fail. However, since the signals are non-negative and most coordinates with errors only have small values, thresholding with ReLU properly can remove those errors while keeping a large fraction of the signals. This leads to large $\widetilde{\boldsymbol{\Sigma}}_{i,i}$ and small $\widetilde{\mathbf{E}}_{j,i}$'s, and then we can choose an r such that $\mathbf{E}_{j,i}$'s keep decreasing while $\boldsymbol{\Sigma}_{i,i}$'s stay in a certain range.

To get a quantitative bound, we divide $\mathbf{E}$ into its positive part $\mathbf{E}_+$ and its negative part $\mathbf{E}_-$:

$$[\mathbf{E}_+]_{i,j} = \max\{\mathbf{E}_{i,j}, 0\}, \qquad\qquad [\mathbf{E}_-]_{i,j} = \max\{-\mathbf{E}_{i,j}, 0\}. \qquad (6)$$

6

The reason to do so is the following: when $\mathbf{E}_{i,j}$ is negative, by the Taylor expansion approximation, $\left[(\boldsymbol{\Sigma} + \mathbf{E})^{-1}x^*\right]_i$ will tend to be more positive and will not be thresholded most of the time. Therefore, $\mathbf{E}_{j,i}$ will turn more positive at next iteration. On the other hand, when $\mathbf{E}_{i,j}$ is positive, $\left[(\boldsymbol{\Sigma} + \mathbf{E})^{-1}x^*\right]_i$ will tend to be more negative and zeroed out by the threshold function. Therefore, $\mathbf{E}_{j,i}$ will *not* be more negative at next iteration. We will show for positive and negative parts of $\mathbf{E}$:

$$\text{postive}^{(t+1)} \leftarrow (1-\eta)\text{positive}^{(t)} + (\eta)\text{negative}^{(t)}, \text{negative}^{(t+1)} \leftarrow (1-\eta)\text{negative}^{(t)} + (\varepsilon\eta)\text{positive}^{(t)}$$

for a small $\varepsilon \ll 1$. Due to ϵ, we can couple the two parts so that a weighted average of them will decrease, which implies that $\|\mathbf{E}\|_s$ is small at the end. This leads to our coupled potential function.[2]

5.2   Analysis: proof sketch

Here we describe a proof sketch for the simplified case while the complete proof is presented in the appendix.

One iteration We focus on one update and omit the superscript (t). Recall the definitions of $\mathbf{E}$, $\boldsymbol{\Sigma}$ and $\mathbf{N}$ in (5.1), and $\widetilde{\mathbf{E}}$, $\widetilde{\boldsymbol{\Sigma}}$ and $\widetilde{\mathbf{N}}$ in (5.1). Our goal is to derive lower and upper bounds for $\widetilde{\mathbf{E}}$, $\widetilde{\boldsymbol{\Sigma}}$ and $\widetilde{\mathbf{N}}$, assuming that $\boldsymbol{\Sigma}_{i,i}$ falls into some range around 1, while $\mathbf{E}$ and $\mathbf{N}$ are small. This will allow doing induction on them.

First, begin with the decoding. Some calculation shows that, the decoding for $y = \mathbf{A}^* x^* + \nu$ is

$$x = \phi_\alpha \left(\mathbf{Z}x^* + \xi\right), \quad \text{where } \mathbf{Z} = (\boldsymbol{\Sigma} + \mathbf{E})^{-1}, \ \xi = -\mathbf{A}^\dagger \mathbf{N}\mathbf{Z}x^* + \mathbf{A}^\dagger \nu. \tag{7}$$

Now, we can present our key lemmas bounding $\widetilde{\mathbf{E}}$, $\widetilde{\boldsymbol{\Sigma}}$, and $\widetilde{\mathbf{N}}$.

Lemma 3 (Simplified bound on $\widetilde{\mathbf{E}}$, informal). *(1) if $\mathbf{Z}_{i,j} < 0$, then $\left|\widetilde{\mathbf{E}}_{j,i}\right| \leq O\left(\frac{1}{n^2}\left(|\mathbf{Z}_{i,j}| + c\right)\right)$,*

(2) if $\mathbf{Z}_{i,j} \geq 0$, then $-O\left(\frac{c}{n^2} + \frac{c}{n}|\mathbf{Z}_{i,j}| + \frac{1}{n^2}|\mathbf{Z}_{i,j}|\right) \leq \left|\widetilde{\mathbf{E}}_{j,i}\right| \leq O\left(\frac{1}{n}\|\mathbf{Z}_{i,j}\|\right)$.

Note that $\mathbf{Z} \approx \boldsymbol{\Sigma}^{-1} - \boldsymbol{\Sigma}^{-1}\mathbf{E}\boldsymbol{\Sigma}^{-1}$, so $\mathbf{Z}_{i,j} < 0$ corresponds roughly to $\mathbf{E}_{i,j} > 0$. In this case, the upper bound on $|\widetilde{\mathbf{E}}_{j,i}|$ is very small and thus $|\mathbf{E}_{j,i}|$ decreases, as described in the intuition. What is most interesting is the case when $\mathbf{Z}_{i,j} \geq 0$ (roughly $\mathbf{E}_{i,j} < 0$). The upper bound is much larger, corresponding to the intuition that negative $\mathbf{E}_{i,j}$ can contribute a large positive value to $\mathbf{E}_{j,i}$. Fortunately, the lower bounds are of much smaller absolute value, which allows us to show that a potential function that couples Case (1) and Case (2) in Lemma 3 actually decreases; see the induction below.

Lemma 4 (Simplified bound on $\widetilde{\boldsymbol{\Sigma}}$, informal). $\widetilde{\boldsymbol{\Sigma}}_{i,i} \geq \Omega(\boldsymbol{\Sigma}_{i,i}^{-1} - \alpha)/n$.

Lemma 5 (Simplified bound on $\widetilde{\mathbf{N}}$, adversarial noise, informal). $\left|\widetilde{\mathbf{N}}_{i,j}\right| \leq O(C_\nu/n)$.

Induction by iterations We now show how to use the three lemmas to prove the theorem for the adversarial noise, and that for the unbiased noise is similar.

Let $a_t := \left\|\mathbf{E}_+^{(t)}\right\|_s$ and $b_t := \left\|\mathbf{E}_-^{(t)}\right\|_s$, and choose $\eta = \ell/6$. We begin with proving the following three claims by induction on t: at the beginning of iteration t,

(1) $(1 - \ell)\mathbf{I} \preceq \boldsymbol{\Sigma}^{(t)}$

(2) $\left\|\mathbf{E}^{(t)}\right\|_s \leq 1/8$, and if $t > 0$, then $a_t + \beta b_t \leq \left(1 - \frac{1}{25}\eta\right)(a_{t-1} + \beta b_{t-1}) + \eta h$, for some $\beta \in (1, 8)$, and some small value h,

(3) $\left\|\mathbf{N}^{(t)}\right\|_s \leq c/10$.

The most interesting part is the second claim. At a high level, by Lemma 3, we can show that

$$a_{t+1} \leq \left(1 - \frac{3}{25}\eta\right)a_t + 7\eta b_t + \eta h, \qquad b_{t+1} \leq \left(1 - \frac{24}{25}\eta\right)b_t + \frac{1}{100}\eta a_t + \eta h.$$

---

[2]Note that since intuitively, $\mathbf{E}_{i,j}$ gets affected by $\mathbf{E}_{j,i}$ after an update, if we have a row which contains negative entries, it is possible that $\|\mathbf{A}_i - \mathbf{A}_i^*\|_1$ increases. So we cannot simply use $\max_i \|\mathbf{A}_i - \mathbf{A}_i^*\|_1$ as a potential function.

7

Notice that the contribution of b_t to a_{t+1} is quite large (due to the larger upper bound in Case (2) in Lemma 3), but the other terms are much nicer, such as the small contribution of a_t to b_{t+1}. This allows to choose a $\beta \in (1, 8)$ so that $a_{t+1} + \beta b_{t+1}$ leads to the desired recurrence in the second claim. In other words, $a_{t+1} + \beta b_{t+1}$ is our potential function which decreases at each iteration up to the level h. The other claims can also be proved by the corresponding lemmas. Then the theorem follows from the induction claims.

6 More general results

More general weight distributions. Our argument holds under more general assumptions on x^*.

Theorem 6 (Adversarial noise). *There exists an absolute constant $\mathcal{G}$ such that when Assumption (A0)-(A3) and (N1) are satisfied with $l = 1/10$, $C_2 \leq 2c_2$, $C_1^3 \leq \mathcal{G}c_2^2 n$, $C_\nu \leq \left\{ \frac{c_2^2 \mathcal{G} c}{C_1^2 m}, \frac{c_2^4 \mathcal{G} c}{C_1^5 n \left\| (\mathbf{A}^*)^\dagger \right\|_\infty} \right\}$ for $0 \leq c \leq 1$, and $\left\| \mathbf{N}^{(0)} \right\|_\infty \leq \frac{c_2^2 \mathcal{G} c}{C_1^3 \left\| (\mathbf{A}^*)^\dagger \right\|_\infty}$, then there exist α, η, r such that for every $0 < \epsilon, \delta < 1$ and $N = \mathrm{poly}(n, m, 1/\epsilon, 1/\delta)$, with probability at least $1 - \delta$ the following holds.*

After $T = O\left(\ln \frac{1}{\epsilon} \right)$ iterations, Algorithm 1 outputs a solution $\mathbf{A} = \mathbf{A}^*(\mathbf{\Sigma} + \mathbf{E}) + \mathbf{N}$ where $\mathbf{\Sigma} \succeq (1 - \ell)\mathbf{I}$ is diagonal, $\|\mathbf{E}\|_1 \leq \epsilon + c/2$ is off-diagonal, and $\|\mathbf{N}\|_1 \leq c/2$.

Theorem 7 (Unbiased noise). *If Assumption (A0)-(A3) and (N2) are satisfied with $C_\nu = \frac{c_2 \mathcal{G} \sqrt{cn}}{C_1 \max\{m, n \| (\mathbf{A}^*)^\dagger \|_\infty\}}$ and the other parameters set as in Theorem 6, then the same guarantee holds.*

The conditions on C_1, c_2, C_2 intuitively mean that each feature needs to appear with reasonable probability. $C_2 \leq 2c_2$ means that their proportions are reasonably balanced. This may be a mild restriction for some applications, and additionally we propose a pre-processing step that can relax this in the next subsection. The conditions allow a rather general family of distributions, so we point out an important special case to provide a more concrete sense of the parameters. For example, for the uniform independent distribution considered in the simplified case, we can actually allow s to be much larger than a constant; our algorithm just requires $s \leq \mathcal{G}n$ for a fixed constant $\mathcal{G}$. So it works for uniform sparse distributions even when the sparsity is linear, which is an order of magnitude larger than in the dictionary learning regime. Furthermore, the distributions of x_i^* can be very different, since we only require $C_1^3 = O(c_2^2 n)$. Moreover, all these can be handled without specific structural assumptions on $\mathbf{A}^*$.

More general proportions. A mild restriction in Theorem 6 and 7 is that $C_2 \leq 2c_2$, that is, $\max_{i \in [n]} \mathbb{E}[(x_i^*)^2] \leq 2 \min_{i \in [n]} \mathbb{E}[(x_i^*)^2]$. To satisfy this, we propose a preprocessing algorithm for balancing $\mathbb{E}[(x_i^*)^2]$. The idea is quite simple: instead of solving $\mathbf{Y} \approx \mathbf{A}^*\mathbf{X}$, we could also solve $\mathbf{Y} \approx [\mathbf{A}^*\mathbf{D}][(\mathbf{D})^{-1}\mathbf{X}]$ for a positive diagonal matrix $\mathbf{D}$, where $\mathbb{E}[(x_i^*)^2]/\mathbf{D}_{i,i}^2$ is with in a factor of 2 from each other. We show in the appendix that this can be done under assumptions as the above theorems, and additionally $\mathbf{\Sigma} \preceq (1 + \ell)\mathbf{I}$ and $\mathbf{E}^{(0)} \geq$ entry-wise. After balancing, one can use Algorithm 1 on the new ground-truth matrix $[\mathbf{A}^*\mathbf{D}]$ to get the final result.

7 Conclusion

A natural algorithm with alternating updates is proposed for non-negative matrix factorization and theoretical guarantees are provided. The algorithm provably recovers a feature matrix close to the ground-truth and is robust to noise. Our analysis provides insights on the effect of the ReLU units in the presence of the non-negativity constraints, and the resulting interesting dynamics of the convergence.

Acknowledgements

This work was supported in part by NSF grants CCF-1527371, DMS-1317308, Simons Investigator Award, Simons Collaboration Grant, and ONR-N00014-16-1-2329.

8

References

[AGH+13] S. Arora, R. Ge, Y. Halpern, D. Mimno, A. Moitra, D. Sontag, Y. Wu, and M. Zhu. A practical algorithm for topic modeling with provable guarantees. In *ICML*, 2013.

[AGKM12] Sanjeev Arora, Rong Ge, Ravindran Kannan, and Ankur Moitra. Computing a nonnegative matrix factorization–provably. In *STOC*, pages 145–162. ACM, 2012.

[AGM12] S. Arora, R. Ge, and A. Moitra. Learning topic models – going beyond svd. In *FOCS*, 2012.

[AGMM15] S. Arora, R. Ge, T. Ma, and A. Moitra. Simple, efficient, and neural algorithms for sparse coding. In *COLT*, 2015.

[AGMS12] Sanjeev Arora, Rong Ge, Ankur Moitra, and Sushant Sachdeva. Provable ica with unknown gaussian noise, with implications for gaussian mixtures and autoencoders. In *NIPS*, pages 2375–2383, 2012.

[AHJK13] A. Anandkumar, D. Hsu, A. Javanmard, and S. Kakade. Learning latent bayesian networks and topic models under expansion constraints. In *ICML*, 2013.

[AKF+12] A. Anandkumar, S. Kakade, D. Foster, Y. Liu, and D. Hsu. Two svds suffice: Spectral decompositions for probabilistic topic modeling and latent dirichlet allocation. Technical report, 2012.

[AR15] Pranjal Awasthi and Andrej Risteski. On some provably correct cases of variational inference for topic models. In *NIPS*, pages 2089–2097, 2015.

[BBK14] T. Bansal, C. Bhattacharyya, and R. Kannan. A provable svd-based algorithm for learning topics in dominant admixture corpus. In *NIPS*, 2014.

[BGKP16] Chiranjib Bhattacharyya, Navin Goyal, Ravindran Kannan, and Jagdeep Pani. Nonnegative matrix factorization under heavy noise. In *Proceedings of the 33nd International Conference on Machine Learning*, 2016.

[Ble12] David M Blei. Probabilistic topic models. *Communications of the ACM*, 2012.

[BNJ03] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *JMLR*, 3:993–1022, 2003.

[Com94] Pierre Comon. Independent component analysis, a new concept? *Signal processing*, 36(3):287–314, 1994.

[DRIS13] W. Ding, M.H. Rohban, P. Ishwar, and V. Saligrama. Topic discovery through data dependent and random projections. *arXiv preprint arXiv:1303.3664*, 2013.

[DRIS14] W. Ding, M.H. Rohban, P. Ishwar, and V. Saligrama. Efficient distributed topic modeling with provable guarantees. In *AISTAT*, pages 167–175, 2014.

[FJK96] Alan Frieze, Mark Jerrum, and Ravi Kannan. Learning linear transformations. In *focs*, page 359. IEEE, 1996.

[HPC14] Tao Hu, Cengiz Pehlevan, and Dmitri B Chklovskii. A hebbian/anti-hebbian network for online sparse dictionary learning derived from symmetric matrix factorization. In *Asilomar Conference on Signals, Systems and Computers*, pages 613–619. IEEE, 2014.

[lda16] Lda-c software. `https://github.com/blei-lab/lda-c/blob/master/readme.txt`, 2016. Accessed: 2016-05-19.

[LS97] Daniel D Lee and H Sebastian Seung. Unsupervised learning by convex and conic coding. *NIPS*, pages 515–521, 1997.

[LS99] Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

9

[LS01] Daniel D Lee and H Sebastian Seung. Algorithms for non-negative matrix factorization. In *NIPS*, pages 556–562, 2001.

[OF97] 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.

[PC14] Cengiz Pehlevan and Dmitri B Chklovskii. A hebbian/anti-hebbian network derived from online non-negative matrix factorization can cluster and discover sparse features. In *Asilomar Conference on Signals, Systems and Computers*, pages 769–775. IEEE, 2014.

[PC15a] Cengiz Pehlevan and Dmitri Chklovskii. A normative theory of adaptive dimensionality reduction in neural networks. In *NIPS*, pages 2260–2268, 2015.

[PC15b] Cengiz Pehlevan and Dmitri B Chklovskii. Optimization theory of hebbian/anti-hebbian networks for pca and whitening. *arXiv preprint arXiv:1511.09468*, 2015.

[ZX12] Jun Zhu and Eric P Xing. Sparse topical coding. *arXiv preprint arXiv:1202.3778*, 2012.

10

A Preliminary

Given a matrix $\mathbf{Y} \in \mathbb{R}^{m \times N}$, the goal of non-negative matrix factorization (NMF) is to find a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and a non-negative matrix $\mathbf{X} \in \mathbb{R}^{n \times N}$, so that $\mathbf{Y} \approx \mathbf{A}\mathbf{X}$. The columns of $\mathbf{Y}$ are called data points, those of $\mathbf{A}$ are features, and those of $\mathbf{X}$ are weights.

The notation $[\mathbf{M}]_j$ denotes the j-th column of $\mathbf{M}$, $[\mathbf{M}]^i$ denotes the i-th row of $\mathbf{M}$, and $\mathbf{M}_{i,j}$ denotes the element of $\mathbf{M}$ at the i-th row and j-th column. Furthermore, let $\mathbf{M}_+ =$ denote the positive part of the matrix, and let $\mathbf{M}_-$ denote the absolute value of the negative part of the matrix:

$$[\mathbf{M}_+]_{i,j} = \begin{cases} \mathbf{M}_{i,j} & \text{if } \mathbf{M}_{i,j} \geq 0, \\ 0 & \text{if } \mathbf{M}_{i,j} < 0, \end{cases}$$

$$[\mathbf{M}_-]_{i,j} = \begin{cases} 0 & \text{if } \mathbf{M}_{i,j} \geq 0, \\ |\mathbf{M}_{i,j}| & \text{if } \mathbf{M}_{i,j} < 0. \end{cases}$$

For analysis, the following norms of the matrices are needed.

Definition (l_1 norm of a matrix (induced column norm))**.** The (induced) l_1 norm of a matrix $\mathbf{E} \in \mathbb{R}^{n \times n}$ is

$$\|\mathbf{E}\|_1 = \max_{i \in [n]} \left\{ \sum_{j=1}^n |\mathbf{E}_{j,i}| \right\}.$$

Definition (l_∞ norm of a matrix (induced row norm))**.** The (induced) l_∞ norm of a matrix $\mathbf{E} \in \mathbb{R}^{n \times n}$ is

$$\|\mathbf{E}\|_\infty = \max_{i \in [n]} \left\{ \sum_{j=1}^n |\mathbf{E}_{i,j}| \right\}.$$

These two norms are related, and they enjoy the sub-multipicity property of the induced norm.

Property 8 (dual norm)**.** *For a matrix* $\mathbf{E} \in \mathbb{R}^{n \times n}$,

$$\|\mathbf{E}\|_1 = \|\mathbf{E}^\top\|_\infty.$$

Note that unlike l_2 norm, it is possible that $\|\mathbf{E}\|_1 \neq \|\mathbf{E}^\top\|_1$ or $\|\mathbf{E}\|_\infty \neq \|\mathbf{E}^\top\|_\infty$.

Property 9 (induced norm of a matrix)**.** *Let* $\mathbf{E}_1, \mathbf{E}_2 \in \mathbb{R}^{n \times n}$ *be two matrices, then*

$$\|\mathbf{E}_1 \mathbf{E}_2\|_1 \leq \|\mathbf{E}_1\|_1 \|\mathbf{E}_2\|_1,$$
$$\|\mathbf{E}_1 \mathbf{E}_2\|_\infty \leq \|\mathbf{E}_1\|_\infty \|\mathbf{E}_2\|_\infty.$$

The following two kinds of norms are also useful for the analysis.

Definition (symmetrized norm of a matrix)**.** The symmetrized norm of a matrix $\mathbf{E} \in \mathbb{R}^{n \times n}$ is

$$\|\mathbf{E}\|_s = \max(\|\mathbf{E}\|_1, \|\mathbf{E}\|_\infty).$$

Note that $\|\mathbf{E}\|_s$ is a norm since it's the maximum of two norms.

Definition (max norm)**.** The max norm of a matrix $\mathbf{E} \in \mathbb{R}^{m \times n}$ is

$$\|\mathbf{E}\|_{\max} = \max_{i,j} |\mathbf{E}_{i,j}|.$$

For the function ϕ_α used in our decoding algorithm, we frequently use the following properties in the analysis.

Property 10 (ReLU)**.** $\phi_\alpha(z) = \max(0, z - \alpha)$ *is non-decreasing. It is 1-Lipschitz, i.e.,*

$$|\phi_\alpha(z_1) - \phi_\alpha(z_2)| \leq |z_1 - z_2|. \tag{8}$$

It satisfies

$$\phi_\alpha(z) \geq z - \alpha, \tag{9}$$
$$\phi_\alpha(z) \leq |z - \alpha|. \tag{10}$$

Furthermore, if $\alpha > 0$,

$$\phi_\alpha(z) \leq |z|. \tag{11}$$

11

B Proofs for main algorithm: Purification

Since NMF is NP-hard in the worst case, some assumptions are needed to make it tractable. In this paper, we consider a generative model for the data point $y = \mathbf{A}^* x^* + \nu$, where $\mathbf{A}^*$ is the ground-truth feature matrix, x^* is the ground-truth non-negative weight from some unknown distribution, and ν is the noise. Our focus is to recover $\mathbf{A}^*$ given access to the data distribution, assuming the following hold for parameters $C_1, c_2, C_2, \ell, C_\nu$ that will be determined in our theorems.

(A1) The columns of $\mathbf{A}^*$ are linearly independent.

(A2) For all $i \in [n]$, $x_i^* \in [0, 1]$, $\mathbb{E}[x_i^*] \leq \frac{C_1}{n}$ and $\frac{c_2}{n} \leq \mathbb{E}[(x_i^*)^2] \leq \frac{C_2}{n}$, and x_i^*'s are independent.

(A3) The initialization $\mathbf{A}^{(0)} = \mathbf{A}^*(\mathbf{\Sigma}^{(0)} + \mathbf{E}^{(0)}) + \mathbf{N}^{(0)}$, where $\mathbf{\Sigma}^{(0)}$ is diagonal, $\mathbf{E}^{(0)}$ is off-diagonal, and
$$\mathbf{\Sigma}^{(0)} \succeq (1 - \ell)\mathbf{I}, \quad \left\|\mathbf{E}^{(0)}\right\|_s \leq \ell.$$

We consider two noise models.

(N1) Adversarial noise: only assume that $\max_i |\nu_i| \leq C_\nu$ almost surely.

(N2) Unbiased noise: $\max_i |\nu_i| \leq C_\nu$ almost surely, and $\mathbb{E}[\nu|x^*] = 0$.

Algorithm 1 Purification

Input: initialization $\mathbf{A}^{(0)}$, threshold α, step size η, scaling factor r, sample size N, iterations T
1: **for** $t = 0, 1, 2, ..., T - 1$ **do**
2:     Draw examples $y_1, \ldots, y_N$.
3:     (Decode) Compute $\mathbf{A}^\dagger$, the pseudo-inverse of $\mathbf{A}^{(t)}$ with minimum $\|(\mathbf{A})^\dagger\|_\infty$.
              Set $x = \phi_\alpha(\mathbf{A}^\dagger y)$ for each example y.   *// ϕ_α is ReLU activation; see (2) for the definition*
4:     (Update) Update the feature matrix
              $\mathbf{A}^{(t+1)} = (1 - \eta)\,\mathbf{A}^{(t)} + r\eta\hat{\mathbb{E}}\left[(y - y')(x - x')^\top\right]$
    where $\hat{\mathbb{E}}$ is over independent uniform y, y' from $\{y_1, \ldots, y_N\}$, and x, x' are their decodings.
Output: $\mathbf{A} = \mathbf{A}^{(T)}$

Our main algorithm is presented in Algorithm 1. It keeps a working feature matrix and operates in iterations. In each iteration, it first compute the weights for N examples (decoding), and then use the computed weights to update the feature matrix (updating).

The decoding is simply multiplying the example by the pseudo-inverse of the current feature matrix and then passing it through a one-sided threshold function ϕ_α. The pseudo-inverse with minimum infinity norm is used so as to maximize the robustness to noise (see the theorems). The one-sided threshold function operates element-wisely on the input vector v, and for an element v_i, it is defined as
$$\phi_\alpha(v_i) = \max\{v_i - \alpha, 0\}.$$
This is just the rectified linear unit (ReLU) with offset α. To get some sense about the decoding, suppose the current feature matrix is the ground-truth. Then $\mathbf{A}^\dagger y = \mathbf{A}^\dagger \mathbf{A}^* x^* + \mathbf{A}^\dagger \nu = x^* + \mathbf{A}^\dagger \nu$. So we would like to use a small $\mathbf{A}^\dagger$ and use threshold to remove the noise term.

In the encoding step, the algorithm move the feature matrix along the direction $\mathbb{E}\left[(y - y')(x - x')^\top\right]$. Suppose we have independent x_i^*'s, perfect decoding and no noise, then $\mathbb{E}\left[(y - y')(x - x')^\top\right] = \mathbf{A}^*$, and thus it is moving towards the ground-truth. Without those ideal conditions, we need to choose a proper step size, which is tune by the parameters η and r.

At the end, the algorithm simply outputs the scaled features with unit norm. The output enjoys the following guarantee in the adversarial noise model.

B.1 Analysis of one update step

In this subsection, we focus on one update step, bounding the changes of $\mathbf{\Sigma}, \mathbf{E}, \mathbf{N}$ and some auxiliary variables, and then in the next subsection we put things together to prove the theorem. So through

12

out this subsection we will focus on a particular iteration t and omit the superscript (t), while in the next subsection we will put back the superscript.

For analysis, denote $\mathbf{A}^{(t)}$ as

$$\mathbf{A} = \mathbf{A}^*(\boldsymbol{\Sigma} + \mathbf{E}) + \mathbf{N}$$

where $\boldsymbol{\Sigma}$ is a diagonal matrix, $\mathbf{E}$ is an off-diagonal matrix, and $\mathbf{N}$ is the component of $\mathbf{A}$ that lies outside the span of $\mathbf{A}^*$ (e.g., the noise caused by the noise in the sample).

Recall the following notations:

$$\mathbf{Z} = (\boldsymbol{\Sigma} + \mathbf{E})^{-1},$$

$$\mathbf{V} = \mathbf{Z} - \boldsymbol{\Sigma}^{-1} = \boldsymbol{\Sigma}^{-1} \sum_{k=1}^{\infty} (-\mathbf{E}\boldsymbol{\Sigma}^{-1})^k,$$

$$\xi = -\mathbf{A}^\dagger \mathbf{N} \mathbf{Z} x^* + \mathbf{A}^\dagger \nu.$$

Consider the update term $\hat{\mathbb{E}}\left[(y - y')(x - x')^\top\right]$ and denote it as

$$\boldsymbol{\Delta} = \hat{\mathbb{E}}\left[(y - y')(x - x')^\top\right] = \mathbf{A}^*(\widetilde{\boldsymbol{\Sigma}} + \widetilde{\mathbf{E}}) + \widetilde{\mathbf{N}}$$

where $\widetilde{\boldsymbol{\Sigma}}$ is a diagonal matrix, $\widetilde{\mathbf{E}}$ is an off-diagonal matrix, and $\mathbf{N}$ is the component of $\boldsymbol{\Delta}$ that lies outside the span of $\mathbf{A}^*$.

Since we now use empirical average, we will have sampling noise. Denote it as

$$\mathbf{N}_s = \hat{\mathbb{E}}[(y - y')(x - x')^\top] - \mathbb{E}[(y - y')(x - x')^\top].$$

Then by definition, for $y = \mathbf{A}^* x^* + \nu$ and $y' = \mathbf{A}^*(x')^* + \nu'$, we have

$$\hat{\mathbb{E}}[(y - y')(x - x')^\top] = \mathbb{E}[(y - y')(x - x')^\top] + \mathbf{N}_s$$
$$= \mathbf{A}^* \underbrace{\mathbb{E}\left[(x^* - (x')^*)(x - x')^\top\right]}_{\widetilde{\boldsymbol{\Sigma}} + \widetilde{\mathbf{E}}} + \underbrace{\mathbb{E}\left[(\nu - \nu')(x - x')^\top\right] + \mathbf{N}_s}_{\widetilde{\mathbf{N}}}.$$

Our goal is then bounding $\widetilde{\boldsymbol{\Sigma}}, \widetilde{\mathbf{E}}, \widetilde{\mathbf{N}}$ in terms of $\boldsymbol{\Sigma}, \mathbf{E}, \mathbf{N}$. Before doing so, we present a lemma for the decoding.

Lemma 11 (Main: Decoding). *Let $m \geq n$ be two positive integers. Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ be a matrix such that $\mathbf{A} = \mathbf{A}^*(\boldsymbol{\Sigma} + \mathbf{E}) + \mathbf{N}$ where $\mathbf{A}^*$ is full rank, $\boldsymbol{\Sigma}$ is a diagonal matrix such that $\boldsymbol{\Sigma} \succeq \frac{1}{2}\mathbf{I}$ and $\|\mathbf{E}\|_1 < \frac{1}{2}$. Then for $y = \mathbf{A}^* x^* + \nu$, the decoding is*

$$x = \phi_\alpha \left(\mathbf{Z} x^* + \xi\right)$$
$$= \phi_\alpha \left(\left(\boldsymbol{\Sigma}^{-1} + \mathbf{V}\right) x^* + \xi\right).$$

Proof of Lemma 11. Since $\mathbf{A} = \mathbf{A}^*(\boldsymbol{\Sigma} + \mathbf{E}) + \mathbf{N}$, we have

$$\mathbf{A}^* = (\mathbf{A} - \mathbf{N})(\boldsymbol{\Sigma} + \mathbf{E})^{-1}$$
$$y = (\mathbf{A} - \mathbf{N})(\boldsymbol{\Sigma} + \mathbf{E})^{-1} x^* + \nu.$$

Plugging into the decoding we get the first statement.

Observing that $\boldsymbol{\Sigma} + \mathbf{E} = (\mathbf{I} + \mathbf{E}\boldsymbol{\Sigma}^{-1})\boldsymbol{\Sigma}$ and $\|\mathbf{E}\boldsymbol{\Sigma}^{-1}\|_1 \leq \|\boldsymbol{\Sigma}^{-1}\|_1 \|\mathbf{E}\|_1 \leq 2\|\mathbf{E}\|_1 < 1$, we have $(\boldsymbol{\Sigma} + \mathbf{E})^{-1} = (\boldsymbol{\Sigma}^{-1} + \mathbf{V})$, resulting in the second statement. $\square$

Lemma 12 (Main: Bound on $\widetilde{\boldsymbol{\Sigma}}$). *Suppose $|\xi_i| \leq \rho < \alpha$ for any example and every $i \in [n]$, and suppose $\boldsymbol{\Sigma} \succeq \frac{1}{2}\mathbf{I}$. Then for any $i \in [n]$,*

$$\widetilde{\boldsymbol{\Sigma}}_{i,i} \geq \mathbb{E}\left[(x_i^*)^2\right]\left(2\boldsymbol{\Sigma}_{i,i}^{-1} - 2|\mathbf{V}_{i,i}|\right) - \frac{2C_1}{n}\left(\alpha + 2\rho + \frac{C_1}{n}\boldsymbol{\Sigma}_{i,i}^{-1} + \frac{2C_1}{n}\left\|[\mathbf{V}]^i\right\|_1\right),$$

$$\widetilde{\boldsymbol{\Sigma}}_{i,i} \leq \mathbb{E}\left[(x_i^*)^2\right]\left(2\boldsymbol{\Sigma}_{i,i}^{-1} + 2|\mathbf{V}_{i,i}|\right) + \frac{2C_1}{n}\left(\rho + \frac{C_1}{n}\left\|[\mathbf{V}]^i\right\|_1\right).$$

13

Proof of Lemma 12. According to the definition, we have
$$\widetilde{\boldsymbol{\Sigma}}_{i,i} = \left[(\mathbf{A}^*)^\dagger \mathbb{E}[(y - y')(x - x')^\top]\right]_{i,i}$$
$$= \mathbb{E}\left[(x_i^* - (x_i')^*)(x_i - x_i')\right]$$
$$= \mathbb{E}\left[(x_i^* - (x_i')^*)x_i\right] + \mathbb{E}\left[((x_i')^* - x_i^*)x_i'\right].$$
Since $(x_i^* - (x_i')^*)x_i$ and $((x_i')^* - x_i^*)x_i'$ has the same distribution, and $(x')^*, x^*$ are i.i.d. , we have
$$\widetilde{\boldsymbol{\Sigma}}_{i,i} = 2\mathbb{E}\left[(x_i^* - (x_i')^*)x_i\right]$$
$$= 2\mathbb{E}[x_i^* x_i] - 2\mathbb{E}[x_i^*]\mathbb{E}[x_i].$$
So it suffices to bound $\mathbb{E}[x_i^* x_i]$ and $\mathbb{E}[x_i]$. To do so, we first take a look at x_i. By the decoding rule,
$$x_i = \left[\phi_\alpha\left(\left(\boldsymbol{\Sigma}^{-1} + \mathbf{V}\right)x^* + \xi\right)\right]_i.$$
Since ϕ_α is 1-Lipschitz, denoting $\Delta = |[\mathbf{V}x^*]_i + \xi_i|$ we have
$$\left[\phi_\alpha\left(\boldsymbol{\Sigma}^{-1}x^*\right)\right]_i - \Delta \le x_i \le \left[\phi_\alpha\left(\boldsymbol{\Sigma}^{-1}x^*\right)\right]_i + \Delta. \tag{12}$$
For $\left[\phi_\alpha\left(\boldsymbol{\Sigma}^{-1}x^*\right)\right]_i$, by the Property 10 of $\phi_\alpha(z)$,
$$\boldsymbol{\Sigma}_{i,i}^{-1}x_i^* - \alpha \le \left[\phi_\alpha\left(\boldsymbol{\Sigma}^{-1}x^*\right)\right]_i = \phi_\alpha\left(\boldsymbol{\Sigma}_{i,i}^{-1}x_i^*\right) \le \boldsymbol{\Sigma}_{i,i}^{-1}x_i^*. \tag{13}$$
For $\Delta = |[\mathbf{V}x^*]_i + \xi_i|$,

$$\mathbb{E}[\Delta] \le \mathbb{E}\left[\left|\sum_j \mathbf{V}_{i,j}x_j^*\right|\right] + \mathbb{E}\left[|\xi_i|\right]$$

$$\le \mathbb{E}\left[\sum_j |\mathbf{V}_{i,j}|\, x_j^*\right] + \rho$$

$$= \sum_j |\mathbf{V}_{i,j}|\, \mathbb{E}\left[x_j^*\right] + \rho$$

$$\le \frac{C_1}{n}\left\|[\mathbf{V}]^i\right\|_1 + \rho \tag{14}$$

where the second step follows from the assumption $|\xi_i| \le \rho$, and the last step follows from Assumption **(A2)**.

Bounding $\mathbb{E}[x_i]$. By (12),(13), and (14), we have
$$\mathbb{E}[x_i] \le \mathbb{E}[\boldsymbol{\Sigma}_{i,i}^{-1}x^*] + \mathbb{E}[\Delta] \le \frac{C_1}{n}\boldsymbol{\Sigma}_{i,i}^{-1} + \frac{C_1}{n}\left\|[\mathbf{V}]^i\right\|_1 + \rho.$$

Bounding $\mathbb{E}[x_i^* x_i]$. First, note that

$$\mathbb{E}[x_i^* \Delta] \le \mathbb{E}\left[x_i^* \left|\sum_j \mathbf{V}_{i,j}x_j^*\right|\right] + \mathbb{E}\left[x_i^* |\xi_i|\right]$$

$$\le \mathbb{E}\left[x_i^* \sum_j x_j^* |\mathbf{V}_{i,j}|\right] + \frac{\rho C_1}{n}$$

$$= \sum_j \mathbb{E}\left[x_i^* x_j^*\right] |\mathbf{V}_{i,j}| + \frac{\rho C_1}{n}$$

$$= \mathbb{E}\left[(x_i^*)^2\right] |\mathbf{V}_{i,i}| + \sum_{j:j\neq i} \mathbb{E}\left[x_i^* x_j^*\right] |\mathbf{V}_{i,j}| + \frac{\rho C_1}{n}$$

$$\le \mathbb{E}\left[(x_i^*)^2\right] |\mathbf{V}_{i,i}| + \frac{C_1^2}{n^2}\sum_{j:j\neq i} |\mathbf{V}_{i,j}| + \frac{\rho C_1}{n}$$

$$\le \mathbb{E}\left[(x_i^*)^2\right] |\mathbf{V}_{i,i}| + \frac{C_1^2}{n^2}\left\|[\mathbf{V}]^i\right\|_1 + + \frac{\rho C_1}{n},$$

14

where the second and the fifth steps follow from Assumption **(A2)**. Therefore,

$$\mathbb{E}[x_i^* x_i] \geq \mathbb{E}\left[x_i^*\left(\Sigma_{i,i}^{-1}x_i^* - \alpha - \Delta\right)\right] \tag{15}$$

$$\geq \Sigma_{i,i}^{-1}\mathbb{E}\left[(x_i^*)^2\right] - \frac{(\alpha+\rho)C_1}{n} - \mathbb{E}\left[(x_i^*)^2\right]|\mathbf{V}_{i,i}| - \frac{C_1^2}{n^2}\left\|[\mathbf{V}]^i\right\|_1. \tag{16}$$

Putting together. For the first statement,

$$\widetilde{\Sigma}_{i,i} = 2\mathbb{E}[x_i^* x_i] - 2\mathbb{E}[x_i^*]\mathbb{E}[x_i]$$

$$\geq 2\Sigma_{i,i}^{-1}\mathbb{E}\left[(x_i^*)^2\right] - 2\frac{(\alpha+\rho)C_1}{n} - 2\mathbb{E}\left[(x_i^*)^2\right]|\mathbf{V}_{i,i}| - 2\frac{C_1^2}{n^2}\left\|[\mathbf{V}]^i\right\|_1$$

$$- 2\frac{C_1^2}{n^2}\Sigma_{i,i}^{-1} - 2\frac{C_1^2}{n^2}\left\|[\mathbf{V}]^i\right\|_1 - 2\frac{\rho C_1}{n}$$

$$\geq \mathbb{E}\left[(x_i^*)^2\right]\left(2\Sigma_{i,i}^{-1} - 2|\mathbf{V}_{i,i}|\right) - \frac{2C_1}{n}\left(\alpha + 2\rho + \frac{C_1}{n}\Sigma_{i,i}^{-1} + \frac{2C_1}{n}\left\|[\mathbf{V}]^i\right\|_1\right).$$

The second statement follows from

$$\widetilde{\Sigma}_{i,i} \leq 2\mathbb{E}[x_i^* x_i] \leq 2\mathbb{E}[x_i^*(\Sigma_{i,i}^{-1}x_i^* + \Delta)]$$

and the bound on $\mathbb{E}[x_i^*\Delta]$. $\qquad\qquad\square$

Lemma 13 (Main: Bound on $\widetilde{\mathbf{E}}$)**.** *Suppose* $|\xi_i| \leq \rho < \alpha$ *for any example and every* $i \in [n]$. *Then for all* $i, j \in [n]$ *such that* $i \neq j$, *the following holds.*
(1) If $\mathbf{Z}_{i,j} < 0$, *then*

$$\left|\widetilde{\mathbf{E}}_{j,i}\right| \leq \frac{4C_1^2\|\mathbf{Z}^i\|_1}{n^2(\alpha-\rho)}\left(|\mathbf{Z}_{i,j}| + \rho\right).$$

(2) If $\mathbf{Z}_{i,j} \geq 0$, *then*

$$-\frac{8C_1\rho}{n(\alpha-\rho)}\left(\frac{C_1\|\mathbf{Z}^i\|_1}{n} + \mathbf{Z}_{i,j}\right) - \frac{2C_1^2}{n^2}\mathbf{Z}_{i,j} \leq \widetilde{\mathbf{E}}_{j,i} \leq \frac{8C_1\rho}{n(\alpha-\rho)}\left(\frac{C_1\|\mathbf{Z}^i\|_1}{n} + \mathbf{Z}_{i,j}\right) + 2\mathbb{E}[(x_j^*)^2]\mathbf{Z}_{i,j}.$$

Proof of Lemma 13. Since $i \neq j$, we know that

$$\widetilde{\mathbf{E}}_{j,i} = \mathbb{E}\left[(x_j^* - (x_j')^*)(x_i - x_i')\right]$$

$$= \mathbb{E}\left[x_j^*(x_i - x_i')\right] + \mathbb{E}\left[(x_j')^*(x_i' - x_i)\right]$$

$$= 2\mathbb{E}\left[x_j^*(x_i - x_i')\right]$$

where the last equality follows from that $x_j^*(x_i - x_i')$ and $(x_j')^*(x_i' - x_i)$ has the same distribution. This quantity can be bounded by a coupling between x_i and x_i'. Define a new variable $\tilde{x}^*$ as

$$[\tilde{x}^*]_i = \begin{cases} x_i^*, & \text{if } i \neq j, \\ (x_j')^*, & \text{if } i = j. \end{cases}$$

By Assumption **(A2)**, conditional on x_j^*, $\tilde{x}^*$ has the same distribution as $(x')^*$. Therefore, consider the variable $\tilde{x}$ given by $\tilde{x} = \phi_\alpha(\mathbf{A}^\dagger(\mathbf{A}^*\tilde{x}^* + \nu'))$, we then have

$$\mathbb{E}\left[x_j^*(x_i - x_i')\right] = \mathbb{E}\left[x_j^*(x_i - \tilde{x}_i)\right].$$

In summary, we have

$$\widetilde{\mathbf{E}}_{j,i} = 2\mathbb{E}[x_j^*(x_i - \tilde{x}_i)]$$

where

$$x_i = [\phi_\alpha\left(\mathbf{Z}x^* + \xi\right)]_i, \quad \xi = -\mathbf{A}^\dagger\mathbf{N}\mathbf{Z}x^* + \mathbf{A}^\dagger\nu,$$

$$\tilde{x}_i = \left[\phi_\alpha\left(\mathbf{Z}x^* + \tilde{\xi}\right)\right]_i, \quad \tilde{\xi} = -\mathbf{A}^\dagger\mathbf{N}\mathbf{Z}\tilde{x}^* + \mathbf{A}^\dagger\nu'.$$

15

Introduce the notation

$$w = \mathbf{Z}_{i,i} x_i^* + \sum_{l \neq i,j} \mathbf{Z}_{i,l} x_l^*.$$

We have

$$x_i = \phi_\alpha \left(w + \mathbf{Z}_{i,j} x_j^* + \xi_i \right),$$
$$\tilde{x}_i = \phi_\alpha \left(w + \mathbf{Z}_{i,j} (x_j')^* + \tilde{\xi}_i \right).$$

(1) Since $\mathbf{Z}_{i,j} < 0$, $|\xi_i| \leq \rho$, and $|\tilde{\xi}_i| \leq \rho$, we know that when $w < \alpha - \rho$, $x_i = \tilde{x}_i = 0$. Then

$$\mathbb{E} \left[x_j^* (x_i - \tilde{x}_i) \right] = \Pr[w \geq \alpha - \rho] \, \mathbb{E} \left[x_j^* (x_i - \tilde{x}_i) | w \geq \alpha - \rho \right]. \tag{17}$$

By Property 10, $\phi_\alpha (\cdot)$ is 1-Lipschitz, so

$$|x_i - \tilde{x}_i| \leq |\mathbf{Z}_{i,j}| \left| x_j^* - (x_j')^* \right| + \left| \xi_i - \tilde{\xi}_i \right|,$$

which implies that

$$
\begin{aligned}
\left| \mathbb{E} \left[x_j^* (x_i - \tilde{x}_i) | w \geq \alpha - \rho \right] \right| &\leq \mathbb{E} \left[x_j^* |\mathbf{Z}_{i,j}| \left| x_j^* - (x_j')^* \right| + x_j^* \left| \xi_i - \tilde{\xi}_i \right| \Big| w \geq \alpha - \rho \right] \\
&\leq |\mathbf{Z}_{i,j}| \max \left\{ \left| x_j^* - (x_j')^* \right| \right\} \mathbb{E} \left[x_j^* | w \geq \alpha - \rho \right] + 2\rho \mathbb{E} \left[x_j^* | w \geq \alpha - \rho \right] \\
&\leq |\mathbf{Z}_{i,j}| \max \left\{ \left| x_j^* - (x_j')^* \right| \right\} \mathbb{E} \left[x_j^* \right] + 2\rho \mathbb{E} \left[x_j^* \right] \\
&\leq 2\mathbb{E} \left[x_j^* \right] \left(|\mathbf{Z}_{i,j}| + \rho \right) \\
&\leq \frac{2C_1}{n} \left(|\mathbf{Z}_{i,j}| + \rho \right). \tag{18}
\end{aligned}
$$

Now consider $\Pr[w \geq \alpha - \rho]$. Since

$$\mathbb{E} |w| \leq |\mathbf{Z}_{i,i}| \, \mathbb{E}[x_i^*] + \sum_{l \neq i,j} |\mathbf{Z}_{i,l}| \, \mathbb{E}[x_j^*] \leq \frac{C_1}{n} \|\mathbf{Z}^i\|_1,$$

we have that

$$\Pr[w \geq \alpha - \rho] \leq \frac{\mathbb{E} |w|}{\alpha - \rho} \leq \frac{C_1 \|\mathbf{Z}^i\|_1}{n(\alpha - \rho)} \tag{19}$$

Combining (17)(18) and (19) together completes the proof for the case when $\mathbf{Z}_{i,j} < 0$.

(2) Now consider the case when $\mathbf{Z}_{i,j} \geq 0$. Again, we have

$$x_i = \phi_\alpha \left(w + \mathbf{Z}_{i,j} x_j^* + \xi_i \right),$$
$$\tilde{x}_i = \phi_\alpha \left(w + \mathbf{Z}_{i,j} (x_j')^* + \tilde{\xi}_i \right).$$

For the analysis, introduce a variable

$$\tilde{u}_i = \phi_\alpha \left(w + \mathbf{Z}_{i,j} x_j^* + \tilde{\xi}_i \right).$$

If $(x_j')^* > x_j^*$, by Property 10 $\phi_\alpha(\cdot)$ is 1-Lipschitz, so

$$\tilde{x}_i \leq \tilde{u}_i + \mathbf{Z}_{i,j} \left((x_j')^* - x_j^* \right).$$

If $(x_j')^* \leq x_j^*$, by Property 10 $\phi_\alpha(\cdot)$ is non-decreasing, then

$$\tilde{x}_i \leq \tilde{u}_i.$$

In any case,

$$\tilde{x}_i \leq \tilde{u}_i + \mathbf{Z}_{i,j} (x_j')^*.$$

16

Therefore,

$$\mathbb{E}\left[x_j^*(x_i - \tilde{x}_i)\right] \geq \mathbb{E}\left[x_j^*(x_i - \tilde{u}_i)\right] - \mathbb{E}\left[x_j^* \mathbf{Z}_{i,j}(x_j')^*\right]$$

$$\geq \mathbb{E}\left[x_j^*(x_i - \tilde{u}_i)\right] - \frac{C_1^2}{n^2}\mathbf{Z}_{i,j}.$$

So we only need to consider $\mathbb{E}\left[x_j^*(x_i - \tilde{u}_i)\right]$. Let $\mathbf{G}$ denote the event that $x_i \neq 0$ or $\tilde{u}_i \neq 0$. Then by conditioning on x_j^*, we have

$$\mathbb{E}\left[x_j^*(x_i - \tilde{u}_i)\right] = \mathbb{E}\left[x_j^* \mathbb{E}\left[x_i - \tilde{u}_i \middle| x_j^*\right]\right]$$

and

$$\mathbb{E}\left[x_i - \tilde{u}_i \middle| x_j^*\right] = \Pr\left[\mathbf{G}\middle| x_j^*\right]\mathbb{E}\left[x_i - \tilde{u}_i \middle| x_j^*, \mathbf{G}\right].$$

By Property 10 $\phi_\alpha(\cdot)$ is 1-Lipschitz, so

$$\left|\mathbb{E}\left[x_i - \tilde{u}_i \middle| x_j^*, \mathbf{G}\right]\right| \leq \mathbb{E}\left[|\xi_i| + \left|\tilde{\xi}_i\right| \middle| x_j^*, \mathbf{G}\right] \leq 2\rho.$$

Now consider $\Pr\left[\mathbf{G}\middle| x_j^*\right]$. We have

$$\mathbb{E}\left[\left|w + \mathbf{Z}_{i,j}x_j^*\right| \middle| x_j^*\right] \leq \mathbb{E}\left[|w| \middle| x_j^*\right] + \mathbf{Z}_{i,j}$$

$$\leq \frac{C_1}{n}\left\|\mathbf{Z}^i\right\|_1 + \mathbf{Z}_{i,j},$$

where the first step follows from $x_j^* \leq 1$ and the second step follows from the conditional independence in Assumption (**A2**). Then by Markov's inequality,

$$\Pr\left[x_i \neq 0 \middle| x_j^*\right] \leq \Pr\left[\left|w + \mathbf{Z}_{i,j}x_j^*\right| \geq \alpha - \rho \middle| x_j^*\right]$$

$$\leq \frac{1}{\alpha - \rho}\left(\frac{C_1}{n}\left\|\mathbf{Z}^i\right\|_1 + \mathbf{Z}_{i,j}\right).$$

A similar argument leads to that

$$\Pr\left[\tilde{u}_i \neq 0 \middle| x_j^*\right] \leq \frac{1}{\alpha - \rho}\left(\frac{C_1}{n}\left\|\mathbf{Z}^i\right\|_1 + \mathbf{Z}_{i,j}\right)$$

and thus

$$\Pr\left[\mathbf{G}\middle| x_j^*\right] \leq \frac{2}{\alpha - \rho}\left(\frac{C_1}{n}\left\|\mathbf{Z}^i\right\|_1 + \mathbf{Z}_{i,j}\right).$$

Putting things together,

$$\left|\mathbb{E}\left[x_j^*(x_i - \tilde{u}_i)\right]\right| \leq \frac{4\rho}{\alpha - \rho}\left(\frac{C_1\|\mathbf{Z}^i\|_1}{n} + \mathbf{Z}_{i,j}\right)\mathbb{E}\left[x_j^*\right]$$

$$\leq \frac{4C_1\rho}{n(\alpha - \rho)}\left(\frac{C_1\|\mathbf{Z}^i\|_1}{n} + \mathbf{Z}_{i,j}\right).$$

This completes the proof for the lower bound.

Similarly, for the upper bound, introduce

$$u_i = \phi_\alpha\left(w + \mathbf{Z}_{i,j}(x_j')^* + \xi_i\right).$$

Then in any case,

$$x_i \leq u_i + \mathbf{Z}_{i,j}x_j^*$$

17

and thus

$$\mathbb{E}\left[x_j^*(x_i - \tilde{x}_i)\right] \leq \mathbb{E}\left[x_j^*(u_i - \tilde{x}_i)\right] + \mathbb{E}\left[(x_j^*)^2\right]\mathbf{Z}_{i,j}.$$

The same argument as above shows that

$$\left|\mathbb{E}\left[x_j^*(u_i - \tilde{x}_i)\right]\right| \leq \frac{4C_1\rho}{n(\alpha - \rho)}\left(\frac{C_1\|\mathbf{Z}^i\|_1}{n} + \mathbf{Z}_{i,j}\right).$$

This completes the whole proof. $\qquad\square$

Lemma 14 (Main: Bound on $\widetilde{\mathbf{N}}$). *Suppose $\|\mathbf{E}\|_{\mathrm{s}} \leq \ell$, $\boldsymbol{\Sigma} \succeq (1-\ell)\mathbf{I}$, and $|\xi_j| \leq \rho < \alpha$.*

(1) If the noise is correlated (Assumption (N1)), then

$$\left|\widetilde{\mathbf{N}}_{i,j}\right| \leq \frac{4C_\nu C_1}{(1-2\ell)^2 n(\alpha - \rho)} + |[\mathbf{N}_s]_{i,j}|$$

(2) If the noise is unbiased (Assumption (N2)) and $\|\mathbf{A}^\dagger\nu\|_\infty \leq \rho' < \alpha$, then

$$\left|\widetilde{\mathbf{N}}_{i,j}\right| \leq \frac{2C_1 C_\nu \rho'(1 + \|\mathbf{A}^\dagger\mathbf{N}\|_\infty)}{(1-2\ell)n(\alpha - \rho')} + \left|[\mathbf{N}_s]_{i,j}\right|.$$

Proof of Lemma 14. (1) By the update rule,

$$\widetilde{\mathbf{N}} = 2\mathbb{E}[\nu(x - x')^\top] + \mathbf{N}_s.$$

Under Assumption (**N1**), we have that for every $i \in [n], j \in [n]$,

$$\begin{aligned}
|\widetilde{\mathbf{N}}_{i,j}| &= |2\mathbb{E}[\nu_i(x_j - x_j')] + [\mathbf{N}_s]_{i,j}| \\
&\leq 4C_\nu\mathbb{E}[x_j] + |[\mathbf{N}_s]_{i,j}| \\
&= 4C_\nu\mathbb{E}\left[\phi_\alpha\left([\mathbf{Z}x^*]_j + \xi_j\right)\right] + |[\mathbf{N}_s]_{i,j}|.
\end{aligned}$$

since $|\nu_i|$ is bounded by C_ν.

Now focus on the term $\mathbb{E}\left[\phi_\alpha\left([\mathbf{Z}x^*]_j + \xi_j\right)\right]$. We have

$$|[\mathbf{Z}x^*]_j| \leq \|\mathbf{Z}\|_\infty\|x^*\|_\infty \leq \|\mathbf{Z}\|_\infty \leq \frac{1}{1-2\ell}$$

by the fact that $\|x^*\|_\infty \leq 1$ in Assumption (**A2**), and the assumptions of the lemma on $\boldsymbol{\Sigma}$ and $\mathbf{E}$. Then when $[\mathbf{Z}x^*]_j + \xi_j \geq \alpha$,

$$\phi_\alpha\left([\mathbf{Z}x^*]_j + \xi_j\right) \leq [\mathbf{Z}x^*]_j + \xi_j - \alpha \leq \frac{1}{1-2\ell} + \rho - \alpha \leq \frac{1}{1-2\ell},$$

and thus

$$\begin{aligned}
\mathbb{E}\left[\phi_\alpha\left([\mathbf{Z}x^*]_j + \xi_j\right)\right] &\leq \frac{1}{1-2\ell}\Pr\left\{[\mathbf{Z}x^*]_j + \xi_j \geq \alpha\right\} \\
&\leq \frac{1}{1-2\ell}\Pr\left\{|[\mathbf{Z}x^*]_j| \geq \alpha - \rho\right\} \\
&\leq \frac{1}{1-2\ell}\frac{\mathbb{E}|[\mathbf{Z}x^*]_j|}{\alpha - \rho} \\
&\leq \frac{1}{1-2\ell}\frac{\|\mathbf{Z}\|_\infty}{\alpha - \rho}\mathbb{E}\left[x_j^*\right] \\
&\leq \frac{C_1}{(1-2\ell)^2 n(\alpha - \rho)}
\end{aligned}$$

where the last step uses the bound on $\mathbb{E}\left[x_j^*\right]$ in Assumption (**A2**). Therefore,

$$|\widetilde{\mathbf{N}}_{i,j}| \leq \frac{4C_\nu C_1}{(1-2\ell)^2 n(\alpha - \rho)} + |[\mathbf{N}_s]_{i,j}|.$$

18

(2) When the noise is unbiased, we have $\mathbb{E}[\nu|x^*] = 0$. Then $\mathbb{E}[\nu_i x_j'] = 0$, and

$$\left|\widetilde{\mathbf{N}}_{i,j}\right| = \left|2\mathbb{E}[\nu_i(x_j - x_j')] + [\mathbf{N}_s]_{i,j}\right| \leq 2\left|\mathbb{E}[\nu_i x_j]\right| + |[\mathbf{N}_s]_{i,j}|. \tag{20}$$

Consider the first term for a fixed x^*, i.e., consider the conditional expectation $\mathbb{E}[\nu_i x_j \mid x^*]$. For notational simplicity, let $\widetilde{\mathbf{Z}} = (\mathbf{Z} - \mathbf{A}^\dagger \mathbf{N}\mathbf{Z})$ and $\tilde{\xi} = \mathbf{A}^\dagger \nu$. Then

$$\mathbb{E}[\nu_i x_j \mid x^*] = \mathbb{E}\left[\nu_i \phi_\alpha\left([\mathbf{Z}x^*]_j + \xi_j\right) \mid x^*\right] = \mathbb{E}\left[\nu_i \phi_\alpha\left([\widetilde{\mathbf{Z}}x^*]_j + \tilde{\xi}_j\right) \mid x^*\right].$$

We consider the following two cases about $[\widetilde{\mathbf{Z}}x^*]_j$.

(a) If $[\widetilde{\mathbf{Z}}x^*]_j \leq \alpha - \rho'$, then $\phi_\alpha\left([\widetilde{\mathbf{Z}}x^*]_j + \tilde{\xi}_j\right) = 0$ always holds, which implies that

$$|\mathbb{E}[\nu_i x_j \mid x^*]| = \mathbb{E}\left[\nu_i \phi_\alpha\left([\widetilde{\mathbf{Z}}x^*]_j + \tilde{\xi}_j\right) \mid x^*\right] = 0.$$

(b) If $[\widetilde{\mathbf{Z}}x^*]_j > \alpha - \rho'$, then

$$\phi_\alpha\left([\widetilde{\mathbf{Z}}x^*]_j + \tilde{\xi}_j\right) \leq \phi_\alpha\left([\widetilde{\mathbf{Z}}x^*]_j + \rho'\right) \leq [\widetilde{\mathbf{Z}}x^*]_j + \rho' - \alpha.$$

On the other side, by Property 10,

$$\phi_\alpha\left([\widetilde{\mathbf{Z}}x^*]_j + \tilde{\xi}_j\right) \geq [\widetilde{\mathbf{Z}}x^*]_j + \tilde{\xi}_j - \alpha \geq [\widetilde{\mathbf{Z}}x^*]_j - \rho' - \alpha.$$

Putting together, we conclude that

$$\nu_i([\widetilde{\mathbf{Z}}x^*]_j - \alpha) - |\nu_i\rho'| \leq \nu_i\phi_\alpha\left([\widetilde{\mathbf{Z}}x^*]_j + \tilde{\xi}_j\right) \leq \nu_i([\widetilde{\mathbf{Z}}x^*]_j - \alpha) + |\nu_i\rho'|.$$

Note that $\mathbb{E}[\nu_i([\widetilde{\mathbf{Z}}x^*]_j - \alpha)|x^*] = 0$, so

$$|\mathbb{E}[\nu_i x_j \mid x^*]| = \left|\mathbb{E}\left[\nu_i \phi_\alpha\left([\widetilde{\mathbf{Z}}x^*]_j + \tilde{\xi}_j\right) \mid x^*\right]\right| \leq \mathbb{E}[|\nu_i\rho'||x^*] \leq C_\nu\rho'.$$

Putting case (a) and case (b) together, we have

$$|\mathbb{E}[\nu_i x_j \mid x^*]| \leq C_\nu\rho' \Pr\left\{[\widetilde{\mathbf{Z}}x^*]_j > \alpha - \rho'\right\} \leq C_\nu\rho' \Pr\left\{\left|[\widetilde{\mathbf{Z}}x^*]_j\right| > \alpha - \rho'\right\}.$$

By definition of $\widetilde{\mathbf{Z}}$ and the assumptions of the lemma on $\mathbf{\Sigma}$ and $\mathbf{E}$,

$$\left|[\widetilde{\mathbf{Z}}x^*]_j\right| \leq (1 + \|\mathbf{A}^\dagger \mathbf{N}\|_\infty)\left|[\mathbf{Z}x^*]_j\right| \leq (1 + \|\mathbf{A}^\dagger \mathbf{N}\|_\infty)|\mathbf{Z}|_\infty x_j^* \leq \frac{1 + \|\mathbf{A}^\dagger \mathbf{N}\|_\infty}{1 - 2\ell}x_j^*. \tag{21}$$

Then

$$\Pr\left\{\left|[\widetilde{\mathbf{Z}}x^*]_j\right| > \alpha - \rho'\right\} \leq \frac{\mathbb{E}\left|[\widetilde{\mathbf{Z}}x^*]_j\right|}{\alpha - \rho'} \leq \frac{C_1(1 + \|\mathbf{A}^\dagger \mathbf{N}\|_\infty)}{(1 - 2\ell)n(\alpha - \rho')}.$$

The lemma then follows from (20) and (21). $\qquad\square$

There are three terms $\mathbf{Z}$, $\mathbf{V}$ and ξ in the above lemmas that need to be bounded. Since $\mathbf{Z} = \mathbf{V} + \mathbf{\Sigma}^{-1}$, we only need to bound $\mathbf{V}$ and ξ in the following two lemmas, respectively.

Lemma 15 (Bound on $\mathbf{V}$). *Suppose* $\|\mathbf{E}\|_s < \ell_e$ *and* $\mathbf{\Sigma} \succeq (1 - \ell)\mathbf{I}$. *Then*

(1) $\|\mathbf{V}_+\|_s \leq \dfrac{1 - \ell_e}{(1 - \ell)(1 - \ell_e - \ell)}\|\mathbf{E}_-\|_s + \dfrac{\ell}{(1 - \ell)^2(1 - \ell_e - \ell)}\|\mathbf{E}_+\|_s,$

(2) $\|\mathbf{V}_-\|_s \leq \dfrac{1 - \ell_e}{(1 - \ell)(1 - \ell_e - \ell)}\|\mathbf{E}_+\|_s + \dfrac{\ell}{(1 - \ell)^2(1 - \ell_e - \ell)}\|\mathbf{E}_-\|_s,$

(3) $\|\mathbf{V}\|_s \leq \dfrac{\ell_e(1 - \ell_e)}{(1 - \ell)^2(1 - \ell_e - \ell)},$

(4) $|\mathbf{V}_{i,i}| \leq \dfrac{\ell\ell_e}{(1 - \ell)^2(1 - \ell_e - \ell)}, \quad \forall i \in [n].$

19

Proof of Lemma 15. Denote $\mathbf{T} = \mathbf{\Sigma}^{-1} \sum_{k=2}^{\infty} (-\mathbf{E}\mathbf{\Sigma}^{-1})^k$, so that

$$\mathbf{V} = -\mathbf{\Sigma}^{-1}\mathbf{E}\mathbf{\Sigma}^{-1} + \mathbf{T}.$$

The following bound on $\|\mathbf{T}\|_1$ will be useful.

$$\|\mathbf{T}\|_1 \leq \left\|\mathbf{\Sigma}^{-1}\right\|_1 \sum_{k=2}^{\infty} \left\|(\mathbf{E}\mathbf{\Sigma}^{-1})^k\right\|_1$$

$$\leq \left\|\mathbf{\Sigma}^{-1}\right\|_1 \sum_{k=2}^{\infty} \left\|\mathbf{E}\mathbf{\Sigma}^{-1}\right\|_1^k$$

$$\leq \left\|\mathbf{\Sigma}^{-1}\right\|_1 \frac{\left\|\mathbf{E}\mathbf{\Sigma}^{-1}\right\|_1^2}{1 - \left\|\mathbf{E}\mathbf{\Sigma}^{-1}\right\|_1}$$

$$\leq \frac{1}{(1-\ell)^3} \times \ell \times \frac{\|\mathbf{E}\|_1}{1 - \frac{\ell_e}{1-\ell}}$$

$$\leq \frac{\ell}{(1-\ell)^2(1-\ell_e-\ell)}\|\mathbf{E}\|_1. \tag{22}$$

(1) We need to show the bound for both $\|\mathbf{V}_+\|_1$ and $\|\mathbf{V}_+\|_\infty$. By definition of $\mathbf{V}$, for any i,

$$\|\mathbf{V}_+\|_1 = \left\|\left[-\mathbf{\Sigma}^{-1}\mathbf{E}\mathbf{\Sigma}^{-1} + \mathbf{T}\right]_+\right\|_1.$$

Since for any $\mathbf{A}$ and $\mathbf{B}$,

$$\|[\mathbf{A}+\mathbf{B}]_+\|_1 \leq \|[\mathbf{A}]_+\|_1 + \|[\mathbf{B}]_+\|_1, \text{ and } \|[\mathbf{A}]_+\|_1 \leq \|\mathbf{A}\|_1,$$

we have

$$\|[\mathbf{V}_+]_i\|_1 \leq \left\|\left[-\mathbf{\Sigma}^{-1}\mathbf{E}\mathbf{\Sigma}^{-1}\right]_+\right\|_1 + \|\mathbf{T}_+\|_1$$

$$\leq \frac{1}{(1-\ell)^2}\|\mathbf{E}_-\|_1 + \|\mathbf{T}\|_1. \tag{23}$$

By (22),

$$\|\mathbf{T}\|_1 \leq \frac{\ell}{(1-\ell)^2(1-\ell_e-\ell)}\|\mathbf{E}\|_1 \leq \frac{\ell}{(1-\ell)^2(1-\ell_e-\ell)}(\|\mathbf{E}_-\|_1 + \|\mathbf{E}_+\|_1).$$

Combined with (23), it implies

$$\|[\mathbf{V}_+]_i\|_1 \leq \frac{1-\ell_e}{(1-\ell)^2(1-\ell_e-\ell)}\|\mathbf{E}_-\|_1 + \frac{\ell}{(1-\ell)^2(1-\ell_e-\ell)}\|\mathbf{E}_+\|_1.$$

Similarly, we have

$$\left\|[\mathbf{V}_+]^i\right\|_1 \leq \frac{1-\ell_e}{(1-\ell)^2(1-\ell_e-\ell)}\|\mathbf{E}_-\|_\infty + \frac{\ell}{(1-\ell)^2(1-\ell_e-\ell)}\|\mathbf{E}_+\|_\infty.$$

Putting things together we have

$$\|\mathbf{V}_+\|_s \leq \frac{1-\ell_e}{(1-\ell)(1-\ell_e-\ell)}\|\mathbf{E}_-\|_s + \frac{\ell}{(1-\ell)^2(1-\ell_e-\ell)}\|\mathbf{E}_+\|_s.$$

(2) The argument for $\|\mathbf{V}_-\|_s$ is similar to that for $\|\mathbf{V}_+\|_s$.

(3) We need to show the bound for both $\|\mathbf{V}\|_1$ and $\|\mathbf{V}\|_\infty$.

$$\|\mathbf{V}\|_1 \leq \left\|-\mathbf{\Sigma}^{-1}\mathbf{E}\mathbf{\Sigma}^{-1}\right\|_1 + \|\mathbf{T}\|_1$$

$$\leq \frac{\ell_e}{(1-\ell)^2} + \frac{\ell}{(1-\ell)^2(1-\ell_e-\ell)}\|\mathbf{E}\|_1$$

$$\leq \frac{\ell_e}{(1-\ell)^2} + \frac{\ell\ell_e}{(1-\ell)^2(1-\ell_e-\ell)}$$

$$= \frac{\ell_e(1-\ell_e)}{(1-\ell)^2(1-\ell_e-\ell)}$$

20

where the second step is by (22).

Similarly, $\|\mathbf{V}\|_\infty \le \frac{\ell_e(1-\ell_e)}{(1-\ell)^2(1-\ell_e-\ell)}$, so $\|\mathbf{V}\|_s \le \frac{\ell_e(1-\ell_e)}{(1-\ell)^2(1-\ell_e-\ell)}$.

(4) Now consider $\mathbf{V}_{i,i}$. By definition of $\mathbf{T}$.

$$\mathbf{V}_{i,i} = \left[-\mathbf{\Sigma}^{-1}\mathbf{E}\mathbf{\Sigma}^{-1}\right]_{i,i} + \mathbf{T}_{i,i}.$$

Note that since $\mathbf{E}_{i,i} = 0$, $\left[-\mathbf{\Sigma}^{-1}\mathbf{E}\mathbf{\Sigma}^{-1}\right]_{i,i} = 0$. Then

$$\begin{aligned}
|\mathbf{V}_{i,i}| &= |\mathbf{T}_{i,i}| \\
&\le \|\mathbf{T}\|_1 \\
&\le \frac{\ell}{(1-\ell)^2(1-\ell_e-\ell)}\|\mathbf{E}\|_1 \\
&\le \frac{\ell\ell_e}{(1-\ell)^2(1-\ell_e-\ell)}
\end{aligned}$$

where the third step is by (22). This completes the proof. $\qquad\square$

Lemma 16 (Bound on ξ). *Suppose* $\|\mathbf{E}\|_s < \ell \le 1/8$ *and* $\mathbf{\Sigma} \succeq (1-\ell)\mathbf{I}$. *Then for any* $i \in [n]$,

$$|\xi_i| \le \gamma := \frac{1}{1-2\ell}\left\|\mathbf{A}^\dagger\right\|_\infty\|\mathbf{N}\|_\infty + C_\nu\left\|\mathbf{A}^\dagger\right\|_\infty.$$

If furthermore, $\|\mathbf{N}\|_\infty\left\|(\mathbf{A}^*)^\dagger\right\|_\infty < 1/8$, *then*

$$\begin{aligned}
\left\|\mathbf{A}^\dagger\right\|_\infty &\le 2\left\|(\mathbf{A}^*)^\dagger\right\|_\infty, \\
\gamma &\le 3\left\|(\mathbf{A}^*)^\dagger\right\|_\infty \left(\|\mathbf{N}\|_\infty + C_\nu\right).
\end{aligned}$$

Proof of Lemma 16. First, we have

$$\|\xi\|_\infty \le \left\|\mathbf{A}^\dagger\mathbf{N}\mathbf{Z}x^*\right\|_\infty + \left\|\mathbf{A}^\dagger\nu\right\|_\infty \le \left\|\mathbf{A}^\dagger\right\|_\infty\|\mathbf{N}\|_\infty\|\mathbf{Z}\|_\infty\|x^*\|_\infty + \left\|\mathbf{A}^\dagger\right\|_\infty\|\nu\|_\infty.$$

Note that $\|x^*\|_\infty \le 1$ and $\|\nu\|_\infty \le C_\nu$. Furthermore,

$$\|\mathbf{Z}\|_\infty \le \frac{1}{1-2\ell}.$$

The first statement follows from combining these terms.

Now consider the second statement. We apply Lemma 17. Since

$$\begin{aligned}
\zeta &= \|\mathbf{E}\mathbf{\Sigma}^{-1} + (\mathbf{A}^*)^\dagger\mathbf{N}\mathbf{\Sigma}^{-1}\|_\infty \\
&\le \|\mathbf{E}\mathbf{\Sigma}^{-1}\|_\infty + \|(\mathbf{A}^*)^\dagger\mathbf{N}\mathbf{\Sigma}^{-1}\|_\infty \\
&\le \frac{1}{7} + \|(\mathbf{A}^*)^\dagger\|_\infty \times \|\mathbf{N}\|_\infty \times \|\mathbf{\Sigma}^{-1}\|_\infty \\
&\le \frac{2}{7},
\end{aligned}$$

Lemma 17 implies that

$$\|\mathbf{A}^\dagger\|_\infty \le \frac{\|\mathbf{\Sigma}^{-1}\|_\infty}{1-\zeta}\|(\mathbf{A}^*)^\dagger\|_\infty \le 2\|(\mathbf{A}^*)^\dagger\|_\infty.$$

Then γ is bounded by

$$\begin{aligned}
\gamma &= \frac{1}{1-2\ell}\left\|\mathbf{A}^\dagger\right\|_\infty\|\mathbf{N}\|_\infty + C_\nu\left\|\mathbf{A}^\dagger\right\|_\infty \\
&\le \frac{1}{1-2\ell} \times \left(2\|(\mathbf{A}^*)^\dagger\|_\infty\right) \times \|\mathbf{N}\|_\infty + C_\nu \times \left(2\|(\mathbf{A}^*)^\dagger\|_\infty\right) \\
&\le 3\left\|(\mathbf{A}^*)^\dagger\right\|_\infty \left(\|\mathbf{N}\|_\infty + C_\nu\right). \qquad\square
\end{aligned}$$

21

The following is the lemma about the norm of the pseudo-inverse, which is used in Lemma 16.

Lemma 17 (Pseudo-inverse). *Let $\mathbf{A}^*, \mathbf{N} \in \mathbb{R}^{m \times n}$ be two matrices with $m \geq n$. Let $(\mathbf{A}^*)^\dagger$ be one pseudo-inverse of $\mathbf{A}^*$ such that $(\mathbf{A}^*)^\dagger \mathbf{A}^* = \mathbf{I}$. Let $\mathbf{A} = \mathbf{A}^*(\mathbf{\Sigma} + \mathbf{E}) + \mathbf{N}$ be another matrix, with $\mathbf{\Sigma}$ being diagonal and*

$$\zeta := \|\mathbf{E}\mathbf{\Sigma}^{-1} + (\mathbf{A}^*)^\dagger \mathbf{N}\mathbf{\Sigma}^{-1}\|_\infty.$$

satisfies $\zeta < 1$. Then there exists a pseudo-inverse $\mathbf{A}^\dagger$ of $\mathbf{A}$ such that $\mathbf{A}^\dagger \mathbf{A} = \mathbf{I}$ and

$$\|\mathbf{A}^\dagger\|_\infty \leq \frac{\|\mathbf{\Sigma}^{-1}\|_\infty}{1 - \zeta} \|(\mathbf{A}^*)^\dagger\|_\infty.$$

Proof of Lemma 17. Consider the matrix

$$\mathbf{A}^\dagger = (\mathbf{\Sigma} + \mathbf{E} + (\mathbf{A}^*)^\dagger \mathbf{N})^{-1}(\mathbf{A}^*)^\dagger.$$

Then by definition,

$$\begin{aligned}
\mathbf{A}^\dagger \mathbf{A} &= (\mathbf{\Sigma} + \mathbf{E} + (\mathbf{A}^*)^\dagger \mathbf{N})^{-1}(\mathbf{A}^*)^\dagger \left(\mathbf{A}^*(\mathbf{\Sigma} + \mathbf{E}) + \mathbf{N}\right) \\
&= (\mathbf{\Sigma} + \mathbf{E} + (\mathbf{A}^*)^\dagger \mathbf{N})^{-1}(\mathbf{\Sigma} + \mathbf{E} + (\mathbf{A}^*)^\dagger \mathbf{N}) \\
&= \mathbf{I}.
\end{aligned}$$

What remains is to bound $\|\mathbf{A}^\dagger\|_\infty$. We have

$$\|\mathbf{A}^\dagger\|_\infty \leq \|(\mathbf{\Sigma} + \mathbf{E} + (\mathbf{A}^*)^\dagger \mathbf{N})^{-1}\|_\infty \|(\mathbf{A}^*)^\dagger\|_\infty.$$

By Taylor expansion rule, the first term on the right-hand side is

$$\begin{aligned}
(\mathbf{\Sigma} + \mathbf{E} + (\mathbf{A}^*)^\dagger \mathbf{N})^{-1} &= \left(\left(\mathbf{I} + \mathbf{E}\mathbf{\Sigma}^{-1} + (\mathbf{A}^*)^\dagger \mathbf{N}\mathbf{\Sigma}^{-1}\right)\mathbf{\Sigma}\right)^{-1} \\
&= \mathbf{\Sigma}^{-1}\left(\mathbf{I} + \mathbf{E}\mathbf{\Sigma}^{-1} + (\mathbf{A}^*)^\dagger \mathbf{N}\mathbf{\Sigma}^{-1}\right)^{-1} \\
&= \sum_{i=0}^{\infty} \mathbf{\Sigma}^{-1}\left(-\mathbf{E}\mathbf{\Sigma}^{-1} - (\mathbf{A}^*)^\dagger \mathbf{N}\mathbf{\Sigma}^{-1}\right)^i
\end{aligned}$$

where we use the assumption that $\|\mathbf{E}\mathbf{\Sigma}^{-1} + (\mathbf{A}^*)^\dagger \mathbf{N}\mathbf{\Sigma}^{-1}\|_\infty = \zeta < 1$. Therefore,

$$\|(\mathbf{\Sigma} + \mathbf{E} + (\mathbf{A}^*)^\dagger \mathbf{N})^{-1}\|_\infty \leq \|\mathbf{\Sigma}^{-1}\|_\infty \sum_{i=0}^{\infty} \zeta^i = \frac{\|\mathbf{\Sigma}^{-1}\|_\infty}{1 - \zeta}. \qquad \square$$

B.2 Putting things together

We are now ready to prove our main theorems.

Theorem 6 (Adversarial noise). *There exists an absolute constant $\mathcal{G}$ such that when Assumption (A0)-(A3) and (N1) are satisfied with $l = 1/10$, $C_2 \leq 2c_2$, $C_1^3 \leq \mathcal{G}c_2^2 n$, $C_\nu \leq \left\{\frac{c_2^2 \mathcal{G}c}{C_1^2 m}, \frac{c_2^4 \mathcal{G}c}{C_1^5 n \|(\mathbf{A}^*)^\dagger\|_\infty}\right\}$ for $0 \leq c \leq 1$, and $\|\mathbf{N}^{(0)}\|_\infty \leq \frac{c_2^2 \mathcal{G}c}{C_1^3 \|(\mathbf{A}^*)^\dagger\|_\infty}$, then there exist α, η, r such that for every $0 < \epsilon, \delta < 1$ and $N = \text{poly}(n, m, 1/\epsilon, 1/\delta)$, with probability at least $1 - \delta$ the following holds.*

After $T = O\left(\ln \frac{1}{\epsilon}\right)$ iterations, Algorithm 1 outputs a solution $\mathbf{A} = \mathbf{A}^*(\mathbf{\Sigma} + \mathbf{E}) + \mathbf{N}$ where $\mathbf{\Sigma} \succeq (1 - \ell)\mathbf{I}$ is diagonal, $\|\mathbf{E}\|_1 \leq \epsilon + c/2$ is off-diagonal, and $\|\mathbf{N}\|_1 \leq c/2$.

Proof of Theorem 6. We consider the following set of parameters

$$\alpha = \frac{c_2}{80C_1}, r = \frac{n}{c_2}, \eta = \frac{\ell}{6}.$$

Furthermore, set $\rho = B_1 \frac{c_2^2 c}{C_1^3}$ for a sufficiently small absolute constant B_1. Since $C_1 \geq n\mathbb{E}[x_i^*] \geq n\mathbb{E}[(x_i^*)^2] \geq c_2$, this is small enough so that

$$\rho \leq \min\left\{\frac{\alpha}{2}, \frac{c_2 \alpha}{2048 C_1}, \frac{c_2 \alpha}{8000 \times 100 C_1^2}, \frac{cc_2 \alpha}{48000 C_1^2}\right\}$$

22

which will be used in the proof. The proof also needs $C_1^2 \leq B_1 c_2 n, C_1^3 \leq B_2 c_2^2 n$ for sufficiently small absolute constants B_1 and B_2. Since $C_1 > c_2$, we only need $C_1^3 \leq \mathcal{G} c_2^2 n$. Similarly, we need

$$C_\nu \leq B_1 \min\left\{ \frac{c(\alpha-\rho)c_2}{mC_1}, \frac{(\alpha-\rho)c_2}{nC_1\|(\mathbf{A}^*)^\dagger\|_\infty}, \frac{(\alpha-\rho)c_2\rho}{nC_1\|(\mathbf{A}^*)^\dagger\|_\infty}, \frac{\rho}{\|(\mathbf{A}^*)^\dagger\|_\infty} \right\}$$

for a sufficiently small absolute constant B_1. This can be satisfied by setting $\mathcal{G}$ small enough in the theorem assumption.

After setting the parameters needed, we now prove the theorem. We prove it by proving the following three claims by induction on t: at the beginning of iteration t,

(1) $(1-\ell)\mathbf{I} \preceq \mathbf{\Sigma}^{(t)}$,

(2) $\left\|\mathbf{E}^{(t)}\right\|_s \leq \frac{1}{8}$, and if $t > 0$

$$\left\|\mathbf{E}_+^{(t)}\right\|_s + \beta\left\|\mathbf{E}_-^{(t)}\right\|_s \leq \left(1 - \frac{1}{25}\eta\right)\left(\left\|[\mathbf{E}]_+^{(t-1)}\right\|_s + \beta\left\|[\mathbf{E}]_-^{(t-1)}\right\|_s\right) + \frac{c}{10},$$

for $\beta = \frac{\sqrt{84^2+2800}-84}{2} \in (1,8)$,

(3) $\left\|\mathbf{N}^{(t)}\right\|_\infty \leq \frac{1}{8\|(\mathbf{A}^*)^\dagger\|_\infty}$, and $\|\xi^{(t)}\|_\infty \leq \rho$.

Claim (1) and (2) are clearly true at $t = 0$ by the assumption on initialization. The first part of Claim (3) is true because of the assumption that $\left\|\mathbf{N}^{(0)}\right\|_\infty \leq \frac{\mathcal{G}c}{8\mu^3\|(\mathbf{A}^*)^\dagger\|_\infty}$ and that $\mu = C_1/c_2 \geq 1$. Then the second part follows from Lemma 16.

Now we assume they are true up to t, and show them for $t + 1$.

(1) First consider the diagonal terms. Combining Lemma 12 and Lemma 15, we have

$$\widetilde{\mathbf{\Sigma}}_{i,i}^{(t)} \geq \mathbb{E}\left[(x_i^*)^2\right]\left(2(\mathbf{\Sigma}_{i,i}^{(t)})^{-1} - 2\left|\mathbf{V}_{i,i}^{(t)}\right|\right) - \frac{2C_1}{n}\left(\alpha + 2\rho + \frac{C_1}{n}(\mathbf{\Sigma}_{i,i}^{(t)})^{-1} + \frac{2C_1}{n}\left\|\left[\mathbf{V}^{(t)}\right]^i\right\|_1\right).$$

$$\geq \frac{2C_2}{n}\left(0 - \frac{\ell^2}{(1-2\ell)(1-\ell)^2}\right) - \frac{2C_1}{n}\left(\alpha + \alpha + \frac{C_1}{n}\frac{1}{1-\ell} + \frac{2C_1}{n}\frac{\ell}{(1-\ell)(1-2\ell)}\right).$$

$$= \frac{2C_2}{n}\left(0 - \frac{\ell^2}{(1-2\ell)(1-\ell)^2}\right) - \frac{2C_1}{n}\left(2\alpha + \frac{C_1}{n(1-\ell)(1-2\ell)}\right)$$

$$> -\frac{c_2}{5n}.$$

The first inequality uses $\rho < \alpha/2$ and the last inequality is due to $\alpha \leq \frac{c_2}{80C_1}$ and $C_1^2 \leq \frac{c_2 n}{80}$. Therefore,

$$\mathbf{\Sigma}_{i,i}^{(t+1)} = (1-\eta)\mathbf{\Sigma}_{i,i}^{(t)} + \eta r\widetilde{\mathbf{\Sigma}}_{i,i}^{(t)} \geq (1-\eta)\mathbf{\Sigma}_{i,i}^{(t)} - \frac{\eta}{5}.$$

Assume for contradiction $\mathbf{\Sigma}_{i,i}^{(t+1)} < 1 - \ell$. Then by the above inequality,

$$1 - \ell > \mathbf{\Sigma}_{i,i}^{(t+1)} \geq (1-\eta)\mathbf{\Sigma}_{i,i}^{(t)} - \frac{\eta}{5}.$$

which implies $\mathbf{\Sigma}_{i,i}^{(t)} \leq 1 - \ell + 2\eta$. In this case, by Lemma 12 and Lemma 15,

$$\widetilde{\mathbf{\Sigma}}_{i,i}^{(t)} \geq \mathbb{E}\left[(x_i^*)^2\right]\left(2(\mathbf{\Sigma}_{i,i}^{(t)})^{-1} - 2\left|\mathbf{V}_{i,i}^{(t)}\right|\right) - \frac{2C_1}{n}\left(\alpha + 2\rho + \frac{C_1}{n}(\mathbf{\Sigma}_{i,i}^{(t)})^{-1} + \frac{2C_1}{n}\left\|\left[\mathbf{V}^{(t)}\right]^i\right\|_1\right).$$

$$\geq \frac{2c_2}{n}\left(\frac{1}{1-\ell+2\eta} - \frac{\ell^2}{(1-2\ell)(1-\ell)^2}\right) - \frac{2C_1}{n}\left(2\alpha + \frac{C_1}{n(1-\ell)(1-2\ell)}\right)$$

$$> \frac{c_2}{n}.$$

Then

$$\mathbf{\Sigma}_{i,i}^{(t+1)} = (1-\eta)\mathbf{\Sigma}_{i,i}^{(t)} + \eta r\widetilde{\mathbf{\Sigma}}_{i,i}^{(t)} = (1-\eta)\mathbf{\Sigma}_{i,i}^{(t)} + \eta > \mathbf{\Sigma}_{i,i}^{(t)},$$

23

which is a contradiction. Therefore, $(1 - \ell)\mathbf{I} \preceq \mathbf{\Sigma}^{(t)}$.

(2) Now consider the off-diagonal terms. We shall split them into the positive part and the negative part. By the update rule, for any $i \in [n]$,

$$\left\| \left[\mathbf{E}_+^{(t+1)} \right]_i \right\|_1 \leq (1 - \eta) \left\| \left[\mathbf{E}_+^{(t)} \right]_i \right\|_1 + \eta r \left\| \left[\widetilde{\mathbf{E}}_+^{(t)} \right]_i \right\|_1.$$

Recall the notations

$$\mathbf{Z}^{(t)} = (\mathbf{\Sigma}^{(t)} + \mathbf{E}^{(t)})^{-1} = (\mathbf{\Sigma}^{(t)})^{-1} + \mathbf{V}^{(t)},$$

$$\mathbf{V}^{(t)} = (\mathbf{\Sigma}^{(t)})^{-1} \sum_{k=1}^{\infty} (-\mathbf{E}^{(t)}(\mathbf{\Sigma}^{(t)})^{-1})^k$$

By Lemma 13, we have

$$\left\| \left[\widetilde{\mathbf{E}}_+^{(t)} \right]_i \right\|_1 \leq \underbrace{\sum_{j \neq i} \frac{4C_1^2}{n^2(\alpha - \rho)} \left\| \left[\mathbf{Z}^{(t)} \right]^i \right\|_1 \left(\left| \left[\mathbf{Z}_-^{(t)} \right]_{i,j} \right| + \rho \right)}_{T1}$$

$$+ \underbrace{\sum_{j \neq i} \frac{8C_1\rho}{n(\alpha - \rho)} \left(\frac{C_1}{n} \left\| \left[\mathbf{Z}^{(t)} \right]^i \right\|_1 + \left| \left[\mathbf{Z}_+^{(t)} \right]_{i,j} \right| \right)}_{T2}$$

$$+ \underbrace{\sum_{j \neq i} 2\mathbb{E}[(x_j^*)^2] \left| \left[\mathbf{Z}_+^{(t)} \right]_{i,j} \right|}_{T3}.$$

First, by Lemma 15,

$$\left\| \left[\mathbf{Z}^{(t)} \right]^i \right\|_1 \leq \left[\left(\mathbf{\Sigma}^{(t)} \right)^{-1} \right]_{i,i} + \left\| \left[\mathbf{V}^{(t)} \right]^i \right\|_1 \leq \frac{1}{1 - 2\ell}$$

Now consider $\mathbf{Z}_+^{(t)}$ and $\mathbf{Z}_-^{(t)}$. We have

$$\sum_{j:j \neq i} \left| \left[\mathbf{Z}_-^{(t)} \right]_{i,j} \right| \leq \left\| \left[\mathbf{V}_-^{(t)} \right]_i \right\|_1, \quad \sum_{j:j \neq i} \left| \left[\mathbf{Z}_+^{(t)} \right]_{i,j} \right| \leq \left\| \left[\mathbf{V}_+^{(t)} \right]_i \right\|_1.$$

Therefore,

$$T1 \leq \frac{8C_1^2}{n^2(\alpha - \rho)} \left\| \left[\mathbf{V}_-^{(t)} \right]_i \right\|_1 + \frac{8C_1^2\rho}{n(\alpha - \rho)},$$

$$T2 \leq \frac{16C_1^2\rho}{n(\alpha - \rho)} + \frac{8C_1\rho}{n(\alpha - \rho)} \left\| \left[\mathbf{V}_+^{(t)} \right]_i \right\|_1,$$

$$T3 \leq \frac{2C_2}{n} \left\| \left[\mathbf{V}_+^{(t)} \right]_i \right\|_1.$$

and thus we have

$$\left\| \left[\widetilde{\mathbf{E}}_+^{(t)} \right]_i \right\|_1 \leq \frac{8C_1^2}{n^2(\alpha - \rho)} \left\| \left[\mathbf{V}_-^{(t)} \right]_i \right\|_1 + \left(\frac{2C_2}{n} + \frac{8C_1\rho}{n(\alpha - \rho)} \right) \left\| \left[\mathbf{V}_+^{(t)} \right]_i \right\|_1 + \frac{24C_1^2\rho}{n(\alpha - \rho)}.$$

Similarly, for any $i \in [n]$,

$$\left\| \left[\widetilde{\mathbf{E}}_+^{(t)} \right]^i \right\|_1 \leq \frac{8C_1^2}{n^2(\alpha - \rho)} \left\| \left[\mathbf{V}_-^{(t)} \right]^i \right\|_1 + \left(\frac{2C_2}{n} + \frac{8C_1\rho}{n(\alpha - \rho)} \right) \left\| \left[\mathbf{V}_+^{(t)} \right]^i \right\|_1 + \frac{24C_1^2\rho}{n(\alpha - \rho)}.$$

Putting the two together, we have

$$\left\| \widetilde{\mathbf{E}}_+^{(t)} \right\|_{\mathrm{s}} \leq \frac{8C_1^2}{n^2(\alpha - \rho)} \left\| \mathbf{V}_-^{(t)} \right\|_{\mathrm{s}} + \left(\frac{2C_2}{n} + \frac{8C_1\rho}{n(\alpha - \rho)} \right) \left\| \mathbf{V}_+^{(t)} \right\|_{\mathrm{s}} + \frac{24C_1^2\rho}{n(\alpha - \rho)}. \qquad (24)$$

24

By Lemma 15 and $\ell \le \frac{1}{8}$, we have:

$$\left\| \mathbf{V}_+^{(t)} \right\|_{\mathrm{s}} \le \frac{32}{21} \left\| \mathbf{E}_-^{(t)} \right\|_{\mathrm{s}} + \frac{32}{147} \left\| \mathbf{E}_+^{(t)} \right\|_{\mathrm{s}},$$

$$\left\| \mathbf{V}_-^{(t)} \right\|_{\mathrm{s}} \le \frac{32}{21} \left\| \mathbf{E}_+^{(t)} \right\|_{\mathrm{s}} + \frac{32}{147} \left\| \mathbf{E}_-^{(t)} \right\|_{\mathrm{s}}$$

So (24) becomes

$$\left\| \widetilde{\mathbf{E}}_+^{(t)} \right\|_{\mathrm{s}} \le \left(\frac{64 C_2}{147 n} + \frac{256 C_1 \rho}{147 n (\alpha - \rho)} + \frac{256 C_1^2}{21 n^2 (\alpha - \rho)} \right) \left\| \mathbf{E}_+^{(t)} \right\|_{\mathrm{s}} \tag{25}$$

$$+ \left(\frac{64 C_2}{21 n} + \frac{256 C_1 \rho}{21 n (\alpha - \rho)} + \frac{256 C_1^2}{147 n^2 (\alpha - \rho)} \right) \left\| \mathbf{E}_-^{(t)} \right\|_{\mathrm{s}} + \frac{24 C_1^2 \rho}{n (\alpha - \rho)}. \tag{26}$$

Now consider the negative part. The same argument as above leads to

$$\left\| \widetilde{\mathbf{E}}_-^{(t)} \right\|_{\mathrm{s}} \le \left(\frac{64 C_1^2}{147 n^2} + \frac{256 C_1 \rho}{147 n (\alpha - \rho)} + \frac{256 C_1^2}{21 n^2 (\alpha - \rho)} \right) \left\| \mathbf{E}_+^{(t)} \right\|_{\mathrm{s}}$$

$$+ \left(\frac{64 C_1^2}{21 n^2} + \frac{256 C_1 \rho}{21 n (\alpha - \rho)} + \frac{256 C_1^2}{147 n^2 (\alpha - \rho)} \right) \left\| \mathbf{E}_-^{(t)} \right\|_{\mathrm{s}} + \frac{24 C_1^2 \rho}{n (\alpha - \rho)}. \tag{27}$$

Note the difference between (26) and (27): $\frac{C_2}{n}$ in the former is replaced by $\frac{C_1^2}{n^2}$ in the latter, which is much smaller. This is crucial for our proof, which will be clear below.

For simplicity, we introduce the following notations:

$$a_t := \left\| \mathbf{E}_+^{(t)} \right\|_{\mathrm{s}}, \quad b_t := \left\| \mathbf{E}_-^{(t)} \right\|_{\mathrm{s}}.$$

Then by the update rule, we have

$$a_{t+1} \le (1 - \eta) a_t + \eta r \left\| \widetilde{\mathbf{E}}_+^{(t)} \right\|_{\mathrm{s}},$$

$$b_{t+1} \le (1 - \eta) b_t + \eta r \left\| \widetilde{\mathbf{E}}_-^{(t)} \right\|_{\mathrm{s}}.$$

Plugging in (26) and since $r = \frac{n}{c_2} \le \frac{2n}{C_2}$, we have

$$a_{t+1} \le (1 - \eta) a_t + \eta \frac{2n}{C_2} \left(\frac{64 C_2}{147 n} + \frac{256 C_1 \rho}{147 n (\alpha - \rho)} + \frac{256 C_1^2}{21 n^2 (\alpha - \rho)} \right) a_t$$

$$+ \eta \frac{2n}{C_2} \left(\frac{64 C_2}{21 n} + \frac{256 C_1 \rho}{21 n (\alpha - \rho)} + \frac{256 C_1^2}{147 n^2 (\alpha - \rho)} \right) b_t + \eta \frac{2n}{C_2} \frac{24 C_1^2 \rho}{n (\alpha - \rho)}$$

$$b_{t+1} \le (1 - \eta) b_t + \eta \frac{2n}{C_2} \left(\frac{64 C_1^2}{147 n^2} + \frac{256 C_1 \rho}{147 n (\alpha - \rho)} + \frac{256 C_1^2}{21 n^2 (\alpha - \rho)} \right) a_t$$

$$+ \eta \frac{2n}{C_2} \left(\frac{64 C_1^2}{21 n^2} + \frac{256 C_1 \rho}{21 n (\alpha - \rho)} + \frac{256 C_1^2}{147 n^2 (\alpha - \rho)} \right) b_t + \eta \frac{2n}{C_2} \frac{24 C_1^2 \rho}{n (\alpha - \rho)}.$$

When $\frac{512 C_1 \rho}{C_2 (\alpha - \rho)} \le \frac{1}{2}$ and $\frac{512 C_1^2}{C_2 n (\alpha - \rho)} \le \frac{1}{14}$,

$$a_{t+1} \le (1 - \eta) a_t + \frac{129}{147} \eta a_t + \frac{129}{21} \eta b_t + \eta \frac{48 C_1^2 \rho}{C_2 (\alpha - \rho)}$$

$$\le \left(1 - \frac{18}{147} \eta \right) a_t + \frac{129}{21} \eta b_t + \eta \frac{48 C_1^2 \rho}{C_2 (\alpha - \rho)}$$

Similarly, when $\frac{512 C_1 \rho}{C_2 (\alpha - \rho)} \le \frac{1}{2}$ and $\frac{512 C_1^2}{C_2 n (\alpha - \rho)} \le \frac{1}{14}$, and furthermore, $\frac{128 C_1^2}{C_2 n} \le \frac{1}{4}$,

$$b_{t+1} \le (1 - \eta) b_t + \frac{1}{100} \eta a_t + \frac{1}{25} \eta b_t + \eta \frac{48 C_1^2 \rho}{C_2 (\alpha - \rho)}$$

$$\le \left(1 - \frac{24}{25} \eta \right) b_t + \frac{1}{100} \eta a_t + \eta \frac{48 C_1^2 \rho}{C_2 (\alpha - \rho)}$$

25

Let $h = \frac{48C_1^2\rho}{C_2(\alpha-\rho)}$, we then have:

$$a_{t+1} \le \left(1 - \frac{3}{25}\eta\right) a_t + 7\eta b_t + \eta h,$$

$$b_{t+1} \le \left(1 - \frac{24}{25}\eta\right) b_t + \frac{1}{100}\eta a_t + \eta h.$$

Now set $\beta = \frac{\sqrt{84^2+2800}-84}{2}$, so that

$$a_{t+1} + \beta b_{t+1} \le \left(1 - \frac{3}{25}\eta\right) a_t + 7\eta b_t + \eta h + \left(\beta - \frac{24}{25}\eta\beta\right) b_t + \frac{\beta}{100}\eta a_t + \eta\beta h$$

$$= \left(1 - \frac{3}{25}\eta + \frac{\beta}{100}\eta\right)(a_t + \beta b_t) + \eta(1+\beta)h$$

$$\le \left(1 - \frac{1}{25}\eta\right)(a_t + \beta b_t) + 9\eta h,$$

where the last inequality follows from that $\beta < 8$.

Note that the recurrence is true up to $t+1$. Using Lemma 29 to solve this recurrence, we obtain

$$a_t + b_t \le a_0 + b_0 + 250h \le \frac{1}{10} + 250h \le \frac{1}{8}$$

when $\frac{4000C_1^2\rho}{C_2(\alpha-\rho)} \le \frac{1}{100}$. Moreover, we know that

$$\left\|\mathbf{E}^{(t+1)}\right\|_s \le a_{t+1} + \beta b_{t+1} \le \left(1 - \frac{1}{25}\eta\right)^t + 250h.$$

(3) Finally, consider the noise term. Set the sample size N to be large enough, so that by Lemma 14, we have

$$\left|\widetilde{\mathbf{N}}_{i,j}^{(t)}\right| \le \frac{4C_\nu C_1}{(1 - 2 \times \ell)^2 n(\alpha-\rho)} + \left|[\mathbf{N}_s^{(t)}]_{i,j}\right|$$

$$\le \frac{8C_\nu C_1}{n(\alpha-\rho)}.$$

Then by the update rule, we have $\left|\mathbf{N}_{i,j}^{(t+1)}\right| \le \frac{8C_\nu C_1}{(\alpha-\rho)c_2}$. Then

$$\left\|\mathbf{N}^{(t+1)}\right\|_\infty \le n \max_{i,j}\left|\mathbf{N}_{i,j}^{(t+1)}\right| \le \frac{8nC_\nu C_1}{(\alpha-\rho)c_2} \le \frac{1}{8\|(\mathbf{A}^*)^\dagger\|_\infty}$$

where the last inequality is due to

$$C_\nu \le \frac{(\alpha-\rho)c_2}{64nC_1\|(\mathbf{A}^*)^\dagger\|_\infty}.$$

On the other hand, by Lemma 16, we have

$$\|\xi^{(t+1)}\|_\infty \le 3\|(\mathbf{A}^*)^\dagger\|_\infty(\|\mathbf{N}^{(t+1)}\|_\infty + C_\nu)$$

$$\le 3\|(\mathbf{A}^*)^\dagger\|_\infty \left(\frac{8nC_\nu C_1}{(\alpha-\rho)c_2} + C_\nu\right) \le \rho$$

where the last inequality is due to

$$C_\nu \le \frac{(\alpha-\rho)c_2\rho}{48nC_1\|(\mathbf{A}^*)^\dagger\|_\infty}, \text{ and } C_\nu \le \frac{\rho}{6\|(\mathbf{A}^*)^\dagger\|_\infty}.$$

We also have (which will be useful in proving the final bound)

$$\left\|\mathbf{N}^{(t+1)}\right\|_1 \le m \max_{i,j}\left|\mathbf{N}_{i,j}^{(t+1)}\right| \le \frac{8mC_\nu C_1}{(\alpha-\rho)c_2} \le \frac{c}{10}$$

26

where the last inequality is due to

$$C_\nu \leq \frac{c(\alpha - \rho)c_2}{80mC_1}.$$

Now, we shall prove the theorem statements. Recall that solving the recurrence about a_t and b_t leads to

$$\left\| \mathbf{E}^{(t+1)} \right\|_{\mathrm{s}} \leq a_{t+1} + \beta b_{t+1} \leq \left(1 - \frac{1}{25}\eta \right)^t + 250h.$$

Since the setting of ρ makes sure $h = O(c)$, when $t = O\left(\ln \frac{1}{\epsilon} \right)$, we have the second statement $\left\| \widehat{\mathbf{E}} \right\|_{\mathrm{s}} \leq \epsilon + \frac{c}{2}$. Note that

$$\mathbf{A}^* \widehat{\boldsymbol{\Sigma}} = \mathbf{A} - \mathbf{A}^* \widehat{\mathbf{E}} - \widehat{\mathbf{N}}$$

and

$$\left\| \left[\mathbf{A}^* \widehat{\boldsymbol{\Sigma}} \right]_i \right\| = \widehat{\boldsymbol{\Sigma}}_{i,i}, \quad \|\mathbf{A}\|_1 = 1, \quad \left\| \mathbf{A}^* \widehat{\mathbf{E}} \right\|_1 = \left\| \widehat{\mathbf{E}} \right\|_1,$$

so we have

$$\widehat{\boldsymbol{\Sigma}}_{i,i} \geq \|\mathbf{A}\|_1 - \left\| \widehat{\mathbf{E}} \right\|_1 - \left\| \widehat{\mathbf{N}} \right\|_1$$
$$\geq 1 - \epsilon - c.$$

Similarly,

$$\widehat{\boldsymbol{\Sigma}}_{i,i} \leq \|\mathbf{A}\|_1 + \left\| \widehat{\mathbf{E}} \right\|_1 + \left\| \widehat{\mathbf{N}} \right\|_1$$
$$\leq 1 + \epsilon + c.$$

Then the final statement of the theorem follows by replacing c with $c/4$. This completes the proof. $\square$

Theorem 7 (Unbiased noise). *If Assumption (A0)-(A3) and (N2) are satisfied with $C_\nu = \frac{c_2 \mathcal{G}\sqrt{cn}}{C_1 \max\{m,n\} \|(\mathbf{A}^*)^\dagger\|_\infty}$ and the other parameters set as in Theorem 6, then the same guarantee holds.*

Proof. The proof is similar to that of Theorem 6, except using the second bound for unbiased noise in Lemma 14. We highlight the different part, that is, the induction on the noise term.

In the induction, by Lemma 14 we have when N is large enough,

$$\left| \widetilde{\mathbf{N}}_{i,j}^{(t)} \right| \leq \frac{2C_1 C_\nu \rho'(1 + \|\mathbf{A}^\dagger \mathbf{N}^{(t)}\|_\infty)}{(1 - 2\ell)n(\alpha - \rho')} + \left| \left[\mathbf{N}_s^{(t)} \right]_{i,j} \right| \leq \frac{3C_1 C_\nu \rho'(1 + \|\mathbf{A}^\dagger \mathbf{N}^{(t)}\|_\infty)}{n(\alpha - \rho')}.$$

By Lemma 16 and the induction, we have $\|\mathbf{A}^\dagger \mathbf{N}^{(t)}\|_\infty \leq 1/4$. Furthermore, $\rho' \leq C_\nu \|\mathbf{A}^\dagger\|_\infty \leq 2C_\nu \|(\mathbf{A}^*)^\dagger\|_\infty$ and the parameter setting makes sure $\rho' \leq \alpha/2$. Then

$$\left| \widetilde{\mathbf{N}}_{i,j}^{(t)} \right| \leq \frac{16C_\nu^2 C_1 \|(\mathbf{A}^*)^\dagger\|_\infty}{n\alpha}.$$

Then by the update rule, we have

$$\left| \mathbf{N}_{i,j}^{(t+1)} \right| \leq \frac{32C_\nu^2 C_1 \|(\mathbf{A}^*)^\dagger\|_\infty}{c_2 \alpha}$$

and

$$\left\| \mathbf{N}^{(t+1)} \right\|_\infty \leq \frac{32nC_\nu^2 C_1 \|(\mathbf{A}^*)^\dagger\|_\infty}{c_2 \alpha} \leq \frac{1}{8\|(\mathbf{A}^*)^\dagger\|_\infty} \tag{28}$$

by the definition of α, and $C_\nu \leq \frac{1}{256} \frac{c_2}{C_1} \frac{\sqrt{n}}{n\|(\mathbf{A}^*)^\dagger\|_\infty}$. This completes the induction for the noise.

Also, in proving the final bounds, we have

$$\left\| \mathbf{N}^{(t+1)} \right\|_1 \leq \frac{32mC_\nu^2 C_1 \|(\mathbf{A}^*)^\dagger\|_\infty}{c_2 \alpha} \leq \frac{c}{10} \tag{29}$$

27

by the definition of α, and

$$C_\nu \leq \frac{c}{320} \frac{c_2}{C_1} \frac{\sqrt{n}}{\max\{m, n\|(\mathbf{A}^*)^\dagger\|_\infty\}} \leq \frac{\sqrt{c}}{320} \frac{c_2}{C_1} \frac{1}{\sqrt{m\|(\mathbf{A}^*)^\dagger\|_\infty}}$$

where the last inequality can be shown by consider the two cases when $\|(\mathbf{A}^*)^\dagger\|_\infty \leq m/n$ and $\|(\mathbf{A}^*)^\dagger\|_\infty \geq m/n$. The rest of the proof is the same as in Theorem 6. $\qquad\square$

C   Results for general proportions: Equilibration

Algorithm 2 ColumnUpdate

Input: A matrix $\mathbf{A}$, a threshold value α, a step size η, ratios $\{r_j : j \in [n]\}$, iteration number T, a subset $S \subseteq [n]$, sample size N
1: Set $\mathbf{A}^{(0)} = \mathbf{A}$
2: **for** $t = 0 \to T - 1$ **do**
3:

$$\forall i \in S, [\mathbf{A}^{(t+1)}]_i = \left[(1 - \eta)\,\mathbf{A}^{(t)} + r_i \eta \tilde{\mathbb{E}}\left[(y - y')(x - x')^\top\right]\right]_i \qquad (30)$$

Output: $\hat{\mathbf{A}} = \mathbf{A}^{(T)}$

Algorithm 3 Rescale

Input: A matrix $\mathbf{A}$, a threshold value α, a step size η, ratios $\{r_j : j \in [n]\}$, iteration number T, and a set $S \subseteq [n]$, $\epsilon \in (0, 1)$.
1: Let $\tilde{\mathbf{A}} = \text{ColumnUpdate}(\mathbf{A}, \alpha, \eta, \{r_j\}_j, T, S, N)$
2: **for** $i \in S$ **do**
3:    Set $[\hat{\mathbf{A}}]_i = \frac{1}{1-\epsilon}[\tilde{\mathbf{A}}]_i$
Output: $\hat{\mathbf{A}}$

Algorithm 4 Equilibration

Input: $\mathbf{A}, \alpha, \eta, T,$ and $\epsilon \in (0, 1), \lambda, N$
1: $S \leftarrow \emptyset, \mathbf{D} \leftarrow \mathbf{I}$
2: **while** $|S| \leq n$ **do**
3:    $m_j \leftarrow \hat{\mathbb{E}}[x_j^2]$ for $j \notin S$ using N examples
4:    **while** $\max_{j \notin S} m_j < \lambda$ **do**
5:       $\mathbf{A} \leftarrow \text{Rescale}(\mathbf{A}, \alpha, \eta, \{3/(5m_j) : j \in [n]\}, T, S, \epsilon, N)$
6:       $\lambda \leftarrow (1 - \epsilon)\lambda, \mathbf{D}_{j,j} \leftarrow \mathbf{D}_{j,j}/(1 - \epsilon)$
7:       $m_j \leftarrow (1 - \epsilon)^2 m_j$ for $j \in S$, and $m_j \leftarrow \hat{\mathbb{E}}[x_j^2]$ for $j \notin S$ using N examples
8:    $S \leftarrow S \cup \{j : m_j \geq \lambda\}$
Output: $\mathbf{A}$

When the feature have various proportions (i.e., $\mathbb{E}[(x_i^*)^2]$ varies for different i), we propose Algorithm 4 for balancing them. The idea is quite simple: instead of solving $\mathbf{Y} \approx \mathbf{A}^*\mathbf{X}$, we could also solve $\mathbf{Y} \approx [\mathbf{A}^*\mathbf{D}][(\mathbf{D})^{-1}\mathbf{X}]$ for a positive diagonal matrix $\mathbf{D}$. Our goal is to find $\mathbf{A} = \mathbf{A}^*\mathbf{D}(\boldsymbol{\Sigma} + \mathbf{E}) + \mathbf{N}$ so that $\boldsymbol{\Sigma}$ is large, $\mathbf{E}, \mathbf{N}$ are small, while $\mathbb{E}[(x_i^*)^2]/\mathbf{D}_{i,i}^2$ is with in a factor of 2 from each other.

The algorithm works at stages and keeps a working set S of column index i such that $\mathbb{E}[(x_i^*)^2]/\mathbf{D}_{i,i}^2$ is above a threshold λ. At each stage, it only updates the columns in S; at the end of the stage, it increases these columns by a small factor so that $\mathbb{E}[(x_i^*)^2]/\mathbf{D}_{i,i}^2$ decreases. Then it decreases the threshold λ, and add more columns to the working set and repeat. In this way, $\mathbb{E}[(x_i^*)^2]/\mathbf{D}_{i,i}^2 (i \in S)$ are always balanced; in particular, they are balanced at the end when $S = [n]$. Formally,

28

Theorem 18 (Main: Equilibration). *If there exists an absolute constant $\mathcal{G}$ such that Assumption (A1)-(A3) and (N1) are satisfied with $l = 1/50$, $C_1^3 \le \mathcal{G}c_2^2 n$, $\max\{C_\nu, \|\mathbf{N}^{(0)}\|_\infty\} \le \frac{\mathcal{G}c_2^4}{C_1^5 n \|(\mathbf{A}^*)^\dagger\|_\infty}$, and additionally $\boldsymbol{\Sigma}^{(0)} \preceq (1 - \ell)\mathbf{I}$, and $\mathbf{E} \ge 0$ entry-wise, then there exist α, η, T, λ such that for sufficiently small $\epsilon > 0$ and sufficiently large $N = \mathrm{poly}(n, m, 1/\epsilon, 1/\delta)$ the following hold with probability at least $1 - \delta$: Algorithm 4 outputs a solution $\mathbf{A} = \mathbf{A}^*\mathbf{D}(\boldsymbol{\Sigma} + \mathbf{E}) + \mathbf{N}$ where $\boldsymbol{\Sigma} \succeq (1 - \ell)\mathbf{I}$ is diagonal, $\|\mathbf{E}\|_s \le \gamma\ell$ is off-diagonal, $\|\mathbf{N}\|_\infty \le 2\|\mathbf{N}^{(0)}\|_\infty$, and $\mathbf{D}$ is diagonal and satisfies*

$$\frac{\max_{i \in [n]} \frac{1}{\mathbf{D}_{i,i}^2}\mathbb{E}[(x_i^*)^2]}{\min_{j \in [n]} \frac{1}{\mathbf{D}_{j,j}^2}\mathbb{E}[(x_j^*)^2]} \le 2.$$

If Assumption (A1)-(A3) and (N2) are satisfied with the same parameters except
$$\max\{C_\nu, \|\mathbf{N}^{(0)}\|_\infty\} \le \min\left\{\sqrt{\frac{\mathcal{G}c_2^4}{C_1^5 n}}\frac{1}{\|(\mathbf{A}^*)^\dagger\|_\infty}, \frac{\mathcal{G}c_2^2}{C_1^3\|(\mathbf{A}^*)^\dagger\|_\infty}\right\}, \text{ then the same guarantees hold.}$$

Now, we can view $\mathbf{A}^*\mathbf{D}$ as the ground-truth feature matrix and $\mathbf{D}^{-1}x^*$ as the weights. Then applying Algorithm 1 with $\mathbf{A}$ can recover $\mathbf{A}^*\mathbf{D}$, and after normalization we get $\mathbf{A}^*$.

The initialization condition of the theorem can be achieved by the popular practical heuristic that sets the columns of $\mathbf{A}^{(0)}$ to reasonable almost pure data points. It is generally believed that it gives $\mathbf{E}_{i,j}^{(0)} \ge 0$ and $\mathbf{N}^{(0)} = 0$. We note that the parameters are not optimized; the algorithm can potentially tolerate much better initialization.

Intuition. Before delving into the specifics of the algorithm, it will be useful to provide a high-level outline of the proof. As described above, the algorithm makes use of the fact that samples from a ground truth matrix $\mathbf{A}^*$ and distribution x^* can equivalently be viewed as coming from the ground truth matrix $\mathbf{A}^*\mathbf{D}$ and distribution $\mathbf{D}^{-1}x^*$, for some diagonal matrix $\mathbf{D}$. Therefore, the goal is to find a $\mathbf{D}$ such that the features are balanced:

$$\frac{\max_{i \in [n]} \frac{\mathbb{E}[(x_i^*)^2]}{\mathbf{D}_{i,i}^2}}{\min_{i \in [n]} \frac{\mathbb{E}[(x_i^*)^2]}{\mathbf{D}_{i,i}^2}} \le \kappa.$$

The algorithm will implicitly calculate such a $\mathbf{D}$ gradually. Namely, at any point in time, the algorithm will have an active set $S \subseteq [n]$ of features, which are balanced, i.e.

$$\frac{\max_{i \in [n]} \frac{\mathbb{E}[(x_i^*)^2]}{\mathbf{D}_{i,i}^2}}{\min_{i \in S} \frac{\mathbb{E}[(x_i^*)^2]}{\mathbf{D}_{i,i}^2}} \le \kappa. \tag{31}$$

It is clear that when $S = [n]$ the algorithm achieves the goal. Our algorithm begins with $S = \emptyset$ and gradually increase S until $S = [n]$.

The mechanism for increasing S will be as follows. Given S, $\mathbf{A}$ is of the form

$$\mathbf{A} = \mathbf{A}^*\mathbf{D}(\boldsymbol{\Sigma} + \mathbf{E}) + \mathbf{N}$$

with

$$\mathbf{E} = \begin{bmatrix} \mathbf{E}_{1,1} & \mathbf{E}_{1,2} \\ \mathbf{E}_{2,1} & \mathbf{E}_{2,2} \end{bmatrix}$$

where the columns of $\mathbf{A}$ are sorted such that the first $|S|$ columns correspond to the features of S, and $\mathbf{E}_{1,1} \in \mathbb{R}^{|S| \times |S|}$, $\mathbf{E}_{2,1} \in \mathbb{R}^{(n-|S|) \times |S|}$, $\mathbf{E}_{1,2} \in \mathbb{R}^{|S| \times (n-|S|)}$, $\mathbf{E}_{2,2} \in \mathbb{R}^{(n-|S|) \times (n-|S|)}$. Then scaling up the columns of $\mathbf{A}$ indexed by S by a factor of $\frac{1}{1-\epsilon}$ is equivalent to

1. scaling up the columns of $\mathbf{D}$ indexed by S by a factor of $\frac{1}{1-\epsilon}$ and
2. scaling up the columns of $\mathbf{E}_{2,1}$ by a factor of $\frac{1}{1-\epsilon}$ and
3. scaling down the columns of $\mathbf{E}_{1,2}$ by a factor of $1 - \epsilon$.

29

Therefore, to increase the set S, the algorithm will scale up the columns of $\mathbf{A}$ indexed by S, until some $j \notin S$ satisfies

$$\max_{i \in [n]} \frac{\mathbb{E}[(x_i^*)^2]}{\mathbf{D}_{i,i}^2} \leq \kappa \frac{\mathbb{E}[(x_j^*)^2]}{\mathbf{D}_{j,j}^2}.$$

Then it can add j into S while keeping the corresponding features balanced as in (31). Note that we do not need to explicitly maintain $\mathbf{D}$, though it can be calculated along with the scaling. Further note that the values of $\mathbb{E}[(x_i^*)^2]$ are not known but they can be estimated using the current $\mathbf{A}$.

However, there is still one caveat: $\mathbf{E}$ should be kept small, so that at the end of the algorithm, we still have a good initialization $\mathbf{A}$. For this reason, the algorithm additionally maintains that for a small constant $1 < \gamma < 2$,

$$\|\mathbf{E}_{1,1}\|_{\mathrm{s}} \leq \gamma\ell, \quad \|\mathbf{E}_{1,2}\|_{\mathrm{s}} \leq \ell,$$
$$\|\mathbf{E}_{2,1}\|_{\mathrm{s}} \leq \gamma\ell, \quad \|\mathbf{E}_{2,2}\|_{\mathrm{s}} \leq \ell. \tag{32}$$

Since scaling up $\mathbf{A}$ will scale up $\mathbf{E}_{2,1}$, we will need to first decrease $\|\mathbf{E}_{2,1}\|_{\mathrm{s}}$ before the scaling step. The key observation is that by applying our training algorithm only on the columns indexed by S, $\|\mathbf{E}_{1,1}\|_{\mathrm{s}}$ and $\|\mathbf{E}_{2,1}\|_{\mathrm{s}}$ will be decreased, while $\|\mathbf{E}_{1,2}\|_{\mathrm{s}}$ and $\|\mathbf{E}_{2,2}\|_{\mathrm{s}}$ unchanged. On a high level, using the fact that the matrix $\mathbf{E}_{1,2}$ has no negative entries (which we get by virtue of our initialization), and the fact that the contribution in the updates to the entry $(\mathbf{E}_{1,1})_{i,j}$ mostly comes from $(\mathbf{E}_{1,1})_{j,i}$ (i.e. the matrix $\mathbf{E}_{1,1}$ in the first order contribution "updates itself"), and the fact that the features in S are balanced, we can show that after sufficiently many updates, the symmetric norm of $\mathbf{E}_{1,1}$ and $\mathbf{E}_{2,1}$ drops by a reasonable amount: $\|\mathbf{E}_{1,1}\|_{\mathrm{s}} \leq (\gamma - 1)\ell$ and $\|\mathbf{E}_{2,1}\|_{\mathrm{s}} \leq (1 - \epsilon)(\gamma - 1)\ell$. Now, we can do the scaling step without hurting the invariant 32.

Organization. The result of the section is as follows. We first prove in Section C.1 that applying our training algorithm only on the columns indexed by S will decrease $\|\mathbf{E}_{1,1}\|_{\mathrm{s}}$ and $\|\mathbf{E}_{2,1}\|_{\mathrm{s}}$. Then in Section C.2 we analyze the scaling step, and show that the invariant (32) is maintained. In Section C.3, we show how to increase S while maintaining the invariant (31), where the main technical details are about how to estimate $\mathbb{E}[(x_i^*)^2]$.

C.1 Equilibration: ColumnUpdate

In this subsection, we focus on the update step, bounding the changes of $\mathbf{\Sigma}, \mathbf{E}$, and $\mathbf{N}$.

First recall some notations. Let $\mathbf{A} = \mathbf{A}^*(\mathbf{\Sigma} + \mathbf{E}) + \mathbf{N}$ where $\mathbf{\Sigma}$ is diagonal, $\mathbf{E}$ is off diagonal, and $\mathbf{N}$ is the component outside the span of $\mathbf{A}^*$.[3] Given the set $S \subseteq [n]$ and a matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$, let $\mathbf{M}_{1,1}$ denote the submatrix indexed by $S \times S$, and $\mathbf{M}_{2,1}$ denote the submatrix indexed by $([n] - S) \times S$, $\mathbf{M}_{1,2}$ denote the submatrix indexed by $S \times ([n] - S)$, and $\mathbf{M}_{2,2}$ denote the submatrix indexed by $([n] - S) \times ([n] - S)$.[4] In the special case when $S = [s]$ where $s = |S|$,

$$\mathbf{M} = \begin{bmatrix} \mathbf{M}_{1,1} & \mathbf{M}_{1,2} \\ \mathbf{M}_{2,1} & \mathbf{M}_{2,2} \end{bmatrix}.$$

Also, let $\mathbf{M}_S$ denote the submatrix formed by the columns indexed by S, and $\mathbf{M}_{-S}$ the submatrix formed by the other columns.[5]

The input $\mathbf{A}^{(0)}$ of Algorithm 2 can be written as $\mathbf{A}^{(0)} = \mathbf{A}^*(\mathbf{\Sigma}^{(0)} + \mathbf{E}^{(0)}) + \mathbf{N}^{(0)}$ where $\mathbf{\Sigma}^{(0)}$ is diagonal, and $\mathbf{E}^{(0)}$ is off diagonal. Define $\mathbf{E}_{1,1}^{(0)}, \mathbf{E}_{1,2}^{(0)}, \mathbf{E}_{2,1}^{(0)}$ and $\hat{\mathbf{E}}_{2,2}$ as described above. Similarly, define $\hat{\mathbf{E}}_{1,1}, \hat{\mathbf{E}}_{1,2}, \hat{\mathbf{E}}_{2,1}$ and $\hat{\mathbf{E}}_{2,2}$ for the output $\hat{\mathbf{A}} = \mathbf{A}^*(\hat{\mathbf{\Sigma}} + \hat{\mathbf{E}}) + \hat{\mathbf{N}}$ of Algorithm 2. Finally, define $\mathbf{N}_S^{(0)}, \mathbf{N}_{-S}^{(0)}, \hat{\mathbf{N}}_S$, and $\hat{\mathbf{N}}_{-S}$ as described above.

The main result of the subsection is Lemma 19.

---

[3]Note that $\mathbf{A}^*$ here can be any ground-truth matrix; in particular, later Lemma 19 will be applied where $\mathbf{A}^*$ in the lemma corresponds to $\mathbf{A}^*\mathbf{D}$ in the intuition described above.

[4]These notations will be used for $\mathbf{M} = \mathbf{E}$, $\mathbf{M} = \widetilde{\mathbf{E}}$, and related matrices.

[5]These notations will be used for $\mathbf{M} = \mathbf{N}$ or $\mathbf{M} = \widetilde{\mathbf{N}}$, and related matrices.

30

Lemma 19 (Main: ColumnUpdate). *Define*

$$R_j = \mathbb{E}[(x_j^*)^2], \quad R = \max_{j \in [n]} R_j, \quad r = \max_{j \in S} r_j, \tag{33}$$

$$h_1 = r \frac{8C_1(C_1 + 1)\rho}{(1 - \ell - \beta\ell)n(\alpha - \rho)} + \frac{4C_1^2}{(1 - \ell - \beta\ell)n^2(\alpha - \rho)} r \left(\frac{1}{(1 - \ell - \beta\ell)} + 1 \right), \tag{34}$$

$$h_2 = r \frac{R\beta^2\ell^2}{(1 - \ell)^2(1 - \ell - \beta\ell)} + \frac{12C_1(C_1 + 1)}{n^2(\alpha - \rho)(1 - \ell - \beta\ell)} \left(\frac{1}{1 - \ell - \beta\ell} + n\rho \right) r, \tag{35}$$

$$h = h_1 + h_2, \tag{36}$$

$$U_a = \frac{8rC_\nu C_1}{\alpha - \rho}, \tag{37}$$

$$U_n = \frac{10rC_1C_\nu^2 \|(\mathbf{A}^*)^\dagger\|_\infty}{(1 - 2\ell)(\alpha - 2C_\nu\|(\mathbf{A}^*)^\dagger\|_\infty)}. \tag{38}$$

Suppose $\ell \le 1/8$, β is a constant with $\beta\ell \le 1/2$, $\gamma \in (1, 2)$, $\epsilon \in (0, 1)$. The initialization satisfies $(1 - \ell)\mathbf{I} \preceq \mathbf{\Sigma}^{(0)}$, $\left\|\mathbf{E}_{1,1}^{(0)}\right\|_s \le \gamma\ell$, $\left\|\mathbf{E}_{2,1}^{(0)}\right\|_s \le \gamma\ell$, $\left\|(\mathbf{E}_{1,2}^{(0)}; \mathbf{E}_{2,2}^{(0)})\right\|_s \le \ell$, $\mathbf{E}_{1,2}^{(0)} \ge 0$ and $\mathbf{E}_{2,2}^{(0)} \ge 0$ *entry-wise, and* $\|\mathbf{N}_{-S}^{(0)}\|_\infty \le U$ *and* $\|\mathbf{N}_S^{(0)}\|_\infty \le 2U \le 1/(16\|(\mathbf{A}^*)^\dagger\|_\infty)$. *Furthermore, the parameters satisfy that for any $i \in S$,*

$$\eta \left(1 + 2r_i R_i \frac{1}{(1 - \ell)^2} \frac{\beta\ell^2}{1 - \beta\ell - \ell} + r_i \frac{2C_1}{n} \left(\alpha + 2\rho + \frac{C_1}{n} + \frac{2C_1}{n} \frac{\beta\ell(1 - \beta\ell)}{(1 - \ell)^2(1 - \beta\ell - \ell)} \right) \right) \le \ell \tag{39}$$

$$r_i R_i \left(2 - 2\frac{1}{(1 - \ell)^2} \frac{\beta\ell^2}{1 - \beta\ell - \ell} \right) - r_i \left(\frac{2C_1}{n} \left(\alpha + 2\rho + \frac{C_1}{n} \frac{1}{1 - \ell} + \frac{2C_1}{n} \frac{\beta\ell(1 - \beta\ell)}{(1 - \ell)^2(1 - \beta\ell - \ell)} \right) \right) \ge 1 - \ell \tag{40}$$

$$h_1 \le \ell, \quad \left(\frac{rR}{(1 - \ell)^2} + 1 \right)(\epsilon + h_1) + (\epsilon + h_2) \le (\gamma - 1)\ell \tag{41}$$

$$\epsilon + h_2 \le (1 - \epsilon)(\gamma - 1)\ell \tag{42}$$

$$h_1 + \ell \le (\beta - 1)\ell, \quad h_2 + \left(\frac{rR}{(1 - \ell)^2} + 1 \right)\ell \le (\beta - 1)\ell \tag{43}$$

$$3\|(\mathbf{A}^*)^\dagger\|_\infty (3U + C_\nu) \le \rho < \alpha. \tag{44}$$

*If we have adversarial noise (Assumption (**N1**)), assume*

$$\epsilon' + U_a \le (1 - \epsilon)U, \text{ and } 3\|(\mathbf{A}^*)^\dagger\|_\infty (2U + U_a + C_\nu) \le \rho < \alpha < 1. \tag{45}$$

*If we have unbiased noise (Assumption (**N2**)), assume*

$$\epsilon' + U_n \le (1 - \epsilon)U. \tag{46}$$

Finally, let $N = poly\,(n, m, 1/\delta, 1/\epsilon)$ sufficiently large.

Then with probability at least $1 - \delta$, after $\frac{2\ln(\epsilon/(\gamma\ell))}{\ln(1-\eta)} + \frac{\ln(\epsilon'/U)}{\ln(1-\eta)}$ iterations, the output of Algorithm 2 is $\hat{\mathbf{A}} = \mathbf{A}^*(\hat{\mathbf{\Sigma}} + \hat{\mathbf{E}}) + \hat{\mathbf{N}}$ satisfying

$$(1 - \ell)\mathbf{I} \preceq \hat{\mathbf{\Sigma}} \preceq u\mathbf{I}, \quad \|\hat{\mathbf{E}}_{1,1}\|_s \le (\gamma - 1)\ell, \quad \|\hat{\mathbf{E}}_{2,1}\|_s \le (1 - \epsilon)(\gamma - 1)\ell, \quad \|(\hat{\mathbf{E}}_{1,2}; \hat{\mathbf{E}}_{2,2})\|_s \le \ell,$$

and $\hat{\mathbf{E}}_{1,2} \ge 0$ and $\hat{\mathbf{E}}_{2,2} \ge 0$ entry-wise. Furthermore, $\|\hat{\mathbf{N}}_{-S}\|_\infty \le U$ and $\|\hat{\mathbf{N}}_S\|_\infty \le (1 - \epsilon)U$.

Proof of Lemma 19. It follows from Lemma 22 and the conditions (41) and (42). ☐

To prove Lemma 22, we will first consider how $\mathbf{E}$ changes after one update step, and then derive the recurrence for all steps in Lemma 22.

31

C.1.1 One update step of E

In this subsection, we focus on one update step, bounding the change of $\mathbf{E}$. So through out this subsection we will focus on a particular iteration t and omit the superscript (t), while in the next subsection we will put back the superscript.

For analysis, denote $\mathbf{A}^{(t)}$ as

$$\mathbf{A} = \mathbf{A}^*(\mathbf{\Sigma} + \mathbf{E}) + \mathbf{N}$$

where $\mathbf{\Sigma}$ is a diagonal matrix, $\mathbf{E}$ is an off-diagonal matrix, and $\mathbf{N}$ is the component of $\mathbf{A}$ that lies outside the span of $\mathbf{A}^*$ (e.g., the noise caused by the noise in the sample).

Recall the following notations:

$$\mathbf{Z} = (\mathbf{\Sigma} + \mathbf{E})^{-1},$$

$$\mathbf{V} = \mathbf{Z} - \mathbf{\Sigma}^{-1} = \mathbf{\Sigma}^{-1} \sum_{k=1}^{\infty} (-\mathbf{E}\mathbf{\Sigma}^{-1})^k,$$

$$\xi = -\mathbf{A}^\dagger \mathbf{N} \mathbf{Z} x^* + \mathbf{A}^\dagger \nu.$$

Consider the update term $\hat{\mathbb{E}}\left[(y - y')(x - x')^\top\right]$ and denote it as

$$\mathbf{\Delta} = \hat{\mathbb{E}}\left[(y - y')(x - x')^\top\right] = \mathbf{A}^*(\widetilde{\mathbf{\Sigma}} + \widetilde{\mathbf{E}}) + \widetilde{\mathbf{N}}$$

where $\widetilde{\mathbf{\Sigma}}$ is a diagonal matrix, $\widetilde{\mathbf{E}}$ is an off-diagonal matrix, and $\mathbf{N}$ is the component of $\mathbf{\Delta}$ that lies outside the span of $\mathbf{A}^*$.

Since we now use empirical average, we will have sampling noise. Denote it as

$$\mathbf{N}_s = \hat{\mathbb{E}}[(y - y')(x - x')^\top] - \mathbb{E}[(y - y')(x - x')^\top].$$

Then by definition, for $y = \mathbf{A}^* x^* + \nu$ and $y' = \mathbf{A}^*(x')^* + \nu'$, we have

$$\hat{\mathbb{E}}[(y - y')(x - x')^\top] = \mathbb{E}[(y - y')(x - x')^\top] + \mathbf{N}_s$$
$$= \mathbf{A}^* \underbrace{\mathbb{E}\left[(x^* - (x')^*)(x - x')^\top\right]}_{\widetilde{\mathbf{\Sigma}} + \widetilde{\mathbf{E}}} + \underbrace{\mathbb{E}\left[(\nu - \nu')(x - x')^\top\right] + \mathbf{N}_s}_{\widetilde{\mathbf{N}}}.$$

Recall the definition of $\mathbf{E}_{1,1}$, i.e., it is the submatrix of $\mathbf{E}$ indexed by $S \times S$. Define $\widetilde{\mathbf{E}}_{1,1}$ similarly, i.e., it is the submatrix of $\widetilde{\mathbf{E}}$ indexed by $S \times S$. Define $\widetilde{\mathbf{E}}_{1,2}, \widetilde{\mathbf{E}}_{2,1}$ and $\widetilde{\mathbf{E}}_{2,2}$ accordingly. So in the special case when $S = [s]$ where $s = |S|$,

$$\widetilde{\mathbf{E}} = \begin{bmatrix} \widetilde{\mathbf{E}}_{1,1} & \widetilde{\mathbf{E}}_{1,2} \\ \widetilde{\mathbf{E}}_{2,1} & \widetilde{\mathbf{E}}_{2,2} \end{bmatrix}.$$

We also use the notation $\mathbf{M}^+$ or $\mathbf{M}^-$ to denote the positive or negative part of a matrix $\mathbf{M}$.

Lemma 20 (Update $\widetilde{\mathbf{E}}_{1,1}$). *Let $\widetilde{\mathbf{E}}_{1,1}$ be defined as above. If $\|\xi\|_\infty \leq \rho < \alpha < 1$ and $\mathbf{\Sigma} \succeq (1 - \ell)\mathbf{I}$, then*
(1). Negative entries:

$$\|\widetilde{\mathbf{E}}_{1,1}^-\|_s \leq \frac{4C_1^2 \|\mathbf{Z}\|_s(\|\mathbf{Z}\|_s + 1)}{n^2(\alpha - \rho)} + \frac{8C_1(C_1 + 1)\rho\|\mathbf{Z}\|_s}{n(\alpha - \rho)}.$$

(2) Positive entries:

$$\|\widetilde{\mathbf{E}}_{1,1}^+\|_s \leq \frac{12C_1(C_1 + 1)\|\mathbf{Z}\|_s}{n^2(\alpha - \rho)} \left(\|\mathbf{Z}\|_s + n\rho\right) + 2\max_{j \in [n]}\{\mathbb{E}[(x_j^*)^2]\} \left(\frac{1}{(1 - \ell)^2}\|\mathbf{E}_{1,1}^-\|_s + \frac{\|\mathbf{E}\|_s^2}{(1 - \ell)^2(1 - \ell - \|\mathbf{E}\|_s)}\right).$$

Proof of Lemma 20. (1) By Lemma 13, we have

$$\|\widetilde{\mathbf{E}}_{1,1}^-\|_s \leq \max\left\{\frac{4C_1^2\|\mathbf{Z}\|_s}{n^2(\alpha - \rho)}\|\mathbf{Z}\|_s + \frac{4C_1^2\|\mathbf{Z}\|_s}{n^2(\alpha - \rho)}n\rho, \frac{8C_1\rho}{n(\alpha - \rho)}(C_1 + 1)\|\mathbf{Z}\|_s + \frac{2C_1^2}{n^2}\|\mathbf{Z}\|_s\right\}.$$

32

Observe that for $\alpha < 1$,

$$\frac{4C_1^2\|\mathbf{Z}\|_s(\|\mathbf{Z}\|_s + 1)}{n^2(\alpha - \rho)} \geq \max\left\{\frac{4C_1^2\|\mathbf{Z}\|_s^2}{n^2(\alpha - \rho)}, \frac{2C_1^2}{n^2}\|\mathbf{Z}\|_s\right\}.$$

Moreover,

$$\frac{8C_1\rho}{n(\alpha - \rho)}(C_1 + 1)\|\mathbf{Z}\|_s \geq \frac{4C_1^2\|\mathbf{Z}\|_s}{n^2(\alpha - \rho)}n\rho.$$

Therefore,

$$\|\widetilde{\mathbf{E}}_{1,1}^-\|_s \leq \frac{4C_1^2\|\mathbf{Z}\|_s}{n^2(\alpha - \rho)} + \frac{8C_1(C_1 + 1)\rho\|\mathbf{Z}\|_s}{n(\alpha - \rho)}.$$

(2) By Lemma 13, when $\mathbf{Z}_{i,j} < 0$,

$$\widetilde{\mathbf{E}}_{j,i} \leq \frac{4C_1^2\|\mathbf{Z}^i\|_1}{n^2(\alpha - \rho)}\left(|\mathbf{Z}_{i,j}| + \rho\right).$$

When $\mathbf{Z}_{i,j} \geq 0$,

$$\widetilde{\mathbf{E}}_{j,i} \leq \frac{8C_1\rho}{n(\alpha - \rho)}\left(\frac{C_1\|\mathbf{Z}^i\|_1}{n} + \mathbf{Z}_{i,j}\right) + 2\mathbb{E}[(x_j^*)^2]\mathbf{Z}_{i,j}$$

Consider a fixed i. Let $G = \{j \in S, \mathbf{Z}_{i,j} \geq 0\}$ and let $G^c = S - G$. We know that

$$
\begin{aligned}
\|[\widetilde{\mathbf{E}}_{1,1}^+]_i\|_1 &= \sum_{j \in [n]} [\widetilde{\mathbf{E}}_{1,1}^+]_{j,i} \\
&\leq \sum_{j \in G^c} \frac{4C_1^2\|\mathbf{Z}^i\|_1}{n^2(\alpha - \rho)}\left(|\mathbf{Z}_{i,j}| + \rho\right) \\
&\quad + \sum_{j \in G}\left(\frac{8C_1\rho}{n(\alpha - \rho)}\left(\frac{C_1\|\mathbf{Z}^i\|_1}{n} + \mathbf{Z}_{i,j}\right) + 2\mathbb{E}[(x_j^*)^2]\mathbf{Z}_{i,j}\right) \\
&\leq \frac{4C_1^2\|\mathbf{Z}\|_s}{n^2(\alpha - \rho)}\left(\|\mathbf{Z}\|_s + n\rho\right) + \frac{8C_1(C_1 + 1)\rho}{n(\alpha - \rho)}\|\mathbf{Z}\|_s + \sum_{j \in G} 2\mathbb{E}[(x_j^*)^2]\mathbf{Z}_{i,j} \\
&\leq \frac{4C_1^2\|\mathbf{Z}\|_s^2}{n^2(\alpha - \rho)} + \frac{4C_1^2\|\mathbf{Z}\|_s}{n^2(\alpha - \rho)}n\rho + \frac{8C_1(C_1 + 1)\rho}{n(\alpha - \rho)}\|\mathbf{Z}\|_s + \sum_{j \in S} 2\mathbb{E}[(x_j^*)^2]\mathbf{Z}_{i,j} \\
&\leq \frac{12C_1(C_1 + 1)\|\mathbf{Z}\|_s}{n^2(\alpha - \rho)}\left(\|\mathbf{Z}\|_s + n\rho\right) + \sum_{j \in G} 2\mathbb{E}[(x_j^*)^2]\mathbf{Z}_{i,j}.
\end{aligned}
$$

A similar bound holds for $\|[\widetilde{\mathbf{E}}_{1,1}^+]^i\|_1$.

By the definition of $\mathbf{Z}$, we know that

$$
\begin{aligned}
\mathbf{Z} &= (\mathbf{\Sigma} + \mathbf{E})^{-1} \\
&= \mathbf{\Sigma}^{-1}\sum_{k=0}^{\infty}(-\mathbf{E}\mathbf{\Sigma}^{-1})^k \\
&= \mathbf{\Sigma}^{-1} - \mathbf{\Sigma}^{-1}\mathbf{E}\mathbf{\Sigma}^{-1} + \mathbf{\Sigma}^{-1}\sum_{k=2}^{\infty}(-\mathbf{E}\mathbf{\Sigma}^{-1})^k.
\end{aligned}
$$

Therefore, we know that for $i \neq j$,

$$\mathbf{Z}_{i,j} \leq -[\mathbf{\Sigma}^{-1}\mathbf{E}\mathbf{\Sigma}^{-1}]_{i,j} + |\sum_{k=2}^{\infty}\mathbf{\Sigma}^{-1}[(-\mathbf{E}\mathbf{\Sigma}^{-1})^k]_{i,j}|.$$

33

This implies that

$$\sum_{j \in G} \mathbf{Z}_{i,j} \le \sum_{j \in G} \left(-[\mathbf{\Sigma}^{-1}\mathbf{E}\mathbf{\Sigma}^{-1}]_{i,j} + \sum_{k=2}^{\infty} \left| \mathbf{\Sigma}^{-1}[(-\mathbf{E}\mathbf{\Sigma}^{-1})^k]_{i,j} \right| \right)$$

$$\le \frac{1}{(1-\ell)^2}\|\mathbf{E}_{1,1}^-\|_s + \frac{1}{1-\ell}\frac{\frac{\|\mathbf{E}\|_s^2}{(1-\ell)^2}}{1 - \frac{\|\mathbf{E}\|_s}{1-\ell}}$$

$$\le \frac{1}{(1-\ell)^2}\|\mathbf{E}_{1,1}^-\|_s + \frac{\|\mathbf{E}\|_s^2}{(1-\ell)^2(1-\ell-\|\mathbf{E}\|_s)}.$$

Putting together, we complete the proof. $\square$

Lemma 21 (Update $\widetilde{\mathbf{E}}_{2,1}$). *Let $\widetilde{\mathbf{E}}_{2,1}$ be defined as above, and suppose $\|\xi\|_\infty \le \rho < \alpha < 1$, $\mathbf{\Sigma} \succeq (1-\ell)\mathbf{I}$ and $\mathbf{E}_{1,2} \ge 0$, then we have*

$$\|\widetilde{\mathbf{E}}_{2,1}\|_s \le \frac{12C_1(C_1+1)\|\mathbf{Z}\|_s}{n^2(\alpha-\rho)}\left(\|\mathbf{Z}\|_s + n\rho\right) + 2\max_{j \in [n]}\{\mathbb{E}[(x_j^*)^2]\}\left(\frac{\|\mathbf{E}\|_s^2}{(1-\ell)^2(1-\ell-\|\mathbf{E}\|_s)}\right).$$

Proof of Lemma 21. The proof is almost the same as that of Lemma 20, combined with the fact that $\mathbf{E}_{1,2} \ge 0$ entry-wise. $\square$

C.1.2 Recurrence

Recall that
$$\mathbf{A} = \mathbf{A}^*(\mathbf{\Sigma}+\mathbf{E}) + \mathbf{N}$$
and recall that $\mathbf{E}_{1,1}$ is the submatrix indexed by $S \times S$, and $\mathbf{E}_{1,2}, \mathbf{E}_{2,1}, \mathbf{E}_{2,2}$ are defined according. Recall that $\mathbf{M}_S$ denote the submatrix of $\mathbf{M}$ formed by columns indexed by S, and let $\mathbf{M}_{-S}$ denote the submatrix formed by the other columns.

Lemma 22 (Recurrence). *Suppose the conditions in Lemma 19 hold. Then with probability at least $1-\delta$, after $\frac{2\ln(\epsilon/(\gamma\ell))}{\ln(1-\eta)}$ iterations,*

$$(1-\ell)\mathbf{I} \preceq \mathbf{\Sigma}^{(t)},$$

$$\|(\mathbf{E}_{1,1}^{(t)})^-\|_s \le \epsilon + h_1,$$

$$\|(\mathbf{E}_{1,1}^{(t)})^+\|_s \le \frac{rR}{(1-\ell)^2}(\epsilon+h_1) + h_2 + \epsilon,$$

$$\|(\mathbf{E}_{2,1}^{(t)})\|_s \le \epsilon + h_2.$$

Also, after $\frac{\ln(\epsilon'/U)}{\ln(1-\eta)}$ iterations, for both adversarial and unbiased noise,

$$\left\|\mathbf{N}_{-S}^{(t)}\right\|_\infty \le U, \quad \left\|\mathbf{N}_S^{(t)}\right\|_\infty \le (1-\epsilon)U.$$

Proof of Lemma 22. We first prove the following claims by induction.
(1) $(1-\ell)\mathbf{I} \preceq \mathbf{\Sigma}^{(t)}$,
(2)

$$\|(\mathbf{E}_{1,1}^-)^{(t)}\|_s \le \gamma\ell$$

$$\|(\mathbf{E}_{1,1}^+)^{(t)}\|_s \le \frac{rR}{(1-\ell)^2}\gamma\ell + h_2$$

$$\|\mathbf{E}_{2,1}^{(t)}\|_s \le \gamma\ell$$

$$\|\mathbf{E}_{1,2}^{(t)}\|_s \le \ell$$

$$\|\mathbf{E}_{2,2}^{(t)}\|_s \le \ell,$$

34

(3) $\|\mathbf{E}^{(t)}\|_s \leq \beta\ell$,

(4) for adversarial noise, $\left\|\mathbf{N}_S^{(t)}\right\|_\infty \leq U + U_a$, and $\|\xi^{(t)}\|_\infty \leq \rho$; or for unbiased noise, $\left\|\mathbf{N}_S^{(t)}\right\|_\infty \leq U + U_u$.

The basis case for $t = 0$ is trivial by assumptions. Now assume they are true for iteration t and show that they are true for iteration $t + 1$.

(1) By the update of $\boldsymbol{\Sigma}$, we have

$$\boldsymbol{\Sigma}^{(t+1)} = (1 - \eta)\boldsymbol{\Sigma}^{(t)} + \eta r \widetilde{\boldsymbol{\Sigma}}^{(t)}.$$

To lower bound $\boldsymbol{\Sigma}_{i,i}^{(t+1)}$, we will consider two cases, $\boldsymbol{\Sigma}_{i,i}^{(t)} \geq 1$ and $\boldsymbol{\Sigma}_{i,i}^{(t)} \leq 1$.

For $\boldsymbol{\Sigma}_{i,i}^{(t)} \geq 1$, by Lemma 12,

$$\widetilde{\boldsymbol{\Sigma}}_{i,i} \geq \mathbb{E}\left[(x_i^*)^2\right]\left(2\boldsymbol{\Sigma}_{i,i}^{-1} - 2\left|\mathbf{V}_{i,i}\right|\right) - \frac{2C_1}{n}\left(\alpha + 2\rho + \frac{C_1}{n}\boldsymbol{\Sigma}_{i,i}^{-1} + \frac{2C_1}{n}\left\|[\mathbf{V}]^i\right\|_1\right)$$

$$\geq -2R_i\left|\mathbf{V}_{i,i}\right| - \left(\frac{2C_1}{n}\left(\alpha + 2\rho + \frac{C_1}{n}\boldsymbol{\Sigma}_{i,i}^{-1} + \frac{2C_1}{n}\left\|[\mathbf{V}]^i\right\|_1\right)\right).$$

Hence,

$$\boldsymbol{\Sigma}_{i,i}^{(t+1)} \geq (1 - \eta)\boldsymbol{\Sigma}_{i,i}^{(t)} - \eta\left(2r_iR_i\left|\mathbf{V}_{i,i}^{(t)}\right| + r_i\left(\frac{2C_1}{n}\left(\alpha + 2\rho + \frac{C_1}{n}(\boldsymbol{\Sigma}_{i,i}^{(t)})^{-1} + \frac{2C_1}{n}\left\|\left[\mathbf{V}^{(t)}\right]^i\right\|_1\right)\right)\right)$$

$$\geq 1 - \eta\left(1 + 2r_iR_i\left|\mathbf{V}_{i,i}^{(t)}\right| + r_i\frac{2C_1}{n}\left(\alpha + 2\rho + \frac{C_1}{n} + \frac{2C_1}{n}\left\|\left[\mathbf{V}^{(t)}\right]^i\right\|_1\right)\right)$$

$$\geq 1 - \eta\left(1 + 2r_iR_i\frac{1}{(1-\ell)^2}\frac{\beta\ell^2}{1 - \beta\ell - \ell} + r_i\frac{2C_1}{n}\left(\alpha + 2\rho + \frac{C_1}{n} + \frac{2C_1}{n}\frac{\beta\ell(1 - \beta\ell)}{(1-\ell)^2(1 - \beta\ell - \ell)}\right)\right).$$

where we use the bound on $\mathbf{V}^{(t)}$. By condition (39), the claim follows.

For $\boldsymbol{\Sigma}_{i,i}^{(t)} \leq 1$, again by Lemma 12,

$$\widetilde{\boldsymbol{\Sigma}}_{i,i}^{(t)} \geq \mathbb{E}\left[(x_i^*)^2\right]\left(2 - 2\left|\mathbf{V}_{i,i}^{(t)}\right|\right) - \left(\frac{2C_1}{n}\left(\alpha + 2\rho + \frac{C_1}{n}(\boldsymbol{\Sigma}_{i,i}^{(t)})^{-1} + \frac{2C_1}{n}\left\|\left[\mathbf{V}^{(t)}\right]^i\right\|_1\right)\right).$$

Hence,

$$\boldsymbol{\Sigma}_{i,i}^{(t+1)} = (1 - \eta)\boldsymbol{\Sigma}_{i,i}^{(t)} + \eta r \widetilde{\boldsymbol{\Sigma}}_{i,i}^{(t)}$$
$$\geq (1 - \eta)(1 - \ell)$$
$$\quad + \eta\left(r_iR_i\left(2 - 2\left|\mathbf{V}_{i,i}^{(t)}\right|\right) - r_i\left(\frac{2C_1}{n}\left(\alpha + 2\rho + \frac{C_1}{n}(\boldsymbol{\Sigma}_{i,i}^{(t)})^{-1} + \frac{2C_1}{n}\left\|\left[\mathbf{V}^{(t)}\right]^i\right\|_1\right)\right)\right)$$
$$\geq (1 - \eta)(1 - \ell) + \eta r_iR_i\left(2 - 2\frac{1}{(1-\ell)^2}\frac{\beta\ell^2}{1 - \beta\ell - \ell}\right)$$
$$\quad - \eta r_i\left(\frac{2C_1}{n}\left(\alpha + 2\rho + \frac{C_1}{n}\frac{1}{1-\ell} + \frac{2C_1}{n}\frac{\beta\ell(1 - \beta\ell)}{(1-\ell)^2(1 - \beta\ell - \ell)}\right)\right).$$

By condition (40), the claim follows.

(2) By Lemma 20,

$$\|(\widetilde{\mathbf{E}}_{1,1}^{(t+1)})^-\|_s \leq \frac{8C_1(C_1 + 1)\rho\|\mathbf{Z}^{(t)}\|_s}{n(\alpha - \rho)} + \frac{4C_1^2\|\mathbf{Z}^{(t)}\|_s(\|\mathbf{Z}^{(t)}\|_s + 1)}{n^2(\alpha - \rho)},$$

$$\|(\widetilde{\mathbf{E}}_{1,1}^{(t+1)})^+\|_s \leq \frac{R}{(1-\ell)^2}\|(\mathbf{E}_{1,1}^-)^{(t)}\|_s + \frac{R\|\mathbf{E}^{(t)}\|_s^2}{(1-\ell)^2(1 - \ell - \|\mathbf{E}^{(t)}\|_s)}$$
$$\quad + \frac{12C_1(C_1 + 1)\|\mathbf{Z}^{(t)}\|_s}{n^2(\alpha - \rho)}\left(\left\|\mathbf{Z}^{(t)}\right\|_s + n\rho\right).$$

35

By the update rule, we have

$$\|(\mathbf{E}_{1,1}^{(t+1)})^-\|_s \leq (1-\eta)\|(\mathbf{E}_{1,1}^{(t)})^-\|_s$$

$$+ r\eta \frac{8C_1(C_1+1)\rho}{(1-\ell-\beta\ell)n(\alpha-\rho)} + \frac{4C_1^2}{(1-\ell-\beta\ell)n^2(\alpha-\rho)}\left(\frac{1}{(1-\ell-\beta\ell)}+1\right)r\eta,$$

$$\leq (1-\eta)\|(\mathbf{E}_{1,1}^{(t)})^-\|_s + \eta h_1 \tag{47}$$

$$\|(\mathbf{E}_{1,1}^{(t+1)})^+\|_s \leq (1-\eta)\|(\mathbf{E}_{1,1}^{(t)})^+\|_s + r\eta \frac{R}{(1-\ell)^2}\|(\mathbf{E}_{1,1}^{(t)})^-\|_s$$

$$+ r\eta \frac{R\beta^2\ell^2}{(1-\ell)^2(1-\ell-\beta\ell)}$$

$$+ \frac{12C_1(C_1+1)}{n^2(\alpha-\rho)(1-\ell-\beta\ell)}\left(\frac{1}{1-\ell-\beta\ell}+n\rho\right)r\eta$$

$$\leq (1-\eta)\|(\mathbf{E}_{1,1}^{(t)})^+\|_s + r\eta \frac{R}{(1-\ell)^2}\|(\mathbf{E}_{1,1}^{(t)})^-\|_s + \eta h_2 \tag{48}$$

where we use $\left\|\mathbf{E}^{(t)}\right\|_{\mathrm{s}} \leq \beta\ell$ and $\|\mathbf{Z}^{(t)}\|_s \leq \frac{1}{1-\ell-\beta\ell}$.

The claim on $\|(\mathbf{E}_{1,1}^{(t+1)})^-\|_s$ follows from (47) and the condition (41).

For $\|(\mathbf{E}_{1,1}^{(t+1)})^+\|_s$, by induction (48) becomes

$$\|(\mathbf{E}_{1,1}^{(t+1)})^+\|_s \leq (1-\eta)\|(\mathbf{E}_{1,1}^{(t)})^+\|_s + r\eta \frac{R}{(1-\ell)^2}\gamma\ell + \eta h_2 \leq \frac{rR}{(1-\ell)^2}\gamma\ell + h_2.$$

Now we consider $\|(\mathbf{E}_{2,1}^{(t+1)})\|_s$. By Lemma 21,

$$\|(\mathbf{E}_{2,1}^{(t+1)})\|_s \leq (1-\eta)\|(\mathbf{E}_{2,1}^{(t)})\|_s$$

$$+ r\eta \frac{R\beta^2\ell^2}{(1-\ell)^2(1-\ell-\beta\ell)}$$

$$+ \frac{12C_1(C_1+1)}{n^2(\alpha-\rho)(1-\ell-\beta\ell)}\left(\frac{1}{1-\ell-\beta\ell}+n\rho\right)r\eta$$

$$= (1-\eta)\|(\mathbf{E}_{2,1}^{(t)})\|_s + \eta h_2 \tag{49}$$

$$\leq \gamma\ell$$

where the last line follows by condition (42) and induction.

Finally, clearly we have $\left\|\mathbf{E}_{1,2}^{(t+1)}\right\|_{\mathrm{s}} \leq \ell$ and $\left\|\mathbf{E}_{2,2}^{(t+1)}\right\|_{\mathrm{s}} \leq \ell$, since they are not updated.

(3) Note that (47) (48) hold for all iterations up to $t+1$. Then by Lemma 28, we have

$$\|(\mathbf{E}_{1,1}^{(t+1)})^-\|_s + \|(\mathbf{E}_{1,1}^{(t+1)})^+\|_s$$

$$\leq \max\left\{\|(\mathbf{E}_{1,1}^{(0)})^-\|_s + \|(\mathbf{E}_{1,1}^{(0)})^+\|_s, \|(\mathbf{E}_{1,1}^{(0)})^+\|_s + h_1, h_2 + \left(\frac{rR}{(1-\ell)^2}+1\right)\|(\mathbf{E}_{1,1}^{(0)})^-\|_s, h_2 + \left(\frac{rR}{(1-\ell)^2}+1\right)h_1\right\}.$$

Since $h_1 \leq \ell$ and $h_2 \leq \ell$ by (41)(42), and $\|(\mathbf{E}_{1,1}^{(0)})^-\|_s + \|(\mathbf{E}_{1,1}^{(0)})^+\|_s \leq \ell$ by assumption, we have

$$\|(\mathbf{E}_{1,1}^{(t+1)})^-\|_s + \|(\mathbf{E}_{1,1}^{(t+1)})^+\|_s \leq \max\left\{\ell + h_1, h_2 + \left(\frac{rR}{(1-\ell)^2}+1\right)\ell\right\}. \tag{50}$$

Then we have by condition (43),

$$\|(\mathbf{E}_{1,1}^{(t+1)})^-\|_s + \|(\mathbf{E}_{1,1}^{(t+1)})^+\|_s \leq (\beta-1)\ell, \quad \|\mathbf{E}^{(t+1)}\|_s \leq \beta\ell.$$

36

(4) Finally, we consider the noise. We first consider the adversarial noise. Set the sample size N to be large enough, so that by Lemma 14, we have

$$\left|\widetilde{\mathbf{N}}_{i,j}^{(t)}\right| \le \frac{4C_\nu C_1}{(1-2\ell)^2 n(\alpha-\rho)} + \left|[\widetilde{\mathbf{N}}_s^{(t)}]_{i,j}\right| \le \frac{8C_\nu C_1}{n(\alpha-\rho)}$$

and thus

$$\left\|\mathbf{N}^{(t+1)}\right\|_\infty \le (1-\eta)\left\|\mathbf{N}^{(t)}\right\|_\infty + \eta\frac{8rC_\nu C_1}{\alpha-\rho}. \tag{51}$$

Then for any $t \ge 0$,

$$\left\|\mathbf{N}^{(t)}\right\|_\infty \le \left\|\mathbf{N}^{(0)}\right\|_\infty + \frac{8rC_\nu C_1}{\alpha-\rho} \le U + \frac{8rC_\nu C_1}{\alpha-\rho} \le 2U + U_a$$

where the last inequality is by the definition of U_a. On the other hand, by Lemma 16, we have

$$\|\xi^{(t)}\|_\infty \le 3\|(\mathbf{A}^*)^\dagger\|_\infty(\|\mathbf{N}^{(t)}\|_\infty + C_\nu)$$
$$\le 3\|(\mathbf{A}^*)^\dagger\|_\infty\left(2U + \frac{8rC_\nu C_1}{\alpha-\rho} + C_\nu\right)$$
$$\le \rho$$

where the last inequality is due to condition (45).

We now consider the unbiased noise, where the proof is similar. Set the sample size N to be large enough, so that by Lemma 14, we have

$$\left|\widetilde{\mathbf{N}}_{i,j}^{(t)}\right| \le \frac{2C_1 C_\nu \rho'(1+\|\mathbf{A}^\dagger\mathbf{N}^{(t)}\|_\infty)}{(1-2\ell)n(\alpha-\rho')} + \left|[\mathbf{N}_s]_{i,j}\right|$$
$$\le \frac{8C_1 C_\nu^2\|(\mathbf{A}^*)^\dagger\|_\infty}{(1-2\ell)n(\alpha-2C_\nu\|(\mathbf{A}^*)^\dagger\|_\infty)} + \left|[\mathbf{N}_s]_{i,j}\right|$$
$$\le \frac{10C_1 C_\nu^2\|(\mathbf{A}^*)^\dagger\|_\infty}{(1-2\ell)n(\alpha-2C_\nu\|(\mathbf{A}^*)^\dagger\|_\infty)},$$

and thus

$$\left\|\mathbf{N}_S^{(t+1)}\right\|_\infty \le (1-\eta)\left\|\mathbf{N}_S^{(t)}\right\|_\infty + \eta\frac{10rC_1 C_\nu^2\|(\mathbf{A}^*)^\dagger\|_\infty}{(1-2\ell)(\alpha-2C_\nu\|(\mathbf{A}^*)^\dagger\|_\infty)}. \tag{52}$$

Then for any $t \ge 0$,

$$\left\|\mathbf{N}_S^{(t)}\right\|_\infty \le \left\|\mathbf{N}_S^{(0)}\right\|_\infty + \frac{10rC_1 C_\nu^2\|(\mathbf{A}^*)^\dagger\|_\infty}{(1-2\ell)(\alpha-2C_\nu\|(\mathbf{A}^*)^\dagger\|_\infty)} \le 2U + \frac{10rC_1 C_\nu^2\|(\mathbf{A}^*)^\dagger\|_\infty}{(1-2\ell)(\alpha-2C_\nu\|(\mathbf{A}^*)^\dagger\|_\infty)} \le 2U + U_n$$

where the last inequality is by the definition of U_n. This completes the proof for the claims.

Now, after proving the claims, we are ready to prove the last statement of the lemma. First, by (47) and Lemma 29, we have that after $\frac{\ln(\epsilon/(\gamma\ell))}{\ln(1-\eta)}$ iterations,

$$\|(\mathbf{E}_{1,1}^{(t)})^-\|_s \le \epsilon + h_1.$$

Now (48) becomes

$$\|(\mathbf{E}_{1,1}^{(t+1)})^+\|_s \le (1-\eta)\|(\mathbf{E}_{1,1}^{(t)})^+\|_s + r\eta\frac{R}{(1-\ell)^2}(\epsilon+h_1) + \eta h_2 \tag{53}$$

After an additional $\frac{\ln(\epsilon/(\gamma\ell))}{\ln(1-\eta)}$ iterations, by Lemma 29,

$$\|(\mathbf{E}_{1,1}^{(t)})^+\|_s \le \frac{rR}{(1-\ell)^2}(\epsilon+h_1) + h_2 + \epsilon$$

37

Similarly, Lemma 29 and (49), after $\frac{\ln(\epsilon/(\gamma\ell))}{\ln(1-\eta)}$ iterations,

$$\|(\mathbf{E}_{2,1}^{(t)})\|_s \leq \epsilon + h_2.$$

$\left\|\mathbf{N}_{-S}^{(t)}\right\|_\infty$ does not change since it is not updated. Now consider $\left\|\mathbf{N}_{S}^{(t)}\right\|_\infty$.

For the adversarial noise, by (51) and Lemma 29, after $\frac{\ln(\epsilon'/U)}{\ln(1-\eta)}$ iterations,

$$\left\|\mathbf{N}_{S}^{(t)}\right\|_\infty \leq \epsilon' + \frac{8rC_\nu C_1}{\alpha - \rho} \leq (1-\epsilon)U$$

where the last inequality is due to condition (45).

For the unbiased noise, by (52) and Lemma 29, after $\frac{\ln(\epsilon'/U)}{\ln(1-\eta)}$ iterations,

$$\left\|\mathbf{N}_{S}^{(t)}\right\|_\infty \leq \epsilon' + \frac{10rC_1C_\nu^2\|(\mathbf{A}^*)^\dagger\|_\infty}{(1-2\ell)(\alpha - 2C_\nu\|(\mathbf{A}^*)^\dagger\|_\infty)} \leq (1-\epsilon)U$$

where the last inequality is due to condition (46).

This completes the proof. □

C.2 Equilibration: Rescale

The input of of Algorithm 3 can be written as $\mathbf{A}^{(0)} = \mathbf{A}^*(\boldsymbol{\Sigma}^{(0)} + \mathbf{E}^{(0)}) + \mathbf{N}^{(0)}$. The output $\hat{\mathbf{A}}$ can be written as $\hat{\mathbf{A}} = (\mathbf{A}^*\mathbf{D})(\hat{\boldsymbol{\Sigma}} + \hat{\mathbf{E}}) + \hat{\mathbf{N}}$ where $\hat{\boldsymbol{\Sigma}}$ is diagonal, and $\hat{\mathbf{E}}$ is off diagonal, and $\mathbf{D}$ is a diagonal matrix with $\mathbf{D}_{i,i} = \frac{1}{1-\epsilon}$ for $i \in S$ and the rest being 1. Recall that for a matrix $\mathbf{M}$, let $\mathbf{M}_{1,1}$ denote the submatrix of $\mathbf{M}$ indexed by $S \times S$, and define $\mathbf{M}_{1,2}, \mathbf{M}_{2,1}$ and $\mathbf{M}_{2,2}$ accordingly. Also recall that $\mathbf{M}_S$ denote the submatrix of $\mathbf{M}$ formed by columns indexed by S, and let $\mathbf{M}_{-S}$ denote the submatrix formed by the other columns.

Lemma 23 (Main: Rescale). *Let* $\mathbf{A}^{(0)} = \mathbf{A}^*(\boldsymbol{\Sigma}^{(0)} + \mathbf{E}^{(0)}) + \mathbf{N}^{(0)}$ *satisfies the condition in Lemma 19 and* ϵ *be defined as in Lemma 19. Then the output of Algorithm 3 is* $\hat{\mathbf{A}} = (\mathbf{A}^*\mathbf{D})(\hat{\boldsymbol{\Sigma}} + \hat{\mathbf{E}}) + \hat{\mathbf{N}}$ *satisfying*

$$(1-\ell)\mathbf{I} \preceq \hat{\boldsymbol{\Sigma}}, \quad \|\hat{\mathbf{E}}_{1,1}\|_s \leq (\gamma-1)\ell, \quad \|\hat{\mathbf{E}}_{2,1}\|_s \leq (\gamma-1)\ell, \quad \|(\hat{\mathbf{E}}_{1,2}, \hat{\mathbf{E}}_{2,2})\|_s \leq \ell, \quad \|\hat{\mathbf{N}}_S\|_\infty \leq U, \quad \|\hat{\mathbf{N}}_{-S}\|_\infty \leq U.$$

Moreover, $\hat{\mathbf{E}}_{1,2} \geq 0$ *and* $\hat{\mathbf{E}}_{2,2} \geq 0$ *entry-wise.*

Proof of Lemma 23. Note that $\tilde{\mathbf{A}} = \mathbf{A}^*(\widetilde{\boldsymbol{\Sigma}} + \widetilde{\mathbf{E}}) + \widetilde{\mathbf{N}}$ for a diagonal matrix $\widetilde{\boldsymbol{\Sigma}}$, off-diagonal matrix $\widetilde{\mathbf{E}}$ and error matrix $\widetilde{\mathbf{N}}$. By lemma 19, we have $\widetilde{\boldsymbol{\Sigma}} \succeq (1-\ell)\mathbf{I}$, error matrix $\|\widetilde{\mathbf{N}}_S\|_\infty \leq (1-\epsilon)U$ and

$$\|\widetilde{\mathbf{E}}_{1,1}\|_s \leq (\gamma-1)\ell, \quad \|\widetilde{\mathbf{E}}_{2,1}\|_s \leq (1-\epsilon)(\gamma-1)\ell, \quad \|(\widetilde{\mathbf{E}}_{1,2}; \widetilde{\mathbf{E}}_{2,2})\|_s \leq \ell$$

and $\widetilde{\mathbf{E}}_{1,2} \geq 0$ and $\widetilde{\mathbf{E}}_{2,2} \geq 0$ entry-wise.

Therefore, by the rescaling rule:

$$\hat{\mathbf{A}} = \tilde{\mathbf{A}}\mathbf{D} = \mathbf{A}^*(\widetilde{\boldsymbol{\Sigma}} + \widetilde{\mathbf{E}})\mathbf{D} + \widetilde{\mathbf{N}}\mathbf{D}$$
$$= \mathbf{A}^*\mathbf{D}(\widetilde{\boldsymbol{\Sigma}} + \mathbf{D}^{-1}\widetilde{\mathbf{E}}\mathbf{D}) + \widetilde{\mathbf{N}}\mathbf{D}.$$

Therefore, $\hat{\boldsymbol{\Sigma}} = \widetilde{\boldsymbol{\Sigma}} \succeq (1-\ell)\mathbf{I}$, $\|\hat{\mathbf{N}}_S\|_\infty \leq \frac{1}{1-\epsilon}\|\widetilde{\mathbf{N}}_S\|_\infty \leq U$. $\|\hat{\mathbf{N}}_{-S}\|_\infty = \|\widetilde{\mathbf{N}}_{-S}\|_\infty \leq U$ since it is not updated.

For the $\hat{\mathbf{E}}$ term, denote $\mathbf{D}_1 = \mathsf{Diag}\left(\frac{1}{1-\epsilon}, \ldots, \frac{1}{1-\epsilon}\right) \in \mathbb{R}^{s \times s}$. We know that

$$\hat{\mathbf{E}}_{1,1} = \mathbf{D}_1^{-1}\widetilde{\mathbf{E}}_{1,1}\mathbf{D}_1 = \widetilde{\mathbf{E}}_{1,1}$$
$$\hat{\mathbf{E}}_{2,1} = \widetilde{\mathbf{E}}_{2,1}\mathbf{D}_1 = \frac{1}{1-\epsilon}\widetilde{\mathbf{E}}_{2,1}$$
$$\hat{\mathbf{E}}_{1,2} = \mathbf{D}_1^{-1}\widetilde{\mathbf{E}}_{1,2} = (1-\epsilon)\widetilde{\mathbf{E}}_{1,2}$$
$$\hat{\mathbf{E}}_{2,2} = \widetilde{\mathbf{E}}_{2,2}.$$

38

This leads to

$$\|\hat{\mathbf{E}}_{1,1}\|_s \le (\gamma - 1)\ell, \quad \|\hat{\mathbf{E}}_{2,1}\|_s \le (\gamma - 1)\ell, \quad \|(\hat{\mathbf{E}}_{1,2}, \hat{\mathbf{E}}_{2,2})\|_s \le \ell,$$

with $\hat{\mathbf{E}}_{1,2}, \hat{\mathbf{E}}_{2,2} \ge 0$. This completes the proof. $\qquad\square$

C.3   Equilibration: Main algorithm

Lemma 24 (Main: Equilibration). *Suppose the conditions in Lemma 23 each time Algorithm 3. Additionally, there exists constant $0 < b < 1$, $\kappa > 1$ and $u > 1$ such that $b\kappa > 1$ such that the initial $\lambda \ge \max_{i \in [n]} \mathbb{E}[(x_i^*)^2]/b$, and the initial $\boldsymbol{\Sigma} \preceq u\mathbf{I}$. Furthermore, for any $\lambda \ge \min_{i \in [n]} \mathbb{E}[(x_i^*)^2]/\kappa$,*

$$\left(\frac{1}{1-\ell} + h_6\right)^2 b\lambda + h_5^2 b\kappa\lambda + h_3 \le \left(1 - \frac{1}{100}\right)\lambda, \tag{54}$$

$$\left(\frac{1}{u} - h_6\right)^2 (1-\epsilon)b\kappa\lambda - h_5^2 b\kappa\lambda - h_4 \ge \left(1 + \frac{1}{100}\right)\lambda, \quad \frac{1}{u} > h_6 \tag{55}$$

$$h_3 \le \frac{1}{200} \min_{i \in [n]} \mathbb{E}[(x_i^*)^2], \tag{56}$$

$$h_4 \le \frac{1}{200} \min_{i \in [n]} \mathbb{E}[(x_i^*)^2], \tag{57}$$

where

$$h_3 = \frac{C_1^2}{n^2} h_5 \left(h_5 + \frac{2}{1-\ell}\right),$$

$$h_4 = \frac{C_1^2}{n^2} h_5 \left(h_5 + \frac{2}{1-\ell}\right) + \frac{2(\alpha + \rho)C_1}{n(1-\ell)},$$

$$h_5 = \frac{(\gamma + 1)\ell(1 - (\gamma + 1)\ell)}{(1-\ell)^2(1 - (\gamma + 2)\ell)},$$

$$h_6 = \frac{(\gamma + 1)\ell^2}{(1-\ell)^2(1 - (\gamma + 2)\ell)}.$$

Finally, set $N = poly(1/\min_{i \in [n]} \mathbb{E}[(x_i^*)^2], n, 1/\delta)$ large enough.

Then with probability at least $1 - \delta$, the following hold. During the execution of the algorithm, for any $j \in S$,

$$\left(\left(\frac{1}{u} - h_6\right)^2 - \kappa h_5^2 - \frac{1}{100}\right) \frac{\mathbb{E}[(x_j^*)^2]}{(\mathbf{D}_{j,j})^2} \le m_j \le \left(\left(\frac{1}{1-\ell} + h_6\right)^2 + \kappa h_5^2 + \frac{1}{100}\right) \frac{\mathbb{E}[(x_j^*)^2]}{(\mathbf{D}_{j,j})^2}.$$

Furthermore, the output of Algorithm 4 is $\mathbf{A} = \mathbf{A}^* \mathbf{D}(\boldsymbol{\Sigma} + \mathbf{E}) + \mathbf{N}$ where $\boldsymbol{\Sigma}$ is diagonal and $(1-\ell)\mathbf{I} \preceq \boldsymbol{\Sigma}$, $\mathbf{E}$ is off diagonal and $\|\mathbf{E}\|_s \le \gamma\ell$, $\mathbf{N}$ satisfies $\|\mathbf{N}\|_\infty \le 2U$, and

$$\frac{\max_{i \in [n]} \frac{1}{\mathbf{D}_{i,i}^2} \mathbb{E}[(x_i^*)^2]}{\min_{j \in [n]} \frac{1}{\mathbf{D}_{j,j}^2} \mathbb{E}[(x_j^*)^2]} \le \kappa.$$

Proof of Lemma 24. We prove the lemma by induction. For notational convenience, let us introduce a counter (p) denoting the number of times the inner while cycle has been executed, and denote $\mathbf{A}$ as $\mathbf{A}^{(p)}$. Recall that for a matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$ and index set $S \subseteq [n]$, let $\mathbf{M}_{1,1}$ denote the submatrix indexed by $S \times S$, and $\mathbf{M}_{1,2}, \mathbf{M}_{2,1}$ and $\mathbf{M}_{2,2}$ are defined accordingly. Also, let $\mathbf{M}_S$ denote the submatrix formed by the columns indexed by S, and $\mathbf{M}_{-S}$ the submatrix formed by the other columns.

Our inductive claims are as follows. At the beginning of each inner while cycle,

$$\mathbf{A}^{(p)} = \mathbf{A}^* \mathbf{D}^{(p)} \left(\boldsymbol{\Sigma}^{(p)} + \mathbf{E}^{(p)}\right) + \mathbf{N}^{(p)}$$

where $\mathbf{D}^{(p)}$ and $\boldsymbol{\Sigma}^{(p)}$ are diagonal, $\mathbf{E}^{(p)}$ are off diagonal satisfying

39

(1) $(1 - \ell)\mathbf{I} \preceq \mathbf{\Sigma}^{(p)}$,

(2) $\mathbf{E}_{1,2}^{(p)} \geq 0$ and $\mathbf{E}_{2,2}^{(p)} \geq 0$ entry-wise and

$$\left\|\mathbf{E}_{1,1}^{(p)}\right\|_s \leq \gamma\ell,$$

$$\left\|\mathbf{E}_{2,1}^{(p)}\right\|_s \leq \gamma\ell,$$

$$\left\|(\mathbf{E}_{1,2}^{(p)}; \mathbf{E}_{2,2}^{(p)})\right\|_s \leq \ell,$$

(3) $\mathbf{N}_{-S}^{(p)} \leq U$ and $\mathbf{N}_S^{(p)} \leq 2U$,

(4) We have

    (a) When $\mathbb{E}[(x_j^*)^2] < b\lambda^{(p)}, j \notin S$, then $m_j \leq \lambda^{(p)}$,

    (b) When $\mathbb{E}[(x_j^*)^2] \geq (1 - \epsilon)b\kappa\lambda^{(p)}, j \notin S$, then $m_j > \lambda^{(p)}$,

and consequently,

    (c) $\forall i \in S, b\lambda^{(p)} \leq \frac{\mathbb{E}[(x_i^*)^2]}{\left(\mathbf{D}_{i,i}^{(p)}\right)^2}$,

    (d) $\forall i \in [n], \frac{\mathbb{E}[(x_i^*)^2]}{\left(\mathbf{D}_{i,i}^{(p)}\right)^2} \leq b\kappa\lambda^{(p)}$.

The claims are trivially true at initialization, so we proceed to the induction. Assume the claim is true at time p, we proceed to show it is true at time $p + 1$.

First, consider (1), (2) and (3). By Lemma 23, after applying the rescaling algorithm, $(1-\ell)\mathbf{I} \preceq \mathbf{\Sigma}^{(p)}$ and

$$\|\mathbf{E}_{1,1}^{(p)}\|_s \leq (\gamma - 1)\ell, \quad \|\mathbf{E}_{2,1}^{(p)}\|_s \leq (\gamma - 1)\ell, \quad \|(\mathbf{E}_{1,2}^{(p)}, \mathbf{E}_{2,2}^{(p)})\|_s \leq \ell, \quad \|\mathbf{N}_S^{(p)}\|_\infty \leq U, \quad \|\mathbf{N}_{-S}^{(p)}\|_\infty \leq U.$$

Moreover, $\mathbf{E}_{1,2}^{(p)} \geq 0$ and $\mathbf{E}_{2,2}^{(p)} \geq 0$ entry-wise. Observe that when moving from time p to $p + 1$, potentially the algorithm includes new elements in S. Then

$$\|\mathbf{E}_{1,1}^{(p+1)}\|_s \leq \|\mathbf{E}_{1,1}^{(p)}\|_s + \max\{\|\mathbf{E}_{2,1}^{(p)}\|_s, \|\mathbf{E}_{1,2}^{(p)}\|_s\} \leq (\gamma - 1)\ell + \ell = \gamma\ell$$

Where the last inequality used the fact that $\gamma < 2$. Similarly,

$$\|\mathbf{E}_{2,1}^{(p+1)}\|_s \leq \|\mathbf{E}_{2,1}^{(p)}\|_s + \|\mathbf{E}_{2,2}^{(p)}\|_s \leq (\gamma - 1)\ell + \ell = \gamma\ell.$$

Also, $\|(\mathbf{E}_{1,2}^{(p+1)}, \mathbf{E}_{2,2}^{(p+1)})\|_s \leq \|(\mathbf{E}_{1,2}^{(p)}, \mathbf{E}_{2,2}^{(p)})\|_s \leq \ell$, and $(\mathbf{E}_{1,2}^{(p+1)}, \mathbf{E}_{2,2}^{(p+1)}) \geq 0$ entry-wise. Furthermore, $\|\mathbf{N}_{-S}^{(p+1)}\|_\infty \leq \|\mathbf{N}_{-S}^{(p)}\|_\infty \leq U$ and

$$\|\mathbf{N}_S^{(p+1)}\|_\infty \leq \|\mathbf{N}_S^{(p)}\|_\infty + \|\mathbf{N}_{-S}^{(p)}\|_\infty \leq 2U.$$

Hence, (1), (2) and (3) are also true at time $(p + 1)$.

Finally, we proceed to (4). Since (a)(b) are true at time p, (c)(d) are true at time $p+1$. [6] Furthermore, when $\lambda \leq \min_{i \in [n]} \mathbb{E}[(x_i^*)^2]/\kappa$, it is guaranteed that all $[n] \subseteq S$, so we only need to prove that when $\lambda \geq \min_{i \in [n]} \mathbb{E}[(x_i^*)^2]/\kappa$, (a)(b) are also true at time $p + 1$.

To prove (a)(b) are true at time $p + 1$, we will use Lemma 25. Note that since $\mathbf{A}$ has been scaled, so $\mathbf{A}^*\mathbf{D}$ should be regarded as the ground truth matrix $\mathbf{A}^*$ in Lemma 25. We first make sure its assumption is satisfied. First, $\|\mathbf{N}\|_\infty \leq 3U$ and $\left\|(\mathbf{A}^*\mathbf{D})^\dagger\right\|_\infty \leq \left\|(\mathbf{A}^*)^\dagger\right\|_\infty$. By Lemma 16 and condition (44), the assumption in Lemma 25 is satisfied.

We are now ready to prove (a). By Lemma 25,

$$\mathbb{E}[x_j^2] \leq \left(\mathbf{\Sigma}_{j,j}^{-1} + |\mathbf{V}_{j,j}|\right)^2 \frac{\mathbb{E}[(x_j^*)^2]}{\mathbf{D}_{j,j}^{(p+1)}} + \|[\mathbf{V}]^j\|_2^2 \max_{k \in [n]} \frac{\mathbb{E}[(x_k^*)^2]}{\mathbf{D}_{k,k}^{(p+1)}} + \|[\mathbf{V}]^j\|_1 \left(\|[\mathbf{V}]^j\|_1 + 2\mathbf{\Sigma}_{j,j}^{-1}\right) \frac{C_1^2}{n^2}.$$

---

[6]Note that in (b), the factor $(1 - \epsilon)$ is needed to ensure (d) is true at time $p + 1$.

40

By Lemma 15, $|\mathbf{V}_{j,j}| \le h_6$, $\|[\mathbf{V}]^j\|_2^2 \le \|[\mathbf{V}]^j\|_1^2 \le h_5^2$, so

$$\mathbb{E}[x_j^2] \le \left(\frac{1}{1-\ell} + h_6\right)^2 \frac{\mathbb{E}[(x_j^*)^2]}{(\mathbf{D}_{j,j}^{(p+1)})^2} + h_5^2 \max_{k \in [n]} \frac{\mathbb{E}[(x_k^*)^2]}{(\mathbf{D}_{k,k}^{(p+1)})^2} + h_3.$$

By (d), $\max_{k \in [n]} \frac{\mathbb{E}[(x_k^*)^2]}{(\mathbf{D}_{k,k}^{(p+1)})^2} \le b\kappa\lambda$, so for any $j \notin S$ with $\mathbb{E}[(x_j^*)^2] = \frac{\mathbb{E}[(x_j^*)^2]}{(\mathbf{D}_{j,j}^{(p+1)})^2} < b\lambda$, we have

$$\mathbb{E}[x_j^2] \le \left(\frac{1}{1-\ell} + h_6\right)^2 b\lambda + h_5^2 b\kappa\lambda + h_3.$$

By using large enough sample, with high probability, the empirical estimation

$$\hat{\mathbb{E}}[x_j^2] \le \mathbb{E}[x_j^2] + \frac{1}{100}\lambda \le \lambda$$

where the last step is by condition (54).

As for (b), by Lemma 25 we have

$$\mathbb{E}[x_j^2] \ge \left(\mathbf{\Sigma}_{j,j}^{-1} - |\mathbf{V}_{j,j}|\right)^2 \frac{\mathbb{E}[(x_j^*)^2]}{(\mathbf{D}_{j,j}^{(p+1)})^2} - \|[\mathbf{V}]^j\|_2^2 \max_{k \in [n]} \frac{\mathbb{E}[(x_k^*)^2]}{(\mathbf{D}_{k,k}^{(p+1)})^2} - \left(\frac{C_1^2}{n^2}\|[\mathbf{V}]^j\|_1(\|[\mathbf{V}]^j\|_1 + 2\mathbf{\Sigma}_{j,j}^{-1}) + \frac{2(\alpha+\rho)C_1}{n}\mathbf{\Sigma}_{j,j}^{-1}\right)$$

$$\ge \left(\frac{1}{u} - h_6\right)^2 \frac{\mathbb{E}[(x_j^*)^2]}{(\mathbf{D}_{j,j}^{(p+1)})^2} - h_5^2 \max_{k \in [n]} \frac{\mathbb{E}[(x_k^*)^2]}{(\mathbf{D}_{k,k}^{(p+1)})^2} - h_4.$$

The last step uses that $\mathbf{\Sigma}_{j,j}^{-1} \le u$, which is by the initial condition assumed and that it is not updated for $j \notin S$. Putting in the bound that $\frac{\mathbb{E}[(x_k^*)^2]}{(\mathbf{D}_{k,k}^{(p+1)})^2} \le b\kappa\lambda$, then for any $j \notin S$ with $\mathbb{E}[(x_j^*)^2] = \frac{\mathbb{E}[(x_j^*)^2]}{(\mathbf{D}_{j,j}^{(p+1)})^2} \ge (1-\epsilon)b\kappa\lambda$, we have

$$\mathbb{E}[x_j^2] \ge \left(\frac{1}{u} - h_6\right)^2 (1-\epsilon)b\kappa\lambda - h_5^2 b\kappa\lambda - h_4.$$

Again, use large enough sample to ensure that with high probability

$$\tilde{\mathbb{E}}[x_j^2] \ge \mathbb{E}[x_j^2] - \frac{1}{100}\lambda \ge \lambda$$

where the last step follows from condition (55). This completes the proof of the induction.

We now prove the statements of the lemma. The statement about the output follows from the above claims. What is left is to prove that $m_j (j \in S)$ approximates $\mathbb{E}[(x_j^*)^2]$ well. Since m_j for $j \in S$ is updated along with $\mathbf{D}_{j,j}$, we only need to check the right after adding j to S, the statement holds. Suppose the time point is p, we have

$$\mathbb{E}[x_j^2] \le \left(\frac{1}{1-\ell} + h_6\right)^2 \frac{\mathbb{E}[(x_j^*)^2]}{(\mathbf{D}_{j,j}^{(p)})^2} + h_5^2 \max_{k \in [n]} \frac{\mathbb{E}[(x_k^*)^2]}{(\mathbf{D}_{k,k}^{(p)})^2} + h_3.$$

Since j is in S, by the claims (c)(d) we have

$$\max_{k \in [n]} \frac{\mathbb{E}[(x_k^*)^2]}{(\mathbf{D}_{k,k}^{(p)})^2} \le b\kappa\lambda^{(p)} \le \kappa\frac{\mathbb{E}[(x_j^*)^2]}{(\mathbf{D}_{j,j}^{(p)})^2}.$$

Since N is large enough so that

$$\mathbb{E}[x_j^2] \le \mathbb{E}[(x_j^*)^2]\left(1 + \frac{1}{200}\right).$$

Combined these with the condition (56), we have

$$m_j \le \left(\left(\frac{1}{1-\ell} + h_6\right)^2 + \kappa h_5^2 + \frac{1}{100}\right)\frac{\mathbb{E}[(x_j^*)^2]}{(\mathbf{D}_{j,j})^2}.$$

The upper bound on m_j can be bounded similarly. This completes the proof of the lemma. $\qquad\square$

41

The following is the lemma used in the proof of Lemma 24.

Lemma 25 (Estimate of feature weight). *Suppose $|\xi_i| \le \rho < \alpha$ for any example and every $i \in [n]$, and suppose $\mathbf{\Sigma} \succeq \frac{1}{2}\mathbf{I}$. Then*

$$\mathbb{E}[x_i^2] \ge \left(\mathbf{\Sigma}_{i,i}^{-1} - |\mathbf{V}_{i,i}|\right)^2 \mathbb{E}[(x_i^*)^2] - \|[\mathbf{V}]^i\|_2^2 \max_{j \in [n]} \mathbb{E}[(x_j^*)^2]$$

$$- \left(\frac{C_1^2}{n^2}\|[\mathbf{V}]^i\|_1(\|[\mathbf{V}]^i\|_1 + 2\mathbf{\Sigma}_{i,i}^{-1}) + \frac{2(\alpha+\rho)C_1}{n}\mathbf{\Sigma}_{i,i}^{-1}\right)$$

$$\mathbb{E}[x_i^2] \le \left(\mathbf{\Sigma}_{i,i}^{-1} + |\mathbf{V}_{i,i}|\right)^2 \mathbb{E}[(x_i^*)^2] + \|[\mathbf{V}]^i\|_2^2 \max_{j \in [n]} \mathbb{E}[(x_j^*)^2] + \|[\mathbf{V}]^i\|_1 \left(\|[\mathbf{V}]^i\|_1 + 2\mathbf{\Sigma}_{i,i}^{-1}\right) \frac{C_1^2}{n^2}.$$

Proof of Lemma 25. By the decoding rule,

$$x_i = \left[\phi_\alpha(\mathbf{A}^\dagger[\mathbf{A}^* x^* + \nu])\right]_i$$
$$= \left[\phi_\alpha\left((\mathbf{\Sigma}^{-1} + \mathbf{V})\, x^* + \xi\right)\right]_i.$$

Let $[\mathbf{V}]^i = v$ and $\Sigma_{i,i}^{-1} = \sigma$, then we can rewrite above as

$$x_i = \phi_\alpha(\sigma x_i^* + \langle v, x^* \rangle + \xi_i)$$

which implies that

$$\sigma x_i^* + \langle v, x^* \rangle - \rho - \alpha \le x_i \le |\sigma x_i^* + \langle v, x^* \rangle|. \tag{58}$$

First, consider the lower bound.

$$\mathbb{E}[x_i^2] \ge \mathbb{E}\left[(\sigma x_i^* + \langle v, x^* \rangle - \rho - \alpha)\, \phi_\alpha(\sigma x_i^* + \langle v, x^* \rangle + \xi_i)\right]$$

The following simple lemma is useful.

Claim 26. *Let χ be a variable such that $|\chi| \le \alpha$, then for every $w \in \mathbb{R}^n$, $k \in [n]$,*

$$\mathbb{E}[x_k^* \phi_\alpha(\langle w, x^* \rangle + \chi)] \le |w_k|\mathbb{E}[(x_k^*)^2] + \frac{C_1^2}{n^2}\sum_{j \ne k} |w_j| \tag{59}$$

$$\le |w_k|\mathbb{E}[(x_k^*)^2] + \frac{C_1^2}{n^2}\|w\|_1. \tag{60}$$

Proof. The proof is a direct observation that when $|\chi| < \alpha$,

$$\phi_\alpha(\langle w, x \rangle + \chi) \le |\langle w, x \rangle| \le \langle |w|, x \rangle$$

where $|w|$ is the entry wise absolute value. □

Therefore, we can obtain the following bounds.

(1). By (16) in Lemma 12, we have

$$\mathbb{E}[x_i^* \phi_\alpha(\sigma x_i^* + \langle v, x^* \rangle + \xi_i)] \ge \mathbf{\Sigma}_{i,i}^{-1}\mathbb{E}\left[(x_i^*)^2\right] - \frac{(\alpha+\rho)C_1}{n} - \mathbb{E}\left[(x_i^*)^2\right]|\mathbf{V}_{i,i}| - \frac{C_1^2}{n^2}\left\|[\mathbf{V}]^i\right\|_1,$$

(2). By (60) in the above claim,

$$\mathbb{E}[x_j^* \phi_\alpha(\sigma x_i^* + \langle v, x^* \rangle + \xi_i)] \le |v_j|\mathbb{E}[(x_j^*)^2] + \frac{C_1^2}{n^2}(\|v\|_1 + \sigma),$$

(3). By (58), for $j \ne i$,

$$\mathbb{E}[\phi_\alpha(\sigma x_i^* + \langle v, x^* \rangle + \xi_i)] \le \mathbb{E}[|\sigma x_i^* + \langle v, x^* \rangle|] \le \frac{(\sigma + \|v\|_1)C_1}{n}.$$

42

Putting together, we can obtain

$$\mathbb{E}[x_i^2] \geq \left(\boldsymbol{\Sigma}_{i,i}^{-1} - |\mathbf{V}_{i,i}|\right)^2 \mathbb{E}[(x_i^*)^2] - \|[\mathbf{V}]^i\|_2^2 \max_{j \in [n]} \mathbb{E}[(x_j^*)^2]$$

$$- \left(\frac{C_1^2}{n^2}\|[\mathbf{V}]^i\|_1(\|[\mathbf{V}]^i\|_1 + 2\boldsymbol{\Sigma}_{i,i}^{-1}) + \frac{2(\alpha+\rho)C_1}{n}\boldsymbol{\Sigma}_{i,i}^{-1}\right).$$

Second, we proceed to the upper bound. Similarly as the lower bound, by (58), we have

$$\mathbb{E}[x_i^2] \leq \mathbb{E}\left[\left((|v_i|+\sigma)x_i^* + \sum_{j\neq i}|v_j|x_j^*\right)\phi_\alpha(\sigma x_i^* + \langle v, x^*\rangle + \xi_i)\right]$$

$$= (|v_i|+\sigma)\mathbb{E}[x_i^*\phi_\alpha(\sigma x_i^* + \langle v, x^*\rangle + \xi_i)] + \sum_{j\neq i}|v_j|\mathbb{E}[x_j^*\phi_\alpha(\sigma x_i^* + \langle v, x^*\rangle + \xi_i)].$$

For the first summand, same as in (2), by (60) in the above claim we get

$$\mathbb{E}[x_i^*\phi_\alpha(\sigma x_i^* + \langle v, x^*\rangle + \xi_i)] \leq (\sigma + |v_i|)\mathbb{E}[(x_i^*)^2] + \frac{C_1^2}{n^2}\|v\|_1,$$

$$\mathbb{E}[x_j^*\phi_\alpha(\sigma x_i^* + \langle v, x^*\rangle + \xi_i)] \leq |v_j|\mathbb{E}[(x_j^*)^2] + \frac{C_1^2}{n^2}(\|v\|_1 + \sigma).$$

Therefore, we get

$$\mathbb{E}[x_i^2] \leq \left(\boldsymbol{\Sigma}_{i,i}^{-1} + |\mathbf{V}_{i,i}|\right)^2 \mathbb{E}[(x_i^*)^2] + \|[\mathbf{V}]^i\|_1(\|[\mathbf{V}]^i\|_1 + 2\boldsymbol{\Sigma}_{i,i}^{-1})\frac{C_1^2}{n^2} + \|[\mathbf{V}]^i\|_2^2 \max_{j\in[n]} \mathbb{E}[(x_j^*)^2].$$

which completes the proof. $\qquad\square$

C.4 Main theorem

Theorem 18 (Main: Equilibration). *If there exists an absolute constant $\mathcal{G}$ such that Assumption (A1)-(A3) and (N1) are satisfied with $l = 1/50$, $C_1^3 \leq \mathcal{G}c_2^2 n$, $\max\left\{C_\nu, \|\mathbf{N}^{(0)}\|_\infty\right\} \leq \frac{\mathcal{G}c_2^4}{C_1^5 n\|(\mathbf{A}^*)^\dagger\|_\infty}$, and additionally $\boldsymbol{\Sigma}^{(0)} \preceq (1-\ell)\mathbf{I}$, and $\mathbf{E} \geq 0$ entry-wise, then there exist α, η, T, λ such that for sufficiently small $\epsilon > 0$ and sufficiently large $N = \mathrm{poly}(n, m, 1/\epsilon, 1/\delta)$ the following hold with probability at least $1 - \delta$: Algorithm 4 outputs a solution $\mathbf{A} = \mathbf{A}^*\mathbf{D}(\boldsymbol{\Sigma}+\mathbf{E})+\mathbf{N}$ where $\boldsymbol{\Sigma} \succeq (1-\ell)\mathbf{I}$ is diagonal, $\|\mathbf{E}\|_s \leq \gamma\ell$ is off-diagonal, $\|\mathbf{N}\|_\infty \leq 2\|\mathbf{N}^{(0)}\|_\infty$, and $\mathbf{D}$ is diagonal and satisfies*

$$\frac{\max_{i\in[n]} \frac{1}{\mathbf{D}_{i,i}^2}\mathbb{E}[(x_i^*)^2]}{\min_{j\in[n]} \frac{1}{\mathbf{D}_{j,j}^2}\mathbb{E}[(x_j^*)^2]} \leq 2.$$

If Assumption (A1)-(A3) and (N2) are satisfied with the same parameters except $\max\left\{C_\nu, \|\mathbf{N}^{(0)}\|_\infty\right\} \leq \min\left\{\sqrt{\frac{\mathcal{G}c_2^4}{C_1^5 n}}\frac{1}{\|(\mathbf{A}^*)^\dagger\|_\infty}, \frac{\mathcal{G}c_2^2}{C_1^3\|(\mathbf{A}^*)^\dagger\|_\infty}\right\}$, then the same guarantees hold.

Proof of Theorem 18. The theorem follows from Lemma 24 (taking union bound over all the iterations and setting a proper δ), if the conditions are satisfied. So in the following, we first specify the parameters and then verify the conditions in Lemma 19 and Lemma 24.

Recall that $\ell = 1/50$. Define $u = 1 + \ell$, $\gamma = 3/2$, $\beta = 4$, $\kappa = 2$, $b = 3/4$, and let $\epsilon < 1/1000$.

Conditions in Lemma 19. For (39), we need to compute $r_i R_i$ and the the third term. Note that by the induction in Lemma 24, the m_j is an good approximation of $\mathbb{E}[(x_j^*)^2]/(\mathbf{D}_{j,j})^2$. Furthermore, when Lemma 19 is applied in Lemma 24, it is applied on the ground-truth matrix $(\mathbf{A}^*)' = \mathbf{A}^*\mathbf{D}$ and $(x_j^*)' = x_j^*/\mathbf{D}_{j,j}$, so m_j is a good approximation of $\mathbb{E}[((x_j^*)')^2]$. Then

$$r_i R_i = \frac{3\mathbb{E}[((x_i^*)')^2]}{5m_i} \leq \frac{3}{5\left(\left(\frac{1}{u} - h_6\right)^2 - \kappa h_5^2 - \frac{1}{100}\right)}.$$

43

For the third term, first note that $C_1^3 \leq \mathcal{G}c_2^2 n$, and thus $C_1^2 \leq \mathcal{G}c_2 n$ by $C_1 > c_2$. Furthermore, $r_i = O(1/m_i) = O(n/c_2)$ for $i \in S$. Plugging in the parameters, we know that the third term is less than $1/1000$ when $\mathcal{G}$ is sufficiently small. Then (39) can be verified by plugging the parameters.

Similarly, for (40), we can compute $r_i R_i$ and let $\mathcal{G}$ small enough so that the second term is less than $1/1000$, and then verify the condition.

For (41) (42) and (43), we need to bound h_1 and h_2, which in turn relies on r and rR. Since for $i \in S$, $r_i = O(n/c_2)$, $r = O(n/c_2)$. Then similar to the argument as above, $h_1 < 2/10000$ when $\mathcal{G}$ is sufficiently small. when Lemma 19 is applied in Lemma 24, it is applied on the ground-truth matrix $(\mathbf{A}^*)' = \mathbf{A}^* \mathbf{D}$ and $(x_j^*)' = x_j^*/\mathbf{D}_{j,j}$. By the induction claims there, $\max_{j \in [n]} \mathbb{E}[((x_j^*)')^2]$ differ from $\min_{j \in S} \mathbb{E}[((x_j^*)')^2]$ by a factor of at most κ, so $rR \leq \frac{3\kappa}{5}$. So the first term can be computed. The second term is less than $1/10000$ when $\mathcal{G}$ is small enough. Then h_2 can be computed. And the conditions can be verified.

Condition (44) is true since $\max\{C_\nu, \|\mathbf{N}\|_\infty\} = O(\frac{c_2^2}{C_1^3 \|(\mathbf{A}^*)^\dagger\|_\infty})$. Condition (45) is true by setting $\epsilon' < U/8$ and by $U_a < U/8$ and $U = \|\mathbf{N}\|_\infty \leq O(\frac{c_2^2}{C_1^3 \|(\mathbf{A}^*)^\dagger\|_\infty})$. Similarly, condition (45) is true by setting $\epsilon' < U/8$ and by $U_n < U/8$ and $\|\mathbf{N}\|_\infty$ is sufficiently small.

Conditions in Lemma 24. First, consider (56) and (57). As mentioned above, since $C_1^3 = O(c_2^2 n)$ and $C_1^2 = O(c_2 n)$, then h_3 and h_4 can be made sufficiently small to satisfy the conditions. (54) and (55) can be verified by plugging (56) and (57) and the assumption that $\lambda \geq \min_{i \in [n]} \mathbb{E}[(x_i^*)^2]/\kappa$.

This completes the proof. $\qquad\square$

D Auxiliary lemmas for solving recurrence

The following lemmas are used when solving recurrence in our analysis.

Lemma 27 (Coupling update rule). *Let $\{a_t\}_{t=0}^\infty, \{b_t\}_{t=0}^\infty$ be sequences of non-negative numbers such that for fixed values $h \geq 0$, $\eta \in [0, 1]$, $R > 4r > 0$:*

$$
\begin{aligned}
a_{t+1} &\leq (1-\eta)a_t + \eta r b_t + \eta h \\
b_{t+1} &\leq (1-\eta)b_t + \frac{\eta}{R}a_t + \eta h
\end{aligned}
$$

Then the following two properties holds:

1.

$$
\forall t \geq 0, a_t + b_t \leq a_0 + b_0 + \frac{Rr + 2R + 1}{R - r}h
$$

2. For all $\epsilon > 0$, when $t \geq \ln \frac{a_0 + b_0}{8\eta\epsilon}$, we have:

$$
a_t \leq \frac{R(r+1)}{R - r}h + \epsilon, \quad b_t \leq \frac{R+1}{R - r}h + \epsilon
$$

Proof of Lemma 27. Observe that the update rule is equivalent to

$$
\begin{aligned}
\left(a_{t+1} - \frac{R(r+1)}{R-r}h\right) &\leq (1-\eta)\left(a_t - \frac{R(r+1)}{R-r}h\right) + \eta r\left(b_t - \frac{R+1}{R-r}h\right) \\
\left(b_{t+1} - \frac{R+1}{R-r}h\right) &\leq (1-\eta)\left(b_t - \frac{R+1}{R-r}h\right) + \frac{\eta}{R}\left(a_t - \frac{R(r+1)}{R-r}h\right)
\end{aligned}
$$

Therefore, define $c_t = a_t - \frac{R(r+1)}{R-r}h$ and $d_t = b_t - \frac{R+1}{R-r}h$, we can rewrite above as:

$$
\begin{aligned}
c_{t+1} &\leq (1-\eta)c_t + \eta r d_t \\
d_{t+1} &\leq (1-\eta)d_t + \frac{\eta}{R}c_t
\end{aligned}
$$

44

Since we just need to upper bound c_t, d_t. without lose of generality, we can assume that

$$
\begin{aligned}
c_{t+1} &= (1 - \eta)c_t + \eta r d_t \\
d_{t+1} &= (1 - \eta)d_t + \frac{\eta}{R}c_t
\end{aligned}
$$

Which implies that

$$
\begin{aligned}
\left(c_{t+1} + \sqrt{\frac{R}{r}}d_{t+1}\right) &= \left(1 - \eta + \eta\sqrt{\frac{r}{R}}\right)\left(c_t + \sqrt{\frac{R}{r}}d_t\right) \\
\left(c_{t+1} - \sqrt{\frac{R}{r}}d_{t+1}\right) &= \left(1 - \eta - \eta\sqrt{\frac{r}{R}}\right)\left(c_t - \sqrt{\frac{R}{r}}d_t\right)
\end{aligned}
$$

Which can be simplified to

$$
\begin{aligned}
\left(c_t + \sqrt{\frac{R}{r}}d_t\right) &= \left(1 - \eta + \eta\sqrt{\frac{r}{R}}\right)^t\left(c_0 + \sqrt{\frac{R}{r}}d_0\right) \\
\left(c_t - \sqrt{\frac{R}{r}}d_t\right) &= \left(1 - \eta - \eta\sqrt{\frac{r}{R}}\right)^t\left(c_0 - \sqrt{\frac{R}{r}}d_0\right)
\end{aligned}
$$

Therefore, we can solve

$$
c_t = \frac{1}{2}\left[\left(1 - \eta + \eta\sqrt{\frac{r}{R}}\right)^t + \left(1 - \eta - \eta\sqrt{\frac{r}{R}}\right)^t\right]c_0 + \frac{1}{2}\sqrt{\frac{R}{r}}\left[\left(1 - \eta + \eta\sqrt{\frac{r}{R}}\right)^t - \left(1 - \eta - \eta\sqrt{\frac{r}{R}}\right)^t\right]d_0
$$

$$
d_t = \frac{1}{2}\sqrt{\frac{r}{R}}\left[\left(1 - \eta + \eta\sqrt{\frac{r}{R}}\right)^t - \left(1 - \eta - \eta\sqrt{\frac{r}{R}}\right)^t\right]c_0 + \frac{1}{2}\left[\left(1 - \eta + \eta\sqrt{\frac{r}{R}}\right)^t + \left(1 - \eta - \eta\sqrt{\frac{r}{R}}\right)^t\right]d_0
$$

Observe that for every $t \geq 0$, $a \geq b \geq 0$, $a^t - b^t \leq (a - b)ta^{t-1}$

Which implies:

$$
\left(1 - \eta + \eta\sqrt{\frac{r}{R}}\right)^t - \left(1 - \eta - \eta\sqrt{\frac{r}{R}}\right)^t \leq 2t\eta\sqrt{\frac{r}{R}}\left(1 - \eta + \eta\sqrt{\frac{r}{R}}\right)^{t-1}
$$

Therefore, when $c_0, d_0 \geq 0$,

$$
c_t \leq \left(1 - \eta + \eta\sqrt{\frac{r}{R}}\right)^t c_0 + t\eta\left(1 - \eta + \eta\sqrt{\frac{r}{R}}\right)^{t-1}d_0
$$

Moreover,

$$
d_t \leq \frac{r}{R}\eta\left(1 - \eta + \eta\sqrt{\frac{r}{R}}\right)^t c_0 + \left(1 - \eta + \eta\sqrt{\frac{r}{R}}\right)^t d_0
$$

Taking the optimal t, we obtain $c_t + d_t \leq c_0 + d_0$, which implies that

$$
a_t + b_t \leq a_0 + b_0 + \frac{Rr + 2R + 1}{R - r}h
$$

On the other hand, when $t \geq \ln\frac{c_0 + d_0}{8\eta\epsilon}$, $c_t, d_t \leq \epsilon$, which implies that

$$
a_t \leq \frac{R(r + 1)}{R - r}h + \epsilon, \quad b_t \leq \frac{R + 1}{R - r}h + \epsilon.
$$

$\square$

45

Lemma 28 (Simple coupling). *Let $\{a_t\}_{t=0}^{\infty}, \{b_t\}_{t=0}^{\infty}$ be sequences of non-negative numbers such that for fixed values $h_1, h_2 \geq 0$, $\eta \in [0,1]$, $r > 0$:*

$$
\begin{aligned}
a_{t+1} &\leq (1-\eta)a_t + \eta h_1 \\
b_{t+1} &\leq (1-\eta)b_t + \eta s a_t + \eta h_2
\end{aligned}
$$

Then

$$
\begin{aligned}
a_t &\leq u_a := \max\{a_0, h_1\}, \\
b_t &\leq \max\{b_0, h_2 + su_a\}.
\end{aligned}
$$

Proof. We have

$$
\begin{aligned}
(a_{t+1} - h_1) &\leq (1-\eta)(a_t - h_1) \\
(b_{t+1} - h_2) &\leq (1-\eta)(b_t - h_2) + \eta s a_t
\end{aligned}
$$

Solving the first one gives

$$a_t \leq u_a := \max\{a_0, h_1\}.$$

Then

$$(b_{t+1} - h_2) \leq (1-\eta)(b_t - h_2) + \eta s u_a$$

leads to

$$b_t \leq \max\{b_0, h_2 + su_a\}.$$

$\square$

Lemma 29 (Simple recursion). *Let $\{a_t\}_{t=0}^{\infty}$ be a sequences of non-negative numbers such that for fixed values $h \geq 0$, $\eta \in [0,1]$,*

$$a_{t+1} \leq (1-\eta)a_t + \eta h.$$

Then,

$$a_t \leq (1-\eta)^t a_0 + h,$$

and thus for $t \geq \frac{\ln(\epsilon/a_0)}{\ln(1-\eta)}$, we have

$$a_t \leq \epsilon + h.$$

Proof. We will prove by induction that $a_t \leq (1-\eta)^t a_0 + h$, which implies the statement of the lemma. The base case is trivial, so we proceed to the induction:

$$a_{t+1} \leq (1-\eta)\left((1-\eta)^t a_0 + h\right) + \eta h \leq (1-\eta)^{t+1} a_0 + h$$

as we need. $\square$

E   Detailed discussion about related work

E.1   Non-negative matrix factorization

The area of non-negative matrix factorization (henceforth NMF) has a rich empirical history, starting with the work of [LS99]. In that paper, the authors propose two algorithms based on alternating minimization, one in KL divergence norm, and the other in Frobenius norm. They observe that these heuristics work quite well in practice, but no theoretical understanding of it is provided.

On the theoretical side, [AGKM12] provide a fixed-parameter tractable algorithm for NMF: namely when if the matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{X} \in \mathbb{R}^{n \times N}$, they provide an algorithm that runs in time $(mN)^n$. This is prohibitive unless n is extremely small. Furthermore, the algorithm is based on routines from algebraic geometry, so its tolerance to noise is fairly weak. More precisely, if there are matrices $\mathbf{A}^*$, $\mathbf{X}^*$, s.t.

$$\|\mathbf{Y} - \mathbf{A}^*\mathbf{X}^*\|_F \leq \epsilon\mathbf{Y}$$

their algorithm produces matrices $\mathbf{A}, \mathbf{X}$, s.t.

$$\|\mathbf{Y} - \mathbf{A}^*\mathbf{X}^*\|_F \leq O(\epsilon^{1/2} n^{1/4})\mathbf{Y}$$

They further provide matching hardness results: namely they show there is no algorithm running in time $(mN)^{o(n)}$ unless there is a sub-exponential running time algorithm for 3-SAT. They also study the problem under separability assumptions about the feature matrix. [BGKP16] studies the problem under heavy noise setting, but also needs assumptions related to separability, such as the existence of dominant features. Also, their noise model is different from ours.

46

E.2 Topic modeling

A closely related problem is topic modeling. Topic models are a generative model for text data, using the common bag-of-words assumption. In this case, the columns of the matrix $\mathbf{A}^*$ (which have norm 1) can naturally be interpreted as topics, with the entries being the emmision probabilities of words in that topic. The vectors x^* in this case also will have norm 1, and can be viewed as distributions over topics. In this way, $y^* = A^*x^*$ can be viewed as the vector describing the emission probabilities of words in a given document: first a topic i is selected according to the distribution x^*, then a word is selected from topic i according to the distribution in column $[\mathbf{A}^*]^i$. There also exist work that assume $x_i^* \in [0, 1]$ and are independent (e.g., [ZX12]), which is closely related to our model.

The distinction from NMF is that when documents are fairly short, the empirical frequencies of the words in the document might be very far from y^*. For this reason, typicall the algorithms with provable guarantees look at the empirical covariance matrix of the words, which will concentrate to the true one when the number of documents grows, even if the documents are very short. This, however, results in algorithms that scale quadratically in the vocabulary size, which often is prohibitive in practice. Also note that since x^* is assumed to have norm 1 in topic modeling, it does not satisfy our assumption (**A2**). However, there also exist work on topic modeling [ZX12] that do not restrict x^* is assumed to have norm 1 and can satisfying our assumption.

There is a rich body of empirical work on topic models, starting from the seminal work on LDA due to [BNJ03]. Typically in empirical papers the matrices $\mathbf{A}^*$, as well as the vectors x^* are learned using *variational inference*, which can be interpreted as a kind of alternating minimization in KL divergence norm, and in the limit of infinite-length documents converges to the [LS99] updates ([AR15]).

From the theoretical side, there was a sequence of works by [AGM12],[AGH+13], as well as [AHJK13], [DRIS13], [DRIS14] and [BBK14]. All of these works are based on either spectral or combinatorial (overlapping clustering) approaches, and need certain "non-overlapping" assumptions on the topics. For example, [AGM12] and [AGH+13] assume that the topic-word matrix contains "anchor words". This means that each topic has a word which appears in that topic, and no other. [AHJK13] on the other hand work with a certain expansion assumption on the word-topic graph, which says that if one takes a subset S of topics, the number of words in the support of these topics should be at least $|S| + s_{max}$, where s_{max} is the maximum support size of any topic.

Finally, in the paper [AR15] a version of the standard variational inference updates is analyzed in the limit of infinite length documents. The algorithm there also involves a step of "decoding", which recovers correctly the support of a given sample, and a "gradient descent" step, which updates $\mathbf{A}^*$ in the direction of the gradient of a KL-divergence based objective function. However, [AR15] requires quite strong assumptions on both the warm start, and the amount of "non-overlapping" of the topics in the topic-word matrix.

E.3 ICA

In the problem of independent component analysis (henceforth ICA, also known as blind-source separation), one is given samples $y = \mathbf{A}^*x^* + \eta$, where the distribution on the samples x^* is independent for each coordinate, the 4-th moment of x_i^* is strictly smaller than that of a Gaussian and $\mathbf{A}^*$ has full rank. The classic papers [Com94] and [FJK96] solved this problem in the noiseless case, with an approach based on cumulants, and [AGMS12] solved it in another special case, when the noise η is Gaussian (albeit with an unknown covariance matrix).

Our approach is significantly more robust to noise than these prior approaches, since it can handle both adversarial noise and zero mean noise. This is extremely important in practice, as often the nature of the noise may not be precisely known, let alone exactly Gaussian.

E.4 Non-convex optimization via gradient descent

The framework of having a "decoding" for the samples, along with performing a gradient descent-like update for the model parameters has proven successful for dictionary learning as well, which is the problem of recovering the matrix $\mathbf{A}^*$ from samples $y = \mathbf{A}^*x^* + \eta$, where the matrix $\mathbf{A}^* \in \mathbb{R}^{m \times n}$ is typically long (i.e. $n \gg m$) and x^* is sparse. (No non-negativity constraints are imposed on either

47

A or x^*.) In this scenario, the columns of $\mathbf{A}^*$ are thought of as a *dictionary*, and each sample y is generated as a (noisy) sparse combination of the columns of the dictionary.

The original empirical work which proposed an algorithm like this (in fact, it suggested that the V1 layer processes visual signals in the same manner) was due to [OF97]. In fact, it is suggest that similar families of algorithms based on "decoding" and gradient-descent are neurally plausible as mechanisms for a variety of tasks like clustering, dimension-reduction, NMF, etc. ([PC15a, PC15b, HPC14, PC14])

A theoretical analysis of it came latter due to [AGMM15]. They showed that with a suitable warm start, the gradient calculated from the "decoding" of the samples is sufficiently correlated with the gradient calculated with the correct value x^*, therefore allowing them to show the algorithm converges to a matrix $\mathbf{A}$ close to the ground truth $\mathbf{A}^*$. However, the assumption made in [AGMM15] is that the columns of $\mathbf{A}^*$ are *incoherent*, which means that they have l_2 norm bounded by 1, and inner products bounded by $O(\frac{1}{\sqrt{m}})$. [7]

The above techniques are not directly applicable to our case, as we don't wish to have any assumptions on the matrix $\mathbf{A}^*$. Additionally, the incoherence assumptions on the matrix $\mathbf{A}^*$ used in [AGMM15], in the case when $\mathbf{A}^*$ needs to be non-negative and has l_1 column-wise norm would effectively imply that the columns of $\mathbf{A}^*$ have very small overlap.

---

[7]This is satisfied when the columns are random unit vectors, and intuitively says the columns of the dictionary are not too correlated.

48