

Nearly minimax-optimal rates for noisy sparse phase retrieval via early-stopped mirror descent

FAN WU AND PATRICK REBESCHINI[†]

Department of Statistics, University of Oxford, 24-29 St Giles', Oxford OX1 3LB, UK

[†]Corresponding author. Email: rebeschini@stats.ox.ac.uk

[Received on 14 June 2021; accepted on 12 September 2022]

This paper studies early-stopped mirror descent applied to noisy sparse phase retrieval, which is the problem of recovering a k -sparse signal $\mathbf{x}^* \in \mathbb{R}^n$ from a set of quadratic Gaussian measurements corrupted by sub-exponential noise. We consider the (non-convex) unregularized empirical risk minimization problem and show that early-stopped mirror descent, when equipped with the hypentropy mirror map and proper initialization, achieves a nearly minimax-optimal rate of convergence, provided the sample size is at least of order k^2 (modulo logarithmic term) and the minimum (in modulus) non-zero entry of the signal is on the order of $\|\mathbf{x}^*\|_2/\sqrt{k}$. Our theory leads to a simple algorithm that does not rely on explicit regularization or thresholding steps to promote sparsity. More generally, our results establish a connection between mirror descent and sparsity in the non-convex problem of noisy sparse phase retrieval, adding to the literature on early stopping that has mostly focused on non-sparse, Euclidean and convex settings via gradient descent. Our proof combines a potential-based analysis of mirror descent with a quantitative control on a variational coherence property that we establish along the path of mirror descent, up to a prescribed stopping time.

Keywords: mirror descent; early stopping; sparsity; minimax rate; implicit regularization; non-convex empirical risk; phase retrieval.

1. Introduction

In many fields of science and engineering, one is tasked with phase retrieval, the problem of recovering an n -dimensional signal \mathbf{x}^* from phaseless measurements

$$Y_j = (\mathbf{A}_j^\top \mathbf{x}^*)^2 + \varepsilon_j, \quad j = 1, \dots, m, \quad (1.1)$$

where the sensing vectors $\mathbf{A}_1, \dots, \mathbf{A}_m$ are observed, and the random variables $\varepsilon_1, \dots, \varepsilon_m$ model a possible contamination of the measurements with noise. This problem arises naturally in applications such as optics, X-ray crystallography and astronomy, where it is often easier to build detectors which measure intensities and not phases [29]. In many applications of interest, the underlying signal \mathbf{x}^* is naturally sparse [9, 15, 22]. (Noisy) sparse phase retrieval represents a cornerstone in high-dimensional statistics. It is more challenging than sparse linear regression due to the missing phase information, which often leads to non-convex problems. A geometric analysis of a natural least-squares formulation of phase retrieval (without sparsity and noise) shows that, with high probability, there are no spurious local minimizers, provided the number of Gaussian measurements m is at least of order $n \log^3 n$ [32]. However, recovery of a sparse signal is possible from potentially far fewer measurements $m \ll n$,

in which case the least-squares formulation studied in [32] has potentially exponentially many local minimizers.

Over the years, a variety of methods have been proposed to tackle the problem of sparse phase retrieval. These methods typically enforce sparsity either by adding explicit regularization via penalty terms [8, 26], sparsity inducing priors [30] or sparsity constraints [16, 25, 28], or by using algorithmic principles such as thresholding steps [5, 37, 43]. The noisy model has been studied and stability results have been established in [8] for a convex relaxation-based formulation, which includes an ℓ_1 -penalty term to promote sparsity, and in [16] for an alternating minimization-based approach that uses an explicit sparsity constraint to enforce sparsity. However, the optimal dependence of the estimation error rate on the sample size, the sparsity of the signal and the noise level was not considered in aforementioned works. Empirical risk minimization with sparsity constraint has been shown to achieve a nearly minimax-optimal rate of convergence in [19], though it does not lead to a tractable algorithm due to the non-convex sparsity constraint. Thresholded Wirtinger flow (TWF) [5] is a non-convex optimization-based algorithm that relies on thresholding steps to enforce sparsity and has been shown to achieve a nearly minimax-optimal rate of convergence for noisy sparse phase retrieval, provided the number of measurements m is of order k^2 modulo logarithmic terms.

Mirror descent was first introduced by Nemirovski and Yudin [24] for solving large-scale convex optimization problems and has since gained in popularity in various optimization and machine learning settings. An appealing property of mirror descent is the fact that it can be adapted to the (possibly non-Euclidean) geometry of the problem at hand by choosing a suitable strictly convex function, the so-called mirror map. In convex optimization, the algorithm often admits a potential-based convergence analysis in terms of the Bregman divergence associated to the mirror map, see e.g. [3, 4, 23]. There is a large body of literature that analyzes mirror descent in optimization (both convex and non-convex) and online learning settings, see e.g. [2, 10, 31]. In particular, an algorithm based on online mirror descent has been studied for an online version of phase retrieval (without sparsity) [18]. In the case of *noiseless* sparse phase retrieval, an approach based on continuous-time mirror descent has been recently shown to lead to recovery of the signal up to an arbitrary precision, without requiring explicit regularization terms or thresholding steps [39]. These results, however, do not apply to noisy measurements where perfect recovery of the signal is no longer possible and where the iterates of mirror descent may cease to be (approximately) sparse as time increases.

In the present work, we build upon the framework considered in [39] and we present a full analysis of *early-stopped* mirror descent (both continuous-time and discrete-time), equipped with the hypentropy mirror map, applied to minimize the (non-convex) unregularized empirical risk to solve *noisy* sparse phase retrieval. We establish the lower bound $\frac{\sigma}{\|\mathbf{x}^*\|_2} \sqrt{k/m}$ for the minimax-optimal rate of convergence for noisy sparse phase retrieval under sub-exponential noise for signals \mathbf{x}^* with minimum non-zero entry (in modulus) on the order of $\|\mathbf{x}^*\|_2/\sqrt{k}$, where σ describes the noise level. We prove that early-stopped mirror descent achieves this rate up to a factor $\sqrt{\log n}$, provided the sample size is on the order of k^2 (modulo logarithmic term) and the sensing vectors $\mathbf{A}_1, \dots, \mathbf{A}_m$ are independent standard Gaussian vectors. Up to logarithmic terms, the sample size requirement matches that of existing results in the literature on (noisy) sparse phase retrieval [5, 8].

The proposed approach yields a simple algorithm for noisy sparse phase retrieval, which, unlike most existing algorithms, does not rely on added regularization terms or algorithmic principles such as thresholding steps to enforce sparsity of the estimates. Running mirror descent involves tuning at most two parameters, namely the step size η (only for the discrete-time version of the algorithm) and the mirror map parameter β . Our analysis shows that η should be chosen smaller than a quantity depending on the signal size $\|\mathbf{x}^*\|_2$, whereas the choice of β depends on $\|\mathbf{x}^*\|_2$ and the (known) ambient dimension

n . As the signal size $\|\mathbf{x}^*\|_2$ can be easily estimated by considering the average size of the observations [6, 36], the tuning of η and β only involves known or easily estimated quantities, and, in particular, does not require knowledge (or estimation) of the sparsity k or the noise level σ . The optimal stopping time, whose knowledge is not required *a priori* to run mirror descent, is the only parameter that depends on the unknown quantities k and σ . Our numerical simulations attest that a commonly used data-dependent stopping rule based on cross-validation, where no knowledge of k or σ is used, yields results that validate our theoretical findings.

The main idea behind the proof of our results is to combine a generic potential-based analysis of mirror descent in terms of the Bregman divergence with a quantitative control of a variational coherence property that we establish along the path of mirror descent, when properly initialized, up to a prescribed stopping time. Variational coherence has been previously used as an *a priori* assumption to establish convergence results for mirror descent in non-convex optimization. As defined in [47], variational coherence corresponds to the requirement that the inner product $\langle \nabla F(\mathbf{x}), \mathbf{x} - \mathbf{x}' \rangle$ is non-negative for any vector \mathbf{x} in a certain region, where F is the objective function to be minimized and \mathbf{x}' is a local minimizer of F . Two versions of variational coherence have been considered in [47]: a global property that holds for any vector \mathbf{x} and any global minimizer \mathbf{x}' , and a local property that only needs to be satisfied for any vector \mathbf{x} in an open neighbourhood of a local minimizer \mathbf{x}' . The global property precludes the existence of local minimizers (that are not also global minimizers) and saddle points and is not satisfied in the problem of sparse phase retrieval when F is taken to be the unregularized empirical risk. Further, both the global and local versions of variational coherence are formulated to yield convergence towards local minimizers of the objective function, which is not the goal in noisy sparse phase retrieval as, almost surely, the signal \mathbf{x}^* does not coincide with a local minimizer \mathbf{x}' of the empirical risk. Moreover, to obtain a direct control on the speed of convergence, the stricter *strong variational coherence* [46] property, namely that $\langle \nabla F(\mathbf{x}), \mathbf{x} - \mathbf{x}' \rangle \geq \frac{c}{2} \|\mathbf{x} - \mathbf{x}'\|_2^2$ for some $c > 0$, needs to be assumed. In contrast, our analysis unveils and exploits a *quantitative* control on variational coherence within a region where the trajectory of mirror descent lies up to the stopping time, upon proper initialization. The object of our analysis is given by the inner product $\langle \nabla F(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^* \rangle$, where $\mathbf{X}(t)$ represents the iterate of mirror descent at time t and where the signal \mathbf{x}^* now replaces the local minimizer \mathbf{x}' of F . In our analysis, we establish a lower bound for this inner product that is strictly positive and directly depends on the iterates of the algorithm. This adaptive lower bound yields a control on the rate of convergence of mirror descent, which is key to investigate early stopping and establish our results. In existing gradient-based approaches to sparse phase retrieval, the inclusion of thresholding steps confines the algorithm to the low-dimensional subspace of sparse vectors, see e.g. [5, 37, 42]. On the other hand, our analysis shows that, when properly initialized, the iterates of mirror descent have negligibly small off-support coordinates until a certain stopping time, allowing to focus on the subspace of (approximately) sparse vectors.

Our work adds to the literature on early stopping for iterative methods, which has been primarily developed in the context of ridge regression and kernel methods [27, 38, 41, 44], i.e. convex problems based on the Euclidean geometry which, in particular, do not involve sparsity. An exception is the work [33], where a connection between early stopping and sparsity has been established for gradient descent with Hadamard parametrization applied to the (convex) problem of sparse recovery. This result was later recovered in [34] using mirror descent with the hypentropy mirror map. Our results extend this line of research by establishing a connection between early stopping and sparsity for mirror descent applied to the non-convex problem of noisy sparse phase retrieval. Our approach shows how to apply the generic potential-based analysis of mirror descent to non-convex problems where a quantitative control on variational coherence is exhibited along the path traced by the algorithm.

The remainder of the paper is organized as follows. In Section 2, we give the required background and a brief literature review on sparse phase retrieval. In Section 3, we describe our approach to solving sparse phase retrieval using mirror descent. In Section 4, we present our main results for mirror descent in continuous time and in discrete time. In Section 5, we perform a numerical study to verify that the dependence of the estimation error achieved by mirror descent with respect to the noise level σ , sample size m and sparsity k can match the behaviour prescribed by our upper bounds. In Section 6, we present the proof of the results for the continuous-time algorithm. The proof of the results for the discrete-time algorithm involves further technicalities, and we defer it to the supplementary material. Code to reproduce our numerical study can be found at https://github.com/fawuuu/esmd_nspr.

Notation. We use boldface letters for vectors and matrices, normal font for real numbers, and, generally, uppercase letters for random and lowercase letters for deterministic quantities. For any number $n \in \mathbb{N}$, we write $[n] = \{1, \dots, n\}$. For any vector $\mathbf{x} \in \mathbb{R}^n$ and $p \geq 0$, we write $\|\mathbf{x}\|_p = (\sum_{i=1}^n |x_i|^p)^{1/p}$, which is the standard ℓ_p -norm for $p \geq 1$. For a random variable X , we write $\|X\|_{\psi_1} = \sup_{p \geq 1} p^{-1} (\mathbb{E}[|X|^p])^{1/p}$ for its sub-exponential norm. We denote by $\mathcal{S} = \{i \in [n] : x_i^* \neq 0\}$ the support of \mathbf{x}^* . For any $\mathbf{x} \in \mathbb{R}^n$ and subset of coordinates $\mathcal{C} \subseteq [n]$, we write $\mathbf{x}_{\mathcal{C}} = (x_i)_{i \in \mathcal{C}} \in \mathbb{R}^{|\mathcal{C}|}$, and, similarly for $i \in [n]$, we write $\mathbf{x}_{-i} = (x_j)_{j \in [n] \setminus \{i\}} \in \mathbb{R}^{n-1}$. The notation $f(n) \lesssim g(n)$ (resp. $f(n) \gtrsim g(n)$, $f(n) \asymp g(n)$) means that there are constants $c_1, c_2 > 0$ such that $f(n) \leq c_1 g(n)$ (resp. $f(n) \geq c_2 g(n)$, $c_2 g(n) \leq f(n) \leq c_1 g(n)$).

2. Sparse phase retrieval

The goal in sparse phase retrieval is to reconstruct an unknown k -sparse signal $\mathbf{x}^* \in \mathbb{R}^n$ from a set of (possibly noisy) quadratic measurements $Y_j = (\mathbf{A}_j^\top \mathbf{x}^*)^2 + \varepsilon_j$, $j = 1, \dots, m$. We consider the standard setting where the sensing vectors $\mathbf{A}_j \sim \mathcal{N}(0, \mathbf{I}_n)$ are independent standard Gaussian vectors, and the noise terms ε_j are independent centered sub-exponential random variables with maximum sub-exponential norm $\sigma = \max_{j \in [m]} \|\varepsilon_j\|_{\psi_1}$. For the clarity of the analytical results, we focus on noisy sparse phase retrieval with real signal and measurement vectors. Our results extend naturally to the complex case.

We follow the well-established approach to estimating the signal \mathbf{x}^* based on non-convex optimization [5, 37, 42, 43]. Given observations $\{\mathbf{A}_j, Y_j\}_{j=1}^m$, we consider the following (non-convex) unregularized empirical risk minimization problem:

$$F(\mathbf{x}) = \frac{1}{4m} \sum_{j=1}^m ((\mathbf{A}_j^\top \mathbf{x})^2 - Y_j)^2. \quad (2.1)$$

It is worth mentioning that a risk function based on amplitude measurements $|\mathbf{A}_j^\top \mathbf{x}^*|$ has also been considered [37, 43]. However, the objective function becomes non-smooth in that case, and an analysis via the mirror descent framework appears more challenging.

Without any restrictions, the non-convex function F in (2.1) could potentially have exponentially many local minimizers and saddle points if $m \lesssim n$, cf. [32]. In the noiseless case, it has been shown in [20] that $m \geq 4k - 1$ Gaussian measurements suffice for $\{\pm \mathbf{x}^*\}$ to be the sparsest minimizer of F with high probability, that is

$$\{\pm \mathbf{x}^*\} = \operatorname{argmin}_{\mathbf{x}: \|\mathbf{x}\|_0 \leq k} F(\mathbf{x}). \quad (2.2)$$

However, solving (2.2) is challenging due to the non-convexity of the objective F and the combinatorial nature of the sparsity constraint $\|\mathbf{x}\|_0 \leq k$. Existing gradient-based methods such as TWF [5], sparse truncated amplitude flow [37], compressive reweighted amplitude flow [43] and sparse Wirtinger flow [42] minimize the non-convex empirical risk by employing a non-trivial diagonal thresholding, orthogonality-promoting or spectral initialization scheme, which produces an initial estimate close to the signal (up to a global sign) $\{\pm \mathbf{x}^*\}$, inside the so-called basin of attraction. This is followed by a refinement procedure via thresholded gradient descent iterations, where the thresholding steps are necessary to enforce sparsity of the iterates.

Numerous other algorithms have been developed to solve sparse phase retrieval, which rely on a strategy other than thresholding steps to exploit sparsity. One such strategy is to include an ℓ_1 -regularization term to the objective function, which is commonly found in convex relaxation-based approaches such as compressive phase retrieval via lifting [26] and SparsePhaseMax [13], but also in the generalized message passing algorithm PR-GAMP [30], which includes a sparsity promoting prior. Alternatively, one can directly restrict the search to k -sparse vectors, either with a preliminary support recovery step as in the alternating minimization algorithm SparseAltMinPhase [25], or by updating the estimated support in every iteration of the search in a greedy fashion as in GESPAR [28].

Hadamard Wirtinger flow (HWF) [40] is an algorithm that performs gradient descent on the unregularized empirical risk (2.1) using the Hadamard parametrization, and it can be recovered as a first-order approximation to mirror descent equipped with the hypentropy mirror map [34, 39]. HWF has been empirically studied in [40], and in some settings it has been shown to exhibit favourable sample complexities compared with other gradient-based approaches to sparse phase retrieval in numerical simulations.

3. Mirror descent algorithms

Our proposed approach to solving noisy sparse phase retrieval consists of running unconstrained mirror descent equipped with the hypentropy mirror map to minimize the empirical risk (2.1). When initialized as described below, we require neither an added regularization term nor thresholding steps to enforce sparsity of the estimates. We begin by giving a brief overview of the main quantities and definitions for mirror descent. The key object defining the geometry of mirror descent is the *mirror map*.

DEFINITION 3.1. Let $\mathcal{D} \subseteq \mathbb{R}^n$ be a convex open set. We say that $\Phi : \mathcal{D} \rightarrow \mathbb{R}$ is a mirror map if it is strictly convex, differentiable and $\{\nabla \Phi(\mathbf{x}) : \mathbf{x} \in \mathcal{D}\} = \mathbb{R}^n$.

A central quantity in the analysis of mirror descent is the *Bregman divergence* associated to a mirror map Φ , which is a measure of distance between two points $\mathbf{x}, \mathbf{y} \in \mathcal{D}$ given by

$$D_\Phi(\mathbf{x}, \mathbf{y}) = \Phi(\mathbf{x}) - \Phi(\mathbf{y}) - \nabla \Phi(\mathbf{y})^\top (\mathbf{x} - \mathbf{y}).$$

As a special case, we obtain the squared Euclidean distance by choosing the mirror map $\Phi(\mathbf{x}) = \|\mathbf{x}\|_2^2$, in which case mirror descent coincides with standard gradient descent.

As we consider unconstrained mirror descent, we have $\mathcal{D} = \mathbb{R}^n$. Let $F : \mathbb{R}^n \rightarrow \mathbb{R}$ be the objective function we seek to minimize. The trajectory of mirror descent is characterized by the mirror map Φ and an initial point $\mathbf{X}(0)$, and is, in continuous time, defined by the ordinary differential equation [24]

$$\frac{d}{dt} \mathbf{X}(t) = -(\nabla^2 \Phi(\mathbf{X}(t)))^{-1} \nabla F(\mathbf{X}(t)). \quad (3.1)$$

In discrete time, the mirror descent update is given by

$$\nabla \Phi(\mathbf{X}^{t+1}) = \nabla \Phi(\mathbf{X}^t) - \eta \nabla F(\mathbf{X}^t), \quad (3.2)$$

where $\eta > 0$ is the step-size. Throughout, we write $\mathbf{X}(t)$ and \mathbf{X}^t for the iterates of mirror descent in continuous and discrete time, respectively.

For the mirror map, we choose the hypentropy mirror map [11] given by

$$\Phi(\mathbf{x}) = \sum_{i=1}^n \left(x_i \operatorname{arcsinh}\left(\frac{x_i}{\beta}\right) - \sqrt{x_i^2 + \beta^2} \right), \quad (3.3)$$

for some parameter $\beta > 0$. We provide a discussion on the choice of the parameter β in Section 4. This choice of mirror map is motivated by the fact that, equipped with this mirror map, mirror descent is approximated by gradient descent with Hadamard parametrization [11, 34], which has been studied in the recovery of low-rank structures in sparse recovery [14, 33, 45] and matrix factorization [1, 12, 21].

For the initialization, we follow the approach outlined in [40], which estimates a single coordinate on the support of the signal \mathbf{x}^* and initializes all other coordinates to zero.

$$X_i(0) \equiv X_i^0 = \begin{cases} \frac{1}{\sqrt{3}} \sqrt{\frac{1}{m} \sum_{j=1}^m Y_j} & i = I_0, \\ 0 & i \neq I_0 \end{cases}, \quad I_0 \in \arg \max_i \left\{ \frac{1}{m} \sum_{j=1}^m Y_j A_{ji}^2 \right\}. \quad (3.4)$$

The term $(\frac{1}{m} \sum_{j=1}^m Y_j)^{1/2}$ is an estimate of $\|\mathbf{x}^*\|_2$ [6, 36]. Lemma 1 of [40] shows in the noiseless case that $x_{I_0}^* \neq 0$ with high probability, provided $m \gtrsim (x_{\max}^*)^{-2} \max\{k \log n, \log^3 n\}$, where $x_{\max}^* = \max_i |x_i^*| / \|\mathbf{x}^*\|_2$. The following is an immediate extension of Lemma 1 of [40] to the noisy observation model.

LEMMA 3.2. Let $\mathbf{x}^* \in \mathbb{R}^n$ be any k -sparse vector, and let the observations $\{Y_j, \mathbf{A}_j\}_{j=1}^m$ be given as in (1.1), where $\mathbf{A}_j \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ i.i.d. and $\{\varepsilon_j\}_{j=1}^m$ are independent centered sub-exponential random variables with maximum sub-exponential norm $\sigma = \max_j \|\varepsilon_j\|_{\psi_1}$. Let $I_0 \in \arg \max_i \{\frac{1}{m} \sum_{j=1}^m Y_j A_{ji}^2\}$. There exist universal constants $c_s, c_p > 0$ such that the following holds. If $m \geq c_s (1 + \frac{\sigma^2}{\|\mathbf{x}^*\|_2^4}) (x_{\max}^*)^{-2} \max\{k \log n, \log^3 n\}$, then, with probability at least $1 - c_p n^{-10}$, we have $|x_{I_0}^*| \geq \frac{1}{2} x_{\max}^* \|\mathbf{x}^*\|_2$.

Proof. By Lemma C.5 in the supplementary material, $|\frac{1}{m} \sum_{j=1}^m \varepsilon_j A_{ji}^2| \leq c\sigma \sqrt{\frac{\log n}{m}}$ holds with probability $1 - c_p n^{-10}$, where $c > 0$ is an absolute constant. The rest of the proof follows the same steps as the proof of Lemma 1 in [40], and we omit the details. \square

4. Main results

We show that, with high probability, early-stopped mirror descent recovers any k -sparse signal $\mathbf{x}^* \in \mathbb{R}^n$ with $x_{\min}^* \gtrsim 1/\sqrt{k}$ from $(1 + \frac{\sigma^2}{\|\mathbf{x}^*\|_2^4})k^2$ (modulo logarithmic term) Gaussian measurements, where we write $x_{\min}^* = \min_{i: x_i^* \neq 0} |x_i^*| / \|\mathbf{x}^*\|_2$.

We begin by characterizing the relationship between the Bregman divergence $D_\Phi(\mathbf{x}^*, \mathbf{x})$ associated to the hypentropy mirror map Φ in (3.3) and the ℓ_2 -norm $\|\mathbf{x} - \mathbf{x}^*\|_2$.

LEMMA 4.1. ([39]) Let $\mathbf{x}^* \in \mathbb{R}^n$ be any k -sparse vector with $x_{\min}^* \geq c_\star/\sqrt{k}$ for a constant $c_\star > 0$. Let $S = \{i \in [n] : x_i^* \neq 0\}$ be its support, and let Φ be as in (3.3) with parameter $\beta > 0$.

- For any vector $\mathbf{x} \in \mathbb{R}^n$, we have

$$\|\mathbf{x} - \mathbf{x}^*\|_2^2 \leq 2\sqrt{\max\{\|\mathbf{x}\|_\infty^2, \|\mathbf{x}^*\|_\infty^2\} + \beta^2 D_\Phi(\mathbf{x}^*, \mathbf{x})}. \quad (4.1)$$

- Let $\mathbf{x} \in \mathbb{R}^n$ be any vector with $x_i x_i^* \geq 0$ (no mismatched sign) and $|x_i| \geq \frac{1}{2}|x_i^*|$ for all $i = 1, \dots, n$. Then, we have

$$D_\Phi(\mathbf{x}^*, \mathbf{x}) \leq \frac{\sqrt{k}}{c_\star \|\mathbf{x}^*\|_2} \|\mathbf{x}_S - \mathbf{x}_S^*\|_2^2 + \|\mathbf{x}_{S^c}\|_1. \quad (4.2)$$

A proof of Lemma 4.1 appears in the appendix of [40]. For completeness, we include a proof of Lemma 4.1 in Section 6.1 below. Lemma 4.1 implies that, if we are interested in convergence with respect to the ℓ_2 -norm, we can consider the Bregman divergence D_Φ as a proxy for it. As it is impossible to distinguish \mathbf{x}^* from $-\mathbf{x}^*$ using phaseless measurements, we consider the ℓ_2 -distance and Bregman divergence from the solution set $\{\pm \mathbf{x}^*\}$ given by $\text{dist}(\mathbf{x}^*, \mathbf{x}) = \min\{\|\mathbf{x} - \mathbf{x}^*\|_2, \|\mathbf{x} + \mathbf{x}^*\|_2\}$ and $\text{dist}_\Phi(\mathbf{x}^*, \mathbf{x}) = \min\{D_\Phi(\mathbf{x}^*, \mathbf{x}), D_\Phi(-\mathbf{x}^*, \mathbf{x})\}$, respectively. Applying the bound in (4.1) to both \mathbf{x}^* and $-\mathbf{x}^*$, we find the upper bound

$$\text{dist}(\mathbf{x}^*, \mathbf{x})^2 \leq 2\sqrt{\max\{\|\mathbf{x}\|_\infty^2, \|\mathbf{x}^*\|_\infty^2\} + \beta^2 \text{dist}_\Phi(\mathbf{x}^*, \mathbf{x})}.$$

The setting for our convergence results is summarized in the following assumption.

Assumption 4.2. For some universal constants $c_\star, c_s > 0$, the following holds. The signal $\mathbf{x}^* \in \mathbb{R}^n$ is k -sparse with $x_{\min}^* \geq c_\star/\sqrt{k}$. The observations $\{Y_j, \mathbf{A}_j\}_{j=1}^m$ are i.i.d. given as in (1.1), with $\mathbf{A}_j \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ and ε_j centred sub-exponential with maximum sub-exponential norm $\sigma = \max_j \|\varepsilon_j\|_{\psi_1}$. The sample size is at least $m \geq c_s(1 + \frac{\sigma^2}{\|\mathbf{x}^*\|_2^4}) \max\{k^2 \log^2 n, \log^5 n\}$.

We now formulate our results for continuous-time mirror descent in terms of the Bregman divergence D_Φ . Lemma 4.1 allows to translate the upper bound on $\text{dist}_\Phi(\mathbf{x}^*, \mathbf{X}(t))$ into an upper bound on the ℓ_2 -distance $\text{dist}(\mathbf{x}^*, \mathbf{X}(t))$. Recall that we denote the support of the signal \mathbf{x}^* by $S = \{i \in [n] : x_i^* \neq 0\}$.

THEOREM 4.3. Let Assumption 4.2 hold. There exist universal constants $c_p, c, c_1, c_2 > 0$ such that the following holds. Let $\mathbf{X}(t)$ be defined by the continuous-time mirror descent algorithm given by (3.1)

with mirror map (3.3) and initialization (3.4) with $\beta \leq c_1 \|\mathbf{x}^\star\|_2/n^3$. Let $\delta = \sqrt{n\beta/\|\mathbf{x}^\star\|_2}$ and define

$$T_2 = \inf \left\{ t > 0 : \frac{\text{dist}(\mathbf{x}^\star, \mathbf{X}(t))}{\|\mathbf{x}^\star\|_2} \leq c \max \left\{ \sqrt{c_\star \sqrt{k} \delta}, \frac{\sigma}{\|\mathbf{x}^\star\|_2^2} \sqrt{\frac{k \log n}{m}} \right\} \right\}.$$

Then, there exists $T_1 \leq \frac{c_2}{\|\mathbf{x}^\star\|_2^3} k \log(\frac{\|\mathbf{x}^\star\|_2}{\beta}) \log(k \log \frac{\|\mathbf{x}^\star\|_2}{\beta})$ such that

$$\frac{\text{dist}_\phi(\mathbf{x}^\star, \mathbf{X}(t))}{\|\mathbf{x}^\star\|_2} \leq \frac{6\sqrt{k}}{c_\star} \exp\left(-\frac{c_\star \|\mathbf{x}^\star\|_2^3}{4\sqrt{k}}(t - T_1)\right) \quad \text{for all } T_1 \leq t \leq T_2, \quad (4.3)$$

with probability at least $1 - c_p n^{-10}$. Furthermore, for all $t \leq T_2$, we have

$$\|\mathbf{X}_{S^c}(t)\|_1 \leq \delta \|\mathbf{x}^\star\|_2. \quad (4.4)$$

Analogous results also hold in discrete time.

THEOREM 4.4. Let Assumption 4.2 hold. There exist universal constants $c_p, c, c_1, c_2, c_3 > 0$ such that the following holds. Let \mathbf{X}^t be defined by the discrete-time mirror descent algorithm given by (3.2) with mirror map (3.3) and initialization (3.4) with $\beta \leq c_1 \|\mathbf{x}^\star\|_2/n^3$ and $\eta \leq c_3/\|\mathbf{x}^\star\|_2^3$. Let $\delta = \sqrt{n\beta/\|\mathbf{x}^\star\|_2}$, and define

$$T_2 = \inf \left\{ t > 0 : \frac{\text{dist}(\mathbf{x}^\star, \mathbf{X}^t)}{\|\mathbf{x}^\star\|_2} \leq c \max \left\{ \sqrt{c_\star \sqrt{k} \delta}, \frac{\sigma}{\|\mathbf{x}^\star\|_2^2} \sqrt{\frac{k \log n}{m}} \right\} \right\}.$$

Then, there exists $T_1 \leq \frac{c_2}{\eta \|\mathbf{x}^\star\|_2^3} k \log(\frac{\|\mathbf{x}^\star\|_2}{\beta}) \log(k \log \frac{\|\mathbf{x}^\star\|_2}{\beta})$ such that

$$\frac{\text{dist}_\phi(\mathbf{x}^\star, \mathbf{X}^t)}{\|\mathbf{x}^\star\|_2} \leq \frac{6\sqrt{k}}{c_\star} \left(1 - \frac{c_\star \eta \|\mathbf{x}^\star\|_2^3}{8\sqrt{k}}\right)^{t-T_1} \quad \text{for all } T_1 \leq t \leq T_2, \quad (4.5)$$

with probability at least $1 - c_p n^{-10}$. Furthermore, for all $t \leq T_2$, we have

$$\|\mathbf{X}_{S^c}^t\|_1 \leq \delta \|\mathbf{x}^\star\|_2. \quad (4.6)$$

In the noiseless case, when $\sigma = 0$, Theorem 4.3 recovers the results in [39] and Theorem 4.4 shows that any desired relative error $\varepsilon > 0$ can be achieved by running mirror descent with the choice $\beta \asymp \varepsilon^2 \|\mathbf{x}^\star\|_2/n^3$ for $k \log \frac{\|\mathbf{x}^\star\|_2}{\beta}$ (modulo double-logarithmic term) iterations.

We now discuss some of the main implications of Theorem 4.3 and Theorem 4.4. For notational simplicity, we focus on discrete-time mirror descent throughout the remainder of this section, but the same considerations also apply to the continuous-time algorithm.

Nearly minimax-optimal rate. Theorem 4.4 implies that mirror descent achieves a relative mean-squared error $\frac{\text{dist}(\mathbf{X}^{T_2}, \mathbf{x}^\star)}{\|\mathbf{x}^\star\|_2}$ on the order of $\frac{\sigma}{\|\mathbf{x}^\star\|_2^2} \sqrt{\frac{k \log n}{m}}$, provided the mirror map parameter β is chosen smaller than $\sigma^4 k \log^2 n / (m^2 n \|\mathbf{x}^\star\|_2^7)$. This is identical to the estimation error rate achieved by TWF [5]

and nearly matches the lower bound on the minimax-optimal rate that we derive below in Theorem 4.5. As we only consider k -sparse signals with minimum non-zero entry on the order of $\|\mathbf{x}^*\|_2/\sqrt{k}$, our lower bound on the minimax rate of convergence differs by a factor $\sqrt{\log(en/k)}$ from the bound of $\frac{\sigma}{\|\mathbf{x}^*\|_2^2} \sqrt{\frac{k \log(en/k)}{m}}$ for k -sparse vectors, which has been stated as Theorem 3.2 in [5].

THEOREM 4.5. Let $\Theta(k, n, r) = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 = r, \|\mathbf{x}\|_0 = k, \min_{i: x_i \neq 0} |x_i| \geq c_\star r/\sqrt{k}\}$ for some constant $c_\star > 0$. Let $\{\mathbf{A}_j\}_{j=1}^m$ be a collection of i.i.d. $\mathcal{N}(0, \mathbf{I}_n)$ random vectors, $\{\varepsilon_j\}_{j=1}^m$ be a collection of i.i.d. $\mathcal{N}(0, \sigma^2)$ random variables, and let $\{Y_j\}_{j=1}^m$ be given as in (1.1). There exist universal constants $c_s, c > 0$ such that if $m \geq c_s(1 + \frac{\sigma^2}{R^4})k \log(en/k)$, then

$$\inf_{\hat{\mathbf{x}}} \sup_{\mathbf{x}^* \in \Theta(k, n, r)} \mathbb{P}_{(\mathbf{A}, \mathbf{Y}|\mathbf{x}^*)} \left[\frac{\text{dist}(\mathbf{x}^*, \hat{\mathbf{x}})}{r} \geq \frac{c\sigma}{r^2} \sqrt{\frac{k}{m}} \right] \geq \frac{1}{5}, \quad (4.7)$$

where the infimum is taken over all procedures $\hat{\mathbf{x}}(\mathbf{A}, \mathbf{Y})$.

Proof. Theorem 4.5 follows from an application of Theorem C in [19]. We first introduce some notation used for Theorem C. For a vector $\mathbf{x} \in \mathbb{R}^n$ and $r > 0$, we write $\mathcal{B}_2(\mathbf{x}, \rho)$ for the ℓ_2 -ball with centre \mathbf{x} and radius ρ . For any set $\mathcal{X} \subset \mathbb{R}^n$ and $\rho > 0$, we denote the packing number with respect to the ℓ_2 -norm by

$$M(\mathcal{X}, \rho) = \sup\{n : \exists \mathbf{x}_1, \dots, \mathbf{x}_n \in \mathcal{X} \text{ s.t. } \forall i \neq j, \|\mathbf{x}_i - \mathbf{x}_j\|_2 > \rho\}.$$

To apply Theorem C of [19], we need to bound the packing number of $\Theta(k, n, r) \cap \mathcal{B}_2(\mathbf{x}_0, c_0\rho)$, where $\mathbf{x}_0 \in \Theta(k, n, r)$ is an arbitrary vector. More precisely, we will show that, for all $\rho > 0$ and $c_0 \geq 2$,

$$\sup_{\mathbf{x}_0 \in \Theta(k, n, r)} \log^{1/2} M(\Theta(k, n, r) \cap \mathcal{B}_2(\mathbf{x}_0, c_0\rho), \rho) \asymp \sqrt{k}. \quad (4.8)$$

Theorem 4.5 then follows from an application of Theorem C of [19] (c.f. Corollary 6.2 in [19]).

In the minimax rate given in Corollary 6.2 in [19] and Theorem 3.2 in [5], the set $\Theta(n, k, r)$ is defined without restricting $\min_{i: x_i \neq 0} |x_i| > c_\star r/\sqrt{k}$. In that case, standard estimates on the packing number yield

$$\log^{1/2} M(\tilde{\Theta}(k, n, r) \cap \mathcal{B}_2(\mathbf{x}_0, c_0\rho), \rho) \asymp \sqrt{k \log(en/k)},$$

where $\tilde{\Theta}(k, n, r) = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 = r, \|\mathbf{x}\|_0 = k\}$, see e.g. Lemma 1.4.2 and Lemma 2.2.17 in [7]. On the other hand, with the restriction $\min_{i: x_i \neq 0} |x_i| > c_\star r/\sqrt{k}$, when we fix any $\mathbf{x}_0 \in \Theta(k, n, r)$, then any vector $\mathbf{x} \in \Theta(k, n, r) \cap \mathcal{B}_2(\mathbf{x}_0, c_0\rho)$ must be supported on the support of \mathbf{x}_0 for $\rho < \sqrt{2}c_\star r/(c_0\sqrt{k})$, because any non-zero coordinate of \mathbf{x}_0 and \mathbf{x} has absolute value at least $c_\star r/\sqrt{k}$, and therefore $\|\mathbf{x} - \mathbf{x}_0\|_2 \geq \sqrt{2}c_\star r/\sqrt{k}$ if \mathbf{x} and \mathbf{x}_0 have different supports. This leads to a packing number that scales exponentially in k (as opposed to a scaling of $\binom{n}{k}$ in the case of $\tilde{\Theta}(n, k, r)$) and therefore the estimate in (4.8).

On early stopping and sparsity. The bound (4.6) in Theorem 4.4 controls the magnitude off-support coordinates can attain while mirror descent runs, up to a time T_2 . To show the bound (4.6), we crucially use the assumption $x_{\min}^* \gtrsim 1/\sqrt{k}$, which is needed to show that all support coordinates

grow faster than any off-support coordinate, which then leads to $\|\mathbf{X}_{\mathcal{S}^c}^t\|_1 < \delta \|\mathbf{x}^*\|_2$ for a sufficiently small δ ; see Section 4.1 for an outline of the proof idea, and Stage (i), part (b) in Section 6.2 for details. The fact that mirror descent iterates stay approximately k -sparse allows us to show bound (4.5): after an initial ‘warm-up’ period of length T_1 , mirror descent converges linearly to the solution set $\{\pm \mathbf{x}^*\}$ up to a time T_2 , at which a precision determined by the mirror map parameter β and the quantity $\frac{\sigma}{\|\mathbf{x}^*\|_2^2} \sqrt{\frac{k \log n}{m}}$ is reached. The necessity for early stopping is explained by the fact that, as elaborated in the proof sketch below, we need to control the size of off-support variables $\|\mathbf{X}_{\mathcal{S}^c}^t\|_1$, which is guaranteed to be sufficiently small up to the time T_2 .

Using the fact that the Bregman divergence decreases linearly for $T_1 \leq t \leq T_2$, a quick calculations shows that $T_2 - T_1 \lesssim \sqrt{k} \log \frac{k \|\mathbf{x}^*\|_2}{\beta}$, see Stage (ii), part (a) in Section 6.2 for details. Therefore, an upper bound for T_2 is given by $k \log \frac{\|\mathbf{x}^*\|_2}{\beta}$ (modulo double-logarithmic term). In applications, we can use a data-dependent stopping rule such as the hold-out method outlined in [27] to estimate T_2 : we can run mirror descent using a fraction $\alpha \in (0, 1)$ of the data and return the iterate \mathbf{X}^t , which minimizes the empirical risk F evaluated on the remaining $(1 - \alpha)$ of the data. Section 5 presents a numerical study that also investigates this data-dependent stopping rule, for which establishing theoretical guarantees is outside the scope of our contribution.

Role of the mirror map parameter. In Theorem 4.4, choosing a small parameter β is needed to guarantee that off-support variables stay sufficiently small in the bound (4.6). The quantity δ in the bound (4.6) depends polynomially on β , while choosing a small β leads to the length of the initial warm-up period T_1 increasing as $\log \frac{\|\mathbf{x}^*\|_2}{\beta}$. In practice, we can therefore simply choose a very small β (e.g. 10^{-20}), so that convergence is linear up to a precision on the order of $\frac{\sigma}{\|\mathbf{x}^*\|_2^2} \sqrt{\frac{k \log n}{m}}$. A similar trade-off between computational cost and statistical accuracy with respect to the size of the initialization has previously been observed in [33] in the setting of noisy sparse linear regression.

4.1 Proof ideas

We give a high-level overview of the main ideas behind the proofs of Theorem 4.3 and Theorem 4.4. We begin with the proof of Theorem 4.3. Without loss of generality, let us assume that the initialization in (3.4) satisfies $x_{i_0}^* > 0$, so that $\mathbf{X}(0)^\top \mathbf{x}^* > 0$. In this case, we will show that $\mathbf{X}(t)$ converges to \mathbf{x}^* . Otherwise, we can show convergence to $-\mathbf{x}^*$.

Employing a potential-based analysis in terms of the Bregman divergence, we will bound the Bregman divergence $D_\phi(\mathbf{x}^*, \mathbf{X}(t))$, which then also yields a bound on the ℓ_2 -distance $\|\mathbf{x}^* - \mathbf{X}(t)\|_2$ via Lemma 4.1. A quick calculation yields

$$\frac{d}{dt} D_\phi(\mathbf{x}', \mathbf{X}(t)) = -\langle \nabla F(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}' \rangle, \quad (4.9)$$

where $\mathbf{x}' \in \mathbb{R}^n$ is any reference point. If F were to be convex, then choosing \mathbf{x}' in equation (4.9) to be any global minimizer of F would imply

$$\langle \nabla F(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}' \rangle \geq F(\mathbf{X}(t)) - F(\mathbf{x}'),$$

which shows that continuous-time mirror descent strictly decreases the Bregman divergence as long as $F(\mathbf{X}(t)) > F(\mathbf{x}')$. As the Bregman divergence is bounded from below by zero by construction due to

the convexity of the mirror map, this implies that $F(\mathbf{X}(t))$ must converge to $\min_{\mathbf{x}} F(\mathbf{x})$, which in turn implies that $\mathbf{X}(t)$ converges to a global minimizer of F , provided the function F is continuous. For non-convex objective functions, the inner product in (4.9) has been used to define the notion of *variational coherence* [47], namely the assumption that

$$\langle \nabla F(\mathbf{x}), \mathbf{x} - \mathbf{x}' \rangle \geq 0 \quad \text{for all } \mathbf{x} \in \mathbb{R}^n, \mathbf{x}' \in \mathcal{X}^*, \quad (4.10)$$

where $\mathcal{X}^* = \arg \min_{\mathbf{z}} F(\mathbf{z})$, and that there exists an $\mathbf{x}' \in \mathcal{X}^*$ such that $\langle \nabla F(\mathbf{x}), \mathbf{x} - \mathbf{x}' \rangle = 0$ implies that also $\mathbf{x} \in \mathcal{X}^*$. Under this assumption, convergence of a stochastic version of mirror descent towards the set of minimizers of F has been shown in [47]. In the batch setting, convergence of mirror descent towards a global minimizer of F can be shown as in the convex case under the assumption of variational coherence. A local version of variational coherence, only assuming that inequality (4.10) holds for a local minimizer \mathbf{x}' and for \mathbf{x} belonging to an open neighbourhood of \mathbf{x}' , was also considered in [47], yielding a local version of the convergence results.

In order to study convergence rates, a more refined notion, namely quantitative control on a strictly positive lower bound in (4.10), would be necessary. For example, if the function F were to satisfy a lower bound of the form

$$\langle \nabla F(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}' \rangle \geq c D_{\Phi}(\mathbf{x}', \mathbf{X}(t))$$

for some constant $c > 0$ and minimizer $\mathbf{x}' \in \mathcal{X}^*$, then this would imply exponential convergence towards \mathbf{x}' , since we have

$$\frac{d}{dt} D_{\Phi}(\mathbf{x}', \mathbf{X}(t)) \leq -c D_{\Phi}(\mathbf{x}', \mathbf{X}(t)) \quad \Rightarrow \quad D_{\Phi}(\mathbf{x}', \mathbf{X}(t)) \leq e^{-ct} D_{\Phi}(\mathbf{x}', \mathbf{X}(0)).$$

By design, both the global and local version of variational coherence can only be used to establish convergence towards local minimizers of the objective function. As we are interested in establishing convergence towards the statistical signal \mathbf{x}^* , we consider the inner product $\langle \nabla F(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^* \rangle$ and bound this quantity from below within a region where the trajectory of mirror descent is confined to stay up to the prescribed stopping time, upon proper initialization. In order to obtain results on the rate of convergence, it is crucial to quantify the lower bound in (4.10). In particular, we derive a lower bound of the form

$$\langle \nabla F(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^* \rangle \geq \begin{cases} \frac{c_1}{\sqrt{k}} & \text{if } \|\mathbf{X}(t)\|_2^2 \leq \frac{2}{3} \|\mathbf{x}^*\|_2^2 \\ c_2 \|\mathbf{X}(t) - \mathbf{x}^*\|_2^2 & \text{else} \end{cases} \quad (4.11)$$

for some constants $c_1, c_2 > 0$. The analysis is divided into two stages: (i) an initial warm-up period of length T_1 , where the Bregman divergence decreases first at a constant rate $1/\sqrt{k}$, after which we can relate the quantity $\|\mathbf{X}(t) - \mathbf{x}^*\|_2^2$ in the second bound in (4.11) to the Bregman divergence $D_{\Phi}(\mathbf{x}^*, \mathbf{X}(t))$ by counting the number of ‘small’ coordinates for which $|X_i(t)| < \frac{1}{2}|x_i^*|$ and (ii) a subsequent stage where we show linear convergence by relating the Bregman divergence to the quantity $\|\mathbf{X}(t) - \mathbf{x}^*\|_2^2$ in the lower bound (4.11) via Lemma 4.1.

The proof of the lower bound (4.11) involves three main steps. Writing $f(\mathbf{x}) = \mathbb{E}[F(\mathbf{x})]$ for the population risk, we assume in the first step that we have access to the population gradient, which, by the

dominated convergence theorem, can be computed as

$$\nabla f(\mathbf{x}) = \mathbb{E}[\nabla F(\mathbf{x})] = (3\|\mathbf{x}\|_2^2 - \|\mathbf{x}^*\|_2)\mathbf{x} - 2(\mathbf{x}^\top \mathbf{x}^*)\mathbf{x}^*.$$

In this setting, we can obtain the lower bounds in (4.11) for the inner product $\langle \nabla f(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^* \rangle$ via simple algebraic manipulations: the first bound in (4.11) relies on Lemma 6.3 and the assumption $x_{\min}^* \geq c_*/\sqrt{k}$ to bound the inner product $\mathbf{X}(t)^\top \mathbf{x}^*$ from below by c_1/\sqrt{k} , whereas the second bound only uses the Cauchy–Schwarz inequality.

In the second step, we need to control the deviation of the empirical quantity

$$\langle \nabla F(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^* \rangle = \frac{1}{m} \sum_{j=1}^m ((\mathbf{A}_j^\top \mathbf{X}(t))^2 - Y_j)(\mathbf{A}_j^\top \mathbf{X}(t))(\mathbf{A}_j^\top (\mathbf{X}(t) - \mathbf{x}^*)) \quad (4.12)$$

from the population quantity $\langle \nabla f(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^* \rangle$. The deviation is bounded in Lemma 6.2, and we summarize the main proof ideas in the following. If the iterate $\mathbf{X}(t)$ were to be independent from the measurement vectors $\{\mathbf{A}_j\}_{j=1}^m$ and noise terms $\{\varepsilon_j\}_{j=1}^m$, then, conditionally on $\mathbf{X}(t)$, the term in (4.12) would be a fourth order Gaussian polynomial, which can be controlled using standard concentration bounds. More precisely, we could use a truncation argument and split it into a bounded and an unbounded part, where the unbounded part is non-zero only with small probability, and use Lipschitz concentration and Hoeffding’s bounds to control the bounded part, and Chebyshev’s inequality for the unbounded part. However, the assumption of independence is not reasonable. Instead, we use an ϵ -net argument followed by above concentration argument for each fixed vector \mathbf{x} in the ϵ -net. If we considered an ϵ -net in a unit ball in \mathbb{R}^n , our argument would require a sample size of $m \gtrsim nk$, with which plenty of algorithms provably recover the signal \mathbf{x}^* regardless of the assumption on sparsity, e.g. [6, 8, 25]. To obtain our sample requirement of k^2 (modulo logarithmic terms), we show that mirror descent exhibits algorithmic regularization in the sense that $\mathbf{X}(t)$ stays ‘almost k -sparse’ (more precisely, we show that $\|\mathbf{X}_{\mathcal{S}^c}(t)\|_1 < \delta \|\mathbf{x}^*\|_2$) and apply a union bound to control the deviation of the term in (4.12) from its expectation for all bounded vectors \mathbf{x} supported on \mathcal{S} .

In the third step, in order to show that off-support coordinates stay negligibly small, we consider the trajectory of a coordinate $X_i(t)$, which is defined by (cf. (3.1))

$$\frac{d}{dt}X_i(t) = -\sqrt{X_i(t)^2 + \beta^2} \nabla F(\mathbf{X}(t))_i.$$

If we imagined that $\beta = 0$ and $\nabla F(\mathbf{X}(t))_i = -c$ was a constant, then the trajectory would simply be given by $X_i(t) = e^{ct}$. Using the assumption $x_{\min}^* \geq c_*/\sqrt{k}$ along with the concentration argument outlined above, we can show that, for any $i \in \mathcal{S}$ and $j \notin \mathcal{S}$,

$$|\nabla F(\mathbf{X}(t))_j| \leq \gamma |\nabla F(\mathbf{X}(t))_i|$$

holds with high probability for some $\gamma < 1$. Informally, the behaviour of support and off-support coordinates is comparable to $\beta(1+c)^t$ and $\beta(1+\gamma c)^t$, respectively, and the gap can be made large by choosing β small and t large. A detailed proof can be found in Section 6.2.

The main ideas of the proof of Theorem 4.4 are similar to those of Theorem 4.3. To show that the Bregman divergence is decreasing in discrete time, we need to bound the difference

$$D_\Phi(\mathbf{x}^\star, \mathbf{X}^{t+1}) - D_\Phi(\mathbf{x}^\star, \mathbf{X}^t) = -\eta \langle \nabla F(\mathbf{X}^t), \mathbf{X}^t - \mathbf{x}^\star \rangle + D_\Phi(\mathbf{X}^t, \mathbf{X}^{t+1}). \quad (4.13)$$

The first term in (4.13) can be bounded as in continuous time, while the second term $D_\Phi(\mathbf{X}^t, \mathbf{X}^{t+1})$ measures the distance between consecutive iterates in terms of the Bregman divergence and is due to the discretization and can be bounded using the strong convexity of the mirror map Φ in (3.3). Unlike in continuous time, we need to establish coordinate-wise upper and lower bounds for how much each coordinate X_i^t changes in one iteration in the discrete-time case in order to characterize the region to which the trajectory of mirror descent is confined. This is not required in continuous time, since in that case we can use the continuity of $\mathbf{X}_i(t)$. A detailed proof can be found in the supplementary material.

5. Numerical study

In this section, we provide numerical experiments demonstrating how the minimum relative estimation error achieved after t_{\max} iterations, that is

$$\min_{1 \leq t \leq t_{\max}} \frac{\text{dist}(\mathbf{x}^\star, \mathbf{X}^t)}{\|\mathbf{x}^\star\|_2}, \quad (5.1)$$

depends on the noise-to-signal ratio $\sigma/\|\mathbf{x}^\star\|_2^2$, the sample size m and the sparsity level k , and how the length of the warm-up period as defined in Section 6.2, that is

$$T_1 = \inf \left\{ t > 0 : \min_{i \in S} \frac{|X_i^t|}{|x_i^\star|} > \frac{1}{2} \right\},$$

depends on the parameter β and the sparsity level k . The main purpose of this section is to validate how the theoretical bounds in Theorem 4.4 on the estimation error and T_1 , as well as the convergence behaviour suggested by Theorem 4.4, can be representative of the actual performance of mirror descent.

The exponentiated gradient algorithm EG \pm without normalization [17] is an equivalent formulation of mirror descent equipped with the hypentropy map (Theorem 24 [11]), which is given by

$$\begin{aligned} \mathbf{X}^t &= \mathbf{U}^t - \mathbf{V}^t, \\ \mathbf{U}^{t+1} &= \mathbf{U}^t \odot \exp(-\eta \nabla F(\mathbf{X}^t)), \quad \mathbf{V}^{t+1} = \mathbf{V}^t \odot \exp(\eta \nabla F(\mathbf{X}^t)), \end{aligned} \quad (5.2)$$

where \odot denotes the elementwise Hadamard product, and the initialization (3.4) becomes

$$U_i^0 = \begin{cases} \frac{\hat{\theta}}{2\sqrt{3}} + \sqrt{\frac{\hat{\theta}^2}{12} + \frac{\beta^2}{4}} & i = I_0, \\ \frac{\beta}{2} & i \neq I_0 \end{cases}, \quad V_i^0 = \begin{cases} -\frac{\hat{\theta}}{2\sqrt{3}} + \sqrt{\frac{\hat{\theta}^2}{12} + \frac{\beta^2}{4}} & i = I_0, \\ \frac{\beta}{2} & i \neq I_0 \end{cases}, \quad (5.3)$$

where we write $\hat{\theta} = (\sum_{j=1}^m Y_j/m)^{1/2}$ for the estimate of $\|\mathbf{x}^\star\|_2$. In our simulations, we found EG \pm to be numerically more stable than the update (3.2) for small values of β due to the division by β in (3.2).




FIG. 1. Relative error $\text{dist}(\mathbf{x}^*, \mathbf{X}^{T_{\text{stop}}})/\|\mathbf{x}^*\|_2$ plus-minus one standard deviation (shaded area) of early-stopped mirror descent, with the stopping time T_{stop} selected with oracle knowledge (red circles), and using the hold-out method (blue crosses). **Left:** relative error against noise-to-signal ratio $\sigma/\|\mathbf{x}^*\|_2^2$, with $n = m = 2000$, $k = 10$. **Center:** log-log plot of relative error against sample size m , with $n = 2000$, $k = 10$, $\sigma/\|\mathbf{x}^*\|_2^2 = 0.1$. **Right:** log-log plot of relative error against sparsity level k , with $n = 2000$, $m = 4000$, $\sigma/\|\mathbf{x}^*\|_2^2 = 0.1$.

In all experiments, we generate a k -sparse signal vector $\mathbf{x}^* \in \mathbb{R}^n$ by sampling x_i^* uniformly from $[-1, -0.15] \cup [0.15, 1]$ for $i = 1, \dots, n$ and setting $(n - k)$ random entries of \mathbf{x}^* to zero. We sample m i.i.d. sensing vectors $\mathbf{A}_j \sim \mathcal{N}(0, \mathbf{I}_n)$ and noise terms $\varepsilon_j \sim \mathcal{N}(0, \sigma^2)$ (note that $\|\varepsilon_j\|_{\psi_1} = \sqrt{2/\pi}\sigma$), and generate observations $Y_j = (\mathbf{A}_j^\top \mathbf{x}^*)^2 + \varepsilon_j$, $j = 1, \dots, m$. For mirror descent, we find that the step size $\eta = 0.3/(\sum_{j=1}^m Y_j/m)^{3/2}$ works well, which is consistent with the value prescribed by Theorem 4.

5.1 Estimation error

First, we examine how the relative estimation error of early-stopped mirror descent depends on the noise-to-signal ratio $\sigma/\|\mathbf{x}^*\|_2^2$, the sample size m and the sparsity level k . Theorem 4.4 provides an upper bound for the relative estimation error of order $\frac{\sigma}{\|\mathbf{x}^*\|_2^2} \sqrt{\frac{k \log n}{m}}$, provided β is chosen small enough so that $\sqrt{c_* \sqrt{k} \delta} \leq \frac{\sigma}{\|\mathbf{x}^*\|_2^2} \sqrt{\frac{k \log n}{m}}$. In the following experiments, we fix $n = 2000$ and $\beta = 10^{-20}$. We run mirror descent for $t_{\text{max}} = 5000$ iterations and consider the minimum relative estimation error defined in (5.1), averaged across 100 independent Monte Carlo trials.

Evaluating the minimum relative error defined in (5.1) requires oracle knowledge of the underlying signal $\|\mathbf{x}^*\|_2$. For the sake of completeness, in our simulations we also consider the hold-out method to implement early stopping, even if this data-dependent rule goes beyond the reach of the theoretical results that we have developed (in this, we follow the same approach considered in prior related literature [27]). The hold-out method only uses observed data to select a stopping time. More precisely, we run mirror descent using 90% of the data and select the stopping time $T_{\text{stop}} \in \{1, \dots, t_{\text{max}}\}$ for which the iterate $\mathbf{X}^{T_{\text{stop}}}$ minimizes the risk F evaluated on the remaining 10% of the data. Our implementation of the early stopped mirror descent algorithm is summarized in Algorithm 1, where F_{train} and F_{val} denote the objective function (2.1) computed on the training and hold-out dataset, respectively.

In the first experiment, we fix $k = 10$, $m = 2000$ and vary the noise-to-signal ratio $\sigma/\|\mathbf{x}^*\|_2^2$. The left plot in Fig. 1 verifies the linear relationship between the relative estimation error and $\sigma/\|\mathbf{x}^*\|_2^2$ suggested by the upper bound in Theorem 4.4.

Algorithm 1 Early-stopped mirror descent (hold-out method)

- 1: **Input:** training dataset $\{Y_j^{train}, \mathbf{A}_j^{train}\}_{j=1}^{m_1}$, hold-out dataset $\{Y_j^{val}, \mathbf{A}_j^{val}\}_{j=1}^{m_2}$, step size η , number of iterations t_{max} , mirror map parameter β
- 2: Set

$$I_0 = \arg \max_i \left\{ \sum_{j=1}^{m_1} Y_j^{train} (A_{ji}^{train})^2 + \sum_{j=1}^{m_2} Y_j^{val} (A_{ji}^{val})^2 \right\}$$
- 3: Initialize $\mathbf{U}^0, \mathbf{V}^0$ as in (5.3)
- 4: **for** $t = 0, \dots, t_{max}$ **do**
- 5: Update

$$\begin{aligned} \mathbf{X}^t &= \mathbf{U}^t \odot \mathbf{U}^t - \mathbf{V}^t \odot \mathbf{V}^t \\ \mathbf{U}^{t+1} &= \mathbf{U}^t \odot \exp(-\eta \nabla F_{train}(\mathbf{X}^t)) \\ \mathbf{V}^{t+1} &= \mathbf{V}^t \odot \exp(\eta \nabla F_{train}(\mathbf{X}^t)) \end{aligned}$$
- 6: **end for**
- 7: Set

$$T_{stop} = \arg \min_t F_{val}(\mathbf{X}^t)$$
- 8: **Return:** $\mathbf{X}^{T_{stop}}$

Next, we fix $k = 10$, $\sigma/\|\mathbf{x}^*\|_2^2 = 0.1$, and increase m from 1500 to 5000. The log-log plot (Fig. 1, center) shows a slope of -0.5137 , which validates the prediction of Theorem 4.4 that the estimation error decreases as $1/\sqrt{m}$ as the sample size increases.

Finally, we fix $m = 4000$, $\sigma/\|\mathbf{x}^*\|_2^2 = 0.1$ and vary the sparsity k . The log-log plot (Fig. 1, right) shows a slope of 0.5775 , which, considering the standard deviation, is in line with the upper bound of Theorem 4.4 that increases as \sqrt{k} as the sparsity level increases.

In all experiments, the relative error and its standard deviation are larger when the stopping time is selected using the hold-out method compared with a stopping time selected using oracle knowledge of the signal \mathbf{x}^* , but the same trends are exhibited.

5.2 Warm-up time T_1

We note that the asymptotic computational complexity of the algorithm is controlled by the growth of T_1 , since $T_2 - T_1 \lesssim \sqrt{k} \log \frac{k\|\mathbf{x}^*\|_2}{\beta} \lesssim T_1$, see Stage (ii) part (a) in Section 6.2. As a result, we focus on T_1 in the following simulations, where we examine how the length of the warm-up period T_1 depends on the mirror map parameter β and the sparsity level k . Theorem 4.4 provides an upper bound of order $k \log \frac{\|\mathbf{x}^*\|_2}{\beta}$ (modulo double-logarithmic terms).

We first fix $n = 2000$, $m = 1500$, $k = 10$, $\sigma/\|\mathbf{x}^*\|_2^2 = 0.1$ and decrease $\beta/\|\mathbf{x}^*\|_2$ from 10^{-4} to 10^{-40} . The left plot in Fig. 2 validates the upper bound in Theorem 4.4, the length of the warm-up period T_1 increases linearly with $\log \frac{\|\mathbf{x}^*\|_2}{\beta}$.

In the second experiment, we set $n = 2000$, $m = 4000$, $\sigma/\|\mathbf{x}^*\|_2^2 = 0.1$, $\beta = 10^{-20}$ and increase k from 5 to 25. Again, the right plot in Fig. 2 shows that T_1 increases linearly with the sparsity level k , validating the upper bound in Theorem 4.4.




FIG. 2. Length of warm-up period T_1 plus/minus one standard deviation (shaded area) of early-stopped mirror descent. **Left:** T_1 against $\log(\|\mathbf{x}^*\|_2/\beta)$, with $n = 1000$, $m = 1500$, $k = 10$, $\sigma/\|\mathbf{x}^*\|_2^2 = 0.1$. **Right:** T_1 against sparsity level k , with $n = 2000$, $m = 4000$, $\sigma/\|\mathbf{x}^*\|_2^2 = 0.1$, $\beta/\|\mathbf{x}^*\|_2 = 10^{-20}$.




FIG. 3. Relative error $\text{dist}(\mathbf{x}^*, \mathbf{X}^t)/\|\mathbf{x}^*\|_2$ (log-scale) of mirror descent for $n = 50000$, $m = 1000$, $k = 10$ and different values of β . **Left:** noiseless case with $\sigma/\|\mathbf{x}^*\|_2^2 = 0$. **Right:** noisy case with $\sigma/\|\mathbf{x}^*\|_2^2 = 0.5$.

5.3 Convergence behaviour of mirror descent

The next experiment examines the convergence behaviour of mirror descent and its dependence on the parameter β . We only present the ℓ_2 -error $\text{dist}(\mathbf{x}^*, \mathbf{X}^t)$. The Bregman divergence term displays the same qualitative behaviour. We set $n = 50000$, $m = 1000$, $k = 10$ and run mirror descent for different values of β . In both the noiseless ($\sigma/\|\mathbf{x}^*\|_2^2 = 0$) and noisy ($\sigma/\|\mathbf{x}^*\|_2^2 = 0.5$) case, we observe the expected dependence on β : as β decreases from 10^{-6} to 10^{-14} , the length of the initial warm-up period increases as $\log \frac{1}{\beta}$, whereas we have linear convergence up to an improved accuracy that depends polynomially on β . In the noisy case, the precision up to which we have linear convergence barely improves as we decrease β from 10^{-10} to 10^{-14} , which we attribute to the noise term $\frac{\sigma}{\|\mathbf{x}^*\|_2^2} \sqrt{\frac{k \log n}{m}}$ dominating the achievable accuracy. Figure 3 suggests that early stopping is necessary in the presence of noise.

6. Proofs

In this section, we prove the results stated in Section 4. Since the proof of Theorem 4.4 largely follows the same steps as the proof of Theorem 4.3, with a few technical difficulties arising due to the discretization, we defer the proof of Theorem 4.4 to the supplementary material.

6.1 Proof of Lemma 4.1

The proof of Lemma 4.1 relies on the fact that the Bregman divergence associated to the hypentropy mirror map can be decomposed into a sum of functions, each of which depends only on a single coordinate x_i and is (locally) bounded from above and below by quadratic functions.

Proof of Lemma 4.1 Using the identity $\operatorname{arcsinh}(x) = \log(x + \sqrt{1 + x^2})$, we can write

$$D_\Phi(\mathbf{x}^\star, \mathbf{x}) = \sum_{i=1}^n g_i(x_i),$$

where

$$g_i(x_i) = \sqrt{x_i^2 + \beta^2} - \sqrt{(x_i^\star)^2 + \beta^2} - x_i^\star \log \frac{x_i + \sqrt{x_i^2 + \beta^2}}{x_i^\star + \sqrt{(x_i^\star)^2 + \beta^2}}.$$

We begin by showing the bound (4.2). For any $i \in [n]$, we have $g_i(x_i^\star) = 0$ and

$$g'_i(x_i) = \frac{x_i}{\sqrt{x_i^2 + \beta^2}} - x_i^\star \frac{1 + x_i/\sqrt{x_i^2 + \beta^2}}{x_i + \sqrt{x_i^2 + \beta^2}} = \frac{x_i - x_i^\star}{\sqrt{x_i^2 + \beta^2}}. \quad (6.1)$$

As we assume $x_i x_i^\star \geq 0$ and $|x_i| \geq \frac{1}{2}|x_i^\star|$, we have $\sqrt{z^2 + \beta^2} \geq \frac{x_{\min}^\star \|\mathbf{x}^\star\|_2}{2} \geq \frac{c_\star \|\mathbf{x}^\star\|_2}{2\sqrt{k}}$ for any $z \in (x_i, x_i^\star)$ (or (x_i^\star, x_i) if $x_i^\star < x_i$). Hence, using the convention $\int_a^b f(x)dx = -\int_b^a f(x)dx$ when $a > b$, we can bound

$$g_i(x_i) = g_i(x_i^\star) + \int_{x_i^\star}^{x_i} g'_i(z)dz < \int_{x_i^\star}^{x_i} \frac{2\sqrt{k}}{c_\star \|\mathbf{x}^\star\|_2} (z - x_i^\star)dz = \frac{\sqrt{k}}{c_\star \|\mathbf{x}^\star\|_2} (x_i - x_i^\star)^2.$$

Summing over all $i \in \mathcal{S}$, this gives

$$\sum_{i \in \mathcal{S}} g_i(x_i) \leq \frac{\sqrt{k}}{c_\star \|\mathbf{x}^\star\|_2} \|\mathbf{x}_\mathcal{S} - \mathbf{x}_\mathcal{S}^\star\|_2^2$$

as claimed. On the other hand, if $x_i^\star = 0$, we have

$$g_i(x_i) = \sqrt{x_i^2 + \beta^2} - \beta \leq |x_i|,$$

which gives

$$\sum_{i \notin \mathcal{S}} g_i(x_i) \leq \|x_{\mathcal{S}^c}\|_1$$

and completes the proof of the bound (4.2) in Lemma 4.1. The other bound (4.1) can be shown similarly using (6.1). We have $\sqrt{z^2 + \beta^2} \leq \sqrt{\max\{\|\mathbf{x}\|_\infty^2, \|\mathbf{x}^\star\|_\infty^2\} + \beta^2}$ for all $z \in (x_i, x_i^\star)$ (or (x_i^\star, x_i) if $x_i^\star < x_i$), which gives

$$\begin{aligned} g(x_i) &= g(x_i^\star) + \int_{x_i}^{x_i^\star} g'(z) dz > \int_{x_i}^{x_i^\star} \left(\max\{\|\mathbf{x}\|_\infty^2, \|\mathbf{x}^\star\|_\infty^2\} + \beta^2 \right)^{-\frac{1}{2}} (z - x_i^\star) dz \\ &= \frac{1}{2} \left(\max\{\|\mathbf{x}\|_\infty^2, \|\mathbf{x}^\star\|_\infty^2\} + \beta^2 \right)^{-\frac{1}{2}} (x_i - x_i^\star)^2. \end{aligned}$$

Summing over all $i \in [n]$ gives the bound (4.1) of Lemma 4.1. \square

6.2 Proof of Theorem 4.1

In this section, we will make the following assumptions, which are purely for notational simplicity.

- We will assume $\|\mathbf{x}^\star\|_2 = 1$. The general case $\|\mathbf{x}^\star\|_2 \neq 1$ follows by writing $\mathbf{x}^\star = \frac{\mathbf{x}^\star}{\|\mathbf{x}^\star\|_2} \|\mathbf{x}^\star\|_2$, $\mathbf{x} = \frac{\mathbf{x}}{\|\mathbf{x}^\star\|_2} \|\mathbf{x}^\star\|_2$ and $\varepsilon_j = \frac{\varepsilon_j}{\|\mathbf{x}^\star\|_2} \|\mathbf{x}^\star\|_2$ in what follows.
- We will assume $\mathbf{X}(0)^\top \mathbf{x}^\star \geq 0$, in which case we show convergence towards \mathbf{x}^\star . Otherwise, we can show convergence towards $-\mathbf{x}^\star$.

The gradient of the objective function (2.1) is given by

$$\nabla F(\mathbf{x}) = \frac{1}{m} \sum_{j=1}^m ((\mathbf{A}_j^\top \mathbf{x})^2 - (\mathbf{A}_j^\top \mathbf{x}^\star)^2 - \varepsilon_j) (\mathbf{A}_j^\top \mathbf{x}) \mathbf{A}_j.$$

A brief calculation gives the following expression for its expectation, the population gradient.

$$\nabla f(\mathbf{x}) = \mathbb{E}[\nabla F(\mathbf{x})] = (3\|\mathbf{x}\|_2^2 - 1)\mathbf{x} - 2(\mathbf{x}^\top \mathbf{x}^\star)\mathbf{x}^\star.$$

First, we provide three supporting lemmas characterizing the behaviour of mirror descent. We defer the proofs of the following lemmas to the supplementary material. Lemmas 6.1 and 6.2 show that the gradient $\nabla F(\mathbf{x})$ is close to its expectation $\nabla f(\mathbf{x})$.

LEMMA 6.1. Let $\mathbf{x}^\star \in \mathbb{R}^n$ be a k -sparse vector, and let $\mathcal{S} = \{i \in [n] : x_i^\star \neq 0\}$ be its support. Let $\{\mathbf{A}_j\}_{j=1}^m$ be a collection of i.i.d. $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ random vectors and $\{\varepsilon_j\}_{j=1}^m$ a collection of independent centered sub-exponential random variables with maximum sub-exponential norm $\sigma = \max_j \|\varepsilon_j\|_{\psi_1}$. There exist universal constants $c, c_p > 0$ such that for any constant $\gamma > 0$, there is a constant $c_s(\gamma) > 0$ such that the following holds. For any $0 \leq \delta \leq \frac{1}{4}$, let

$$\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 \leq \|\mathbf{x}^\star\|_2, \|\mathbf{x}_{\mathcal{S}^c}\|_1 \leq \delta \|\mathbf{x}^\star\|_2\}.$$

If $m \geq c_s(\gamma)(1 + \frac{\sigma^2}{\|\mathbf{x}^*\|_2^4}) \max\{k^2 \log^2 n, \log^5 n\}$, then, with probability at least $1 - c_p n^{-10}$,

$$\begin{aligned} \frac{\|\nabla F(\mathbf{x}) - \nabla f(\mathbf{x})\|_\infty}{\|\mathbf{x}^*\|_2^2} &\leq \gamma \min\left\{\frac{\|\mathbf{x}_S\|_1}{k}, \frac{\|\mathbf{x}_S - \mathbf{x}_S^*\|_2}{\sqrt{k}}\right\} + c\delta\|\mathbf{x}^*\|_2 \\ &\quad + \frac{c\sigma}{\|\mathbf{x}^*\|_2^2} \sqrt{\frac{\log n}{m}} \min\left\{\|\mathbf{x}\|_1, (1 + \delta)\|\mathbf{x}^*\|_2 + \sqrt{k}\|\mathbf{x}_S - \mathbf{x}_S^*\|_2\right\} \end{aligned}$$

holds for all $\mathbf{x} \in \mathcal{X}$. In particular, if $\frac{\|\mathbf{x} - \mathbf{x}^*\|_2}{\|\mathbf{x}^*\|_2} \geq \max\left\{\frac{c\sqrt{k}\delta}{\gamma}, \frac{c(1+\delta)\sigma}{\gamma\|\mathbf{x}^*\|_2^2} \sqrt{\frac{k \log n}{m}}\right\}$, $\|\mathbf{x}_S\|_1 \geq \frac{1}{2}\|\mathbf{x}^*\|_2$, $\delta \leq \frac{\gamma}{2ck}$ and $c_s(\gamma) \geq \frac{c^2}{\gamma^2}$, then we have the simplified bound

$$\frac{\|\nabla F(\mathbf{x}) - \nabla f(\mathbf{x})\|_\infty}{\|\mathbf{x}^*\|_2^2} \leq 4\gamma \min\left\{\frac{\|\mathbf{x}_S\|_1}{k}, \frac{\|\mathbf{x} - \mathbf{x}^*\|_2}{\sqrt{k}}\right\} \quad \text{for all } \mathbf{x} \in \mathcal{X}.$$

LEMMA 6.2. Let $\mathbf{x}^* \in \mathbb{R}^n$ be a k -sparse vector, and let $\mathcal{S} = \{i \in [n] : x_i^* \neq 0\}$ be its support. Let $\{\mathbf{A}_j\}_{j=1}^m$ be a collection of i.i.d. $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ random vectors and $\{\varepsilon_j\}_{j=1}^m$ a collection of independent centered sub-exponential random variables with maximum sub-exponential norm $\sigma = \max_j \|\varepsilon_j\|_{\psi_1}$. There exist universal constants $c, c_p > 0$ such that for any constant $\gamma > 0$, there is a constant $c_s(\gamma) > 0$ such that the following holds. For any $0 \leq \delta \leq \frac{1}{4} \min\{\frac{\gamma}{c}, 1\}$, let

$$\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 \leq \|\mathbf{x}^*\|_2, \|\mathbf{x}_{\mathcal{S}^c}\|_1 \leq \delta\|\mathbf{x}^*\|_2\}.$$

If $m \geq c_s(\gamma)(1 + \frac{\sigma^2}{\|\mathbf{x}^*\|_2^4}) \max\{k^2 \log^2 n, \log^5 n\}$, then, with probability at least $1 - c_p n^{-10}$,

$$\begin{aligned} \frac{|\langle \nabla F(\mathbf{x}) - \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle|}{\|\mathbf{x}^*\|_2^3} &\leq \gamma \min\left\{\frac{\|\mathbf{x}_S\|_1}{\sqrt{k}} + \delta\|\mathbf{x}^*\|_2, \frac{\|\mathbf{x}_S - \mathbf{x}_S^*\|_2^2}{\|\mathbf{x}^*\|_2} + \delta^2\|\mathbf{x}^*\|_2\right\} \\ &\quad + \frac{c\sigma}{\|\mathbf{x}^*\|_2^2} \sqrt{\frac{\log n}{m}} \left(\delta\|\mathbf{x}^*\|_2 + \min\{\|\mathbf{x}_S\|_1, \sqrt{k}\|\mathbf{x}_S - \mathbf{x}_S^*\|_2\}\right) \end{aligned}$$

holds for all $\mathbf{x} \in \mathcal{X}$. In particular, if $\frac{\|\mathbf{x} - \mathbf{x}^*\|_2}{\|\mathbf{x}^*\|_2} \geq \max\left\{\delta, \frac{c\sigma}{\gamma\|\mathbf{x}^*\|_2^2} \sqrt{\frac{k \log n}{m}}\right\}$, $\|\mathbf{x}_S\|_1 \geq \frac{1}{2}\|\mathbf{x}^*\|_2$, $\delta \leq \frac{1}{2\sqrt{k}}$ and $c_s(\gamma) \geq \frac{c^2}{\gamma^2}$, then we have the simplified bound

$$\frac{|\langle \nabla F(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle - \langle \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle|}{\|\mathbf{x}^*\|_2^3} \leq 4\gamma \min\left\{\frac{\|\mathbf{x}_S\|_1}{\sqrt{k}}, \frac{\|\mathbf{x} - \mathbf{x}^*\|_2^2}{\|\mathbf{x}^*\|_2}\right\} \quad \text{for all } \mathbf{x} \in \mathcal{X}.$$

The following lemma characterizes the region to which the trajectory of early-stopped mirror descent (both in continuous and discrete time) is confined.

LEMMA 6.3. Let Assumption 4.2 hold. There exist universal constants $c_p, c_1, c_2, c_3 > 0$ such that the following holds. Let $\delta \leq c_1/n$ and let $\mathbf{X}(t)$ be defined by the continuous-time mirror descent algorithm given by (3.1) with mirror map (3.3) and initialization (3.4) with $\beta \leq \delta$. Assume that there is a $T > 0$ such that

$$\|\mathbf{X}_{\mathcal{S}^c}(t)\|_1 \leq \delta \|\mathbf{x}^*\|_2 \quad \text{and} \quad \frac{\|\mathbf{X}(t) - \mathbf{x}^*\|_2}{\|\mathbf{x}^*\|_2} \geq c_2 \max \left\{ \sqrt{k}\delta, \frac{\sigma}{\|\mathbf{x}^*\|_2^2} \sqrt{\frac{k \log n}{m}} \right\}$$

are satisfied for all $t \leq T$. Then, there is a $\xi \in \{-1, +1\}$, such that, with probability at least $1 - c_p n^{-10}$, the following holds for all $t \leq T$:

$$\xi X_i(t) x_i^* \geq 0 \quad \text{for all } i \in [n], \quad (6.2)$$

$$\frac{1}{3} - (3 + 9\sigma) \sqrt{\frac{\log n}{m}} \leq \frac{\|\mathbf{X}(t)\|_2^2}{\|\mathbf{x}^*\|_2^2} \leq 2, \quad (6.3)$$

$$\sqrt{3} \cdot 2 |\mathbf{X}(t)^\top \mathbf{x}^*| \geq 3 \|\mathbf{X}(t)\|_2^2 - \|\mathbf{x}^*\|_2^2. \quad (6.4)$$

The same result holds for \mathbf{X}^t defined by the discrete-time mirror descent algorithm given by (3.2) with $\eta \leq c_3/\|\mathbf{x}^*\|_2^3$.

The following inequalities will be useful throughout the following proofs. Under the assumptions of Lemma 6.3, we can use (6.3) to bound

$$\|\mathbf{X}_{\mathcal{S}}(t)\|_1 = \|\mathbf{X}(t)\|_1 - \|\mathbf{X}_{\mathcal{S}^c}(t)\|_1 \geq \left(\frac{1}{3} - (3 + 9\sigma) \sqrt{\frac{\log n}{m}} \right)^{-\frac{1}{2}} - \delta \geq \frac{1}{2}, \quad (6.5)$$

and, using (6.2) of Lemma 6.3, we can bound the inner product

$$\mathbf{X}(t)^\top \mathbf{x}^* \geq \|\mathbf{X}_{\mathcal{S}}(t)\|_1 x_{\min}^* \geq \frac{c_*}{2\sqrt{k}}. \quad (6.6)$$

Proof of Theorem 4.3. We first consider the population dynamics to illustrate the main ideas of the proof of Theorem 4.3. That is, we first assume that we had access to the population gradient ∇f , or in other words that $m = \infty$. The proof of Theorem 4.3 relies on the identity

$$\frac{d}{dt} D_\phi(\mathbf{x}^*, \mathbf{X}(t)) = -\langle \nabla F(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^* \rangle, \quad (6.7)$$

and we bound the inner product on the right hand side to show that the Bregman divergence $D_\phi(\mathbf{x}^*, \mathbf{X}(t))$ decreases as claimed.

Initial Bregman divergence

Writing $\hat{\theta} = (\sum_{j=1}^m Y_j/m)^{1/2}$ for the estimate of the signal size $\|\mathbf{x}^*\|_2$, we can use standard concentration bounds for sub-exponential random variables (see e.g. Prop. 5.16 of [35]) to bound, with

probability $1 - 4n^{10}$,

$$\hat{\theta}^2 = \frac{1}{m} \sum_{j=1}^m Y_j = \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_j^\top \mathbf{x}^*)^2 + \frac{1}{m} \sum_{j=1}^m \varepsilon_j < \left(1 + (9 + 25\sigma) \sqrt{\frac{\log n}{m}}\right) \|\mathbf{x}^*\|_2^2.$$

Using the definition of the initialization (3.4), we can bound the initial Bregman divergence

$$\begin{aligned} D_\Phi(\mathbf{x}^*, \mathbf{X}(0)) &= \sum_{i \neq I_0} \beta - \sqrt{(x_i^*)^2 + \beta^2} - x_i^* \log \frac{\beta}{x_i^* + \sqrt{(x_i^*)^2 + \beta^2}} \\ &\quad + \sqrt{\hat{\theta}^2/3 + \beta^2} - \sqrt{(x_{I_0}^*)^2 + \beta^2} - x_{I_0}^* \log \frac{\hat{\theta}/\sqrt{3} + \sqrt{\hat{\theta}^2/3 + \beta^2}}{x_{I_0}^* + \sqrt{(x_{I_0}^*)^2 + \beta^2}} \\ &\leq \sum_{i \neq I_0} |x_i^*| \log \frac{1}{\beta} + \beta - \sqrt{(x_i^*)^2 + \beta^2} + |x_i^*| \log \left(|x_i^*| + \sqrt{(x_i^*)^2 + \beta^2} \right) + 1 \\ &\leq \|\mathbf{x}^*\|_1 \log \frac{1}{\beta} + 1, \end{aligned} \quad (6.8)$$

where we used the assumption $\|\mathbf{x}^*\|_2 = 1$ and the fact that the function $x \log \frac{\beta}{x + \sqrt{x^2 + \beta^2}}$ is symmetric, that is $x \log \frac{\beta}{x + \sqrt{x^2 + \beta^2}} = -x \log \frac{\beta}{-x + \sqrt{x^2 + \beta^2}}$.

Bounding $\langle \nabla f(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^* \rangle$

We can compute

$$\langle \nabla f(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^* \rangle = -[3(\|\mathbf{X}(t)\|_2^2 - 1) + 2(\mathbf{X}(t)^\top \mathbf{x}^*)](\mathbf{X}(t)^\top \mathbf{x}^*) + (3\|\mathbf{X}(t)\|_2^2 - 1)\|\mathbf{X}(t)\|_2^2.$$

To bound this quantity, we distinguish two cases.

- **Case 1:** $\|\mathbf{X}(t)\|_2^2 \leq \frac{2}{5}$

In this case, we use the fact that $3\|\mathbf{X}(t)\|_2^2 - 1 \geq -(9 + 27\sigma) \sqrt{\frac{\log n}{m}}$ by Lemma 6.3 to bound

$$\langle \nabla f(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^* \rangle \geq \left(\frac{9}{5} - 2\sqrt{\frac{2}{5}}\right)(\mathbf{X}(t)^\top \mathbf{x}^*) - \frac{18 + 54\sigma}{5} \sqrt{\frac{\log n}{m}} \geq \frac{1}{2}(\mathbf{X}(t)^\top \mathbf{x}^*), \quad (6.9)$$

where for the first inequality we used that $\mathbf{X}(t)^\top \mathbf{x}^* \leq \|\mathbf{X}(t)\|_2 \|\mathbf{x}^*\|_2$ by the Cauchy–Schwarz inequality, and the second inequality holds by (6.6) since $m \geq c_s k^2 \log^2 n$.

- **Case 2:** $\|\mathbf{X}(t)\|_2^2 > \frac{2}{5}$

In this case, we can write $2(\mathbf{X}(t)^\top \mathbf{x}^*) = \|\mathbf{X}(t)\|_2^2 + \|\mathbf{x}^*\|_2^2 - \|\mathbf{X}(t) - \mathbf{x}^*\|_2^2$ to compute

$$\langle \nabla f(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^* \rangle = (2 - 4\|\mathbf{X}(t)\|_2^2 + \Delta)(\mathbf{X}(t)^\top \mathbf{x}^*) + (3\|\mathbf{X}(t)\|_2^2 - 1)\|\mathbf{X}(t)\|_2^2,$$

where we denote $\Delta = \|\mathbf{X}(t) - \mathbf{x}^\star\|_2^2$. If $2 - 4\|\mathbf{X}(t)\|_2^2 + \Delta \geq 0$, this is lower bounded by

$$\langle \nabla f(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^\star \rangle \geq \frac{2}{25}. \quad (6.10)$$

Otherwise, again using $\mathbf{X}(t)^\top \mathbf{x}^\star \leq \|\mathbf{X}(t)\|_2 \|\mathbf{x}^\star\|_2$, we have

$$\begin{aligned} \langle \nabla f(\mathbf{X}(t)), \mathbf{x}^\star - \mathbf{X}(t) \rangle &\geq \underbrace{3\|\mathbf{X}(t)\|_2^4 - 4\|\mathbf{X}(t)\|_2^3 - \|\mathbf{X}(t)\|_2^2 + 2\|\mathbf{X}(t)\|_2 + \Delta}_{\geq 0} \|\mathbf{X}(t)\|_2 \\ &\geq \sqrt{\frac{2}{5}} \Delta. \end{aligned} \quad (6.11)$$

The inequalities (6.9)–(6.11) can be used to show that the Bregman divergence $D_\phi(\mathbf{x}^\star, \mathbf{X}(t))$ decreases as claimed. In particular, we can use Lemma 4.1 to bound $\|\mathbf{X}(t) - \mathbf{x}^\star\|_2^2$ in terms of $D_\phi(\mathbf{x}^\star, \mathbf{X}(t))$, and inequality (6.11) yields a bound leading to linear convergence as long as $\|\mathbf{X}_{\mathcal{S}^c}(t)\|_1$ is sufficiently small compared to $\|\mathbf{X}_{\mathcal{S}}(t) - \mathbf{x}_{\mathcal{S}}^\star\|_2^2$. In order to make the proof of Theorem 4.3 rigorous, we need to replace the population gradient ∇f by the empirical gradient ∇F in the outline we provided above. To this end, we divide the analysis of the convergence of mirror descent into two stages bounded by

$$\begin{aligned} T_1 &= \inf \left\{ t > 0 : \min_{i \in \mathcal{S}} \frac{|X_i(t)|}{|x_i^\star|} > \frac{1}{2} \right\} \text{ and} \\ T_2 &= \inf \left\{ t > 0 : \frac{\|\mathbf{X}(t) - \mathbf{x}^\star\|_2^2}{\|\mathbf{x}^\star\|_2^2} \leq c^2 \max \left\{ c_\star \sqrt{k} \delta, \frac{\sigma^2}{\|\mathbf{x}^\star\|_2^4} \frac{k \log n}{m} \right\} \right\}, \end{aligned}$$

respectively, where $\delta = \sqrt{n\beta}$. We consider the stages (i) $t \leq T_1$ and (ii) $T_1 < t \leq T_2$. Note that we have $T_2 > T_1$, because if there is an index $i \in \mathcal{S}$ with $|X_i(t)| < \frac{1}{2}|x_i^\star|$, then we also have $\|\mathbf{X}(t) - \mathbf{x}^\star\|_2 > \frac{c_\star}{2\sqrt{k}}$. In both stages, we will (a) bound the length of the stage by using (6.7) to bound T_i and (b) show that off-support coordinates satisfy $\|\mathbf{X}_{\mathcal{S}^c}(t)\|_1 \leq \delta$. Throughout the proof we will assume that the inequalities in Lemmas 6.1–6.3 are satisfied, which happens with probability at least $1 - c_p n^{-10}$. Note that as $\|\mathbf{X}_{\mathcal{S}}(t)\|_1 \geq \|\mathbf{X}(t)\|_2 - \delta \geq \frac{1}{2}$ by the lower bound in (6.3) of Lemma 6.3 and since $\delta \lesssim n^{-1}$, the conditions for the simplified bounds in Lemmas 6.1 and 6.2 are satisfied for $t \leq T_2$.

Stage (i), part (a): $t \leq T_1$, bound T_1

Assume for now that we have already shown $\|\mathbf{X}_{\mathcal{S}^c}(t)\|_1 < \delta_1 = n\beta^{3/4}$ for all $t \leq T_1$. Note that we have $\delta_1 \leq \delta$ since $\beta \leq n^{-2}$. We have already computed the rate (6.7) at which the Bregman divergence $D_\phi(\mathbf{x}^\star, \mathbf{X}(t))$ decreases if we had access to the population gradient ∇f in (6.9) and (6.11). Using these bounds, we can bound the inner product in (6.7) with the empirical gradient ∇F .

- **Case 1:** $\|\mathbf{X}(t)\|_2^2 \leq \frac{2}{5}$

By Lemma 6.2 and the definitions of δ_1 and T_2 , we have

$$\left| \langle \nabla F(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^\star \rangle - \langle \nabla f(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^\star \rangle \right| \leq \frac{c_\star}{4\sqrt{k}} \|\mathbf{X}_{\mathcal{S}^c}(t)\|_1 \leq \frac{1}{4} (\mathbf{X}(t)^\top \mathbf{x}^\star).$$

Together with (6.9), this bound leads to

$$\begin{aligned} \frac{d}{dt} D_\phi(\mathbf{x}^\star, \mathbf{X}(t)) &\leq -\langle \nabla f(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^\star \rangle + \left| \langle \nabla F(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^\star \rangle - \langle \nabla f(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^\star \rangle \right| \\ &\leq -\frac{1}{4} (\mathbf{X}(t)^\top \mathbf{x}^\star) \leq -\frac{c_\star}{8\sqrt{k}}, \end{aligned} \quad (6.12)$$

where for the last line we used (6.6).

- **Case 2:** $\|\mathbf{X}(t)\|_2^2 \geq \frac{2}{5}$

As in the previous case, we can bound the difference $\langle \nabla F(\mathbf{X}(t)) - \nabla f(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^\star \rangle$ using Lemma 6.2. Recalling (6.10) and (6.11), we obtain

$$\frac{d}{dt} D_\phi(\mathbf{x}^\star, \mathbf{X}(t)) = -\langle \nabla F(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^\star \rangle \leq -\frac{1}{25} \quad (6.13)$$

if $2 - 4\|\mathbf{X}(t)\|_2^2 + \|\mathbf{X}(t) - \mathbf{x}^\star\|_2^2 \geq 0$, and

$$\frac{d}{dt} D_\phi(\mathbf{x}^\star, \mathbf{X}(t)) = -\langle \nabla F(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^\star \rangle \leq -\frac{1}{2} \|\mathbf{X}(t) - \mathbf{x}^\star\|_2^2 \quad (6.14)$$

if $2 - 4\|\mathbf{X}(t)\|_2^2 + \|\mathbf{X}(t) - \mathbf{x}^\star\|_2^2 < 0$, where we used that

$$\left| \langle \nabla F(\mathbf{X}(t)) - \nabla f(\mathbf{X}(t)), \mathbf{X}(t) - \mathbf{x}^\star \rangle \right| \leq \left(\sqrt{\frac{2}{5}} - \frac{1}{2} \right) \|\mathbf{X}(t) - \mathbf{x}^\star\|_2^2,$$

for all $t \leq T_2$ by the simplified bound in Lemma 6.2.

We can now bound T_1 . Define

$$T_0 = \inf \left\{ t > 0 : \|\mathbf{X}(t)\|_2^2 > \frac{2}{5} \right\},$$

as the time until which Case 1 holds. Then, the bound (6.12) from Case 1 implies

$$\min\{T_1, T_0\} \leq \frac{8\sqrt{k}D_\phi(\mathbf{x}^\star, \mathbf{X}(0))}{c_\star} \leq \frac{8}{c_\star} k \log \frac{1}{\beta} + \frac{8\sqrt{k}}{c_\star} \leq \frac{c_2}{2} k \log \frac{1}{\beta},$$

provided that $c_2 \geq \frac{16}{c_\star} + \frac{16}{c_\star \sqrt{k} \log \frac{1}{\beta}}$, where we used (6.8) to bound the Bregman divergence $D_\phi(\mathbf{x}^\star, \mathbf{X}(0))$.

If $T_1 < T_0$, then we have bounded T_1 as desired. Otherwise, we can use the bounds (6.13) and (6.14) from Case 2 to control $T_1 - T_0$. The first bound (6.13) can apply at most for $t \leq 25D_\phi(\mathbf{x}^\star, \mathbf{X}(0)) \leq 25\sqrt{k} \log \frac{1}{\beta} + 25$, where we again used the bound (6.8). As the bound (6.14) depends on the ℓ_2 -distance $\|\mathbf{X}(t) - \mathbf{x}^\star\|_2^2$, the idea is to show that this quantity is sufficiently large as long as the Bregman divergence

$D_\phi(\mathbf{x}^\star, \mathbf{X}(t))$ is large. To this end, define

$$S(t) = \left\{ i \in \mathcal{S} : \frac{|X_i(t)|}{|x_i^\star|} < \frac{1}{2} \right\}.$$

With this, we can bound

$$\|\mathbf{X}(t) - \mathbf{x}^\star\|_2^2 \geq \sum_{i \in S(t)} \left(\frac{1}{2}x_i^\star\right)^2 + \sum_{i \notin S(t)} (X_i(t) - x_i^\star)^2 \geq \frac{c_\star^2}{4k} |S(t)| + \sum_{i \notin S(t)} (X_i(t) - x_i^\star)^2. \quad (6.15)$$

As in the computation for (6.8) and the proof of Lemma 4.1, we can bound

$$D_\phi(\mathbf{x}^\star, \mathbf{X}(t)) \leq \sum_{i \in S(t)} |x_i^\star| \log \frac{1}{\beta} + \frac{\sqrt{k}}{c_\star} \sum_{i \notin S(t)} (X_i(t) - x_i^\star)^2 + \|\mathbf{X}_{S^c}(t)\|_1. \quad (6.16)$$

Since $\|\mathbf{x}^\star\|_\infty \leq \|\mathbf{x}^\star\|_2 = 1$, we have the bound $\sum_{i \in S(t)} |x_i^\star| \leq |S(t)|$. Further, since $t < T_1$, we also have $\|\mathbf{X}_{S^c}(t)\|_1 \leq \delta_1 \leq \frac{1}{4}(x_{\min}^\star)^2 \leq \sum_{i \notin S(t)} (X_i(t) - x_i^\star)^2$. With this, we can combine (6.14), (6.15) and (6.16) to bound

$$\frac{d}{dt} D_\phi(\mathbf{x}^\star, \mathbf{X}(t)) \leq -\frac{1}{2} \|\mathbf{X}(t) - \mathbf{x}^\star\|_2^2 \leq -\frac{c_\star^2}{8k \log \frac{1}{\beta}} D_\phi(\mathbf{x}^\star, \mathbf{X}(t)),$$

which shows that $D_\phi(\mathbf{x}^\star, \mathbf{X}(t))$ decreases linearly at the rate $c_\star^2/(8k \log \frac{1}{\beta})$ for $T_0 < t < T_1$. We have

$$D_\phi(\mathbf{x}^\star, \mathbf{X}(T_0)) \leq D_\phi(\mathbf{x}^\star, \mathbf{X}(0)) \leq \sqrt{k} \log \frac{1}{\beta} + 1,$$

and, for $t < T_1$,

$$\frac{c_\star^2}{4k} \leq \|\mathbf{X}(t) - \mathbf{x}^\star\|_2^2 \leq 2\sqrt{2 + \beta^2} D_\phi(\mathbf{x}^\star, \mathbf{X}(t)) \leq 3D_\phi(\mathbf{x}^\star, \mathbf{X}(t)),$$

where for the second inequality we used that $\|\mathbf{X}(t)\|_\infty \leq \sqrt{2}$ by Lemma 6.3 and the bound (4.1) of Lemma 4.1. This implies

$$T_1 - T_0 \leq \frac{8k \log \frac{1}{\beta}}{c_\star^2} \log \left(\frac{12}{c_\star^2} k^{\frac{3}{2}} \log \frac{1}{\beta} + \frac{12}{c_\star^2} k \right) \leq \frac{c_2}{2} k \log \left(\frac{1}{\beta} \right) \log \left(k \log \frac{1}{\beta} \right),$$

provided that $c_2 \geq \frac{24}{c_\star^2} + \frac{16}{c_\star^2} \log \frac{24}{c_\star^2}$.

Stage (i), part (b): $t \leq T_1$, bound $\|\mathbf{X}_{S^c}(t)\|_1$

The idea to controlling $\|\mathbf{X}_{S^c}(t)\|_1$ is as follows: we will show that for coordinates $j \notin \mathcal{S}$ and $i \in \mathcal{S}$, $X_j(t)$ can only grow at a comparatively slower rate than $X_i(t)$. We will show that the growth of both coordinates is bounded by exponentials, and use the fact that, for any fixed $\epsilon > 0$, the gap between

$\beta(1 + 2\epsilon)^t$ and $\beta(1 + \epsilon)^t$ can be made arbitrarily large by choosing t large and β small enough. Recall that

$$\frac{d}{dt}X_i(t) = -\sqrt{X_i(t)^2 + \beta^2} \nabla F(\mathbf{X}(t))_i.$$

By Lemma 6.3, we have $\sqrt{3} \cdot 2(\mathbf{X}(t)^\top \mathbf{x}^*) \geq 3\|\mathbf{X}(t)\|_2^2 - 1$. For any $i \in \mathcal{S}$ with $x_i^* > 0$ and $X_i(t) \leq \frac{1}{2}x_i^*$, we have

$$\nabla f(\mathbf{X}(t))_i \leq \left(\frac{\sqrt{3}}{2} - 1\right) \cdot 2(\mathbf{X}(t)^\top \mathbf{x}^*)x_i^* \leq -\frac{c_\star^2 \|\mathbf{X}_\mathcal{S}(t)\|_1}{4k}.$$

As before, Lemma 6.1 gives

$$|\nabla F(\mathbf{X}(t))_i - \nabla f(\mathbf{X}(t))_i| \leq \frac{c_\star^2 \|\mathbf{X}_\mathcal{S}(t)\|_1}{8k},$$

for $t \leq T_2$. This allows us to bound

$$\nabla F(\mathbf{X}(t))_i \leq -\frac{c_\star^2 \|\mathbf{X}_\mathcal{S}(t)\|_1}{8k},$$

from which it follows that $\frac{d}{dt}X_i(t) > 0$. The analogous result holds for coordinates $i \in \mathcal{S}$ with $x_i^* < 0$. In other words, once we have $|X_i(t_0)| \geq \frac{1}{2}|x_i^*|$ for some $t_0 > 0$, then $|X_i(t)| \geq \frac{1}{2}|x_i^*|$ continues to hold for $t \geq t_0$. With this, we can define $I_1 \in \mathcal{S}$ to be the last coordinate that crosses this threshold, that is for which $\frac{|X_i(t)|}{|x_i^*|} \geq \frac{1}{2}$ and assume without loss of generality that $x_{I_1}^* > 0$. By definition, we have $|X_{I_1}(t)| \leq \frac{1}{2}x_{I_1}^*$ for all $t \leq T_1$. For any $j \notin \mathcal{S}$ with $X_j(t) \geq 0$, we have

$$\nabla f(\mathbf{X}(t))_j \geq -(9 + 27\sigma)\sqrt{\frac{\log n}{m}}X_j(t).$$

As before, we can use Lemma 6.1 to obtain

$$\nabla F(\mathbf{X}(t))_j \geq -\frac{1}{4\sqrt{2}} \frac{c_\star^2 \|\mathbf{X}_\mathcal{S}(t)\|_1}{8k}.$$

The analogous result holds for $X_j(t) < 0$, which shows that, for any $j \notin \mathcal{S}$,

$$|\nabla F(\mathbf{X}(t))_j| \leq \frac{1}{4\sqrt{2}} |\nabla F(\mathbf{X}(t))_{I_1}| \quad (6.17)$$

holds for all $t \leq T_1$. The idea is that $X_{I_1}(t)$ and $X_j(t)$ both grow (approximately) exponentially, but at different (time-varying) rates. By the definition of I_1 , we have $X_{I_1}(T_1) = \frac{1}{2}x_{I_1}^*$, and $X_j(T_1)$ can be made arbitrarily small by choosing a sufficiently small parameter β .

To make this rigorous, let $T_\beta = \inf\{t > 0 : \mathbf{X}_{I_1}(t) \geq \beta\}$ be the time when $\mathbf{X}_{I_1}(t)$ first reaches β . Then, since $\sqrt{X_{I_1}(t)^2 + \beta^2} \geq X_{I_1}(t)$, we can bound

$$X_{I_1}(t) \geq \beta \exp\left(-\int_{T_\beta}^t \nabla F(\mathbf{X}(s)) ds\right) \Rightarrow \exp\left(-\int_{T_\beta}^t \nabla F(\mathbf{X}(s)) ds\right) \leq \frac{X_{I_1}(t)}{\beta}.$$

Recalling the bound (6.17), we have for $j \notin S$,

$$|X_j(T_\beta)| \leq \beta, \quad \left|\frac{d}{dt}X_j(t)\right| \leq \frac{1}{4\sqrt{2}}\sqrt{X_j(t)^2 + \beta^2} |\nabla F(\mathbf{X}(t))_{I_1}|.$$

Then, since $\sqrt{X_j(t)^2 + \beta^2} \leq \sqrt{2}|X_j(t)|$ for $X_j(t) \geq \beta$, we can bound

$$X_j(t) \leq \beta \exp\left(-\frac{1}{4} \int_{T_\beta}^t \nabla F(\mathbf{X}(s)) ds\right) \leq \beta \left(\frac{X_{I_1}(t)}{\beta}\right)^{1/4} \leq \beta^{3/4},$$

for all $t \leq T_1$, where we used the fact that $X_{I_1}(t) \leq \frac{1}{2}x_{I_1}^* \leq 1$ for $t \leq T_1$. As this holds for every $j \notin S$, we have, for all $t \leq T_1$,

$$\|\mathbf{X}_{S^c}(t)\|_1 \leq n\beta^{3/4} = \delta_1$$

Stage (ii), part (a): $T_1 < t \leq T_2$, bound $D_\phi(\mathbf{x}^*, \mathbf{X}(t))$

As in Stage (i), we first assume that we have already shown $\|\mathbf{X}_{S^c}(t)\|_1 \leq \delta$ for all $t \leq T_2$. As before, we can use Lemma 6.2 to derive the bound (6.14). By the definition of T_1 , we have $|X_i(t)| \geq \frac{1}{2}|x_i^*|$ for all $i \in [n]$. By Lemma 6.3 we also have $X_i(t)x_i^* \geq 0$ (since we assumed $x_{I_0}^* > 0$), so the assumptions for inequality (4.2) of Lemma 4.1 are satisfied, and we can bound

$$D_\phi(\mathbf{x}^*, \mathbf{X}(t)) \leq \frac{\sqrt{k}}{c_\star} \|\mathbf{X}_S(t) - \mathbf{x}_S^*\|_2^2 + \|\mathbf{X}_{S^c}(t)\|_1 \leq \frac{2\sqrt{k}}{c_\star} \|\mathbf{X}(t) - \mathbf{x}^*\|_2^2,$$

where we used that $\|\mathbf{X}_{S^c}(t)\|_1 \leq \delta \leq \frac{\sqrt{k}}{c_\star} \|\mathbf{X}(t) - \mathbf{x}^*\|_2^2$. With this, inequality (6.14) reads

$$\frac{d}{dt}D_\phi(\mathbf{x}^*, \mathbf{X}(t)) \leq -\frac{c_\star}{4\sqrt{k}}D_\phi(\mathbf{x}^*, \mathbf{X}(t)).$$

Using the fact that $\|\mathbf{X}(t) - \mathbf{x}^*\|_2^2 \leq 3$ by Lemma 6.3, we can bound $D_\phi(\mathbf{x}^*, \mathbf{X}(T_1)) \leq \frac{6\sqrt{k}}{c_\star}$. Hence, we have for $T_1 \leq t \leq T_2$,

$$D_\phi(\mathbf{x}^*, \mathbf{X}(t)) \leq D_\phi(\mathbf{x}^*, \mathbf{X}(T_1)) \exp\left(-\frac{c_\star}{4\sqrt{k}}(t - T_1)\right) \leq \frac{6\sqrt{k}}{c_\star} \exp\left(-\frac{c_\star}{4\sqrt{k}}(t - T_1)\right).$$

Together with the fact that $\|\mathbf{X}(t)\|_\infty^2 \leq 2$ by Lemma 6.3, bound (4.1) of Lemma 4.1 implies

$$D_\phi(\mathbf{x}^*, \mathbf{X}(t)) > \frac{1}{3} \|\mathbf{X}(t) - \mathbf{x}^*\|_2^2 \geq \frac{c^2 c_* \sqrt{k} \delta}{3}$$

for $t \leq T_2$, so we have $T_2 - T_1 \leq \frac{4\sqrt{k}}{c_*} \log \frac{18}{c^2 c_*^2 \delta}$.

Stage (ii), part (b): $T_1 < t \leq T_2$, bound $\|\mathbf{X}_{\mathcal{S}^c}(t)\|_1$

Recall that $|X_j(t)| \leq \beta^{3/4}$ for all $j \notin \mathcal{S}$ and $t \leq T_1$. Further, as in the previous stage we can use Lemma 6.1 to show

$$|\nabla F(\mathbf{X}(t))_j| \leq \frac{c_* \|\mathbf{X}_{\mathcal{S}^c}(t)\|_1}{32k} \leq \frac{c_*}{16\sqrt{2k}},$$

where we used that $\|\mathbf{X}_{\mathcal{S}^c}(t)\|_1 \leq \sqrt{k} \|\mathbf{X}_{\mathcal{S}^c}(t)\|_2 \leq \sqrt{2k}$. Noting $\sqrt{x^2 + \beta^2} \leq \sqrt{2}x$ for $x \geq \beta$, we can bound, for $t \leq T_2$,

$$|X_j(t)| \leq \beta^{3/4} \exp\left(\frac{c_*}{16\sqrt{k}}(T_2 - T_1)\right) \leq \beta^{3/4} \exp\left(\frac{c_*}{16\sqrt{k}} \frac{4\sqrt{k}}{c_*} \log \frac{18}{c^2 c_*^2 \delta}\right) \leq \beta^{3/4} \left(\frac{18}{c^2 c_*^2 \delta}\right)^{1/4} \leq \frac{\delta}{n},$$

provided that $\beta \leq (c^2 c_*^2 / 18)^{2/n^3}$, where we used the definition $\delta = \sqrt{n\beta}$. This completes the proof that $\|\mathbf{X}_{\mathcal{S}^c}(t)\|_1 \leq \delta$ for all $t \leq T_2$. \square

7. Conclusion

In this paper, we establish a general theory of mirror descent with early stopping for the problem of noisy sparse phase retrieval. We provide a full convergence analysis in both continuous and discrete time and we show that, when equipped with the hypentropy mirror map and with proper initialization, early-stopped mirror descent achieves a nearly minimax-optimal rate of convergence, provided the number of measurements m is sufficiently large and the minimum (in modulus) non-zero signal component is on the order of $\|\mathbf{x}^*\|_2 / \sqrt{k}$. These conditions have been previously considered to establish theoretical results for sparse phase retrieval [25, 37, 43].

Previous procedures achieving a nearly minimax-optimal rate of convergence include empirical risk minimization with sparsity constraint [19], which does not lead to a tractable algorithm, and TWF [5], which relies on thresholding steps to promote sparsity. In contrast, unlike most existing algorithms designed to solve sparse phase retrieval, mirror descent does not rely on added regularization terms or thresholding steps to promote sparsity, and no *a priori* knowledge of the sparsity k or the noise level σ is needed to run the algorithm. Our numerical simulations attest that a standard data-dependent stopping rule that does not require knowledge of k or σ yields results that validate our theoretical findings.

Most of the literature on early stopping for iterative gradient-based methodologies has focused on the convex setting of ridge regression and kernel methods via Euclidean gradient descent and boosting. Our results establish connections between early stopping and sparsity in the non-convex setting of sparse phase retrieval via the general framework of mirror descent. The potential-based analysis that we present unveils and exploits a quantitative version of variational coherence that is satisfied along the path traced by the iterates of mirror descent, and it might inspire similar approaches to investigate implicit regularization via early stopping in other non-convex settings in statistical inference.

Data Availability Statement

No new data were generated or analyzed in support of this research.

Funding

EPSRC and MRC OxWaSP CDT Programme (grant number EP/L016710/1 to F.W.); Alan Turing Institute (EPSRC grant EP/N510129/1 to P.R.).

REFERENCES

1. ARORA, S., COHEN, N., HU, W. & LUO, Y. (2019) Implicit regularization in deep matrix factorization. *Adv. Neural Inform. Process. Syst.*, **32**, 7411–7422.
2. AUDIBERT, J.-Y., BUBECK, S. & LUGOSI, G. (2013) Regret in online combinatorial optimization. *Math. Oper. Res.*, **39**, 31–45.
3. BECK, A. & TEBoulLE, M. (2003) Mirror descent and nonlinear projected subgradient methods for convex optimization. *Oper. Res. Lett.*, **31**, 167–175.
4. BUBECK, S. (2015) Convex optimization: Algorithms and complexity. *Found. Trends Mach. Learn.*, **8**, 231–358.
5. CAI, T., LI, X. & MA, Z. (2016) Optimal rates of convergence for noisy sparse phase retrieval via thresholded Wirtinger flow. *Ann. Statist.*, **44**, 2221–2251.
6. CANDÈS, E. J., LI, X. & SOLTANOLKOTABI, M. (2015) Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Trans. Inform. Theory*, **61**, 1985–2007.
7. CHAFAÏ, D., GUÉDON, O., LECUÉ, G. & PAJOR, A. (2012) *Interactions Between Compressed Sensing Random Matrices and High Dimensional Geometry*. Paris: Société Mathématique de France.
8. CHEN, Y., CHI, Y. & GOLDSMITH, A. J. (2015) Exact and stable covariance estimation from quadratic sampling via convex programming. *IEEE Trans. Inform. Theory*, **61**, 4034–4059.
9. DAINTY, J. C. & FIENUP, J. R. (1987) Phase retrieval and image reconstruction for astronomy. *Image Recovery: Theory and Application*, (H. STARK ed.), chap. 7. New York, Academic Press, pp. 231–275.
10. GHADIMI, S. & LAN, G. (2013) Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM J. Optim.*, **23**, 2341–2368.
11. GHAI, U., HAZAN, E. & SINGER, Y. (2020) Exponentiated gradient meets gradient descent. *International Conference on Algorithmic Learning Theory*, vol. **117**. Proceedings of Machine Learning Research, pp. 386–407.
12. GUNASEKAR, S., WOODWORTH, B., BHOJANAPALLI, S., NEYSHABUR, B. & SREBRO, N. (2017) Implicit regularization in matrix factorization. *Adv. Neural Inform. Process. Syst.*, **30**, 6151–6159.
13. HAND, P. & VORONINSKI, V. (2016) Compressed sensing from phaseless Gaussian measurements via linear programming in the natural parameter spaces. arXiv preprint arXiv:1611.05985.
14. HOFF, P. D. (2017) Lasso, fractional norm and structured sparse estimation using a Hadamard product parametrization. *Comput. Statist. Data Anal.*, **115**, 186–198.
15. JAGANATHAN, K., ELDAR, Y. C. & HASSIBI, B. (2016) Phase retrieval: An overview of recent developments. *Optical Compressive Imaging*, (A. STERN ed.), chap. **13**. Boca Raton, FL: Taylor Francis Group, pp. 263–296.
16. JAGATAP, G. & HEGDE, C. (2019) Sample-efficient algorithms for recovering structured signals from magnitude-only measurements. *IEEE Trans. Inform. Theory*, **65**, 4434–4456.
17. KIVINEN, J. & WARMUTH, M. K. (1997) Exponentiated gradient versus gradient descent for linear predictors. *Inform. and Comput.*, **132**, 1–63.
18. KOTÉŁOWSKI, W. & NEU, G. (2019) Bandit principal component analysis. *Conference on Learning Theory*, vol. **99**. Proceedings of Machine Learning Research, pp. 1994–2024.
19. LECUÉ, G. & MENDELSON, S. (2015) Minimax rate of convergence and the performance of empirical risk minimization in phase recovery. *Electron. J. Probab.*, **20**, 1–29.

20. LI, X. & VORONINSKI, V. (2013) Sparse signal recovery from quadratic measurements via convex programming. *SIAM J. Math. Anal.*, **45**, 3019–3033.
21. LI, Y., MA, T. & ZHANG, H. (2018) Algorithmic regularization in over-parametrized matrix sensing and neural networks with quadratic activation. *Conference on Learning Theory*, vol. **75**. Proceedings of Machine Learning Research, pp. 2–47.
22. MILLANE, R. (1990) Phase retrieval in crystallography and optics. *J. Opt. Soc. Am. A*, **7**, 394–411.
23. NEMIROVSKI, A., JUDITSKY, A., LAN, G. & SHAPIRO, A. (2009) Robust stochastic approximation approach to stochastic programming. *SIAM J. Optim.*, **19**, 1574–1609.
24. NEMIROVSKI, A. & YUDIN, D. B. (1983) *Problem Complexity and Method Efficiency in Optimization*. New York: Wiley.
25. NETRAPALLI, P., JAIN, P. & SANGHAVI, S. (2015) Phase retrieval using alternating minimization. *IEEE Trans. Signal Process.*, **63**, 4814–4826.
26. OHLSSON, H., YANG, A. Y., DONG, R. & SASTRY, S. S. (2012) CPRL—An extension of compressive sensing to the phase retrieval problem. *Adv. Neural Inform. Process. Syst.*, **25**, 1367–1375.
27. RASKUTTI, G., WAINWRIGHT, M. J. & YU, B. (2014) Early stopping and non-parametric regression: an optimal data-dependent stopping rule. *J. Mach. Learn. Res.*, **15**, 335–366.
28. SCHECHTMAN, Y., BECK, A. & ELDAR, Y. C. (2014) GESPAR: efficient phase retrieval of sparse signals. *IEEE Trans. Signal Process.*, **62**, 928–938.
29. SCHECHTMAN, Y., ELDAR, Y. C., COHEN, O., CHAPMAN, H. N., MIAO, J. & SEGEV, M. (2014) Phase retrieval with application to optical imaging: A contemporary overview. *IEEE Signal Process. Mag.*, **32**, 87–109.
30. SCHNITER, P. & RANGAN, S. (2015) Compressive phase retrieval via generalized approximate message passing. *IEEE Trans. Signal Process.*, **63**, 1043–1055.
31. SHALEV-SCHWARTZ, S. (2011) Online learning and online convex optimization. *Found. Trends Mach. Learn.*, **4**, 107–194.
32. SUN, J., QU, Q. & WRIGHT, J. (2018) A geometric analysis of phase retrieval. *Found. Comput. Math.*, **18**, 1131–1198.
33. VAŠKEVIČIUS, T., KANADE, V. & REBESCHINI, P. (2019) Implicit regularization for optimal sparse recovery. *Adv. Neural Inform. Process. Syst.*, **32**, 2968–2979.
34. VAŠKEVIČIUS, T., KANADE, V. & REBESCHINI, P. (2020) The statistical complexity of early-stopped mirror descent. *Adv. Neural Inform. Process. Syst.*, **33**, 253–264.
35. VERSHYNIN, R. (2012) Introduction to the non-asymptotic analysis of random matrices. in *Compressed Sensing, Theory and Applications*, ed. by Y. ELDAR, & G. KUTYNIOK, chap. **5**, pp. 210–268. Cambridge University Press, Cambridge.
36. WANG, G., GIANNAKIS, G. B. & ELDAR, Y. C. (2017) Solving systems of random quadratic equations via truncated amplitude flow. *IEEE Trans. Inform. Theory*, **64**, 773–794.
37. WANG, G., ZHANG, L., GIANNAKIS, G. B., AKÇAKAYA, M. & CHEN, J. (2018) Sparse phase retrieval via truncated amplitude flow. *IEEE Trans. Signal Process.*, **66**, 479–491.
38. WEI, Y., YANG, F. & WAINWRIGHT, M. J. (2019) Early stopping for kernel boosting algorithms: a general analysis with localized complexities. *IEEE Trans. Inform. Theory*, **65**, 6685–6703.
39. WU, F. & REBESCHINI, P. (2020) A continuous-time mirror descent approach to sparse phase retrieval. *Adv. Neural Inform. Process. Syst.*, **33**, 20192–20203.
40. WU, F. & REBESCHINI, P. (2021) Hadamard Wirtinger flow for sparse phase retrieval. *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, vol. **130**. Proceedings of Machine Learning Research, pp. 982–990.
41. YAO, Y., ROSASCO, L. & CAPONNETTO, A. (2007) On early stopping in gradient descent learning. *Constr. Approx.*, **26**, 289–315.
42. YUAN, Z., WANG, H. & WANG, Q. (2019) Phase retrieval via sparse Wirtinger flow. *J. Comput. Appl. Math.*, **355**, 162–173.
43. ZHANG, L., WANG, G., GIANNAKIS, G. B. & CHEN, J. (2018) Compressive phase retrieval via reweighted amplitude flow. *IEEE Trans. Signal Process.*, **66**, 5029–5040.

- 44. ZHANG, T. & YU, B. (2005) Boosting with early stopping: Convergence and consistency. *Ann. Stat.*, **33**, 1538–1579.
- 45. ZHAO, P., YANG, Y. & HE, Q.-C. (2019) Implicit regularization via Hadamard product over-parametrization in high-dimensional linear regression. arXiv preprint arXiv:1903.09367.
- 46. ZHOU, Z., MERTIKOPOULOS, P., BAMBOS, N., BOYD, S. P. & GLYNN, P. W. (2017) Stochastic mirror descent in variationally coherent optimization problems. *Adv. Neural Inform. Process. Syst.*, **30**, 7040–7049.
- 47. ZHOU, Z., MERTIKOPOULOS, P., BAMBOS, N., BOYD, S. P. & GLYNN, P. W. (2020) On the convergence of mirror descent beyond stochastic convex programming. *SIAM J. Optim.*, **30**, 687–716.

The supplementary article is organized as follows. We prove Theorem 4.4 in Section A. In Section B, we prove Lemmas 6.1–6.3 from Section 6.2. Further technical lemmas are provided in Section C.

A. Proof of Theorem 4.4

As in the proof of Theorem 4.3, we will assume $\|\mathbf{x}^*\|_2 = 1$ and $(\mathbf{X}^0)^\top \mathbf{x}^* \geq 0$ for notational simplicity. The proof of Theorem 4.4 largely follows the same steps as the proof of Theorem 4.3, and we can reuse many of the bounds shown before. We first derive some generic bounds on how much X_i^t can increase or decrease in one iteration. The update (5.2) can be written as (see e.g. Theorem 24 in [11])

$$\begin{aligned} X_i^{t+1} - X_i^t &= \frac{\sqrt{(X_i^t)^2 + \beta^2} + X_i^t}{2} (\exp(-\eta \nabla F(\mathbf{X}^t)_i) - 1) \\ &\quad - \frac{\sqrt{(X_i^t)^2 + \beta^2} - X_i^t}{2} (\exp(\eta \nabla F(\mathbf{X}^t)_i) - 1). \end{aligned} \quad (\text{A.1})$$

Assuming that the conditions for the simplified bound in Lemma 6.1 are satisfied, we can bound

$$|\nabla F(\mathbf{X}^t)_i - \nabla f(\mathbf{X}^t)_i| \leq \frac{0.1c_\star((\mathbf{X}^t)^\top \mathbf{x}^*)}{\sqrt{k}} \leq 0.1((\mathbf{X}^t)^\top \mathbf{x}^*)|x_i^*|, \quad (\text{A.2})$$

with probability $1 - c_p n^{-10}$, where we used (6.6) to bound $\|\mathbf{X}_S^t\|_1 \leq \frac{\sqrt{k}}{c_\star}((\mathbf{X}^t)^\top \mathbf{x}^*)$ in the first inequality, and for the second inequality we used the assumption $x_{\min}^* \leq \frac{c_\star}{\sqrt{k}}$. Recalling the definition of the population gradient

$$\nabla f(\mathbf{x}) = (3\|\mathbf{x}\|_2^2 - 1)\mathbf{x} - 2(\mathbf{x}^\top \mathbf{x}^*)\mathbf{x}^*,$$

an application of the triangle inequality yields

$$|\nabla F(\mathbf{X}^t)_i| \leq |(3\|\mathbf{X}^t\|_2^2 - 1)X_i^t| + |2.1((\mathbf{X}^t)^\top \mathbf{x}^*)x_i^*| \leq 7.1\sqrt{2}, \quad (\text{A.3})$$

provided that $\|\mathbf{X}^t\|_2 \leq \sqrt{2}$, where we used that $(\mathbf{X}^t)^\top \mathbf{x}^* \leq \|\mathbf{X}^t\|_2 \|\mathbf{x}^*\|_2$ by the Cauchy–Schwarz inequality. Hence, we have $|\eta \nabla F(\mathbf{X}^t)_i| \leq \frac{1}{2}$, provided that $\eta \leq \frac{1}{14.2\sqrt{2}}$.

To bound the exponentials in (A.1), we use that $e^x \leq 1 + \frac{3}{2}x$ for $0 \leq x \leq \frac{1}{2}$, and $1 + x \leq e^x$ for $-\frac{1}{2} \leq x \leq 0$. Without loss of generality, we assume $X_i^t \geq 0$; analogous bounds can be derived for $X_i^t < 0$. If $\nabla F(\mathbf{X}^t)_i < 0$ (i.e. $X_i^{t+1} > X_i^t$), then we can bound

$$\begin{aligned} X_i^{t+1} - X_i^t &\leq \frac{\sqrt{(X_i^t)^2 + \beta^2} + X_i^t}{2} \left(-\frac{3}{2}\eta \nabla F(\mathbf{X}^t)_i\right) - \frac{\sqrt{(X_i^t)^2 + \beta^2} - X_i^t}{2} \eta \nabla F(\mathbf{X}^t)_i \\ &\leq -\frac{3}{2}\eta \nabla F(\mathbf{X}^t)_i \sqrt{(X_i^t)^2 + \beta^2}. \end{aligned}$$

If $\nabla F(\mathbf{X}^t)_i > 0$, we obtain the analogous lower bound

$$X_i^{t+1} - X_i^t \geq -\frac{5}{4}\eta \nabla F(\mathbf{X}^t)_i \sqrt{(X_i^t)^2 + \beta^2}.$$

Combining both bounds yields an upper bound on how much X_i^t can increase or decrease in one iteration:

$$|X_i^{t+1} - X_i^t| \leq \frac{3}{2}\eta |\nabla F(\mathbf{X}^t)_i| \sqrt{(X_i^t)^2 + \beta^2}. \quad (\text{A.4})$$

Similarly, we can use the bounds $1 + x \leq e^x$ for $0 \leq x \leq \frac{1}{2}$ and $e^x \leq 1 + \frac{x}{2}$ for $-\frac{1}{2} \leq x \leq 0$ to obtain the lower bound

$$|X_i^{t+1} - X_i^t| \geq \frac{1}{2}\eta |\nabla F(\mathbf{X}^t)_i| |X_i^t|. \quad (\text{A.5})$$

Proof of Theorem 4.4 The proof of Theorem 4.4 relies on the following identity, which follows from the definition of the Bregman divergence and the mirror descent update (3.2), to show that $D_\Phi(\mathbf{x}^*, \mathbf{X}^t)$ decreases as claimed:

$$D_\Phi(\mathbf{x}^*, \mathbf{X}^{t+1}) - D_\Phi(\mathbf{x}^*, \mathbf{X}^t) = -\eta \langle \nabla F(\mathbf{X}^t), \mathbf{X}^t - \mathbf{x}^* \rangle + D_\Phi(\mathbf{X}^t, \mathbf{X}^{t+1}). \quad (\text{A.6})$$

The first term in (A.6) has been bounded in the continuous-time case, and we can reuse those bounds. The second term measures the distance between the iterates \mathbf{X}^t and \mathbf{X}^{t+1} in terms of the Bregman divergence due to the discretization, and we will derive bounds for the second term similar to those shown in the continuous-time case for the first term. The function Φ is $1/2$ -strongly convex with respect to the ℓ_2 -norm on the ℓ_2 -ball $\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2^2 \leq 2\}$ provided $\beta \leq \sqrt{2}$, since the Hessian $\nabla^2 \Phi(\mathbf{x})$ is a diagonal matrix with entries

$$\nabla^2 \Phi(\mathbf{x})_{ii} = \frac{1}{\sqrt{x_i^2 + \beta^2}} \geq \frac{1}{\sqrt{2 + \beta^2}} \geq \frac{1}{2}.$$

Hence, we can bound

$$\begin{aligned} D_\Phi(\mathbf{X}^t, \mathbf{X}^{t+1}) &= \Phi(\mathbf{X}^t) - \Phi(\mathbf{X}^{t+1}) - \langle \nabla \Phi(\mathbf{X}^{t+1}), \mathbf{X}^t - \mathbf{X}^{t+1} \rangle \\ &\leq \langle \nabla \Phi(\mathbf{X}^t) - \nabla \Phi(\mathbf{X}^{t+1}), \mathbf{X}^t - \mathbf{X}^{t+1} \rangle - \frac{1}{4} \|\mathbf{X}^t - \mathbf{X}^{t+1}\|_2^2 \\ &\leq \langle \eta \nabla F(\mathbf{X}^t), \mathbf{X}^t - \mathbf{X}^{t+1} \rangle - \frac{1}{4} \|\mathbf{X}^t - \mathbf{X}^{t+1}\|_2^2 \\ &\leq \eta \|\nabla F(\mathbf{X}^t)\|_2 \|\mathbf{X}^t - \mathbf{X}^{t+1}\|_2 - \frac{1}{4} \|\mathbf{X}^t - \mathbf{X}^{t+1}\|_2^2 \\ &\leq \eta^2 \|\nabla F(\mathbf{X}^t)\|_2^2, \end{aligned} \quad (\text{A.7})$$

where the first inequality follows from strong convexity of Φ , the second inequality from the definition of the mirror descent update, the third inequality from the Cauchy–Schwarz inequality, and for the last inequality we optimized the quadratic function in $\|\mathbf{X}^t - \mathbf{X}^{t+1}\|_2$.

As in continuous time, we divide the proof of Theorem 4.4 into two stages bounded by

$$T_1 = \inf \left\{ t > 0 : \min_{i \in \mathcal{S}} \frac{|X_i^t|}{|x_i^*|} > \frac{1}{2} \right\} \text{ and}$$

$$T_2 = \inf \left\{ t > 0 : \frac{\|\mathbf{X}^t - \mathbf{x}^*\|_2^2}{\|\mathbf{x}^*\|_2^2} \leq c^2 \max \left\{ c_* \sqrt{k} \delta, \frac{\sigma^2}{\|\mathbf{x}^*\|_2^4} \frac{k \log n}{m} \right\} \right\},$$

respectively, where $\delta = \sqrt{n\beta}$, and (a) bound the length of the stage by using (A.6) to bound T_i and (b) show that off-support coordinates $\|\mathbf{X}_{\mathcal{S}^c}^t\|_1$ stay sufficiently small. Throughout the proof we will assume that the inequalities in Lemmas 6.1–6.3 are satisfied, which happens with probability at least $1 - c_p n^{-10}$. As in the continuous-time case, the conditions for the simplified bounds in Lemmas 6.1 and 6.2 are satisfied for $t \leq T_2$.

Stage (i), part (a): $t \leq T_1$, bound T_1

Assume for now that we have already shown $\|\mathbf{X}_{\mathcal{S}^c}^t\|_1 < \delta_1 = 2n\beta^{3/4}$ for all $t \leq T_1$. As in continuous time, we consider the following two cases:

- **Case 1:** $\|\mathbf{X}^t\|_2^2 \leq \frac{2}{5}$

In this case, we have shown in (6.12) that

$$-\eta \langle \nabla F(\mathbf{X}^t), \mathbf{X}^t - \mathbf{x}^* \rangle \leq -\frac{\eta}{4} ((\mathbf{X}^t)^\top \mathbf{x}^*).$$

It remains to bound the term $D_\phi(\mathbf{X}^t, \mathbf{X}^{t+1})$. Using the bounds (A.3), (A.7) and the inequality $(a + b)^2 \leq 2a^2 + 2b^2$, we can bound

$$\begin{aligned} D_\phi(\mathbf{X}^t, \mathbf{X}^{t+1}) &\leq 2\eta^2 \sum_{i=1}^n (3\|\mathbf{X}^t\|_2^2 - 1)^2 (X_i^t)^2 + 2.1^2 ((\mathbf{X}^t)^\top \mathbf{x}^*)^2 (x_i^*)^2 \\ &\leq 2\eta^2 ((\mathbf{X}^t)^\top \mathbf{x}^*)^2 \sum_{i=1}^n (12(X_i^t)^2 + 2.1^2 (x_i^*)^2) \\ &\leq \frac{\eta}{8} ((\mathbf{X}^t)^\top \mathbf{x}^*) \end{aligned}$$

for $\eta \leq \frac{1}{94}$, where we used (6.4) of Lemma 6.3 to bound $3\|\mathbf{X}^t\|_2^2 - 1 \leq \sqrt{3} \cdot 2((\mathbf{X}^t)^\top \mathbf{x}^*)$ for the second inequality, and for the last inequality we used the assumptions $\|\mathbf{X}^t\|_2^2 \leq 2/5$ and $\|\mathbf{x}^*\|_2 = 1$, and the Cauchy–Schwarz inequality to bound $(\mathbf{X}^t)^\top \mathbf{x}^* \leq \sqrt{2/5}$. Plugging this bound into (A.6), we have

$$D_\phi(\mathbf{x}^*, \mathbf{X}^{t+1}) - D_\phi(\mathbf{x}^*, \mathbf{X}^t) \leq -\frac{\eta}{8} ((\mathbf{X}^t)^\top \mathbf{x}^*) \leq -\frac{c_* \eta}{16\sqrt{k}}, \quad (\text{A.8})$$

where for the second inequality we used (6.6).

- **Case 2:** $\|\mathbf{X}^t\|_2^2 \geq \frac{2}{5}$

In this case, we have already shown that

$$-\eta \langle \nabla F(\mathbf{X}^t), \mathbf{X}^t - \mathbf{x}^* \rangle \leq -\frac{\eta}{25} \quad (\text{A.9})$$

if $4\|\mathbf{X}^t\|_2^2 - 2 - \|\mathbf{X}^t - \mathbf{x}^\star\|_2^2 \leq 0$, and

$$-\eta \langle \nabla F(\mathbf{X}^t), \mathbf{X}^t - \mathbf{x}^\star \rangle \leq -\frac{\eta}{2} \|\mathbf{X}^t - \mathbf{x}^\star\|_2^2 \quad (\text{A.10})$$

if $4\|\mathbf{X}^t\|_2^2 - 2 - \|\mathbf{X}^t - \mathbf{x}^\star\|_2^2 > 0$. It remains to bound $D_\phi(\mathbf{X}^t, \mathbf{X}^{t+1})$ in terms of the expressions in (A.9) and (A.10). As the bound corresponding to (A.9) can be obtained the same (but easier) way as for (A.10), we only show the computation for the latter case.

Substituting $\mathbf{Z} = \mathbf{X}^t - \mathbf{x}^\star$, we can write

$$\begin{aligned} |\nabla f(\mathbf{X}(t))_i| &= |(3\|\mathbf{X}^t\|_2^2 - 1)Z_i + (3\|\mathbf{Z} + \mathbf{x}^\star\|_2^2 - 1 - 2((\mathbf{Z} + \mathbf{x}^\star)^\top \mathbf{x}^\star))x_i^\star| \\ &= |(3\|\mathbf{X}^t\|_2^2 - 1)Z_i + (\mathbf{Z}^\top (3\mathbf{X}^t + \mathbf{x}^\star))x_i^\star| \\ &\leq 5|Z_i| + (1 + 3\sqrt{2})\|\mathbf{Z}\|_2|x_i^\star|, \end{aligned}$$

where we used that $\|\mathbf{X}^t\|_2 \leq \sqrt{2}$ by (6.3) of Lemma 6.3. By Lemma 6.1, we can bound

$$|\nabla F(\mathbf{X}^t)_i - \nabla f(\mathbf{X}^t)_i| \leq \frac{0.1c_\star \|\mathbf{X}^t - \mathbf{x}^\star\|_2}{\sqrt{k}} \leq 0.1\|\mathbf{X}^t - \mathbf{x}^\star\|_2|x_i^\star|,$$

for $t \leq T_2$, so that we can bound

$$|\nabla F(\mathbf{X}^t)_i| \leq 5.4\|\mathbf{X}^t - \mathbf{x}^\star\|_2|x_i^\star| + 5|X_i^t - x_i^\star|.$$

As in the previous case, we can use this bound together with (A.3) and (A.7) to obtain

$$\begin{aligned} D_\phi(\mathbf{X}^t, \mathbf{X}^{t+1}) &\leq 2\eta^2 \sum_{i=1}^n (5.4^2 \|\mathbf{X}^t - \mathbf{x}^\star\|_2^2 (x_i^\star)^2 + 5^2 (X_i^t - x_i^\star)^2) \\ &\leq 2\eta^2 \|\mathbf{X}^t - \mathbf{x}^\star\|_2^2 \left(5^2 + \sum_{i=1}^n 5.4^2 (x_i^\star)^2 \right) \\ &\leq \frac{\eta}{4} \|\mathbf{X}^t - \mathbf{x}^\star\|_2^2 \end{aligned}$$

for $\eta \leq \frac{1}{434}$, where we used that $\|\mathbf{x}^\star\|_2 = 1$. Plugging this bound into (A.6), we have

$$D_\phi(\mathbf{x}^\star, \mathbf{X}^{t+1}) - D_\phi(\mathbf{x}^\star, \mathbf{X}^t) \leq -\frac{\eta}{4} \|\mathbf{X}^t - \mathbf{x}^\star\|_2^2. \quad (\text{A.11})$$

We can now bound T_1 . Define

$$T_0 = \inf \left\{ t > 0 : \|\mathbf{X}^t\|_2^2 > \frac{2}{5} \right\},$$

as the time until which Case 1 holds. Then, the bound (A.8) from Case 1 shows that

$$\min\{T_1, T_0\} \leq \frac{16\sqrt{k}D_\phi(\mathbf{x}^\star, \mathbf{X}^0)}{c_\star\eta} \leq \frac{16}{c_\star\eta} k \log \frac{1}{\beta} + \frac{16\sqrt{k}}{c_\star\eta} \leq \frac{c_2}{2\eta} k \log \frac{1}{\beta},$$

provided that $c_2 \geq \frac{32}{c_\star} + \frac{32}{c_\star\sqrt{k} \log \frac{1}{\beta}}$, where we used (6.8) to bound the Bregman divergence $D_\phi(\mathbf{x}^\star, \mathbf{X}^0)$.

If $T_1 < T_0$, then we have bounded T_1 as desired. Otherwise, we can use (A.11) and the bound

corresponding to (A.9) to control $T_1 - T_0$. As in the continuous-time case, the bound corresponding to (A.9) can apply for at most $t \leq 50\sqrt{k} \log \frac{1}{\beta} + 50$ iterations. Following the same steps as in the continuous-time case, we can derive the bounds (6.15) and (6.16), which allow us to bound the ℓ_2 -distance $\|\mathbf{X}^t - \mathbf{x}^*\|_2^2$ in terms of the Bregman divergence $D_\phi(\mathbf{x}^*, \mathbf{X}^t)$. Combined with (A.11), this gives

$$D_\phi(\mathbf{x}^*, \mathbf{X}^{t+1}) - D_\phi(\mathbf{x}^*, \mathbf{X}^t) \leq -\frac{\eta}{4} \|\mathbf{X}^t - \mathbf{x}^*\|_2^2 \leq -\frac{c_\star^2 \eta}{16k \log \frac{1}{\beta}} D_\phi(\mathbf{x}^*, \mathbf{X}^t),$$

which shows that $D_\phi(\mathbf{x}^*, \mathbf{X}^t)$ decreases at least linearly at the rate $(1 - c_\star^2 \eta / (16k \log \frac{1}{\beta}))$ for $T_0 \leq t < T_1$. We have

$$D_\phi(\mathbf{x}^*, \mathbf{X}^{T_0}) \leq D_\phi(\mathbf{x}^*, \mathbf{X}^0) \leq \sqrt{k} \log \frac{1}{\beta} + 1,$$

and, for $t \leq T_1$,

$$\frac{c_\star^2}{4k} \leq \|\mathbf{X}^t - \mathbf{x}^*\|_2^2 \leq 2\sqrt{2 + \beta^2} D_\phi(\mathbf{x}^*, \mathbf{X}^t) \leq 3D_\phi(\mathbf{x}^*, \mathbf{X}^t),$$

where for the second inequality we used the bound (4.1) of Lemma 4.1 and $\|\mathbf{X}^t\|_\infty \leq \sqrt{2}$ by Lemma 6.3. This implies

$$T_1 - T_0 \leq \frac{16k \log \frac{1}{\beta}}{c_\star^2 \eta} \log \left(\frac{12}{c_\star^2} k^{\frac{3}{2}} \log \frac{1}{\beta} + \frac{12}{c_\star^2} k \right) \leq \frac{c_2}{2\eta} k \log \left(\frac{1}{\beta} \right) \log \left(k \log \frac{1}{\beta} \right)$$

provided that $c_2 \geq \frac{48}{c_\star^2} + \frac{32}{c_\star^2} \log \frac{24}{c_\star^2}$.

Stage (i), part (b): $t \leq T_1$, **bound** $\|\mathbf{X}_{\mathcal{S}^c}^t\|_1$

As in the continuous-time case, let $I_1 \in \mathcal{S}$ be the last coordinate, which satisfies $|X_{I_1}^t| \geq \frac{1}{2}|x_{I_1}^*|$ for some $t > 0$, and let $j \notin \mathcal{S}$. To simplify notation, assume $X_{I_1}^t, X_j^t > 0$, in which case we can show $\nabla F(\mathbf{X}^t)_{I_1} < 0$ for all $t \leq T_1$ as in the continuous-time case.

As shown in (6.17), we have (using different constants in Lemma 6.1)

$$|\nabla F(\mathbf{X}^t)_j| \leq \frac{\sqrt{2}}{30} |\nabla F(\mathbf{X}^t)_{I_1}| \quad (\text{A.12})$$

for all $t \leq T_1$. The idea of the proof is the same as in continuous time: the growth of both $X_{I_1}^t$ and X_j^t is bounded by exponentials with different (time-varying) rates. Using (A.5), we have

$$X_{I_1}^{t+1} \geq X_{I_1}^t - \frac{1}{2} \eta \nabla F(\mathbf{X}^t)_{I_1} X_{I_1}^t \geq \left(1 - \frac{1}{2} \eta \nabla F(\mathbf{X}^t)_{I_1}\right) X_{I_1}^t, \quad (\text{A.13})$$

and similarly, using (A.4), we have, for $X_j^t \geq \beta$,

$$X_j^{t+1} \leq \left(1 + \frac{3}{2} \sqrt{2} \eta |\nabla F(\mathbf{X}^t)_j|\right) X_j^t. \quad (\text{A.14})$$

Denote $G_t = -\nabla F(\mathbf{X}^t)_{I_1}$, and let $T_\beta = \min\{t : |X_{I_1}^t| \geq \beta\}$. Then, we use (A.13) to bound

$$X_{I_1}^{T_1} \geq X_{I_1}^{T_\beta} \prod_{t=T_\beta}^{T_1} \left(1 + \frac{1}{2} \eta G_t\right),$$

and, since $|X_j^{T_\beta}| \leq 2\beta$ and $\sqrt{(X_j^{T_\beta})^2 + \beta^2} \leq \sqrt{2}|X_j^{T_\beta}|$ for $|X_j^{T_\beta}| \geq \beta$, we can use (A.12) and (A.14) to bound

$$X_j^{T_1} \leq 2\beta \prod_{t=T_\beta}^{T_1} \left(1 + \frac{1}{10}\eta G_t\right) \leq 2\beta \left(\prod_{t=T_\beta}^{T_1} 1 + \frac{1}{2}\eta G_t\right)^{\frac{1}{4}} \leq 2\beta \left(\frac{X_{i_1}^{T_1}}{X_{i_1}^{T_\beta}}\right)^{\frac{1}{4}} \leq 2\beta \left(\frac{1}{2\beta}\right)^{\frac{1}{4}} \leq 2\beta^{\frac{3}{4}},$$

where for the second inequality we used the inequality $\frac{x}{x+1} \leq \log(1+x) \leq x$ for $x \geq 0$, with which we have, for $\eta G_t \leq \frac{1}{2}$,

$$\log\left(1 + \frac{1}{10}\eta G_t\right) \leq \frac{1}{10}\eta G_t \leq \frac{1}{4} \frac{\eta G_t/2}{1 + \eta G_t/2} \leq \frac{1}{4} \log\left(1 + \frac{1}{2}\eta G_t\right).$$

Stage (ii), part (a): $T_1 < t \leq T_2$, bound $D_\phi(\mathbf{x}^*, \mathbf{X}^t)$

As in the continuous-time case, we can use Lemma 4.1 to bound

$$D_\phi(\mathbf{x}^*, \mathbf{X}^t) \leq \frac{2\sqrt{k}}{c_\star} \|\mathbf{X}^t - \mathbf{x}^*\|_2^2,$$

and Lemma 6.3 to bound $D_\phi(\mathbf{x}^*, \mathbf{X}^{T_1}) \leq \frac{6\sqrt{k}}{c_\star}$. With this, inequality (A.11) reads

$$D_\phi(\mathbf{x}^*, \mathbf{X}^{t+1}) - D_\phi(\mathbf{x}^*, \mathbf{X}^t) \leq -\frac{c_\star \eta}{8\sqrt{k}} D_\phi(\mathbf{x}^*, \mathbf{X}^t).$$

Hence, we have for $T_1 \leq t \leq T_2$,

$$D_\phi(\mathbf{x}^*, \mathbf{X}^t) \leq D_\phi(\mathbf{x}^*, \mathbf{X}^{T_1}) \left(1 - \frac{c_\star \eta}{8\sqrt{k}}\right)^{t-T_1} \leq \frac{6\sqrt{k}}{c_\star} \left(1 - \frac{c_\star \eta}{8\sqrt{k}}\right)^{t-T_1}.$$

Together with the fact that $\|\mathbf{X}^t\|_\infty^2 \leq 2$ by Lemma 6.3, bound (4.1) of Lemma 4.1 implies

$$D_\phi(\mathbf{x}^*, \mathbf{X}^t) \geq \frac{1}{3} \|\mathbf{X}^t - \mathbf{x}^*\|_2^2 \geq \frac{c^2 c_\star \sqrt{k} \delta}{3}$$

for $t \leq T_2$, so we can bound

$$T_2 - T_1 \leq \frac{8\sqrt{k}}{c_\star \eta} \log \frac{18}{c^2 c_\star^2 \delta}.$$

Stage (ii), part (b): $T_1 < t \leq T_2$, bound $\|\mathbf{X}_{\mathcal{S}^c}^t\|_1$

Recall that we have shown $|X_j^t| \leq 2\beta^{\frac{3}{4}}$ for all $j \notin \mathcal{S}$ and $t \leq T_1$. As in the continuous-time case, we can use Lemma 6.1 to show

$$|\nabla F(\mathbf{X}^t)_j| \leq \frac{c_\star \|\mathbf{X}_{\mathcal{S}}^t\|_1}{64k} \leq \frac{c_\star}{32\sqrt{2k}},$$

where we used that $\|\mathbf{X}_{\mathcal{S}}^t\|_1 \leq \sqrt{k} \|\mathbf{X}_{\mathcal{S}}^t\|_2 \leq \sqrt{2k}$. As in Stage (i), we can use the fact that $\sqrt{x^2 + \beta^2} \leq \sqrt{2}x$ for $x \geq \beta$ to bound, for $t \leq T_2$,

$$|X_j(t)| \leq 2\beta^{\frac{3}{4}} \left(1 + \frac{3}{\sqrt{2}} \eta \frac{c_\star}{32\sqrt{k}}\right)^{T_2-T_1} \leq 2\beta^{\frac{3}{4}} \exp\left(\frac{\eta c_\star}{32\sqrt{k}} \frac{8\sqrt{k}}{c_\star \eta} \log \frac{18}{c^2 c_\star^2 \delta}\right) \leq 2\beta^{\frac{3}{4}} \left(\frac{18}{c^2 c_\star^2 \delta}\right)^{\frac{1}{4}} \leq \frac{\delta}{n},$$

provided that $\beta \leq (c^2 c_\star^2 / 288)^2 / n^3$, where we used the definition $\delta = \sqrt{n\beta}$. This completes the proof that $\|\mathbf{X}_{\mathcal{S}^c}^t\|_1 \leq \delta$ for all $t \leq T_2$. \square

B. Proof of supporting lemmas

In this section, we provide the proofs of the supporting lemmas stated in Section 6.2. Throughout this section, we will assume $\|\mathbf{x}^\star\|_2 = 1$ for notational simplicity's sake. The general case $\|\mathbf{x}^\star\|_2 \neq 1$ follows by writing $\mathbf{x}^\star = \frac{\mathbf{x}^\star}{\|\mathbf{x}^\star\|_2} \|\mathbf{x}^\star\|_2$, $\mathbf{x} = \frac{\mathbf{x}}{\|\mathbf{x}^\star\|_2} \|\mathbf{x}^\star\|_2$ and $\varepsilon_j = \frac{\varepsilon_j}{\|\mathbf{x}^\star\|_2} \|\mathbf{x}^\star\|_2$ in what follows.

Proof of Lemma 6.1 To prove Lemma 6.1, we will bound the difference $|\nabla F(\mathbf{x})_i - \nabla f(\mathbf{x})_i|$ for any $i \in [n]$. The lemma then follows by taking a union bound. First, we write $\mathbf{w} \in \mathbb{R}^n$ for the vector $\mathbf{x}_{\mathcal{S}}$ padded with zeroes, that is $w_i = x_i$ for $i \in \mathcal{S}$ and $w_i = 0$ otherwise. For convenience, we denote by

$$\nabla \tilde{F}(\mathbf{x}) = \frac{1}{m} \sum_{j=1}^m ((\mathbf{A}_j^\top \mathbf{x})^2 - (\mathbf{A}_j^\top \mathbf{x}^\star)^2) (\mathbf{A}_j^\top \mathbf{x}) \mathbf{A}_j$$

the gradient if the measurements were noiseless. Then, as $\nabla f(\mathbf{x}) = \mathbb{E}[\nabla F(\mathbf{x})] = \mathbb{E}[\nabla \tilde{F}(\mathbf{x})]$ since the random variables $\{\varepsilon_j\}_{j=1}^m$ are centered and independent of $\{\mathbf{A}_j\}_{j=1}^m$, we can decompose and bound the difference by

$$\begin{aligned} |\nabla F(\mathbf{x})_i - \nabla f(\mathbf{x})_i| &\leq |\nabla F(\mathbf{x})_i - \nabla \tilde{F}(\mathbf{x})_i| + |\nabla \tilde{F}(\mathbf{x})_i - \nabla \tilde{F}(\mathbf{w})_i| \\ &\quad + |\nabla \tilde{F}(\mathbf{w})_i - \mathbb{E}[\nabla \tilde{F}(\mathbf{w})_i]| + |\mathbb{E}[\nabla \tilde{F}(\mathbf{w})_i] - \mathbb{E}[\nabla \tilde{F}(\mathbf{x})_i]|, \end{aligned} \quad (\text{B.1})$$

and we will bound the four terms separately.

Step 1: Bound the term $|\nabla F(\mathbf{x})_i - \nabla \tilde{F}(\mathbf{x})_i|$

We begin by bounding the first term of (B.1) by

$$c\sigma \sqrt{\frac{\log n}{m}} \min\left\{\|\mathbf{x}\|_1, 1 + \delta + \sqrt{k}\|\mathbf{x}_{\mathcal{S}} - \mathbf{x}_{\mathcal{S}}^\star\|_2\right\}.$$

To obtain the second bound in the minimum, we write $\mathbf{z} = \mathbf{x} - \mathbf{x}^\star$ and bound

$$\begin{aligned} |\nabla F(\mathbf{x})_i - \nabla \tilde{F}(\mathbf{x})_i| &= \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_j^\top \mathbf{x}) A_{ji} \right| \\ &\leq \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_{j,\mathcal{S}}^\top \mathbf{x}_{\mathcal{S}}) A_{ji} \right| + \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_{j,\mathcal{S}^c}^\top \mathbf{x}_{\mathcal{S}^c}) A_{ji} \right| \\ &\leq \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_{j,\mathcal{S}}^\top \mathbf{z}_{\mathcal{S}}) A_{ji} \right| + \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_{j,\mathcal{S}}^\top \mathbf{x}_{\mathcal{S}}^\star) A_{ji} \right| + \left| \sum_{l \notin \mathcal{S}} x_l \frac{1}{m} \sum_{j=1}^m \varepsilon_j A_{jl} A_{ji} \right| \\ &\leq (\|\mathbf{z}_{\mathcal{S}}\|_1 + \|\mathbf{x}_{\mathcal{S}^c}\|_1) \max_{l \in [n]} \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j A_{jl} A_{ji} \right| + \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_{j,\mathcal{S}}^\top \mathbf{x}_{\mathcal{S}}^\star) A_{ji} \right| \\ &\leq c\sigma \sqrt{\frac{\log n}{m}} (\|\mathbf{z}_{\mathcal{S}}\|_1 + 1 + \delta), \end{aligned}$$

with probability $1 - \frac{c_p}{3}n^{-11}$ for any constant $c > 0$ that is at least the universal constant from Lemma C.5, where we used Hölder's inequality in the penultimate line. The second bound in the minimum follows as $\|\mathbf{z}_S\|_1 \leq \sqrt{k}\|\mathbf{z}_S\|_2$. The first bound in the minimum follows by directly bounding $\frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_j^\top \mathbf{x}) A_{ji}$ using Hölder's inequality and Lemma C.5 as above, without substituting $\mathbf{z} = \mathbf{x} - \mathbf{x}^*$.

Step 2: Bound the term $|\nabla \tilde{F}(\mathbf{x})_i - \nabla \tilde{F}(\mathbf{w})_i|$

Next, we bound the second term of (B.1) by $\frac{c}{2}\delta$. We can write

$$\nabla \tilde{F}(\mathbf{x})_i = \frac{1}{m} \sum_{j=1}^m ((\mathbf{A}_{j,S}^\top \mathbf{x}_S + \mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^3 - (\mathbf{A}_{j,S}^\top \mathbf{x}_S + \mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})(\mathbf{A}_j^\top \mathbf{x}^*)^2) A_{ji}.$$

Then, we have

$$\begin{aligned} \nabla \tilde{F}(\mathbf{x})_i - \nabla \tilde{F}(\mathbf{w})_i &= \frac{3}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^2 (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c}) + \frac{3}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,S}^\top \mathbf{x}_S) (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^2 \\ &\quad + \frac{1}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^3 - \frac{1}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c}) (\mathbf{A}_j^\top \mathbf{x}^*)^2. \end{aligned} \quad (\text{B.2})$$

These four terms can be bounded as follows: for the first term, we have

$$\begin{aligned} \left| \frac{1}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^2 (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c}) \right| &= \left| \sum_{l \notin S} x_l \frac{1}{m} \sum_{j=1}^m A_{ji} A_{jl} (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^2 \right| \\ &\leq \|\mathbf{x}_{S^c}\|_1 \max_{l \notin S} \left| \frac{1}{m} \sum_{j=1}^m A_{ji} A_{jl} (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^2 \right| \\ &\leq \|\mathbf{x}_{S^c}\|_1 \max_{l \notin S} \sqrt{\frac{1}{m} \sum_{j=1}^m A_{ji}^2 A_{jl}^2} \sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^4}, \end{aligned}$$

where we used Hölder's inequality to obtain both inequalities. The first sum is bounded by Lemma C.5: recalling $m \geq c_s(\gamma)k^2 \log^2 n$, we have with probability at least $1 - c_2 n^{-13}$, where c_2 is the universal constant from Lemma C.5,

$$\max_{l \notin S} \left| \frac{1}{m} \sum_{j=1}^m A_{ji}^2 A_{jl}^2 \right| \leq 1 + \frac{1}{k} \leq 2.$$

By Lemma C.3 with $t = 5\sqrt{\log n}$, we can bound, with probability $1 - 4n^{-12.5}$,

$$\sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^4} \leq \frac{1}{\sqrt{m}} \left((3m)^{\frac{1}{4}} + \sqrt{k} + 5\sqrt{\log n} \right)^2 \leq 11$$

for all $\mathbf{x} \in \mathcal{X}$, where we used $5\sqrt{\log n} \leq m^{\frac{1}{4}}$, which holds if $c_s(\gamma) \geq 5^4 \min\{k^{-2}, \log^{-3} n\}$. Put together, this gives

$$\frac{1}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^2 (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c}) \leq c' \delta,$$

where $c' = 11\sqrt{2}$. The other terms in (B.2) can be bounded the same way. For instance, we can write

$$\left| \frac{1}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,S}^\top \mathbf{x}_S) (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^2 \right| = \left| \sum_{l \notin S} x_l \sum_{r \notin S} x_r \frac{1}{m} \sum_{j=1}^m A_{ji} A_{jl} A_{js} (\mathbf{A}_{j,S}^\top \mathbf{x}_S) \right|$$

and, following the same steps as above, we obtain the bounds

$$\begin{aligned} \left| \frac{1}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c}) (\mathbf{A}_j^\top \mathbf{x}^*)^2 \right| &\leq c' \delta, \\ \left| \frac{1}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,S}^\top \mathbf{x}_S) (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^2 \right| &\leq c'' \delta^2, \\ \left| \frac{1}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^3 \right| &\leq c''' \delta^3, \end{aligned}$$

for all $\mathbf{x} \in \mathcal{X}$ with probability $1 - 3(c_2 + 4)n^{-12.5}$. Recalling that $\delta \leq \frac{1}{4}$, we have for any constant $c > 0$ satisfying $4c'\delta + 3c''\delta^2 + c'''\delta^3 \leq \frac{c}{2}\delta$,

$$|\nabla \tilde{F}(\mathbf{x})_i - \nabla \tilde{F}(\mathbf{w})_i| \leq \frac{c}{2} \delta \quad \text{for all } \mathbf{x} \in \mathcal{X}$$

with probability at least $1 - \frac{c_p}{3} n^{-11}$.

Step 3: Bound the term $|\mathbb{E}[\nabla \tilde{F}(\mathbf{x})_i] - \mathbb{E}[\nabla \tilde{F}(\mathbf{w})_i]|$

We use the Cauchy–Schwarz inequality to bound each of the four terms in (B.2) in expectation. The first term in (B.2) can be bounded by

$$\begin{aligned} \mathbb{E} \left[\frac{1}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^2 (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c}) \right] &\leq \mathbb{E}[A_{1i}^2]^{\frac{1}{2}} \mathbb{E}[(\mathbf{A}_{1,S}^\top \mathbf{x}_S)^4 (\mathbf{A}_{1,S^c}^\top \mathbf{x}_{S^c})^2]^{\frac{1}{2}} \\ &= \left(\mathbb{E}[(\mathbf{A}_{1,S}^\top \mathbf{x}_S)^4] \mathbb{E}[(\mathbf{A}_{1,S^c}^\top \mathbf{x}_{S^c})^2] \right)^{\frac{1}{2}} \\ &\leq \sqrt{3 \|\mathbf{x}_{S^c}\|_2^2} \\ &\leq \sqrt{3} \delta \end{aligned}$$

for all $\mathbf{x} \in \mathcal{X}$, where we used that $\mathbf{A}_{1,S}^\top \mathbf{x}_S \sim \mathcal{N}(0, \|\mathbf{x}_S\|_2^2)$ and $\mathbf{A}_{1,S^c}^\top \mathbf{x}_{S^c} \sim \mathcal{N}(0, \|\mathbf{x}_{S^c}\|_2^2)$ are independent and that $\|\mathbf{x}_S\|_2 \leq 1$ and $\|\mathbf{x}_{S^c}\|_2 \leq \delta$ by the definition of \mathcal{X} . Similarly, we can bound

the other terms:

$$\begin{aligned}\mathbb{E}\left[\frac{1}{m}\sum_{j=1}^m A_{ji}(\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})(\mathbf{A}_j^\top \mathbf{x}^\star)^2\right] &\leq \sqrt{3}\delta, \\ \mathbb{E}\left[\frac{1}{m}\sum_{j=1}^m A_{ji}(\mathbf{A}_{j,S}^\top \mathbf{x}_S)(\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^2\right] &\leq \sqrt{3}\delta^2, \\ \mathbb{E}\left[\frac{1}{m}\sum_{j=1}^m A_{ji}(\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^3\right] &\leq \sqrt{15}\delta^3.\end{aligned}$$

This completes the proof that

$$|\mathbb{E}[\nabla \tilde{F}(\mathbf{x})_i] - \mathbb{E}[\nabla \tilde{F}(\mathbf{w})_i]_i| \leq \frac{c}{2}\delta.$$

Step 4, part (a): Bound the term $|\nabla \tilde{F}(\mathbf{w})_i - \mathbb{E}[\nabla \tilde{F}(\mathbf{w})_i]|$ by $\gamma \frac{\|\mathbf{x}_S\|_1}{k}$

Finally, we need to bound the term $|\nabla \tilde{F}(\mathbf{w})_i - \mathbb{E}[\nabla \tilde{F}(\mathbf{w})_i]|$ in (B.1) for all $\mathbf{x} \in \mathcal{X}$ and $i \in [n]$ with probability $1 - \frac{c_p}{3}n^{-10}$, which then completes the proof of Lemma 6.1. We begin with the bound $\gamma \frac{\|\mathbf{x}_S\|_1}{k}$.

We decompose $\nabla \tilde{F}(\mathbf{w})_i$ in a straightforward, albeit somewhat lengthy manner. We have

$$\begin{aligned}\nabla \tilde{F}(\mathbf{w})_i &= \frac{1}{m}\sum_{j=1}^m ((\mathbf{A}_j^\top \mathbf{w})^3 - (\mathbf{A}_j^\top \mathbf{w})(\mathbf{A}_j^\top \mathbf{x}^\star)^2)A_{ji} \\ &= (w_i^3 - w_i(x_i^\star)^2)\frac{1}{m}\sum_{j=1}^m A_{ji}^4 + (3w_i^2 - (x_i^\star)^2)\frac{1}{m}\sum_{j=1}^m A_{ji}^3(\mathbf{A}_{j,-i}^\top \mathbf{w}_{-i}) \\ &\quad - 2w_i x_i^\star \frac{1}{m}\sum_{j=1}^m A_{ji}^3(\mathbf{A}_{j,-i}^\top \mathbf{x}_{-i}^\star) + 3w_i \frac{1}{m}\sum_{j=1}^m A_{ji}^2(\mathbf{A}_{j,-i}^\top \mathbf{w}_{-i})^2 \\ &\quad - 2x_i^\star \frac{1}{m}\sum_{j=1}^m A_{ji}^2(\mathbf{A}_{j,-i}^\top \mathbf{w}_{-i})(\mathbf{A}_{j,-i}^\top \mathbf{x}_{-i}^\star) - w_i \frac{1}{m}\sum_{j=1}^m A_{ji}^2(\mathbf{A}_{j,-i}^\top \mathbf{x}_{-i}^\star)^2 \\ &\quad + \frac{1}{m}\sum_{j=1}^m A_{ji}(\mathbf{A}_{j,-i}^\top \mathbf{w}_{-i})^3 - \frac{1}{m}\sum_{j=1}^m A_{ji}(\mathbf{A}_{j,-i}^\top \mathbf{w}_{-i})(\mathbf{A}_{j,-i}^\top \mathbf{x}_{-i}^\star)^2 \\ &=: B_1 + B_2 + B_3 + B_4 + B_5 + B_6 + B_7 + B_8,\end{aligned}$$

and we will show that $|B_l - \mathbb{E}[B_l]|$ is small for all $l = 1, \dots, 8$. By Lemmas C.5 and C.9, all the following statements hold with probability $1 - \frac{c_p}{3}n^{-10}$ for all $\mathbf{w} \in \mathbb{R}^n$ with $\|\mathbf{w}\|_2 \leq 1$ and $\mathbf{w}_{S^c} = \mathbf{0}$, and for all $i = 1, \dots, n$. We write the bounds in terms of k instead of m using the assumption $m \geq c_s(\gamma)k^2 \log^2 n$.

- For the first term, we have $\mathbb{E}[B_1] = 3(w_i^3 - w_i(x_i^\star)^2)$, and

$$|B_1 - \mathbb{E}[B_1]| = \left| (w_i^3 - w_i(x_i^\star)^2) \left(\frac{1}{m}\sum_{j=1}^m A_{ji}^4 - 3 \right) \right| \leq \frac{\gamma \|\mathbf{w}\|_1}{8k}$$

by (C.3) of Lemma C.5, where we used $|w_i^3 - w_i(x_i^*)^2| \leq \|\mathbf{w}\|_1$.

- For the second term, we have $\mathbb{E}[B_2] = 0$, and

$$|B_2| \leq |3w_i^2 - (x_i^*)^2| \|\mathbf{w}\|_1 \max_{l \neq i} \left| \frac{1}{m} \sum_{j=1}^m A_{ji}^3 A_{jl} \right| \leq \frac{\gamma \|\mathbf{w}\|_1}{8k}$$

by (C.3) of Lemma C.5.

- For the third term, we have $\mathbb{E}[B_3] = 0$ and

$$|B_3| \leq |2w_i x_i^*| \frac{\gamma}{16k} \leq \frac{\gamma \|\mathbf{w}\|_1}{8k}$$

by (C.4) of Lemma C.5.

- For the fourth term, we have $\mathbb{E}[B_4] = 3w_i \|\mathbf{w}_{-i}\|_2^2$, and

$$\begin{aligned} |B_4 - \mathbb{E}[B_4]| &\leq 3|w_i| \left| \sum_{l \neq i} w_l^2 \left(\frac{1}{m} \sum_{j=1}^m A_{ji}^2 A_{jl}^2 - 1 \right) + \sum_{l \neq i} w_l \sum_{r \neq l, i} w_r \frac{1}{m} \sum_{j=1}^m A_{ji}^2 A_{jl} A_{jr} \right| \\ &\leq \frac{\gamma \|\mathbf{w}\|_1}{16k} + \frac{\gamma \|\mathbf{w}\|_1}{16k} \\ &= \frac{\gamma \|\mathbf{w}\|_1}{8k}, \end{aligned}$$

where we used Lemma (C.3) of C.5 to bound the first and (C.14) of Lemma C.9 to bound the second term.

- For the fifth term, we have $\mathbb{E}[B_5] = -2x_i^* \sum_{l \neq i} w_l x_l^*$ and

$$|B_5 - \mathbb{E}[B_5]| \leq |2x_i^*| \left| \sum_{l \neq i} w_l \left(\frac{1}{m} \sum_{j=1}^m A_{ji}^2 A_{jl} (\mathbf{A}_{j,-i}^\top \mathbf{x}_{-i}^*) - x_l^* \right) \right| \leq \frac{\gamma \|\mathbf{w}\|_1}{8k}$$

where we used Hölder's inequality and (C.4) of Lemma C.5.

- For the sixth term, we have $\mathbb{E}[B_6] = -w_i \|\mathbf{x}_{-i}^*\|_2^2$, and

$$|B_6 - \mathbb{E}[B_6]| \leq |w_i| \frac{\gamma}{8k} \leq \frac{\gamma \|\mathbf{w}\|_1}{8k}$$

by (C.5) of Lemma C.5.

- For the seventh term, we have $\mathbb{E}[B_7] = 0$ and

$$|B_7| \leq \frac{\gamma \|\mathbf{w}\|_1}{8k}$$

by (C.13) of Lemma C.9.

- Finally, for the eighth term, we have $\mathbb{E}[B_8] = 0$ and

$$|B_8| \leq \left| \sum_{l \neq i} w_l \frac{1}{m} \sum_{j=1}^m A_{ji} A_{jl} (\mathbf{A}_{j,-i}^\top \mathbf{x}_{-i}^*)^2 \right| \leq \frac{\gamma \|\mathbf{w}\|_1}{8k}$$

where we used Hölder's inequality and (C.5) of Lemma C.5.

Step 4, part (b): Bound the term $|\nabla\tilde{F}(\mathbf{w})_i - \mathbb{E}[\nabla\tilde{F}(\mathbf{w})_i]|$ by $\gamma \frac{\|\mathbf{x}_S - \mathbf{x}_S^*\|_2}{\sqrt{k}}$

To show this bound, we parametrize $\mathbf{z} = \mathbf{w} - \mathbf{x}^*$ and write

$$\begin{aligned}\nabla\tilde{F}(\mathbf{w})_i &= \frac{1}{m} \sum_{j=1}^m ((\mathbf{A}_j^\top (\mathbf{z} + \mathbf{x}^*))^3 - (\mathbf{A}_j^\top (\mathbf{z} + \mathbf{x}^*))(\mathbf{A}_j^\top \mathbf{x}^*)^2) A_{ji} \\ &= \frac{1}{m} \sum_{j=1}^m ((\mathbf{A}_j^\top \mathbf{z})^3 + 3(\mathbf{A}_j^\top \mathbf{z})^2(\mathbf{A}_j^\top \mathbf{x}^*) + 2(\mathbf{A}_j^\top \mathbf{z})(\mathbf{A}_j^\top \mathbf{x}^*)^2) A_{ji}.\end{aligned}$$

We can decompose this expression as before and observe that each term depends at least linearly on $\|\mathbf{z}\|_2$. Further, because $\mathbf{z}_{S^c} = \mathbf{w}_{S^c} = \mathbf{0}$, we have $\|\mathbf{z}\|_1 \leq \sqrt{k}\|\mathbf{z}\|_2$. Hence, we obtain the bound $\gamma \frac{\|\mathbf{z}\|_1}{\sqrt{k}} \leq \gamma \frac{\|\mathbf{z}\|_2}{\sqrt{k}}$ the same way as above using the bounds in Lemmas C.5 and C.9. We omit the details to avoid repetition.

All in all, putting everything together we have, with probability $1 - \frac{c_p}{3}n^{-10}$,

$$|\nabla\tilde{F}(\mathbf{w})_i - \mathbb{E}[\nabla\tilde{F}(\mathbf{w})_i]| \leq \gamma \min\left\{\frac{\|\mathbf{w}\|_1}{k}, \frac{\|\mathbf{z}\|_2}{\sqrt{k}}\right\} = \gamma \min\left\{\frac{\|\mathbf{x}_S\|_1}{k}, \frac{\|\mathbf{x}_S - \mathbf{x}_S^*\|_2}{\sqrt{k}}\right\},$$

for all $\mathbf{w} \in \mathbb{R}^n$ with $\|\mathbf{w}\|_2 \leq 1$ and $\mathbf{w}_{S^c} = \mathbf{0}$, and all $i = 1, \dots, n$, which completes the proof of Lemma 6.1. Finally, the simplified bound can be obtained directly by plugging in the additional assumptions. \square

Proof of Lemma 6.2 As in the proof of Lemma 6.1, we assume $\|\mathbf{x}^*\|_2 = 1$ for notational simplicity, and denote the gradient corresponding to noiseless measurements by ∇F . Then,

$$|\langle \nabla F(\mathbf{x}) - \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle| \leq |\langle \nabla F(\mathbf{x}) - \nabla\tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle| + |\langle \nabla\tilde{F}(\mathbf{x}) - \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle|,$$

and we begin by bounding the first of the two terms.

Step 1: Bound the term $|\langle \nabla F(\mathbf{x}) - \nabla\tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle|$

We begin by bounding the first term $|\langle \nabla F(\mathbf{x}) - \nabla\tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle|$ by

$$c\sigma \sqrt{\frac{\log n}{m}} \left(\delta + \min\{\|\mathbf{x}_S\|_1, \sqrt{k}\|\mathbf{x}_S - \mathbf{x}_S^*\|_2\} \right).$$

To obtain the second bound in the minimum, we can substitute $\mathbf{z} = \mathbf{x} - \mathbf{x}^*$ and write

$$\begin{aligned}|\langle \nabla F(\mathbf{x}) - \nabla\tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle| &= \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_j^\top \mathbf{x}) (\mathbf{A}_j^\top (\mathbf{x} - \mathbf{x}^*)) \right| \\ &\leq \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_j^\top \mathbf{z})^2 \right| + \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_j^\top \mathbf{z}) (\mathbf{A}_j^\top \mathbf{x}^*) \right|.\end{aligned}$$

Using Lemmas C.5 and C.9 we can bound the first term, with probability $1 - \frac{c_p}{3}n^{10}$, by

$$\begin{aligned}
\left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_j^\top \mathbf{z})^2 \right| &\leq \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_{j,S}^\top \mathbf{z}_S)^2 \right| + 2 \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_{j,S}^\top \mathbf{z}_S) (\mathbf{A}_{j,S^c}^\top \mathbf{z}_{S^c}) \right| + \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_{j,S^c}^\top \mathbf{z}_{S^c})^2 \right| \\
&\leq \frac{c}{4} \sigma \|\mathbf{z}_S\|_1 \sqrt{\frac{\log n}{m}} + 2 \|\mathbf{z}_{S^c}\|_1 \max_{l \notin S} \left| \sum_{s \in S} z_s \frac{1}{m} \sum_{j=1}^m \varepsilon_j A_{js} A_{jl} \right| \\
&\quad + \|\mathbf{z}_{S^c}\|_1^2 \max_{l, s \notin S} \left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j A_{jl} A_{js} \right| \\
&\leq \frac{c}{4} \sigma \|\mathbf{z}_S\|_1 \sqrt{\frac{\log n}{m}} + 2 \frac{c}{4} \sigma \|\mathbf{z}_S\|_1 \sqrt{\frac{\log n}{m}} \delta + \frac{c}{4} \sigma \sqrt{\frac{\log n}{m}} \delta^2,
\end{aligned}$$

provided that $\frac{c}{4}$ is at least as large as the universal constants from Lemmas C.5 and C.9, where we used

$$\frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_{j,S}^\top \mathbf{z}_S)^2 = \sum_{l \in S} z_l^2 \frac{1}{m} \sum_{j=1}^m \varepsilon_j A_{jl}^2 + \sum_{\substack{l, s \in S \\ l \neq s}} z_l z_s \frac{1}{m} \sum_{j=1}^m \varepsilon_j A_{jl} A_{js}$$

and Lemmas C.5 and C.9 for the second inequality and Lemma C.5 to bound the other terms. Similarly, we can bound

$$\left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j (\mathbf{A}_j^\top \mathbf{z}) (\mathbf{A}_j^\top \mathbf{x}^*) \right| \leq \left| \sum_{l=1}^n z_l \frac{1}{m} \sum_{j=1}^m \varepsilon_j A_{jl} (\mathbf{A}_j^\top \mathbf{x}^*) \right| \leq \frac{c}{2} \sigma (\|\mathbf{z}_S\|_1 + \delta) \sqrt{\frac{\log n}{m}},$$

where we used Lemma C.5. All in all, noting that $\|\mathbf{z}_S\|_1 \leq \sqrt{k} \|\mathbf{x}_S - \mathbf{x}_S^*\|_2$, this yields, with probability $1 - \frac{c_p}{3}n^{-10}$,

$$|\langle \nabla F(\mathbf{x}) - \nabla \tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle| \leq c\sigma \sqrt{\frac{\log n}{m}} \left(\delta + \sqrt{k} \|\mathbf{x}_S - \mathbf{x}_S^*\|_2 \right).$$

The first bound in the minimum can be obtained following exactly the same steps as above without substituting $\mathbf{z} = \mathbf{x} - \mathbf{x}^*$.

Step 2: Bound the term $|\langle \nabla \tilde{F}(\mathbf{x}) - \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle|$ by $\gamma(\frac{\|\mathbf{x}_S\|_1}{\sqrt{k}} + \delta)$

Writing $\mathbf{w} \in \mathbb{R}^n$ for the vector \mathbf{x}_S padded with zeroes, i.e. $w_i = x_i$ for $i \in S$ and $w_i = 0$; otherwise, we can bound

$$\begin{aligned}
|\langle \nabla \tilde{F}(\mathbf{x}) - \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle| &\leq |\langle \nabla \tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle - \langle \nabla \tilde{F}(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle| \\
&\quad + |\langle \nabla \tilde{F}(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle - \mathbb{E}[\langle \nabla \tilde{F}(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle]| \\
&\quad + |\mathbb{E}[\langle \nabla \tilde{F}(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle] - \mathbb{E}[\langle \nabla \tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle]|, \tag{B.3}
\end{aligned}$$

where we used that $\nabla f(\mathbf{x}) = \mathbb{E}[\nabla F(\mathbf{x})] = \mathbb{E}[\nabla \tilde{F}(\mathbf{x})]$. To bound these three terms, we write

$$\begin{aligned} \langle \nabla \tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle &= \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_j^\top \mathbf{x})^4 - \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_j^\top \mathbf{x})^3 (\mathbf{A}_j^\top \mathbf{x}^*) \\ &\quad - \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_j^\top \mathbf{x})^2 (\mathbf{A}_j^\top \mathbf{x}^*)^2 + \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_j^\top \mathbf{x}) (\mathbf{A}_j^\top \mathbf{x}^*)^3 \\ &= g_{1,\mathbf{x}}(\mathbf{A}) - g_{2,\mathbf{x}}(\mathbf{A}) - g_{3,\mathbf{x}}(\mathbf{A}) + g_{4,\mathbf{x}}(\mathbf{A}), \end{aligned}$$

and note that we have analogous expressions for the terms $\langle \tilde{F}(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle$, $\mathbb{E}[\langle \tilde{F}(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle]$ and $\mathbb{E}[\langle \tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle]$. We will show that, with probability $1 - \frac{c_p}{3} n^{-10}$, the four terms $g_{i,\mathbf{x}}(\mathbf{A})$ deviate by at most $\frac{\gamma}{4} (\|\mathbf{x}\|_1 / \sqrt{k} + \delta)$ from their means by bounding each difference as in (B.3).

Step 2, part (a): Split $g_{i,\mathbf{x}}(\mathbf{A})$ into a term depending on \mathbf{x}_S and a residual term

We first show that the first and last terms in (B.3) are both bounded by $\frac{\gamma}{2} \delta$. To this end, we split each of the four terms $g_{i,\mathbf{x}}(\mathbf{A})$ into a term that only depends on \mathbf{x}_S (which corresponds to $\langle \nabla \tilde{F}(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle$) and a residual term that also depends on \mathbf{x}_{S^c} (which corresponds to the difference $\langle \nabla \tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle - \langle \nabla \tilde{F}(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle$). We only present the computation for $g_{1,\mathbf{x}}(\mathbf{A})$, since the other three terms can be bounded following the same steps. We have

$$\begin{aligned} g_{1,\mathbf{x}}(\mathbf{A}) &= \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S + \mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^4 \\ &= \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^4 + \frac{4}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^3 (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c}) + \frac{6}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^2 (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^2 \\ &\quad + \frac{4}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S) (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^3 + \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^4. \end{aligned}$$

The first term only depends on $\mathbf{x}_S = \mathbf{w}_S$, and we denote it by $h_{1,\mathbf{w}}(\mathbf{A}) = \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{w}_S)^4$. The other terms can be bounded as follows:

$$\begin{aligned} &\frac{4}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^3 (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c}) \\ &= 4 \sum_{l \notin S} x_l \frac{1}{m} \sum_{j=1}^m A_{jl} (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^3 \leq 4 \|\mathbf{x}_{S^c}\|_1 \max_{l \notin S} \left| \frac{1}{m} \sum_{j=1}^m A_{jl} (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^3 \right| \leq \frac{\gamma}{11} \delta, \end{aligned}$$

provided that $m \geq (\frac{44\tilde{c}}{\gamma})^2 k \log^2 n$, where we used that $\|\mathbf{x}_S\|_1 \leq \sqrt{k}$ and \tilde{c} is the universal constant from Lemma C.9. For the first inequality we used Hölder's inequality, and the second inequality holds by

(C.13) of Lemma C.9 with probability $1 - \frac{c_p}{9}n^{-10}$. Similarly, we can bound

$$\begin{aligned} \frac{6}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^2 (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^2 &= 6 \sum_{l \notin S} x_l \sum_{r \notin S} x_r \frac{1}{m} \sum_{j=1}^m A_{jl} A_{jr} (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^2 \\ &\leq 6 \|\mathbf{x}_{S^c}\|_1^2 \max_{l,r \notin S} \sqrt{\frac{1}{m} \sum_{j=1}^m A_{jl}^2 A_{jr}^2} \sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^4} \\ &\leq \frac{c}{11} \delta^2, \end{aligned}$$

provided that $c > 0$ is large enough, where we used Hölder's inequality in the second line and the last inequality holds by Lemmas C.3 and C.5 with probability $1 - \frac{2c_p}{9}n^{-10}$. The same computation yields

$$\begin{aligned} \frac{4}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S) (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^3 &\leq \frac{c}{11} \delta^3, \\ \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^4 &\leq \frac{c}{11} \delta^4. \end{aligned}$$

Putting everything together, we can bound the residual term by $\frac{c}{8}\delta$ since $\delta \leq \frac{1}{4} \min\{\frac{\gamma}{c}, 1\}$. Bounding $g_{2,\mathbf{x}}(\mathbf{A})$, $g_{3,\mathbf{x}}(\mathbf{A})$ and $g_{4,\mathbf{x}}(\mathbf{A})$ the same way, we have, with probability $1 - \frac{c_p}{3}n^{-10}$,

$$|\langle \nabla \tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle - \langle \nabla \tilde{F}(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle| \leq \frac{\gamma}{2} \delta.$$

For the difference in expectation, we can bound, using $\|\mathbf{x}\|_2 \leq 1$ and $\|\mathbf{x}_{S^c}\|_2 \leq \|\mathbf{x}_{S^c}\|_1 \leq \delta$,

$$\begin{aligned} \mathbb{E} \left[\frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^3 (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c}) \right] &= \mathbb{E}[(\mathbf{A}_{1,S}^\top \mathbf{x}_S)^3] \mathbb{E}[(\mathbf{A}_{1,S^c}^\top \mathbf{x}_{S^c})] = 0, \\ \mathbb{E} \left[\frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S)^2 (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^2 \right] &= \mathbb{E}[(\mathbf{A}_{1,S}^\top \mathbf{x}_S)^2] \mathbb{E}[(\mathbf{A}_{1,S^c}^\top \mathbf{x}_{S^c})^2] \leq \delta^2, \\ \mathbb{E} \left[\frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{x}_S) (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^3 \right] &= \mathbb{E}[(\mathbf{A}_{1,S}^\top \mathbf{x}_S)] \mathbb{E}[(\mathbf{A}_{1,S^c}^\top \mathbf{x}_{S^c})^3] = 0, \\ \mathbb{E} \left[\frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S^c}^\top \mathbf{x}_{S^c})^4 \right] &= \mathbb{E}[(\mathbf{A}_{1,S^c}^\top \mathbf{x}_{S^c})^4] \leq 3\delta^4. \end{aligned}$$

Repeating this for $g_{2,\mathbf{x}}(\mathbf{A})$, $g_{3,\mathbf{x}}(\mathbf{A})$ and $g_{4,\mathbf{x}}(\mathbf{A})$ shows

$$|\mathbb{E}[\langle \nabla F(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle] - \mathbb{E}[\langle \nabla F(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle]| \leq \frac{\gamma}{2} \delta.$$

Step 2, part (b): Bound the term $|\langle \nabla F(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle - \mathbb{E}[\langle \nabla F(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle]|$

To bound the second term in (B.3), we need to show that the terms $h_{i,\mathbf{w}}(\mathbf{A})$ concentrate around their respective expectations. Let N_ϵ be a smallest ϵ -net of the set $\mathcal{X} = \{\mathbf{w} \in \mathbb{R}^n : \mathbf{w}_{S^c} = \mathbf{0}, \|\mathbf{w}\|_2 = 1\}$ (which corresponds to the unit sphere in \mathbb{R}^k), where $\epsilon = c_1 \gamma n^{-3}$ for some constant $c_1 > 0$. We will first show that we can bound the difference $|h_{i,\mathbf{w}}(\mathbf{A}) - \mathbb{E}[h_{i,\mathbf{w}}(\mathbf{A})]| \leq \frac{\gamma}{8} \frac{\|\mathbf{w}\|_1}{\sqrt{k}}$ for every $\mathbf{w} \in N_\epsilon$ via concentration of Lipschitz functions for Gaussian random variables. Then, we will extend this bound to every $\mathbf{w} \in \mathcal{X}$.

As the functions $h_{i,\mathbf{w}}$ are not globally Lipschitz continuous, we cannot directly apply Theorem C.1. We will first bound the Lipschitz constant of $h_{i,\mathbf{w}}$ restricted to the set \mathcal{A} defined as the intersection of the sets defined in Lemma C.4 and Lemma C.6 and then extend this restricted function to a function $\tilde{h}_{i,\mathbf{w}}$ on the entire space such that $\tilde{h}_{i,\mathbf{w}}$ is globally Lipschitz continuous. We can then apply Theorem C.1 to $\tilde{h}_{i,\mathbf{w}}$, which also provides a high probability bound for $h_{i,\mathbf{w}}$, since by construction $h_{i,\mathbf{w}}(\mathbf{A}) = \tilde{h}_{i,\mathbf{w}}(\mathbf{A})$ with high probability.

Step 2, part (b.1): Bound the Lipschitz constant of $h_{i,\mathbf{w}}$ restricted to

\mathcal{A} Let the subset $\mathcal{A} \subset \mathbb{R}^{m \times n}$ be defined as the intersection of the sets defined in Lemma C.4 with $t = 5\sqrt{\log n}$ and Lemma C.6. More precisely, we require the projection of \mathcal{A} onto $\mathbb{R}^{m \times \mathcal{S}}$ to be as in Lemma C.4. By the aforementioned lemmas, we have $\mathbb{P}[\mathcal{A}^c] \leq c_2 n^{-12}$ for a universal constant $c_2 > 0$. Since \mathcal{A} is convex, it follows from the mean-value theorem that the Lipschitz constant of $h_{i,\mathbf{w}}$ restricted to \mathcal{A} is bounded by the norm of its gradient $\|\nabla h_{i,\mathbf{w}}\|_2$. For any $\{\mathbf{a}_j\}_{j=1}^m \in \mathcal{A}$, we can compute the following.

- $\frac{\partial}{\partial a_{jl}} h_{1,\mathbf{w}}(\mathbf{a}) = \frac{4}{m} w_l (\mathbf{a}_j^\top \mathbf{w})^3$. We can bound the Lipschitz constant by the squareroot of

$$\begin{aligned} \|\nabla h_{1,\mathbf{w}}(\mathbf{a})\|_2^2 &= \frac{16}{m^2} \sum_{l=1}^n w_l^2 \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{w})^6 \\ &\leq \frac{16}{m^2} ((15m)^{\frac{1}{6}} + \sqrt{8 \log(2k)} \|\mathbf{w}\|_1 + 5\sqrt{\log n})^6 \\ &\leq c_3 \frac{\|\mathbf{w}\|_1^2 \log k}{m}, \end{aligned}$$

where we used $\|\mathbf{w}\|_2 = 1$, $\|\mathbf{w}\|_1 \leq \sqrt{k}$, $m \geq c_s(\gamma) \max\{k^2 \log^2 n, \log^5 n\}$ and Lemma C.4.

- $\frac{\partial}{\partial a_{jl}} h_{2,\mathbf{w}}(\mathbf{a}) = \frac{1}{m} (x_l^* (\mathbf{a}_j^\top \mathbf{w})^3 + 3w_l (\mathbf{a}_j^\top \mathbf{w})^2 (\mathbf{a}_j^\top \mathbf{w}^*))$. Hence, we can bound the Lipschitz constant by the squareroot of

$$\begin{aligned} \|\nabla h_{2,\mathbf{w}}(\mathbf{a})\|_2^2 &= \frac{2}{m^2} \sum_{l=1}^n \sum_{j=1}^m (x_l^*)^2 (\mathbf{a}_j^\top \mathbf{w})^6 + 9w_l^2 (\mathbf{a}_j^\top \mathbf{w})^4 (\mathbf{a}_j^\top \mathbf{x}^*)^2 \\ &\leq \frac{2}{m^2} ((15m)^{\frac{1}{6}} + 2\sqrt{2 \log(2k)} \|\mathbf{w}\|_1 + 5\sqrt{\log n})^6 \\ &\quad + \frac{18}{m} \sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{w})^8} \sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x}^*)^4} \\ &\leq c_4 \frac{\|\mathbf{w}\|_1^2 \log k}{m}, \end{aligned}$$

where we use Lemma C.4 to bound $\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{w})^8$ and Lemma C.6 to bound $\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x}^*)^4$.

- $\frac{\partial}{\partial a_{ji}} h_{3,\mathbf{w}}(\mathbf{a}) = \frac{2}{m} (x_l^* (\mathbf{a}_j^\top \mathbf{w})^2 (\mathbf{a}_j^\top \mathbf{x}^*) + w_l (\mathbf{a}_j^\top \mathbf{w}) (\mathbf{a}_j^\top \mathbf{x}^*)^2)$. Hence, we can bound the Lipschitz constant by the squareroot of

$$\begin{aligned} \|\nabla h_{3,\mathbf{w}}(\mathbf{a})\|_2^2 &= \frac{4}{m^2} \sum_{l=1}^n \sum_{j=1}^m 2(x_l^*)^2 (\mathbf{a}_j^\top \mathbf{w})^4 (\mathbf{a}_j^\top \mathbf{x}^*)^2 + 2w_l^2 (\mathbf{a}_j^\top \mathbf{w})^2 (\mathbf{a}_j^\top \mathbf{x}^*)^4 \\ &\leq \frac{8}{m} \sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{w})^8} \sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x}^*)^4} + \frac{8}{m} \sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{w})^4} \sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x}^*)^8} \\ &\leq c_5 \frac{\|\mathbf{w}\|_1^2 \log k}{m}, \end{aligned}$$

again by Lemmas C.4 and C.6.

- $\frac{\partial}{\partial a_{ji}} h_{4,\mathbf{w}}(\mathbf{a}) = \frac{1}{m} (3x_l^* (\mathbf{a}_j^\top \mathbf{w}) (\mathbf{a}_j^\top \mathbf{x}^*)^2 + w_l (\mathbf{a}_j^\top \mathbf{x}^*)^3)$. Hence, we can bound the Lipschitz constant by the squareroot of

$$\begin{aligned} \|\nabla h_{4,\mathbf{w}}(\mathbf{a})\|_2^2 &= \frac{2}{m^2} \sum_{l=1}^n \sum_{j=1}^m 9(x_l^*)^2 (\mathbf{a}_j^\top \mathbf{w})^2 (\mathbf{a}_j^\top \mathbf{x}^*)^4 + w_l^2 (\mathbf{a}_j^\top \mathbf{x}^*)^6 \\ &\leq \frac{18}{m} \sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{w})^4} \sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x}^*)^8} + \frac{2}{m} \frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x}^*)^6 \\ &\leq c_6 \frac{\|\mathbf{w}\|_1^2 \log k}{m}, \end{aligned}$$

again by Lemmas C.4 and C.6.

Step 2, part (b.2): Construct a globally Lipschitz continuous extension of $h_{i,\mathbf{w}}$

We only present the following steps for the first term $h_{1,\mathbf{w}}$, as the proofs for the other three terms follow the exact same steps. Consider the following Lipschitz extension of $h_{1,\mathbf{w}}$:

$$\tilde{h}_{1,\mathbf{w}}(\mathbf{a}) = \inf_{\mathbf{a}' \in \mathcal{A}} (h_{1,\mathbf{w}}(\mathbf{a}') + \text{Lip}(h_{1,\mathbf{w}}) \|\mathbf{a} - \mathbf{a}'\|_2),$$

where we write $\text{Lip}(h_{1,\mathbf{w}}) = \sqrt{c_3 \frac{\|\mathbf{w}\|_1^2 \log k}{m}}$. By definition, we have $\tilde{h} = h$ on \mathcal{A} , and it follows from an application of the triangle inequality that \tilde{h} is globally Lipschitz continuous with Lipschitz constant $\text{Lip}(h)$ (see e.g. Theorem 7.2 of ([6])).

We will show that $\tilde{h}_{1,\mathbf{w}}$ concentrates around its mean, which can potentially differ from the mean of $h_{1,\mathbf{w}}$. Since $h_{1,\mathbf{w}}$ and $\tilde{h}_{1,\mathbf{w}}$ differ only on \mathcal{A}^c (which has probability less than $c_2 n^{-12}$), we can bound, using the Cauchy–Schwarz inequality,

$$\mathbb{E}[|h_{1,\mathbf{w}}(\mathbf{A})| \mathbf{1}_{\mathcal{A}^c}(\mathbf{A})] \leq \sqrt{\mathbb{E}[h_{1,\mathbf{w}}(\mathbf{A})^2]} \sqrt{\mathbb{E}[\mathbf{1}_{\mathcal{A}^c}(\mathbf{A})]} \leq \frac{c'}{n^6},$$

where we used $\mathbb{E}[h_{1,\mathbf{w}}(\mathbf{A})^2] \leq 3 + 105/m$. Similarly, since $\tilde{h}_{1,\mathbf{w}}(\mathbf{a}) \leq \text{Lip}(h_{1,\mathbf{w}})\|\mathbf{a}\|_2$,

$$\mathbb{E}[|\tilde{h}_{1,\mathbf{w}}(\mathbf{A})|\mathbf{1}_{\mathcal{A}^c}(\mathbf{A})] \leq \text{Lip}(h_{1,\mathbf{w}})\sqrt{\mathbb{E}[\|\mathbf{A}\|_2^2]}\sqrt{\mathbb{E}[\mathbf{1}_{\mathcal{A}^c}(\mathbf{A})]} \leq \frac{c''}{n^5}.$$

All in all, this shows that

$$|\mathbb{E}[h_{1,\mathbf{w}}(\mathbf{A})] - \mathbb{E}[\tilde{h}_{1,\mathbf{w}}(\mathbf{A})]| \leq \frac{c_7}{n^5}$$

for a constant $c_7 > 0$. Finally, using the triangle inequality and Theorem C.1, we have

$$\mathbb{P}\left[|\tilde{h}_{1,\mathbf{w}}(\mathbf{A}) - \mathbb{E}[h_{1,\mathbf{w}}(\mathbf{A})]| > \frac{\gamma}{8} \frac{\|\mathbf{w}\|_1}{\sqrt{k}}\right] \leq 2 \exp\left(-\frac{(\frac{\gamma}{8} \frac{\|\mathbf{w}\|_1}{\sqrt{k}} - \frac{c_7}{n^5})^2}{2c_3 \frac{\|\mathbf{w}\|_1^2 \log k}{m}}\right) \leq 2 \exp(-c_8 k \log n)$$

for a constant $c_8 \leq \frac{\gamma^2 c_5(\gamma)}{128c_3} - \frac{\gamma c_5(\gamma)c_7}{8c_3 n^5}$.

Step 2, part (b.3): Union bound over $\mathbf{w} \in N_\epsilon$

Taking the union bound over all $\mathbf{w} \in N_\epsilon$, which has cardinality bounded by $(3/\epsilon)^k$, we have

$$\mathbb{P}\left[|\tilde{h}_{1,\mathbf{w}}(\mathbf{A}) - \mathbb{E}[h_{1,\mathbf{w}}(\mathbf{A})]| > \frac{\gamma}{8} \frac{\|\mathbf{w}\|_1}{\sqrt{k}} \text{ for some } \mathbf{w} \in N_\epsilon\right] \leq 2 \exp\left(-c_8 k \log n + k \log \frac{3}{\epsilon}\right).$$

Since $h_{1,\mathbf{w}}(\mathbf{a}) = \tilde{h}_{1,\mathbf{w}}(\mathbf{a})$ for all \mathbf{w} provided that $\mathbf{a} \in \mathcal{A}$, this implies

$$\mathbb{P}\left[|h_{1,\mathbf{w}}(\mathbf{A}) - \mathbb{E}[h_{1,\mathbf{w}}(\mathbf{A})]| > \frac{\gamma}{8} \frac{\|\mathbf{w}\|_1}{\sqrt{k}} \text{ for some } \mathbf{w} \in N_\epsilon\right] \leq 2 \exp(-c_9 k \log n) + c_2 n^{-12},$$

for a constant $c_9 \leq c_8 - 3 - \frac{1}{\log n} \log \frac{3}{c_1 \gamma}$, as we have $\epsilon = c_1 \gamma n^{-3}$.

Step 2, part (b.4): From ϵ -net to the full sphere

Next, we show that $|h_{1,\mathbf{w}}(\mathbf{A}) - \mathbb{E}[h_{1,\mathbf{w}}(\mathbf{A})]| \leq \frac{\gamma}{4} \frac{\|\mathbf{w}\|_1}{\sqrt{k}}$ for any $\mathbf{w} \in \mathcal{X}$. The case $\|\mathbf{w}\|_2 < 1$ follows by considering the vector $\mathbf{w}/\|\mathbf{w}\|_2$ and rescaling. For any $\mathbf{w} \in \mathcal{X}$, let $\mathbf{w}' \in N_\epsilon$ with $\|\mathbf{w} - \mathbf{w}'\|_2 \leq \epsilon$. Then,

$$\begin{aligned} |h_{1,\mathbf{w}}(\mathbf{A}) - \mathbb{E}[h_{1,\mathbf{w}}(\mathbf{A})]| &\leq |h_{1,\mathbf{w}}(\mathbf{A}) - h_{1,\mathbf{w}'}(\mathbf{A})| + |h_{1,\mathbf{w}'}(\mathbf{A}) - \mathbb{E}[h_{1,\mathbf{w}'}(\mathbf{A})]| \\ &\quad + |\mathbb{E}[h_{1,\mathbf{w}'}(\mathbf{A})] - \mathbb{E}[h_{1,\mathbf{w}}(\mathbf{A})]|. \end{aligned}$$

The first and third term can be bounded using the identity $a^4 - b^4 = (a^2 + b^2)(a + b)(a - b)$:

$$\begin{aligned} |h_{1,\mathbf{w}}(\mathbf{A}) - h_{1,\mathbf{w}'}(\mathbf{A})| &= \left| \frac{1}{m} \sum_{j=1}^m ((\mathbf{A}_j^\top \mathbf{w})^2 + (\mathbf{A}_j^\top \mathbf{w}')^2)(\mathbf{A}_j^\top (\mathbf{w} + \mathbf{w}'))(\mathbf{A}_j^\top (\mathbf{w} - \mathbf{w}')) \right| \\ &\leq \max_j 2\|\mathbf{A}_j\|_2^2 \cdot 2\|\mathbf{A}_j\|_2 \|\mathbf{A}_j\|_2 \|\mathbf{w} - \mathbf{w}'\|_2 \\ &\leq 4(\sqrt{k} + 5\sqrt{\log n})^4 \|\mathbf{w} - \mathbf{w}'\|_2 \\ &\leq \frac{\gamma}{16} \frac{\|\mathbf{w}\|_1}{\sqrt{k}}, \end{aligned}$$

with probability $1 - mn^{-12.5}$, provided c_1 is sufficiently large, where we used the fact that the norm $\|\cdot\|_2$ is 1-Lipschitz continuous and applied Theorem C.1 to bound the term $\|\mathbf{A}_j\|_2$ and used $\|\mathbf{w} - \mathbf{w}'\|_2 \leq \epsilon$

for the last inequality. For the expectation, the same argument yields

$$\begin{aligned} |\mathbb{E}[h_{1,\mathbf{w}}(\mathbf{A})] - \mathbb{E}[h_{1,\mathbf{w}'}(\mathbf{A})]| &\leq \mathbb{E}[2\|\mathbf{A}_j\|_2^2 \cdot 2\|\mathbf{A}_j\|_2\|\mathbf{A}_j\|_2\|\mathbf{w} - \mathbf{w}'\|_2] \\ &= 4(3k + k(k-1))\|\mathbf{w} - \mathbf{w}'\|_2 \\ &\leq \frac{\gamma}{16} \frac{\|\mathbf{w}\|_1}{\sqrt{k}}. \end{aligned}$$

This completes the proof that

$$\mathbb{P}\left[|h_{1,\mathbf{w}}(\mathbf{A}) - \mathbb{E}[h_{1,\mathbf{w}}(\mathbf{A})]| < \frac{\gamma}{4} \frac{\|\mathbf{w}\|_1}{\sqrt{k}} \text{ for all } \mathbf{w} \in \mathcal{X}\right] \geq 1 - \frac{c_p}{24} n^{-10},$$

provided that $c_p \geq 24 + 24c_2n^{-2}$, and if $m \leq n^{2.5}$. The other case $m > n^{2.5}$ is simpler and can be shown following the same steps, writing the probabilities in terms of m instead of n . We omit the details to avoid repetition. Repeating the same steps for the terms $h_{2,\mathbf{w}}$, $h_{3,\mathbf{w}}$ and $h_{4,\mathbf{w}}$ shows that

$$\mathbb{P}\left[|\langle \nabla \tilde{F}(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle - \mathbb{E}[\langle \nabla \tilde{F}(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle]| \leq \gamma \frac{\|\mathbf{w}\|_1}{\sqrt{k}} \text{ for all } \mathbf{w} \in \mathcal{X}\right] \geq 1 - \frac{c_p}{6} n^{-10}.$$

Finally, the bound also holds for any vector $\mathbf{w} \in \mathbb{R}^n$ with $\mathbf{w}_{S^c} = \mathbf{0}$ and $\|\mathbf{w}\|_2 < 1$ by considering $\mathbf{w}/\|\mathbf{w}\|_2$ and noting that each of the four terms that make up $\langle \nabla \tilde{F}(\mathbf{w}), \mathbf{w} - \mathbf{x}^* \rangle$ scale at least linearly in $\|\mathbf{w}\|_2$.

Step 3: Bound the term $|\langle \nabla \tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle|$ by $\gamma(\|\mathbf{x}_S - \mathbf{x}_S^*\|_2^2 + \delta^2)$

Substituting $\mathbf{z} = \mathbf{x} - \mathbf{x}^*$, we have

$$\langle \nabla \tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle = \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_j^\top \mathbf{z})^4 + \frac{3}{m} \sum_{j=1}^m (\mathbf{A}_j^\top \mathbf{z})^3 (\mathbf{A}_j^\top \mathbf{x}^*) + \frac{2}{m} \sum_{j=1}^m (\mathbf{A}_j^\top \mathbf{z})^2 (\mathbf{A}_j^\top \mathbf{x}^*)^2. \quad (\text{B.4})$$

The proof follows the same steps as the above proof of the bound $\gamma(\frac{\|\mathbf{x}_S\|_1}{\sqrt{k}} + \delta)$, writing $\mathbf{z}_S = \|\mathbf{z}_S\|_2 \frac{\mathbf{z}_S}{\|\mathbf{z}_S\|_2}$ and using the fact that $\|\mathbf{z}_{S^c}\|_1 = \|\mathbf{x}_{S^c}\|_1 \leq \delta$ and $\frac{\|\mathbf{z}_S\|_1}{\sqrt{k}} \leq \|\mathbf{z}_S\|_2$. We then obtain the desired bound since each of the three terms depends at least quadratically on \mathbf{z} .

We illustrate this argument for the last term in (B.4). The other two terms can be controlled similarly (they are easier to control because of the higher order dependence on \mathbf{z}). We have

$$\begin{aligned} \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_j^\top \mathbf{z})^2 (\mathbf{A}_j^\top \mathbf{x}^*)^2 &= \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{z}_S)^2 (\mathbf{A}_j^\top \mathbf{x}^*)^2 + \frac{2}{m} \sum_{j=1}^m (\mathbf{A}_{j,S}^\top \mathbf{z}_S) (\mathbf{A}_{j,S^c}^\top \mathbf{z}_{S^c}) (\mathbf{A}_j^\top \mathbf{x}^*)^2 \\ &\quad + \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_{j,S^c}^\top \mathbf{z}_{S^c})^2 (\mathbf{A}_j^\top \mathbf{x}^*)^2 \\ &=: B_1 + B_2 + B_3. \end{aligned}$$

The same computation as in Step 2, part (b) shows that, with probability $1 - \frac{c_p}{27} n^{-10}$,

$$|B_1 - \mathbb{E}[B_1]| \leq \frac{\gamma}{12} \|\mathbf{z}_S\|_2^2.$$

As in Step 2, part (a), we can bound with probability $1 - \frac{c_p}{27}n^{-10}$,

$$\begin{aligned} |B_2 - \mathbb{E}[B_2]| &= 2 \left| \sum_{i \notin \mathcal{S}} z_i \sum_{l \in \mathcal{S}} z_l \frac{1}{m} \sum_{j=1}^m A_{ji} A_{jl} (\mathbf{A}_j^\top \mathbf{x}^*)^2 \right| \\ &\leq 2 \|\mathbf{z}_{\mathcal{S}^c}\|_1 \|\mathbf{z}_{\mathcal{S}}\|_1 \max_{i \notin \mathcal{S}, l \in \mathcal{S}} \left| \frac{1}{m} \sum_{j=1}^m A_{ji} A_{jl} (\mathbf{A}_j^\top \mathbf{x}^*)^2 \right| \\ &\leq \frac{\gamma}{6} \|\mathbf{z}_{\mathcal{S}}\|_2 \delta, \end{aligned}$$

where we used Hölder's inequality, (C.5) of Lemma C.5 together with $m \geq c_s(\gamma)k^2 \log^2 n$, and the fact that $\|\mathbf{z}_{\mathcal{S}}\|_1 \leq \sqrt{k} \|\mathbf{z}_{\mathcal{S}}\|_2$. The same argument gives

$$|B_3 - \mathbb{E}[B_3]| \leq \frac{\gamma}{12} \delta^2$$

with probability $1 - \frac{c_p}{27}n^{-10}$. Combining these bounds, we have, with probability $1 - \frac{c_p}{9}n^{-10}$,

$$\left| \frac{2}{m} \sum_{j=1}^m (\mathbf{A}_j^\top \mathbf{z})^2 (\mathbf{A}_j^\top \mathbf{x}^*)^2 - \mathbb{E} \left[\frac{2}{m} \sum_{j=1}^m (\mathbf{A}_j^\top \mathbf{z})^2 (\mathbf{A}_j^\top \mathbf{x}^*)^2 \right] \right| \leq \frac{\gamma}{3} (\|\mathbf{z}_{\mathcal{S}}\|_2^2 + \delta^2),$$

where we used the inequality $2ab \leq a^2 + b^2$. Repeating these steps for the other two terms in (B.4) completes the proof that, with probability $1 - \frac{c_p}{3}n^{-10}$, we have

$$|\langle \nabla \tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle - \mathbb{E}[\langle \nabla \tilde{F}(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle]| \leq \gamma (\|\mathbf{x}_{\mathcal{S}} - \mathbf{x}_{\mathcal{S}}^*\|_2^2 + \delta^2) \quad \text{for all } \mathbf{x} \in \mathcal{X}.$$

Finally, the simplified bound can be obtained directly by plugging in the additional assumptions. \square

As the proof of Lemma 6.3 involves additional technical challenges in discrete time, we present the proof of Lemma 6.3 separately for the continuous-time and discrete-time cases.

Proof of Lemma 6.3 in the continuous-time case We show that the three inequalities (6.2), (6.3) and (6.4) are satisfied by showing that, as long as all inequalities are satisfied, neither can be violated first.

Let I_0 be the index for which we have the non-zero initialization $X_{I_0}(0) > 0$. As we can only recover the signal \mathbf{x}^* up to a global sign from phaseless measurements, we can assume without loss of generality that $x_{I_0}^* > 0$, since we can otherwise replace \mathbf{x}^* by $-\mathbf{x}^*$ in the proof below. That is, we need to show (6.2) with $\xi = +1$.

At $t = 0$, (6.2) is satisfied by the definition of the initialization. Using standard concentration bounds for sub-exponential random variables (e.g. Prop. 5.16 of [8]), we can bound with probability $1 - 4n^{-10}$,

$$\frac{1}{m} \sum_{j=1}^m Y_j = \frac{1}{m} \sum_{j=1}^m (\mathbf{A}_j^\top \mathbf{x}^*)^2 + \frac{1}{m} \sum_{j=1}^m \varepsilon_j > 1 - (9 + 25\sigma) \sqrt{\frac{\log n}{m}}.$$

Hence, the initialization (3.4) satisfies $\|\mathbf{X}(0)\|_2^2 \geq \frac{1}{3} - (3 + 9\sigma) \sqrt{\frac{\log n}{m}}$. Similarly, we have $\|\mathbf{X}(0)\|_2^2 \leq 2$ and therefore (6.3). Finally, inequality (6.4) is satisfied at $t = 0$ since we have

$$\mathbf{X}(0)^\top \mathbf{x}^* \geq \|\mathbf{X}_{\mathcal{S}}(0)\|_1 x_{\min}^* \geq \frac{c_*}{2\sqrt{k}} \geq (9 + 25\sigma) \sqrt{\frac{\log n}{m}} \geq 3\|\mathbf{X}(0)\|_2^2 - 1.$$

Step 1: (6.2) continues to hold as long as (6.3) holds

We prove this inequality by contradiction. Define $T_1 = \inf\{t \geq 0 : X_i(t)x_i^* < 0 \text{ for some } i\}$ as the first time inequality (6.2) is violated. Assume that $T_1 < T$ and that (6.3) holds for all $t \leq T_1$. Let i be the index for which $X_i(t)x_i^* < 0$ first occurs. This is only possible for a coordinate $i \in \mathcal{S}$, and by continuity we must have $X_i(T_1) = 0$. Without loss of generality, assume that $x_i^* > 0$. We will show that

$$\frac{d}{dt}X_i(T_1) = -\sqrt{X_i(T_1)^2 + \beta^2} \nabla F(\mathbf{X}(T_1))_i > 0,$$

that is $X_i(t)$ must become positive for t close enough to T_1 , which is a contradiction to the definition of T_1 , and we hence must have $T_1 \geq T$. We can bound, since both (6.2) and (6.3) hold at T_1 ,

$$\mathbf{X}(T_1)^\top \mathbf{x}^* \geq \|\mathbf{X}_{\mathcal{S}}(T_1)\|_1 x_{\min}^* \geq \|\mathbf{X}_{\mathcal{S}}(t_1)\|_1 \frac{c_\star^*}{\sqrt{k}}, \quad (\text{B.5})$$

and hence,

$$\nabla f(\mathbf{X}(T_1))_i = -2(\mathbf{X}(T_1)^\top \mathbf{x}^*)x_i^* \leq -2\|\mathbf{X}_{\mathcal{S}}(T_1)\|_1 \frac{c_\star^2}{k},$$

where we used $x_i^* \geq x_{\min}^* \geq \frac{c_\star^*}{\sqrt{k}}$. As we assume $T_1 < T$, we can use the simplified bound in Lemma 6.1 to bound

$$|\nabla F(\mathbf{X}(T_1))_i - \nabla f(\mathbf{X}(T_1))_i| \leq 0.1 \frac{c_\star^2 \|\mathbf{X}_{\mathcal{S}}(T_1)\|_1}{k} \leq \frac{1}{20} |\nabla f(\mathbf{X}(T_1))_i|$$

with probability $1 - c_p n^{-10}$ if c_s is sufficiently large, since we have $\|\mathbf{X}_{\mathcal{S}^c}(T_1)\|_1 \leq \delta \leq c_1/n$ by assumption. Hence, we have $\nabla F(\mathbf{X}(T_1))_i < 0$, which implies $\frac{d}{dt}X_i(T_1) > 0$ and contradicts the definition of T_1 . Therefore, we must have $T_1 \geq T$.

Step 2: (6.3) continues to hold as long as (6.2) and (6.4) hold

We begin by showing the lower bound in (6.3).

Define $T_2 = \inf\{t \geq 0 : \|\mathbf{X}(t)\|_2^2 < \frac{1}{3} - (3+9\sigma)\sqrt{\frac{\log n}{m}}\}$ and assume that $T_2 < T$ and that inequalities (6.2) and (6.4) are satisfied for all $t \leq T_2$. By continuity, we must have $\|\mathbf{X}(T_2)\|_2^2 = \frac{1}{3} - (3+9\sigma)\sqrt{\frac{\log n}{m}}$, and we will show that

$$\frac{d}{dt}\|\mathbf{X}(T_2)\|_2^2 = -2 \sum_{i=1}^n X_i(T_2) \frac{d}{dt}X_i(T_2) = -2 \sum_{i=1}^n X_i(T_2) \sqrt{X_i(T_2)^2 + \beta^2} \nabla F(\mathbf{X}(T_2))_i$$

is positive, which implies $\|\mathbf{X}(t)\|_2^2 > \frac{1}{3} - (3+9\sigma)\sqrt{\frac{\log n}{m}}$ for t close enough to T_2 and contradicts the definition of T_2 . Hence, we must have $T_2 \geq T$. Since $3\|\mathbf{X}(T_2)\|_2^2 - 1 < 0$, we can bound the population gradient using (B.5) and the assumption $\|\mathbf{X}_{\mathcal{S}^c}(t)\|_1 \leq \delta$:

$$\nabla f(\mathbf{X}(T_2))_i \begin{cases} \leq -2\|\mathbf{X}_{\mathcal{S}}(T_2)\|_1 \frac{c_\star^2}{k} & x_i^* > 0 \\ \in (-\delta, \delta) & x_i^* = 0 \\ \geq 2\|\mathbf{X}_{\mathcal{S}}(T_2)\|_1 \frac{c_\star^2}{k} & x_i^* < 0. \end{cases} \quad (\text{B.6})$$

Recall that in the previous step we have shown, with probability $1 - c_p n^{-10}$ and for all $i \in [n]$,

$$|\nabla F(\mathbf{X}(T_1))_i - \nabla f(\mathbf{X}(T_1))_i| \leq 0.1 \frac{c_\star^2 \|\mathbf{X}_S(T_1)\|_1}{k}, \quad (\text{B.7})$$

which results in bounds analogous to (B.6) for the gradient $\nabla F(\mathbf{X}(T_2))_i$. In order to bound $\frac{d}{dt} \|\mathbf{X}(T_2)\|_2^2$, we write

$$\left| \sum_{i \notin S} X_i(T_2) \sqrt{X_i(T_2)^2 + \beta^2} \nabla F(\mathbf{X}(T_2))_i \right| \leq \|\mathbf{X}_{S^c}(T_2)\|_1 \leq \delta,$$

where we used $|\sqrt{X_i(T_2)^2 + \beta^2} \nabla F(\mathbf{X}(T_2))_i| \leq 1$. Using $\|\mathbf{X}_S(T_2)\|_1 \geq \frac{1}{2}$, we have

$$-\sum_{i \in S} X_i(T_2) \sqrt{X_i(T_2)^2 + \beta^2} \nabla F(\mathbf{X}(T_2))_i \geq \sum_{i \in S} X_i(T_2)^2 \frac{c_\star^2}{2k} \geq 0.15 \frac{c_\star^2}{k},$$

where we used $\|\mathbf{X}_S(T_2)\|_2^2 = \|\mathbf{X}(T_2)\|_2^2 - \|\mathbf{X}_{S^c}(T_2)\|_2^2 \geq 0.3$. This shows that $\frac{d}{dt} \|\mathbf{X}(T_2)\|_2^2 > 0$, which contradicts the definition of T_2 . Therefore, we must have $T_2 \geq T$.

The upper bound in (6.3) is an immediate consequence of (6.4): by the Cauchy–Schwarz inequality we have $\mathbf{X}(t)^\top \mathbf{x}^\star \leq \|\mathbf{X}(t)\|_2$, so, for $3\|\mathbf{X}(t)\|_2^2 - 1 > 0$, we can bound

$$\frac{2\|\mathbf{X}(t)\|_2}{3\|\mathbf{X}(t)\|_2^2 - 1} \geq \frac{2(\mathbf{X}(t)^\top \mathbf{x}^\star)}{3\|\mathbf{X}(t)\|_2^2 - 1} \geq \frac{1}{\sqrt{3}} \Rightarrow \|\mathbf{X}(t)\|_2 \leq \frac{1 + \sqrt{2}}{\sqrt{3}} < \sqrt{2}$$

by solving the quadratic form.

Step 3: (6.4) continues to hold as long as (6.2) and (6.3) hold

The proof of (6.4) follows the same recipe as the two previous steps, although the calculations are more complicated. When $3\|\mathbf{X}(t)\|_2^2 - 1 \leq 0$, there is nothing to show. Otherwise, we can consider the ratio

$$R(t) = \frac{2(\mathbf{X}(t)^\top \mathbf{x}^\star)}{3\|\mathbf{X}(t)\|_2^2 - 1}$$

and show that it is bounded from below by $1/\sqrt{3}$ for all $t \leq T$.

Let $T_3 = \inf\{t \geq 0 : R(t) < 1/\sqrt{3}\}$ and assume $T_3 < T$ as before. For notational simplicity, we will omit the argument T_3 in $X_i(T_3)$ in what follows. We can compute

$$\frac{d}{dt} R(T_3) = \sum_{i=1}^n -\sqrt{X_i^2 + \beta^2} \nabla F(\mathbf{X})_i \frac{2x_i^\star (3\|\mathbf{X}\|_2^2 - 1) - 2(\mathbf{X}^\top \mathbf{x}^\star) \cdot 6X_i}{(3\|\mathbf{X}\|_2^2 - 1)^2}. \quad (\text{B.8})$$

By continuity, we have $R(T_3) = 1/\sqrt{3}$. Then, for $X_i, x_i^\star > 0$, we have

$$\begin{aligned} 2x_i^\star (3\|\mathbf{X}\|_2^2 - 1) - 2(\mathbf{X}^\top \mathbf{x}^\star) \cdot 6X_i &> 0 \\ \Leftrightarrow X_i &< \frac{1}{\sqrt{3}} x_i^\star, \end{aligned} \quad (\text{B.9})$$

and the analogous result for $X_i, x_i^\star < 0$. The idea to showing $\frac{d}{dt} R(t_3) > 0$ is to show that coordinates with small magnitude $|X_i| < |x_i^\star|/\sqrt{3}$ are increasing in magnitude, and conversely coordinates with large

magnitude $|X_i| > |x_i^*|/\sqrt{3}$ are decreasing in magnitude. To this end, we split the coordinates $i \in [n]$ into five subsets:

$$\begin{aligned}\mathcal{S}^c &= \{i \in [n] : x_i^* = 0\}, \\ \mathcal{S}_1 &= \left\{i \in \mathcal{S} : |X_i| < \left(\frac{1}{\sqrt{3}} - 0.1\right)|x_i^*|\right\}, \\ \mathcal{S}_2 &= \left\{i \in \mathcal{S} : \left(\frac{1}{\sqrt{3}} - 0.1\right)|x_i^*| \leq |X_i| < \left(\frac{1}{\sqrt{3}} + 0.1\right)|x_i^*|\right\}, \\ \mathcal{S}_3 &= \left\{i \in \mathcal{S} : \left(\frac{1}{\sqrt{3}} + 0.1\right)|x_i^*| \leq |X_i| < \frac{2}{\sqrt{3}}|x_i^*|\right\}, \\ \mathcal{S}_4 &= \left\{i \in \mathcal{S} : |X_i| \geq \frac{2}{\sqrt{3}}|x_i^*|\right\}.\end{aligned}$$

We bound the sum (B.8) on each of these five sets.

- For \mathcal{S}^c , we have

$$\left| \sum_{i \in \mathcal{S}^c} \sqrt{X_i^2 + \beta^2} \nabla F(\mathbf{X})_i \cdot 2(\mathbf{X}^\top \mathbf{x}^*) \cdot 6X_i \right| \leq \frac{2c_\star^2 \delta^2}{k} \|\mathbf{X}_\mathcal{S}\|_1 (\mathbf{X}^\top \mathbf{x}^*),$$

where we used that $\sqrt{X_i^2 + \beta^2} \leq \sqrt{2}\delta$, $\|\mathbf{X}_{\mathcal{S}^c}\|_1^2 \leq \delta^2$ and that we can bound $|\nabla F(\mathbf{X})_i| \leq 0.1\|\mathbf{X}_\mathcal{S}\|_1 \frac{c_\star^2}{k}$ using Lemma 6.1 as in (B.7).

- For an $i \in \mathcal{S}_1$ with $x_i^* > 0$, we have

$$\nabla f(\mathbf{X})_i = (3\|\mathbf{X}\|_2^2 - 1)X_i - 2(\mathbf{X}^\top \mathbf{x}^*)x_i^* \leq -0.2 \cdot \sqrt{3}\|\mathbf{X}_\mathcal{S}\|_1 \frac{c_\star^2}{k},$$

where we used that $3\|\mathbf{X}\|_2^2 - 1 = \sqrt{3} \cdot 2(\mathbf{X}^\top \mathbf{x}^*)$, $X_i < (1/\sqrt{3} - 0.1)x_i^*$ and (B.5). Together with the bound (B.7), this shows that $\nabla F(\mathbf{X})_i < 0$. Similarly, we can show $\nabla F(\mathbf{X})_i > 0$ for $i \in \mathcal{S}_1$ with $x_i^* < 0$. Hence, recalling (B.9),

$$\sum_{i \in \mathcal{S}_1} -\sqrt{X_i^2 + \beta^2} \nabla F(\mathbf{X})_i (2x_i^* (3\|\mathbf{X}\|_2^2 - 1) - 2(\mathbf{X}^\top \mathbf{x}^*) \cdot 6X_i) \geq 0,$$

as each summand is non-negative.

- For \mathcal{S}_3 , we can use the same argument as for \mathcal{S}_1 to show that

$$\sum_{i \in \mathcal{S}_3} -\sqrt{X_i^2 + \beta^2} \nabla F(\mathbf{X})_i (2x_i^* (3\|\mathbf{X}\|_2^2 - 1) - 2(\mathbf{X}^\top \mathbf{x}^*) \cdot 6X_i) \geq 0.$$

- For \mathcal{S}_2 , we need to show that the sum

$$\sum_{i \in \mathcal{S}_2} -\sqrt{X_i^2 + \beta^2} \nabla F(\mathbf{X})_i (2x_i^* (3\|\mathbf{X}\|_2^2 - 1) - 2(\mathbf{X}^\top \mathbf{x}^*) \cdot 6X_i)$$

is bounded from below. Let $i \in \mathcal{S}_2$ with $x_i^* > 0$. If $X_i < x_i^*/\sqrt{3}$ and $\nabla F(\mathbf{X})_i < 0$, or if $X_i > x_i^*/\sqrt{3}$ and $\nabla F(\mathbf{X})_i > 0$, then the summand is non-negative, i.e. bounded from below by zero. For $i \in \mathcal{S}_2$ with $X_i < x_i^*/\sqrt{3}$ and $\nabla F(\mathbf{X})_i > 0$, we can bound

$$\sqrt{X_i^2 + \beta^2} \leq (1 + \beta)|X_i|,$$

$$\nabla f(\mathbf{X})_i < 0,$$

$$\nabla F(\mathbf{X})_i \leq \nabla f(\mathbf{X})_i + |\nabla f(\mathbf{X})_i - \nabla F(\mathbf{X})_i| \leq 0.1 \|\mathbf{X}_{\mathcal{S}}\|_1 \frac{c_\star^2}{k},$$

where we used the bound (B.7). Recalling the definition of \mathcal{S}_2 , we have

$$2x_i^*(3\|\mathbf{X}\|_2^2 - 1) - 2(\mathbf{X}^\top \mathbf{x}^\star) \cdot 6X_i = 12(\mathbf{X}^\top \mathbf{x}^\star) \left(\frac{1}{\sqrt{3}}x_i^* - X_i \right) \leq 3(\mathbf{X}^\top \mathbf{x}^\star)X_i.$$

Putting this together, we can bound

$$-\sqrt{X_i^2 + \beta^2} \nabla F(\mathbf{X})_i (2x_i^*(3\|\mathbf{X}\|_2^2 - 1) - 2(\mathbf{X}^\top \mathbf{x}^\star) \cdot 6X_i) \geq -\frac{0.3(1 + \beta)c_\star^2}{k} \|\mathbf{X}_{\mathcal{S}}\|_1 (\mathbf{X}^\top \mathbf{x}^\star) X_i^2.$$

Together with the analogous bounds for the cases $X_i > x_i^*/\sqrt{3}$ and $\nabla F(\mathbf{X})_i < 0$, and for $i \in \mathcal{S}_2$ with $x_i^* < 0$, this yields

$$\begin{aligned} & \sum_{i \in \mathcal{S}_2} -\sqrt{X_i^2 + \beta^2} \nabla F(\mathbf{X})_i (2x_i^*(3\|\mathbf{X}\|_2^2 - 1) - 2(\mathbf{X}^\top \mathbf{x}^\star) \cdot 6X_i) \\ & \geq -\frac{0.3(1 + \beta)c_\star^2}{k} \|\mathbf{X}_{\mathcal{S}}\|_1 (\mathbf{X}^\top \mathbf{x}^\star) \|\mathbf{X}_{\mathcal{S}_2}\|_2^2. \end{aligned}$$

- Finally, for $i \in \mathcal{S}_4$ we have

$$\sqrt{X_i^2 + \beta^2} \geq |X_i|,$$

$$|\nabla f(\mathbf{X})_i| = 2(\mathbf{X}^\top \mathbf{x}^\star) |\sqrt{3}X_i - x_i^*| \geq 2(\mathbf{X}^\top \mathbf{x}^\star) |x_i^*| \geq 2\|\mathbf{X}_{\mathcal{S}}\|_1 \frac{c_\star^2}{k},$$

$$|\nabla F(\mathbf{X})_i| \geq |\nabla f(\mathbf{X})_i| - |\nabla f(\mathbf{X})_i - \nabla F(\mathbf{X})_i| \geq 1.9\|\mathbf{X}_{\mathcal{S}}\|_1 \frac{c_\star^2}{k},$$

where we used the bound (B.7). Recalling the definition of \mathcal{S}_4 , we can bound

$$|2x_i^*(3\|\mathbf{X}\|_2^2 - 1) - 2(\mathbf{X}^\top \mathbf{x}^\star) \cdot 6X_i| = 12(\mathbf{X}^\top \mathbf{x}^\star) \left| \frac{1}{\sqrt{3}}x_i^* - X_i \right| \geq 6(\mathbf{X}^\top \mathbf{x}^\star)X_i.$$

Putting everything together, we can bound

$$\sum_{i \in \mathcal{S}_2} -\sqrt{X_i^2 + \beta^2} \nabla F(\mathbf{X})_i (2x_i^*(3\|\mathbf{X}\|_2^2 - 1) - 2(\mathbf{X}^\top \mathbf{x}^\star) \cdot 6X_i) \frac{11.4c_\star^2}{k} \|\mathbf{X}_{\mathcal{S}}\|_1 (\mathbf{X}^\top \mathbf{x}^\star) \|\mathbf{X}_{\mathcal{S}_4}\|_2^2.$$

Putting these five sums together, we have shown that $\frac{d}{dt}R(T_3) > 0$ if we can show that

$$\frac{11.4c_*^2}{k}\|\mathbf{X}_S\|_1(\mathbf{X}^\top \mathbf{x}^*)\|\mathbf{X}_{S_4}\|_2^2 \geq \left(\frac{0.3(1+\beta)c_*^2}{k}\|\mathbf{X}_{S_2}\|_2^2 + \frac{2c_*^2\delta^2}{k}\right)\|\mathbf{X}_S\|_1(\mathbf{X}^\top \mathbf{x}^*).$$

Since $\delta \leq c_1/n$ is sufficiently small compared to $\|\mathbf{X}_{S_4}\|_2$, this reduces to showing

$$11\|\mathbf{X}_{S_4}\|_2^2 \geq 0.3(1+\beta)\|\mathbf{X}_{S_2}\|_2^2. \quad (\text{B.10})$$

Now, we can rearrange the equality $R(t) = 1/\sqrt{3}$ to obtain

$$\begin{aligned} & 3\left(\|\mathbf{X}_{S_1}\|_2^2 + \|\mathbf{X}_{S_2}\|_2^2 + \|\mathbf{X}_{S_3}\|_2^2 + \|\mathbf{X}_{S_4}\|_2^2 + \|\mathbf{X}_{S^c}\|_2^2\right) - 1 \\ &= \sqrt{3} \cdot 2\left(\mathbf{X}_{S_1}^\top \mathbf{x}_{S_1}^* + \mathbf{X}_{S_2}^\top \mathbf{x}_{S_2}^* + \mathbf{X}_{S_3}^\top \mathbf{x}_{S_3}^* + \mathbf{X}_{S_4}^\top \mathbf{x}_{S_4}^*\right). \end{aligned}$$

By definition, we have $3X_i^2 < \sqrt{3} \cdot 2X_i x_i^*$ for $i \in S_1 \cup S_2 \cup S_3 \cup S_4$, which gives

$$3\|\mathbf{X}_{S_j}\|_2^2 < \sqrt{3} \cdot 2\mathbf{X}_{S_j}^\top \mathbf{x}_{S_j}^*, \quad \text{for } j = 1, \dots, 4.$$

Further, we have $\|\mathbf{X}_{S^c}\|_2^2 \leq \delta$, so

$$3\|\mathbf{X}_{S_4}\|_2^2 - 2\sqrt{3} \cdot \mathbf{X}_{S_4}^\top \mathbf{x}_{S_4}^* > 1 - 3\delta + \left(2\sqrt{3} \cdot \mathbf{X}_{S_2}^\top \mathbf{x}_{S_2}^* - 3\|\mathbf{X}_{S_2}\|_2^2\right).$$

By the definition of S_2 , we have for $i \in S_2$,

$$2\sqrt{3}X_i x_i^* - 3X_i^2 \geq \left(\frac{2\sqrt{3}}{1/\sqrt{3} + 0.1} - 3\right)X_i^2 \geq 2.2X_i^2.$$

Since $\mathbf{X}_{S_4}^\top \mathbf{x}_{S_4}^* \geq 0$ by (6.2), this gives

$$3\|\mathbf{X}_{S_4}\|_2^2 > 1 - 3\delta + 2.2\|\mathbf{X}_{S_2}\|_2^2,$$

which shows that (B.10) holds, thus completing the proof of (6.4). \square

Proof of Lemma 6.3 in the discrete-time case. We will prove the inequalities (6.2), (6.3) and (6.4) via induction. Let I_0 be the index for which we have the non-zero initialization $X_{I_0}^0 > 0$. As in the continuous-time case, we can assume without loss of generality that $x_{I_0}^* > 0$ and show (6.2) with $\xi = +1$. The inequalities (6.2), (6.3) and (6.4) hold at $t = 0$ as in the continuous-time case.

Step 1: (6.2) continues to hold as long as (6.3) holds

For $t \geq 0$, let $i \in S$ with $x_i^* \geq 0$. Negative coordinates can be treated the same way, and for $i \notin S$ there is nothing to show. Rearranging (A.1) shows that $X_i^{t+1} \geq 0$ if and only if

$$\frac{\beta^2}{(\sqrt{(X_i^t)^2 + \beta^2} + X_i^t)^2} = \frac{\sqrt{(X_i^t)^2 + \beta^2} - X_i^t}{\sqrt{(X_i^t)^2 + \beta^2} + X_i^t} \leq \exp(-2\eta \nabla F(\mathbf{X}^t)_i). \quad (\text{B.11})$$

In order to show (B.11), we need to bound the gradient $\nabla F(\mathbf{X}^t)_i$ from above. Since by the induction hypothesis we have $\|\mathbf{X}_S^t\|_1 = \|\mathbf{X}^t\|_1 - \|\mathbf{X}_{S^c}^t\|_1 \geq \frac{1}{2}$, we can apply the simplified bound in Lemma 6.1

to bound, for $t \leq T$,

$$|\nabla F(\mathbf{X}^t)_i - \nabla f(\mathbf{X}^t)_i| \leq \frac{c_\star^2 \|\mathbf{X}_S^t\|_1}{k}$$

with probability $1 - c_p n^{-10}$, provided that c_s is sufficiently large, where we used the assumption that $\delta \leq c_1/n$. Further, we can bound, since (6.2) holds at time t by induction hypothesis,

$$(\mathbf{X}^t)^\top \mathbf{x}^\star \geq \|\mathbf{X}_S^t\|_1 x_{\min}^\star \geq \|\mathbf{X}_S^t\|_1 \frac{c_\star}{\sqrt{k}},$$

which leads to

$$\nabla F(\mathbf{X}^t)_i \leq \nabla f(\mathbf{X}^t)_i + |\nabla F(\mathbf{X}^t)_i - \nabla f(\mathbf{X}^t)_i| \leq (3\|\mathbf{X}^t\|_2^2 - 1)X_i^t - ((\mathbf{X}^t)^\top \mathbf{x}^\star)x_i^\star. \quad (\text{B.12})$$

If $X_i^t \leq \frac{c_\star}{2\sqrt{3k}}$, then $\nabla F(\mathbf{X}^t)_i \leq 0$ by (6.4). In this case, (B.11) holds as

$$\frac{\beta^2}{(\sqrt{(X_i^t)^2 + \beta^2} + X_i^t)^2} \leq 1 \leq \exp(-2\eta \nabla F(\mathbf{X}^t)_i).$$

If on the other hand $X_i^t > \frac{c_\star}{2\sqrt{3k}}$, then we can bound $\nabla F(\mathbf{X}^t)_i \leq 5X_i^t$ since $\|\mathbf{X}^t\|_2^2 \leq 2$ and

$$\frac{\beta^2}{(\sqrt{(X_i^t)^2 + \beta^2} + X_i^t)^2} \leq \frac{\beta^2}{4(X_i^t)^2} \leq \exp(-10\eta X_i^t)$$

holds for all $\eta \leq \frac{1}{\sqrt{200}}$, where we used that $X_i^t \leq \sqrt{2}$ by (6.3) and $\log \frac{c_\star^2}{3k\beta^2} \geq 1$ since by assumption $\beta \leq c_1/n$. This completes the proof that (B.11) holds. Repeating the same argument for $i \in \mathcal{S}$ with $x_i^\star, X_i^t \leq 0$ shows that (6.2) is satisfied at $t+1$.

Step 2: (6.3) continues to hold as long as (6.2) and (6.4) hold

We begin by inductively showing the lower bound in (6.3). For $t \geq 0$, we write

$$\|\mathbf{X}^{t+1}\|_2^2 - \|\mathbf{X}^t\|_2^2 = \sum_{i=1}^n (X_i^{t+1} + X_i^t)(X_i^{t+1} - X_i^t)$$

and consider the partial sums containing all $i \in \mathcal{S}$ and $i \notin \mathcal{S}$, respectively. For $i \notin \mathcal{S}$, we have

$$\sum_{i \notin \mathcal{S}} (X_i^{t+1} + X_i^t)(X_i^{t+1} - X_i^t) \geq -\sum_{i \notin \mathcal{S}} (X_i^t)^2 \geq -n\delta^2,$$

where we used $\|\mathbf{X}_{\mathcal{S}^c}^t\|_2 \leq \sqrt{n}\|\mathbf{X}_{\mathcal{S}^c}^t\|_1 \leq \sqrt{n}\delta$.

To bound the partial sum containing $i \in \mathcal{S}$, we distinguish the two cases $\|\mathbf{X}^t\|_2^2 - \frac{1}{3} \geq 2n\delta^2$ and $\|\mathbf{X}^t\|_2^2 - \frac{1}{3} < 2n\delta^2$. In the first case, let $i \in \mathcal{S}$ and without loss of generality $x_i^\star > 0$. If $X_i^{t+1} \geq X_i^t$, then the summand $(X_i^{t+1} + X_i^t)(X_i^{t+1} - X_i^t)$ is non-negative. If on the other hand $X_i^{t+1} < X_i^t$, then $\nabla F(\mathbf{X}^t)_i > 0$, and we can bound, recalling (A.4),

$$X_i^{t+1} - X_i^t \geq -\frac{3}{2}\eta \nabla F(\mathbf{X}^t)_i \sqrt{(X_i^t)^2 + \beta^2} \geq -\frac{3}{2}\eta (3\|\mathbf{X}^t\|_2^2 - 1)X_i^t \sqrt{(X_i^t)^2 + \beta^2},$$

where we used (B.12) for the last inequality. An analogous upper bound can be derived for negative coordinates with $x_i^\star < 0$, so writing $\mathcal{S}' = \{i \in \mathcal{S} : |X_i^{t+1}| < |X_i^t|\}$ for the coordinates that decrease in

magnitude, we can bound

$$\begin{aligned} \sum_{i \in \mathcal{S}'} (X_i^{t+1} + X_i^t)(X_i^{t+1} - X_i^t) &\geq -3\eta \sum_{i \in \mathcal{S}'} X_i^t (3\|\mathbf{X}^t\|_2^2 - 1) X_i^t \sqrt{(X_i^t)^2 + \beta^2} \\ &\geq -\left(9\eta \sum_{i \in \mathcal{S}'} (X_i^t)^2 (|X_i^t| + \beta)\right) \left(\|\mathbf{X}^t\|_2^2 - \frac{1}{3}\right), \end{aligned}$$

where for the last inequality we used that $\sqrt{x^2 + \beta^2} \leq |x| + \beta$ for $\beta \geq 0$. Hence, we have

$$\|\mathbf{X}^{t+1}\|_2^2 - \|\mathbf{X}^t\|_2^2 \geq \frac{1}{3} - \|\mathbf{X}^t\|_2^2$$

for $\eta \leq \frac{1}{54}$, where we used that $\sum_{i \in \mathcal{S}'} (X_i^t)^2 (|X_i^t| + \beta) \leq 3$ since the upper bound in (6.3) holds at time t by induction hypothesis.

Next, we consider the case $\|\mathbf{X}^t\|_2^2 - \frac{1}{3} < 2n\delta^2$. In this case, it follows from (B.12) that every coordinate $i \in \mathcal{S}$ must be increasing in magnitude, i.e. $|X_i^{t+1}| > |X_i^t|$ for all $i \in \mathcal{S}$. Assuming without loss of generality that $x_i^* > 0$, we can use (A.5) to bound by how much X_i^t must at least increase:

$$\begin{aligned} X_i^{t+1} - X_i^t &\geq -\frac{1}{2}\eta \nabla F(\mathbf{X}^t)_i \sqrt{(X_i^t)^2 + \beta^2} \\ &\geq -\frac{1}{2}\eta \left((3\|\mathbf{X}^t\|_2^2 - 1) X_i^t - (\mathbf{X}^t)^\top \mathbf{x}^* \right) \sqrt{(X_i^t)^2 + \beta^2} \\ &\geq -\frac{1}{2}\eta \left(6n\delta^2 X_i^t - \frac{c_\star^2}{2k} \right) \sqrt{(X_i^t)^2 + \beta^2} \\ &\geq \frac{c_\star^2 \eta}{5k} X_i^t, \end{aligned}$$

where we used (6.6) together with the assumption $x_{\min}^* \geq \frac{c_\star}{\sqrt{k}}$, and $\delta \leq \frac{c_1}{n}$. Hence, we have

$$\sum_{i \in \mathcal{S}} (X_i^{t+1} + X_i^t)(X_i^{t+1} - X_i^t) \geq \sum_{i \in \mathcal{S}} 2X_i^t \frac{c_\star^2 \eta}{5k} X_i^t \geq \eta \frac{2c_\star^2}{5k} \left(\frac{1}{3} - (3 + 9\sigma) \sqrt{\frac{\log n}{m}} - \delta^2 \right),$$

where we used the induction hypothesis that the lower bound in (6.3) is satisfied at time t and the bound $\|\mathbf{X}_{\mathcal{S}^c}^t\|_2^2 \leq \|\mathbf{X}_{\mathcal{S}^c}^t\|_1^2 \leq \delta^2$. Thus, we have $\|\mathbf{X}^{t+1}\|_2^2 - \|\mathbf{X}^t\|_2^2 > 0$ if $\delta \leq c_1/n$ for $c_1 > 0$ sufficiently small. The upper bound in (6.3) can be shown the same way as in the continuous-time case.

Step 3: (6.4) continues to hold as long as (6.2) and (6.3) hold

As in the continuous-time case, we can restrict our attention to the case $3\|\mathbf{X}^t\|_2^2 - 1 > 0$ (since otherwise there is nothing to show), and consider the ratio

$$R_t = \frac{2((\mathbf{X}^t)^\top \mathbf{x}^*)}{3\|\mathbf{X}^t\|_2^2 - 1}.$$

As in the continuous-time case, we can show $R_0 \geq 1/\sqrt{3}$.

To show $R_{t+1} \geq 1/\sqrt{3}$ for any $t \geq 0$, we separately consider the two cases $R_t \geq 1/\sqrt{3} + 0.05$ and $1/\sqrt{3} \leq R_t < 1/\sqrt{3} + 0.05$. We assume $3\|\mathbf{X}^t\|_2^2 - 1 > 0$ for convenience's sake, as the other case

$3\|\mathbf{X}^t\|_2^2 - 1 < 0$ can be treated the same way, but requires distinguishing cases at various points in the following argument, which we omit to avoid repetition.

Step 3, Case 1: $R_t \geq 1/\sqrt{3} + 0.05$

We consider four types of coordinates, defined by

$$\mathcal{S}^c = \{i \in [n] : x_i^* = 0\} \quad (\text{off-support})$$

$$\mathcal{S}_1 = \left\{i \in \mathcal{S} : |X_i^t| \leq \frac{1}{3}|x_i^*|\right\} \quad (\text{small coordinates})$$

$$\mathcal{S}_2 = \left\{i \in \mathcal{S} : |X_i^t| > \frac{1}{3}|x_i^*|, |X_i^{t+1}| > |X_i^t|\right\} \quad (\text{large increasing coordinates})$$

$$\mathcal{S}_3 = \left\{i \in \mathcal{S} : |X_i^t| > \frac{1}{3}|x_i^*|, |X_i^{t+1}| \leq |X_i^t|\right\} \quad (\text{large decreasing coordinates}).$$

We will show $R_{t+1} \geq 1/\sqrt{3}$ by considering the following sequence of vectors defined by

$$\begin{aligned} \mathbf{X}^{(1)} &= \mathbf{X}^t, \quad X_i^{(2)} = \begin{cases} X_i^{(1)} & i \notin \mathcal{S}^c \\ X_i^{t+1} & i \in \mathcal{S}^c \end{cases}, \quad X_i^{(3)} = \begin{cases} X_i^{(2)} & i \notin \mathcal{S}_2 \\ X_i^{t+1} & i \in \mathcal{S}_2 \end{cases}, \\ X_i^{(4)} &= \begin{cases} X_i^{(3)} & i \notin \mathcal{S}_3 \\ X_i^{t+1} & i \in \mathcal{S}_3 \end{cases}, \quad \mathbf{X}^{(5)} = \mathbf{X}^{t+1}, \end{aligned}$$

and bound the ratio $\frac{2((\mathbf{X}^{(i)})^\top \mathbf{x}^*)}{3\|\mathbf{X}^{(i)}\|_2^2 - 1}$ for each $i = 2, \dots, 5$.

Exchanging off-support coordinates

By assumption, we have $R_t \geq 1/\sqrt{3} + 0.05$, which can also be written as

$$\left(\frac{1}{\sqrt{3}} + 0.05\right)^{-1} \cdot 2((\mathbf{X}^{(1)})^\top \mathbf{x}^*) \geq 3\|\mathbf{X}^{(1)}\|_2^2 - 1.$$

As shown in Step 2, $\|\mathbf{X}^{(2)}\|_2^2 \leq \|\mathbf{X}^{(1)}\|_2^2 + n\delta^2$, so we have, since $((\mathbf{X}^{(1)})^\top \mathbf{x}^*) \geq \frac{c_*}{2\sqrt{k}}$,

$$\left(\frac{1}{\sqrt{3}} + 0.04\right)^{-1} \cdot 2((\mathbf{X}^{(2)})^\top \mathbf{x}^*) \geq 3\|\mathbf{X}^{(2)}\|_2^2 - 1, \quad (\text{B.13})$$

provided that $\delta \leq c_1/n$ is sufficiently small.

Exchanging large increasing coordinates

Let $i \in \mathcal{S}_2$ with $x_i^* > 0$. Analogous bounds for the case $x_i^* < 0$ can be derived the same way. As $X_i^{t+1} > X_i^t$, we must have $\nabla F(\mathbf{X}^t)_i < 0$. First, we bound by how much X_i^t can increase in one iteration. We can write

$$\nabla f(\mathbf{X}^t)_i = \frac{2((\mathbf{X}^t)^\top \mathbf{x}^*)}{R_t} X_i^t - 2((\mathbf{X}^t)^\top \mathbf{x}^*) x_i^* = -2((\mathbf{X}^t)^\top \mathbf{x}^*) \left(x_i^* - \frac{1}{R_t} X_i^t\right).$$

Using the simplified bound of Lemma 6.1, we have with probability $1 - c_p n^{-10}$,

$$|\nabla F(\mathbf{X}^t)_i - \nabla f(\mathbf{X}^t)_i| \leq 0.1 \|\mathbf{X}_S^t\|_1 \frac{c_*^2}{k} \leq 0.1 ((\mathbf{X}^t)^\top \mathbf{x}^*) x_i^*,$$

where we used (B.5). This gives the lower bound

$$\nabla F(\mathbf{X}^t)_i \geq -2((\mathbf{X}^t)^\top \mathbf{x}^*) \left(1.05x_i^* - \frac{1}{R_t} X_i^t \right) \geq -1.05 \cdot 2((\mathbf{X}^t)^\top \mathbf{x}^*) x_i^*, \quad (\text{B.14})$$

since we assumed $x_i^* > 0$, and hence $X_i^t \geq 0$ by (6.2). With this, we can use (A.4) to bound

$$X_i^{t+1} - X_i^t \leq -\frac{3}{2} \eta \nabla F(\mathbf{X}^t)_i \sqrt{(X_i^t)^2 + \beta^2} \leq 1.7\eta \cdot 2((\mathbf{X}^t)^\top \mathbf{x}^*) x_i^* |X_i^t|, \quad (\text{B.15})$$

where we used that $\sqrt{(X_i^t)^2 + \beta^2} \leq 1.05|X_i^t|$ for $i \in \mathcal{S}_2$. Further, we can use the upper bound in (6.3) to bound $X_i^{t+1} + X_i^t \leq 3$. With this, we obtain

$$\begin{aligned} \frac{2((\mathbf{X}^{(3)})^\top \mathbf{x}^*)}{3\|\mathbf{X}^{(3)}\|_2^2 - 1} &\geq \frac{2((\mathbf{X}^{(2)})^\top \mathbf{x}^*)}{3\|\mathbf{X}^{(2)}\|_2^2 - 1 + \sum_{i \in \mathcal{S}_2} 3(X_i^{t+1} + X_i^t)(X_i^{t+1} - X_i^t)} \\ &\geq \frac{2((\mathbf{X}^{(2)})^\top \mathbf{x}^*)}{3\|\mathbf{X}^{(2)}\|_2^2 - 1 + 3 \cdot 3 \cdot 1.7 \cdot 2\eta((\mathbf{X}^t)^\top \mathbf{x}^*)^2}, \end{aligned}$$

which is bounded from below by $1/\sqrt{3}$, provided that

$$\eta \leq \frac{1}{193} \leq \frac{1}{15.3((\mathbf{X}^t)^\top \mathbf{x}^*)} \left(\sqrt{3} - \frac{3\|\mathbf{X}^{(2)}\|_2^2 - 1}{2((\mathbf{X}^{(2)})^\top \mathbf{x}^*)} \right),$$

where we used that $(\mathbf{X}^{(2)})^\top \mathbf{x}^* = (\mathbf{X}^t)^\top \mathbf{x}^* \leq \sqrt{2}$ by (6.3) and inequality (B.13).

Exchanging large decreasing coordinates

Let $i \in \mathcal{S}_3$ with $x_i^* > 0$. Analogous bounds for the case $x_i^* < 0$ can be derived the same way. We first bound by how much X_i^t can decrease in one iteration. Following the same steps as in the derivation of the lower bound (B.14), we obtain the analogous upper bound

$$\nabla F(\mathbf{X}^t)_i \leq 2((\mathbf{X}^t)^\top \mathbf{x}^*) \left(\frac{1}{R_t} X_i^t - 0.95x_i^* \right). \quad (\text{B.16})$$

Since we are considering a positive decreasing coordinate, we must have $\nabla F(\mathbf{X}^t)_i > 0$, and therefore, $X_i^t \geq 0.95R_t x_i^*$. As before, we can use (A.4) to bound

$$X_i^{t+1} - X_i^t \geq -\frac{3}{2} \eta \nabla F(\mathbf{X}^t)_i \sqrt{(X_i^t)^2 + \beta^2} \geq -3.2\eta((\mathbf{X}^t)^\top \mathbf{x}^*) \frac{1}{R_t} X_i^t (X_i^t - 0.95R_t x_i^*), \quad (\text{B.17})$$

where we used that $\sqrt{(X_i^t)^2 + \beta^2} \leq 1.05|X_i^t|$ for $i \in \mathcal{S}_3$. Further, we can write

$$X_i^{t+1} \geq \left(1 - 3.2\eta((\mathbf{X}^t)^\top \mathbf{x}^*) \frac{1}{R_t} X_i^t \right) X_i^t + 3.2\eta((\mathbf{X}^t)^\top \mathbf{x}^*) \frac{1}{R_t} X_i^t \cdot 0.95R_t x_i^* \geq 0.95R_t x_i^*,$$

provided that $\eta \leq \frac{1}{6.4\sqrt{3}} \leq \frac{R_t}{3.2((\mathbf{X}^t)^\top \mathbf{x}^*)X_i^t}$, where we used inequality (6.3) to bound $((\mathbf{X}^t)^\top \mathbf{x}^*)X_i^t \leq 2$ and $X_i^t \geq 0.95R_t x_i^*$. We can now bound

$$\begin{aligned} 3\|\mathbf{X}^{(4)}\|_2^2 - 1 &= 3\|\mathbf{X}^{(3)}\|_2^2 - 1 + \sum_{i \in \mathcal{S}_3} 3(X_i^{t+1} + X_i^t)(X_i^{t+1} - X_i^t) \\ &\leq 3\|\mathbf{X}^{(3)}\|_2^2 - 1 + \sum_{i \in \mathcal{S}_3} 3 \cdot 2 \cdot 0.95R_t x_i^* (X_i^{t+1} - X_i^t) \\ &\leq \sqrt{3} \left(2((\mathbf{X}^{(3)})^\top \mathbf{x}^*) + \sum_{i \in \mathcal{S}_3} 2x_i^* (X_i^{t+1} - X_i^t) \right) \\ &= \sqrt{3} \cdot 2((\mathbf{X}^{(4)})^\top \mathbf{x}^*), \end{aligned}$$

where the second inequality holds because $3 \cdot 0.95R_t \geq \sqrt{3}$ for $R_t \geq 1/\sqrt{3} + 0.05$, and $x_i^* (X_i^{t+1} - X_i^t) < 0$.

Exchanging small coordinates

Let $i \in \mathcal{S}_1$ with $x_i^* > 0$. Analogous bounds for the case $x_i^* < 0$ can be derived the same way. Then, we have $\nabla F(\mathbf{X}^t)_i < 0$ by (B.16) as $X_i^t \leq \frac{1}{3}x_i^*$, that is all small coordinates must increase in magnitude. Further, we can use (A.4) to bound

$$X_i^{t+1} + X_i^t \leq |X_i^{t+1} - X_i^t| + 2X_i^t \leq -\frac{3}{2}\eta \nabla F(\mathbf{X}^t)_i \sqrt{(X_i^t)^2 + \beta^2} + \frac{2}{3}x_i^* \leq \frac{2}{\sqrt{3}}x_i^*,$$

if $\eta \leq 0.3$, where we used that $(\mathbf{X}^t)^\top \mathbf{x}^* \leq \sqrt{2}$ by (6.3), $\sqrt{(X_i^t)^2 + \beta^2} \leq 0.35$ by the definition of \mathcal{S}_1 , and the lower bound (B.14) for the gradient $\nabla F(\mathbf{X}^t)_i$. Hence, we can bound

$$\begin{aligned} 3\|\mathbf{X}^{(5)}\|_2^2 - 1 &= 3\|\mathbf{X}^{(4)}\|_2^2 - 1 + \sum_{i \in \mathcal{S}_1} 3(X_i^{t+1} + X_i^t)(X_i^{t+1} - X_i^t) \\ &\leq \sqrt{3} \left(2((\mathbf{X}^{(4)})^\top \mathbf{x}^*) + \sum_{i \in \mathcal{S}_1} 2x_i^* (X_i^{t+1} - X_i^t) \right) \\ &= \sqrt{3} \cdot 2((\mathbf{X}^{(5)})^\top \mathbf{x}^*). \end{aligned}$$

This completes the proof of $R_{t+1} \geq 1/\sqrt{3}$ if $R_t \geq 1/\sqrt{3} + 0.05$.

Step 3, Case 2: $1/\sqrt{3} \leq R_t < 1/\sqrt{3} + 0.05$

The proof that $R_{t+1} \geq 1/\sqrt{3}$ largely follows the same steps as the previous case, although it requires a different sequence of vectors interpolating between \mathbf{X}^t and \mathbf{X}^{t+1} . Define $\tilde{\mathbf{X}}^{t+1}$ by

$$\nabla \Phi(\tilde{\mathbf{X}}^{t+1}) = \nabla \Phi(\mathbf{X}^t) - \eta \nabla f(\mathbf{X}^t),$$

that is the vector obtained by applying one step of mirror descent with the population gradient. Replacing X_i^{t+1} by \tilde{X}_i^{t+1} in the definitions of the sets \mathcal{S}_i , we consider the sequence

$$\begin{aligned} \mathbf{X}^{(1)} &= \mathbf{X}^t, \quad X_i^{(2)} = \begin{cases} X_i^{(1)} & i \notin \mathcal{S}_2 \\ \tilde{X}_i^{t+1} & i \in \mathcal{S}_2 \end{cases}, \quad X_i^{(3)} = \begin{cases} X_i^{(2)} & i \notin \mathcal{S}_3 \\ \tilde{X}_i^{t+1} & i \in \mathcal{S}_3 \end{cases}, \\ X_i^{(4)} &= \begin{cases} X_i^{(3)} & i \notin \mathcal{S}_2 \cup \mathcal{S}_3 \\ X_i^{t+1} & i \in \mathcal{S}_2 \cup \mathcal{S}_3 \end{cases}, \quad X_i^{(5)} = \begin{cases} X_i^{(4)} & i \notin \mathcal{S}^c \\ X_i^{t+1} & i \in \mathcal{S}^c \end{cases}, \quad \mathbf{X}^{(6)} = \mathbf{X}^{t+1}. \end{aligned}$$

Exchanging large increasing coordinates with population gradient

Let $i \in \mathcal{S}_2$ with $x_i^* > 0$. Analogous bounds for the case $x_i^* < 0$ can be established the same way. Since we are considering an increasing coordinate, the population gradient

$$\nabla f(\mathbf{X}^t)_i = -(3\|\mathbf{X}^t\|_2^2 - 1)(R_t x_i^* - X_i^t)$$

must be negative, so we must have $X_i^t \leq R_t x_i^*$. The bound (B.15) becomes

$$\tilde{X}_i^{t+1} - X_i^t \leq 1.6\eta(3\|\mathbf{X}^t\|_2^2 - 1)(R_t x_i^* - X_i^t)X_i^t,$$

and we can bound

$$\tilde{X}_i^{t+1} \leq \left(1 + 1.6\eta(3\|\mathbf{X}^t\|_2^2 - 1)(R_t x_i^* - X_i^t)\right)X_i^t \leq R_t x_i^*$$

if $\eta \leq \frac{1}{8\sqrt{2}} \leq \frac{1}{1.6(3\|\mathbf{X}^t\|_2^2 - 1)X_i^t}$, as the expression is increasing in η . We can bound the ratio

$$\frac{2((\mathbf{X}^{(2)})^\top \mathbf{x}^*)}{3\|\mathbf{X}^{(2)}\|_2^2 - 1} = \frac{2((\mathbf{X}^t)^\top \mathbf{x}^*) + \sum_{i \in \mathcal{S}_2} 2x_i^*(\tilde{X}_i^{t+1} - X_i^t)}{3\|\mathbf{X}^t\|_2^2 - 1 + \sum_{i \in \mathcal{S}_2} 3(\tilde{X}_i^{t+1} + X_i^t)(\tilde{X}_i^{t+1} - X_i^t)} \geq \frac{2((\mathbf{X}^t)^\top \mathbf{x}^*) + B}{3\|\mathbf{X}^t\|_2^2 - 1 + 3R_t B},$$

where we write $B = \sum_{i \in \mathcal{S}_2} 2x_i^*(\tilde{X}_i^{t+1} - X_i^t)$. This ratio is bounded from below by $1/\sqrt{3}$ if

$$B \leq \frac{2((\mathbf{X}^t)^\top \mathbf{x}^*) - (3\|\mathbf{X}^t\|_2^2 - 1)/\sqrt{3}}{\sqrt{3}R_t - 1} = \frac{3\|\mathbf{X}^t\|_2^2 - 1}{\sqrt{3}},$$

which is satisfied as, using $|X_i^t| \leq R_t |x_i^*|$, we can bound

$$\begin{aligned} B &= \sum_{i \in \mathcal{S}_2} 2x_i^*(\tilde{X}_i^{t+1} - X_i^t) \\ &\leq \sum_{i \in \mathcal{S}_2} 2x_i^* \cdot 1.6\eta(3\|\mathbf{X}^t\|_2^2 - 1)(R_t x_i^* - X_i^t)|X_i^t| \\ &\leq (3\|\mathbf{X}^t\|_2^2 - 1) \cdot 3.2\eta R_t^2 \sum_{i \in \mathcal{S}_2} |x_i^*|^3 \\ &\leq \frac{3\|\mathbf{X}^t\|_2^2 - 1}{\sqrt{3}} \end{aligned}$$

for $\eta \leq 0.45 \leq \frac{1}{3.2\sqrt{3}R_t^2}$, where we used that $\sum_{i \in \mathcal{S}_2} |x_i^*|^3 \leq 1$ and $R_t \leq 1/\sqrt{3} + 0.05$.

Exchanging large decreasing coordinates with population gradient

Let $i \in \mathcal{S}_3$ with $x_i^* > 0$. Analogous bounds for the case $x_i^* < 0$ can be established the same way. Following the same steps as before, we can bound by how much X_i^t can decrease after one step of mirror descent using the population gradient, and the bound (B.17) becomes

$$X_i^{(3)} - X_i^t \geq -3.2\eta((\mathbf{X}^t)^\top \mathbf{x}^*) \left(\frac{1}{R_t} X_i^t - x_i^* \right) X_i^t.$$

As in the previous step, rearranging this inequality shows that $X_i^{(3)} \geq R_t x_i^*$ provided the step size satisfies $\eta \leq 0.39 \leq \frac{R_t}{3.2((\mathbf{X}^t)^\top \mathbf{x}^*) X_i^t}$. We will show

$$3\|\mathbf{X}^{(3)}\|_2^2 - 1 \leq \sqrt{3} \cdot 2((\mathbf{X}^{(3)})^\top \mathbf{x}^*) - \Delta,$$

where $\Delta > 0$ is a buffer term due to very large coordinates decreasing. To this end, define the set $\mathcal{S}'_3 = \{i \in \mathcal{S}_3 : |X_i^t| \geq \frac{2}{\sqrt{3}+0.15} |x_i^*|\}$. Then, since $R_t \leq 1/\sqrt{3} + 0.05$, we have

$$\begin{aligned} & \frac{1}{1/\sqrt{3} + 0.05} \cdot 2((\mathbf{X}^t)^\top \mathbf{x}^*) \leq 3\|\mathbf{X}^t\|_2^2 - 1 \\ \Rightarrow & 3 \sum_{i \in \mathcal{S}'_3} (X_i^t)^2 - \frac{2}{1/\sqrt{3} + 0.05} X_i^t x_i^* \geq 1 + \sum_{i \notin \mathcal{S}'_3} \frac{2}{1/\sqrt{3} + 0.05} X_i^t x_i^* - 3(X_i^t)^2 \geq 1 - 3\delta^2, \end{aligned}$$

since $\frac{2}{1/\sqrt{3}+0.05} X_i^t x_i^* - 3(X_i^t)^2 \geq 0$ if $i \notin \mathcal{S}'_3 \cup \mathcal{S}^c$ and $\|\mathbf{X}_{\mathcal{S}^c}^t\|_2^2 \leq \delta^2$. This implies

$$\sum_{i \in \mathcal{S}'_3} (X_i^t)^2 \geq \frac{1}{3}.$$

Next, we bound by how much coordinates in \mathcal{S}'_3 must at least decrease. Let $i \in \mathcal{S}'_3$ with $x_i^* > 0$. We have

$$\begin{aligned} \nabla f(\mathbf{X}^t)_i &= 2((\mathbf{X}^t)^\top \mathbf{x}^*) \left(\frac{1}{R_t} X_i^t - x_i^* \right) \\ &\geq 2((\mathbf{X}^t)^\top \mathbf{x}^*) \left(\frac{2}{(\sqrt{3} + 0.15)(\frac{1}{\sqrt{3}} + 0.05)} - 1 \right) x_i^* \\ &\geq 2((\mathbf{X}^t)^\top \mathbf{x}^*) \frac{2c_\star}{3\sqrt{k}}. \end{aligned}$$

With this, we can bound, using (A.5),

$$\tilde{X}_i^{t+1} - X_i^t \leq -\frac{1}{2}\eta \nabla f(\mathbf{X}^t)_i X_i^t \leq -\eta \cdot 2((\mathbf{X}^t)^\top \mathbf{x}^*) \frac{c_\star}{3\sqrt{k}} X_i^t.$$

Following the same steps as in Case 1, we can show that $\tilde{X}_i^{t+1}, X_i^t \geq R_t x_i^* \geq x_i^*/\sqrt{3}$ provided the step size satisfies $\eta \leq \frac{1}{6.4\sqrt{3}}$, where we used that $R_t \geq 1/\sqrt{3}$. With this, we can bound

$$\begin{aligned} 3\|\mathbf{X}^{(3)}\|_2^2 - 1 &= 3\|\mathbf{X}^{(2)}\|_2^2 - 1 + \sum_{i \in \mathcal{S}_3 \setminus \mathcal{S}'_3} 3(\tilde{X}_i^{t+1} + X_i^t)(\tilde{X}_i^{t+1} - X_i^t) + \sum_{i \in \mathcal{S}'_3} 3(\tilde{X}_i^{t+1} + X_i^t)(\tilde{X}_i^{t+1} - X_i^t) \\ &\leq \sqrt{3} \left(2((\mathbf{X}^{(2)})^\top \mathbf{x}^*) + \sum_{i \in \mathcal{S}_3 \setminus \mathcal{S}'_3} 2x_i^*(\tilde{X}_i^{t+1} - X_i^t) \right) + \sum_{i \in \mathcal{S}'_3} 3(R_t + 1)X_i^t(\tilde{X}_i^{t+1} - X_i^t) \\ &\leq \sqrt{3} \cdot 2((\mathbf{X}^{(3)})^\top \mathbf{x}^*) + \sum_{i \in \mathcal{S}'_3} \left(\sqrt{3} + 3 - (\sqrt{3} + 0.15) \right) X_i^t(\tilde{X}_i^{t+1} - X_i^t), \end{aligned}$$

where in the last line we used that $2|x_i^*| \leq (\sqrt{3} + 0.15)|X_i^t|$ for $i \in \mathcal{S}'_3$. Finally, we have

$$\begin{aligned} \sum_{i \in \mathcal{S}'_3} (3 - 0.15)X_i^t(\tilde{X}_i^{t+1} - X_i^t) &\leq - \sum_{i \in \mathcal{S}'_3} 2.85X_i^t \eta \cdot 2((\mathbf{X}^t)^\top \mathbf{x}^*) \frac{c_\star}{3\sqrt{k}} X_i^t \\ &\leq -2.85\eta \cdot 2((\mathbf{X}^t)^\top \mathbf{x}^*) \frac{c_\star}{3\sqrt{k}} \cdot \frac{1}{3} \\ &=: -\Delta, \end{aligned}$$

that is

$$3\|\mathbf{X}^{(3)}\|_2^2 - 1 \leq \sqrt{3} \cdot 2((\mathbf{X}^{(3)})^\top \mathbf{x}^*) - \Delta. \quad (\text{B.18})$$

Exchanging large coordinates with empirical gradient

Let $i \in \mathcal{S}_2 \cup \mathcal{S}_3$ with $x_i^* > 0$. Analogous bounds for the case $x_i^* < 0$ can be established the same way.

In the following, we write $a = \frac{1}{2}(X_i^t + \sqrt{(X_i^t)^2 + \beta^2})$, $b = \frac{1}{2}(-X_i^t + \sqrt{(X_i^t)^2 + \beta^2})$, $G = \nabla F(\mathbf{X}^t)_i$ and $g = \nabla f(\mathbf{X}^t)_i$ for notational brevity. We can bound the ratio

$$\begin{aligned} \frac{X_i^{t+1}}{\tilde{X}_i^{t+1}} &= \frac{ae^{-\eta G} - be^{\eta G}}{ae^{-\eta g} - be^{\eta g}} \\ &\leq \frac{ae^{-\eta G}}{ae^{-\eta g} - be^{\eta g}} \\ &= \frac{ae^{-\eta G}}{ae^{-\eta g}} + \frac{ae^{-\eta G}be^{\eta g}}{ae^{-\eta g}(ae^{-\eta g} - be^{\eta g})} \\ &\leq \exp\left(\eta \frac{0.01c_\star((\mathbf{X}^t)^\top \mathbf{x}^*)}{\sqrt{k}}\right) + \frac{3\sqrt{k}}{c_\star}\beta \\ &\leq 1 + 0.02c_\star\eta \frac{((\mathbf{X}^t)^\top \mathbf{x}^*)}{\sqrt{k}}, \end{aligned}$$

where for the penultimate line we used the inequality $-x + \sqrt{x^2 + \beta^2} \leq \beta$ for $x > 0$, the fact that $|X_i^t| \geq |x_i^*|/3 \geq c_*/(3\sqrt{k})$ for $i \in \mathcal{S}_2 \cup \mathcal{S}_3$, and that, as in (A.2), we can use Lemma 6.1 to bound, with probability $1 - c_p n^{-10}$,

$$|\nabla F(\mathbf{X}^t)_i - \nabla f(\mathbf{X}^t)_i| \leq \frac{0.01c_*((\mathbf{X}^t)^\top \mathbf{x}^*)}{\sqrt{k}}.$$

Similarly, we can also bound $X_i^{t+1}/\tilde{X}_i^{t+1} \geq 1 - 0.02c_*\eta \frac{((\mathbf{X}^t)^\top \mathbf{x}^*)}{\sqrt{k}}$. With this, we have the bounds

$$\|\mathbf{X}^{(4)}\|_2^2 \leq \left(1 + 0.02c_*\eta \frac{((\mathbf{X}^t)^\top \mathbf{x}^*)}{\sqrt{k}}\right)^2 \|\mathbf{X}^{(3)}\|_2^2 \leq \left(1 + 2.1 \cdot 0.02c_*\eta \frac{((\mathbf{X}^t)^\top \mathbf{x}^*)}{\sqrt{k}}\right) \|\mathbf{X}^{(3)}\|_2^2,$$

provided that $\eta \leq 5/(\sqrt{2}c_*)$, and similarly

$$((\mathbf{X}^{(4)})^\top \mathbf{x}^*) \geq \left(1 - 0.02c_*\eta \frac{((\mathbf{X}^t)^\top \mathbf{x}^*)}{\sqrt{k}}\right) ((\mathbf{X}^{(3)})^\top \mathbf{x}^*).$$

Recalling the definition of Δ and using (B.18), we can now bound

$$\begin{aligned} 3\|\mathbf{X}^{(4)}\|_2^2 - 1 &\leq 3\|\mathbf{X}^{(3)}\|_2^2 - 1 + 3\|\mathbf{X}^{(3)}\|_2^2 \cdot 2.1 \cdot 0.02c_*\eta \frac{((\mathbf{X}^t)^\top \mathbf{x}^*)}{\sqrt{k}} \\ &\leq 3\|\mathbf{X}^{(3)}\|_2^2 - 1 + \frac{\Delta}{2} \\ &\leq \sqrt{3} \cdot 2((\mathbf{X}^{(3)})^\top \mathbf{x}^*) - \frac{\Delta}{2} \\ &\leq \sqrt{3} \cdot 2((\mathbf{X}^{(4)})^\top \mathbf{x}^*) + \sqrt{3} \cdot 2((\mathbf{X}^{(3)})^\top \mathbf{x}^*) \cdot 0.02c_*\eta \frac{((\mathbf{X}^t)^\top \mathbf{x}^*)}{\sqrt{k}} - \frac{\Delta}{2} \\ &\leq \sqrt{3} \cdot 2((\mathbf{X}^{(4)})^\top \mathbf{x}^*) - \frac{\Delta}{4}, \end{aligned}$$

where the last inequality holds because we can show $\|\mathbf{X}^{(3)}\|_2^2 \leq 2$ the same way as in the proof of the upper bound of (6.3).

Exchanging off-support coordinates

We have $\|\mathbf{X}^{(5)}\|_2^2 - \|\mathbf{X}^{(4)}\|_2^2 \leq n\delta^2 \leq \Delta/4$ for $\delta \leq c_1/n$ small enough, so

$$3\|\mathbf{X}^{(5)}\|_2^2 - 1 \leq \sqrt{3} \cdot 2((\mathbf{X}^{(5)})^\top \mathbf{x}^*).$$

Exchanging small coordinates

Following the same steps for small coordinates in the previous case, we can show that also

$$3\|\mathbf{X}^{(6)}\|_2^2 - 1 \leq \sqrt{3} \cdot 2((\mathbf{X}^{(6)})^\top \mathbf{x}^*),$$

which completes the proof that $R_{t+1} \geq 1/\sqrt{3}$ if $1/\sqrt{3} \leq R_t < 1/\sqrt{3} + 0.05$. \square

C. Technical lemmas

In this section, we collect technical lemmas and concentration bounds used in the proofs of Theorems 4.3 and 4.4 and the supporting Lemmas 6.1–6.3.

THEOREM C.1. (Proposition 34 [8]) Let $g : \mathbb{R}^n \rightarrow \mathbb{R}$ be a Lipschitz continuous function with Lipschitz constant λ , i.e. $|g(\mathbf{x}) - g(\mathbf{y})| \leq \lambda \|\mathbf{x} - \mathbf{y}\|_2$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$. Let $\mathbf{A} \in \mathbb{R}^n$ be a standard normal random vector. Then, for any $\epsilon > 0$, we have

$$\mathbb{P}[|g(\mathbf{A}) - \mathbb{E}[g(\mathbf{A})]| \geq \epsilon] \leq 2 \exp\left(-\frac{\epsilon^2}{2\lambda^2}\right).$$

THEOREM C.2. (Theorems 3.6, 3.7 ([3])) Let X_i be independent random variables satisfying $|X_i| \leq \lambda$ for all $i \in [n]$. Let $X = \sum_{i=1}^n X_i$ and $\|X\| = \sqrt{\sum_{i=1}^n \mathbb{E}[X_i^2]}$. Then, for any $\epsilon > 0$, we have

$$\mathbb{P}[|X - \mathbb{E}[X]| \geq \epsilon] \leq \exp\left(-\frac{\epsilon^2}{2(\|X\|^2 + \lambda\epsilon/3)}\right).$$

We state the following Lemma from [2] without proof. While the first and last inequality were not shown in [2], it can be done the same way as in the proof of Lemma A.5 in [2]. Convexity follows from the convexity of the operator norm.

LEMMA C.3. (Lemma A.5 [2]) Let $\{\mathbf{A}_j\}_{j=1}^m$ be a collection of i.i.d. $\mathcal{N}(0, \mathbf{I}_k)$ random vectors. For any $t > 0$, let $\mathcal{A} \subseteq \mathbb{R}^{m \times k}$ be the set consisting of all $\{\mathbf{a}_j\}_{j=1}^m \in \mathbb{R}^{m \times k}$ satisfying

$$\|\mathbf{a}\|_{2 \rightarrow p} \leq (p! \, m)^{\frac{1}{p}} + \sqrt{k} + t, \quad (\text{C.1})$$

for $p = 2, 4, 6, 8$, where we write $p!$ for the double factorial (i.e. the product of all numbers between 1 and p with the same parity as p) and

$$\|\mathbf{a}\|_{2 \rightarrow p} = \sup_{\|\mathbf{x}\|_2 \leq 1} \|\mathbf{a}\mathbf{x}\|_p.$$

Then, we have $P[\{\mathbf{A}_j\}_{j=1}^m \in \mathcal{A}] \geq 1 - 4 \exp(-t^2/2)$. Further, the set \mathcal{A} is convex.

The following lemma is a slight modification of the previous result.

LEMMA C.4. Let $\{\mathbf{A}_j\}_{j=1}^m$ be a collection of i.i.d. $\mathcal{N}(0, \mathbf{I}_k)$ random vectors. For any $t > 0$, let $\mathcal{A} \subseteq \mathbb{R}^{m \times k}$ be the set consisting of all $\{\mathbf{a}_j\}_{j=1}^m \in \mathbb{R}^{m \times k}$ satisfying the following: for all $\mathbf{x} \in \mathbb{R}^k$ with $\|\mathbf{x}\|_2 = 1$ and $p = 4, 6, 8$,

$$\left(\sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x})^4\right)^{\frac{1}{p}} \leq (p! \, m)^{\frac{1}{p}} + \sqrt{8 \log(2k)} \|\mathbf{x}\|_1 + t. \quad (\text{C.2})$$

Then, we have $\mathbb{P}[\{\mathbf{A}_j\}_{j=1}^m \in \mathcal{A}] \geq 1 - 3 \lceil \log_2 \sqrt{k} \rceil \exp(-t^2/2)$. Further, the set \mathcal{A} is convex.

LEMMA C.5. Let $\mathbf{x}^* \in \mathbb{R}^n$ be a k -sparse vector with $\|\mathbf{x}^*\|_2 = 1$, $\{\mathbf{A}_j\}_{j=1}^m$ a collection of i.i.d. $\mathcal{N}(0, \mathbf{I}_n)$ random vectors and $\{\varepsilon_j\}_{j=1}^m$ a collection of independent centered sub-exponential random variables with maximum sub-exponential norm $\sigma = \max_j \|\varepsilon_j\|_{\psi_1}$. There exist universal constants $c, c_s, c_p > 0$ such

that if $m \geq c_s \max\{k^2 \log^2 n, \log^5 n\}$, then, with probability at least $1 - c_p n^{-13}$,

$$\left| \frac{1}{m} \sum_{j=1}^m A_{ji}^2 A_{js} A_{jl} - \mathbb{E}[A_{1i}^2 A_{1s} A_{1l}] \right| \leq c \sqrt{\frac{\log n}{m}} \quad \text{for all } i, l, s \in [n], \quad (\text{C.3})$$

$$\left| \frac{1}{m} \sum_{j=1}^m A_{ji}^2 A_{jl} (\mathbf{A}_{j,-i}^\top \mathbf{x}_{-i}^*) - \mathbb{E}[A_{1i}^2 A_{1l} (\mathbf{A}_{1,-i}^\top \mathbf{x}_{-i}^*)] \right| \leq c \sqrt{\frac{\log n}{m}} \quad \text{for all } i, l \in [n], \quad (\text{C.4})$$

$$\left| \frac{1}{m} \sum_{j=1}^m A_{ji} A_{jl} (\mathbf{A}_{j,-i}^\top \mathbf{x}_{-i}^*)^2 - \mathbb{E}[A_{1i} A_{1l} (\mathbf{A}_{1,-i}^\top \mathbf{x}_{-i}^*)^2] \right| \leq c \sqrt{\frac{\log n}{m}} \quad \text{for all } i, l \in [n], \quad (\text{C.5})$$

$$\left| \frac{1}{m} \sum_{j=1}^m A_{ji}^8 - 105 \right| \leq c \sqrt{\frac{\log^5 n}{m}} \quad \text{for all } i \in [n], \quad (\text{C.6})$$

$$\left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j A_{ji} A_{jl} \right| \leq c \sigma \sqrt{\frac{\log n}{m}} \quad \text{for all } i, l \in [n], \quad (\text{C.7})$$

$$\left| \frac{1}{m} \sum_{j=1}^m \varepsilon_j A_{jl} (\mathbf{A}_j^\top \mathbf{x}^*) \right| \leq c \sigma \sqrt{\frac{\log n}{m}} \quad \text{for all } l \in [n]. \quad (\text{C.8})$$

LEMMA C.6. Let $\mathbf{x}^* \in \mathbb{R}^n$ be a k -sparse vector with $\|\mathbf{x}^*\|_2 = 1$, and let $\{\mathbf{A}_j\}_{j=1}^m$ be a collection of i.i.d. $\mathcal{N}(0, \mathbf{I}_n)$ random vectors. There exist universal constants $c, c_s, c_p > 0$ such that the following holds. Let $\mathcal{A} \subseteq \mathbb{R}^{m \times n}$ be the set consisting of all $\{\mathbf{a}_j\}_{j=1}^m \in \mathbb{R}^{m \times n}$ satisfying the following:

$$\begin{aligned} \left| \frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x}^*)^4 - 3 \right| &\leq c \sqrt{\frac{\log n}{m}} \\ \left| \frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x}^*)^6 - 15 \right| &\leq c \sqrt{\frac{\log^3 n}{m}} \\ \left| \frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x}^*)^8 - 105 \right| &\leq c \sqrt{\frac{\log^5 n}{m}}. \end{aligned}$$

Then, if $m \geq c_s \log^5 n$, we have $\mathbb{P}[\{\mathbf{A}_j\}_{j=1}^m \in \mathcal{A}] \geq 1 - c_p n^{-12}$. Further, the set \mathcal{A} is convex.

LEMMA C.7. Let $\{\mathbf{A}_j\}_{j=1}^m$ be a collection of i.i.d. $\mathcal{N}(0, \mathbf{I}_k)$ random vectors. There exist universal constants $c, c_s, c_p > 0$ such that if $m \geq c_s \max\{k^2 \log^2 n, \log^5 n\}$, where $n \geq k$ is any natural number, then the set $\mathcal{A} \subseteq \mathbb{R}^{m \times k}$ defined by

$$\mathcal{A} = \left\{ \{\mathbf{a}_j\}_{j=1}^m \in \mathbb{R}^{m \times k} : \frac{1}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} a_{jl} \right)^4 \leq c |\mathcal{L}|^2 \text{ for all } \mathcal{L} \subseteq [k] \right\}$$

satisfies $\mathbb{P}[\{\mathbf{A}_j\}_{j=1}^m \in \mathcal{A}] \geq 1 - c_p n^{-11}$. Further, the set \mathcal{A} is convex.

LEMMA C.8. Let $\mathbf{x}^* \in \mathbb{R}^n$ be a k -sparse vector, and let $\mathcal{S} = \{i \in [n] : x_i^* \neq 0\}$ be its support. Let $\{\mathbf{A}_j\}_{j=1}^m$ be a collection of i.i.d. $\mathcal{N}(0, \mathbf{I}_n)$ random variables and let $\{\varepsilon_j\}_{j=1}^m$ be a collection of independent centered sub-exponential random variables with maximum sub-exponential norm $\sigma = \max_j \|\varepsilon_j\|_{\psi_1}$. There exist universal constants $c_s, c_p > 0$ such that, for any constant $c_1 > 0$, there is a constant $c > 0$ such that if $m \geq c_s \max\{k^2 \log n, \log^5 n\}$, then, with probability at least $1 - c_p n^{-10}$, the following holds: for any $i \in [n]$, $\mathcal{K} \subseteq \mathcal{S} \setminus \{i\}$ and $\epsilon \geq c_1/m$, let $N_\epsilon^\mathcal{K}$ be a smallest ϵ -net of $\{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}_{\mathcal{K}^c} = \mathbf{0}, \|\mathbf{x}\|_2 = 1\}$. Then, for any $\mathcal{L} \subseteq \mathcal{K}$ with $|\mathcal{L}| \geq \frac{1}{2}|\mathcal{K}|$, and $\mathbf{x} \in N_\epsilon^\mathcal{K}$,

$$\left| \sum_{l \in \mathcal{L}} \frac{1}{m} \sum_{j=1}^m A_{ji} A_{jl} (\mathbf{A}_j^\top \mathbf{x})^2 \right| \leq c |\mathcal{L}| \frac{\log n}{\sqrt{m}}, \quad (\text{C.9})$$

$$\left| \sum_{l \in \mathcal{L}} \frac{1}{m} \sum_{j=1}^m A_{ji}^2 A_{jl} (\mathbf{A}_{j,-l}^\top \mathbf{x}_{-l}) \right| \leq c |\mathcal{L}| \sqrt{\frac{\log n}{m}}, \quad (\text{C.10})$$

$$\left| \sum_{l \in \mathcal{L}} \frac{1}{m} \sum_{j=1}^m \varepsilon_j A_{jl} (\mathbf{A}_{j,-l}^\top \mathbf{x}_{-l}) \right| \leq c \sigma |\mathcal{L}| \sqrt{\frac{\log n}{m}}, \quad (\text{C.11})$$

$$\left| \sum_{l \in \mathcal{L}} \frac{1}{m} \sum_{j=1}^m A_{ji} A_{jl} (\mathbf{A}_j^\top \mathbf{x}) (\mathbf{A}_{j,-i}^\top \mathbf{x}_{-i}^*) \right| \leq c |\mathcal{L}| \frac{\log n}{\sqrt{m}}. \quad (\text{C.12})$$

LEMMA C.9. Let $\mathbf{x}^* \in \mathbb{R}^n$ be a k -sparse vector, and let $\mathcal{S} = \{i \in [n] : x_i^* \neq 0\}$ be its support. Let $\{\mathbf{A}_j\}_{j=1}^m$ be a collection of i.i.d. $\mathcal{N}(0, \mathbf{I}_n)$ random vectors and let $\{\varepsilon_j\}_{j=1}^m$ be a collection of independent centered sub-exponential random variables with maximum sub-exponential norm $\sigma = \max_j \|\varepsilon_j\|_{\psi_1}$. There exist universal constants $c, c_s, c_p > 0$ such that if $m \geq c_s \max\{k^2 \log^2 n, \log^5 n\}$, then the following holds with probability at least $1 - c_p n^{-10}$: For all $i \in [n]$ and $\mathbf{x} \in \mathcal{X} = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}_{\mathcal{S}^c} = \mathbf{0}, \|\mathbf{x}\|_2 \leq 1\}$,

$$\left| \frac{1}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,-i}^\top \mathbf{x}_{-i})^3 \right| \leq c \|\mathbf{x}\|_1 \frac{\log n}{\sqrt{m}}, \quad (\text{C.13})$$

$$\left| \sum_{l \neq i} x_l \sum_{r \neq l, i} x_r \frac{1}{m} \sum_{j=1}^m A_{ji}^2 A_{jl} A_{jr} \right| \leq c \|\mathbf{x}\|_1 \sqrt{\frac{\log n}{m}}, \quad (\text{C.14})$$

$$\left| \sum_{l \neq i} x_l \sum_{r \neq l, i} x_r \frac{1}{m} \sum_{j=1}^m \varepsilon_j A_{jl} A_{jr} \right| \leq c \sigma \|\mathbf{x}\|_1 \sqrt{\frac{\log n}{m}}, \quad (\text{C.15})$$

$$\left| \frac{1}{m} \sum_{j=1}^m A_{ji} (\mathbf{A}_{j,-i}^\top \mathbf{x}_{-i})^2 (\mathbf{A}_{j,-i}^\top \mathbf{x}_{-i}^*) \right| \leq c \|\mathbf{x}\|_1 \frac{\log n}{\sqrt{m}}. \quad (\text{C.16})$$

In the following, we prove Lemmas C.4–C.9.

Proof of Lemma C.4. The proof follows that of Lemma A.5 of [2] closely. Since inequality (C.2) can be shown the same way for $p = 4, 6, 8$, we only present the proof for the case $p = 8$. For any fixed $\lambda \geq 1$, define

$$\|\mathbf{A}\|_{2 \rightarrow 8, \lambda} = \max\{\|\mathbf{A}\mathbf{x}\|_8 : \|\mathbf{x}\|_2 = 1, \|\mathbf{x}\|_1 \leq \lambda\}.$$

We will show that $\|\mathbf{A}\|_{2 \rightarrow 8, \lambda} \leq (105m)^{\frac{1}{8}} + \sqrt{2 \log(2k)\lambda} + t$ with probability $1 - 2 \exp(-t^2/2)$. The result then follows by taking the union bound over $\lambda = 2^1, 2^2, \dots, 2^{\lceil \log_2 \sqrt{k} \rceil}$ for inequality (C.2) with each $p = 4, 6, 8$, since we have $1 \leq \|\mathbf{x}\|_1 \leq \sqrt{k}$ and the union bound is chosen such that there always exists a λ with $\|\mathbf{x}\|_1 \leq \lambda \leq 2\|\mathbf{x}\|_1$. Define $X_{\mathbf{u}, \mathbf{v}} = \langle \mathbf{A}\mathbf{u}, \mathbf{v} \rangle$ on the set

$$T_\lambda = \{(\mathbf{u}, \mathbf{v}) : \mathbf{u} \in \mathbb{R}^k, \|\mathbf{u}\|_2 = 1, \|\mathbf{u}\|_1 \leq \lambda, \mathbf{v} \in \mathbb{R}^m, \|\mathbf{v}\|_{8/7} \leq 1\}.$$

Then, by Hölder's inequality, we have $\|\mathbf{A}\|_{2 \rightarrow 8, \lambda} = \max_{(\mathbf{u}, \mathbf{v}) \in T_\lambda} X_{\mathbf{u}, \mathbf{v}}$.

Define $Y_{\mathbf{u}, \mathbf{v}} = \langle \mathbf{G}, \mathbf{u} \rangle + \langle \mathbf{H}, \mathbf{v} \rangle$, where $\mathbf{G} \in \mathbb{R}^k$ and $\mathbf{H} \in \mathbb{R}^m$ are independent standard normal random vectors. We have, for any $(\mathbf{u}, \mathbf{v}), (\mathbf{u}', \mathbf{v}') \in T_\lambda$,

$$\mathbb{E}[|X_{\mathbf{u}, \mathbf{v}} - X_{\mathbf{u}', \mathbf{v}'}|^2] = \|\mathbf{v}\|_2^2 + \|\mathbf{v}'\|_2^2 - 2\langle \mathbf{u}, \mathbf{u}' \rangle \langle \mathbf{v}, \mathbf{v}' \rangle,$$

and

$$\mathbb{E}[|Y_{\mathbf{u}, \mathbf{v}} - Y_{\mathbf{u}', \mathbf{v}'}|^2] = 2 + \|\mathbf{v}\|_2^2 + \|\mathbf{v}'\|_2^2 - 2\langle \mathbf{u}, \mathbf{u}' \rangle - 2\langle \mathbf{v}, \mathbf{v}' \rangle,$$

where we used $\|\mathbf{u}\|_2 = \|\mathbf{u}'\|_2 = 1$. Therefore, we have

$$\mathbb{E}[|Y_{\mathbf{u}, \mathbf{v}} - Y_{\mathbf{u}', \mathbf{v}'}|^2] - \mathbb{E}[|X_{\mathbf{u}, \mathbf{v}} - X_{\mathbf{u}', \mathbf{v}'}|^2] = 2(1 - \langle \mathbf{u}, \mathbf{u}' \rangle)(1 - \langle \mathbf{v}, \mathbf{v}' \rangle) \geq 0,$$

where we used $\|\mathbf{v}\|_2 \leq \|\mathbf{v}\|_{8/7} \leq 1$ and $\|\mathbf{v}'\|_2 \leq \|\mathbf{v}'\|_{8/7} \leq 1$. By Proposition 33 of [8], this implies

$$\mathbb{E}\left[\max_{(\mathbf{u}, \mathbf{v}) \in T_\lambda} X_{\mathbf{u}, \mathbf{v}}\right] \leq \mathbb{E}\left[\max_{(\mathbf{u}, \mathbf{v}) \in T_\lambda} Y_{\mathbf{u}, \mathbf{v}}\right],$$

and hence,

$$\mathbb{E}[\|\mathbf{A}\|_{2 \rightarrow 8, \lambda}] \leq \mathbb{E}\left[\max_{(\mathbf{u}, \mathbf{v}) \in T_\lambda} Y_{\mathbf{u}, \mathbf{v}}\right] \leq \mathbb{E}\left[\max_{(\mathbf{u}, \mathbf{v}) \in T_\lambda} \|\mathbf{G}\|_\infty \|\mathbf{u}\|_1 + \|\mathbf{H}\|_8 \|\mathbf{v}\|_{8/7}\right] \leq \sqrt{2 \log(2k)\lambda} + (105m)^{\frac{1}{8}},$$

where we used Hölder's inequality in the second line, and for the last inequality the bound

$$\exp(t\mathbb{E}[\|\mathbf{G}\|_\infty]) \leq \mathbb{E}[\exp(t\|\mathbf{G}\|_\infty)] \leq \sum_{i=1}^k \mathbb{E}[\exp(t|G_i|)] \leq 2k \exp(t^2/2)$$

with $t = \sqrt{2 \log(2k)}$. Finally, $\|\cdot\|_{2 \rightarrow 8, \lambda}$ is a 1-Lipschitz function: let $\mathbf{a}, \mathbf{b} \in \mathbb{R}^{m \times k}$ and, without loss of generality, $\|\mathbf{a}\|_{2 \rightarrow 8, \lambda} \geq \|\mathbf{b}\|_{2 \rightarrow 8, \lambda}$. Then,

$$\begin{aligned} \|\mathbf{a}\|_{2 \rightarrow 8, \lambda} - \|\mathbf{b}\|_{2 \rightarrow 8, \lambda} &= \max_{\|\mathbf{x}\|_2=1, \|\mathbf{x}\|_1 \leq \lambda} \|\mathbf{a}\mathbf{x}\|_8 - \max_{\|\mathbf{y}\|_2=1, \|\mathbf{y}\|_1 \leq \lambda} \|\mathbf{b}\mathbf{y}\|_8 \\ &\leq \max_{\|\mathbf{x}\|_2=1, \|\mathbf{x}\|_1 \leq \lambda} \|\mathbf{a}\mathbf{x}\|_8 - \|\mathbf{b}\mathbf{x}\|_8 \\ &\leq \max_{\|\mathbf{x}\|_2=1, \|\mathbf{x}\|_1 \leq \lambda} \|(\mathbf{a} - \mathbf{b})\mathbf{x}\|_8 \\ &\leq \max_{\|\mathbf{x}\|_2=1, \|\mathbf{x}\|_1 \leq \lambda} \|(\mathbf{a} - \mathbf{b})\mathbf{x}\|_2 \\ &\leq \|\mathbf{a} - \mathbf{b}\|_F, \end{aligned}$$

where we write $\|\cdot\|_F$ for the Frobenius norm and used the fact that it is an upper bound to the ℓ_2 -operator norm. Hence, an application of Theorem C.1 yields

$$\mathbb{P}\left[\|\mathbf{A}\|_{2 \rightarrow 8, \lambda} < \sqrt{2 \log(2k)}\lambda + (105m)^{\frac{1}{8}} + t\right] \geq 1 - \exp(-t^2/2).$$

Taking the union bound over $\lambda = 2^1, \dots, 2^{\lceil \log_2 \sqrt{k} \rceil}$ completes the proof that (C.2) holds for $p = 8$ with probability $1 - \lceil \log_2 \sqrt{k} \rceil \exp(-t^2/2)$, since the union bound is chosen such that for any $\mathbf{x} \in \mathbb{R}^k$ with $\|\mathbf{x}\|_2 = 1$ there is a λ satisfying $\|\mathbf{x}\|_1 \leq \lambda \leq 2\|\mathbf{x}\|_1$. The inequalities (C.2) can be treated the same way for $p = 4, 6$, and another union bound gives the desired result.

Finally, convexity of \mathcal{A} follows from an application of the Minkowski inequality. Let $\mathbf{a}, \mathbf{b} \in \mathbb{R}^{m \times k}$ satisfy (C.2) and $\alpha \in (0, 1)$. Then, for any $\mathbf{x} \in \mathbb{R}$ with $\|\mathbf{x}\|_2 = 1$ and $p = 4, 6, 8$, we have

$$\begin{aligned} \left(\sum_{j=1}^m (\alpha \mathbf{a}^\top \mathbf{x} + (1 - \alpha) \mathbf{b}^\top \mathbf{x})^p\right)^{\frac{1}{p}} &\leq \alpha \left(\sum_{j=1}^m (\mathbf{a}^\top \mathbf{x})^p\right)^{\frac{1}{p}} + (1 - \alpha) \left(\sum_{j=1}^m (\mathbf{b}^\top \mathbf{x})^p\right)^{\frac{1}{p}} \\ &\leq (p! \cdot m)^{\frac{1}{p}} + \sqrt{8 \log(2k)} \|\mathbf{x}\|_1 + t, \end{aligned}$$

which means that $\alpha \mathbf{a} + (1 - \alpha) \mathbf{b} \in \mathcal{A}$. \square

Proof of Lemma C.5. The proof of Lemma C.5 relies on Theorem C.2 and the following truncation, which allows us to consider bounded random variables. We begin by showing the inequality (C.3).

Let $i, l, s \in [n]$. Writing $B_j = \{\max\{|A_{ji}|, |A_{jl}|, |A_{js}|\} \leq \sqrt{64 \log n}\}$, we can decompose the term $A_{ji}^2 A_{jl} A_{js} = A_{ji}^2 A_{jl} A_{js} \mathbf{1}(B_j) + A_{ji}^2 A_{jl} A_{js} (1 - \mathbf{1}(B_j)) = Y_j + Z_j$, where by $\mathbf{1}(\cdot)$ we denote the indicator function. We will show that each of the two terms concentrates around its mean.

Since $\frac{1}{m} |Y_j| \leq \frac{1}{m} 64^2 \log^2 n$ is bounded, we can apply Theorem C.2. We can bound

$$\sqrt{\sum_{j=1}^m \frac{1}{m^2} \mathbb{E}[Y_j^2]} = \sqrt{\sum_{j=1}^m \frac{1}{m^2} \mathbb{E}[A_{ji}^4 A_{jl}^2 A_{js}^2 \mathbf{1}(B_j)]} \leq \sqrt{\frac{105}{m}}.$$

Then, since $m \geq c_s \log^5 n$, we have, for $c > 0$ sufficiently large,

$$\mathbb{P}\left[\left|\frac{1}{m} \sum_{j=1}^m Y_j - \mathbb{E}[Y_j]\right| > \frac{c}{2} \sqrt{\frac{\log n}{m}}\right] \leq 2 \exp\left(-\frac{c^2 \log n / (4m)}{2(\frac{105}{m} + 64^2 c \frac{\log^{5/2} n}{6m^{3/2}})}\right) \leq 2n^{-16}.$$

For the second term Z_j , we can use the Chebyshev inequality to bound

$$\begin{aligned} \text{Var}\left(\frac{1}{m} \sum_{j=1}^m Z_j\right) &\leq \frac{1}{m} \mathbb{E}[A_{1i}^4 A_{1l}^2 A_{1s}^2 \mathbf{1}(B_j^c)] \\ &\leq \frac{1}{m} \sqrt{\mathbb{E}[A_{1i}^8 A_{1l}^4 A_{1s}^4] \mathbb{P}\left[\max\{|A_{1i}|, |A_{1l}|, |A_{1s}|\} > \sqrt{64 \log n}\right]} \\ &\leq \frac{c'}{mn^{16}}, \end{aligned}$$

for an absolute constant $c' > 0$, and hence, by the Chebyshev inequality,

$$\mathbb{P}\left[\left|\frac{1}{m} \sum_{j=1}^m Z_j - \mathbb{E}[Z_j]\right| > \frac{c}{2} \sqrt{\frac{\log n}{m}}\right] \leq \frac{\frac{c'}{mn^{16}}}{\frac{c^2 \log n}{4m}} = \frac{4c'}{c^2 \log n} n^{-16} \leq n^{-16},$$

for $c \geq \sqrt{\frac{4c'}{\log n}}$. This completes the proof that

$$\mathbb{P}\left[\left|\frac{1}{m} \sum_{j=1}^m A_{ji}^2 A_{jl} A_{js} - \mathbb{E}[A_{ji}^2 A_{jl} A_{js}]\right| > c \sqrt{\frac{\log n}{m}}\right] \leq 3n^{-16}.$$

Taking the union bound over all $i, l, s \in [n]$ shows that (C.3) holds with probability at least $1 - 3n^{-13}$.

The inequalities (C.4)–(C.8) can be shown following the same steps. To show (C.6), we need to control higher order terms, which is the reason for the additional logarithmic factor. The bounds (C.7) and (C.8) can be shown the same way, as, for sub-exponential random variables, we have the standard tail bound $\mathbb{P}[|\varepsilon| > t] \leq \exp(1 - \tilde{c}t/\sigma)$ for all $t \geq 0$, where $\tilde{c} > 0$ is an absolute constant, and the bound on the second moment $\mathbb{E}[\varepsilon^2] \leq 4\sigma^2$. As the proof of each bound follows exactly the same steps, we omit the details to avoid repetition. \square

Proof of Lemma C.6. Using the fact that $\mathbf{A}_j^\top \mathbf{x}^* \sim \mathcal{N}(0, \|\mathbf{x}^*\|_2^2)$ are independent random variables, this lemma can be shown following the same steps as in the proof of Lemma C.5. The convexity of \mathcal{A} follows, as in Lemma C.4, from the Minkowski inequality. \square

Proof of Lemma C.7. We prove Lemma C.7 via induction over the size $|\mathcal{L}|$. For any $1 \leq \lambda \leq k$, define

$$\mathcal{A}_\lambda = \left\{ \{\mathbf{a}_j\}_{j=1}^m \in \mathbb{R}^{m \times k} : \frac{1}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} a_{jl} \right)^4 \leq c|\mathcal{L}|^2 \text{ for all } \mathcal{L} \subseteq [k] \text{ with } |\mathcal{L}| = \lambda \right\}.$$

We will show that $\mathbb{P}[\{\mathbf{A}_j\}_{j=1}^m \in \mathcal{A}_\lambda] \geq 1 - c_1 \lambda n^{-13}$ for some constant $c_1 > 0$, from which the result $\mathbb{P}[\{\mathbf{A}_j\}_{j=1}^m \in \mathcal{A}] \geq 1 - c_p n^{-11}$ follows by taking the union bound over all possible sizes λ . As in the proof of Lemma C.4, convexity of \mathcal{A}_λ follows from the Minkowski inequality. As the intersection of convex sets, \mathcal{A} is also convex.

The base case $\lambda = 1$ follows from the bound (C.4) of Lemma C.5. For the induction step, assume that we have already shown $\mathbb{P}[\{\mathbf{A}_j\}_{j=1}^m \in \mathcal{A}_\lambda] \geq 1 - c_1 \lambda n^{-13}$ for some $1 \leq \lambda < k$. Let $r \in [k]$ and

$\mathcal{L} \subseteq [k] \setminus \{r\}$ be a subset of coordinates with $|\mathcal{L}| = \lambda$. We have

$$\begin{aligned} \frac{1}{m} \sum_{j=1}^m \left(A_{jr} + \sum_{l \in \mathcal{L}} A_{jl} \right)^4 &= \frac{1}{m} \sum_{j=1}^m A_{jr}^4 + \frac{4}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} A_{jl} \right) A_{jr}^3 + \frac{6}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} A_{jl} \right)^2 A_{jr}^2 \\ &\quad + \frac{4}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} A_{jl} \right)^3 A_{jr} + \frac{1}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} A_{jl} \right)^4. \end{aligned} \quad (\text{C.17})$$

We need to show that, with probability $1 - c_1(\lambda + 1)n^{-13}$, the sum (C.17) is bounded by $c(|\mathcal{L}| + 1)^2$ for all $r \in [k]$ and $\mathcal{L} \subseteq [k] \setminus \{r\}$ with $|\mathcal{L}| = \lambda$.

By the induction hypothesis, we can bound the last term with probability $1 - c_1 \lambda n^{-13}$,

$$\frac{1}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} A_{jl} \right)^4 \leq c|\mathcal{L}|^2.$$

The first, second and third term in (C.17) can be bounded using Lemma C.5. The following holds with probability at least $1 - c_2 n^{-13}$, where c_2 is the universal constant from Lemma C.5. Using the assumption $m \geq c_s k^2 \log^2 n$, we can bound the first term by

$$\frac{1}{m} \sum_{j=1}^m A_{jr}^4 \leq 4,$$

the second term by

$$\frac{4}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} A_{jl} \right) A_{jr}^3 = 4 \sum_{l \in \mathcal{L}} \frac{1}{m} \sum_{j=1}^m A_{jl} A_{jr}^3 \leq 4 \left(4 + \frac{|\mathcal{L}| - 1}{k} \right) \leq 20,$$

and the third term in (C.17) can be bounded, using the Cauchy–Schwarz inequality, by

$$\frac{6}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} A_{jl} \right)^2 A_{jr}^2 \leq 6 \sqrt{\frac{1}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} A_{jl} \right)^4} \sqrt{\frac{1}{m} \sum_{j=1}^m A_{jr}^4} \leq 6\sqrt{c|\mathcal{L}|^2} \sqrt{4} \leq 0.4c|\mathcal{L}|,$$

for $c \geq 900$, where we used the induction hypothesis in the second line.

To bound the fourth term in (C.17), fix any $r \in [k]$ and subset $\mathcal{L} \subseteq [k] \setminus \{r\}$ with cardinality $|\mathcal{L}| = \lambda$, and consider the function

$$h_{r,\mathcal{L}}(\mathbf{a}) = \frac{4}{m} \sum_{j=1}^m a_{jr} \left(\sum_{l \in \mathcal{L}} a_{jl} \right)^3.$$

Then, we need to bound the probability of the event

$$B_r = \{ |h_{r,\mathcal{L}}(\mathbf{A})| \leq 1.6c|\mathcal{L}| \text{ for all } \mathcal{L} \subseteq [k] \setminus \{r\} \text{ with } |\mathcal{L}| = \lambda \}$$

for all $r \in [k]$. To this end, we condition on $\mathbf{A}_{\cdot r} = \mathbf{a}_{\cdot r}$ using the formula

$$\mathbb{P}[B] = \int \mathbb{P}[B | \mathbf{A}_{\cdot r} = \mathbf{a}_{\cdot r}] \mu(\mathbf{a}_{\cdot r}) d\mathbf{a}_{\cdot r},$$

which holds for any event B , where we write μ for the standard normal density.

To show that the conditional probability $\mathbb{P}[B_r | \mathbf{A}_{\cdot r} = \mathbf{a}_{\cdot r}]$ is close to one, we begin by showing that, for any fixed \mathcal{L} , $h_{r,\mathcal{L}}(\mathbf{A})|_{\mathbf{A}_{\cdot r}=\mathbf{a}_{\cdot r}}$ (making explicit the fact that we are conditioning on $\mathbf{A}_{\cdot r} = \mathbf{a}_{\cdot r}$) concentrates around its expectation $\mathbb{E}[h_{r,\mathcal{L}}(\mathbf{A}) | \mathbf{A}_{\cdot r} = \mathbf{a}_{\cdot r}] = 0$. For the sake of brevity, we will omit explicitly writing the condition $\mathbf{A}_{\cdot r} = \mathbf{a}_{\cdot r}$ in what follows.

Step 1: Bound the fourth term in (C.17) conditioned on $r, \mathbf{A}_{\cdot r}$ and \mathcal{L}

Fix an $r \in [k]$, a vector $\mathbf{a}_{\cdot r}$ satisfying $\max_j |a_{jr}| \leq 6\sqrt{\log n}$ and a subset $\mathcal{L} \subseteq [k] \setminus \{r\}$ with cardinality $|\mathcal{L}| = \lambda$. The idea is to apply concentration of Lipschitz functions of Gaussian random variables to show that the fourth term in (C.17), $h_{r,\mathcal{L}}(\mathbf{A})$, is close to its expectation $\mathbb{E}[h_{r,\mathcal{L}}(\mathbf{A})] = 0$. However, as $h_{r,\mathcal{L}}$ is not globally Lipschitz continuous, we cannot directly apply Theorem C.1. Similar to the proof of Theorem 3 in [4], we will first restrict $h_{r,\mathcal{L}}$ to a high probability event on which $h_{r,\mathcal{L}}$ is Lipschitz continuous. Then, we extend this restricted function to a function $\tilde{h}_{r,\mathcal{L}}$ on the entire space in a way such that $\tilde{h}_{r,\mathcal{L}}$ is globally Lipschitz continuous, and apply Theorem C.1 to $\tilde{h}_{r,\mathcal{L}}$. This also provides a high probability bound for $h_{r,\mathcal{L}}(\mathbf{A})$, since, by construction, $\tilde{h}_{r,\mathcal{L}}(\mathbf{A}) = h_{r,\mathcal{L}}(\mathbf{A})$ with high probability.

Step 1, part (a): Bound the Lipschitz constant of $h_{r,\mathcal{L}}$ restricted to \mathcal{A}_λ

Restricted to \mathcal{A}_λ and conditioned on $\mathbf{A}_{\cdot r} = \mathbf{a}_{\cdot r}$, we can bound the Lipschitz constant of $h_{r,\mathcal{L}}$ by the norm of its gradient. The norm of the gradient is an upper bound for the Lipschitz constant by the mean-value theorem, since \mathcal{A}_λ is a convex set. For any $l \in \mathcal{L}$, we have

$$\frac{\partial}{\partial a_{jl}} h_{r,\mathcal{L}}(\mathbf{a}) = \frac{12}{m} a_{jr} \left(\sum_{l' \in \mathcal{L}} a_{jl'} \right)^2.$$

Hence, we can bound

$$\|\nabla h_{r,\mathcal{L}}(\mathbf{a})\|_2^2 = \sum_{l \in \mathcal{L}} \sum_{j=1}^m \left(\frac{\partial}{\partial a_{jl}} h_{r,\mathcal{L}}(\mathbf{a}) \right)^2 \leq \frac{144|\mathcal{L}| \max_j a_{jr}^2}{m} \frac{1}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} a_{jl} \right)^4 \leq \frac{5184|\mathcal{L}| \log n}{m} c |\mathcal{L}|^2$$

on \mathcal{A}_λ , where we used the assumption $\max_j a_{jr}^2 \leq 36 \log n$ and the induction hypothesis.

Step 1, part (b): Construct a globally Lipschitz continuous extension of $h_{r,\mathcal{L}}$

Consider the following Lipschitz extension of the function $h_{r,\mathcal{L}}$:

$$\tilde{h}_{r,\mathcal{L}}(\mathbf{a}) = \inf_{\mathbf{a}' \in \mathcal{A}_\lambda} ((h_{r,\mathcal{L}}(\mathbf{a}') + \text{Lip}(h_{r,\mathcal{L}}) \|\mathbf{a} - \mathbf{a}'\|_2),$$

where we write $\text{Lip}(h_{r,\mathcal{L}}) = \sqrt{\frac{5184c|\mathcal{L}|^3 \log n}{m}}$. By definition, we have $\tilde{h}_{r,\mathcal{L}} = h_{r,\mathcal{L}}$ on \mathcal{A}_λ , and it follows from an application of the triangle inequality that $\tilde{h}_{r,\mathcal{L}}$ is globally Lipschitz continuous with Lipschitz constant $\text{Lip}(h_{r,\mathcal{L}})$ (see e.g. Theorem 7.2 of [6]).

We will show that $\tilde{h}_{r,\mathcal{L}}$ concentrates around its mean, which can potentially differ from the mean of $h_{r,\mathcal{L}}$. Since $h_{r,\mathcal{L}}$ and $\tilde{h}_{r,\mathcal{L}}$ differ only on \mathcal{A}_λ^c , which has probability less than $c_1 \lambda n^{-13}$, we can bound, using the Cauchy–Schwarz inequality,

$$\mathbb{E}[h_{r,\mathcal{L}}(\mathbf{A}) | \mathbf{1}_{\mathcal{A}_\lambda^c}(\mathbf{A})] \leq \sqrt{\mathbb{E}[h_{r,\mathcal{L}}(\mathbf{A})^2]} \sqrt{\mathbb{E}[\mathbf{1}_{\mathcal{A}_\lambda^c}(\mathbf{A})]} \leq \frac{c'}{n^5},$$

where we used that $\mathbb{E}[h_{r,\mathcal{L}}(\mathbf{A})^2] = 15|\mathcal{L}|^3/m$. Using $\tilde{h}_{r,\mathcal{L}}(\mathbf{a}) \leq \text{Lip}(h_{r,\mathcal{L}}) \|\mathbf{a}\|_2$, we have

$$\mathbb{E}[\tilde{h}_{r,\mathcal{L}}(\mathbf{A}) | \mathbf{1}_{\mathcal{A}_\lambda^c}(\mathbf{A})] \leq \text{Lip}(h_{r,\mathcal{L}}) \sqrt{\mathbb{E}[\|\mathbf{A}\|_2^2]} \sqrt{\mathbb{E}[\mathbf{1}_{\mathcal{A}_\lambda^c}(\mathbf{A})]} \leq \frac{c''}{n^4}.$$

All in all, this shows that

$$|\mathbb{E}[h_{r,\mathcal{L}}(\mathbf{A})] - \mathbb{E}[\tilde{h}_{r,\mathcal{L}}(\mathbf{A})]| \leq \frac{c_3}{n^4}$$

for a constant $c_3 > 0$. Hence, we can apply Theorem C.1 to obtain

$$\mathbb{P}[|\tilde{h}_{r,\mathcal{L}}(\mathbf{A})| > 1.6c|\mathcal{L}| \mid \mathbf{A}_{\cdot r} = \mathbf{a}_{\cdot r}] \leq 2 \exp\left(-\frac{(1.6c|\mathcal{L}| - c_3/n^4)^2}{2 \frac{5184c|\mathcal{L}|^3 \log n}{m}}\right) \leq 2 \exp\left(-c_4 \frac{m}{|\mathcal{L}| \log n}\right),$$

for a constant $c_4 \leq \frac{c}{4050} - \frac{c_3}{3240n^4}$.

Step 2: Unravel the conditions: take union bounds and integrate over $\mathbf{a}_{\cdot r}$

Let

$$\tilde{B}_r = \left\{ |\tilde{h}_{r,\mathcal{L}}(\mathbf{A})| \leq 1.6c|\mathcal{L}| \text{ for all } \mathcal{L} \subseteq [k] \setminus \{r\} \text{ with } |\mathcal{L}| = \lambda \right\}.$$

Since we assume $m \geq c_s k^2 \log^2 n$, we can take the union bound over all possible subsets $\mathcal{L} \subseteq [k] \setminus \{r\}$ with $|\mathcal{L}| = \lambda$ to obtain, using the upper bound $\binom{k}{\lambda} \leq \left(\frac{ek}{\lambda}\right)^\lambda$,

$$\mathbb{P}[\tilde{B}_r \mid \mathbf{A}_{\cdot r} = \mathbf{a}_{\cdot r}] \geq 1 - 2 \exp\left(-c_4 \frac{m}{|\mathcal{L}| \log n} + |\mathcal{L}| \log \frac{ek}{|\mathcal{L}|}\right) \geq 1 - 2 \exp(-(c_s c_4 - 1)k \log n).$$

Next, we integrate over all $\mathbf{a}_{\cdot r}$ satisfying $\max_j |a_{jr}| \leq 6\sqrt{\log n}$:

$$\begin{aligned} \mathbb{P}[\tilde{B}_r] &\geq \int_{\{\max_j |a_{jr}| \leq 6\sqrt{\log n}\}} \mathbb{P}[B_r \mid \mathbf{A}_{\cdot r} = \mathbf{a}_{\cdot r}] \mu(\mathbf{a}_{\cdot r}) d\mathbf{a}_{\cdot r} \\ &\geq (1 - 2 \exp(-(c_s c_4 - 1)k \log n)) \mathbb{P}\left[\max_j |a_{jr}| \leq 6\sqrt{\log n}\right] \\ &\geq 1 - (m + 2)n^{-18}, \end{aligned}$$

where we write μ for the standard normal density and the last line follows from standard Gaussian tail bounds and a union bound, provided that $(c_s c_4 - 1)k \geq 18$.

By construction, we have $h_{r,\mathcal{L}}(\mathbf{a}) = \tilde{h}_{r,\mathcal{L}}(\mathbf{a})$ for all $r \in [k]$ and subsets $\mathcal{L} \subseteq [k] \setminus \{r\}$ with cardinality $|\mathcal{L}| = \lambda$ on the set $\mathcal{A}_\lambda \cap \{\max_{j,r} |a_{jr}| \leq 6\sqrt{\log n}\}$. Hence, a union bound gives a lower bound on the probability that the fourth term in (C.17) is bounded by $1.6c|\mathcal{L}|$ for all $r \in [k]$ and $\mathcal{L} \subseteq [k] \setminus \{r\}$ with $|\mathcal{L}| = \lambda$:

$$\mathbb{P}\left[\bigcap_{r=1}^k B_r\right] \geq \mathbb{P}\left[\bigcap_{r=1}^k \tilde{B}_r \cap \mathcal{A}_\lambda \cap \left\{\max_{j,r} |a_{jr}| \leq 6\sqrt{\log n}\right\}\right] \geq 1 - (m + 2)n^{-17} - c_1 \lambda n^{-13}.$$

Putting everything together, this completes the induction step (C.17),

$$\frac{1}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} A_{jl} + A_{jr} \right)^4 \leq c(|\mathcal{L}| + 1)^2 \quad \text{for all } r \in [k], \mathcal{L} \subseteq [k] \setminus \{r\} \text{ with } |\mathcal{L}| = \lambda,$$

holds with probability at least $1 - c_1(\lambda + 1)n^{-13}$, provided that $c_1 \geq c_2 + (m + 2)n^{-4}$, which is satisfied for a universal constant c_1 if $m \lesssim n^4$. The case $m \gtrsim n^4$ is simpler and can be shown following the same steps, writing probabilities in terms of m instead of n . \square

Proof of Lemma C.8. We will show that each of the inequalities (C.9)–(C.12) is satisfied with probability at least $1 - \frac{c_p}{4}n^{-10}$. The lemma then follows by taking a union bound. Throughout this proof, we will assume $m \leq n^{3/2}$. The other case $m > n^{3/2}$ is simpler and can be shown following the same steps, writing probabilities in terms of m instead of n .

Proof that (C.9) is satisfied with high probability

In order to show that (C.9) holds with high probability, we proceed as in the induction step of the proof of Lemma C.7: for any index $i \in [n]$, any subset $\mathcal{K} \subseteq \mathcal{S} \setminus \{i\}$, any subset $\mathcal{L} \subseteq \mathcal{K}$ with $|\mathcal{L}| \geq \frac{1}{2}|\mathcal{K}|$ and any vector $\mathbf{x} \in N_\epsilon^\mathcal{K}$, we consider the function

$$h_{i,\mathcal{K},\mathcal{L},\mathbf{x}}(\mathbf{a}) = \sum_{l \in \mathcal{L}} \frac{1}{m} \sum_{j=1}^m a_{ji} a_{jl} (\mathbf{a}_j^\top \mathbf{x})^2.$$

Then, we define B_i as the following event:

$$B_i = \left\{ |h_{i,\mathcal{K},\mathcal{L},\mathbf{x}}(\mathbf{A})| \leq c|\mathcal{L}| \frac{\log n}{\sqrt{m}} \text{ for all subsets } \mathcal{K} \subseteq \mathcal{S} \setminus \{i\}, \mathcal{L} \subseteq \mathcal{K} \text{ with } |\mathcal{L}| \geq \frac{1}{2}|\mathcal{K}| \text{ and } \mathbf{x} \in N_\epsilon^\mathcal{K} \right\}.$$

To bound the probability of B_i , we condition on $\mathbf{A}_{\cdot i} = \mathbf{a}_{\cdot i}$ using the formula

$$\mathbb{P}[B] = \int \mathbb{P}[B | \mathbf{A}_{\cdot i} = \mathbf{a}_{\cdot i}] \mu(\mathbf{a}_{\cdot i}) d\mathbf{a}_{\cdot i},$$

which holds for any event B , where we write μ for the standard normal density.

To show that the conditional probability $\mathbb{P}[B_i | \mathbf{A}_{\cdot i} = \mathbf{a}_{\cdot i}]$ is close to one, we begin by showing that, for any fixed \mathcal{K} , \mathcal{L} and \mathbf{x} as described above, $h_{i,\mathcal{K},\mathcal{L},\mathbf{x}}(\mathbf{A})|_{\mathbf{A}_{\cdot i}=\mathbf{a}_{\cdot i}}$ (making explicit the fact that we are conditioning on $\mathbf{A}_{\cdot i} = \mathbf{a}_{\cdot i}$) concentrates around its expectation $\mathbb{E}[h_{i,\mathcal{K},\mathcal{L},\mathbf{x}}(\mathbf{A}) | \mathbf{A}_{\cdot i} = \mathbf{a}_{\cdot i}] = 0$. To simplify notation, we will omit the subscripts and write $h(\mathbf{A})$ in what follows.

Step 1: Bound $h(\mathbf{A})$ conditioned on $i, \mathcal{K}, \mathcal{L}, \mathbf{x}$ and $\mathbf{A}_{\cdot i}$

Let \mathcal{A}_1 be as in Lemma C.7, and \mathcal{A}_2 as in Lemma C.3 with $t = 5\sqrt{\log n}$ (more precisely, we have $\mathcal{A}_1, \mathcal{A}_2 \subset \mathbb{R}^{m \times n}$, and we require the projections onto $\mathbb{R}^{m \times \mathcal{S}}$ to be as in the respective lemmas). Then, by these two Lemmas, we have $\mathbb{P}[\mathcal{A}_1 \cap \mathcal{A}_2] \geq 1 - c_2 n^{-11}$ for a universal constant $c_2 > 0$, and as the intersection of two convex sets, $\mathcal{A}_1 \cap \mathcal{A}_2$ is also convex.

Fix an $i \in [n]$, a subset $\mathcal{K} \subseteq \mathcal{S} \setminus \{i\}$, a subset $\mathcal{L} \subseteq \mathcal{K}$ with $|\mathcal{L}| \geq \frac{1}{2}|\mathcal{K}|$ and a vector $\mathbf{x} \in N_\epsilon^\mathcal{K}$. We begin by conditioning on $\mathbf{A}_{\cdot i} = \mathbf{a}_{\cdot i}$ for a vector $\mathbf{a}_{\cdot i}$ satisfying $\max_j |a_{ji}| \leq 5\sqrt{\log n}$. As in the proof of Lemma C.7, the idea is use Theorem C.1 to show that $h(\mathbf{A})$ is close to its expectation. However, $h(\mathbf{A})$ is not globally Lipschitz continuous. We will show that the function h restricted to $\mathcal{A}_1 \cap \mathcal{A}_2$ is Lipschitz continuous. Then, we extend this restricted function to a globally Lipschitz continuous function \tilde{h} on the entire space, and apply Theorem C.1 to \tilde{h} . By construction, we have $h(\mathbf{A}) = \tilde{h}(\mathbf{A})$ with high probability, which therefore also yields a high probability bound for $h(\mathbf{A})$.

Step 1, part (a): Bound the Lipschitz constant of h restricted to $\mathcal{A}_1 \cap \mathcal{A}_2$

Restricted to $\mathcal{A}_1 \cap \mathcal{A}_2$, we can bound the Lipschitz constant of h by the norm of its gradient. The norm of the gradient is an upper bound for the Lipschitz constant by the mean-value theorem, since $\mathcal{A}_1 \cap \mathcal{A}_2$ is a convex set. Note that h only depends on a_{jr} for $r \in \mathcal{K}$, since we have $x_r = 0$ for $r \notin \mathcal{K}$. We have

$$\frac{\partial}{\partial a_{jr}} h(\mathbf{a}) = \begin{cases} \frac{2}{m} a_{ji} x_r (\mathbf{a}_j^\top \mathbf{x}) \sum_{l \in \mathcal{L}} a_{jl} & r \in \mathcal{K} \setminus \mathcal{L} \\ \frac{2}{m} a_{ji} x_r (\mathbf{a}_j^\top \mathbf{x}) \sum_{l \in \mathcal{L}} a_{jl} + \frac{1}{m} a_{ji} (\mathbf{a}_j^\top \mathbf{x})^2 & r \in \mathcal{L}. \end{cases}$$

Using the inequality $(a + b)^2 \leq 2(a^2 + b^2)$, we can bound

$$\|\nabla h(\mathbf{a})\|_2^2 = \sum_{r \in \mathcal{K}} \sum_{j=1}^m \left(\frac{\partial}{\partial a_{jr}} h(\mathbf{a}) \right)^2 \leq 8 \sum_{r \in \mathcal{K}} \sum_{j=1}^m \frac{1}{m^2} x_r^2 a_{ji}^2 \left(\sum_{l \in \mathcal{L}} a_{jl} \right)^2 (\mathbf{a}_j^\top \mathbf{x})^2 + 2 \sum_{r \in \mathcal{L}} \sum_{j=1}^m \frac{1}{m^2} a_{ji}^2 (\mathbf{a}_j^\top \mathbf{x})^4. \quad (\text{C.18})$$

We bound the two sums separately. For the first sum, we have

$$\begin{aligned} \sum_{r \in \mathcal{K}} \sum_{j=1}^m \frac{1}{m^2} x_r^2 a_{ji}^2 \left(\sum_{l \in \mathcal{L}} a_{jl} \right)^2 (\mathbf{a}_j^\top \mathbf{x})^2 &= \frac{1}{m} \sum_{r \in \mathcal{K}} x_r^2 \sum_{j=1}^m \frac{1}{m} a_{ji}^2 \left(\sum_{l \in \mathcal{L}} a_{jl} \right)^2 (\mathbf{a}_j^\top \mathbf{x})^2 \\ &\leq \left(\max_j a_{ji}^2 \right) \frac{1}{m} \sum_{r \in \mathcal{K}} x_r^2 \sum_{j=1}^m \frac{1}{m} \left(\sum_{l \in \mathcal{L}} a_{jl} \right)^2 (\mathbf{a}_j^\top \mathbf{x})^2 \\ &\leq \frac{25 \log n}{m} \sum_{r \in \mathcal{K}} x_r^2 \sqrt{\frac{1}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} a_{jl} \right)^4} \sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x})^4}, \end{aligned}$$

where we used Hölder's inequality in the second and the Cauchy–Schwarz inequality in the last line. Since $\mathbf{a} \in \mathcal{A}_1$, we can apply Lemma C.7 to bound

$$\sqrt{\frac{1}{m} \sum_{j=1}^m \left(\sum_{l \in \mathcal{L}} a_{jl} \right)^4} \leq c_3 |\mathcal{L}|,$$

where c_3 is the universal constant from Lemma C.7. Since also $\mathbf{a} \in \mathcal{A}_2$, Lemma C.3 yields

$$\sqrt{\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x})^4} \leq \sqrt{\frac{1}{m} \|\mathbf{a}_{\cdot \mathcal{K}}\|_{2 \rightarrow 4}^4} \leq 11,$$

as $\|\mathbf{x}\|_2 = 1$. Putting this together, we can bound the first sum in (C.18),

$$\sum_{r \in \mathcal{K}} \sum_{j=1}^m \frac{1}{m^2} x_r^2 a_{ji}^2 \left(\sum_{l \in \mathcal{L}} a_{jl} \right)^2 (\mathbf{a}_j^\top \mathbf{x})^2 \leq \frac{5 \log n}{m} \sum_{r \in \mathcal{K}} 11 c_3 |\mathcal{L}| x_r^2 \leq 55 c_3 \frac{|\mathcal{L}| \log n}{m},$$

where we again used $\|\mathbf{x}\|_2 = 1$. For the second sum in (C.18), we can write

$$\sum_{r \in \mathcal{L}} \sum_{j=1}^m \frac{1}{m^2} a_{ji}^2 (\mathbf{a}_j^\top \mathbf{x})^4 \leq \frac{5 \log n}{m} |\mathcal{L}| \frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x})^4 \leq 605 \frac{|\mathcal{L}| \log n}{m},$$

where we again used Lemma C.3 to bound $\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_j^\top \mathbf{x})^4 \leq 121$. All in all, we have shown

$$\|\nabla h(\mathbf{a})\|_2 = \left(\sum_{r \in \mathcal{K}} \sum_{j=1}^m \left(\frac{\partial}{\partial a_{jr}} h(\mathbf{a}) \right)^2 \right)^{\frac{1}{2}} \leq \sqrt{c_4 \frac{|\mathcal{L}| \log n}{m}}$$

for $\mathbf{a} \in \mathcal{A}_1 \cap \mathcal{A}_2$ and $c_4 = 440c_3 + 1210$.

Step 1, part (b): Construct a globally Lipschitz continuous extension of h

Consider the following Lipschitz extension of the function h :

$$\tilde{h}(\mathbf{a}) = \inf_{\mathbf{a}' \in \mathcal{A}_1 \cap \mathcal{A}_2} (h(\mathbf{a}') + \text{Lip}(h) \|\mathbf{a} - \mathbf{a}'\|_2),$$

where we write $\text{Lip}(h) = \sqrt{c_4 \frac{|\mathcal{L}| \log n}{m}}$. By definition, we have $\tilde{h} = h$ on $\mathcal{A}_1 \cap \mathcal{A}_2$, and it follows from an application of the triangle inequality that \tilde{h} is globally Lipschitz continuous with Lipschitz constant $\text{Lip}(h)$ (see e.g. Theorem 7.2 of [6]).

We will show that \tilde{h} concentrates around its mean, which can potentially differ from the mean of h . Since h and \tilde{h} differ only on $(\mathcal{A}_1 \cap \mathcal{A}_2)^c$, which has probability at most $c_2 n^{-11}$, we can bound, using the Cauchy–Schwarz inequality,

$$\mathbb{E}[|h(\mathbf{A})| \mathbf{1}_{(\mathcal{A}_1 \cap \mathcal{A}_2)^c}(\mathbf{A})] \leq \sqrt{\mathbb{E}[h(\mathbf{A})^2]} \sqrt{\mathbb{E}[\mathbf{1}_{(\mathcal{A}_1 \cap \mathcal{A}_2)^c}(\mathbf{A})]} \leq \frac{c'}{n^5},$$

where we used that $\mathbb{E}[h(\mathbf{A})^2] \leq \sqrt{315} k^2 / m$. Using $\tilde{h}(\mathbf{a}) \leq \text{Lip}(h) \|\mathbf{a}\|_2$, we can bound

$$\mathbb{E}[|\tilde{h}(\mathbf{A})| \mathbf{1}_{(\mathcal{A}_1 \cap \mathcal{A}_2)^c}(\mathbf{A})] \leq \text{Lip}(h) \sqrt{\mathbb{E}[\|\mathbf{A}\|_2^2]} \sqrt{\mathbb{E}[\mathbf{1}_{(\mathcal{A}_1 \cap \mathcal{A}_2)^c}(\mathbf{A})]} \leq \frac{c''}{n^3}.$$

All in all, this shows that

$$|\mathbb{E}[h(\mathbf{A})] - \mathbb{E}[\tilde{h}(\mathbf{A})]| \leq \frac{c_5}{n^3}.$$

for a constant $c_5 > 0$. Hence, we can apply Theorem C.1 to obtain

$$\mathbb{P}\left[|\tilde{h}(\mathbf{A})| > c \frac{|\mathcal{L}| \log n}{\sqrt{m}} \mid \mathbf{A}_{\cdot i} = \mathbf{a}_{\cdot i}\right] \leq 2 \exp\left(-\frac{(c \frac{|\mathcal{L}|}{\sqrt{m}} \log n - c_5/n^3)^2}{2c_4 \frac{|\mathcal{L}| \log n}{m}}\right) \leq 2 \exp(-c_6 |\mathcal{L}| \log n),$$

for a universal constant $c_6 \leq \frac{c}{2c_4} - \frac{cc_5}{c_4} \frac{\sqrt{m}}{n^3}$, where we used the assumption $m \leq n^{3/2}$.

Step 2: Unravel the conditions: take union bounds and integrate over $\mathbf{a}_{\cdot i}$

Let

$$\begin{aligned} \tilde{B}_{i,\lambda} = \left\{ |\tilde{h}_{i,\mathcal{K},\mathcal{L},\mathbf{x}}(\mathbf{A})| \leq c \frac{|\mathcal{L}| \log n}{\sqrt{m}} \quad \text{for all subsets } \mathcal{K} \subseteq \mathcal{S} \setminus \{i\} \text{ with } |\mathcal{K}| = \lambda, \right. \\ \left. \mathcal{L} \subseteq \mathcal{K} \text{ with } |\mathcal{L}| \geq \frac{1}{2} |\mathcal{K}| \text{ and } \mathbf{x} \in N_\epsilon^\mathcal{K} \right\}, \end{aligned}$$

and let \tilde{B}_i be the same event without the restriction $|\mathcal{K}| = \lambda$. The cardinality of the ϵ -net can be bounded by $|N_\epsilon^\mathcal{K}| \leq (3/\epsilon)^\lambda \leq (3m/c_1)^\lambda$. Taking union bounds over all $\mathbf{x} \in N_\epsilon^\mathcal{K}$, $\mathcal{L} \subseteq \mathcal{K}$ with $|\mathcal{L}| \geq \frac{1}{2} |\mathcal{K}|$ and $\mathcal{K} \subseteq \mathcal{S} \setminus \{i\}$ with fixed cardinality $|\mathcal{K}| = \lambda$, we obtain (using the upper bound $\binom{k}{\lambda} \leq (\frac{ek}{\lambda})^\lambda$)

$$\mathbb{P}[\tilde{B}_{i,\lambda} \mid \mathbf{A}_{\cdot i} = \mathbf{a}_{\cdot i}] \geq 1 - 2 \exp\left(-\frac{c_6}{2} \lambda \log n + \lambda \log \frac{3m}{c_1} + \lambda \log 2 + \lambda \log \frac{ek}{\lambda}\right) \geq 1 - 2 \exp(-c_7 \lambda \log n)$$

for a constant $c_7 > 0$ if $c_6 \geq 5 + 2 \frac{\log(6e/c_1)}{\log n}$, where we used the assumption $m \leq n^{3/2}$. Taking the union bound over all possible choices for the cardinalities $\lambda = 1, \dots, k$ gives

$$\mathbb{P}[\tilde{B}_i \mid \mathbf{A}_{\cdot i} = \mathbf{a}_{\cdot i}] \geq 1 - \frac{2}{1 - 2 \exp(-c_7 \log n)} \exp(-c_7 \log n) \geq 1 - 2.1 \exp(-c_7 \log n),$$

provided that $n^{-c_7} \leq 0.1/4.2$. Next, we integrate over all \mathbf{a}_i satisfying $\max_j |a_{ji}| \leq 5\sqrt{\log n}$:

$$\begin{aligned} \mathbb{P}[\tilde{B}_i] &\geq \int_{\{\max_j |a_{ji}| \leq 5\sqrt{\log n}\}} \mathbb{P}[\tilde{B}_i \mid \mathbf{A}_i = \mathbf{a}_i] \mu(\mathbf{a}_i) d\mathbf{a}_i \\ &\geq \left(1 - 2.1e^{-c_7 \log n}\right) \mathbb{P}\left[\max_j |a_{ji}| \leq 5\sqrt{\log n}\right] \\ &\geq 1 - 4n^{-11}, \end{aligned}$$

where we write μ for the Gaussian density. The last line follows from standard Gaussian tail bounds and a union bound where we used the assumption $m \leq n^{3/2}$, provided that $c_7 \geq 11$.

On the set $\mathcal{A}_1 \cap \mathcal{A}_2 \cap \{\max_j |a_{ji}| \leq 5\sqrt{\log n}\}$, which is independent of the choices for \mathcal{K} , \mathcal{L} and \mathbf{x} , we have $h_{i,\mathcal{K},\mathcal{L},\mathbf{x}} = \tilde{h}_{i,\mathcal{K},\mathcal{L},\mathbf{x}}$ for all subsets $\mathcal{K} \subseteq \mathcal{S} \setminus \{i\}$, $\mathcal{L} \subseteq \mathcal{K}$ with $|\mathcal{L}| \geq \frac{1}{2}|\mathcal{K}|$ and $\mathbf{x} \in N_\epsilon^\mathcal{K}$. In this case, the bound on $\tilde{h}_{\mathcal{K},\mathcal{L},\mathbf{x}}(\mathbf{A})$ also applies to $h_{\mathcal{K},\mathcal{L},\mathbf{x}}(\mathbf{A})$, so that

$$\mathbb{P}[B_i] \geq \mathbb{P}[\tilde{B}_i \cap \mathcal{A}_1 \cap \mathcal{A}_2 \cap \{\max_j |a_{ji}| \leq 5\sqrt{\log n}\}] \geq 1 - 4n^{-11} - c_2 n^{-11} - n^{-11} \geq 1 - \frac{c_p}{4} n^{-11},$$

for $c_p = 4c_2 + 20$. Finally, taking the union bound over all $i \in [n]$ establishes the bound (C.9) with probability at least $1 - \frac{c_p}{4} n^{-10}$.

Proof that (C.10)–(C.12) are satisfied with high probability

Inequalities (C.10) and (C.12) can be shown following the exact same steps as above for (C.9). The former is slightly simpler because of the lower order of the expression we need to control, which is also the reason for the term $\sqrt{\log n}$ in the bound (C.10) instead of $\log n$. The proof of inequality (C.11) is identical to the one of (C.10), as we can condition on the random variables $\{\varepsilon_j\}_{j=1}^m$ (the same way we conditioned on $\mathbf{A}_i = \mathbf{a}_i$ in Step 1 above) and use the standard tail bound $\mathbb{P}[|\varepsilon_j| > t] \leq \exp(1 - \tilde{c}t/\sigma^2)$ for sub-exponential random variables, where $\tilde{c} > 0$ is an absolute constant. We omit the details of the proof to avoid repetition. \square

Proof of Lemma C.9. We will show that each of the inequalities (C.13)–(C.15) is satisfied with probability at least $1 - \frac{c_p}{4} n^{-10}$. The lemma then follows by taking a union bound.

We begin by showing the bound (C.13). To show that (C.13) holds with high probability, we will show that it is satisfied for all $\{\mathbf{a}_j\}_{j=1}^m \in \mathcal{A} \subseteq \mathbb{R}^{m \times n}$, where \mathcal{A} is defined as the set where the bounds from Lemmas C.5 and C.8 are satisfied, and such that the projection of \mathcal{A} onto $\mathbb{R}^{m \times \mathcal{S}}$ satisfies Lemma C.3 with $t = 5\sqrt{\log n}$. By the aforementioned lemmas, we have $\mathbb{P}[\{\mathbf{A}_j\}_{j=1}^m \in \mathcal{A}] \geq 1 - \frac{c_p}{4} n^{-10}$ for a universal constant $c_p > 0$.

Proof that (C.13) is satisfied with high probability

Let $\{\mathbf{a}_j\}_{j=1}^m \in \mathcal{A}$ and fix any $i \in [n]$. We denote the term that we need to bound in (C.13) by

$$g(\mathbf{x}) = \frac{1}{m} \sum_{j=1}^m a_{ji} (\mathbf{a}_{j,-i}^\top \mathbf{x}_{-i})^3,$$

and we consider the constrained optimization problem

$$\max_{\mathbf{x}} g(\mathbf{x}) \quad \text{s.t.} \quad \mathbf{x} \in \mathcal{X}_b, \tag{C.19}$$

where we write $\mathcal{X}_b = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}_{S^c} = \mathbf{0}, \|\mathbf{x}\|_2^2 \leq 1, \|\mathbf{x}\|_1 \leq b\}$ for some $0 \leq b \leq \sqrt{k}$. We will show that any optimizer $\mathbf{x}' \in \mathcal{X}_b$ satisfies $g(\mathbf{x}') \leq cb\sqrt{\frac{\log n}{m}}$. By symmetry, we can also obtain the lower bound $g(\mathbf{x}') \geq -cb\sqrt{\frac{\log n}{m}}$ by considering the corresponding minimization problem. Since the above holds for an arbitrary $0 \leq b \leq \sqrt{k}$, this would then complete the proof that inequality (C.13) is satisfied for all $\{\mathbf{a}_j\}_{j=1}^m \in \mathcal{A}$ and $i \in [n]$.

Since we are maximizing a continuous function over a compact set, a global maximum is attained at some feasible point $\mathbf{x}' \in \mathcal{X}_b$. Using the KKT conditions at \mathbf{x}' , we will bound the Lagrange multiplier μ_2^* corresponding to the constraint $\|\mathbf{x}\|_1 \leq b$. This controls how the maximum attainable value $g(\mathbf{x}')$ can increase if we relax the constraint $\|\mathbf{x}\|_1 \leq b$ (for details see e.g. Section 5.6 of [1]) and integrating over b gives the desired result.

Step 1: Establish KKT conditions

Let $\mathbf{x}' \in \mathcal{X}_b$ be an optimal point of the constrained optimization problem (C.19). In order to establish the KKT conditions at \mathbf{x}' , we verify that the linear independence constraint qualification (LICP) (see e.g. [7]) holds in this problem: the gradients of all active inequality constraints are linearly independent at any point \mathbf{x} . We restrict our attention to $b \neq \sqrt{s}$ for $s = 1, \dots, k$, which is a set of measure zero and can be ignored when integrating. If only one inequality constraint is binding, there is nothing to show. Otherwise, the gradients of the constraints, $\frac{\partial}{\partial \mathbf{x}} \|\mathbf{x}\|_2^2 = 2\mathbf{x}$ and $\frac{\partial}{\partial \mathbf{x}} \|\mathbf{x}\|_1 = \text{sign}(\mathbf{x})$ can only be linearly dependent if $x_i \in \{-\tilde{c}, 0, \tilde{c}\}$ for all coordinates $i \in [n]$, where $\tilde{c} > 0$ is some constant. But then $\|\mathbf{x}\|_2^2 = s\tilde{c}^2$ and $\|\mathbf{x}\|_1 = s\tilde{c}$, where s is the number of non-zero coordinates of \mathbf{x} . Since $b \neq \sqrt{s}$, the two constraints cannot be simultaneously binding.

From LICP it follows that the following KKT conditions are satisfied at \mathbf{x}' : writing

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = g(\mathbf{x}) - \mu_1(\|\mathbf{x}\|_2^2 - 1) - \mu_2(\|\mathbf{x}\|_1 - b) + \sum_{r \in S^c} \lambda_r x_r$$

for the Lagrangian, there exist Lagrange multipliers $\boldsymbol{\lambda}^*$ and $\boldsymbol{\mu}^*$ satisfying

$$\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}', \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) = \mathbf{0} \quad \text{stationarity} \quad (\text{C.20})$$

$$\mathbf{x}_{S^c} = \mathbf{0}, \quad \|\mathbf{x}'\|_2^2 \leq 1, \quad \|\mathbf{x}'\|_1 \leq b \quad \text{primal feasibility} \quad (\text{C.21})$$

$$\mu_1^* \geq 0, \quad \mu_2^* \geq 0 \quad \text{dual feasibility} \quad (\text{C.22})$$

$$\mu_1^*(\|\mathbf{x}'\|_2^2 - 1) = 0, \quad \mu_2^*(\|\mathbf{x}'\|_1 - b) = 0 \quad \text{complementary slackness.} \quad (\text{C.23})$$

Step 2: Bounding the Lagrange multiplier μ_2^*

Using the stationarity condition (C.20), we can compute, for any $l \in \mathcal{S} \setminus \{i\}$,

$$\begin{aligned} \frac{\partial}{\partial x_l} \mathcal{L}(\mathbf{x}', \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) &= \frac{3}{m} \sum_{j=1}^m a_{ji} a_{jl} (\mathbf{a}_{j,-i}^\top \mathbf{x}'_{-i})^2 - 2\mu_1^* x'_l - \mu_2^* \text{sign}(x'_l) = 0 \\ \Rightarrow \quad h_l(\mathbf{x}') &:= \frac{1}{m} \sum_{j=1}^m a_{ji} a_{jl} (\mathbf{a}_{j,-i}^\top \mathbf{x}'_{-i})^2 = \frac{2}{3} \mu_1^* x'_l + \frac{1}{3} \mu_2^* \text{sign}(x'_l). \end{aligned}$$

This identity implies that $h_l(\mathbf{x}')$ has the same sign as x'_l because of dual feasibility (C.22), and that $x'_l = 0$ if and only if $h_l(\mathbf{x}') = 0$. Rearranging for μ_2^* , we obtain, for any l with $x'_l \neq 0$,

$$\mu_2^* = 3|h_l(\mathbf{x}')| - 2\mu_1^* |x'_l| \leq 3|h_l(\mathbf{x}')|, \quad (\text{C.24})$$

where we again used that $\mu_1^* \geq 0$. Next, we show that

$$\min_{l: x'_l \neq 0} |h_l(\mathbf{x}')| \leq \frac{c \log n}{3\sqrt{m}}. \quad (\text{C.25})$$

To this end, let $\mathcal{K} = \{r \in \mathcal{S} : x'_r \neq 0\}$ and let $\mathcal{L} = \{l \in \mathcal{K} : x'_l > 0\}$. Note that we must have $x'_i = 0$, as x_i does not contribute to the value of $g(\mathbf{x})$. We can assume without loss of generality that $|\mathcal{L}| \geq \frac{1}{2}|\mathcal{K}|$, as we can otherwise consider the set defined by $x'_l < 0$. Further, assume for notational simplicity that $\|\mathbf{x}'\|_2 = 1$; otherwise, we can replace \mathbf{x}' by $\mathbf{x}'/\|\mathbf{x}'\|_2$, which makes every $|h_l(\mathbf{x}')|$ larger, since we have $\|\mathbf{x}'\|_2 \leq 1$ by primal feasibility (C.21).

Let $N_\epsilon^\mathcal{K}$ be an ϵ -net of the set $\{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}_{\mathcal{K}^c} = \mathbf{0}, \|\mathbf{x}\|_2 = 1\}$, where $\epsilon = c \log n / (465\sqrt{m})$. Let $\mathbf{x} \in N_\epsilon^\mathcal{K}$ with $\|\mathbf{x} - \mathbf{x}'\|_2 \leq \epsilon$. Then, we can use inequality (C.9) of Lemma C.8 to bound

$$\sum_{l \in \mathcal{L}} h_l(\mathbf{x}) \leq c_1 |\mathcal{L}| \frac{\log n}{\sqrt{m}} \leq \frac{c |\mathcal{L}| \log n}{6\sqrt{m}},$$

provided that $c \geq 6c_1$, where c_1 is the universal constant from Lemma C.8. Since $h_l(\mathbf{x}) > 0$ for all $l \in \mathcal{L}$, there must be an index $l \in \mathcal{L}$ with $h_l(\mathbf{x}) \leq c \log n / (6\sqrt{m})$ by the pigeonhole principle. By the

Cauchy–Schwarz inequality, we can bound

$$\begin{aligned}
|h_l(\mathbf{x}) - h_l(\mathbf{x}')| &= \left| \frac{1}{m} \sum_{j=1}^m a_{ji} a_{jl} (\mathbf{a}_{j,-i}^\top (\mathbf{x}_{-i} + \mathbf{x}'_{-i})) (\mathbf{a}_{j,-i}^\top (\mathbf{x}_{-i} - \mathbf{x}'_{-i})) \right| \\
&\leq \left(\frac{1}{m} \sum_{j=1}^m a_{ji}^2 a_{jl}^2 (\mathbf{a}_{j,-i}^\top (\mathbf{x}_{-i} + \mathbf{x}'_{-i}))^2 \right)^{\frac{1}{2}} \left(\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_{j,-i}^\top (\mathbf{x}_{-i} - \mathbf{x}'_{-i}))^2 \right)^{\frac{1}{2}} \\
&\leq \left(\frac{1}{m} \sum_{j=1}^m a_{ji}^8 \right)^{\frac{1}{8}} \left(\frac{1}{m} \sum_{j=1}^m a_{jl}^8 \right)^{\frac{1}{8}} \left(\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_{j,-i}^\top (\mathbf{x}_{-i} + \mathbf{x}'_{-i}))^4 \right)^{\frac{1}{4}} \\
&\quad \cdot \left(\frac{1}{m} \sum_{j=1}^m (\mathbf{a}_{j,-i}^\top (\mathbf{x}_{-i} - \mathbf{x}'_{-i}))^2 \right)^{\frac{1}{2}} \\
&\leq 106^{\frac{1}{4}} \frac{(3m)^{\frac{1}{4}} + \sqrt{k} + 5\sqrt{\log n}}{m^{\frac{1}{4}}} \|\mathbf{x} + \mathbf{x}'\|_2 \frac{\sqrt{m} + \sqrt{k} + 5\sqrt{\log n}}{\sqrt{m}} \|\mathbf{x} - \mathbf{x}'\|_2 \\
&\leq \frac{c \log n}{6\sqrt{m}},
\end{aligned}$$

where we used Lemmas C.3 and C.5, and $\|\mathbf{x} - \mathbf{x}'\| \leq c \log n / (465\sqrt{m})$. Since $h_l(\mathbf{x}') \geq 0$ is non-negative, this completes the proof of (C.25). Combining (C.24) with (C.25), we obtain

$$\mu_2^\star \leq c \frac{\log n}{\sqrt{m}}.$$

Step 3: Showing the inequality (C.13)

In order to show the inequality (C.13), define the value function

$$v(b) = \max_{\mathbf{x}} \{g(\mathbf{x}) : \mathbf{x}_{S^c} = \mathbf{0}, \|\mathbf{x}\|_2^2 \leq 1, \|\mathbf{x}\|_1 \leq b\}.$$

We use that μ_2^\star is the shadow price of the constraint $\|\mathbf{x}\|_1 \leq b$ (see e.g. Section 5.6 of [1]):

$$\frac{\partial}{\partial b} v(b) = \mu_2^\star(b).$$

By definition, we have $v(0) = 0$. For any $0 \leq b_0 \leq \sqrt{k}$, the fundamental theorem of calculus yields

$$v(b_0) = \int_0^{b_0} \frac{\partial}{\partial b} v(b) db = \int_0^{b_0} \mu_2^\star(b) db \leq \int_0^{b_0} c \frac{\log n}{\sqrt{m}} db < c b_0 \frac{\log n}{\sqrt{m}}.$$

By definition, we have $g(\mathbf{x}) \leq v(b_0)$ for all $\mathbf{x} \in \mathcal{X}$ with $\|\mathbf{x}\|_1 \leq b_0$. Hence,

$$g(\mathbf{x}) \leq c b \frac{\log n}{\sqrt{m}}$$

must hold for any \mathbf{x} with $\|\mathbf{x}\|_2^2 \leq 1$ and $\|\mathbf{x}\|_1 \leq b$, which completes the proof of (C.13).

Proof that (C.14)–(C.16) are satisfied with high probability

The proofs of (C.14), (C.15) and (C.16) follow the same steps as the proof of (C.13), using the bounds (C.10), (C.11) and (C.12) of Lemma C.8, respectively, to bound the shadow price μ_2^* . We omit the details of the proof to avoid repetition. \square