



HAL
open science

Stochastic gradient descent in high dimensions for multi-spiked tensor PCA

G rard Ben Arous, C dric Gerbelot, Vanessa Piccolo

► **To cite this version:**

G rard Ben Arous, C dric Gerbelot, Vanessa Piccolo. Stochastic gradient descent in high dimensions for multi-spiked tensor PCA. 2024. <hal-04850396>

HAL Id: hal-04850396

<https://hal.science/hal-04850396v1>

Preprint submitted on 20 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destin e au d p t et   la diffusion de documents scientifiques de niveau recherche, publi s ou non,  manant des  tablissements d'enseignement et de recherche fran ais ou  trangers, des laboratoires publics ou priv s.



HAL Authorization

STOCHASTIC GRADIENT DESCENT IN HIGH DIMENSIONS FOR MULTI-SPIKED TENSOR PCA

GÉRARD BEN AROUS¹, CÉDRIC GERBELOT¹, AND VANESSA PICCOLO²

ABSTRACT. We study the dynamics in high dimensions of online stochastic gradient descent for the multi-spiked tensor model. This multi-index model arises from the tensor principal component analysis (PCA) problem with multiple spikes, where the goal is to estimate r unknown signal vectors within the N -dimensional unit sphere through maximum likelihood estimation from noisy observations of a p -tensor. We determine the number of samples and the conditions on the signal-to-noise ratios (SNRs) required to efficiently recover the unknown spikes from natural random initializations. We show that full recovery of all spikes is possible provided a number of sample scaling as N^{p-2} , matching the algorithmic threshold identified in the rank-one case [9, 10]. Our results are obtained through a detailed analysis of a low-dimensional system that describes the evolution of the correlations between the estimators and the spikes, while controlling the noise in the dynamics. We find that the spikes are recovered sequentially in a process we term “sequential elimination”: once a correlation exceeds a critical threshold, all correlations sharing a row or column index become sufficiently small, allowing the next correlation to grow and become macroscopic. The order in which correlations become macroscopic depends on their initial values and the corresponding SNRs, leading to either exact recovery or recovery of a permutation of the spikes. In the matrix case, when $p = 2$, if the SNRs are sufficiently separated, we achieve exact recovery of the spikes, whereas equal SNRs lead to recovery of the subspace spanned by the spikes.

CONTENTS

1. Introduction	2
1.1. Model and algorithm	3
1.2. Main results	4
1.3. Related works	10
1.4. Overview	11
2. Outline of proofs	12
2.1. Full recovery of spikes	12
2.2. Subspace recovery	16
3. Main results	16
4. Comparison inequalities	19
4.1. Preliminary results	19
4.2. Comparison inequalities for correlations	22
4.3. Comparison inequalities for eigenvalues	26
5. Proofs of full recovery of spikes	32
5.1. Proofs for $p \geq 3$	32
5.2. Proofs for $p = 2$ and separated SNRs	41
6. Proof of subspace recovery	48
Appendix A. Comparison inequalities for discrete sequences	53
Appendix B. Matrix concentration inequalities	54
References	55

¹COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY

²UNITÉ DE MATHÉMATIQUES PURES ET APPLIQUÉES (UMPA), ENS LYON

E-mail addresses: benarous@cims.nyu.edu, cedric.gerbelot@cims.nyu.edu, vanessa.piccolo@ens-lyon.fr.

Date: December 20, 2024.

2020 Mathematics Subject Classification. 68Q87, 62F30, 60G42.

Key words and phrases. Multi-index model, Multi-spiked tensor PCA, Tensor estimation, Online stochastic gradient descent, Sequential recovery, Subspace recovery.

1. INTRODUCTION

Understanding the dynamics of gradient-based algorithms for optimizing high-dimensional, non-convex random functions is crucial for advancing data science. Despite the challenges posed by non-convex landscapes, first-order optimization methods such as stochastic gradient descent (SGD) have proven remarkably successful in practice [54], especially in the context of deep learning applications, see e.g. [41, 18]. The time and sample complexity guarantees for SGD are well-understood in convex and quasi-convex settings [19, 47, 48, 27]; however, achieving similar results for non-convex cases remains challenging. A promising approach to address these challenges is to focus on specific instances where the target function is simple enough for closed-form analysis, yet retains the complexities found in larger models. A key observation when training such models with gradient-based methods is the emergence of a few preferred degrees of freedom that govern the behavior of the algorithm, as illustrated for instance by the *neural collapse phenomenon* [49, 50]. Recent work by the first author, in collaboration with Gheissari and Jagannath [9–12] and with Gheissari, Huang, and Jagannath [7], has shown that, for a broad range of estimation and classification tasks, the high-dimensional dynamics of (stochastic) gradient descent can be reduced to low-dimensional, autonomous *effective dynamics* involving a set of *summary statistics*. This allows for a precise understanding of how key factors such as initialization, step-size, sample size, and runtime affect the algorithm’s trajectory toward fixed points in the optimization landscape.

Building on this line of work, the present paper focuses on the benchmark problem of *multi-spiked tensor estimation*, which results in effective dynamics exhibiting a rich phenomenology at the level of the summary statistics that govern the recovery process. The multi-spiked tensor principal component analysis (PCA) problem was introduced by Johnstone [39] for $p = 2$ and by Richard and Montanari [53] for $p \geq 3$. The goal is to estimate r unknown vectors on the N -dimensional unit sphere from noisy observations of a corresponding rank r p -tensor. Detection thresholds for the rank-one spiked tensor model have been established in [43, 52, 51], and for the finite-rank case in [42, 23]. Moreover, recovery thresholds for the rank-one spiked tensor model have been identified for first-order optimization methods in [53, 9], and for Sum-of-Squares and spectral methods in [52, 51]. In the multi-spiked setting, recovery thresholds have been determined for power iteration in [37]. For a more detailed discussion on detection and recovery thresholds, we refer the reader to Subsection 1.3. Our goal is to determine the number of samples required for SGD to efficiently recover the hidden vectors. Additionally, in the extended version of our preprint [6], we also study continuous high-dimensional dynamics, such as Langevin dynamics and gradient flow, which yield suboptimal thresholds compared to SGD.

Our work serves as a benchmark for general multi-index models, which have garnered significant attention in recent years [59, 28, 70, 32, 17, 46, 25, 3, 15, 2, 16, 26]. They consist of functions that are parameterized by a finite number of (unknown) relevant directions $\mathbf{v}_1, \dots, \mathbf{v}_r \in \mathbb{R}^N$ and a multivariate link function $\sigma^*: \mathbb{R}^r \rightarrow \mathbb{R}$, which takes as input the inner products between input vectors in \mathbb{R}^N and the relevant directions. Here, N denotes the ambient dimension, and $r \ll N$ represents the number of latent coordinates. When $r = 1$, these models are referred to as single-index models. From a statistical perspective, multi-index models are canonical examples of target functions in which neural networks, able to exploit the latent low-dimensional structure, outperform kernel methods, which suffer from the curse of dimensionality [4]. These models also naturally arise in dimension reduction methods in high-dimensional statistics, as seen in [69, 68] and more recently in [31, 30]. In particular, [30] makes progress towards reducing the curse of dimensionality in non-parametric methods. From an optimization perspective, multi-index models serve as a typical example of systems leading to low-dimensional effective dynamics. For instance, when considering a set of r orthonormal estimators $\mathbf{x}_1, \dots, \mathbf{x}_r \in \mathbb{R}^N$, as will be the focus of this paper, the effective dynamics are governed by the r^2 inner products $\langle \mathbf{v}_i, \mathbf{x}_j \rangle$. In this paper, we study this r^2 -dimensional dynamical system within the multi-spiked PCA problem. Our goal is to determine the sample complexity required for the high-dimensional empirical dynamics of SGD to approximate the effective low-dimensional population dynamics.

For the single-spiked tensor PCA model, [9] has shown that the difficulty of recovery using gradient flow is due to the strength of the curvature of the loss landscape around the initial condition, which is governed by the order p of the tensor. In the context of single-index models, the rank-one tensor PCA model can be viewed as a single monomial in the Taylor expansion of a more general link function $\sigma^*: \mathbb{R} \rightarrow \mathbb{R}$. Building on this insight, it was shown in [10] that when learning an unknown direction from a single-index model with a known link function using online SGD, the difficulty of recovery again depends on the curvature of the landscape near the initial condition. This curvature is quantified by

the first non-zero coefficient in the Taylor expansion of σ^* , referred to as the *information exponent* [10]. Building on this insight, various extensions have been proposed, notably [17, 15], which address the task of learning both the hidden direction and an unknown link function using tailored algorithms (for more details, see Subsection 1.3). Multi-index models present a significant challenge compared to single-index models: the effective dynamics are multidimensional and thus may exhibit a nonmonotone behavior and multiple fixed points, making it difficult to establish global convergence guarantees. Important progress in this direction has been made in [1, 2], where in particular the latter work provides a complete description of the learning dynamics for a two-layer neural network using a tailored, two-timescale modification of SGD for specific instances of multi-index models. Subsequent work by [26, 16] proposed further results using a similar algorithm, with the former exploring varying batch sizes and the latter implementing a fully non-parametric step to allow for generality at the level of the link function in the population case. The two-timescale framework employed in these works is crucial to their analysis, as it results in monotone dynamics or their superposition and provides guarantees for subspace recovery. Finally, single- and multi-index models have been extensively studied in the context of statistical physics of learning, see e.g. [56, 57, 34], and related rigorous methods in the proportional limit of dimension and number of samples [21, 33]. However, applying these results to the questions considered here seems challenging.

In light of the above discussion, the following questions naturally arise and motivate our work:

- (i) Can we characterize the global convergence of the effective population dynamics for gradient descent and its variants without relying on a two-timescale algorithm?
- (ii) Can we identify which fixed points are reached and determine their probabilities?
- (iii) Can we establish sample complexity guarantees for the empirical dynamics related to the previous two questions for SGD?

We address these questions in the context of the multi-spiked tensor PCA problem. To simplify the analysis, we assume that the unknown vectors are orthogonal. Our study reveals a much richer phenomenology compared to the single-spiked case. Specifically, we focus on three distinct notions of recovery which naturally arise: recovery of the spikes in the exact order, recovery of a permutation of the spikes, and recovery of a rotation of the hidden subspace. The main challenge of our analysis lies in simultaneously controlling the empirical noise and the interactions among the r^2 inner products $\langle \mathbf{v}_i, \mathbf{x}_j \rangle$ which arise in the population dynamics. This is crucial for quantifying the probability of reaching each fixed point during the recovery process. Our proof relies on systems of discrete-time difference inequalities, derived from the signal-plus-martingale decomposition method from [10, 60].

1.1. MODEL AND ALGORITHM

The multi-spiked tensor model is formalized as follows. Fix integers $p \geq 2$ and $r \geq 1$, and suppose that we are given M i.i.d. observations \mathbf{Y}^ℓ of a rank r p -tensor on \mathbb{R}^N in the presence of additive noise, i.e.,

$$\mathbf{Y}^\ell = \mathbf{W}^\ell + \sqrt{N} \sum_{i=1}^r \lambda_i \mathbf{v}_i^{\otimes p}, \quad (1.1)$$

where $(\mathbf{W}^\ell)_\ell$ are i.i.d. samples of a noise p -tensor with i.i.d. sub-Gaussian entries W_{i_1, \dots, i_p}^ℓ , $\lambda_1 \geq \dots \geq \lambda_r \geq 0$ are the signal-to-noise ratios (SNRs), assumed to be of order 1, and $\mathbf{v}_1, \dots, \mathbf{v}_r$ are *unknown, orthogonal vectors* lying on the N -dimensional unit sphere \mathbb{S}^{N-1} . Our goal is to estimate the unknown signal vectors using maximum likelihood estimation (MLE) based on the sample of i.i.d. training data $(\mathbf{Y}^\ell)_{\ell=1}^M$. In our setting, MLE amounts to solving the optimization problem:

$$\min_{\mathbf{X}: \mathbf{X}^\top \mathbf{X} = \mathbf{I}_r} \left\{ - \sum_{i=1}^r \lambda_i \langle \mathbf{Y}, \mathbf{x}_i^{\otimes p} \rangle \right\}, \quad (1.2)$$

where $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_r] \in \mathbb{R}^{N \times r}$ and the feasible set $\text{St}(N, r) = \{\mathbf{X} \in \mathbb{R}^{N \times r}: \mathbf{X}^\top \mathbf{X} = \mathbf{I}_r\}$ is known as the *Stiefel manifold*, once equipped with its natural sub-manifold structure from $\mathbb{R}^{N \times r}$. The optimization problem (1.2) is equivalent in law to minimizing the loss function $\mathcal{L}_{N,r}: \text{St}(N, r) \times (\mathbb{R}^N)^{\otimes p} \rightarrow \mathbb{R}$ given by

$$\mathcal{L}_{N,r}(\mathbf{X}; \mathbf{Y}) = H_{N,r}(\mathbf{X}) + \Phi_{N,r}(\mathbf{X}), \quad (1.3)$$

where $H_{N,r}$ denotes the noise part given by

$$H_{N,r}(\mathbf{X}) = \sum_{i=1}^r \lambda_i \langle \mathbf{W}, \mathbf{x}_i^{\otimes p} \rangle, \quad (1.4)$$

and $\Phi_{N,r}$ denotes the population loss given by

$$\Phi_{N,r}(\mathbf{X}) = \mathbb{E}[\mathcal{L}_{N,r}(\mathbf{X}; \mathbf{Y})] = - \sum_{1 \leq i, j \leq r} \sqrt{N} \lambda_i \lambda_j m_{ij}^p(\mathbf{X}), \quad (1.5)$$

where $m_{ij}(\mathbf{X}) = \langle \mathbf{v}_i, \mathbf{x}_j \rangle$.

Definition 1.1 (Correlation). For every $1 \leq i, j \leq r$, we call $m_{ij}(\mathbf{X})$ the *correlation* of \mathbf{x}_j with \mathbf{v}_i .

The model described in (1.1), also known as matrix and tensor PCA, was introduced by Johnstone [39] for the matrix case and by Richard and Montanari [53] for the tensor case, specifically in the case of a single unknown spike. Our work aims to solve the optimization problem (1.2) in the large- N limit while the rank r is fixed. As previously mentioned, to tackle this tensor estimation problem, we use a first-order optimization method, namely *online stochastic gradient descent (SGD)*. This local algorithm has been studied by the first author, joint with Gheissari and Jagannath in [10], to determine the algorithmic thresholds for the rank-one spiked tensor model. We define the online SGD algorithm as follows.

Let $\mathbf{X}_\ell \in \text{St}(N, r)$ denote the output of the algorithm at time ℓ . The online SGD algorithm with initialization $\mathbf{X}_0 \in \text{St}(N, r)$, which is possibly random, and with step-size parameter $\delta > 0$, will be run using the observations $(\mathbf{Y}^\ell)_{\ell=1}^M$ according to the following update rule:

$$\mathbf{X}_\ell = R_{\mathbf{X}_{\ell-1}} \left(-\frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}_{N,r}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell) \right), \quad (1.6)$$

where $R_{\mathbf{X}}: T_{\mathbf{X}}\text{St}(N, r) \rightarrow \text{St}(N, r)$ is a retraction map that moves a point from the tangent space $T_{\mathbf{X}}\text{St}(N, r)$ back onto the Stiefel manifold $\text{St}(N, r)$, while preserving its geometric properties such as orthonormality (see e.g. [20, Section 7.3]). The tangent space $T_{\mathbf{X}}\text{St}(N, r)$ at a point $\mathbf{X} \in \text{St}(N, r)$ is given by

$$T_{\mathbf{X}}\text{St}(N, r) = \{ \mathbf{V} \in \mathbb{R}^{N \times r} : \mathbf{X}^\top \mathbf{V} + \mathbf{V}^\top \mathbf{X} = 0 \}.$$

There are multiple ways to define the retraction map. Here, we use the polar retraction, defined as

$$R_{\mathbf{X}}(\mathbf{U}) = (\mathbf{X} + \mathbf{U}) \left(\mathbf{I}_r + \mathbf{U}^\top \mathbf{U} \right)^{-1/2}. \quad (1.7)$$

In the update rule (1.6), ∇_{St} denotes the Riemannian gradient on $\text{St}(N, r)$. In particular, for a function $F: \text{St}(N, r) \rightarrow \mathbb{R}$, the Riemannian gradient is given by

$$\nabla_{\text{St}} F(\mathbf{X}) = \nabla F(\mathbf{X}) - \frac{1}{2} \mathbf{X} (\mathbf{X}^\top \nabla F(\mathbf{X}) + \nabla F(\mathbf{X})^\top \mathbf{X}), \quad (1.8)$$

where ∇ denotes the Euclidean gradient. The algorithm runs for M steps, and the final output \mathbf{X}_M is taken as the estimator.

1.2. MAIN RESULTS

Our goal is to determine the *sample complexity* required to recover the unknown orthogonal vectors $\mathbf{v}_1, \dots, \mathbf{v}_r \in \mathbb{S}^{N-1}$. We formalize the notion of recovering all spikes $\mathbf{v}_1, \dots, \mathbf{v}_r$ as achieving macroscopic correlations with these vectors, which is defined as follows.

Definition 1.2 (Recovery of all spikes). We say that we *recover a permutation* of the r unknown vectors $\mathbf{v}_1, \dots, \mathbf{v}_r$ if there exists a permutation $\sigma \in \mathfrak{S}_r$ over r elements such that for every $i \in [r]$,

$$|m_{\sigma(i)i}(\mathbf{X})| = |\langle \mathbf{v}_{\sigma(i)}, \mathbf{x}_i \rangle| = 1 - o(1),$$

with high probability. If σ is the identity permutation, we say that we achieve *exact recovery* of the r spikes $\mathbf{v}_1, \dots, \mathbf{v}_r$ if for every $i \in [r]$,

$$|m_{ii}(\mathbf{X})| = |\langle \mathbf{v}_i, \mathbf{x}_i \rangle| = 1 - o(1),$$

with high probability.

From this point onward, we focus on the sequence of outputs $(\mathbf{X}_\ell)_{\ell \in \mathbb{N}}$ generated by the online SGD algorithm with step-size parameter $\delta > 0$, as described in (1.6). The sequence is initialized randomly, with \mathbf{X}_0 sampled from the uniform distribution $\mu_{N \times r}$ on the Stiefel manifold $\text{St}(N, r)$. The distribution $\mu_{N \times r}$ is the unique probability measure on $\text{St}(N, r)$ that is invariant under both the left- and right-orthogonal transformations. More precisely, if $\mathbf{X} \sim \mu_{N \times r}$, then for any orthogonal matrices $\mathbf{U} \in \mathcal{O}(N)$ and $\mathbf{V} \in \mathcal{O}(r)$, the matrix $\mathbf{U} \mathbf{X} \mathbf{V}$ has the same distribution as \mathbf{X} . Sampling from this distribution is straightforward in practice: if $\mathbf{Z} \in \mathbb{R}^{N \times r}$ is a Gaussian matrix with i.i.d. standard normal entries,

then the random matrix $\mathbf{Z}(\mathbf{Z}^\top \mathbf{Z})^{-1/2}$ is uniformly distributed on $\text{St}(N, r)$, as shown in [24, Theorem 2.2.1]. Furthermore, we denote by $(\Omega, \mathcal{F}, \mathbb{P})$ the probability space on which the order- p tensors $(\mathbf{W}^\ell)_{\ell}$ are defined. We let $\mathbb{P}_{\mathbf{X}_0}$ denote the law of the process \mathbf{X}_t initiated at $\mathbf{X}_0 \sim \mu_{N \times r}$. Similarly, $\mathbb{P}_{\mathbf{X}_0^+}$ denotes the law of the process \mathbf{X}_t initiated at $\mathbf{X}_0 \sim \mu_{N \times r}$ subject to the condition $m_{ij}(\mathbf{X}_0) > 0$ for all $1 \leq i, j \leq r$. Additionally, for two sequences x_N and y_N , we write $x_N \ll y_N$ if $x_N/y_N \rightarrow 0$ as $N \rightarrow \infty$.

We first present our main results for $p \geq 3$, followed by those for $p = 2$. Although stated in asymptotic form here, a stronger, nonasymptotic formulation of our results, including explicit constants and convergence rates, is provided in Section 3. Additionally, we recall that the SNRs $\lambda_1 \geq \dots \geq \lambda_r \geq 0$ are assumed to be of order 1 in N .

1.2.1. MAIN RESULTS FOR $p \geq 3$. Our first main result for $p \geq 3$ shows that exact recovery of all spikes is possible given sufficiently separated SNRs and a sample complexity scaling as N^{p-2} .

Theorem 1.3 (Exact recovery for $p \geq 3$). *Let $(\mathbf{X}_\ell)_{\ell \in \mathbb{N}}$ be online SGD initialized from $\mathbf{X}_0 \sim \mu_{N \times r}$ with step-size $\delta \ll \log(N)^{-1} N^{-\frac{p-3}{2}}$. Let $M \gg \log(N) N^{p-2}$. Assume that for every $\eta > 1$, $\lambda_i \geq \left(\frac{24}{\sqrt{2\pi}} \eta \sqrt{\log(\eta)}\right)^{p-2} \lambda_{i+1}$ for all $1 \leq i \leq r-1$. Then, for every $\varepsilon > 0$ and every $i \in [r]$,*

$$\lim_{N \rightarrow \infty} \mathbb{P}_{\mathbf{X}_0^+} (|m_{ii}(\mathbf{X}_M)| \geq 1 - \varepsilon) \geq 1 - \frac{1}{\eta}.$$

Theorem 1.3 shows that for every $i \in [r]$, the correlation $|m_{ii}(\mathbf{X})| = |\langle \mathbf{x}_i, \mathbf{v}_i \rangle|$ approaches 1 with high probability. This probability depends on the ratio between the SNRs, with higher ratios leading to a higher probability of exact recovery. If the SNRs are not sufficiently separated, exact recovery might not always be achievable and it may occur that \mathbf{x}_i recovers \mathbf{v}_j for $i \neq j$. In such cases, we refer to *recovery of a permutation*, as introduced by Definition 1.2. To define the permutation of the spikes being recovered, we introduce a procedure that identifies r distinct pairs of indices from a given matrix.

Definition 1.4 (Greedy maximum selection). For a matrix $\mathbf{A} \in \mathbb{R}^{r \times r}$ with distinct entries, we define the pairs $\{(i_k^*, j_k^*)\}_{k=1}^r$ recursively by

$$(i_k^*, j_k^*) = \arg \max_{1 \leq i, j \leq r - (k-1)} |\mathbf{A}^{(k-1)}|_{ij},$$

where $|\mathbf{A}^{(k-1)}|$ denotes the absolute value of the entries in the submatrix $\mathbf{A}^{(k-1)} \in \mathbb{R}^{(r-(k-1)) \times (r-(k-1))}$, obtained by removing the i_1^*, \dots, i_{k-1}^* th rows and j_1^*, \dots, j_{k-1}^* th columns from \mathbf{A} . We refer to the pairs $(i_1^*, j_1^*), \dots, (i_r^*, j_r^*)$ obtained through this procedure as the *greedy maximum selection* of \mathbf{A} .

We apply this operation to the matrix

$$\mathbf{I}_0 = \left(\lambda_i \lambda_j m_{ij}^{p-2}(\mathbf{X}_0) \mathbf{1}_{\{m_{ij}^{p-2}(\mathbf{X}_0) \geq 0\}} \right)_{1 \leq i, j \leq r}. \quad (1.9)$$

The initialization matrix $\mathbf{I}_0 \in \mathbb{R}^{r \times r}$ is a random matrix. While its entries could potentially be identical due to randomness, Lemma A.3 of [6] ensures that they are distinct with probability $1 - o(1)$. We can therefore define the greedy maximum selection of \mathbf{I}_0 . In the following, we let $(i_1^*, j_1^*), \dots, (i_r^*, j_r^*)$ denote the greedy maximum selection of \mathbf{I}_0 . Our second main result shows that recovery of a permutation of all spikes is possible, provided a sample complexity of order N^{p-2} .

Theorem 1.5 (Recovery of a permutation for $p \geq 3$). *Let $(\mathbf{X}_\ell)_{\ell \in \mathbb{N}}$ be online SGD initialized from $\mathbf{X}_0 \sim \mu_{N \times r}$ with step-size $\delta \ll \log(N)^{-1} N^{-\frac{p-3}{2}}$. Let $M \gg \log(N) N^{p-2}$. Then, for every $\varepsilon > 0$ and every $k \in [r]$,*

$$\lim_{N \rightarrow \infty} \mathbb{P}_{\mathbf{X}_0} \left(|m_{i_k^* j_k^*}(\mathbf{X}_M)| \geq 1 - \varepsilon \right) = 1.$$

Theorem 1.5 shows that, regardless of the values of the SNRs, recovery of a permutation of the spikes is always possible, provided a sample complexity of order N^{p-2} . The permutation of the recovered spikes is determined by the greedy maximum selection of \mathbf{I}_0 , which, in turn, depends on the values of the correlations at initialization and on the corresponding signal sizes. In both Theorems 1.3 and 1.5, the sample complexity threshold aligns with the sharp threshold obtained using first-order methods in the rank-one spiked tensor model [9, 10]. We remark that the factor $\log(N)$ in both the step-size and sample complexity is due to the control of the retraction term in (1.6), and can be removed in the case of a single spike.




FIGURE 1. Evolution of the correlations m_{ij} under the population dynamics in the case where $p = 3$, $r = 2$, $\lambda_1 = 3$ and $\lambda_2 = 1$. The SNRs are sufficiently separated to ensure exact recovery of both spikes \mathbf{v}_1 and \mathbf{v}_2 . Once m_{11} reaches a sufficiently large microscopic threshold, m_{12} and m_{21} begin to decrease. Recovery of \mathbf{v}_2 is then possible after m_{12} and m_{21} have decreased sufficiently, becoming negligible in the evolution of m_{22} .

The dynamics underlying Theorems 1.3 and 1.5 reveals a richer structure beyond the stated results. In particular, the correlation $m_{i_1^* j_1^*}$ becomes macroscopic first, followed by $m_{i_2^* j_2^*}$, and so on. Theorem 1.3 involves an additional assumption on the ratio between the signal sizes, which ensures that variations due to random initialization are mitigated by differences in the SNRs. As a result, the greedy maximum selection of \mathbf{I}_0 results in $i_k^* = j_k^* = k$ for every $k \in [r]$, allowing m_{11}, \dots, m_{rr} to become macroscopic in sequence. More precisely, as shown by the numerical simulations in Figures 1, 2 and 6, when $m_{i_1^* j_1^*}$ exceeds a critical microscopic threshold (which depends on the order p of the tensor and on the SNRs), the correlations $m_{i_1^* j}$ and $m_{i j_1^*}$ for $i \neq i_1^*, j \neq j_1^*$ decrease to sufficiently small levels. This allows $m_{i_2^* j_2^*}$ to grow and become macroscopic, and the process continues similarly for the remaining correlations $m_{i_k^* j_k^*}$. This phenomenon is referred to as the *sequential elimination process*: the correlations $\{m_{i_k^* j_k^*}\}_{k=1}^r$ increase one by one, sequentially eliminating correlations that share a row or column index. This is formalized as follows.

Definition 1.6 (Sequential elimination). Let $S = \{(i_1, j_1), \dots, (i_r, j_r)\}$ be a set with distinct $i_1, \dots, i_r \in [r]$ and distinct $j_1, \dots, j_r \in [r]$. We say that the correlations $\{m_{ij}(t)\}_{1 \leq i, j \leq r}$ follow a *sequential elimination with ordering* S if there exist r stopping times $T_1 \leq \dots \leq T_r$ such that for every $k \in [r]$, every $\varepsilon, \varepsilon' > 0$, and every $T \geq T_k$,

$$|m_{i_k j_k}(\mathbf{X}_T)| \geq 1 - \varepsilon \quad \text{and} \quad |m_{i_k j}(\mathbf{X}_T)| \leq \varepsilon', |m_{i j_k}(\mathbf{X}_T)| \leq \varepsilon' \quad \text{for } i \neq i_k, j \neq j_k.$$

Based on Definition 1.6, we have the following result, which serves as a foundation for Theorem 1.5.

Theorem 1.7 (Theorem 1.5 revisited). Let $(\mathbf{X}_\ell)_{\ell \in \mathbb{N}}$ be online SGD initialized from $\mathbf{X}_0 \sim \mu_{N \times r}$ with step-size $\delta \ll \log(N)^{-1} N^{-\frac{p-3}{2}}$. Let $M \gg \log(N) N^{p-2}$. Then, the correlations $\{m_{ij}\}_{1 \leq i, j \leq r}$ follow a sequential elimination with ordering $\{(i_k^*, j_k^*)\}_{k=1}^r$ and stopping times of order $\log(N) N^{p-2}$, with \mathbb{P} -probability 1 in the large- N limit.

We can formulate an analogous version of Theorem 1.7 for Theorem 1.3 by adding the corresponding assumption on the ratio between the SNRs.

Remark 1.8. It is important to note that in the above results, namely Theorems 1.3, 1.5 and 1.7, the behavior of online SGD depends on the parity of p in the following sense. When p is odd, then each estimator $\mathbf{x}_{j_k^*}$ recovers the spike $\mathbf{v}_{i_k^*}$ with \mathbb{P} -probability $1 - o(1)$, since the correlations that are negative at initialization get trapped at the equator. Conversely, when p is even, we have that each estimator $\mathbf{x}_{j_k^*}$ recovers $\text{sgn}(m_{i_k^* j_k^*}(\mathbf{X}_0)) \mathbf{v}_{i_k^*}$ with probability $1 - o(1)$. This means that if the correlation at initialization is positive, then $\mathbf{x}_{j_k^*}$ recovers $\mathbf{v}_{i_k^*}$; otherwise, $\mathbf{x}_{j_k^*}$ recovers $-\mathbf{v}_{i_k^*}$.

1.2.2. MAIN RESULTS FOR $p = 2$. We now present our main results for $p = 2$. Our first result shows that exact recovery of all spikes is possible, provided sufficiently separated SNRs and a sample complexity of order $\log(N)^2 N^{\frac{\xi}{2}}$, where $\xi \in (0, 1)$ depends on the relative sizes of the signals.




FIGURE 2. Evolution of the correlations m_{ij} under the population dynamics in the case where $p = 3$, $r = 4$, $\lambda_1 = \dots = \lambda_4 = 1$. Since the SNRs are all equal, the order in which the correlations become macroscopic is determined by the values of the correlations at initialization, see Definition 1.4. This numerical simulation shows the recovery of a permutation of the four spikes $\mathbf{v}_1, \dots, \mathbf{v}_4$ according to the sequential elimination phenomenon.

Theorem 1.9 (Exact recovery for $p = 2$). *Let $(\mathbf{X}_\ell)_{\ell \in \mathbb{N}}$ be online SGD initialized from $\mathbf{X}_0 \sim \mu_{N \times r}$ with step-size $\delta \ll \frac{1}{\log(N)} N^{\frac{1}{2}} \frac{\lambda_r^2}{\lambda_1^2}$. Let $M \gg \log(N)^2 N^{\frac{1}{2}} \left(1 - \frac{\lambda_r^2}{\lambda_1^2}\right)$. Assume that $\lambda_i = \lambda_{i+1}(1 + \kappa_i)$ for every $1 \leq i \leq r - 1$, with $\kappa_i > 0$ being constants of order 1. Then, for every $\varepsilon > 0$ and every $i \in [r]$,*

$$\lim_{N \rightarrow \infty} \mathbb{P}_{\mathbf{X}_0} (|m_{ii}(\mathbf{X}_M)| \geq 1 - \varepsilon) = 1.$$

Similar to Theorem 1.3, variations in the correlations at initialization are compensated by differences in SNRs, leading to exact recovery of all spikes. However, unlike Theorem 1.3, for $p = 2$, exact recovery is achievable as soon as the SNRs differ by factors of order 1. In terms of sample complexity, there is a significant difference compared to the result obtained in [10] for the rank-one matrix PCA problem, where a number of $\log(N)^2$ samples is sufficient for efficiently recovering the spike. In our case, this sample complexity threshold is not sufficient to ensure the stability of the evolution of the subsequent correlations. The difference between both thresholds arises from the control of the retraction map, which is inherent to the online SGD algorithm on the Stiefel manifold. This is illustrated in the numerical simulation shown in Figure 8, and the term responsible for this discrepancy is rigorously quantified in the proof of Lemma 5.5.

As with the case of $p \geq 3$, the correlations exhibit a sequential elimination with ordering $\{(k, k) : k \in [r]\}$. Thus, an analogous version of Theorem 1.7 can be stated for Theorem 1.9. The evolution of the correlations is illustrated by the numerical simulations shown in Figures 3 and 4. As noted in Remark 1.8, the behavior of online SGD depends on the parity of the order p of the tensor. Specifically, for $p = 2$, the recovery of the unknown vector \mathbf{v}_i or its negative counterpart $-\mathbf{v}_i$ is determined by the sign of the correlation m_{ii} at initialization.

Our second main result for $p = 2$ addresses the case where the SNRs are all equal. Due to the invariance of the loss function $\mathcal{L}_{N,r}$ under both right and left rotations, our focus shifts from recovering each individual signal to recovering the subspace spanned by the signal vectors $\mathbf{v}_1, \dots, \mathbf{v}_r$. This notion of recovery is referred to as *subspace recovery* and is defined as follows.

Definition 1.10 (Subspace recovery). We say that we *recover the subspace* spanned by the orthogonal signal vectors $\mathbf{v}_1, \dots, \mathbf{v}_r$ if

$$\|\mathbf{X}\mathbf{X}^\top - \mathbf{V}\mathbf{V}^\top\|_F = o(1),$$

with high probability, where $\|\cdot\|_F$ stands for the Frobenius norm and $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_r]$, $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_r] \in \mathbb{R}^{N \times r}$.

From Definition 1.10, we note that subspace recovery is equivalent to recovering the eigenvalues of the matrix $\mathbf{G} \in \mathbb{R}^{r \times r}$ defined by

$$\mathbf{G} = \mathbf{M}\mathbf{M}^\top, \quad \text{with} \quad \mathbf{M} = \mathbf{V}^\top \mathbf{X} = (m_{ij}(\mathbf{X}))_{1 \leq i, j \leq r} \in \mathbb{R}^{r \times r},$$




FIGURE 3. Evolution of the correlations m_{ij} under the population dynamics in the case where $p = 2$, $r = 2$, $\lambda_1 = 3$ and $\lambda_2 = 1$. The SNRs are sufficiently separated to ensure exact recovery of both spikes, following the elimination of m_{12}, m_{21} by m_{11} .




FIGURE 4. Evolution of the correlations m_{ij} under the population dynamics in the case where $p = 2$, $r = 4$, $\lambda_1 = 10$, $\lambda_2 = 5$, $\lambda_3 = 2$ and $\lambda_4 = 1$. We observe a sequential elimination phenomenon, although less pronounced than in the case $p \geq 3$, in the sense that correlations that eventually become negligible can still exhibit significant increases.

since it holds that $\|\mathbf{X}\mathbf{X}^\top - \mathbf{V}\mathbf{V}^\top\|_F = 2(r - \text{Tr}(\mathbf{G}))$. Hereafter, we let $\theta_1(\mathbf{X}), \dots, \theta_r(\mathbf{X})$ denote the eigenvalues of the matrix-valued function $\mathbf{G}(\mathbf{X})$. Our main result is as follows.

Theorem 1.11 (Subspace recovery for $p = 2$). *Let $\lambda_1 = \dots = \lambda_r$. Let $(\mathbf{X}_\ell)_{\ell \in \mathbb{N}}$ be online SGD initialized from $\mathbf{X}_0 \sim \mu_{N \times r}$ with step-size $\delta \ll \frac{1}{\log(N)^2} N^{\frac{1}{2}}$. Let $M \gg \log(N)^3$. Then, for every $\varepsilon > 0$ and $i \in [r]$,*

$$\lim_{N \rightarrow \infty} \mathbb{P}_{\mathbf{X}_0}(\theta_i(\mathbf{X}_M) \geq 1 - \varepsilon) = 1.$$

The recovery dynamics in this case differs significantly from those in the previous cases: the correlations evolve at similar rates, thereby preventing the sequential elimination phenomenon, as shown in the left-hand side of Figure 5. However, as illustrated on the right-hand side of Figure 5, the eigenvalues of \mathbf{G} establish a natural ordering at initialization, which persists throughout their evolution until recovery, resulting in a monotone evolution.

Unlike the tensor case, where we analyze the dynamics for any value of the signal sizes, for $p = 2$ we have focused on two cases: when the parameters are separated by constants of order 1 (i.e., Theorem 1.9) and when they are all equal (i.e., Theorem 1.11). The intermediate case, where the SNRs are separated by a sufficiently small ϵ -factor, possibly depending on N , requires further effort and will be addressed in future work.

We conclude this section by presenting simulations of the high-dimensional online SGD algorithm in the case with three hidden directions. We refer to Figure 6 for the case where $p = 3$ and to Figure 7 for the case where $p = 2$. Additionally, Figure 8 illustrates the numerical instability observed for $p = 2$ due to the retraction map, as discussed in the paragraph following Theorem 1.9.




FIGURE 5. Evolution of the correlations m_{ij} (left) and eigenvalues (right) of $\mathbf{G} = \mathbf{M}\mathbf{M}^\top$ under the population dynamics. The relative speeds of the correlations are too close to each other for a sequential elimination phenomenon to take place. However, the eigenvalues of \mathbf{G} are monotone and it is possible to characterize subspace recovery.




FIGURE 6. Evolution of the correlations m_{ij} under the online SGD algorithm for the case where $p = 3, r = 3$, with SNRs $\lambda_1 = 3, \lambda_2 = 2, \lambda_3 = 1$. The simulation is performed with $M = 3200$ samples, a step-size $\delta/N = 0.0003$, and a dimension of $N = 500$. The simulation shows recovery of a permutation of the spikes.

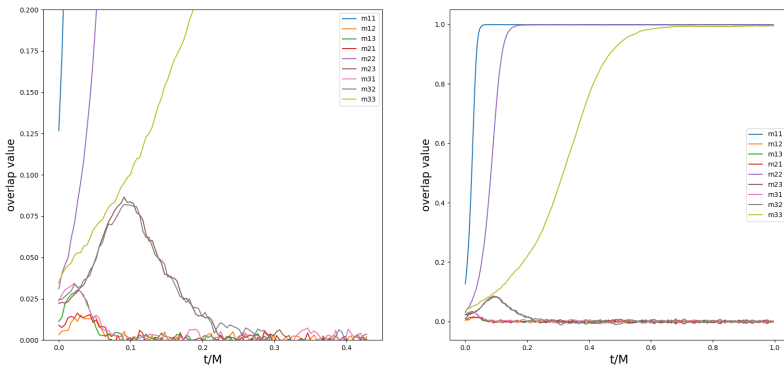


FIGURE 7. Evolution of the correlations m_{ij} under the online SGD algorithm for the case where $p = 2, r = 3$, with SNRs $\lambda_1 = 3, \lambda_2 = 2, \lambda_3 = 1$. The simulation is performed with $M = 230$ samples, a step-size $\delta/N = 0.002$, and a dimension of $N = 500$. The simulation shows exact recovery of the spikes.




FIGURE 8. Evolution of the correlations m_{ij} under the online SGD algorithm for the case where $p = 2, r = 2$, with SNRs $\lambda_1 = 2, \lambda_2 = 1$. The simulation is performed with $M = 70$ samples, a step-size $\delta/N = 0.006$, and a dimension of $N = 500$. While the correlations m_{11} and m_{22} clearly overcome the effect of the noise martingale and become macroscopic, the algorithm becomes unstable. This illustrates the fact that the sample complexity prescribed by Theorem 1.9 is larger than what is required to just ensure that the martingale term is negligible compared to the signal drift.

1.3. RELATED WORKS

We now provide a detailed overview on detection and recovery thresholds within the spiked tensor model, along with a discussion of main results for multi-index models.

Montanari and Richard [53] conjectured that for the single-spiked model, local algorithms with random initialization require a sample complexity scaling as N^{p-2} to efficiently recover the unknown spike, while a sharper threshold scaling as $N^{\frac{p-2}{2}}$ is sufficient using Sum-of-Squares (SoS) and spectral methods. The computational threshold N^{p-2} was achieved using various local methods: gradient flow and Langevin dynamics [9], online SGD [10], and tensor power iteration [37, 67]. For local algorithms, the order p of the tensor determines the curvature of the landscape near the initial condition, which influences the recovery difficulty. The sharp threshold of $N^{\frac{p-2}{2}}$ was achieved using SoS-based algorithms [36, 35, 40], spectral methods based on the Kikuchi Hessian [66], and tensor unfolding [13]. Beyond computational thresholds for efficiently estimating the unknown direction, significant research has focused on two other critical thresholds: the *information-theoretic threshold* for detection of the signal [43, 52, 51, 22, 38, 5], and the *statistical threshold* validating MLE as a reliable statistical method [14, 55, 38]. Detection and recovery thresholds for the multi-rank tensor PCA model have also been explored. On the information-theoretic side, it has been shown in [42] for $p = 2$ and in [23] for $p \geq 3$ that there is an order-1 critical threshold for the SNRs, above which it is possible to detect the unseen low-rank signal tensor $\sqrt{N} \sum_{i=1}^r \lambda_i \mathbf{v}_i^{\otimes p}$. On the algorithmic side, [37] studied the power iteration algorithm and determined the local threshold for efficiently recovering the finite-rank signal component. In the extended version of our preprint [6], we studied Langevin and gradient flow dynamics for the multi-spiked tensor PCA. We demonstrated that the sample complexity required to efficiently recover all spikes scales as N^{p-1} for $p \geq 3$, while for $p = 2$ as $N^{1+\delta-\lambda_r^2/\lambda_1^2}$ for every $\delta > 0$ not depending on N . These thresholds are suboptimal compared to those obtained here using online SGD due to challenges involved by controlling the noise generator under Langevin and gradient flow dynamics.

As mentioned earlier at the beginning of this section, the single/multi-spiked tensor model belongs to the broader class of single/multi-index models. In the context of single-index models, when the hidden link function σ^* is known, the behavior of online SGD exhibits similar phenomenology to the single-spiked tensor PCA model [10]. On the other hand, when the hidden link function σ^* is unknown, recent works by Bietti, Bruna, and Sanford [17], as well as by Berthier, Montanari, and Zhou [15], have studied the ability of two-layer neural networks to jointly learn the hidden link function and the hidden vector. They achieved this using various instances of gradient flow, either coupled with non-parametric regression steps or by employing a separation of timescales between the inner and outer layers. In the context of multi-index models, several recent works have provided upper bounds on the sample complexity thresholds required for shallow neural networks trained with first-order optimization methods to learn multi-index

models. For instance, [46, 25, 3] analyzed the benefits given by a single gradient step coupled with various choices for the second layer weights. Although these works determine the sample complexity required for the gradient to achieve macroscopic correlation with the target subspace and illustrate how neural networks improve over kernel methods, they offer limited insight into the global convergence when training the first layer weights. Using an online layer-wise training approach and regression steps on the second layer weights, Abbe, Boix-Adsera, and Misiakiewicz [2] introduced generalizations of the information exponent from [10] for multi-index functions, referred to as the *leap-exponent*, and proved that their dynamics successively visits several saddle-points before recovering the unknown subspace. A similar analysis was carried out by Dandi et al. [26], providing further sample complexity guarantees when using large batches of data over a few gradient steps. Moreover, Bietti, Bruna, and Pillaud-Vivien studied in [16] the population, continuous time limit of a two step procedure, where each gradient step on the estimator of the hidden subspace is coupled with a multivariate non-parametric regression step to estimate the link function. They showed that this two-step procedure makes the landscape benign, leading to the recovery of both the hidden link function and the correct subspace after *saddle-to-saddle dynamics* similar to the one exhibited in [2]. These works extensively rely on the separation of timescales and regression steps, resulting in a superposition of monotone dynamics, allowing for generality at the level of the link function while focusing on subspace recovery. In contrast, our work provides a complete description of the dynamics for online SGD in the context of a specific multi-index model, yielding recovery of permutations and rotations (when $p = 2$ and the SNRs are equal).

In the field of statistical physics of learning and the related results in high-dimensional probability, single- and multi-index models can be viewed as examples of the *teacher-student model*, see e.g. [29, 71], in which a specific target function is learned using a chosen architecture, typically with knowledge of the form of the hidden model. The dynamics of gradient-based methods are usually studied using closed-form, low-dimensional, and asymptotically exact dynamical systems towards which the gradient trajectories are shown to concentrate as the input dimension and number of sample diverge while remaining proportional. In the case of online learning, these dynamics lead to sets of ordinary differential equations (ODEs). For instance, such ODEs have been analyzed for committee machines [56] and two-layer neural networks [34, 63]. When a full-batch is used, the correlations throughout the entire trajectory result in a set of integro-differential equations, known as *dynamical mean-field theory* in statistical physics, for which there has been renewed interest in recent years in the context of classification problems and learning single and multi-index functions in the proportional limit of dimension and number of samples, using both powerful heuristics [57, 58, 45], and rigorous methods [21, 33]. Despite their theoretical significance, the resulting equations are often challenging to analyze beyond a few steps and are limited to studying linear sample complexity regimes. As a consequence, applying these tools to address questions related to global convergence, characterization of fixed points, or achieving polynomial sample complexity remains difficult.

Finally, the limiting dynamics of online SGD has recently been studied in [11, 12, 7] for broad range of estimation and classification tasks, including learning an XOR function using a two-layer neural network. In particular, [11, 12] shows the existence of critical regimes in the step-size that lead to an additional term in the population dynamics, referred to as the *corrector*, which also appears in the statistical physics literature, e.g. [56, 34]. By appropriately rescaling the step size, they also obtain diffusive limits for online SGD. In [7], a spectral analysis is performed to study the evolution of the Hessian of the loss landscape, showing alignment of the SGD trajectories with emerging eigenspaces. We defer the exploration of critical step sizes, diffusive limits, and spectral phenomena to future work, focusing here on obtaining non-asymptotic results and enhancing our understanding of the population dynamics.

1.4. OVERVIEW

An overview of the paper is given as follows. Section 2 provides an outline of the proofs of our main results. In Section 3, we present the nonasymptotic versions of our main results, from which we derive the main asymptotic results of Subsection 1.2. Sections 4-6 are devoted to the proofs of the main results in their nonasymptotic form. Specifically, Section 4 presents comparison inequalities for both the correlations $m_{ij}(\mathbf{X})$ and the eigenvalues of $\mathbf{M}\mathbf{M}^\top$ with $\mathbf{M} = (m_{ij}(\mathbf{X}))_{i,j \in [r]}$. Section 5 provides the proofs of the main results related to the full recovery of the spikes, i.e., it addresses the case where $p \geq 3$ with arbitrary values for the SNRs, as well as the case where $p = 2$ with sufficiently separated SNRs. Finally, Section 6 focuses on the case where $p = 2$ and the spikes have equal SNRs, thereby proving subspace

recovery.

Acknowledgements. G.B. and C.G. acknowledge the support of the NSF grant DMS-2134216. V.P. acknowledges the support of the ERC Advanced Grant LDRaM No. 884584.

2. OUTLINE OF PROOFS

In this section, we outline the proofs of our main results. To simplify the discussion, we focus on the spiked tensor model with $r = 2$. First, we address the cases where $p \geq 3$ with arbitrary values of λ_1 and λ_2 , as well as the case where $p = 2$ with sufficiently separated SNRs. In these cases, we follow the evolution of the correlations under the online SGD algorithm and show full recovery of the spikes. Next, we consider the case where $p = 2$ and the spikes have equal strength, i.e., $\lambda_1 = \lambda_2$, and show that the algorithm recovers the subspace spanned by the two spikes.

2.1. FULL RECOVERY OF SPIKES

We begin by analyzing the evolution of the correlations m_{ij} under the online SGD algorithm, following the update rule given by (1.6). We assume an initial random start with a completely uninformative prior, specifically the invariant distribution on the Stiefel manifold. As a consequence, all correlations m_{ij} have the typical scale of order $N^{-\frac{1}{2}}$ at initialization. For simplicity, we assume that all correlations are positive at initialization.

According to (1.6), the evolution equation for the correlations $m_{ij}(\mathbf{X}_t)$ at time t is given by

$$m_{ij}(\mathbf{X}_t) = \left\langle \mathbf{v}_i, R_{\mathbf{X}_{t-1}} \left(-\frac{\delta}{N} (\nabla_{\text{St}} \mathcal{L}_{N,r}(\mathbf{X}_{t-1}; \mathbf{Y}^t))_j \right) \right\rangle.$$

To simplify notation, we write $m_{ij}(t) = m_{ij}(\mathbf{X}_t)$. When the step-size $\delta > 0$ is sufficiently small, the matrix inverse in the polar retraction term $R_{\mathbf{X}_{t-1}}$ (see (1.7)) can be expanded as a Neumann series. Thus, up to controlled error terms, we have that

$$m_{ij}(t) \approx m_{ij}(t-1) - \langle \mathbf{v}_i, \frac{\delta}{N} (\nabla_{\text{St}} \mathcal{L}_{N,r}(\mathbf{X}_{t-1}; \mathbf{Y}^t))_j \rangle.$$

Iterating this expression yields

$$m_{ij}(t) \approx m_{ij}(0) - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} \mathcal{L}_{N,r}(\mathbf{X}_{\ell-1}))_j \rangle. \quad (2.1)$$

The condition on the step-size δ required for the Neumann series to converge is not restrictive and does not present a bottleneck in our proof. However, the retraction map introduces additional interaction terms among the correlations, which demand a careful analysis.

Having the evolution equations for the correlations at hand, we decompose the gradient of the loss function $\mathcal{L}_{N,r}$ into noise and population components, i.e.,

$$\nabla_{\text{St}} \mathcal{L}_{N,r}(\mathbf{X}_{t-1}, \mathbf{Y}^t) = \nabla_{\text{St}} H_{N,r}(\mathbf{X}_{t-1}) + \nabla_{\text{St}} \Phi_{N,r}(\mathbf{X}_{t-1}),$$

where $H_{N,r}$ and $\Phi_{N,r}$ are given by (1.4) and (1.5), respectively. By balancing these two terms, we determine the sample complexity thresholds for efficient recovery. Following the approach used to prove [10, Proposition 4.1], we control the magnitude of the noise part by appropriately adjusting the step-size δ , ensuring the signal term dominates over a sufficiently large time scale for efficient recovery. However, unlike the single-spiked case in [10, Proposition 4.1], the contribution from the population loss leads to a significantly more complicated dynamics, requiring careful analysis to control the nonmonotonic interactions among the correlations. At the same time, our control needs to be sufficiently sharp to meet sample complexities and discretization guarantees that are statistically and algorithmically relevant.

2.1.1. CONTROL OF THE NOISE TERM. The evolution equation (2.1) can be rewritten as

$$m_{ij}(t) \approx m_{ij}(0) - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi_{N,r}(\mathbf{X}_{\ell-1}))_j \rangle - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} H_{N,r}(\mathbf{X}_{\ell-1}))_j \rangle,$$

where a simple computation shows that

$$\begin{aligned} -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi_{N,r}(t))_j \rangle &= p\sqrt{N} \lambda_i \lambda_j m_{ij}^{p-1}(t) \\ &\quad - \frac{p}{2} \sqrt{N} \sum_{1 \leq k, \ell \leq 2} \lambda_k m_{i\ell}(t) m_{kj}(t) m_{k\ell}(t) \left(\lambda_j m_{kj}^{p-2}(t) + \lambda_\ell m_{k\ell}^{p-2}(t) \right). \end{aligned}$$

The noise term corresponds to the sum of the martingale increments $-\langle \mathbf{v}_i, (\nabla_{\text{St}} H_{N,r}(\mathbf{X}_{\ell-1}))_j \rangle$, and since the noise tensors $(\mathbf{W}^\ell)_\ell$ are sub-Gaussian, classical concentration inequalities for the sub-Gaussian distribution can be applied to bound this term. The population part consists of two terms: a drift term $p\sqrt{N} \lambda_i \lambda_j m_{ij}^{p-1}$ which dominates the dynamics, particularly near initialization, and a correction term $\frac{p}{2} \sqrt{N} \sum_{1 \leq k, \ell \leq r} \lambda_k m_{i\ell} m_{kj} m_{k\ell} \left(\lambda_j m_{kj}^{p-2} + \lambda_\ell m_{k\ell}^{p-2} \right)$ which arises due to the constraint of the estimator \mathbf{X} being on an orthogonal manifold. For the correlation m_{ij} to increase, the drift term needs to be larger than both the noise and correction terms. Near initialization, the correlations typically scale as $N^{-\frac{1}{2}}$ and the drift term dominates the correction term. Moreover, near initialization, the noise term can be absorbed by $m_{ij}(0)$. This gives that

$$m_{ij}(t) \approx m_{ij}(0) + \frac{\delta}{\sqrt{N}} p \lambda_i \lambda_j \sum_{\ell=1}^t m_{ij}^{p-1}(\ell - 1). \quad (2.2)$$

At this point, we need to show that the drift term continues to dominate the noise term over a sufficiently long time scale, allowing the correlations m_{ij} to *escape mediocrity*. Once the first correlation reaches a critical threshold, the previously negligible correction term in the population loss becomes significant, and the dynamics shifts to a more delicate phase where the correlations begin to interact with one another, as explained below.

2.1.2. ANALYSIS OF THE POPULATION DYNAMICS. We now analyze the evolution of the correlations by studying the population dynamics, first for the case $p \geq 3$, and then for $p = 2$.

Recovery for $p \geq 3$. The solution of (2.2) is given by

$$m_{ij}(t) \approx m_{ij}(0) \left(1 - \frac{\delta}{\sqrt{N}} p \lambda_i \lambda_j m_{ij}(0)^{p-2} t \right)^{-\frac{1}{p-2}}, \quad (2.3)$$

where $m_{ij}(0) = \frac{\gamma_{ij}}{\sqrt{N}}$. The typical time for m_{ij} to reach a macroscopic threshold $\varepsilon > 0$, denoted by $T_\varepsilon^{(ij)}$, is therefore given by

$$T_\varepsilon^{(ij)} \approx \frac{1 - \left(\frac{\gamma_{ij}}{\varepsilon \sqrt{N}} \right)^{p-2}}{\delta \lambda_i \lambda_j p \gamma_{ij}^{p-2}} N^{\frac{p-2}{2}}.$$

In particular, the first correlation to become macroscopic is the one associated with the largest value of $\lambda_i \lambda_j \gamma_{ij}^{p-2}$, where $(\gamma_{ij})_{1 \leq i, j \leq 2}$ approximately follow independent standard normal distributions. According to Definition 1.4, this corresponds to the pair (i_1^*, j_1^*) of the greedy maximum selection of the matrix $\mathbf{I}_0 = (\lambda_i \lambda_j m_{ij}(0))^{p-2} \mathbf{1}_{\{m_{ij}^{p-2}(0) \geq 0\}}_{ij}$. Once the first correlation m_{ij} reaches a macroscopic value $\varepsilon > 0$, the other correlations are still trapped near the equator. This can be verified by plugging the value of $T_\varepsilon^{(i_1^* j_1^*)}$ into (2.3). Note that if the SNRs are sufficiently separated to compensate for initialization differences, with high probability, m_{11} is the first correlation to escape mediocrity. In the following, without loss of generality we may assume that $i_1^* = j_1^* = 1$, i.e., m_{11} is the first correlation to reach the macroscopic threshold ε .

Next, we describe the evolution of m_{12}, m_{21} and m_{22} as $t \geq T_\varepsilon^{(11)}$. We first note that the correction part of the population term, i.e.,

$$\sum_{1 \leq k, \ell \leq 2} \lambda_k m_{kj}(t) m_{k\ell}(t) m_{i\ell}(t) \left(\lambda_j m_{kj}^{p-2}(t) + \lambda_\ell m_{k\ell}^{p-2}(t) \right),$$

becomes dominant in the evolution of the correlations m_{12}, m_{21} , as soon as m_{11} reaches the microscopic threshold $N^{-\frac{p-2}{2p}}$. In other words, m_{12} and m_{21} are non-decreasing and they begin to decrease as soon as m_{11} reaches this microscopic threshold. On the other hand, the correction part of the population generator may become dominant in the evolution of m_{22} , as soon as m_{11} reaches another microscopic threshold of order $N^{-\frac{p-3}{2(p-1)}}$, leading to a potential decrease of m_{22} . Through a careful analysis, we show that this potential decrease is at most of order $\frac{\log(N)}{N}$, so that m_{22} remains stable at the scale

$\frac{1}{\sqrt{N}}$ during the ascending phase of m_{11} . The evolution of the correlations m_{12}, m_{21} and m_{22} as m_{11} becomes macroscopic is illustrated by the numerical experiment in Figure 9. Once m_{12}, m_{21} become sufficiently small, we show that the evolution of m_{22} can still be described by (2.3), thus ensuring the recovery of the second direction. The phenomenon we observe here is referred to as *sequential elimination phenomenon*, as introduced in Definition 1.6. The correlations increase sequentially, one after another. When a correlation (e.g., m_{11}) exceeds a certain threshold, the correlations sharing a row and column index (e.g., m_{12}, m_{21}) start decreasing until they become sufficiently small to be negligible, thereby enabling the subsequent correlations (e.g., m_{22}) to increase. This phenomenon is illustrated in Figure 1. We remark that if the SNRs are sufficiently separated as described earlier, then exact recovery of the unknown signal vectors is achieved with high probability, which depends on the ratio between the signal sizes, as stated in Theorem 1.3. Otherwise, recovery results in a permutation of the spikes, determined by the greedy maximum selection of the matrix $\mathbf{I}_0 = (\lambda_i \lambda_j m_{ij}(0)^{p-2} \mathbf{1}_{\{m_{ij}^{p-2}(0) \geq 0\}})_{ij}$, as illustrated in Figure 2.




FIGURE 9. Evolution of the correlations m_{ij} under the population dynamics in the case where $p = 3$, $r = 2$, $\lambda_1 = 3$ and $\lambda_2 = 1$. The SNRs are sufficiently separated to ensure exact recovery of both spikes. The figure consists of four subfigures, each zooming in on the evolution of one of the four correlations from initialization to the moment when the first correlation m_{11} becomes macroscopic. In particular, m_{12} and m_{21} decrease rapidly once m_{11} becomes macroscopic, whereas m_{22} remains stable overall. The short decrease in m_{22} is bounded by $\frac{\log(N)}{N}$ with high probability.

To determine the sample complexity threshold under online SGD algorithm, we control the retraction and martingale terms so that the discrete dynamics is sufficiently close to the continuous one studied in [6] over a time scale of order $\frac{\sqrt{N}}{\delta} N^{\frac{p-2}{2}}$. We show that this requires a step-size parameter of order $\log(N)^{-1} N^{-\frac{p-3}{2}}$, leading to a sample complexity of order $\log(N) N^{p-2}$. This threshold is required for efficient recovery of both spikes.

Recovery for $p = 2$. In this case, the solution of (2.2) is given by

$$m_{ij} \approx m_{ij}(0) \exp\left(2\delta\lambda_i\lambda_j N^{-\frac{1}{2}}t\right) = \frac{\gamma_{ij}}{\sqrt{N}} \exp\left(2\delta\lambda_i\lambda_j N^{-\frac{1}{2}}t\right), \quad (2.4)$$




FIGURE 10. Evolution of the correlations $\overline{m_{ij}}$ under the population dynamics in the case where $p = 2, r = 2, \lambda_1 = 2$ and $\lambda_2 = 1$. We observe that as m_{11} reaches a certain threshold, the correlations m_{12} and m_{21} start to decrease, despite rising higher than in the case where $p \geq 3$.

so that the typical time for m_{ij} to reach a macroscopic threshold $\varepsilon > 0$ is given by

$$T_\varepsilon^{(ij)} \approx \frac{\log\left(\frac{\varepsilon\sqrt{N}}{\gamma_{ij}}\right)}{2\delta\lambda_i\lambda_j}\sqrt{N}.$$

In contrast to the case where $p \geq 3$, where the quantity $\lambda_i\lambda_j\gamma_{ij}^{p-2}$ determines which correlation first escapes mediocrity, the initialization γ_{ij} has less influence here and the time differences for the correlations to reach a macroscopic threshold are determined by the separation between the SNRs. When $\lambda_1 > \lambda_2$, a sequential elimination phenomenon akin to the one observed for $p \geq 3$ occurs (see Figures 3 and 4). However, there is an important difference compared to the tensor case: once m_{11} reaches the macroscopic threshold $\varepsilon > 0$, the correlations m_{12}, m_{21} and m_{22} scale as $N^{-\delta_1/2}, N^{-\delta_1/2}$ and $N^{-\delta_2/2}$, respectively, where $\delta_1 = 1 - \frac{\lambda_2}{\lambda_1}$ and $\delta_2 = 1 - \frac{\lambda_2^2}{\lambda_1^2}$. We then show that m_{11} must reach approximately the critical value $(\frac{\lambda_2}{\lambda_1})^{1/2}$ for m_{12}, m_{21} to decrease. This value is achieved in an order one time once m_{11} has reached ε , while m_{12}, m_{21} still require a time of order $\log(N)$ to escape their scale of $N^{-\frac{\delta_1}{2}}$. In addition, we show the stability of m_{22} during this first phase by controlling that it remains increasing throughout the ascending phase of m_{11} and until m_{12}, m_{21} begin to decrease. This stability comes with an order 1 lower bound on the logarithm of the dimension (see Theorem 3.6), involving the inverse of the gap between the SNRs, so that the case where the SNRs are close becomes sensitive to finite size effects. Weak recovery of the second direction is then possible as m_{12}, m_{21} descend below the scale $\frac{1}{\sqrt{N}}$. This is illustrated in Figure 10.

Regarding the sample complexity, similar to the case $p \geq 3$, we control the martingale and retraction terms so that the discrete dynamics is sufficiently close to the continuous one studied in [6] over a time scale of order $\frac{\sqrt{N}}{\delta} \log(N)$. We show that this requires a step-size δ of order $\log(N)^{-1} N^{\lambda_r^2/(2\lambda_1^2)}$, leading

to a sample complexity for exact recovery of order $\log(N)^2 N^{\frac{1}{2}(1-\lambda_r^2/\lambda_1^2)}$. The factor $N^{\lambda_r^2/(2\lambda_1^2)}$ is required to control the stability of the retraction map, which is necessary for all directions, including the first one.

2.2. SUBSPACE RECOVERY

As discussed in Subsection 1.2, when $\lambda_1 = \lambda_2 \equiv \lambda$, the spiked covariance model becomes isotropic and we shift our focus to the evolution of the eigenvalues of the matrix $\mathbf{M}\mathbf{M}^\top$, which we denote by \mathbf{G} . To analyze the evolution of the eigenvalues, we use eigenvalue perturbation bounds and matrix generalizations of martingale concentration inequalities. In particular, once the noise is controlled, the population dynamics take the form

$$\mathbf{G}(t) \approx \mathbf{G}(0) + \frac{4\delta\lambda^2}{\sqrt{N}} \sum_{\ell=1}^t \mathbf{G}(\ell-1)(\mathbf{I} - \mathbf{G}(\ell-1)),$$

leading to the following evolution equation for the eigenvalues $\theta_1, \dots, \theta_r$,

$$\theta_i(t) \approx \theta_i(0) + \frac{4\delta\lambda^2}{\sqrt{N}} \sum_{\ell=1}^t \theta_i(\ell-1)(1 - \theta_i(\ell-1)).$$

This equation then prescribes a time horizon of order $\log(N)$ for the various bounds needed to control the noise process, leading to the sample complexities as stated in Theorem 1.11. An important difference with the full recovery case is that, at initialization, the eigenvalues of \mathbf{G} are of order $\frac{1}{N}$. Simply adapting the bounds derived for the correlations m_{ij} is not sufficient here. In particular, we develop time-dependent bounds for the martingale terms using a matrix generalization of martingale concentration inequalities [61].

3. MAIN RESULTS

This section presents the nonasymptotic versions of our main results stated in Subsection 1.2.

As previously discussed, we consider the online SGD algorithm starting from a random initialization, precisely $\mathbf{X}_0 \sim \mu_{N \times r}$, where $\mu_{N \times r}$ denotes the uniform measure on the Stiefel manifold $\text{St}(N, r)$. Our recovery guarantees, however, applies to a broader class of initial data that meets two natural conditions, that we introduce hereafter. The first condition ensures that the initial correlation is on the typical scale $\Theta(N^{-\frac{1}{2}})$.

Definition 3.1 (Condition 1). For every $\gamma_1 > \gamma_2 > 0$, we let $\mathcal{C}_1(\gamma_1, \gamma_2)$ denote the sequence of events given by

$$\mathcal{C}_1(\gamma_1, \gamma_2) = \left\{ \mathbf{X} \in \text{St}(N, r) : \frac{\gamma_2}{\sqrt{N}} \leq m_{ij}(\mathbf{X}) < \frac{\gamma_1}{\sqrt{N}} \text{ for every } 1 \leq i, j \leq r \right\}.$$

We say that $\mathbf{X} \in \text{St}(N, r)$ satisfies *Condition 1* if for every $\gamma_1 > \gamma_2 > 0$,

$$\mathbb{P}(\mathbf{X} \in \mathcal{C}_1(\gamma_1, \gamma_2)^c) \leq C_1 e^{-c_1 \gamma_1^2} + C_2 e^{c_2 \gamma_2 \sqrt{N}} + C_3 \gamma_2,$$

for absolute constants $C_1, c_1, C_2, c_2, C_3 > 0$.

In the special case $p = 2$ and $\lambda_1 = \dots = \lambda_r$, we analyze the evolution of the eigenvalues of the matrix $\mathbf{G} = \mathbf{M}\mathbf{M}^\top$, where $\mathbf{M} = (m_{ij})_{1 \leq i, j \leq r}$. To ensure that the eigenvalues at initialization are on the typical scale of $\Theta(N^{-1})$, we require a modified version of Condition 1 from Definition 3.1.

Definition 3.2 (Condition 1'). For every $\gamma_1 > \gamma_2 > 0$, we let $\mathcal{C}'_1(\gamma_1, \gamma_2)$ denote the sequence of events given by

$$\mathcal{C}'_1(\gamma_1, \gamma_2) = \left\{ \mathbf{X} \in \text{St}(N, r) : \frac{\gamma_2}{N} \leq \theta_i(\mathbf{X}) < \frac{\gamma_1}{N} \text{ for every } i \in [r] \right\}.$$

We say that $\mathbf{X} \in \text{St}(N, r)$ satisfies *Condition 1'* if for every $\gamma_1 > \gamma_2 > 0$,

$$\mathbb{P}(\mathbf{X} \in \mathcal{C}'_1(\gamma_1, \gamma_2)^c) \leq C_1 e^{-c_1 \gamma_1^2} + C_2 e^{-c_2 \gamma_2 \sqrt{N}} + C_3 \gamma_2,$$

for absolute constants $C_1, c_1, C_2, c_2, C_3 > 0$.

The second condition concerns the separation of the initial correlations, which is needed for the recovery when $p \geq 3$.

Definition 3.3 (Condition 2). For every $\gamma_1 > \gamma > 0$, we let $\mathcal{C}_2(\gamma_1, \gamma)$ denote the sequence of events given by

$$\mathcal{C}_2(\gamma_1, \gamma) = \left\{ \mathbf{X} \in \text{St}(N, r) : \left| \frac{\lambda_i \lambda_j m_{ij}^{p-2}(\mathbf{X})}{\lambda_k \lambda_\ell m_{k\ell}^{p-2}(\mathbf{X})} - 1 \right| > \frac{\gamma}{\gamma_1} \text{ for every } 1 \leq i, j, k, \ell \leq r \right\}.$$

We say that $\mathbf{X} \in \text{St}(N, r)$ satisfies *Condition 2* if for every $\gamma_1 > \gamma > 0$,

$$\mathbb{P}(\mathbf{X} \in \mathcal{C}_2(\gamma_1, \gamma)^c) \leq C_1 e^{-c_1 \gamma_1^2} + C_2 e^{-c_2 \sqrt{N} \gamma} + C_3 \left(1 + \left(\frac{\lambda_i \lambda_j}{\lambda_k \lambda_\ell} \right)^{\frac{2}{p-2}} \right)^{-\frac{1}{2}} \gamma,$$

for absolute constants $C_1, c_1, C_2, c_2, C_3 > 0$.

A random matrix \mathbf{X} uniformly distributed on $\text{St}(N, r)$ satisfies the above conditions, as stated in the following lemma.

Lemma 3.4. *Every random matrix $\mathbf{X} \sim \mu_{N \times r}$ satisfies Condition 1, Condition 1', and Condition 2.*

The proof of Lemma 3.4 is deferred to Appendix A of the extended version of our preprint [6]. In particular, in [6] we consider the normalized Stiefel manifold $\mathcal{M}_{N,r} = \{\mathbf{X} \in \mathbb{R}^{N \times r} : \mathbf{X}^\top \mathbf{X} = N \mathbf{I}_r\}$ and define correlations $m_{ij}(\mathbf{X})$ by $m_{ij}^{(N)}(\mathbf{X}) = \frac{\langle \mathbf{v}_i, \mathbf{x}_j \rangle}{N}$, where $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_r] \in \mathcal{M}_{N,r}$. We then show that the invariant measure on $\mathcal{M}_{N,r}$ satisfies the three aforementioned conditions. This framework is equivalent to the one considered here, and therefore, the proofs in Appendix A of [6] can be straightforwardly generalized to $\mu_{N \times r}$.

We are now in the position to state our main results in a nonasymptotic form. Let T denote the number of iterations after which we terminate the online SGD algorithm. For any set \mathcal{E} , let $\mathcal{T}_{\mathcal{E}}$ denote the hitting time of the set \mathcal{E} , i.e.,

$$\mathcal{T}_{\mathcal{E}} = \min\{t \geq 0 : \mathbf{X}_t \in \mathcal{E}\}.$$

Additionally, we point out that the constants K_1, K_2, \dots and c_1, c_2, \dots appearing in the following statements may depend on $\sigma^2, p, r, \{\lambda_i\}_{i=1}^r$.

Our first main result shows that for $p \geq 3$, the sample complexity threshold for efficiently recovering a permutation of all spikes $\mathbf{v}_1, \dots, \mathbf{v}_r$ must scale as $\log(N)N^{p-2}$.

Theorem 3.5 (Recovery of a permutation of all spikes for $p \geq 3$). *Let $p \geq 3$ and $\lambda_1 \geq \dots \geq \lambda_r$. Consider the sequence of online SGD outputs $(\mathbf{X}_t)_{t \in \mathbb{N}}$ started from \mathbf{X}_0 with step-size δ . Let $(i_1^*, j_1^*), \dots, (i_r^*, j_r^*)$ denote the greedy maximum selection of \mathbf{I}_0 (see (1.9) and Definition 1.4). Further, for every $\varepsilon > 0$, let $R(\varepsilon)$ denote*

$$R(\varepsilon) = \left\{ \mathbf{X} : m_{i_k^* j_k^*}(\mathbf{X}) \geq 1 - \varepsilon \text{ for all } k \in [r] \text{ and } m_{ij}(\mathbf{X}) \lesssim \log(N)^{-\frac{1}{2}} N^{-\frac{p-1}{4}} \text{ otherwise} \right\}.$$

Then, the following holds. For every $\gamma_1 > \gamma_2, \gamma > 0$ and every $\varepsilon > 0$, there exist $c_0 \in (0, 1)$ and $d_0 = d_0(N)$ satisfying $d_0 \rightarrow 0$ as $N \rightarrow \infty$ such that if $\mathbf{X}_0 \in \mathcal{C}_1(\gamma_1, \gamma_2) \cap \mathcal{C}_2(\gamma_1, \gamma)$ and $\delta > 0$ verifies

$$\delta = C_\delta d_0 \log(N)^{-1} N^{-\frac{p-3}{2}},$$

for a constant $C_\delta > 0$, then

$$\mathbb{P}_{\mathbf{X}_0}(\mathcal{T}_{R(\varepsilon)} \lesssim T) \geq 1 - \eta,$$

where T and η are given by

$$T = \frac{1 - \left(\frac{\gamma_2}{\varepsilon \sqrt{N}}\right)^{p-2}}{C_\delta (1 - c_0)^{p-1} d_0 \gamma_2^{p-2}} \log(N) N^{p-2} + \frac{1 - \left(\frac{\varepsilon}{1 - \varepsilon}\right)^{p-2}}{C_\delta (1 - c_0)^{p-1} d_0 \varepsilon^{p-1}} \log(N) N^{\frac{p-3}{2}},$$

and

$$\begin{aligned} \eta &= K_1 e^{-c_1 \gamma_1^2} + K_2 e^{-c_2 (\gamma_2 + \gamma) \sqrt{N}} + K_3 (\gamma_2 + \gamma) \\ &+ K_4 \frac{\log(N)^2 N^{p-2}}{d_0} e^{-c_4 N} + K_5 \log(N) e^{-c_5 c_0^2 \log(N)/d_0} + K_6 \frac{d_0^2}{\log(N)} N^{-\frac{p-2}{2}} e^{-c_6/d_0^2}, \end{aligned}$$

respectively.

Our second main result is for $p = 2$ with sufficiently separated SNRs. We show that the sample complexity threshold for strong recovery of all spikes must diverge as $\log(N)^2 N^{\frac{1}{2}} \left(1 - \frac{\lambda_r^2}{\lambda_1^2}\right)$.

Theorem 3.6 (Exact recovery of all spikes for $p = 2$). *Let $p = 2$ and $\lambda_1 \geq \dots \geq \lambda_r$ such that $\lambda_i = \lambda_{i+1}(1 + \kappa_i)$ for every $1 \leq i \leq r - 1$, with $\kappa_i > 0$ being constants of order 1. Let κ denote $\kappa = \min_{1 \leq i \leq r-1} \kappa_i$. Consider the sequence of online SGD outputs $(\mathbf{X}_t)_{t \in \mathbb{N}}$ started from \mathbf{X}_0 with step-size δ . Further, for every $\varepsilon > 0$ and $c_0 \in (0, \frac{1}{2} \wedge \frac{\kappa}{2+\kappa})$, let $R(\varepsilon, c_0)$ denote*

$$R(\varepsilon, c_0) = \left\{ \mathbf{X} : m_{ii}(\mathbf{X}) \geq 1 - \varepsilon \text{ for all } i \in [r] \text{ and } m_{k\ell}(\mathbf{X}) \lesssim N^{-\frac{1}{2}} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right) \text{ for all } k \neq \ell \right\}.$$

Then, the following holds. For every $\gamma_1 > \gamma_2 > 0$ and every $\varepsilon, \varepsilon' > 0$, there exist $c_0 \in (0, 1)$ and $d_0 = d_0(N)$ satisfying $d_0 \rightarrow 0$ as $N \rightarrow \infty$ such that if $\mathbf{X}_0 \in \mathcal{C}_1(\gamma_1, \gamma_2)$ and $\delta > 0$ verifies

$$\delta = \frac{C_\delta d_0 N^{\frac{1}{2} \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}}}{\log\left(\frac{2\varepsilon\sqrt{N}}{\gamma_2}\right)},$$

for a constant $C_\delta > 0$, and if $\log(N) \gtrsim \frac{1}{1+\kappa}$, then

$$\mathbb{P}_{\mathbf{X}_0}(\mathcal{T}_{R(\varepsilon, c_0)} \lesssim T) \geq 1 - \eta,$$

where T and η are given by

$$T = \frac{N^{\frac{1}{2}} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right) \log\left(\frac{2\varepsilon\sqrt{N}}{\gamma_2}\right)^2}{C_\delta(1-c_0)d_0\gamma_2^2},$$

and

$$\begin{aligned} \eta &= K_1 e^{-c_1 \gamma_1^2} + K_2 e^{-c_2 \gamma_2 \sqrt{N}} + K_3 \gamma_2 + K_4 \frac{\log(N)^2 N^{\frac{1}{2}} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right)}{d_0} e^{-c_4 N} \\ &\quad + K_5 e^{-c_5 c_0^2 N^{\frac{1}{2}} \left(\frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right) / d_0} + K_6 \frac{d_0^2 N^{-\left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right)}}{c_0 \log(N)} e^{-c_6 / d_0^2}, \end{aligned}$$

respectively.

We finally state our result when the SNRs are all equal. As discussed in Subsection 1.1, we focus here on the evolution of the eigenvalues θ_i of the symmetric matrix $\mathbf{M}\mathbf{M}^\top \in \mathbb{R}^{r \times r}$, where $\mathbf{M} = (m_{ij})_{1 \leq i, j \leq r}$ is the correlation matrix.

Theorem 3.7 (Subspace recovery for $p = 2$). *Let $p = 2$ and $\lambda_1 = \dots = \lambda_r \equiv \lambda$. Consider the sequence of online SGD outputs $(\mathbf{X}_t)_{t \in \mathbb{N}}$ started from \mathbf{X}_0 with step-size δ . Then, the following holds. For every $\gamma_1 > \gamma_2 > 0$ and every $\varepsilon, \varepsilon' > 0$, there exists $d_0 = d_0(N)$ satisfying $d_0 \rightarrow 0$ as $N \rightarrow \infty$ such that if $\mathbf{X}_0 \in \mathcal{C}'_1(\gamma_1, \gamma_2)$ and $\delta > 0$ verifies*

$$\delta = C_\delta \frac{d_0 \sqrt{N}}{\log(N)^2},$$

for a constant $C_\delta > 0$, then for all $i \in [r]$,

$$\mathbb{P}_{\mathbf{X}_0}(\theta_i(\mathbf{X}_T) \geq 1 - \varepsilon') \geq 1 - \eta,$$

where T and η are given by

$$T = \frac{\log(N)^2 \log\left(\frac{4\varepsilon N}{\gamma_2}\right)}{C_\delta d_0 \lambda^2} + \frac{\log(N)^2 \log\left(\frac{1-\varepsilon'}{\varepsilon}\right)}{C_\delta d_0 \lambda^2 \varepsilon'},$$

and

$$\eta = K_1 e^{-c_1 \gamma_1^2} + K_2 e^{-c_2 \gamma_2 \sqrt{N}} + K_3 \gamma_2 + K_4 \frac{\log(N)^3}{d_0} e^{-c_4 N} + K_5 e^{-c_5 / d_0} + K_6 \frac{\log(N)^3}{N d_0},$$

respectively.

The proofs of the above results follow similar arguments to those in [9]. The two main steps involve controlling both the retraction term and the first-order term, which carries the phenomenology of the problem for a sufficiently small step size. Specifically, the first-order term is decomposed into a drift term and a martingale term, with the martingale term required to be negligible compared to the initialization. This requirement, along with the equivalence of the number of steps and the number of data points

in online SGD, determines the step-size parameter and the sample complexity threshold. Compared to the single-spiked tensor model studied in [10], our primary challenge is that the retraction term is more complex to control. We show that it can be controlled by expanding the matrix inverse in the polar retraction as a Neumann series. Furthermore, under the chosen scaling, the signal term is larger than in [9], thus necessitating a more careful control of the higher-order signal terms in the expansion of the retraction term. In particular, some of these terms cannot be compared to the initialization scale to achieve the desired comparison inequalities. We thus compare them to the first-order signal part, up to a truncation threshold, where the upper part decays exponentially fast and can be compared to the initialization. Our second difference with respect to [10] is related to the control of the first-order noise term. In our case, under the sub-Gaussian assumption of the order- p tensor, exponential forms of Doob's inequality may be used [44, 65], resulting in faster convergence rates.

Before addressing the proof of the three main nonasymptotic results, we first provide the proofs of their asymptotic counterparts, namely Theorems 1.3, 1.5, 1.9 and 1.11.

Proof of Theorems 1.3, 1.5, 1.9 and 1.11. According to Lemma 3.4, the uniform measure $\mu_{N \times r}$ on $\text{St}(N, r)$ satisfies Condition 1 and Condition 2. In particular, we note that for every $\gamma_1 > \gamma_2 > 0$, $\mu_{N \times r}(\mathcal{C}_1(\gamma_1, \gamma_2))$ converges towards one in the large- N limit, provided γ_1 diverges arbitrarily slowly and γ_2 decreases to zero arbitrarily slowly. Similarly, for every $\gamma_1 > \gamma > 0$, $\mu_{N \times r}(\mathcal{C}_2(\gamma_1, \gamma))$ converges towards one in the large- N limit, provided γ_1 diverges arbitrarily slowly and γ decreases to zero arbitrarily slowly. Therefore, Theorems 1.5, 1.9 and 1.11 follow straightforwardly from their nonasymptotic counterparts, namely Theorems 3.5, 3.6 and 3.7, respectively. Finally, Theorem 1.3 follows straightforwardly from Theorem 5.2 of [6], where the condition on the separation of the SNRs is expressed taking into account the value η . \square

4. COMPARISON INEQUALITIES

This section presents the comparison inequalities for both the correlations and the eigenvalues of the matrix $\mathbf{G} = \mathbf{M}\mathbf{M}^\top$, which are needed to prove our main results. Before delving into these, we provide some preliminary results.

Notations. Throughout, we write $\|\cdot\|$ for the operator norm on elements of $\mathbb{R}^{r \times r}$ induced by the L^2 -distance on \mathbb{R}^r . In particular, the operator norm of $\mathbf{A} \in \mathbb{R}^{r \times r}$ is defined as

$$\|\mathbf{A}\| = \max_{\mathbf{x} \in \mathbb{R}^r \setminus \{0\}} \frac{\|\mathbf{A}\mathbf{x}\|_2}{\|\mathbf{x}\|_2}.$$

In terms of its spectrum, the operator norm of $\mathbf{A} \in \mathbb{R}^{r \times r}$ corresponds to the largest singular value of \mathbf{A} , i.e., $\|\mathbf{A}\| = s_{\max}(\mathbf{A}) = \sqrt{\lambda_{\max}(\mathbf{A}^\top \mathbf{A})}$. Moreover, we let $\|\mathbf{A}\|_{\text{F}}$ denote the Frobenius norm of \mathbf{A} :

$$\|\mathbf{A}\|_{\text{F}} = \left(\sum_{1 \leq i, j \leq r} |\mathbf{A}_{ij}|^2 \right)^{1/2}.$$

For any symmetric matrices \mathbf{A} and \mathbf{B} , we write $\mathbf{A} \succeq \mathbf{B}$ if $\mathbf{A} - \mathbf{B} \succeq 0$, meaning that $\mathbf{A} - \mathbf{B}$ is positive semi-definite.

4.1. PRELIMINARY RESULTS

We first recall that the loss function $\mathcal{L}_{N,r}: \text{St}(N, r) \times (\mathbb{R}^N)^{\otimes p} \rightarrow \mathbb{R}$ is given by

$$\mathcal{L}_{N,r}(\mathbf{X}; \mathbf{Y}) = H_{N,r}(\mathbf{X}) + \Phi_{N,r}(\mathbf{X}),$$

where $H_{N,r}$ and $\Phi_{N,r}$ denote the noise and deterministic parts, respectively, which are defined by (1.4) and (1.5). To simplify notation slightly, in the remainder of the article we will write \mathcal{L}, H, Φ in place of $\mathcal{L}_{N,r}, H_{N,r}, \Phi_{N,r}$, respectively. We assume that $\mathbf{W} \in (\mathbb{R}^N)^{\otimes p}$ is an order- p tensor with i.i.d. centered sub-Gaussian entries W_{i_1, \dots, i_p} satisfying

$$\mathbb{E} [e^{tW_{i_1, \dots, i_p}}] \leq e^{Kt^2\sigma^2} \text{ for all } t \in \mathbb{R} \quad \text{and} \quad \mathbb{E} [|W_{i_1, \dots, i_p}|^k] \leq Ck^{k/2}\sigma^k \text{ for all } k \geq 1, \quad (4.1)$$

where $K, C > 0$ are absolute constants. Equivalently, $(W_{i_1, \dots, i_p})_{i_1, \dots, i_p}$ are i.i.d. random variables with sub-Gaussian norm $\|W_{i_1, \dots, i_p}\|_{\psi_2} \equiv \sigma > 0$, where for any sub-Gaussian random variable Z , the sub-Gaussian norm $\|Z\|_{\psi_2}$ is defined by $\|Z\|_{\psi_2} = \inf\{t > 0: \mathbb{E} [Z^2/t^2] \leq 2\}$ (see e.g. [64, Definition 2.5.6]).

Since \mathbf{W} is a sub-Gaussian p -tensor, the random variable $\langle \mathbf{v}, (\nabla_{\text{St}} H(\mathbf{X}))_j \rangle$ is also sub-Gaussian, as shown in the following lemma.

Lemma 4.1. *For every $\mathbf{v} \in \mathbb{S}^{N-1}$ and $i \in [r]$, the random variable $\langle \mathbf{v}, (\nabla_{\text{St}} H(\mathbf{X}))_i \rangle$ is sub-Gaussian with zero mean and sub-Gaussian norm satisfying*

$$\|\langle \mathbf{v}, (\nabla_{\text{St}} H(\mathbf{X}))_i \rangle\|_{\psi_2} \leq \alpha_1 \sigma,$$

for a constant $\alpha_1 = \alpha_1(p, \{\lambda_i\}_{i=1}^r) > 0$.

Proof. Using (1.8) the i th column vector of $\nabla_{\text{St}} H(\mathbf{X})$ is given by

$$(\nabla_{\text{St}} H(\mathbf{X}))_i = \lambda_i \nabla_{\mathbf{x}_i} \langle \mathbf{W}, \mathbf{x}_i^{\otimes p} \rangle - \frac{1}{2} \sum_{j=1}^r (\lambda_i \mathbf{x}_j^\top \nabla_{\mathbf{x}_i} \langle \mathbf{W}, \mathbf{x}_i^{\otimes p} \rangle + \lambda_j (\nabla_{\mathbf{x}_j} \langle \mathbf{W}, \mathbf{x}_j^{\otimes p} \rangle)^\top \mathbf{x}_i) \mathbf{x}_j,$$

where $\nabla_{\mathbf{x}_i} \langle \mathbf{W}, \mathbf{x}_i^{\otimes p} \rangle$ has coordinates given for every $k \in [N]$ by

$$(\nabla_{\mathbf{x}_i} \langle \mathbf{W}, \mathbf{x}_i^{\otimes p} \rangle)_k = p \sum_{1 \leq k_1, \dots, k_{p-1} \leq N} W_{k, k_1, \dots, k_{p-1}} (x_i)_{k_1} \cdots (x_i)_{k_{p-1}}.$$

Now, for every $\mathbf{v} \in \mathbb{S}^{N-1}$ the inner product $\langle \mathbf{v}, (\nabla_{\text{St}} H(\mathbf{X}))_i \rangle$ can be expanded as

$$\langle \mathbf{v}, (\nabla_{\text{St}} H(\mathbf{X}))_i \rangle = p \sum_{k=1}^N \sum_{1 \leq k_1, \dots, k_{p-1} \leq N} W_{k, k_1, \dots, k_{p-1}} f_i(\mathbf{v}, \mathbf{X}, \{\lambda_i\}_{i=1}^r),$$

where

$$\begin{aligned} f_i(\mathbf{v}, \mathbf{X}, \{\lambda_i\}_{i=1}^r) &= \lambda_i v_k (x_i)_{k_1} \cdots (x_i)_{k_{p-1}} \\ &\quad - \frac{1}{2} \sum_{j=1}^r (\lambda_i (x_j)_k (x_i)_{k_1} \cdots (x_i)_{k_{p-1}} + \lambda_j (x_i)_k (x_j)_{k_1} \cdots (x_j)_{k_{p-1}}) \langle \mathbf{v}, \mathbf{x}_j \rangle. \end{aligned}$$

This implies that the moment generating function of $\langle \mathbf{v}, (\nabla_{\text{St}} H(\mathbf{X}))_i \rangle$ is given by

$$\begin{aligned} \mathbb{E} \left[e^{t \langle \mathbf{v}, (\nabla_{\text{St}} H(\mathbf{X}))_i \rangle} \right] &= \prod_{k, k_1, \dots, k_{p-1}} \mathbb{E} \left[e^{t p W_{k, k_1, \dots, k_{p-1}} f_i(\mathbf{v}, \mathbf{X}, \{\lambda_i\}_{i=1}^r)} \right] \\ &\leq e^{C t^2 p^2 \sigma^2 \sum_{1 \leq k, k_1, \dots, k_{p-1} \leq N} f_i(\mathbf{v}, \mathbf{X}, \{\lambda_i\}_{i=1}^r)^2}, \end{aligned}$$

where we used the independence of the random variables W_{i_1, \dots, i_p} and the sub-Gaussian property (4.1). Since the bound on the moment generating function characterizes the sub-Gaussian distribution (see e.g. [64, Proposition 2.5.2]), it follows that the random variable $\langle \mathbf{v}, (\nabla_{\text{St}} H(\mathbf{X}))_i \rangle$ is sub-Gaussian with zero mean and sub-Gaussian norm bounded above by

$$\begin{aligned} \|\langle \mathbf{v}, (\nabla_{\text{St}} H(\mathbf{X}))_i \rangle\|_{\psi_2}^2 &\leq p^2 \sigma^2 \sum_{1 \leq k, k_1, \dots, k_{p-1} \leq N} f_i(\mathbf{v}, \mathbf{X}, \{\lambda_i\}_{i=1}^r)^2 \\ &= p^2 \sigma^2 \left(\frac{1}{4} \sum_{j=1}^r \lambda_j^2 \langle \mathbf{v}, \mathbf{x}_j \rangle^2 - \frac{1}{2} \lambda_i^2 \langle \mathbf{v}, \mathbf{x}_i \rangle^2 - \frac{3}{4} \lambda_i^2 \sum_{j=1}^r \langle \mathbf{v}, \mathbf{x}_j \rangle^2 \right) \\ &\leq \frac{1}{4} p^2 \sigma^2 \sum_{i=1}^r \lambda_i^2, \end{aligned}$$

where the equality follows by the fact that $\mathbf{v} \in \mathbb{S}^{N-1}$ and $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_r] \in \text{St}(N, r)$. The result follows straightforwardly. \square

According to Subsection 1.1, the output \mathbf{X}_t of the online SGD algorithm can be expanded as

$$\mathbf{X}_t = \left(\mathbf{X}_{t-1} - \frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \right) \mathbf{P}_t. \quad (4.2)$$

Here, \mathbf{P}_t denotes the projection term given by

$$\mathbf{P}_t = \left(\mathbf{I}_r + \frac{\delta^2}{N^2} \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \right)^{-1/2}, \quad (4.3)$$

where $\mathbf{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \in \mathbb{R}^{r \times r}$ is the Gram matrix given by

$$\mathbf{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t) = \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t)^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t). \quad (4.4)$$

We have the following estimate for the Frobenius norm of $\mathbf{G}(\mathbf{X}; \mathbf{Y})$.

Lemma 4.2. *There exists a constant $\alpha_2 = \alpha_2(p, \{\lambda_i\}_{i=1}^r)$ such that $\sup_{\mathbf{X} \in \text{St}(N, r)} \mathbb{E} [\|\mathbf{G}(\mathbf{X}; \mathbf{Y})\|_{\text{F}}] \leq \alpha_2 N$.*

In particular, if $\delta > 0$ satisfies $\delta < \left(\frac{N}{\alpha_2}\right)^{1/2}$, we obtain that

$$\mathbb{P} \left(\frac{\delta^2}{N^2} \|\mathbf{G}(\mathbf{X}; \mathbf{Y})\|_{\text{F}} < 1 \right) \geq 1 - \exp(-c(p)N).$$

Proof. We first notice that

$$\|\mathbf{G}(\mathbf{X}; \mathbf{Y})\|_{\text{F}} = \|\nabla_{\text{St}} \mathcal{L}(\mathbf{X}; \mathbf{Y})^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}; \mathbf{Y})\|_{\text{F}} \leq \|\nabla_{\text{St}} \mathcal{L}(\mathbf{X}; \mathbf{Y})\|_{\text{F}}^2.$$

Since by definition $\nabla_{\text{St}} \mathcal{L}(\mathbf{X}; \mathbf{Y}) = \nabla \mathcal{L}(\mathbf{X}; \mathbf{Y}) - \frac{1}{2} \mathbf{X} (\nabla \mathcal{L}(\mathbf{X}; \mathbf{Y}))^\top \mathbf{X} + \mathbf{X}^\top \nabla \mathcal{L}(\mathbf{X}; \mathbf{Y})$, it then follows that

$$\|\nabla_{\text{St}} \mathcal{L}(\mathbf{X}; \mathbf{Y})\|_{\text{F}} \leq \|\nabla \mathcal{L}(\mathbf{X}; \mathbf{Y})\|_{\text{F}} + \|\mathbf{X}^\top \nabla \mathcal{L}(\mathbf{X}; \mathbf{Y})\|_{\text{F}} \leq 2 \|\nabla \mathcal{L}(\mathbf{X}; \mathbf{Y})\|_{\text{F}},$$

where for the first inequality we used the triangle inequality and the fact that $\|\mathbf{X} \nabla \mathcal{L}(\mathbf{X}; \mathbf{Y})^\top \mathbf{X}\|_{\text{F}} = \|\mathbf{X}^\top \nabla \mathcal{L}(\mathbf{X}; \mathbf{Y})\|_{\text{F}}$ and similarly $\|\mathbf{X} \mathbf{X}^\top \nabla \mathcal{L}(\mathbf{X}; \mathbf{Y})\|_{\text{F}} = \|\mathbf{X}^\top \nabla \mathcal{L}(\mathbf{X}; \mathbf{Y})\|_{\text{F}}$, while the second inequality follows by the property $\|\mathbf{X}^\top \mathbf{A}\|_{\text{F}} \leq \|\mathbf{A}\|_{\text{F}}$ for $\mathbf{X} \in \text{St}(N, r)$ and $\mathbf{A} \in \mathbb{R}^{N \times r}$. We can write the Frobenius norm of the Euclidean gradient of \mathcal{L} as

$$\|\nabla \mathcal{L}(\mathbf{X}; \mathbf{Y})\|_{\text{F}}^2 = \sum_{i=1}^r \|\nabla_{\mathbf{x}_i} \mathcal{L}(\mathbf{X}; \mathbf{Y})\|_2^2,$$

where $\nabla_{\mathbf{x}_i} \mathcal{L}(\mathbf{X}; \mathbf{Y})$ denotes the i th column of the Euclidean gradient $\nabla \mathcal{L}(\mathbf{X}; \mathbf{Y})$ given by

$$\nabla_{\mathbf{x}_i} \mathcal{L}(\mathbf{X}; \mathbf{Y}) = \lambda_i \nabla_{\mathbf{x}_i} \langle \mathbf{W}, \mathbf{x}_i^{\otimes p} \rangle + \sum_{j=1}^r \sqrt{N} p \lambda_i \lambda_j \langle \mathbf{v}_j, \mathbf{x}_i \rangle^{p-1} \mathbf{v}_j.$$

This implies that

$$\begin{aligned} \mathbb{E} [\|\mathbf{G}(\mathbf{X}; \mathbf{Y})\|_{\text{F}}] &\leq 4 \sum_{i=1}^r \left(p \lambda_i \mathbb{E} [\|\mathbf{W}\{\mathbf{x}_i\}\|_2] + \left\| \sum_{j=1}^r \sqrt{N} p \lambda_i \lambda_j \langle \mathbf{v}_j, \mathbf{x}_i \rangle^{p-1} \mathbf{v}_j \right\|_2 \right)^2 \\ &\leq 8 \sum_{i=1}^r \left(p^2 \lambda_i^2 \mathbb{E} [\|\mathbf{W}\{\mathbf{x}_i\}\|_2^2] + N \sum_{j=1}^r p^2 \lambda_i^2 \lambda_j^2 \right), \end{aligned}$$

where we used the fact that $(x+y)^2 \leq 2(x^2+y^2)$, Jensen's inequality, and the bound $\max_{i,j} |\langle \mathbf{v}_j, \mathbf{x}_i \rangle| \leq 1$. Furthermore, an explicit computation of the term $\mathbb{E} [\|\mathbf{W}\{\mathbf{x}_i\}\|_2^2]$ gives the bound $\mathbb{E} [\|\mathbf{W}\{\mathbf{x}_i\}\|_2^2] \leq 2K\sigma^2 N$. This shows the first part of the statement. For the second part, we note that

$$\|\mathbf{G}(\mathbf{X}; \mathbf{Y})\|_{\text{F}} \leq 8 \sum_{i=1}^r \left(p^2 \lambda_i^2 \|\mathbf{W}\{\mathbf{x}_i\}\|_2^2 + N \sum_{j=1}^r p^2 \lambda_i^2 \lambda_j^2 \right) \leq 8 \sum_{i=1}^r \left(p^2 \lambda_i^2 \|\mathbf{W}\|^2 + N \sum_{j=1}^r p^2 \lambda_i^2 \lambda_j^2 \right).$$

From Lemma 4.7 of [8] we then have that for every $\gamma > 0$ there exists an absolute constant $c > 0$ such that

$$\mathbb{P} \left(\|\mathbf{W}\| \geq \gamma \sqrt{N} \right) \leq 2(8p)^{Np} e^{-c\gamma^2 N/8}.$$

Choosing $\gamma > \sqrt{2K}\sigma$, we then obtain that there is $c(p) > 0$ such that

$$\|\mathbf{G}(\mathbf{X}; \mathbf{Y})\|_{\text{F}} \leq 8p^2 \sum_{i=1}^r \left(\lambda_i^2 2K\sigma^2 N + N \sum_{j=1}^r \lambda_i^2 \lambda_j^2 \right) = N\alpha_2(p, \{\lambda_i\}_{i=1}^r).$$

with \mathbb{P} -probability at least $1 - e^{-c(p)N}$. This completes the proof of Lemma 4.2. \square

4.2. COMPARISON INEQUALITIES FOR CORRELATIONS

The goal of this subsection is to obtain a two-sided difference inequality for the evolution of the correlations $m_{ij}(t) = m_{ij}(\mathbf{X}_t)$. This will indeed be needed to prove both Theorems 3.5 and 3.6.

Proposition 4.3. *Let $\alpha_1 = \alpha_1(p, \{\lambda_i\}_{i=1}^r)$ and $\alpha_2 = \alpha_2(p, \{\lambda_i\}_{i=1}^r)$ denote the constants given by Lemmas 4.1 and 4.2, respectively. For constants $c_0 \in (0, 1)$ and $K_N > 0$, and for every $i, j \in [r]$, we define the comparison functions ℓ_{ij} and u_{ij} by*

$$\begin{aligned}\ell_{ij}(t) &= (1 - c_0)m_{ij}(0) + \frac{\delta}{N} \sum_{\ell=1}^t \left\{ -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle - |h_i(\ell-1)| \right\}, \\ u_{ij}(t) &= (1 + c_0)m_{ij}(0) + \frac{\delta}{N} \sum_{\ell=1}^t \left\{ -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle + |h_i(\ell-1)| \right\},\end{aligned}$$

where

$$|h_i(\ell-1)| = \frac{\delta \alpha_2}{2} \sum_{k=1}^r \left(|m_{ik}(\ell-1)| + \frac{\delta}{N} |\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_k \rangle| + \frac{K_N \delta}{N} \right),$$

and

$$\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}))_j \rangle = -\sqrt{N} p \lambda_i \lambda_j m_{ij}^{p-1} + \sqrt{N} \frac{p}{2} \sum_{1 \leq k, \ell \leq r} \lambda_k m_{i\ell} m_{kj} m_{k\ell} \left(\lambda_j m_{kj}^{p-2} + \lambda_\ell m_{k\ell}^{p-2} \right).$$

Consider the sequence of outputs $(\mathbf{X}_t)_{t \in \mathbb{N}}$ given by (1.6) with step-size $\delta > 0$ satisfying $\delta < (N/\alpha_2)^{1/2}$. Let $T > 0$ denote a fixed time horizon. Then, for every $\gamma_1 > \gamma_2 > 0$ there exist $c_0 \in (0, 1)$ and $K_N > 0$ such that for every $i, j \in [r]$ and every $t \leq T$,

$$\inf_{\mathbf{X}_0 \in \mathcal{C}_1(\gamma_1, \gamma_2)} \mathbb{P}_{\mathbf{X}_0} (\ell_{ij}(t) \leq m_{ij}(t) \leq u_{ij}(t)) \geq 1 - \eta(\delta, c_0, T, K_N),$$

where $\eta = \eta(\delta, c_0, T, K_N)$ is given by

$$\eta = T \exp(-c_1 N) - 2 \exp\left(-c_2 \frac{c_0^2 \gamma_2^2 N}{8 \delta^2 \alpha_1^2 \sigma^2 T}\right) - C \frac{2 \delta^3 r \sigma \alpha_1 \alpha_2}{N^{3/2} c_0 \gamma_2} T \exp\left(-c_3 \frac{K_N^2}{2 \alpha_1^2 \sigma^2}\right). \quad (4.5)$$

The proof of Proposition 4.3 will be the main focus of this subsection. Our strategy involves first expanding the polar retraction term $R_{\mathbf{X}_{t-1}}$ from the algorithm (see (1.7)) as a Neumann series and then decompose the evolution of $m_{ij}(t)$ in three parts: the drift induced by the population loss, a martingale induced by the gradient of the noise Hamiltonian H , and high-order terms in δ . By controlling the latter two parts, the proposition will follow.

Lemma 4.4. *Let $\alpha_2 = \alpha_2(p, \{\lambda_i\}_{i=1}^r)$ be the constant of Lemma 4.2. Assume that $\delta < (N/\alpha_2)^{1/2}$. For every $T > 0$, we let $\mathcal{E}_{\mathcal{G}} = \mathcal{E}_{\mathcal{G}}(T)$ denote the event*

$$\mathcal{E}_{\mathcal{G}}(T) = \bigcap_{1 \leq t \leq T} \left\{ \left\| \frac{\delta^2}{N^2} \mathcal{G}(\mathbf{X}_{t-1}, \mathbf{Y}^t) \right\|_{\text{F}} < 1 \right\}. \quad (4.6)$$

Then, there exists a constant $c(p)$ such that $\mathbb{P}(\mathcal{E}_{\mathcal{G}}) \geq 1 - T \exp(-c(p)N)$. Moreover, on the event $\mathcal{E}_{\mathcal{G}}$, for every $i, j \in [r]$ and every $t \leq T$, it holds that

$$\left| m_{ij}(t) - \left(m_{ij}(0) - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_j \rangle \right) \right| \leq \frac{\delta}{N} \sum_{\ell=1}^t |a_i(\ell-1)|, \quad (4.7)$$

where

$$|a_i(\ell-1)| \leq \frac{\delta \alpha_2}{2} \sum_{k=1}^r \left(|m_{ik}(\ell-1)| + \frac{\delta}{N} |\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_k \rangle| + \frac{\delta}{N} |\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| \right). \quad (4.8)$$

Proof. We begin by expanding the projection term \mathbf{P}_t in (4.2) as a Neumann series. To do so, we require that $\left\| \frac{\delta^2}{N^2} \mathcal{G}(\mathbf{X}_{t-1}, \mathbf{Y}^t) \right\|_{\text{F}} < 1$ for every $t \leq T$. According to Lemma 4.2, this condition holds if δ satisfies $\delta < \left(\frac{N}{\alpha_2} \right)^{1/2}$, leading to

$$\mathbb{P}(\mathcal{E}_{\mathcal{G}}^c) \leq T \mathbb{P} \left(\left\| \frac{\delta^2}{N^2} \mathcal{G}(\mathbf{X}_{t-1}, \mathbf{Y}^t) \right\|_{\text{F}} \geq 1 \right) \leq T e^{-c(p)N}.$$

We therefore place ourselves on the event $\mathcal{E}_{\mathcal{G}}$ for the rest of the proof. Expanding \mathbf{P}_t as a Neumann series yields

$$\mathbf{P}_t = \sum_{k=0}^{\infty} (-1)^k \binom{\frac{1}{2}}{k} \frac{\delta^{2k}}{N^{2k}} \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)^k = \mathbf{I}_r + \sum_{k=1}^{\infty} (-1)^k \binom{\frac{1}{2}}{k} \frac{\delta^{2k}}{N^{2k}} \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)^k.$$

Plugging this expansion into (4.2), the online SGD output \mathbf{X}_t at time t can be rewritten as

$$\begin{aligned} \mathbf{X}_t &= \mathbf{X}_{t-1} - \frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \\ &\quad + \sum_{k=1}^{\infty} (-1)^k \binom{\frac{1}{2}}{k} \frac{\delta^{2k}}{N^{2k}} \left(\mathbf{X}_{t-1} - \frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \right) \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)^k. \end{aligned}$$

For every $i, j \in [r]$ and every $1 \leq t \leq T$, the evolution of the correlation $m_{ij}(\mathbf{X}_t) = m_{ij}(t)$ under (1.6) is therefore given by

$$\begin{aligned} m_{ij}(t) &= \langle \mathbf{v}_i, (\mathbf{X}_t)_j \rangle \\ &= m_{ij}(t-1) - \frac{\delta}{N} \langle \mathbf{v}_i, (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t))_j \rangle \\ &\quad + \sum_{k=1}^{\infty} (-1)^k \binom{\frac{1}{2}}{k} \frac{\delta^{2k}}{N^{2k}} \text{Tr} \left([\mathbf{v}_i]_j^\top \left(\mathbf{X}_{t-1} - \frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \right) \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)^k \right), \end{aligned}$$

where $(\mathbf{X}_t)_j$ denotes the j th column of \mathbf{X}_t , $(\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t))_j$ the j th column of $\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t)$, and $[\mathbf{v}_i]_j = [\mathbf{0}, \dots, \mathbf{v}_i, \dots, \mathbf{0}] \in \mathbb{R}^{N \times r}$ the matrix in which all columns are the zero vector except column j which is given by \mathbf{v}_i . Iterating the above expression, we then obtain that

$$\begin{aligned} m_{ij}(t) &= m_{ij}(0) - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell))_j \rangle \\ &\quad + \sum_{\ell=1}^t \sum_{k=1}^{\infty} (-1)^k \binom{\frac{1}{2}}{k} \frac{\delta^{2k}}{N^{2k}} \text{Tr} \left([\mathbf{v}_i]_j^\top \left(\mathbf{X}_{\ell-1} - \frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell) \right) \mathcal{G}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell)^k \right). \end{aligned}$$

In the following, for every $i, j \in [r]$ and every $1 \leq \ell \leq t$ we let $a_{ij}(\ell-1)$ denote

$$a_{ij}(\ell-1) = \sum_{k=1}^{\infty} (-1)^k \binom{\frac{1}{2}}{k} \frac{\delta^{2k}}{N^{2k}} \text{Tr} \left([\mathbf{v}_i]_j^\top \left(\mathbf{X}_{\ell-1} - \frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell) \right) \mathcal{G}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell)^k \right).$$

By Cauchy-Schwarz we then have that

$$\begin{aligned} |a_{ij}(\ell-1)| &\leq \sum_{k=1}^{\infty} \binom{\frac{1}{2}}{k} \frac{\delta^{2k}}{N^{2k}} \left| \text{Tr} \left([\mathbf{v}_i]_j^\top \left(\mathbf{X}_{\ell-1} - \frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell) \right) \mathcal{G}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell)^k \right) \right| \\ &\leq \sum_{k=1}^{\infty} \binom{\frac{1}{2}}{k} \frac{\delta^{2k}}{N^{2k}} \|\mathcal{G}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell)^k\|_{\text{F}} \|[\mathbf{v}_i]_j^\top \left(\mathbf{X}_{\ell-1} - \frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell) \right)\|_{\text{F}} \\ &\leq \sum_{k=1}^{\infty} \binom{\frac{1}{2}}{k} C_{\mathcal{G}}^{k-1} \frac{\delta^2}{N^2} \|\mathcal{G}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell)\|_{\text{F}} \|[\mathbf{v}_i]_j^\top \left(\mathbf{X}_{\ell-1} - \frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell) \right)\|_{\text{F}}, \end{aligned} \quad (4.9)$$

where $C_{\mathcal{G}} = \max \|\frac{\delta^2}{N^2} \mathcal{G}(\mathbf{X}; \mathbf{Y})\| \in (0, 1)$ under the event $\mathcal{E}_{\mathcal{G}}$ for $\delta < (N/\alpha_2)^{1/2}$. From the Newton Binomial Theorem we find that

$$\sum_{k=1}^{\infty} \binom{\frac{1}{2}}{k} C_{\mathcal{G}}^{k-1} = \frac{1}{C_{\mathcal{G}}} \left(\sqrt{1 + C_{\mathcal{G}}} - 1 \right) \leq \frac{1}{C_{\mathcal{G}}} \frac{C_{\mathcal{G}}}{2} = \frac{1}{2}. \quad (4.10)$$

Moreover, a simple computation shows that

$$\begin{aligned} \|[\mathbf{v}_i]_j^\top \left(\mathbf{X}_{\ell-1} - \frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell) \right)\|_{\text{F}} &= \sqrt{\sum_{k=1}^r \left(m_{ik}(\ell-1) - \frac{\delta}{N} \langle \mathbf{v}_i, (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell))_k \rangle \right)^2} \\ &\leq \sum_{k=1}^r \left(m_{ik}(\ell-1) - \frac{\delta}{N} \langle \mathbf{v}_i, (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell))_k \rangle \right). \end{aligned} \quad (4.11)$$

Combining (4.9)-(4.11) yields

$$|a_{ij}(\ell-1)| \leq \frac{\delta^2}{2N^2} \|\mathcal{G}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell)\|_F \sum_{k=1}^r \left(m_{ik}(\ell-1) - \frac{\delta}{N} \langle \mathbf{v}_i, (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell))_k \rangle \right),$$

which gives the desired bound since $\|\mathcal{G}(\mathbf{X}_{\ell-1}; \mathbf{Y}^\ell)\|_F \leq \alpha_2 N$ on the event $\mathcal{E}_{\mathcal{G}}$. \square

According to (4.7) from Lemma 4.4, we now provide an estimate for the first-order (in δ) noise Hamiltonian H , i.e., $\frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_j \rangle$.

Lemma 4.5. *Let $\alpha_1 = \alpha_1(p, \{\lambda_i\}_{i=1}^r)$ be the constant from Lemma 4.1. Then, for every $a > 0$, there is an absolute constant $c > 0$ such that*

$$\sup_{\mathbf{X}_0 \in \text{St}(N,r)} \mathbb{P}_{\mathbf{X}_0} \left(\max_{t \leq T} \frac{\delta}{N} \left| \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_j \rangle \right| \geq a \right) \leq 2 \exp \left(-c \frac{a^2 N^2}{2\delta^2 \alpha_1^2 \sigma^2 T} \right).$$

Proof. We let \mathcal{F}_t denote the natural filtration associated to $\mathbf{X}_1, \dots, \mathbf{X}_t$, namely $\mathcal{F}_t = \sigma(\mathbf{X}_\ell, \ell \leq t)$. We let M_t denote the process

$$M_t = \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_j \rangle.$$

Since

$$\mathbb{E}[M_t - M_{t-1} | \mathcal{F}_{t-1}] = \mathbb{E} \left[\langle \mathbf{v}_i, (\nabla_{\text{St}} H^t(\mathbf{X}_{t-1}))_j \mid \mathcal{F}_{t-1} \right] = 0,$$

we have that the process M_t is adapted to the filtration $(\mathcal{F}_t)_{t \geq 0}$. Moreover, since $M_s = \mathbb{E}[M_t | \mathcal{F}_s]$, we have that $(M_t)_{t \geq 0}$ is a martingale with respect to $(\mathcal{F}_t)_{t \geq 0}$. Since by Lemma 4.1 the distribution of the increments $M_t - M_{t-1} = \langle \mathbf{v}_i, (\nabla_{\text{St}} H^t(\mathbf{X}_{t-1}))_j \rangle$ given \mathcal{F}_{t-1} is sub-Gaussian with sub-Gaussian norm bounded by $\alpha_1 \sigma$, we have for every $b > 0$, there exists an absolute constants $C > 0$ such that

$$\mathbb{E} \left[e^{b(M_t - M_{t-1})} | \mathcal{F}_{t-1} \right] \leq \exp(Cb^2 \alpha_1^2 \sigma^2).$$

The corresponding Chernoff bound applied to $M_t = \sum_{\ell=1}^t (M_\ell - M_{\ell-1})$ (see e.g. Theorem 2.19 of [65]) gives the tail bound: for every $t \leq T$ and every $a > 0$,

$$\sup_{\mathbf{X}_0 \in \text{St}(N,r)} \mathbb{P}_{\mathbf{X}_0} \left(\frac{\delta}{N} |M_t| \geq a \right) \leq 2 \exp \left(-c \frac{a^2 N^2}{2\delta^2 t \alpha_1^2 \sigma^2} \right).$$

To improve this bound to the maximum $\max_t |M_t|$, we follow the remark at the end of section 3.5 from [44], noting that the random process e^{M_t} is a submartingale. The other side of the tail bound is obtained by noting that $-M_t$ is also a martingale. \square

It remains to estimate the third-order noise term in (4.8) from Lemma 4.4. To this end, we introduce a truncation level $K_N > 0$ to be chosen later and write

$$|\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| \leq K_N + |\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| \mathbf{1}_{\{|\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| > K_N\}}.$$

Lemma 4.6. *Let α_1 and α_2 be the constants from Lemmas 4.1 and 4.2, respectively. Then, for every $a > 0$, there are absolute constants $C, c > 0$ such that*

$$\begin{aligned} & \sup_{\mathbf{X}_0 \in \text{St}(N,r)} \mathbb{P}_{\mathbf{X}_0} \left(\max_{t \leq T} \frac{\delta^3 \alpha_2}{2N^2} \sum_{\ell=1}^t \sum_{k=1}^r |\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| \mathbf{1}_{\{|\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| > K_N\}} \geq a \right) \\ & \leq C \frac{\delta^3 r \sigma \alpha_1 \alpha_2}{N^2 a} T \exp \left(-c \frac{K_N^2}{2\alpha_1^2 \sigma^2} \right). \end{aligned}$$

Proof. Since the summand is positive, the maximum over $t \leq T$ is attained at $t = T$, so it suffices to consider that case. Moreover, by Lemma 4.1 we have that the distribution of $\langle \mathbf{v}_i, (\nabla_{\text{St}} H(\mathbf{X}))_j \rangle$ is sub-Gaussian with sub-Gaussian norm bounded by $\alpha_1 \sigma$. It therefore follows that

$$\sup_{\mathbf{X} \in \text{St}(N,r)} \mathbb{E} [|\langle \mathbf{v}_i, (\nabla_{\text{St}} H(\mathbf{X}))_j \rangle|^2] \leq 2C \alpha_1^2 \sigma^2,$$

and that

$$\sup_{\mathbf{X} \in \text{St}(N,r)} \mathbb{P} (|\langle \mathbf{v}_i, (\nabla_{\text{St}} H(\mathbf{X}))_j \rangle| > K_N) \leq 2 \exp \left(-c \frac{K_N^2}{\alpha_1^2 \sigma^2} \right),$$

where we used the equivalence of the sub-Gaussian properties (see e.g. [64, Proposition 2.5.2]) for both inequalities. Fix $\mathbf{X}_0 \in \text{St}(N, r)$. We then obtain for every $a > 0$,

$$\begin{aligned} & \mathbb{P}_{\mathbf{X}_0} \left(\frac{\delta^3 \alpha_2}{2N^2} \sum_{\ell=1}^T \sum_{k=1}^r |\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| \mathbf{1}_{\{|\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| > K_N\}} \geq a \right) \\ & \leq \frac{\delta^3 r \alpha_2}{2N^2 a} T \left(\sup_{\ell \leq T} \sup_{k \in [r]} \mathbb{E}_{\mathbf{X}_0} [|\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_j \rangle| \mathbf{1}_{\{|\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| > K_N\}}] \right) \\ & \leq \frac{\delta^3 r \alpha_2}{2N^2 a} T \sqrt{\sup_{\mathbf{X}} \mathbb{E} [|\langle \mathbf{v}_i, (\nabla_{\text{St}} H(\mathbf{X}))_k \rangle|^2]} \sqrt{\sup_{\mathbf{X}} \mathbb{E} [\mathbf{1}_{\{|\langle \mathbf{v}_i, (\nabla_{\text{St}} H(\mathbf{X}))_k \rangle| > K_N\}}]} \\ & \leq \frac{C \delta^3 r \sigma \alpha_1 \alpha_2}{N^2 a} T \exp \left(-c \frac{K_N^2}{2\alpha_1^2 \sigma^2} \right), \end{aligned}$$

where we used Markov for the first inequality and Cauchy-Schwarz for the second one. \square

Proposition 4.3 now follows straightforwardly by combining the above three lemmas.

Proof of Proposition 4.3. Assume that the step-size δ satisfies $\delta < \sqrt{\frac{N}{\alpha_2}}$, where $\alpha_2 = \alpha_2(p, \{\lambda_i\}_{i=1}^r)$ is the constant from Lemma 4.2. From Lemma 4.4, for every $i, j \in [r]$ and every $t \leq T$ we have the following comparison inequalities for the correlations m_{ij} :

$$\begin{aligned} m_{ij}(t) & \leq m_{ij}(0) - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle + \frac{\delta}{N} \left| \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_j \rangle \right| + \frac{\delta}{N} \sum_{\ell=1}^t |a_i(\ell-1)|, \\ m_{ij}(t) & \geq m_{ij}(0) - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle - \frac{\delta}{N} \left| \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_j \rangle \right| - \frac{\delta}{N} \sum_{\ell=1}^t |a_i(\ell-1)|, \end{aligned}$$

with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - T \exp(-c_1 N)$, where $|a_i(\ell-1)|$ is given by (4.8). Next, we bound the noise terms by $\frac{c_0}{2} m_{ij}(0)$ for some constant $c_0 \in (0, 1)$ using Lemmas 4.5 and 4.6. We first note that under the event $\mathcal{C}_1(\gamma_1, \gamma_2)$, for every $i, j \in [r]$ there exists $\gamma_{ij} \in (\gamma_2, \gamma_1)$ such that $m_{ij}(0) = \gamma_{ij} N^{-\frac{1}{2}}$. It then follows from Lemma 4.5 that for all $t \leq T$,

$$\frac{\delta}{N} \left| \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_j \rangle \right| \leq \frac{c_0 \gamma_{ij}}{2\sqrt{N}},$$

with \mathbb{P} -probability at least $1 - 2 \exp\left(-c_2 \frac{c_0^2 \gamma_2^2 N}{8\delta^2 \alpha_1^2 \sigma^2 T}\right)$. Moreover, we have from Lemma 4.6 that for all $t \leq T$,

$$\frac{\delta^3 \alpha_2}{2N^2} \sum_{\ell=1}^t \sum_{k=1}^r |\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| \mathbf{1}_{\{|\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| > K_N\}} \leq \frac{c_0 \gamma_{ij}}{2\sqrt{N}},$$

with \mathbb{P} -probability at least $1 - \frac{2C_1 \delta^3 r \sigma \alpha_1 \alpha_2}{N^{\frac{3}{2}} c_0 \gamma_2} T \exp\left(-\frac{c_3 K_N^2}{2\alpha_1^2 \sigma^2}\right)$. The result then follows by standard properties of conditional probabilities.

It remains to compute $\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}))_j \rangle$ for every $\mathbf{X} \in \text{St}(N, r)$. According to (1.5), the j th column of the Euclidean gradient of the population loss $\Phi(\mathbf{X})$ is given by

$$(\nabla \Phi(\mathbf{X}))_j = - \sum_{k=1}^r p \sqrt{N} \lambda_k \lambda_j m_{kj}^{p-1}(\mathbf{X}) \mathbf{v}_k,$$

and the j th column of the Riemannian gradient is given by

$$\begin{aligned} (\nabla_{\text{St}} \Phi(\mathbf{X}))_j & = -(\nabla \Phi(\mathbf{X}))_j + \frac{1}{2} \sum_{\ell=1}^r \left(\mathbf{X}^\top \nabla \Phi(\mathbf{X}) + (\nabla \Phi(\mathbf{X}))^\top \mathbf{X} \right)_{\ell j} \mathbf{x}_\ell \\ & = - \sum_{k=1}^r p \sqrt{N} \lambda_k \lambda_j m_{kj}^{p-1} \mathbf{v}_k + \frac{1}{2} \sum_{1 \leq k, \ell \leq r} p \sqrt{N} \lambda_k m_{kj} m_{k\ell} \left(\lambda_j m_{kj}^{p-2} + \lambda_\ell m_{k\ell}^{p-2} \right) \mathbf{x}_\ell. \end{aligned}$$

We therefore have for every $i, j \in [r]$,

$$\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}))_j \rangle = -\sqrt{N} p \lambda_i \lambda_j m_{ij}^{p-1} + \sqrt{N} \frac{p}{2} \sum_{1 \leq k, \ell \leq r} \lambda_k m_{i\ell} m_{kj} m_{k\ell} \left(\lambda_j m_{kj}^{p-2} + \lambda_\ell m_{k\ell}^{p-2} \right).$$

□

4.3. COMPARISON INEQUALITIES FOR EIGENVALUES

In this subsection, we derive a two-sided difference inequality for the evolution of the eigenvalues $\theta_1(t), \dots, \theta_r(t)$ of $\mathbf{G}_t = \mathbf{G}(\mathbf{X}_t)$, where \mathbf{X}_t denotes the online SGD output at time t . In the following, we let $\Pi_{\mathbf{X}}^{\text{St}}$ denote the orthogonal projection on the tangent space $T_{\mathbf{X}} \text{St}(N, r)$, i.e.,

$$\Pi_{\mathbf{X}}^{\text{St}}(\mathbf{A}) = \mathbf{A} - \mathbf{X} \text{sym}(\mathbf{X}^\top \mathbf{A}), \quad \text{with } \text{sym}(\mathbf{A}) = \frac{1}{2} (\mathbf{A} + \mathbf{A}^\top).$$

In particular, according to (1.8), for a function $F: \text{St}(N, r) \rightarrow \mathbb{R}$, we can write $\nabla_{\text{St}} F(\mathbf{X}) = \Pi_{\mathbf{X}}^{\text{St}}(\nabla F(\mathbf{X}))$. To avoid cumbersome phrasing, we will often omit stating explicitly that bounds on conditional expectations hold almost surely.

Lemma 4.7. *Let $p = 2$ and $\alpha_2 = \alpha_2(r, \{\lambda_i\}_{i=1}^r)$ be the constant of Lemma 4.2. Consider the sequence of outputs $(\mathbf{X}_t)_{t \in \mathbb{N}}$ given by (1.6) with step-size δ satisfying $\delta < (N/\alpha_2)^{1/2}$. Let $T > 0$ be a fixed time horizon. Then, there exist constants $K_N, C, c > 0$ such that for every $t \leq T$,*

$$-\mathbf{A}_t \preceq \mathbf{G}_t - \left(\mathbf{G}_0 + \frac{4\lambda^2 \delta}{\sqrt{N}} \sum_{\ell=1}^t \mathbf{G}_{\ell-1} (\mathbf{I}_r - \mathbf{G}_{\ell-1}) - \mathbf{S}_t \right) \preceq \mathbf{A}_t,$$

with \mathbb{P} -probability at least $1 - T \exp(-c_1 N) - 4r^2 T \exp(-c_2 \frac{K_N^2}{\sigma^2})$. Here, $\mathbf{S}_t \in \mathbb{R}^{r \times r}$ is defined by

$$\mathbf{S}_t = \frac{4\lambda \delta}{N} \sum_{\ell=1}^t \text{sym} \left(\mathbf{V}^\top \Pi_{\mathbf{X}_{\ell-1}}^{\text{St}} \left(\mathbf{W}^\ell \mathbf{X}_{\ell-1} \right) \mathbf{M}_{\ell-1}^\top \right), \quad (4.12)$$

and $\mathbf{A}_t \in \mathbb{R}^{r \times r}$ is given by

$$\begin{aligned} \mathbf{A}_t = \sum_{\ell=1}^t & \left\{ \frac{8\lambda^2 \delta^2}{N^2} (4r^2 K_N^2 + \lambda^2 N \|\mathbf{G}_{\ell-1}\|) + \frac{\alpha_2 \delta^2}{N} \|\mathbf{G}_{\ell-1}\| \right. \\ & \left. + \frac{4\lambda \alpha_2 \delta^3}{N^2} (2r K_N + \lambda \sqrt{N} \|\mathbf{G}_{\ell-1}\|^{1/2}) \|\mathbf{G}_{\ell-1}\|^{1/2} \right\} \mathbf{I}_r. \end{aligned} \quad (4.13)$$

Proof. From (4.2) and (4.3), the iteration for the matrix $\mathbf{G}_t = \mathbf{V}^\top \mathbf{X}_t \mathbf{X}_t^\top \mathbf{V}$ results in

$$\mathbf{G}_t = \mathbf{V}^\top \left(\mathbf{X}_{t-1} - \frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \right) \mathbf{P}_t^2 \left(\mathbf{X}_{t-1} - \frac{\delta}{N} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \right)^\top \mathbf{V},$$

where we used the fact that $\mathbf{P}_t^\top = \mathbf{P}_t$. We let \mathbf{B}_t denote

$$\mathbf{B}_t = \left(\mathbf{X}_{t-1}^\top - \frac{\delta}{N} (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t))^\top \right) \mathbf{V}.$$

In the following, we place ourselves on the event $\mathcal{E}_{\mathcal{G}}(T) = \bigcap_{1 \leq t \leq T} \left\{ \left\| \frac{\delta^2}{N^2} \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \right\|_{\mathbb{F}} < 1 \right\}$. According to Lemma 4.4, we have that $\mathbb{P}(\mathcal{E}_{\mathcal{G}}) \geq 1 - T \exp(-c_1 N)$ for every $p \geq 2$. As in the proof of Lemma 4.4, we can therefore expand the projection term \mathbf{P}_t^2 as a Neumann series:

$$\mathbf{P}_t^2 = \left(\mathbf{I}_r + \frac{\delta^2}{N^2} \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \right)^{-1} = \sum_{k=0}^{\infty} (-1)^k \frac{\delta^{2k}}{N^{2k}} \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)^k.$$

We write the output \mathbf{G}_t at time t as

$$\mathbf{G}_t = \mathbf{B}_t^\top \mathbf{P}_t^2 \mathbf{B}_t = \mathbf{B}_t^\top \mathbf{B}_t + \sum_{k=1}^{\infty} (-1)^k \frac{\delta^{2k}}{N^{2k}} \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)^k.$$

We now want to expand the matrix \mathbf{B}_t by explicitly computing the Riemannian gradient. Since $p = 2$ and $\lambda_1 = \dots = \lambda_r \equiv \lambda$, the loss function $\mathcal{L}(\mathbf{X}; \mathbf{Y})$ given by (1.3) reduces to

$$\mathcal{L}(\mathbf{X}; \mathbf{Y}) = \lambda \sum_{i=1}^r \mathbf{x}_i^\top \mathbf{W} \mathbf{x}_i - \sqrt{N} \lambda^2 \sum_{i,j=1}^r m_{ij}^2(\mathbf{X}).$$

The Euclidean gradient $\nabla \mathcal{L}(\mathbf{X}; \mathbf{Y})$ is then given by

$$\nabla \mathcal{L}(\mathbf{X}; \mathbf{Y}) = 2\lambda \mathbf{W} \mathbf{X} - 2\sqrt{N} \lambda^2 \mathbf{V} \mathbf{M},$$

where $\mathbf{M} = (m_{ij})_{1 \leq i, j \leq r} = \mathbf{V}^\top \mathbf{X} \in \mathbb{R}^{r \times r}$, and using (1.8), the Riemannian gradient is given by

$$\nabla_{\text{St}} \mathcal{L}(\mathbf{X}; \mathbf{Y}) = 2\lambda \Pi_{\mathbf{X}}^{\text{St}}(\mathbf{W} \mathbf{X}) - 2\sqrt{N} \lambda^2 \mathbf{V} \mathbf{M} + 2\sqrt{N} \lambda^2 \mathbf{X} \mathbf{M}^\top \mathbf{M}, \quad (4.14)$$

where we recall that $\Pi_{\mathbf{X}}^{\text{St}}$ denotes the projection onto $T_{\mathbf{X}} \text{St}(N, r)$. This implies that

$$\begin{aligned} \mathbf{G}_t &= \mathbf{G}_{t-1} + \frac{4\delta\lambda^2}{\sqrt{N}} \mathbf{G}_{t-1} (\mathbf{I}_r - \mathbf{G}_{t-1}) - \frac{4\delta\lambda}{N} \text{sym} \left(\mathbf{V}^\top \Pi_{\mathbf{X}_{t-1}}^{\text{St}} (\mathbf{W}^t \mathbf{X}_{t-1}) \mathbf{M}_{t-1}^\top \right) \\ &\quad + \frac{\delta^2}{N^2} \mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t))^\top \mathbf{V} \\ &\quad + \sum_{k=1}^{\infty} (-1)^k \frac{\delta^{2k}}{N^{2k}} \mathbf{B}_t^\top \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)^k \mathbf{B}_t. \end{aligned} \quad (4.15)$$

The remainder of the proof focuses on estimating the terms in (4.15) that are of order 2 or higher in δ .

First, we consider the summand

$$\frac{\delta^2}{N^2} \mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t))^\top \mathbf{V},$$

and observe that

$$\begin{aligned} \frac{\delta^2}{N^2} \mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t))^\top \mathbf{V} &\preceq \frac{\delta^2}{N^2} \|\mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t)\|^2 \mathbf{I}_r, \\ \frac{\delta^2}{N^2} \mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t))^\top \mathbf{V} &\succeq -\frac{\delta^2}{N^2} \|\mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t)\|^2 \mathbf{I}_r. \end{aligned}$$

We therefore want to estimate the operator norm of $\mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t)$. From (4.14), we have that

$$\mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) = 2\lambda \mathbf{V}^\top \Pi_{\mathbf{X}_{t-1}}^{\text{St}}(\mathbf{W}^t \mathbf{X}_{t-1}) - 2\sqrt{N} \lambda^2 \mathbf{M}_{t-1} (\mathbf{I}_r - \mathbf{M}_{t-1}^\top \mathbf{M}_{t-1}),$$

which gives the bound

$$\|\mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t)\| \leq 2\lambda \|\mathbf{V}^\top \Pi_{\mathbf{X}_{t-1}}^{\text{St}}(\mathbf{W}^t \mathbf{X}_{t-1})\| + 2\sqrt{N} \lambda^2 \|\mathbf{M}_{t-1}\|, \quad (4.16)$$

where we used the fact that $\|\mathbf{I}_r - \mathbf{M}_{t-1}^\top \mathbf{M}_{t-1}\| \leq 1$, since $\mathbf{M}_{t-1}^\top \mathbf{M}_{t-1}$ is positive semi-definite with its operator norm bounded above by 1. Additionally, we can bound the operator norm of $\mathbf{V}^\top \Pi_{\mathbf{X}_{t-1}}^{\text{St}}(\mathbf{W}^t \mathbf{X}_{t-1})$ as follows:

$$\begin{aligned} \|\mathbf{V}^\top \Pi_{\mathbf{X}_{t-1}}^{\text{St}}(\mathbf{W}^t \mathbf{X}_{t-1})\| &= \|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1} - \mathbf{V}^\top \mathbf{X}_{t-1} \text{sym}(\mathbf{X}_{t-1}^\top \mathbf{W}^t \mathbf{X}_{t-1})\| \\ &\leq \|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\| + \|\mathbf{X}_{t-1}^\top \mathbf{W}^t \mathbf{X}_{t-1}\|, \end{aligned} \quad (4.17)$$

where we used the fact that $\|\mathbf{V}^\top \mathbf{X}_{t-1}\| = \|\mathbf{M}_{t-1}\| \leq 1$ and the fact that the matrix \mathbf{W}^t is symmetric. Combining (4.16) and (4.17) yields

$$\begin{aligned} \|\mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t)\|^2 &\leq 2 \left(4\lambda^2 \|\mathbf{V}^\top \Pi_{\mathbf{X}_{t-1}}^{\text{St}}(\mathbf{W}^t \mathbf{X}_{t-1})\|^2 + 4N\lambda^4 \|\mathbf{M}_{t-1}\|^2 \right) \\ &\leq 8\lambda^2 \left(2\|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\|^2 + 2\|\mathbf{X}_{t-1}^\top \mathbf{W}^t \mathbf{X}_{t-1}\|^2 + N\lambda^2 \|\mathbf{G}_{t-1}\| \right), \end{aligned} \quad (4.18)$$

where we used the fact that $(a+b)^2 \leq 2(a^2 + b^2)$.

Now, consider the higher order terms in δ from (4.15), i.e.,

$$\sum_{k=1}^{\infty} (-1)^k \frac{\delta^{2k}}{N^{2k}} \mathbf{B}_t^\top \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)^k \mathbf{B}_t.$$

Note first that for every $k \geq 1$,

$$-\|\mathcal{G}\|^{k-1} \mathbf{B}_t^\top \mathcal{G} \mathbf{B}_t \preceq \mathbf{B}_t^\top \mathcal{G}^k \mathbf{B}_t \preceq \|\mathcal{G}\|^{k-1} \mathbf{B}_t^\top \mathcal{G} \mathbf{B}_t.$$

Indeed, since \mathcal{G} is positive semi-definite, for every $\mathbf{y} \in \mathbb{R}^r$ we can write

$$\begin{aligned} \mathbf{y}^\top \mathbf{B}_t^\top \left(\mathcal{G}^k - \|\mathcal{G}\|^{k-1} \mathcal{G} \right) \mathbf{B}_t \mathbf{y} &= \mathbf{y}^\top \mathbf{B}_t^\top \mathcal{G}^{1/2} \left(\mathcal{G}^{k-1} - \|\mathcal{G}\|^{k-1} \mathbf{I}_r \right) \mathcal{G}^{1/2} \mathbf{B}_t \mathbf{y} \\ &\leq \mathbf{z}^\top \left(\mathcal{G}^{k-1} - \|\mathcal{G}\|^{k-1} \mathbf{I}_r \right) \mathbf{z} \leq 0, \end{aligned}$$

where $\mathbf{z} = \mathcal{G}^{1/2} \mathbf{B}_t \mathbf{y}$. The other inequality follows similarly. This implies that

$$-\sum_{k=1}^{\infty} (-1)^k \frac{\delta^{2k}}{N^{2k}} \|\mathcal{G}\|^{k-1} \mathbf{B}_t^\top \mathcal{G} \mathbf{B}_t \preceq \sum_{k=1}^{\infty} (-1)^k \frac{\delta^{2k}}{N^{2k}} \mathbf{B}_t^\top \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)^k \mathbf{B}_t \preceq \sum_{k=1}^{\infty} (-1)^k \frac{\delta^{2k}}{N^{2k}} \|\mathcal{G}\|^{k-1} \mathbf{B}_t^\top \mathcal{G} \mathbf{B}_t,$$

where

$$\sum_{k=1}^{\infty} (-1)^k \frac{\delta^{2k}}{N^{2k}} \|\mathcal{G}\|^{k-1} \mathbf{B}_t^\top \mathcal{G} \mathbf{B}_t = - \left(\sum_{k=0}^{\infty} (-1)^k \|\frac{\delta^2}{N^2} \mathcal{G}\|^k \right) \frac{\delta^2}{N^2} \mathbf{B}_t^\top \mathcal{G} \mathbf{B}_t = - \frac{1}{1 + \alpha_2 \frac{\delta^2}{N}} \frac{\delta^2}{N^2} \mathbf{B}_t^\top \mathcal{G} \mathbf{B}_t,$$

where we used the fact that $\max \|\frac{\delta^2}{N^2} \mathcal{G}\| \leq \alpha_2 \frac{\delta^2}{N}$ on the event $\mathcal{E}_{\mathcal{G}}$. Therefore, we have that

$$\begin{aligned} \sum_{k=1}^{\infty} (-1)^k \frac{\delta^{2k}}{N^{2k}} \mathbf{B}_t^\top \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)^k \mathbf{B}_t &\preceq - \frac{1}{1 + \alpha_2 \frac{\delta^2}{N}} \frac{\delta^2}{N^2} \mathbf{B}_t^\top \mathcal{G} \mathbf{B}_t, \\ \sum_{k=1}^{\infty} (-1)^k \frac{\delta^{2k}}{N^{2k}} \mathbf{B}_t^\top \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)^k \mathbf{B}_t &\succeq \frac{1}{1 + \alpha_2 \frac{\delta^2}{N}} \frac{\delta^2}{N^2} \mathbf{B}_t^\top \mathcal{G} \mathbf{B}_t, \end{aligned} \tag{4.19}$$

It remains to bound the operator norm of $\mathbf{B}_t^\top \mathcal{G} \mathbf{B}_t$:

$$\begin{aligned} &\frac{\delta^2}{N^2} \mathbf{B}_t^\top \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \mathbf{B}_t \\ &= \frac{\delta^2}{N^2} \mathbf{M}_{t-1} \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \mathbf{M}_{t-1}^\top - \frac{2\delta^3}{N^3} \text{sym} \left(\mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \mathbf{M}_{t-1}^\top \right) \\ &\quad + \frac{\delta^4}{N^4} \mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t) (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t))^\top \mathbf{V}. \end{aligned}$$

The operator norm of the first summand is bounded above by

$$\begin{aligned} \frac{\delta^2}{N^2} \|\mathbf{M}_{t-1} \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \mathbf{M}_{t-1}^\top\| &= \frac{\delta^2}{N^2} \|\mathbf{M}_{t-1} \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t)\|^2 \\ &\leq \frac{\delta^2}{N^2} \|\mathbf{M}_{t-1}\|^2 \|\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t)\|^2 \\ &\leq \alpha_2 \frac{\delta^2}{N} \|\mathbf{G}_{t-1}\|, \end{aligned}$$

where we used the fact that $\|\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t)\|^2 = \|\mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)\|$. For the second summand, we have that

$$\begin{aligned} &\frac{2\delta^3}{N^3} \|\text{sym} \left(\mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \mathbf{M}_{t-1}^\top \right)\| \\ &\leq \frac{2\delta^3}{N^3} \|\mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)\| \|\mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t)\| \|\mathbf{M}_{t-1}\| \\ &\leq \frac{2\alpha_2 \delta^3}{N^2} \left(2\lambda \|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\| + 2\lambda \|\mathbf{X}_{t-1}^\top \mathbf{W}^t \mathbf{X}_{t-1}\| + 2\sqrt{N} \lambda^2 \|\mathbf{M}_{t-1}\| \right) \|\mathbf{M}_{t-1}\|, \end{aligned}$$

where the last inequality follows by (4.16) and (4.17). Finally, the third term is bounded in operator norm as follows:

$$\begin{aligned} &\frac{\delta^4}{N^4} \|\mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) \mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t) (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t))^\top \mathbf{V}\| \\ &= \frac{\delta^4}{N^4} \|\mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t) (\nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t))^\top\|^2 \\ &\leq \frac{\delta^4}{N^4} \|\mathbf{V}^\top \nabla_{\text{St}} \mathcal{L}(\mathbf{X}_{t-1}; \mathbf{Y}^t)\|^2 \|\mathcal{G}(\mathbf{X}_{t-1}; \mathbf{Y}^t)\| \\ &\leq 8\lambda^2 \frac{\alpha_2 \delta^4}{N^3} \left(2\|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\|^2 + 2\|\mathbf{X}_{t-1}^\top \mathbf{W}^t \mathbf{X}_{t-1}\|^2 + N\lambda^2 \|\mathbf{G}_{t-1}\| \right), \end{aligned}$$

where we used (4.18) for the last inequality.

Combining the last four estimates with (4.19) and from (4.15) and (4.18), we finally obtain the partial ordering

$$\begin{aligned} \mathbf{G}_t &\preceq \mathbf{G}_{t-1} + \frac{4\lambda^2\delta}{\sqrt{N}}\mathbf{G}_{t-1}(\mathbf{I}_r - \mathbf{G}_{t-1}) - \frac{4\lambda\delta}{N}\text{sym}\left(\mathbf{V}^\top \Pi_{\mathbf{X}_{t-1}}^{\text{St}}(\mathbf{W}^t \mathbf{X}_{t-1}) \mathbf{M}_{t-1}^\top\right) \\ &\quad + \frac{8\lambda^2\delta^2}{N^2}\left(1 + \alpha_2 \frac{\delta^2}{N} \frac{1}{1 + \alpha_2 \frac{\delta^2}{N}}\right)\left(2\|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\|^2 + 2\|\mathbf{X}_{t-1}^\top \mathbf{W}^t \mathbf{X}_{t-1}\|^2 + N\lambda^2\|\mathbf{G}_{t-1}\|\right)\mathbf{I}_r \\ &\quad + \frac{1}{1 + \alpha_2 \frac{\delta^2}{N}}\alpha_2 \frac{\delta^2}{N}\|\mathbf{G}_{t-1}\|\mathbf{I}_r \\ &\quad + \frac{4\lambda\delta^3}{N^2}\frac{\alpha_2}{1 + \alpha_2 \frac{\delta^2}{N}}\left(\|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\| + \|\mathbf{X}_{t-1}^\top \mathbf{W}^t \mathbf{X}_{t-1}\| + \sqrt{N}\lambda\|\mathbf{G}_{t-1}\|^{1/2}\right)\|\mathbf{G}_{t-1}\|^{1/2}\mathbf{I}_r. \end{aligned}$$

Iterating the above expression yields

$$\mathbf{G}_t \preceq \mathbf{G}_0 + \frac{4\lambda^2\delta}{\sqrt{N}}\sum_{\ell=1}^t(\mathbf{G}_{\ell-1}(\mathbf{I}_r - \mathbf{G}_{\ell-1})) - \frac{4\lambda\delta}{N}\sum_{\ell=1}^t\text{sym}\left(\mathbf{V}^\top \Pi_{\mathbf{X}_{\ell-1}}^{\text{St}}(\mathbf{W}^\ell \mathbf{X}_{\ell-1}) \mathbf{M}_{\ell-1}^\top\right) + \sum_{\ell=1}^t \mathbf{A}'_\ell,$$

where \mathbf{A}'_ℓ is given by

$$\begin{aligned} \mathbf{A}'_\ell &= \frac{8\lambda^2\delta^2}{N^2}\left(1 + \alpha_2 \frac{\delta^2}{N} \frac{1}{1 + \alpha_2 \frac{\delta^2}{N}}\right)\left(2\|\mathbf{V}^\top \mathbf{W}^\ell \mathbf{X}_{\ell-1}\|^2 + 2\|\mathbf{X}_{\ell-1}^\top \mathbf{W}^\ell \mathbf{X}_{\ell-1}\|^2 + N\lambda^2\|\mathbf{G}_{\ell-1}\|\right)\mathbf{I}_r \\ &\quad + \frac{1}{1 + \alpha_2 \frac{\delta^2}{N}}\alpha_2 \frac{\delta^2}{N}\|\mathbf{G}_{\ell-1}\|\mathbf{I}_r \\ &\quad + \frac{4\lambda\delta^3}{N^2}\frac{\alpha_2}{1 + \alpha_2 \frac{\delta^2}{N}}\left(\|\mathbf{V}^\top \mathbf{W}^\ell \mathbf{X}_{\ell-1}\| + \|\mathbf{X}_{\ell-1}^\top \mathbf{W}^\ell \mathbf{X}_{\ell-1}\| + \sqrt{N}\lambda\|\mathbf{G}_{\ell-1}\|^{1/2}\right)\|\mathbf{G}_{\ell-1}\|^{1/2}\mathbf{I}_r. \end{aligned}$$

The lower bound follows similarly.

The remaining of the proof focuses on deriving a high-probability estimate for $\|\mathbf{V}^\top \mathbf{W}^\ell \mathbf{X}_{\ell-1}\|$ and $\|\mathbf{X}_{\ell-1}^\top \mathbf{W}^\ell \mathbf{X}_{\ell-1}\|$ for every $\ell \leq t$. Let $\mathcal{F}_t = \sigma(\mathbf{X}_\ell, \ell \leq t)$ denote the canonical filtration. Then, for every $\eta > 0$ it holds that

$$\begin{aligned} \mathbb{E}\left[\exp\left(\eta\left(\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\right)_{ij}\right) \middle| \mathcal{F}_{t-1}\right] &= \mathbb{E}\left[\exp\left(\eta\sum_{k=1}^r\sum_{\ell=1}^r(\mathbf{v}_i)_k \mathbf{W}_{k\ell}^t(\mathbf{x}_j^{t-1})_\ell\right) \middle| \mathcal{F}_{t-1}\right] \\ &= \prod_{k,\ell} \mathbb{E}\left[\exp\left(\eta(\mathbf{v}_i)_k \mathbf{W}_{k\ell}^t(\mathbf{x}_j^{t-1})_\ell\right) \middle| \mathcal{F}_{t-1}\right] \\ &\leq \prod_{k,\ell} \exp\left(C\eta^2\sigma^2(\mathbf{v}_i)_k^2(\mathbf{x}_j^{t-1})_\ell^2\right) \quad \text{for some absolute constant } C \\ &\leq \exp\left(C\eta^2\sigma^2\right), \end{aligned}$$

where we used the fact that the entries of \mathbf{W}^{t-1} are i.i.d. sub-Gaussian satisfying (4.1) and independent from \mathcal{F}_{t-1} , and the fact that $\mathbf{v}_i, \mathbf{x}_j \in \text{St}(N, r)$ for every $i, j \in [r]$. The entries of the matrix $\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}$ given \mathcal{F}_{t-1} are therefore sub-Gaussian with sub-Gaussian norm σ^2 . We have the following sub-Gaussian tail bound: for every $\gamma > 0$,

$$\mathbb{P}\left(\max_{1 \leq i, j \leq r} |(\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1})_{ij}| \geq \gamma \middle| \mathcal{F}_{t-1}\right) \leq 2r^2 \exp\left(-\frac{\gamma^2}{C\sigma^2}\right).$$

Since $\|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\| \leq r \max_{1 \leq i, j \leq r} |(\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1})_{ij}|$, for every $K_N \geq 0$ it holds that

$$\mathbb{P}\left(\|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\| \geq rK_N \middle| \mathcal{F}_{t-1}\right) \leq 2r^2 \exp\left(-\frac{K_N^2}{C\sigma^2}\right).$$

Finally, an union bound over the time horizon T yields

$$\mathbb{P}\left(\bigcup_{1 \leq t \leq T} \|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\| \geq rK_N \middle| \mathcal{F}_{t-1}\right) \leq 2r^2 T \exp\left(-\frac{K_N^2}{C\sigma^2}\right). \quad (4.20)$$

We derive the same bound for $\|\mathbf{X}_{t-1}^\top \mathbf{W}^t \mathbf{X}_{t-1}\|$ using a similar argument. Thus, on the event $\mathcal{E}_{\mathcal{G}}$, we have for every $\ell \leq t$,

$$\begin{aligned} \mathbf{A}'_\ell &\leq \frac{8\lambda^2\delta^2}{N^2} (4r^2K_N^2 + N\lambda^2\|\mathbf{G}_{\ell-1}\|) \mathbf{I}_r + \alpha_2 \frac{\delta^2}{N} \|\mathbf{G}_{\ell-1}\| \mathbf{I}_r \\ &\quad + \frac{4\lambda\delta^3}{N^2} \alpha_2 \left(2rK_N + \sqrt{N}\lambda\|\mathbf{G}_{\ell-1}\|^{1/2} \right) \|\mathbf{G}_{\ell-1}\|^{1/2} \mathbf{I}_r, \end{aligned}$$

with \mathbb{P} -probability at least $1 - 4r^2T \exp(-K_N^2/(C\sigma^2))$. This completes the proof of Lemma 4.7. \square

The following lemma gives a high-probability estimate for \mathbf{S}_t .

Lemma 4.8. *There exists a constant c depending only on r and σ^2 such that for every scalar, positive, continuous functions a, b and any fixed time horizon T ,*

$$\mathbb{P}(\exists t \in [0, T]: \|\mathbf{S}_t\| \geq a(t) \text{ and } \lambda_{\max}(\mathbf{Q}_t) \leq b(t)) \leq \inf_{\theta > 0} \sup_{t \in [0, T]} 2r \exp\left(-\theta a(t) + c \frac{\theta^2}{2} b(t)\right),$$

where \mathbf{Q}_t is given by $\mathbf{Q}_t = \left(\frac{4\lambda\delta}{N}\right)^2 \sum_{\ell=1}^t \|\mathbf{G}_{\ell-1}\| \mathbf{I}_r$, thus $\lambda_{\max}(\mathbf{Q}_t) = \left(\frac{4\lambda\delta}{N}\right)^2 \sum_{\ell=1}^t \|\mathbf{G}_{\ell-1}\|$. In particular, if a is continuously differentiable with non-vanishing derivative and $b(t) = Ka^2(t)$ for some constant K , then it holds that

$$\mathbb{P}(\exists t \in [0, T]: \|\mathbf{S}_t\| \geq a(t) \text{ and } \lambda_{\max}(\mathbf{Q}_t) \leq b(t)) \leq 2r \exp\left(-\frac{c}{2K}\right).$$

Proof. We let \mathcal{F}_t denote the natural filtration associated to the online SGD outputs $\mathbf{X}_1, \dots, \mathbf{X}_t$. Since for every $t \geq 0$,

$$\mathbb{E}[\mathbf{S}_t - \mathbf{S}_{t-1} | \mathcal{F}_{t-1}] = \frac{4\lambda\delta}{N} \mathbb{E}\left[\text{sym}\left(\mathbf{V}^\top \Pi_{\mathbf{X}_{t-1}}^{\text{St}}(\mathbf{W}^t \mathbf{X}_{t-1}) \mathbf{M}_{t-1}^\top\right) | \mathcal{F}_{t-1}\right] = 0,$$

the process \mathbf{S}_t is an \mathcal{F}_t -adapted matrix martingale. Moreover, from (4.17) it follows that

$$\|\mathbf{V}^\top \Pi_{\mathbf{X}_{t-1}}^{\text{St}}(\mathbf{W}^t \mathbf{X}_{t-1}) \mathbf{M}_{t-1}^\top\| \leq \left(\|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\| + \|\mathbf{X}_{t-1}^\top \mathbf{W}^t \mathbf{X}_{t-1}\|\right) \|\mathbf{M}_{t-1}^\top\|.$$

Using the sub-Gaussian tail for the entries of $\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}$ established in the proof of Lemma 4.7, we have for every $k \geq 1$,

$$\begin{aligned} \mathbb{E}\left[\|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\|^k | \mathcal{F}_{t-1}\right] &= \int_0^\infty ky^{k-1} \mathbb{P}\left(\|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\| \geq y | \mathcal{F}_{t-1}\right) dy \\ &\leq 2r^2k \int_0^\infty y^{k-1} \exp\left(-\frac{y^2}{C\sigma^2r^2}\right) dy \\ &= 2r^2k \left(\frac{Cr^2\sigma^2}{2}\right)^{k/2} 2^{k/2-1} \Gamma\left(\frac{k}{2}\right) \\ &\leq 3r^2(Cr^2\sigma^2)^{k/2} k \left(\frac{k}{2}\right)^{k/2}, \end{aligned}$$

where we used the fact that $\Gamma(x) \leq 3x^x$ for every $x \geq 1/2$. We obtain the same upper bound for $\mathbb{E}\left[\|\mathbf{X}_{t-1}^\top \mathbf{W}^t \mathbf{X}_{t-1}\|^k | \mathcal{F}_{t-1}\right]$. Using the fact that $(a+b)^k \leq 2^{k-1}(a^k + b^k) \leq 2^k(a^k + b^k)$, it therefore follows that

$$\begin{aligned} \mathbb{E}\left[\|(\mathbf{S}_t - \mathbf{S}_{t-1})^k\| | \mathcal{F}_{t-1}\right] &= \left(\frac{4\lambda\delta}{N}\right)^k \mathbb{E}\left[\|\text{sym}\left(\mathbf{V}^\top \Pi_{\mathbf{X}_{t-1}}^{\text{St}}(\mathbf{W}^t \mathbf{X}_{t-1}) \mathbf{M}_{t-1}^\top\right)\|^k | \mathcal{F}_{t-1}\right] \\ &\leq \left(\frac{8\lambda\delta}{N}\right)^k 3r^2(Cr^2\sigma^2)^{k/2} k \left(\frac{k}{2}\right)^{k/2} \|\mathbf{G}_{t-1}\|^{k/2}. \end{aligned}$$

According to the Taylor series expansion of the exponential function, we obtain for every $z \geq 0$ that

$$\begin{aligned} \mathbb{E} \left[\exp \left(z^2 (\mathbf{S}_t - \mathbf{S}_{t-1})^2 \right) | \mathcal{F}_{t-1} \right] &= \mathbf{I}_r + \sum_{k=1}^{\infty} \frac{z^{2k}}{k!} \mathbb{E} \left[(\mathbf{S}_t - \mathbf{S}_{t-1})^{2k} | \mathcal{F}_{t-1} \right] \\ &\preceq \mathbf{I}_r + \sum_{k=1}^{\infty} \frac{z^{2k}}{k!} \mathbb{E} \left[\| (\mathbf{S}_t - \mathbf{S}_{t-1})^{2k} \| \mathbf{I}_r | \mathcal{F}_{t-1} \right] \\ &\preceq \left(1 + \sum_{k=1}^{\infty} \frac{z^{2k}}{k!} \left(\frac{8\lambda\delta}{N} \right)^{2k} 12r^2 (Cr^2\sigma^2)^k k^{k+1} \| \mathbf{G}_{t-1} \|^k \right) \mathbf{I}_r. \end{aligned}$$

We then apply similar arguments to that used for the proof of [64, Proposition 2.5.2], in particular $k! \geq \left(\frac{k}{e}\right)^k$, to obtain that

$$\begin{aligned} 1 + \sum_{k=1}^{\infty} \frac{z^{2k}}{k!} \left(\frac{8\lambda\delta}{N} \right)^{2k} 12r^2 (Cr^2\sigma^2)^k k^{k+1} \| \mathbf{G}_{t-1} \|^k &\leq \sum_{k=0}^{\infty} \left(z^2 \tilde{C}(r, \sigma^2) \left(\frac{4\lambda\delta}{N} \right)^2 \| \mathbf{G}_{t-1} \| \right)^k \\ &= \frac{1}{1 - z^2 \tilde{C}(r, \sigma^2) \left(\frac{8\lambda\delta}{N} \right)^2 \| \mathbf{G}_{t-1} \|} \end{aligned}$$

where the last equality is valid for every $z^2 < \frac{1}{\tilde{C}(r, \sigma^2) \left(\frac{8\lambda\delta}{N} \right)^2}$ since $\| \mathbf{G}_{t-1} \| \leq 1$. Then, since $\frac{1}{1-x} \leq e^{2x}$ for every $x \in [0, \frac{1}{2}]$, we have for every $z^2 \leq \frac{1}{2\tilde{C}(r, \sigma^2) \left(\frac{8\lambda\delta}{N} \right)^2}$,

$$\mathbb{E} \left[\exp \left(z^2 (\mathbf{S}_t - \mathbf{S}_{t-1})^2 \right) | \mathcal{F}_{t-1} \right] \preceq \exp \left(2z^2 \tilde{C}(r, \sigma^2) \left(\frac{8\lambda\delta}{N} \right)^2 \| \mathbf{G}_{t-1} \| \mathbf{I}_r \right).$$

Using again a similar argument in the proof of Propostion 2.5.2 in [64] along with a use of the transfer rule for matrix functions to obtain that the numeric inequality $e^x \leq x + e^{2x}$ for all $x \in \mathbb{R}$ (see e.g. Section 2 of [62]) gives

$$\begin{aligned} \mathbb{E} [\exp (z (\mathbf{S}_t - \mathbf{S}_{t-1})) | \mathcal{F}_{t-1}] &\preceq \mathbb{E} [z (\mathbf{S}_t - \mathbf{S}_{t-1}) | \mathcal{F}_{t-1}] + \mathbb{E} \left[\exp \left(z^2 (\mathbf{S}_t - \mathbf{S}_{t-1})^2 \right) | \mathcal{F}_{t-1} \right] \\ &\preceq \exp \left(2z^2 \tilde{C}(r, \sigma^2) \left(\frac{8\lambda\delta}{N} \right)^2 \| \mathbf{G}_{t-1} \| \mathbf{I}_r \right), \end{aligned}$$

for every $|z| \leq \left(\frac{1}{2\tilde{C}(r, \sigma^2) \left(\frac{8\lambda\delta}{N} \right)^2} \right)^{1/2}$, where we used the fact that \mathbf{S}_t is a \mathcal{F}_t -adapted matrix martingale.

The validity of the bound for $|z| > \left(\frac{1}{2\tilde{C}(r, \sigma^2) \left(\frac{8\lambda\delta}{N} \right)^2} \right)^{1/2}$ is obtained by adapting similarly the argument of [64] mentioned above, up to an absolute constant that we absorb in $\tilde{C}(r, \sigma^2)$. Finally, for every $t \geq 0$, we have that

$$\mathbb{E} [\exp (z (\mathbf{S}_t - \mathbf{S}_{t-1})) | \mathcal{F}_{t-1}] \preceq \exp \left(\tilde{C}(r, \sigma^2) z^2 \left(\frac{4\lambda\delta}{N} \right)^2 \| \mathbf{G}_{t-1} \| \mathbf{I}_r \right),$$

where we absorbed a factor 4 in the constant $\tilde{C}(r, \sigma^2)$ for convenience. According to Lemma B.1, for every positive scalar-valued, continuous functions a and b , we obtain that

$$\mathbb{P} (\exists t \geq 0: \lambda_{\max} (\mathbf{S}_t) \geq a(t) \text{ and } \lambda_{\max} (\mathbf{Q}_t) \leq b(t)) \leq r \inf_{\theta > 0} \sup_{t \in [0, T]} \exp \left(-\theta a(t) + \tilde{C}(r, \sigma^2) \frac{\theta^2}{2} b(t) \right).$$

The other side of the inequality is obtained in a similar fashion. To prove the second statement, we consider the optimization problem

$$\sup_{t \in [0, T]} -\theta a(t) + \tilde{C}(r, \sigma^2) \frac{\theta^2}{2} K a^2(t).$$

Setting the derivative with respect to t to zero and using the assumption that a has non-vanishing derivative, we obtain that

$$a(t) = \frac{1}{\tilde{C}(r, \sigma^2) \theta K}.$$

By the regularity assumptions on a , the above equation has a unique solution given by

$$t_* = a^{-1} \left(\frac{1}{\tilde{C}(r, \sigma^2) \theta K} \right).$$

We therefore deduce that for every value of $\theta > 0$, the function $t \mapsto -\theta a(t) + \tilde{C}(r, \sigma^2) \frac{\theta^2}{2} K a^2(t)$ attains its maximum at either $t_{\max} = t_*$ or at the boundaries $t_{\max} = 0$ or $t_{\max} = T$. In both cases, plugging the corresponding value of t_{\max} and solving in θ yields the desired result. \square

5. PROOFS OF FULL RECOVERY OF SPIKES

In this section, we present the proofs of Theorems 3.5 and 3.6. We first address the case $p \geq 3$, followed by $p = 2$ with SNRs separated by constants of order 1. Once the noise from the martingale terms is controlled, we rely on the continuous time analysis carried out for Langevin dynamics and gradient flow in [6]. To do this, we need to ensure that the step-size is sufficiently small for the discrete dynamics to stay close, in a sense that we quantify, to the continuous dynamics, while also controlling the additional interaction between the correlations induced by the retraction step. After detailing these steps, we outline the remainder of the proof and refer to the extensive version of our preprint [6].

Throughout this section, for every $\varepsilon > 0$ and $i, j \in [r]$, we define the hitting times $\mathcal{T}_\varepsilon^{(ij,+)}$ and $\mathcal{T}_\varepsilon^{(ij,-)}$ as follows:

$$\begin{aligned} \mathcal{T}_\varepsilon^{(ij,+)} &= \min\{t \geq 0 : m_{ij}(\mathbf{X}_t) \geq \varepsilon\}, \\ \mathcal{T}_\varepsilon^{(ij,-)} &= \min\{t \geq 0 : m_{ij}(\mathbf{X}_t) \leq \varepsilon\}. \end{aligned} \quad (5.1)$$

5.1. PROOFS FOR $p \geq 3$

In this subsection, we let $p \geq 3$ and assume $\lambda_1 \geq \dots \geq \lambda_r$ to be of order 1. Additionally, let $(i_1^*, j_1^*), \dots, (i_r^*, j_r^*)$ denote the greedy maximum selection of $\mathbf{I}_0 = (\lambda_i \lambda_j m_{ij}(\mathbf{X}_0)^{p-2} \mathbf{1}_{\{m_{ij}(\mathbf{X}_0)^{p-2} \geq 0\}})_{1 \leq i, j \leq r}$ according to Definition 1.4. Before proving Theorem 3.5, we first focus on the recovery of the first direction.

Lemma 5.1. *For every $\gamma_1, \gamma_2, \gamma > 0$ with $\gamma_1 > \gamma_2$, $\gamma_1 > \gamma$, and for every $\varepsilon > 0$, there exist $c_0 < \frac{(1 + \frac{\gamma}{\gamma_1})^{\frac{1}{p-1}} - 1}{(1 + \frac{\gamma}{\gamma_1})^{\frac{1}{p-1}} + 1}$ and $d_0 = d_0(N)$, which goes to zero arbitrarily slowly as $N \rightarrow \infty$, such that if*

$$\delta = C_\delta d_0 \frac{N^{-\frac{p-3}{2}}}{\log(N)}, \quad (5.2)$$

for some constant $C_\delta > 0$, then for sufficiently large N ,

$$\inf_{\mathbf{X}_0 \in \mathcal{C}_1(\gamma_1, \gamma_2) \cap \mathcal{C}_2(\gamma_1, \gamma)} \mathbb{P}_{\mathbf{X}_0} \left(\mathcal{T}_{1-\varepsilon}^{(i_1^* j_1^*)} \lesssim \frac{(N^{p-2} + N^{\frac{p-2}{2}}) \log(N)}{d_0} \right) \geq 1 - \eta,$$

where $\eta = K_1 \frac{\log(N)^2}{d_0} N^{p-2} e^{-c_1 N} + K_2 \log(N) e^{-c_2 c_0^2 \log(N)/d_0} + K_3 \frac{d_0^2}{\log(N)} N^{-\frac{p-2}{2}} e^{-c_3/d_0^2}$.

The proof strategy for Lemma 5.1 involves two steps. First, we prove weak recovery of the correlation $m_{i_1^* j_1^*}$ and show that, when this occurs, the other correlations are still on the scale $\Theta(N^{-1/2})$. Second, we prove strong recovery starting from weak recovery. Lemma 5.1 then follows by combining these two results together with the strong Markov property, as outlined later.

Lemma 5.2. *For every $\varepsilon > 0$, let $E(\varepsilon)$ denote the event*

$$E(\varepsilon) = \left\{ \mathbf{X} : m_{i_1^* j_1^*}(\mathbf{X}) \geq \varepsilon \text{ and } m_{ij}(\mathbf{X}) \in \Theta(N^{-1/2}) \text{ for all } (i, j) \neq (i_1^*, j_1^*) \right\}. \quad (5.3)$$

Then, the following holds: For every $\gamma_1 > 1, \gamma_2 > 0, \gamma > 0$ with $\gamma_1 > \gamma_2$, $\gamma_1 > \gamma$, and for every $\varepsilon > 0$, there exist $c_0 < \frac{(1 + \frac{\gamma}{\gamma_1})^{\frac{1}{p-1}} - 1}{(1 + \frac{\gamma}{\gamma_1})^{\frac{1}{p-1}} + 1}$ and $d_0 = d_0(N)$, which goes to zero arbitrarily slowly as $N \rightarrow \infty$, such that if

$$\delta = C_\delta d_0 \frac{N^{-\frac{p-3}{2}}}{\log(N)},$$

for some constant $C_\delta > 0$, then for sufficiently large N ,

$$\inf_{\mathbf{X}_0 \in \mathcal{C}_1(\gamma_1, \gamma_2) \cap \mathcal{C}_2(\gamma_1, \gamma)} \mathbb{P}_{\mathbf{X}_0} \left(\mathcal{T}_{E(\varepsilon)} \lesssim \frac{N^{p-2} \log(N)}{d_0} \right) \geq 1 - \eta,$$

where η is given as in Lemma 5.1.

Proof. Assume that \mathbf{X}_0 satisfies Condition 1 and Condition 2. In particular, for every $i, j \in [r]$, there exists $\gamma_{ij} \in (\gamma_2, \gamma_1)$ such that $m_{ij}(0) = \gamma_{ij} N^{-\frac{1}{2}}$. Let $T_0 > 0$ be a fixed time horizon to be determined later. From Proposition 4.3 with the step-size prescribed by (5.2) and the time horizon T_0 , we obtain for every $t \leq T_0$,

$$\mathbb{P}_{\mathbf{X}_0} (\ell_{ij}(t) \leq m_{ij}(t) \leq u_{ij}(t)) \geq 1 - \eta(\delta, c_0, T_0, K_N),$$

where $\eta(\delta, c_0, T_0, K_N)$ is given by (4.5). The comparison functions ℓ_{ij} and u_{ij} are given by

$$\begin{aligned} \ell_{ij}(t) &= (1 - c_0) \frac{\gamma_{ij}}{\sqrt{N}} + \frac{\delta}{N} \sum_{\ell=1}^t \{ -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle - |h_i(\ell-1)| \} \\ u_{ij}(t) &= (1 + c_0) \frac{\gamma_{ij}}{\sqrt{N}} + \frac{\delta}{N} \sum_{\ell=1}^t \{ -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle + |h_i(\ell-1)| \} \end{aligned}$$

respectively, where for all $0 \leq t \leq T_0 - 1$,

$$-\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_j \rangle = \sqrt{N} p \lambda_i \lambda_j m_{ij}^{p-1}(t) - \sqrt{N} \frac{p}{2} \sum_{1 \leq k, \ell \leq r} \lambda_k m_{i\ell}(t) m_{kj}(t) m_{k\ell}(t) \left(\lambda_j m_{kj}^{p-2}(t) + \lambda_\ell m_{k\ell}^{p-2}(t) \right),$$

and

$$|h_i(t)| = \frac{\delta \alpha_2}{2} \sum_{k=1}^r |m_{ik}(t)| + \frac{\delta^2 \alpha_2}{2N} \sum_{k=1}^r \langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_j \rangle + \frac{\delta^2 \alpha_2 r K_N}{2N}.$$

In the following, we introduce the microscopic thresholds $\varepsilon_N^+ = \frac{\gamma_+}{\sqrt{N}}$ and $\varepsilon_N^- = \frac{\gamma_-}{\sqrt{N}}$ for some constants $\gamma_+ > \gamma_1$ and $\gamma_- < \gamma_2$. We then refer to (5.1) for the definition of the hitting times $\mathcal{T}_{\varepsilon_N^+}^{(ij,+)}$ and $\mathcal{T}_{\varepsilon_N^-}^{(ij,-)}$. Let $\mathcal{T}_{\tilde{E}}$ denote the time to reach the event

$$\tilde{E} = \left\{ \mathbf{X} : m_{i_1^* j_1^*}(\mathbf{X}) \geq \varepsilon_N^+ > \frac{1}{2} \max_{(i,j) \neq (i_1^*, j_1^*)} m_{ij}(\mathbf{X}) \right\}.$$

We claim that

$$\inf_{\mathbf{X}_0 \in \mathcal{C}_1(\gamma_1, \gamma_2) \cap \mathcal{C}_2(\gamma_1, \gamma)} \mathbb{P}_{\mathbf{X}_0} \left(\mathcal{T}_{\tilde{E}} \leq \frac{1 - \left(\frac{(1-c_0)\gamma_{i_1^* j_1^*}}{\gamma_+} \right)^{p-2}}{C_\delta d_0 (1-c_0)^{p-1} p \lambda_{i_1^*} \lambda_{j_1^*} \gamma_2^{p-2}} \log(N) N^{p-2} \right) \geq 1 - \tilde{\eta}, \quad (5.4)$$

and that

$$\inf_{\mathbf{X} \in \tilde{E}} \mathbb{P}_{\mathbf{X}} \left(\mathcal{T}_{E(\varepsilon)} \leq \frac{1 - \left(\frac{(1-c_0)\gamma_+}{\varepsilon \sqrt{N}} \right)^{p-2}}{C_\delta d_0 (1-c_0)^{p-1} p \lambda_{i_1^*} \lambda_{j_1^*} \gamma_+^{p-2}} \log(N) N^{p-2} \right) \geq 1 - \log(N) \tilde{\eta}, \quad (5.5)$$

where $\tilde{\eta} = K_1 \frac{\log(N)}{d_0} N^{p-2} e^{-c_1 N} + K_2 e^{-c_2 c_0^2 \log(N)/d_0} + K_3 \frac{d_0^2}{\log(N)^2} N^{-\frac{p-2}{2}} e^{-c_3/d_0^2}$. By combining both estimates (5.4) and (5.5) through the strong Markov property, we obtain the desired result. We refer to the proof of Lemma 5.1 for an application of the strong Markov property. The rest of the proof focuses on the proof of both claims.

Step 1: Proof of (5.4). We observe that for every $1 \leq \ell \leq \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)}$ and every $i, j \in [r]$,

$$\begin{aligned} -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle &\leq \sqrt{N} p \lambda_i \lambda_j m_{ij}^{p-1}(\ell-1) - \sqrt{N} p \lambda_1^2 r^2 \left(\frac{\gamma_-}{\sqrt{N}} \right)^{p+1}, \\ -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle &\geq \sqrt{N} p \lambda_i \lambda_j m_{ij}^{p-1}(\ell-1) - \sqrt{N} p \lambda_r^2 r^2 \left(\frac{\gamma_+}{\sqrt{N}} \right)^{p+1}, \end{aligned}$$

and

$$|h_i(\ell-1)| \leq \frac{\delta\alpha_2}{2} \sum_{k=1}^r |m_{ik}(\ell-1)| + \frac{\delta^2\alpha_2}{2N} \left(\sqrt{N}pr\lambda_i\lambda_1 \left(\frac{\gamma_+}{\sqrt{N}} \right)^{p-1} + \sqrt{N}pr^3\lambda_1^2 \left(\frac{\gamma_+}{\sqrt{N}} \right)^{p+1} \right) + \frac{\delta^2r\alpha_2K_N}{2N}.$$

In particular, we have the following estimates for $u_{ij}(t)$ and $\ell_{ij}(t)$:

$$\begin{aligned} u_{ij}(t) &\leq (1+c_0)\frac{\gamma_{ij}}{\sqrt{N}} + \frac{\delta}{N} \sum_{\ell=1}^t \left\{ \sqrt{N}p\lambda_i\lambda_j m_{ij}^{p-1}(\ell-1) - \sqrt{N}p\lambda_r^2 r^2 \left(\frac{\gamma_-}{\sqrt{N}} \right)^{p+1} \right. \\ &\quad \left. + \frac{\delta\alpha_2}{2} \sum_{k=1}^r |m_{ik}(\ell-1)| + \frac{\delta^2\alpha_2}{2N} \left(\sqrt{N}pr\lambda_i\lambda_1 \left(\frac{\gamma_+}{\sqrt{N}} \right)^{p-1} + \sqrt{N}pr^3\lambda_1^2 \left(\frac{\gamma_+}{\sqrt{N}} \right)^{p+1} \right) + \frac{\delta^2r\alpha_2K_N}{2N} \right\}, \end{aligned} \quad (5.6)$$

and

$$\begin{aligned} \ell_{ij}(t) &\geq (1-c_0)\frac{\gamma_{ij}}{\sqrt{N}} + \frac{\delta}{N} \sum_{\ell=1}^t \left\{ \sqrt{N}p\lambda_i\lambda_j m_{ij}^{p-1}(\ell-1) - \sqrt{N}p\lambda_r^2 r^2 \left(\frac{\gamma_+}{\sqrt{N}} \right)^{p+1} \right. \\ &\quad \left. - \frac{\delta\alpha_2}{2} \sum_{k=1}^r |m_{ik}(\ell-1)| - \frac{\delta^2\alpha_2}{2N} \left(\sqrt{N}pr\lambda_i\lambda_1 \left(\frac{\gamma_+}{\sqrt{N}} \right)^{p-1} + \sqrt{N}pr^3\lambda_1^2 \left(\frac{\gamma_+}{\sqrt{N}} \right)^{p+1} \right) - \frac{\delta^2r\alpha_2K_N}{2N} \right\}, \end{aligned} \quad (5.7)$$

for every $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)} \wedge T_0$, with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1-\eta(\delta, c_0, T_0, K_N)$.

We now consider (5.7) and claim that

$$\ell_{ij}(t) \geq (1-c_0)\frac{\gamma_{ij}}{\sqrt{N}} + \frac{\delta}{\sqrt{N}}(1-c_0)p\lambda_i\lambda_j \sum_{\ell=1}^t m_{ij}^{p-1}(\ell-1), \quad (5.8)$$

for every $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)} \wedge T_0$. To see the claim, we first note that, over the considered time interval,

$$\ell_{ij}(t) \geq (1-c_0)\frac{\gamma_{ij}}{\sqrt{N}} + \frac{\delta}{N} \sum_{\ell=1}^t \left\{ \sqrt{N}p\lambda_i\lambda_j m_{ij}^{p-1}(\ell-1) - 3\frac{\delta\alpha_2}{2} \sum_{k=1}^r |m_{ik}(\ell-1)| - \frac{\delta^2\alpha_2rK_N}{2N} \right\},$$

where we used the fact that $\delta = C_\delta d_0 N^{-\frac{p-3}{2}} / \log(N)$ by assumption. Claim (5.8) then follows by the following two sufficient conditions. First,

$$\frac{c_0}{2} \sqrt{N}p\lambda_i\lambda_j m_{ij}^{p-1}(\ell-1) \geq \frac{c_0}{2} p\lambda_r^2 \frac{\gamma_-^{p-1}}{N^{\frac{p-2}{2}}} \geq \frac{3\delta\alpha_2r}{2} \frac{\gamma_+}{\sqrt{N}} \geq \frac{3\delta\alpha_2}{2} \sum_{k=1}^r |m_{ik}(\ell-1)|,$$

for every $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)}$, provided $d_0 \leq c_0 \frac{p\lambda_r^2 \gamma_-^{p-1}}{3\alpha_2 C_\delta r \gamma_+}$ which certainly holds by choosing $d_0 \ll c_0^2$. Second,

$$\frac{c_0}{2} \sqrt{N}p\lambda_r^2 m_{ij}^{p-1}(\ell-1) \geq \frac{c_0}{2} p\lambda_r^2 \frac{\gamma_-^{p-1}}{N^{\frac{p-2}{2}}} \geq \frac{\delta^2\alpha_2rK_N}{2N},$$

for $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)}$, provided $d_0^2 K_N \leq c_0 \frac{p\lambda_r^2 \gamma_-^{p-1}}{C_\delta^2 \alpha_2 r} N^{\frac{p-2}{2}}$. We therefore can choose any value $K_N = o\left(N^{\frac{p-2}{2}}\right)$. Since the convergence speed in (4.5) is primarily determined by the ratio $\frac{c_0^2}{d_0}$, where d_0 can be chosen to go to zero arbitrarily slowly, we can simply pick $K_N = \frac{1}{d_0}$. This proves claim (5.8). A similar reasoning for (5.6) yields

$$u_{ij}(t) \leq (1+c_0)\frac{\gamma_{ij}}{\sqrt{N}} + \frac{\delta}{\sqrt{N}}(1+c_0)p\lambda_i\lambda_j \sum_{\ell=1}^t m_{ij}^{p-1}(\ell-1), \quad (5.9)$$

on the same time interval. Since both upper- and lower-bounding functions are increasing, we have that $\min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} < \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)}$. From item (b) of Lemma A.1 we then obtain that

$$\frac{(1-c_0)\gamma_{ij}}{\sqrt{N} \left(1 - \delta(1-c_0)^{p-1} p \lambda_i \lambda_j \gamma_{ij}^{p-2} N^{-\frac{p-1}{2}} t\right)^{\frac{1}{p-2}}} \leq m_{ij}(t) \leq \frac{(1+c_0)\gamma_{ij}}{\sqrt{N} \left(1 - \delta(1+c_0)^{p-1} p \lambda_i \lambda_j \gamma_{ij}^{p-2} N^{-\frac{p-1}{2}} t\right)^{\frac{1}{p-2}}}, \quad (5.10)$$

for $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)}$. In particular, for every $\varepsilon \leq \varepsilon_N^+$ it follows that

$$T_{u,\varepsilon}^{(ij)} \leq \mathcal{T}_{\varepsilon}^{(ij,+)} \leq T_{\ell,\varepsilon}^{(ij)},$$

where

$$T_{\ell,\varepsilon}^{(ij)} = \frac{1 - \left(\frac{(1-c_0)\gamma_{ij}}{\varepsilon\sqrt{N}}\right)^{p-2}}{\delta(1-c_0)^{p-1} p \lambda_i \lambda_j \gamma_{ij}^{p-2}} N^{\frac{p-1}{2}},$$

$$T_{u,\varepsilon}^{(ij)} = \frac{1 - \left(\frac{(1+c_0)\gamma_{ij}}{\varepsilon\sqrt{N}}\right)^{p-2}}{\delta(1+c_0)^{p-1} p \lambda_i \lambda_j \gamma_{ij}^{p-2}} N^{\frac{p-1}{2}}.$$

According to Definition 1.4, we have that $\lambda_{i_1^*} \lambda_{j_1^*} \gamma_{i_1^* j_1^*}^{p-2} = \max_{1 \leq i, j \leq r} \{\lambda_i \lambda_j \gamma_{ij}^{p-2}\}$. As a consequence, since by assumption $\mathbf{X}_0 \in \mathcal{C}_2(\gamma_1, \gamma)$, for every $i, j \in [r]$ such that $(i, j) \neq (i_1^*, j_1^*)$:

$$\lambda_{i_1^*} \lambda_{j_1^*} \gamma_{i_1^* j_1^*}^{p-2} > \left(1 + \frac{\gamma}{\gamma_1}\right) \lambda_i \lambda_j \gamma_{ij}^{p-2}. \quad (5.11)$$

Let $c_0 < \frac{(1+\frac{\gamma}{\gamma_1})^{\frac{1}{p-1}} - 1}{(1+\frac{\gamma}{\gamma_1})^{\frac{1}{p-1}} + 1}$ and $\gamma_+ > \gamma_1$ be a constant of order 1 satisfying

$$\gamma_+^{p-2} \geq (1+c_0)^{p-2} \gamma_1^{p-2} + \frac{(1+c_0)^{p-2} \gamma_1^{p-2} - (1-c_0)^{p-2} \gamma_2^{p-2}}{\frac{(1-c_0)^{p-1}}{(1+c_0)^{p-1}} \left(1 + \frac{\gamma}{\gamma_1}\right) - 1}.$$

Then, we see that $T_{\ell,\varepsilon_N^+}^{(i_1^* j_1^*)} \wedge T_0 \leq \min_{(i,j) \neq (i_1^*, j_1^*)} T_{u,\varepsilon_N^+}^{(ij)}$. That is, the correlation $m_{i_1^* j_1^*}$ exceeds the microscopic threshold ε_N^+ first, i.e., $\min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} = \mathcal{T}_{\varepsilon_N^+}^{(i_1^* j_1^*,+)}$. We can choose γ_+ such that $m_{ij}(\mathcal{T}_{\varepsilon_N^+}^{(i_1^* j_1^*,+)}) \leq 2m_{i_1^* j_1^*}(\mathcal{T}_{\varepsilon_N^+}^{(i_1^* j_1^*,+)}) = 2\varepsilon_N^+$ for every $(i, j) \neq (i_1^*, j_1^*)$. This shows that $\mathcal{T}_{\tilde{E}} = \mathcal{T}_{\varepsilon_N^+}^{(i_1^* j_1^*,+)}$ and we have that

$$\mathcal{T}_{\varepsilon_N^+}^{(i_1^* j_1^*,+)} \leq T_{\ell,\varepsilon_N^+}^{(i_1^* j_1^*)} \leq \frac{1 - \left(\frac{(1-c_0)\gamma_{i_1^* j_1^*}}{\gamma_+}\right)^{p-2}}{C_\delta d_0 (1-c_0)^{p-1} p \lambda_{i_1^*} \lambda_{j_1^*} \gamma_2^{p-2}} \log(N) N^{p-2},$$

with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - \tilde{\eta}$, where we substituted the prescribed value (5.2) for δ , as well as K_N and $T_0 = T_{\ell,\varepsilon_N^+}^{(i_1^* j_1^*)}$, into $\eta(\delta, c_0, T_0, K_N)$.

Step 2: Proof of (5.5). We now assume that $\mathbf{X}_0 \in \tilde{E}$ and study the evolution of the correlations over the time interval $[0, \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon}^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)} \wedge T_0]$, with $T_0 > 0$ to be determined. To this end, we let \mathcal{T}_{bad} denote the hitting time given by

$$\mathcal{T}_{\text{bad}} = \min \left\{ t \geq 0 : \max_{1 \leq i, j \leq r} m_{ij}(t) \geq 2m_{i_1^* j_1^*}(t) \right\}.$$

In particular, $\mathcal{T}_{\text{bad}} > 0$. By applying Proposition 4.3, we obtain that

$$m_{ij}(t) \geq (1-c_0)m_{ij}(0) + \frac{\delta}{N} \sum_{\ell=1}^t \left\{ -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle - |h_i(\ell-1)| \right\},$$

$$m_{ij}(t) \leq (1+c_0)m_{ij}(0) + \frac{\delta}{N} \sum_{\ell=1}^t \left\{ -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle + |h_i(\ell-1)| \right\},$$

for every $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_\varepsilon^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)} \wedge \mathcal{T}_{\text{bad}} \wedge T_0$, with probability at least $1 - \eta$, where η is given by (4.5). We then note that, on the considered time interval,

$$|h_i(t)| \leq r\delta\alpha_2 \max_{1 \leq k \leq r} |m_{ik}(t)| \leq r\delta\alpha_2 m_{i_1^* j_1^*}(t) \lesssim \frac{d_0}{\log(N)} \sqrt{N} p \lambda_{i_1^*} \lambda_{j_1^*} m_{i_1^* j_1^*}^{p-1}(t),$$

provided N is sufficiently large, where we used the chosen value (5.2) for δ and the definition of $\mathcal{T}_{\varepsilon_N^-}^{(i_1^* j_1^*, -)}$. It then follows that

$$m_{i_1^* j_1^*}(t) \geq (1 - c_0)\varepsilon_N^+ + (1 - c_0) \frac{\delta}{\sqrt{N}} p \lambda_{i_1^*} \lambda_{j_1^*} \sum_{\ell=1}^t m_{i_1^* j_1^*}^{p-1}(\ell - 1),$$

provided $c_0 \gtrsim d_0/\log(N)$, yielding

$$m_{i_1^* j_1^*}(t) \geq (1 - c_0)\varepsilon_N^+ \left(1 - \frac{\delta}{\sqrt{N}} (1 - c_0)^{p-1} p \lambda_{i_1^*} \lambda_{j_1^*} (\varepsilon_N^+)^{p-2} t \right)^{-\frac{1}{p-2}}, \quad (5.12)$$

according to Lemma A.1. Since $m_{i_1^* j_1^*}$ is lower bounded by an increasing function, it follows that $\min_{1 \leq i, j \leq r} \mathcal{T}_\varepsilon^{(ij,+)} \wedge \mathcal{T}_{\text{bad}} = \mathcal{T}_\varepsilon^{(i_1^* j_1^*, +)} \wedge \mathcal{T}_{\text{bad}}$. We therefore need to show that $\mathcal{T}_\varepsilon^{(i_1^* j_1^*, +)} \wedge T_0 \leq \mathcal{T}_{\text{bad}}$. To this end, note first that for every $(i, j) \neq (i_1^*, j_1^*)$, we have the following upper bound on the correlation $m_{ij}(t)$:

$$m_{ij}(t) \leq (1 + c_0) \frac{\varepsilon_N^+}{2} + p \frac{\delta}{\sqrt{N}} \sum_{\ell=1}^t \left\{ \lambda_i \lambda_j m_{ij}^{p-1}(\ell - 1) + \frac{\delta}{\sqrt{N}} r \alpha_2 m_{i_1^* j_1^*}(\ell - 1) \right\}. \quad (5.13)$$

We now claim that the term $\frac{\delta r \alpha_2}{\sqrt{N}} \sum_{\ell=1}^t m_{i_1^* j_1^*}(\ell - 1)$ on the right-hand side of (5.13) is sufficiently small compared to the drift term $\sum_{\ell=1}^t \lambda_i \lambda_j m_{ij}^{p-1}(\ell - 1)$. To bound this term, we divide the time interval $[0, \mathcal{T}_\varepsilon^{(i_1^* j_1^*, +)}]$ into a union of subintervals:

$$[0, \mathcal{T}_\varepsilon^{(i_1^* j_1^*, +)}] = \cup_{k=1}^{n_\varepsilon} [\mathcal{T}_{\varepsilon_k}^{(i_1^* j_1^*, +)}, \mathcal{T}_{\varepsilon_{k+1}}^{(i_1^* j_1^*, +)}],$$

where $\varepsilon_k = \frac{\gamma_+^k}{\sqrt{N}}$ for $1 \leq k \leq n_\varepsilon$, and n_ε denotes the number of intervals that are required to partition $[\varepsilon_N^+, \varepsilon]$. By definition, $\varepsilon_1 = \varepsilon_N^+$ and $\varepsilon_{n_\varepsilon+1} = \varepsilon$, so that $n_\varepsilon = \lceil \frac{\log(\varepsilon) + \frac{1}{2} \log(N)}{\log(\gamma_+)} - 1 \rceil$. Using this partition, we have that

$$\sum_{\ell=1}^t m_{i_1^* j_1^*}(\ell - 1) = \sum_{k=1}^{n_\varepsilon} \sum_{\ell=\mathcal{T}_{\varepsilon_k}^{(i_1^* j_1^*, +)}^{\mathcal{T}_{\varepsilon_{k+1}}^{(i_1^* j_1^*, +)}} m_{i_1^* j_1^*}(\ell - 1) \leq n_\varepsilon \max_{1 \leq k \leq n_\varepsilon} \left\{ \varepsilon_{k+1} \left(\mathcal{T}_{\varepsilon_{k+1}}^{(i_1^* j_1^*, +)} - \mathcal{T}_{\varepsilon_k}^{(i_1^* j_1^*, +)} \right) \right\}.$$

To bound the time difference in the previous display, we use (5.12). Indeed, since $m_{i_1^* j_1^*}$ is lower-bounded by a strictly increasing function, we can upper-bound the first time interval $\mathcal{T}_{\varepsilon_2}^{(i_1^* j_1^*, +)} - \mathcal{T}_{\varepsilon_1}^{(i_1^* j_1^*, +)} = \mathcal{T}_{\varepsilon_2}^{(i_1^* j_1^*, +)}$ by solving

$$(1 - c_0)\varepsilon_N^+ \left(1 - \frac{\delta}{\sqrt{N}} (1 - c_0)^{p-1} p \lambda_{i_1^*} \lambda_{j_1^*} (\varepsilon_N^+)^{p-2} t \right)^{-\frac{1}{p-2}} = \varepsilon_2,$$

yielding

$$\mathcal{T}_{\varepsilon_2}^{(i_1^* j_1^*, +)} \leq \frac{1 - \frac{(1-c_0)^{p-2}}{\gamma_+^{p-2}}}{\frac{\delta}{\sqrt{N}} (1 - c_0)^{p-1} p \lambda_{i_1^*} \lambda_{j_1^*} (\varepsilon_N^+)^{p-2}}.$$

Using the strong Markov property, we can apply Proposition 4.3 for each time interval $[\mathcal{T}_{\varepsilon_k}^{(i_1^* j_1^*, +)}, \mathcal{T}_{\varepsilon_{k+1}}^{(i_1^* j_1^*, +)}]$ and the comparison inequality (5.12) holds with ε_N^+ replaced by ε_k . Hence, for every k , we have that

$$\mathcal{T}_{\varepsilon_{k+1}}^{(i_1^* j_1^*, +)} - \mathcal{T}_{\varepsilon_k}^{(i_1^* j_1^*, +)} \leq \frac{1 - \frac{(1-c_0)^{p-2}}{\gamma_+^{p-2}}}{\frac{\delta}{\sqrt{N}} (1 - c_0)^{p-1} p \lambda_{i_1^*} \lambda_{j_1^*} \varepsilon_k^{p-2}}. \quad (5.14)$$

As a result, we obtain that

$$\frac{\delta^2 p r \alpha_2}{N} \sum_{\ell=1}^{\mathcal{T}_\varepsilon^{(i_1^* j_1^*, +)}} m_{i_1^* j_1^*}(\ell-1) \leq \frac{\delta^2 p r \alpha_2}{N} n_\varepsilon \max_{1 \leq k \leq n_\varepsilon} \left\{ \varepsilon_{k+1} \left(\mathcal{T}_{\varepsilon_{k+1}}^{(i_1^* j_1^*, +)} - \mathcal{T}_{\varepsilon_k}^{(i_1^* j_1^*, +)} \right) \right\} \lesssim \gamma_+ \frac{d_0}{\sqrt{N}},$$

Substituting this estimate into (5.13) leads to

$$m_{ij}(t) \leq (1 + c_0 + d_0) \frac{\varepsilon_N^+}{2} + p \frac{\delta}{\sqrt{N}} \sum_{\ell=1}^t \lambda_i \lambda_j m_{ij}^{p-1}(\ell-1),$$

and from Lemma A.1, it follows that

$$m_{ij}(t) \leq (1 + c_0 + d_0) \frac{\gamma_+}{2\sqrt{N}} \left(1 - p \delta \lambda_i \lambda_j (1 + c_0 + d_0)^{p-2} \gamma_+^{p-2} N^{-\frac{p-1}{2}} t \right)^{-\frac{1}{p-2}},$$

for every $t \leq \mathcal{T}_\varepsilon^{(i_1^* j_1^*, +)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N}^{(ij, -)} \wedge \mathcal{T}_{\text{bad}}$. Finally, in a similar way to (5.14), we have that

$$\mathcal{T}_\varepsilon^{(i_1^* j_1^*, +)} \leq \frac{1 - \left(\frac{(1-c_0)\gamma_+}{\varepsilon\sqrt{N}} \right)^{p-2}}{C_\delta d_0 (1-c_0)^{p-1} p \lambda_{i_1^*} \lambda_{j_1^*} \gamma_+^{p-2}} \log(N) N^{p-2}.$$

Plugging this bound into the above estimate for $m_{ij}(t)$ and performing some computation gives that

$$m_{ij}(\mathcal{T}_\varepsilon^{(i_1^* j_1^*, +)}) \lesssim \gamma_+ N^{-1/2},$$

where \lesssim hides a constant which depends only on $\varepsilon, \gamma_1, \gamma_2, \gamma, c_0, \{\lambda_i\}_{i=1}^r$. This shows that $\mathcal{T}_\varepsilon^{(i_1^* j_1^*, +)} \wedge T_0 < \mathcal{T}_{\text{bad}}$. Since Proposition 4.3 is applied n_ε times and $n_\varepsilon \approx \log(N)$, it follows that $\mathcal{T}_{E(\varepsilon)} \leq \mathcal{T}_\varepsilon^{(i_1^* j_1^*, +)}$ with probability at least $1 - \tilde{\eta} \log(N)$. \square

Given the event $E(\varepsilon)$, strong recovery of the first direction follows straightforwardly.

Lemma 5.3. *Consider the step-size $\delta = C_\delta d_0 \frac{N^{-\frac{p-3}{2}}}{\log(N)}$ from Lemma 5.2, with $d_0 = d_0(N)$ going arbitrarily slowly to 0 as $N \rightarrow \infty$. For every $\varepsilon > 0$, let $E(\varepsilon)$ denote the event defined by (5.3). Then, for sufficiently large N ,*

$$\inf_{\mathbf{X}_0 \in E(\varepsilon)} \mathbb{P}_{\mathbf{X}_0} \left(\mathcal{T}_{1-\varepsilon}^{(i_1^* j_1^*)} \lesssim \frac{1}{d_0} \log(N) N^{\frac{p-2}{2}} \right) \geq 1 - K \frac{\log(N)}{d_0} N^{\frac{p-2}{2}} e^{-cN}.$$

Proof. Fix $\varepsilon > 0$ and assume that $\mathbf{X}_0 \in E(\varepsilon)$. Proposition 4.3 applies to the correlations $m_{ij}(t)$ for all $(i, j) \neq (i_1^*, j_1^*)$, since $m_{ij}(0) \in \Theta(N^{-\frac{1}{2}})$. However, we focus here on the evolution of $m_{i_1^* j_1^*}(t)$, and since $m_{i_1^* j_1^*}(0) = \varepsilon$, a modification of Proposition 4.3 is needed. Fix a time horizon $T > 0$ to be determined later. It then follows from Lemma 4.4 that for every $t \leq T$,

$$\left| m_{i_1^* j_1^*}(t) - \left(\varepsilon - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_{j_1^*} \rangle - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_{j_1^*} \rangle \right) \right| \leq \frac{\delta}{N} \sum_{\ell=1}^t |a_{i_1^*}(\ell-1)|,$$

with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - T e^{-c_1 N}$, where $|a_{i_1^*}(\ell-1)|$ is bounded according to (4.8). The martingale term $\frac{\delta}{N} \sum_{\ell=1}^T \langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_{j_1^*} \rangle$ can be upper-bounded according to Lemma 4.5, i.e., for every $c_0 \in (0, \frac{1}{2})$,

$$\sup_{\mathbf{X}_0 \in \text{St}(N, r)} \mathbb{P}_{\mathbf{X}_0} \left(\max_{t \leq T} \frac{\delta}{N} \left| \sum_{\ell=1}^t \langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_{j_1^*} \rangle \right| \geq \frac{c_0}{2} \varepsilon \right) \leq 2 \exp \left(-c_2 \frac{c_0^2 N^{p-1} \log(N)^2}{d_0^2 T} \right),$$

where we used the prescribed value of δ . To estimate the term $\frac{\delta}{N} |\langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle|$ within $|a_{i_1^*}(\ell-1)|$, we introduce a truncation level K_N , defined such that $K_N \leq \frac{N^{\frac{p-1}{2}} \log(N)}{2C_\delta d_0}$. In particular,

using Lemma 4.6 and simple bounds on the correlations (by 1), for sufficiently large N , we have that

$$\begin{aligned} \frac{\delta}{N} \sum_{\ell=1}^T |a_{i_1^*}(\ell-1)| &\leq \frac{\delta^2 \alpha_2 r}{N} T + \frac{\delta^2 \alpha_2}{2N} \sum_{\ell=1}^T \sum_{k=1}^r |\langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} H^\ell(\mathbf{X}_\ell - 1))_k \rangle| \mathbf{1}_{\{|\langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} H^\ell(\mathbf{X}_\ell - 1))_k \rangle| > K_N\}} \\ &\leq \frac{2\delta^2 \alpha_2 r}{N} \\ &\leq C \frac{d_0^2}{\log(N)^2 N^{p-2}} T, \end{aligned}$$

with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - K \frac{d_0^3}{\log(N)^3 N^{\frac{3p-5}{2}}} T e^{-c \log(N) N^{p-1}/d_0^2}$, where we have substituted the chosen value of δ . Here, the constants C, K, c may depend only on $p, \sigma^2, \{\lambda_i\}_{i=1}^r$. We now observe that, for every time horizon $T \ll \frac{\log(N)^2 N^{p-2}}{d_0^2}$, the sum $\frac{\delta}{N} \sum_{\ell=1}^T |a_{i_1^*}(\ell-1)|$ can be set to any arbitrary fraction of the initial condition ε , i.e.,

$$\frac{\delta}{N} \sum_{\ell=1}^T |a_{i_1^*}(\ell-1)| \leq \frac{c_0}{2} \varepsilon \leq \frac{c_0}{2} m_{ij}(0).$$

Thus, we have established that

$$(1 - c_0)\varepsilon - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_{j_1^*} \rangle \leq m_{i_1^* j_1^*}(t) \leq (1 + c_0)\varepsilon - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_{j_1^*} \rangle, \quad (5.15)$$

with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - T e^{-c_1 N} - 2e^{-c_2 \frac{c_0^2 N^{p-1}}{d_0^2 T}} - K \frac{d_0^3}{\log(N)^3 N^{\frac{3p-5}{2}}} T e^{-c_3 \log(N) N^{p-1}/d_0^2}$. Recalling the expression for $\langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_{j_1^*} \rangle$ from Proposition 4.3, it follows that, for every initialization $\mathbf{X}_0 \in E(\varepsilon)$ and every $t \leq T \wedge \mathcal{T}_{1-\varepsilon}^{(i_1^* j_1^*, +)}$,

$$-\langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_{j_1^*} \rangle \geq p\sqrt{N} \lambda_{i_1^*} \lambda_{j_1^*} m_{i_1^* j_1^*}^{p-1}(t) (1 - m_{i_1^* j_1^*}^2(t)).$$

Hence, for sufficiently large N , we obtain the following comparison inequality:

$$m_{i_1^* j_1^*}(t) \geq (1 - c_0)\varepsilon + (1 - c_0)p\lambda_{i_1^*} \lambda_{j_1^*} \frac{\delta}{\sqrt{N}} \sum_{\ell=1}^t m_{i_1^* j_1^*}^{p-1}(\ell-1) (1 - m_{i_1^* j_1^*}^2(\ell-1)).$$

Since $1 - m_{i_1^* j_1^*}^2 \geq \varepsilon$ for all $t \leq T \wedge \mathcal{T}_{1-\varepsilon}^{(i_1^* j_1^*, +)}$, Lemma A.1 implies that

$$m_{i_1^* j_1^*}(t) \geq (1 - c_0)\varepsilon \left(1 - \frac{\delta}{\sqrt{N}} (1 - c_0)^{p-2} p \lambda_{i_1^*} \lambda_{j_1^*} \varepsilon^{p-1} t \right)^{-\frac{1}{p-2}}.$$

Thus, we obtain the desired result:

$$\mathcal{T}_{1-\varepsilon}^{(i_1^* j_1^*, +)} \leq \frac{1 - \left(\frac{(1-c_0)\varepsilon}{1-\varepsilon} \right)^{p-2}}{C_\delta d_0 (1 - c_0)^{p-2} p \lambda_{i_1^*} \lambda_{j_1^*} \varepsilon^{p-1}} N^{\frac{p-2}{2}} \ll \frac{\log(N)^2 N^{p-2}}{d_0^2},$$

with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - \frac{K_1}{d_0} N^{\frac{p-2}{2}} e^{-c_1 N} - K_2 e^{-c_2 N^{p/2}/d_0} - K_3 \frac{d_0^2}{\log(N)^3 N^{\frac{2p-3}{2}}} e^{-c_3 \log(N) N^{p-1}/d_0^2}$. To obtain the convergence rate stated in the lemma, we absorb the last two terms in the probability of occurrence given by Lemma 4.4. \square

By combining Lemmas 5.2 and 5.3 using the strong Markov property, we prove strong recovery of the first direction.

Proof of Lemma 5.1. Assume that $\mathbf{X}_0 \in \mathcal{C}_1(\gamma_1, \gamma_2) \cap \mathcal{C}_2(\gamma_1, \gamma)$. To simplify notation slightly, we write $T_\varepsilon^* = \mathcal{T}_\varepsilon^{(i_1^* j_1^*, +)}$ and $T_{1-\varepsilon}^* = \mathcal{T}_{1-\varepsilon}^{(i_1^* j_1^*, +)}$. We then note that

$$\begin{aligned} & \mathbb{P}_{\mathbf{X}_0} \left(T_{1-\varepsilon}^* \lesssim \frac{\log(N) \left(N^{p-2} + N^{\frac{p-2}{2}} \right)}{d_0} \right) \\ & \geq \mathbb{P}_{\mathbf{X}_0} \left(\left\{ T_{1-\varepsilon}^* \lesssim \frac{\log(N) \left(N^{p-2} + N^{\frac{p-2}{2}} \right)}{d_0} \right\} \cap \left\{ T_\varepsilon^* \lesssim \frac{\log(N)}{d_0} N^{p-2}, \mathbf{X}_{T_\varepsilon^*} \in E(\varepsilon) \right\} \right) \\ & = \mathbb{P}_{\mathbf{X}_0} \left(T_{1-\varepsilon}^* \lesssim \frac{\log(N) \left(N^{p-2} + N^{\frac{p-2}{2}} \right)}{d_0} \mid T_\varepsilon^* \lesssim \frac{\log(N)}{d_0} N^{p-2}, \mathbf{X}_{T_\varepsilon^*} \in E(\varepsilon) \right) \\ & \quad \times \mathbb{P}_{\mathbf{X}_0} \left(T_\varepsilon^* \lesssim \frac{\log(N)}{d_0} N^{p-2}, \mathbf{X}_{T_\varepsilon^*} \in E(\varepsilon) \right). \end{aligned}$$

The second factor is given according to Lemma 5.2. For the first factor we use the strong Markov property:

$$\begin{aligned} & \mathbb{P}_{\mathbf{X}_0} \left(T_{1-\varepsilon}^* \lesssim \frac{\log(N) \left(N^{p-2} + N^{\frac{p-2}{2}} \right)}{d_0} \mid T_\varepsilon^* \lesssim \frac{\log(N)}{d_0} N^{p-2}, \mathbf{X}_{T_\varepsilon^*} \in E(\varepsilon) \right) \\ & = \mathbb{P}_{\{\mathbf{X}_{T_\varepsilon^*} \in E(\varepsilon)\}} \left(T_{1-\varepsilon}^* \lesssim \frac{\log(N)}{d_0} N^{\frac{p-2}{2}} \right) \\ & \geq \inf_{\mathbf{X} \in E(\varepsilon)} \mathbb{P}_{\mathbf{X}} \left(T_{1-\varepsilon}^* \lesssim \frac{\log(N)}{d_0} N^{\frac{p-2}{2}} \right). \end{aligned}$$

The estimate of the last line is provided by Lemma 5.3. Combining both results proves the statement. \square

To prove the recovery of multiple directions, as stated in Theorem 3.5, we divide the proof into several steps. For every $\varepsilon > 0$, we consider the following sets:

$$\begin{aligned} E_1(\varepsilon) &= R_1(\varepsilon) \cap \left\{ \mathbf{X} : m_{ij}(\mathbf{X}) \in \Theta(N^{-\frac{1}{2}}) \quad \forall i \neq i_1^*, j \neq j_1^* \right\}, \\ E_2(\varepsilon) &= R_1(\varepsilon) \cap R_2(\varepsilon) \cap \left\{ \mathbf{X} : m_{ij}(\mathbf{X}) \in \Theta(N^{-\frac{1}{2}}) \text{ for } i \neq i_1^*, i_2^* \text{ and } j \neq j_1^*, j_2^* \right\}, \\ & \dots \\ E_{r-1}(\varepsilon) &= \cap_{1 \leq i \leq r-1} R_i(\varepsilon) \cap \left\{ \mathbf{X} : m_{i_r^* j_r^*}(\mathbf{X}) \in \Theta(N^{-\frac{1}{2}}) \right\}, \\ E_r(\varepsilon) &= \cap_{1 \leq i \leq r-1} R_i(\varepsilon) \cap \left\{ \mathbf{X} : m_{i_r^* j_r^*}(\mathbf{X}) \geq 1 - \varepsilon \right\}, \end{aligned}$$

where $R_k(\varepsilon)$ denotes the set of strong recovery of the spike $\mathbf{v}_{i_k^*}$, i.e.,

$$R_k(\varepsilon) = \left\{ \mathbf{X} : m_{i_k^* j_k^*}(\mathbf{X}) \geq 1 - \varepsilon \text{ and } m_{i_k^* j}(\mathbf{X}), m_{i j_k^*}(\mathbf{X}) \lesssim \log(N)^{-\frac{1}{2}} N^{-\frac{p-1}{4}} \quad \forall i \neq i_k^*, j \neq j_k^* \right\}.$$

Note that once the set $E_r(\varepsilon)$ is attained, we achieve recovery of a permutation of all spikes.

Lemma 5.4. Assume that $\delta = C_\delta d_0 \frac{N^{-\frac{p-3}{2}}}{\log(N)}$, where $C_\delta > 0$ is a constant and $d_0 = d_0(N)$ is a control parameter which goes to zero arbitrarily slowly as $N \rightarrow \infty$. Then, the following holds.

(a) For every $\gamma_1, \gamma_2, \gamma > 0$ with $\gamma_1 > \gamma_2, \gamma_1 > \gamma$, and for every $\varepsilon > 0$, there exists $c_0 < \frac{(1 + \frac{\gamma}{\gamma_1})^{\frac{1}{p-1}} - 1}{(1 + \frac{\gamma}{\gamma_1})^{\frac{1}{p-1}} + 1}$ such that for sufficiently large N ,

$$\inf_{\mathbf{X}_0 \in \mathcal{C}_1(\gamma_1, \gamma_2) \cap \mathcal{C}_2(\gamma_1, \gamma)} \mathbb{P}_{\mathbf{X}_0} (\mathcal{T}_{E_1(\varepsilon)} \lesssim T_1) \geq 1 - \eta,$$

where T_1 and η are given by

$$T_1 = \frac{1 - \left(\frac{(1-c_0)\gamma_2}{\varepsilon\sqrt{N}} \right)^{p-2}}{(1-c_0)^{p-1} d_0} \log(N) N^{p-2} + \frac{1 - \left(\frac{\varepsilon}{1-\varepsilon} \right)^{p-2}}{C_\delta d_0 \varepsilon^{p-1}} \log(N) N^{\frac{p-2}{2}}.$$

and

$$\eta = K_1 \frac{\log(N)^2}{d_0} N^{p-2} e^{-c_1 N} + K_2 \log(N) e^{-c_2 c_0^2 \log(N)/d_0} + K_3 \frac{d_0^2}{\log(N)} N^{-\frac{p-2}{2}} e^{-c_3/d_0^2}.$$

(b) For every $\varepsilon > 0$, there exists $c_0 \in (0, \frac{1}{2})$ such that for every $2 \leq k \leq r-1$ and sufficiently large N , there exists T_k of order $\log(N)N^{p-2}$ such that

$$\inf_{\mathbf{X} \in E_{k-1}(\varepsilon)} \mathbb{P}_{\mathbf{X}} (\mathcal{T}_{E_k(\varepsilon)} \lesssim T_k) \geq 1 - K\eta.$$

The proof of Theorem 3.5 follows from Lemma 5.4 by combining statements (a) and (b) via the strong Markov property, similarly to the proof of Lemma 5.1, and integrating with respect to the initial conditions 1 and 2. For further details, we refer the reader to the proof of Theorem 5.2 in [6]. It remains to prove Lemma 5.4.

Proof. First, we focus on statement (a). Fix $\varepsilon > 0$. From the strong Markov property and Lemma 5.2, it suffices to show that

$$\inf_{\mathbf{X}_0 \in E(\varepsilon)} \mathbb{P}_{\mathbf{X}_0} \left(\mathcal{T}_{E_1(\varepsilon)} \lesssim \frac{1 - (\varepsilon/(1-\varepsilon))^{p-2}}{C_\delta d_0 \varepsilon^{p-1}} \log(N) N^{\frac{p-2}{2}} \right) \geq 1 - \frac{K}{d_0} \log(N) N^{\frac{p-2}{2}} e^{-cN},$$

where $E(\varepsilon)$ denotes the set defined by (5.3). For every $(i, j) \neq (i_1^*, j_1^*)$ we can apply Proposition 4.3 since the correlations $m_{ij}(0) \in \Theta(N^{-\frac{1}{2}})$. In particular, for every $t \leq T_1$,

$$(1 - c_0)m_{ij}(0) - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle \leq m_{ij}(t) \leq (1 + c_0)m_{ij}(0) - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle, \quad (5.16)$$

with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - T_1 e^{-c_1 N} - 2e^{-c_2 c_0^2 \log(N)^2 N^{p-2}/(d_0^2 T_1)} - \frac{K d_0^3 T_1}{\log(N)^3 N^{\frac{3(p-2)}{2}}} e^{-c_3/d_0^2}$. According to (5.15) from Lemma 5.3, we also have the following two-sided comparison inequality for $m_{i_1^* j_1^*}$:

$$(1 - c_0)\varepsilon - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_{j_1^*} \rangle \leq m_{i_1^* j_1^*}(t) \leq (1 + c_0)\varepsilon - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_{j_1^*} \rangle, \quad (5.17)$$

for every $t \leq T_1$ with $T_1 \ll \frac{\log(N)^2}{d_0^2} N^{p-2}$, with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - T_1 e^{-c_4 N} - 2e^{-c_5 c_0^2 N^{p-1}/(d_0^2 T_1)} - \frac{C d_0^3 T_1}{\log(N)^3 N^{\frac{3p-5}{2}}} e^{-c_6 \log(N) N^{p-1}/d_0^2}$. As in the proof of Lemma 5.3, the lower bound in (5.17) is used to show that $m_{i_1^* j_1^*}$ reaches $1 - \varepsilon$ in $\mathcal{T}_{1-\varepsilon}^{(i_1^* j_1^*, +)}$ time, where

$$\mathcal{T}_{1-\varepsilon}^{(i_1^* j_1^*, +)} \leq \frac{1 - \left(\frac{(1-c_0)\varepsilon}{1-\varepsilon} \right)^{p-2}}{C_\delta d_0 (1-c_0)^{p-2} p \lambda_{i_1^*} \lambda_{j_1^*} \varepsilon^{p-1}} N^{\frac{p-2}{2}} \ll \frac{\log(N)^2}{d_0^2} N^{p-2}.$$

To reach the set $E_1(\varepsilon)$, we need to prove the following two steps using the two-sided inequality (5.16). First, we need to show that for every $i \neq i_1^*, j \neq j_1^*$, the correlations m_{ij}^* and $m_{i_1^* j}^*$ decrease below a level of order $\log(N)^{-\frac{1}{2}} N^{-\frac{p-1}{4}}$ within a time T_{lower} such that $T_{\text{lower}} \ll \frac{\log(N)^2}{d_0^2} N^{p-2}$. Second, we need to show that for every $(i, j) \neq (i_1^*, j_1^*)$, the correlations $m_{ij}(t)$ remain within $\Theta(N^{-\frac{1}{2}})$ for $t \leq T_{\text{lower}} \vee \mathcal{T}_{1-\varepsilon}^{(i_1^* j_1^*, +)}$. These two steps imply that we can pick $T_1 = T_{\text{lower}} \vee \mathcal{T}_{1-\varepsilon}^{(i_1^* j_1^*, +)}$.

To prove both steps, we therefore need to control the term $\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_j \rangle$:

$$\langle \mathbf{v}_i, \nabla_{\text{St}} \Phi(\mathbf{X}_t)_j \rangle = -\sqrt{N} p \lambda_i \lambda_j m_{ij}^{p-1}(t) + \sqrt{N} \frac{p}{2} \sum_{1 \leq k, \ell \leq r} \lambda_k m_{i\ell}(t) m_{kj}(t) m_{k\ell}(t) \left(\lambda_j m_{kj}^{p-2}(t) + \lambda_\ell m_{k\ell}^{p-2}(t) \right).$$

Let us now identify the dominant terms in this expression depending on the choice of indices. For every $i \neq i_1^*$ and $j \neq j_1^*$, the following bounds follow straightforwardly: for every $t \geq 0$,

$$\begin{aligned} -\langle \mathbf{v}_{i_1^*}, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_j \rangle &\leq \sqrt{N} p \lambda_{i_1^*} \left(\lambda_j m_{i_1^* j}^{p-2}(t) - \lambda_{j_1^*} m_{i_1^* j_1^*}^p(t) \right) m_{i_1^* j}(t), \\ -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_{j_1^*} \rangle &\leq \sqrt{N} p \lambda_i \left(\lambda_{j_1^*} m_{ij_1^*}^{p-2}(t) - \lambda_{j_1^*} m_{i_1^* j_1^*}^p(t) \right) m_{ij_1^*}(t), \\ -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_j \rangle &\geq \sqrt{N} p \left(\lambda_i \lambda_j m_{ij}^{p-1}(t) - r^2 \lambda_1^2 m_{ij_1^*}(t) m_{i_1^* j}(t) m_{i_1^* j_1^*}^p(t) \right). \end{aligned}$$

Once these comparison inequalities are established, it becomes straightforward to replicate the analysis conducted for Langevin dynamics in the extended version of our article [6]. By employing a similar decomposition argument as in the proof of Lemma 5.2, we can proceed almost verbatim. For the full details and completion of the proof for statement (a), we therefore refer the reader to the proof of [6, Lemma 5.8]. Statement (b) can be proved in a similar manner by following the same steps as in the proof of Lemmas 5.2 and 5.3, as well as statement (a). \square

5.2. PROOFS FOR $p = 2$ AND SEPARATED SNRS

In this subsection, we focus on the case where $p = 2$ and consider SNRs $\lambda_1 \geq \dots \geq \lambda_r$ which are separated by constants of order 1. Specifically, we assume that for every $1 \leq i \leq r-1$, the SNRs satisfy $\lambda_i = \lambda_{i+1}(1 + \kappa_i)$, where $\kappa_i > 0$, and define $\kappa = \min_{1 \leq i \leq r-1} \kappa_i$. Following a similar strategy to the one used in the proof of Theorem 3.5 from the previous subsection, we divide the proof of Theorem 3.6 into several steps, which will then be combined via the strong Markov property to complete the proof.

For every $\varepsilon > 0$, $c_0 \in \left(0, \frac{1}{2} \wedge \frac{\kappa}{2+\kappa}\right)$, and $i \in [r]$, we define the sets $W_i(\varepsilon, c_0)$ and $S_i(\varepsilon, c_0)$, which represent weak and strong recovery of the spike \mathbf{v}_i , respectively, as follows:

$$\begin{aligned} W_i(\varepsilon, c_0) &= \left\{ \mathbf{X} : m_{ii}(\mathbf{X}) \geq \varepsilon \text{ and } m_{ij}(\mathbf{X}), m_{ji}(\mathbf{X}) \lesssim N^{-\frac{1}{2}\delta_{ij}^{(i)}} \text{ for all } j > i \right\}, \\ S_i(\varepsilon, c_0) &= \left\{ \mathbf{X} : m_{ii}(\mathbf{X}) \geq 1 - \varepsilon \text{ and } m_{ij}(\mathbf{X}), m_{ji}(\mathbf{X}) \lesssim N^{-\frac{1}{2}\xi_{rr}^{(1)}} \text{ for all } j > i \right\}, \end{aligned}$$

where for every $1 \leq k \leq r-1$ and every $k \leq i, j \leq r$, $\delta_{ij}^{(k)}$ and $\xi_{ij}^{(k)}$ are given by

$$\delta_{ij}^{(k)} = 1 - \frac{1 + c_0}{1 - c_0} \frac{\lambda_i \lambda_j}{\lambda_k^2} \quad \text{and} \quad \xi_{ij}^{(k)} = 1 - \frac{1 - c_0}{1 + c_0} \frac{\lambda_i \lambda_j}{\lambda_k^2}.$$

Note that $\delta_{ij}^{(k)} < \xi_{ij}^{(k)}$, and both $\delta_{ij}^{(k)}$ and $\xi_{ij}^{(k)}$ are symmetric, meaning $\delta_{ij}^{(k)} = \delta_{ji}^{(k)}$ and $\xi_{ij}^{(k)} = \xi_{ji}^{(k)}$. Next, for every $\varepsilon > 0$ and $c_0 < \frac{1}{2} \wedge \frac{\kappa}{2+\kappa}$, we define the sets $E_1(\varepsilon, c_0), \dots, E_r(\varepsilon, c_0)$ as follows:

$$\begin{aligned} E_1(\varepsilon, c_0) &= W_1(\varepsilon, c_0) \cap \left\{ \mathbf{X} : N^{-\frac{1}{2}\xi_{ij}^{(1)}} \lesssim m_{ij}(\mathbf{X}) \lesssim N^{-\frac{1}{2}\delta_{ij}^{(1)}} \text{ for all } 2 \leq i, j \leq r \right\}, \\ E_2(\varepsilon, c_0) &= S_1(\varepsilon, c_0) \cap W_2(\varepsilon, c_0) \cap \left\{ \mathbf{X} : N^{-\frac{1}{2}\xi_{ij}^{(2)}} \lesssim m_{ij}(\mathbf{X}) \lesssim N^{-\frac{1}{2}\delta_{ij}^{(2)}} \text{ for all } 3 \leq i, j \leq r \right\}, \\ &\dots \\ E_{r-1}(\varepsilon, c_0) &= (\cap_{1 \leq i \leq r-2} S_i(\varepsilon, c_0)) \cap W_{r-1}(\varepsilon, c_0) \cap \left\{ \mathbf{X} : N^{-\frac{1}{2}\xi_{rr}^{(r-1)}} \lesssim m_{rr}(\mathbf{X}) \lesssim N^{-\frac{1}{2}\delta_{rr}^{(r-1)}} \right\}, \\ E_r(\varepsilon, c_0) &= (\cap_{1 \leq i \leq r-1} S_i(\varepsilon, c_0)) \cap W_r(\varepsilon, c_0). \end{aligned}$$

We observe that the recovery set $R(\varepsilon, c_0)$ in Theorem 3.6 corresponds to $R(\varepsilon, c_0) = \cap_{1 \leq i \leq r} S_i(\varepsilon, c_0)$.

The first result of this subsection shows that the event E_1 can be achieved with high probability, assuming the initial data \mathbf{X}_0 satisfies Condition 1 and the sample complexity is of order $\log(N)^2 N^{\frac{1}{2}} \left(1 - \frac{(1-c_0)\lambda_r^2}{(1+c_0)\lambda_1^2}\right)$.

Lemma 5.5. *For every $\gamma_1 > 1 > \gamma_2 > 0$ and every $\varepsilon > 0$, there exists $c_0 \in \left(0, \frac{1}{2} \wedge \frac{\kappa_1}{2+\kappa_1}\right)$ and $d_0 = d_0(N)$, which goes to zero arbitrarily slowly as $N \rightarrow \infty$, such that if*

$$\delta = \frac{C_\delta d_0 \gamma_2^2 N^{\frac{1}{2} \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}}}{\log\left(\frac{2\varepsilon\sqrt{N}}{\gamma_2}\right)}, \quad (5.18)$$

for some constant $C_\delta > 0$, and if $\log(N) \gtrsim \frac{1}{\kappa_1}$, then

$$\inf_{\mathbf{X}_0 \in \mathcal{C}_1(\gamma_1, \gamma_2)} \mathbb{P}_{\mathbf{X}_0} \left(\mathcal{T}_{E_1(\varepsilon, c_0)} \lesssim T_1 \right) \geq 1 - \eta,$$

where

$$T_1 = \frac{\log(2\varepsilon\sqrt{N}/\gamma_2)^2}{C_\delta d_0 (1-c_0) \gamma_2^2} N^{\frac{1}{2}} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right),$$

and

$$\eta = K_1 \frac{\log(N)^2 N^{\frac{1}{2}} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right)}{d_0} e^{-c_1 N} + K_2 e^{-c_2 c_0^2 N^{\frac{1}{2}} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right)} / d_0 + K_3 \frac{d_0^2 N^{-\left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right)}}{c_0 \log(N)} e^{-c_3/d_0^2}.$$

We now show that once E_1 is attained, we can reach E_2 . More generally, the following result shows that if the set E_k is achieved, then we can attain the set E_{k+1} with high probability.

Lemma 5.6. *For every $\varepsilon > 0$, there exists $c_0 \in (0, \frac{1}{2} \wedge \frac{\kappa}{2+\kappa})$ and $d_0 = d_0(N)$, which goes to zero arbitrarily slowly as $N \rightarrow \infty$, such that if*

$$\delta = \frac{C_\delta d_0 \gamma_2^2 N^{\frac{1}{2} \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}}}{\log\left(\frac{2\varepsilon\sqrt{N}}{\gamma_2}\right)},$$

for some constant $C_\delta > 0$, and if $\log(N) \gtrsim \frac{1}{\kappa}$, then

$$\inf_{\mathbf{X}_0 \in E_{k-1}(\varepsilon, c_0)} \mathbb{P}_{\mathbf{X}_0} \left(\mathcal{T}_{E_k(\varepsilon, c_0)} \lesssim T_k \right) \geq 1 - K\eta,$$

where

$$T_k = \log(N) \log(N)^{\frac{1}{2} \varepsilon_{kk}^{(k-1)}} N^{\frac{1}{2} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right)},$$

and η is given by Lemma 5.5.

Theorem 3.6 is proved by combining Lemmas 5.5 and 5.6 through the strong Markov property. Thus, it remains to prove the intermediate steps. We begin by proving Lemma 5.5, following the strategy used in the proof of Lemma 5.2.

Proof of Lemma 5.5. Since \mathbf{X}_0 satisfies Condition 1, for every $i, j \in [r]$, there exists $\gamma_{ij} \in (\gamma_2, \gamma_1)$ such that $m_{ij}(0) = \gamma_{ij} N^{-\frac{1}{2}}$. Let δ denote the step-size given by (5.18), and let T_1 be a fixed time horizon to be determined later. According to Proposition 4.3, for every $t \leq T_1$, we have that

$$\ell_{ij}(t) \leq m_{ij}(t) \leq u_{ij}(t),$$

with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - \eta$, where $\eta = \eta(\delta, c_0, T_1, K_N)$ is given by (4.5). The comparison functions ℓ_{ij} and u_{ij} are defined by

$$\begin{aligned} \ell_{ij}(t) &= (1 - c_0) \frac{\gamma_{ij}}{\sqrt{N}} + \frac{\delta}{N} \sum_{\ell=1}^t \{-\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle - |h_i(\ell - 1)|\} \\ u_{ij}(t) &= (1 + c_0) \frac{\gamma_{ij}}{\sqrt{N}} + \frac{\delta}{N} \sum_{\ell=1}^t \{-\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle + |h_i(\ell - 1)|\}, \end{aligned}$$

respectively, where for all $0 \leq t \leq T_1 - 1$,

$$-\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_j \rangle = 2\sqrt{N} \lambda_i \lambda_j m_{ij}(t) - \sqrt{N} \sum_{1 \leq k, \ell \leq r} \lambda_k (\lambda_j + \lambda_\ell) m_{i\ell}(t) m_{kj}(t) m_{k\ell}(t),$$

and

$$|h_i(t)| = \frac{\delta \alpha_2}{2} \sum_{k=1}^r |m_{ik}(t)| + \frac{\delta^2 \alpha_2}{2N} \sum_{k=1}^r |\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_k \rangle| + \frac{\delta^2 r \alpha_2 K_N}{2N}. \quad (5.19)$$

Next, we introduce the microscopic thresholds $\varepsilon_N^+ = \frac{\gamma_+}{\sqrt{N}}$ and $\varepsilon_N^- = \frac{\gamma_-}{\sqrt{N}}$ for some constants $\gamma_+ > \gamma_1$ and $\gamma_- < \gamma_2$. Recall that the hitting times $\mathcal{T}_{\varepsilon_N^+}^{(ij,+)}$ and $\mathcal{T}_{\varepsilon_N^-}^{(ij,-)}$ are defined in (5.1). Let $\mathcal{T}_{\tilde{E}}$ denote the time to reach the event

$$\tilde{E} = \left\{ \mathbf{X} : m_{11}(\mathbf{X}) \geq \varepsilon_N^+ > \frac{1}{2} \max_{(i,j) \neq (1,1)} m_{ij}(\mathbf{X}) \right\}.$$

To prove our result, we prove the following two estimates. First, we show that

$$\inf_{\mathbf{X}_0 \in \mathcal{C}_1(\gamma_1, \gamma_2)} \mathbb{P}_{\mathbf{X}_0} \left(\mathcal{T}_{\tilde{E}} \leq \frac{\log\left(\frac{\gamma_+}{(1-c_0)\gamma_{11}}\right) \log\left(\frac{2\varepsilon\sqrt{N}}{\gamma_2}\right)}{2C_\delta d_0 (1-c_0) \gamma_2^2 \lambda_1^2} N^{\frac{1}{2} \left(1 - \frac{(1-c_0)\lambda_r^2}{(1+c_0)\lambda_1^2}\right)} \right) \geq 1 - \tilde{\eta}_1, \quad (5.20)$$

where

$$\tilde{\eta}_1 = \frac{C_1 \log(N) N^{\frac{1}{2} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right)}}{d_0} e^{-c_1 N} + C_2 e^{-c_2 c_0^2 \log(N) N^{\frac{1}{2} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right)}} / d_0 + \frac{C_3 d_0^2 N^{-\left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right)}}{c_0 \log(N)^2} e^{-c_3 / d_0^2}.$$

Next, we show that

$$\inf_{\mathbf{X} \in \tilde{E}} \mathbb{P}_{\mathbf{X}} \left(\mathcal{T}_{E_1(\varepsilon, c_0)} \leq \frac{\log \left(\frac{\varepsilon \sqrt{N}}{(1-c_0)\gamma_+} \right) \log \left(\frac{2\varepsilon \sqrt{N}}{\gamma_2} \right)}{2C_\delta d_0 (1-c_0) \gamma_2^2 \lambda_1^2} N^{\frac{1}{2}} \left(1 - \frac{(1-c_0)\lambda_r^2}{(1+c_0)\lambda_1^2} \right) \right) \geq 1 - \tilde{\eta}_2, \quad (5.21)$$

where

$$\tilde{\eta}_2 = \frac{C_4 \log(N)^2 N^{\frac{1}{2}} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2} \right)}{d_0} e^{-c_4 N} + C_5 e^{-c_5 c_0^2 N^{\frac{1}{2}} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2} \right)} / d_0 + \frac{C_6 d_0^2 N^{-\left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2} \right)}}{c_0 \log(N)} e^{-c_6/d_0^2}.$$

Combining these estimates via the strong Markov property yields the desired result. The remainder of the proof will focus on proving these two claims.

Step 1: Proof of (5.20). We observe that for all $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)}$ and every $i, j \in [r]$, it holds that

$$2\sqrt{N} \lambda_i \lambda_j m_{ij}(t) - 2\sqrt{N} \lambda_r^2 r^2 \left(\frac{\gamma_+}{\sqrt{N}} \right)^3 \leq -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_j \rangle \leq 2\sqrt{N} \lambda_i \lambda_j m_{ij}(t) - 2\sqrt{N} \lambda_1^2 r^2 \left(\frac{\gamma_-}{\sqrt{N}} \right)^3.$$

This leads to the following estimates for the comparison functions $u_{ij}(t)$ and $\ell_{ij}(t)$:

$$\begin{aligned} u_{ij}(t) &\leq (1+c_0) \frac{\gamma_{ij}}{\sqrt{N}} + \frac{\delta}{N} \sum_{\ell=1}^t \left\{ 2\sqrt{N} \lambda_i \lambda_j m_{ij}(\ell-1) - 2\sqrt{N} \lambda_r^2 r^2 \left(\frac{\gamma_-}{\sqrt{N}} \right)^3 \right. \\ &\quad \left. + \frac{\delta \alpha_2 r}{2} \max_{1 \leq k \leq r} |m_{ik}(\ell-1)| + \frac{\delta^2 \alpha_2 r}{2N} \left(2\gamma_+ \lambda_i \lambda_1 + 2\sqrt{N} \lambda_1^2 r^2 \left(\frac{\gamma_+}{\sqrt{N}} \right)^3 \right) + \frac{\delta^2 r \alpha_2 K_N}{2N} \right\}, \end{aligned} \quad (5.22)$$

and

$$\begin{aligned} \ell_{ij}(t) &\geq (1-c_0) \frac{\gamma_{ij}}{\sqrt{N}} + \frac{\delta}{N} \sum_{\ell=1}^t \left\{ 2\sqrt{N} \lambda_i \lambda_j m_{ij}(\ell-1) - 2\sqrt{N} \lambda_r^2 r^2 \left(\frac{\gamma_+}{\sqrt{N}} \right)^3 \right. \\ &\quad \left. - \frac{\delta \alpha_2 r}{2} \max_{1 \leq k \leq r} |m_{ik}(\ell-1)| - \frac{\delta^2 \alpha_2 r}{2N} \left(2\gamma_+ \lambda_i \lambda_1 + 2\sqrt{N} \lambda_1^2 r^2 \left(\frac{\gamma_+}{\sqrt{N}} \right)^3 \right) - \frac{\delta^2 r \alpha_2 K_N}{2N} \right\}. \end{aligned} \quad (5.23)$$

Both bounds hold for every $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)} \wedge T_1$, with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - \eta(\delta, c_0, T_1, K_N)$. We now focus on (5.23) and claim that

$$\ell_{ij}(t) \geq (1-c_0) \frac{\gamma_{ij}}{\sqrt{N}} + 2(1-c_0) \lambda_i \lambda_j \frac{\delta}{\sqrt{N}} \sum_{\ell=1}^t m_{ij}(\ell-1), \quad (5.24)$$

for every $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)} \wedge T_1$. To see the claim, we first note that, on the considered time interval,

$$\ell_{ij}(t) \geq (1-c_0) \frac{\gamma_{ij}}{\sqrt{N}} + \frac{\delta}{N} \sum_{\ell=1}^t \left\{ 2\sqrt{N} \lambda_i \lambda_j m_{ij}(\ell-1) - \frac{3\delta \alpha_2 r}{2} \max_{1 \leq k \leq r} |m_{ik}(\ell-1)| - \frac{\delta^2 r \alpha_2 K_N}{2N} \right\},$$

where we used the prescribed value (5.18) for δ . Claim (5.24) then follows by two sufficient conditions. First, note that over the considered time interval,

$$c_0 \sqrt{N} \lambda_i \lambda_j m_{ij}(\ell-1) \geq c_0 \lambda_i \lambda_j \gamma_- \geq \frac{3\delta \alpha_2 r}{2} \frac{\gamma_+}{\sqrt{N}} \geq \frac{3\delta \alpha_2 r}{2} \max_{1 \leq k \leq r} |m_{ik}(\ell-1)|,$$

where the second inequality holds, provided

$$d_0 \leq \frac{2c_0 \lambda_i \lambda_j \gamma_-}{3C_\delta \alpha_2 r \gamma_+ \gamma_2^2} N^{\frac{1}{2}} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2} \right) \log \left(\frac{2\varepsilon \sqrt{N}}{\gamma_2} \right), \quad (5.25)$$

which certainly holds by choosing $d_0 \ll c_0^2$. Second, we have that

$$c_0 \sqrt{N} \lambda_i \lambda_j m_{ij}(\ell-1) \geq c_0 \lambda_i \lambda_j \gamma_- \geq \frac{\delta^2 r \alpha_2 K_N}{2N},$$

for every $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)}$, where the second inequality holds provided

$$d_0^2 K_N \leq \frac{2c_0 \lambda_i \lambda_j \gamma_-}{C_\delta^2 \alpha_{2r} \gamma_2^4} N^{1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}} \log \left(\frac{2\varepsilon \sqrt{N}}{\gamma_2} \right)^2. \quad (5.26)$$

According to the sufficient conditions (5.25) and (5.26), we can choose the truncation level $K_N = \frac{1}{d_0}$. Claim (5.24) then easily follows. A similar reasoning for (5.22) yields

$$u_{ij}(t) \leq (1+c_0) \frac{\gamma_{ij}}{\sqrt{N}} + 2(1+c_0) \lambda_i \lambda_j \frac{\delta}{\sqrt{N}} \sum_{\ell=1}^t m_{ij}(\ell-1), \quad (5.27)$$

for every $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)} \wedge T_1$. According to (5.24) and (5.27) both lower- and upper-bounding functions are increasing, thus each correlation m_{ij} is increasing on the considered time interval, and $\min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} \wedge T_1 \leq \min_{1 \leq \ell \leq r} \mathcal{T}_{\varepsilon_N^-}^{(ij,-)}$. By applying item (a) of Lemma A.1, we then obtain the following two-sided comparison inequality:

$$\frac{(1-c_0)\gamma_{ij}}{\sqrt{N}} \left(1 + \frac{2\delta}{\sqrt{N}} (1-c_0) \lambda_i \lambda_j \right)^t \leq m_{ij}(t) \leq \frac{(1+c_0)\gamma_{ij}}{\sqrt{N}} \left(1 + \frac{2\delta}{\sqrt{N}} (1+c_0) \lambda_i \lambda_j \right)^t, \quad (5.28)$$

for all $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)}$. In particular, from (5.28), we derive the following upper and lower bounds for $\mathcal{T}_{\varepsilon_N^+}^{(ij,+)}$:

$$T_{u, \varepsilon_N^+}^{(ij)} \leq \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} \leq T_{\ell, \varepsilon_N^+}^{(ij)},$$

where

$$T_{u, \varepsilon_N^+}^{(ij)} = \frac{\log \left(\frac{\gamma_+}{(1+c_0)\gamma_{ij}} \right)}{\log \left(1 + \frac{2\delta}{\sqrt{N}} (1+c_0) \lambda_i \lambda_j \right)}$$

$$T_{\ell, \varepsilon_N^+}^{(ij)} = \frac{\log \left(\frac{\gamma_+}{(1-c_0)\gamma_{ij}} \right)}{\log \left(1 + \frac{2\delta}{\sqrt{N}} (1-c_0) \lambda_i \lambda_j \right)}.$$

We now claim that $T_{\ell, \varepsilon_N^+}^{(11)} \leq T_{u, \varepsilon_N^+}^{(ij)}$ for all $(i, j) \neq (1, 1)$, provided γ_+ is a sufficiently large constant of order 1 in N . Since by assumption $c_0 < \frac{1}{2} \wedge \frac{\kappa}{2+\kappa}$, it follows $(1-c_0)(1+\kappa_1) > (1+c_0)$. Thus, there exists a constant $c(\kappa_1) > 1$ such that

$$\log \left(1 + \frac{2\delta}{\sqrt{N}} (1-c_0)(1+\kappa_1) \lambda_i \lambda_j \right) = c(\kappa_1) \log \left(1 + \frac{2\delta}{\sqrt{N}} (1+c_0) \lambda_i \lambda_j \right).$$

In particular, choosing γ_+ such that

$$\log(\gamma_+) \geq \frac{c(\kappa_1) \log((1+c_0)\gamma_1) - \log((1-c_0)\gamma_2)}{c(\kappa_1) - 1},$$

implies that

$$T_{\ell, \varepsilon_N^+}^{(11)} \leq \frac{\log \left(\frac{\gamma_+}{(1-c_0)\gamma_2} \right)}{\log \left(1 + \frac{2\delta}{\sqrt{N}} (1-c_0) \lambda_i \lambda_j (1+\kappa_1) \right)} \leq \frac{\log \left(\frac{\gamma_+}{(1+c_0)\gamma_1} \right)}{\log \left(1 + \frac{2\delta}{\sqrt{N}} (1+c_0) \lambda_i \lambda_j \right)} \leq T_{u, \varepsilon_N^+}^{(ij)},$$

for all $(i, j) \neq (1, 1)$. The inequality $T_{\ell, \varepsilon_N^+}^{(11)} \leq T_{u, \varepsilon_N^+}^{(ij)}$ ensures that the correlation m_{11} exceeds the microscopic threshold ε_N^+ first, i.e., $\min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N^+}^{(ij,+)} = \mathcal{T}_{\varepsilon_N^+}^{(11,+)}$. Finally, if we choose γ_+ such that $\max_{(i,j) \neq (1,1)} m_{ij}(\mathcal{T}_{\varepsilon_N^+}^{(11,+)}) \leq 2m_{11}(\mathcal{T}_{\varepsilon_N^+}^{(11,+)}) = 2\varepsilon_N^+$, we conclude that $\mathcal{T}_{\bar{E}} = \mathcal{T}_{\varepsilon_N^+}^{(11,+)}$, and we find the bound:

$$\mathcal{T}_{\varepsilon_N^+}^{(11,+)} \leq T_{\ell, \varepsilon_N^+}^{(11)} \leq \frac{\log \left(\frac{\gamma_+}{(1-c_0)\gamma_{11}} \right) \log \left(\frac{2\varepsilon \sqrt{N}}{\gamma_2} \right)}{2C_\delta d_0 (1-c_0) \gamma_2^2 \lambda_1^2} N^{\frac{1}{2} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2} \right)},$$

with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - \tilde{\eta}_1$. Here, we substituted the prescribed value (5.18) for δ , and we used the fact that $\log(1+x)$ is approximately equal to x when x is near zero.

Step 2: Proof of (5.21). We now assume that $\mathbf{X}_0 \in \tilde{E}$. As in the proof of Lemma 5.2, we introduce the hitting time \mathcal{T}_{bad} given by

$$\mathcal{T}_{\text{bad}} = \min \left\{ t \geq 0 : \max_{1 \leq i, j \leq r} m_{ij}(t) \geq 2m_{11}(t) \right\},$$

and study the evolution of the correlations for $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_\varepsilon^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N}^{(ij,-)} \wedge \mathcal{T}_{\text{bad}} \wedge T_1$. We observe that on this time interval, m_{11} satisfies the two-sided comparison inequality given by

$$\frac{(1-c_0)\gamma_+}{\sqrt{N}} \left(1 + \frac{2\delta}{\sqrt{N}}(1-c_0)\lambda_1^2 \right)^t \leq m_{11}(t) \leq \frac{(1+c_0)\gamma_+}{\sqrt{N}} \left(1 + \frac{2\delta}{\sqrt{N}}(1+c_0)\lambda_1^2 \right)^t, \quad (5.29)$$

which is in the same flavor as (5.28). This follows by applying Proposition 4.3 and by using similar bounds as in the first step of this proof. For instance, on this time interval, we have that

$$c_0\sqrt{N}\lambda_1^2 m_{11}(t) \geq \frac{3\delta\alpha_2 r}{2} m_{11}(t) \geq \frac{3\delta\alpha_2 r}{2} \max_{1 \leq k \leq r} |m_{1k}(t)|,$$

where the first inequality holds trivially by the fact that $\delta \in \mathcal{O}(d_0 N^{\frac{1}{2} \frac{1-c_0}{1+c_0}} \frac{\lambda_r^2}{\lambda_1^2} / \log(N))$. This, together with Lemma A.1, gives the comparison inequality (5.29). Furthermore, from (5.29), we obtain the following lower and upper bounds for $\mathcal{T}_\varepsilon^{(11,+)}$:

$$\frac{\log \left(\frac{\varepsilon\sqrt{N}}{(1+c_0)\gamma_+} \right)}{\log \left(1 + \frac{2\delta}{\sqrt{N}}(1+c_0)\lambda_1^2 \right)} \leq \mathcal{T}_\varepsilon^{(11,+)} \leq \frac{\log \left(\frac{\varepsilon\sqrt{N}}{(1-c_0)\gamma_+} \right)}{\log \left(1 + \frac{2\delta}{\sqrt{N}}(1-c_0)\lambda_1^2 \right)}. \quad (5.30)$$

To derive the comparison inequalities for $(i, j) \neq (1, 1)$, we begin with Proposition 4.3, which gives that

$$\begin{aligned} m_{ij}(t) &\geq (1-c_0)\frac{\varepsilon_N^+}{2} + \frac{\delta}{N} \sum_{\ell=1}^t \{ -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle - |h_i(\ell-1)| \}, \\ m_{ij}(t) &\leq (1+c_0)\frac{\varepsilon_N^+}{2} + \frac{\delta}{N} \sum_{\ell=1}^t \{ -\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle + |h_i(\ell-1)| \}, \end{aligned}$$

for every $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_\varepsilon^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N}^{(ij,-)} \wedge \mathcal{T}_{\text{bad}} \wedge T_1$, with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - \eta(\delta, c_0, T_1, K_N)$. According to (5.19), we derive the following bound for the third term in $|h_i(t)|$:

$$c_0\sqrt{N}\lambda_i\lambda_j m_{ij}(t) \geq c_0\lambda_i\lambda_j\gamma_- \geq \frac{\delta^2 r \alpha_2 K_N}{2N},$$

provided $d_0^2 K_N \leq \frac{2c_0\lambda_i\lambda_j\gamma_-}{C_\delta^2 \gamma_2^4 r \alpha_2} N^{1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}} \log \left(\frac{2\varepsilon\sqrt{N}}{\gamma_2} \right)^2$. This holds when choosing $K_N = \frac{1}{d_0}$. Moreover, the second summand in $|h_i(t)|$ (see (5.19)) can be absorbed into the first term in $|h_i(t)|$ by noting that

$$\frac{\delta}{N} |\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_k \rangle| \leq (1+c_0) \frac{\delta}{\sqrt{N}} 2\lambda_i\lambda_k |m_{ik}(t)| \leq |m_{ik}(t)|,$$

provided $d_0 \leq \frac{\log \left(\frac{2\varepsilon\sqrt{N}}{\gamma_2} \right)}{2C_\delta \gamma_2^2 (1+c_0)\lambda_i\lambda_k} N^{\frac{1}{2} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2} \right)}$, which holds trivially by choosing $d_0 \ll c_0^2$. Thus, we obtain the following bounds for m_{ij} :

$$\begin{aligned} m_{ij}(t) &\geq (1-c_0)\frac{\varepsilon_N^+}{2} + \frac{\delta}{N} \sum_{\ell=1}^t \left\{ 2\left(1 + \frac{c_0}{2}\right)\sqrt{N}\lambda_i\lambda_j m_{ij}(\ell-1) - \delta\alpha_2 r \max_{1 \leq k \leq r} |m_{ik}(\ell-1)| \right\}, \\ m_{ij}(t) &\leq (1+c_0)\frac{\varepsilon_N^+}{2} + \frac{\delta}{N} \sum_{\ell=1}^t \left\{ 2\left(1 + \frac{c_0}{2}\right)\sqrt{N}\lambda_i\lambda_j m_{ij}(\ell-1) + \delta\alpha_2 r \max_{1 \leq k \leq r} |m_{ik}(\ell-1)| \right\}, \end{aligned}$$

for every $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_\varepsilon^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N}^{(ij,-)} \wedge \mathcal{T}_{\text{bad}} \wedge T_1$. We now need to estimate the term $\delta\alpha_2 r \max_{1 \leq k \leq r} |m_{ik}(t)|$. To do this, for every $i, j \in [r]$, we introduce the hitting time $\mathcal{T}_\alpha^{(ij)}$ for some $\alpha \in (0, \frac{1}{2})$, which will be determined later:

$$\mathcal{T}_\alpha^{(ij)} = \min \left\{ t \geq 0 : C_{ij} N^\alpha m_{ij}(t) \leq \max_{1 \leq k \leq r} |m_{ik}(t)| \right\},$$

where C_{ij} is a positive constant. We easily see that $\mathcal{T}_\alpha^{(ij)} > 0$ for every $\alpha > 0$. From the definition of $\mathcal{T}_\alpha^{(ij)}$, for all $t \leq \mathcal{T}_\alpha^{(ij)}$, we have that

$$\delta\alpha_2 r \max_{1 \leq k \leq r} |m_{ik}(t)| \leq \delta\alpha_2 r C_{ij} N^\alpha m_{ij}(t) = \frac{C_\delta C_{ij} \alpha_2 r d_0 \gamma_2^2}{\log\left(\frac{2\varepsilon\sqrt{N}}{\gamma_2}\right)} N^{\frac{1}{2} \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2} + \alpha} m_{ij}(t).$$

To ensure the estimate

$$c_0 \sqrt{N} \lambda_i \lambda_j m_{ij}(t) \geq \delta\alpha_2 r \max_{1 \leq k \leq r} |m_{ik}(t)|,$$

the following inequality must hold:

$$c_0 \lambda_i \lambda_j N^{\frac{1}{2}} \geq \frac{C_\delta C_{ij} \alpha_2 r d_0 \gamma_2^2}{\log\left(\frac{2\varepsilon\sqrt{N}}{\gamma_2}\right)} N^{\frac{1}{2} \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2} + \alpha},$$

which is satisfied if

$$\alpha \leq \frac{1}{2} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right). \quad (5.31)$$

Thus, we obtain the two-sided inequality:

$$(1-c_0) \frac{\varepsilon_N^+}{2} + 2(1-c_0) \lambda_i \lambda_j \frac{\delta}{\sqrt{N}} \sum_{\ell=1}^t m_{ij}(\ell-1) \leq m_{ij}(t) \leq (1+c_0) \frac{\varepsilon_N^+}{2} + 2(1+c_0) \lambda_i \lambda_j \frac{\delta}{\sqrt{N}} \sum_{\ell=1}^t m_{ij}(\ell-1),$$

which, by applying Lemma A.1, leads to

$$\frac{(1-c_0)\gamma_+}{2\sqrt{N}} \left(1 + 2(1-c_0) \lambda_i \lambda_j \frac{\delta}{\sqrt{N}}\right)^t \leq m_{ij}(t) \leq \frac{(1+c_0)\gamma_+}{2\sqrt{N}} \left(1 + 2(1+c_0) \lambda_i \lambda_j \frac{\delta}{\sqrt{N}}\right)^t, \quad (5.32)$$

for every $t \leq \min_{1 \leq i, j \leq r} \mathcal{T}_\varepsilon^{(ij,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_{\varepsilon_N}^{(ij,-)} \wedge \mathcal{T}_{\text{bad}} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_\alpha^{(ij)} \wedge T_1$. From (5.32), we derive the following lower and upper bounds on $\mathcal{T}_\varepsilon^{(ij,+)}$ for all $(i, j) \neq (1, 1)$:

$$\frac{\log\left(\frac{2\varepsilon\sqrt{N}}{(1+c_0)\gamma_+}\right)}{\log\left(1 + \frac{2\delta}{\sqrt{N}}(1+c_0)\lambda_i\lambda_j\right)} \leq \mathcal{T}_\varepsilon^{(ij,+)} \leq \frac{\log\left(\frac{2\varepsilon\sqrt{N}}{(1-c_0)\gamma_+}\right)}{\log\left(1 + \frac{2\delta}{\sqrt{N}}(1-c_0)\lambda_i\lambda_j\right)}. \quad (5.33)$$

Finally, combining (5.30) and (5.33), we see that $\mathcal{T}_\varepsilon^{(11,+)} \wedge \mathcal{T}_{\text{bad}} \wedge \min_{i,j} \mathcal{T}_\alpha^{(ij)} \wedge T_1 \leq \mathcal{T}_\varepsilon^{(ij,+)}$ for all $(i, j) \neq (1, 1)$, provided

$$\log(\varepsilon\sqrt{N}) \geq \frac{(1+c_0)\log((1-c_0)\gamma_+) - (1-c_0)(1+\kappa_1)\log((1+c_0)\gamma_+)}{(1-c_0)(1+\kappa_1) - (1+c_0)}.$$

Note that the discriminator of the right-hand side of the above expression is positive under the assumption on c_0 . Simplifying this gives $\log(N) \gtrsim 1/\kappa_1$. It remains to show that $\mathcal{T}_\varepsilon^{(11,+)} \wedge T_1 \leq \mathcal{T}_{\text{bad}} \wedge \min_{i,j} \mathcal{T}_\alpha^{(ij)}$.

Now, consider the upper bound in (5.32). Using $\log(1+x) \leq x$, we have that

$$m_{ij}(t) \leq \frac{(1+c_0)\gamma_+}{2\sqrt{N}} \exp\left(2(1+c_0)\lambda_i\lambda_j \frac{\delta}{\sqrt{N}} t\right).$$

In particular, we find that

$$m_{ij}(\mathcal{T}_\varepsilon^{(11,+)}) \leq \frac{(1+c_0)\gamma_+}{2} \left(\frac{\varepsilon}{(1-c_0)\gamma_+}\right)^{\frac{1+c_0}{1-c_0} \frac{\lambda_i\lambda_j}{\lambda_1^2}} N^{-\frac{1}{2}} \left(1 - \frac{1+c_0}{1-c_0} \frac{\lambda_i\lambda_j}{\lambda_1^2}\right) \lesssim N^{-\frac{1}{2}\delta_{ij}^{(1)}}, \quad (5.34)$$

where we used the fact that

$$\mathcal{T}_\varepsilon^{(11,+)} \leq \frac{\log\left(\frac{\varepsilon\sqrt{N}}{(1-c_0)\gamma_+}\right)}{\log\left(1 + \frac{2\delta}{\sqrt{N}}(1-c_0)\lambda_1^2\right)} \leq \frac{\log\left(\frac{\varepsilon\sqrt{N}}{(1-c_0)\gamma_+}\right) \log\left(\frac{2\varepsilon\sqrt{N}}{\gamma_2}\right)}{2C_\delta d_0 (1-c_0)\gamma_2^2 \lambda_1^2} N^{\frac{1}{2}} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_2^2}{\lambda_1^2}\right),$$

which holds because $\log(1+x) \approx x$ for small x , which is true for large N and by assumption on δ . For the lower bound, since $\log(1+x) \geq x - \frac{x^2}{2}$, it follows from (5.32) that

$$m_{ij}(t) \geq \frac{(1-c_0)\gamma_+}{2\sqrt{N}} \exp\left(2(1-c_0)\lambda_i\lambda_j \frac{\delta}{\sqrt{N}} t\right) \exp\left(-\frac{1}{2}(1-c_0)^2 \lambda_i^2 \lambda_j^2 \frac{\delta^2}{N} t\right).$$

Thus, for any time horizon $T_1 \leq \frac{2 \log(2\varepsilon\sqrt{N}/\gamma_2)^2}{C_\delta^2 d_0 \gamma_2^4 (1-c_0)^2 \lambda_r^2} N^{1-\frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}}$, we get that

$$m_{ij}(t) \geq \frac{(1-c_0)\gamma_+}{2\sqrt{N}} \exp(-d_0) \exp\left(2(1-c_0)\lambda_i\lambda_j \frac{\delta}{\sqrt{N}} t\right).$$

Since $d_0 < 1$, we find that

$$m_{ij}(t) \geq \frac{(1-c_0)\gamma_+}{2e\sqrt{N}} \exp\left(2(1-c_0)\lambda_i\lambda_j \frac{\delta}{\sqrt{N}} t\right).$$

Using the lower bound of $\mathcal{T}_\varepsilon^{(11,+)}$ from (5.30), we obtain that

$$m_{ij}(\mathcal{T}_\varepsilon^{(11,+)}) \geq \frac{(1-c_0)\gamma_+}{2e} \left(\frac{\varepsilon}{(1+c_0)\gamma_+}\right)^{\frac{1-c_0}{1+c_0} \frac{\lambda_i\lambda_j}{\lambda_1^2}} N^{-\frac{1}{2} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_i\lambda_j}{\lambda_1^2}\right)} \lesssim N^{-\frac{1}{2} \xi_{ij}^{(1)}}, \quad (5.35)$$

From (5.34) and (5.35), we deduce that $\mathcal{T}_\varepsilon^{(11,+)} \wedge \min_{1 \leq i, j \leq r} \mathcal{T}_\alpha^{(ij)} \wedge T_1 \leq \mathcal{T}_{\text{bad}}$.

The final step is to show that $\mathcal{T}_\varepsilon^{(11,+)} \wedge T_1 \leq \min_{1 \leq i, j \leq r} \mathcal{T}_\alpha^{(ij)}$. To this end, it suffices to show that there exists a constant C_{ij} such that for some $\alpha \leq \frac{1}{2} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right)$ and for every $k \neq j$, it holds that $|m_{ik}(\mathcal{T}_\varepsilon^{(11,+)})| \leq C_{ij} N^\alpha m_{ij}(\mathcal{T}_\varepsilon^{(11,+)})$. From (5.34) and (5.35) we see that for $i \neq 1$,

$$|m_{i1}(\mathcal{T}_\varepsilon^{(11,+)})| \leq C_{i1} N^{-\frac{1}{2} \delta_{i1}^{(1)}} \leq C_{i1} N^\alpha N^{-\frac{1}{2} \xi_{ir}^{(1)}} \leq C_{ir} N^\alpha m_{ir}(t),$$

where the second inequality holds provided $\alpha \geq \frac{1}{2} \left(\xi_{ir}^{(1)} - \delta_{i1}^{(1)}\right) = \frac{\lambda_i}{2\lambda_1} \left(\frac{1+c_0}{1-c_0} - \frac{1-c_0}{1+c_0} \frac{\lambda_r}{\lambda_1}\right)$. Since this must hold for every $i \neq 1$, we need $\alpha \geq \frac{\lambda_r}{2\lambda_1} \left(\frac{1+c_0}{1-c_0} - \frac{1-c_0}{1+c_0} \frac{\lambda_r}{\lambda_1}\right)$. Comparing with the upper bound in (5.31), we see that c_0 must satisfy $(1+c_0)\lambda_r < (1-c_0)\lambda_1$, i.e., $c_0 < \frac{\lambda_1 - \lambda_r}{\lambda_1 + \lambda_r}$. For $i = 1$ a similar computation shows that it is required $\alpha \geq \frac{1}{2} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r}{\lambda_1}\right)$. Thus, we can choose $\alpha = \frac{1}{2} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right)$. This shows that $\mathcal{T}_{E_1(\varepsilon, c_0)} = \mathcal{T}_\varepsilon^{(11,+)}$ with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - \tilde{\eta}_2$. \square

We now prove Lemma 5.6.

Proof of Lemma 5.6. We provide the proof for the first step of the sequential elimination phenomenon, i.e., we consider $k = 2$ and assume that $\mathbf{X}_0 \in E_1(\varepsilon, c_0)$. In particular, we have that $m_{11}(0) = \varepsilon$, and for all $(i, j) \neq (1, 1)$, the initial values satisfy

$$\gamma_-^{(ij)} N^{-\frac{1}{2} \xi_{ij}^{(1)}} \leq m_{ij}(0) \leq \gamma_+^{(ij)} N^{-\frac{1}{2} \delta_{ij}^{(1)}},$$

where $\gamma_-^{(ij)}$ and $\gamma_+^{(ij)}$ are constants of order 1. As in the proof of Lemma 5.3, a modification of Proposition 4.3 is required, since the correlations at initialization are not in the scale $\Theta(N^{-\frac{1}{2}})$. Let T_2 be a fixed time horizon. Since $\delta < \sqrt{N/\alpha_2}$ by assumption (5.18), it follows from Lemma 4.4 that

$$\left| m_{ij}(t) - \left(m_{ij}(0) - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_j \rangle \right) \right| \leq \frac{\delta}{N} \sum_{\ell=1}^t |a_i(\ell-1)|,$$

for all $t \leq T_2$, with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - T_2 e^{-cN}$. The noise term $\frac{\delta}{N} \left| \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_j \rangle \right|$ can be upper-bounded according to Lemma 4.5, i.e., for every $a > 0$,

$$\frac{\delta}{N} \left| \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_j \rangle \right| \leq a,$$

with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - 2e^{-a^2 N^2 / (\delta^2 T_2)}$. For all $t \leq T_2$ and constant $K_N > 0$, we have that

$$\frac{\delta}{N} \sum_{\ell=1}^{T_2} |a_i(\ell-1)| \leq \frac{\delta}{N} \sum_{\ell=1}^{T_2} |h_i(\ell-1)| + \frac{\delta^2 \alpha_2}{2N} \sum_{\ell=1}^{T_2} \sum_{k=1}^r \frac{\delta}{N} |\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| \mathbf{1}_{\{|\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| > K_N\}},$$

where

$$|h_i(t)| = \frac{\delta \alpha_2}{2} \sum_{k=1}^r \left(|m_{ik}(t)| + \frac{\delta}{N} |\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_k \rangle| + \frac{\delta K_N}{N} \right).$$

We now set the truncation value K_N so that the term $\frac{\delta K_N}{N}$ is at most 1, which implies that $K_N \lesssim \frac{\log(N)N^{1-\frac{1}{2}} \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}}{C_\delta d_0}$. This ensures that for sufficiently large N , $|h_i(t)| \lesssim \delta \alpha_2 r$ for every $i \in [r]$ and every $t \leq T_2$. Therefore, we have that

$$\frac{\delta}{N} \sum_{\ell=1}^{T_2} |h_i(\ell-1)| \lesssim \frac{d_0^2 N^{-\left(1-\frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right)}}{\log(N)^2} T_2.$$

For every time horizon $T_2 \ll \frac{\log(N)^2 N^{1-\frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}}}{d_0^2}$, this can be set to any arbitrary fraction of the initial value ε . Similarly, for every $T_2 \ll \frac{\log(N)^2}{d_0^2} N^{\frac{1}{2}+\frac{1}{2}\frac{1-c_0}{1+c_0}-\frac{1-c_0}{1+c_0}\frac{\lambda_r^2}{\lambda_1^2}}$, this can be set to any arbitrary fraction of the initial value $N^{-\frac{1}{2}\delta_{ij}^{(1)}}$. Finally, applying Lemma 4.6 to the second summand of $|a_i(t)|$, we obtain that for every $a > 0$,

$$\frac{\delta^3 \alpha_2}{2N^2} \left| \sum_{\ell=1}^t \sum_{k=1}^r \frac{\delta}{N} \langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle \mathbf{1}_{\{|\langle \mathbf{v}_i, (\nabla_{\text{St}} H^\ell(\mathbf{X}_{\ell-1}))_k \rangle| > K_N\}} \right| \leq a$$

with \mathbb{P} -probability at least $1 - \frac{\delta^3}{N^2 a} T_2 e^{-\log(N)^2 N^{2-\frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}} / d_0^2}$. Combining all estimates we find that

$$(1-c_0)\varepsilon - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_1, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_1 \rangle \leq m_{11}(t) \leq (1+c_0)\varepsilon - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_1, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_1 \rangle, \quad (5.36)$$

with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - T_2 e^{-c_1 N} - 2e^{-c_2 c_0^2 \varepsilon^2 N^2 / (\delta^2 T_2)} - \frac{C \delta^3}{N^2 c_0 \varepsilon} T_2 e^{-c_3 \log(N)^2 N^{2-\frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}} / d_0^2}$. Similarly, for every $(i, j) \neq (1, 1)$, we obtain that

$$(1-c_0)N^{-\frac{1}{2}\xi_{ij}^{(1)}} - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle \leq m_{ij}(t) \leq (1+c_0)N^{-\frac{1}{2}\delta_{ij}^{(1)}} - \frac{\delta}{N} \sum_{\ell=1}^t \langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_{\ell-1}))_j \rangle, \quad (5.37)$$

with $\mathbb{P}_{\mathbf{X}_0}$ -probability at least $1 - T_2 e^{-c_1 N} - 2e^{-c_2 c_0^2 N^{1+\frac{1-c_0}{1+c_0} \frac{\lambda_2^2}{\lambda_1^2}} / (\delta^2 T_2)} - \frac{K \delta^3}{c_0 N^{\frac{3}{2}+\frac{1}{2}\frac{1+c_0}{1-c_0} \frac{\lambda_2^2}{\lambda_1^2}}} T_2 e^{-c_3 \log(N)^2 N^{2-\frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}} / d_0^2}$.

At this point, note that (5.36) and (5.37) only depend on the population term:

$$-\langle \mathbf{v}_i, (\nabla_{\text{St}} \Phi(\mathbf{X}_t))_j \rangle = 2\sqrt{N} \lambda_i \lambda_j m_{ij}^{p-1}(t) - \sqrt{N} \sum_{1 \leq k, \ell \leq 2} \lambda_k (\lambda_j + \lambda_\ell) m_{i\ell}(t) m_{kj}(t) m_{k\ell}(t).$$

The rest of the proof involves a detailed analysis of this term by distinguishing between the cases $i = j = 1$, $i = 1, j \neq 1$, $i \neq 1, j = 1$, and $i \neq 1, j \neq 1$. For this, we refer the reader to the proof of Lemma 5.13 of [6], where a detailed analysis using Langevin dynamics is carried out. We can apply the same strategy using (5.36) and (5.37) along with discrete difference inequalities, see e.g. Lemmas A.1 and A.2. Following this analysis, we obtain that

$$\mathcal{T}_{E_2(\varepsilon, c_0)} \lesssim \frac{\log(\varepsilon N^{\xi_{22}^{(1)}/2} / (1-c_0)) \log(\varepsilon \sqrt{N})}{C_\delta d_0 (1-c_0)} N^{\frac{1}{2}} \left(1 - \frac{1-c_0}{1+c_0} \frac{\lambda_r^2}{\lambda_1^2}\right),$$

which satisfies the prescribed condition on the time horizon. \square

6. PROOF OF SUBSPACE RECOVERY

In this section, we consider the case where $p = 2$ and $\lambda_1 = \dots = \lambda_r \equiv \lambda > 0$ and prove Theorem 3.7. In particular, we focus on the evolution of the eigenvalues of the matrix-valued function $\mathbf{G}(\mathbf{X}) = \mathbf{V}^\top \mathbf{X} \mathbf{X}^\top \mathbf{V}$. We let $\theta_1(t), \dots, \theta_r(t)$ denote the eigenvalues of $\mathbf{G}_t = \mathbf{G}(\mathbf{X}_t)$, where \mathbf{X}_t denotes the online SGD output at time t . Note that $\|\mathbf{G}_t\| = \theta_{\max}(t)$.

Our first result of this section is about the weak recovery of all the eigenvalues.

Lemma 6.1. For every $\gamma_1 > 1 > \gamma_2 > 0$, every $\varepsilon \in (0, \frac{1}{2})$, and every $C_0 \in (0, \frac{1}{4})$, there exist $\varepsilon_0 < \varepsilon$, $c_0 < C_0$, and a parameter $d_0 = d_0(N)$, which goes to zero arbitrarily slowly as $N \rightarrow \infty$, such that if

$$\delta = C_\delta \frac{d_0 \sqrt{N}}{\log(N)^2},$$

for some constant $C_\delta > 0$, then

$$\inf_{\mathbf{X}_0 \in C'_1(\gamma_1, \gamma_2)} \mathbb{P}_{\mathbf{X}_0}(\theta_{\min}(T) \geq \varepsilon_0) \geq 1 - K_1 \frac{\log(N)^3}{d_0} e^{-c_1 N} - K_2 e^{-c_2/d_0} - K_3 \frac{\log(N)^3}{d_0 N},$$

where T is given by

$$T = \frac{\log(N)^2 \log\left(\frac{4\varepsilon_0 N}{(1-c_0)\gamma_2}\right)}{4C_\delta \lambda^2 (1-2c_0)d_0}.$$

Proof. Since $\mathbf{X}_0 \in C'_1(\gamma_1, \gamma_2)$, it holds that $\frac{\gamma_2}{N} \leq \theta_{\min}(0) \leq \theta_{\max}(0) \leq \frac{\gamma_1}{N}$. For a constant $\xi > 2\frac{\gamma_1}{\gamma_2}$, we define the hitting time \mathcal{T}_ξ as

$$\mathcal{T}_\xi = \min\{t \geq 0: \theta_{\max}(t) \geq \xi \theta_{\min}(t)\}.$$

We easily note that $\mathcal{T}_\xi > 0$. According to Lemma 4.7 and using the fact that $-\|\mathbf{S}_t\| \mathbf{I}_r \preceq \mathbf{S}_t \preceq \|\mathbf{S}_t\| \mathbf{I}_r$, for every time horizon T and every $t \leq T$, we have the partial ordering

$$-\|\mathbf{S}_t\| \mathbf{I}_r - \mathbf{A}_t \preceq \mathbf{G}_t - \left(\mathbf{G}_0 + \frac{4\lambda^2 \delta}{\sqrt{N}} \sum_{\ell=1}^t \mathbf{G}_{\ell-1} (\mathbf{I}_r - \mathbf{G}_{\ell-1}) \right) \preceq \|\mathbf{S}_t\| \mathbf{I}_r + \mathbf{A}_t, \quad (6.1)$$

with probability at least $1 - T \exp(-c_1 N) - 4r^2 T \exp\left(-c_2 \frac{K_N^2}{\sigma^2}\right)$, where \mathbf{S}_t is given by (4.12) and \mathbf{A}_t by (4.13). Let \mathcal{T}_{bad} denote the hitting time

$$\mathcal{T}_{\text{bad}} = \min\left\{t \geq 1: c_0 \lambda_{\min} \left(\mathbf{G}_0 + \frac{4\lambda^2 \delta}{\sqrt{N}} \sum_{\ell=1}^t \mathbf{G}_{\ell-1} (\mathbf{I}_r - \mathbf{G}_{\ell-1}) \right) \leq \|\mathbf{S}_t\| \right\},$$

where $c_0 \in (0, \frac{1}{4})$ is a control parameter, chosen as $c_0 = \frac{C_1}{\log(N)}$, for a constant $C_1 = C_1(\lambda, \sigma^2, r, \gamma_1, \gamma_2)$ to be determined later. Moreover, we set $K_N = C_2 \sqrt{\log(N)}$ for a constant C_2 which may depend only on $\lambda, \sigma^2, r, \gamma_1, \gamma_2$. We first show that $\mathcal{T}_{\text{bad}} > 1$. Recalling that \mathbf{S}_t is given by

$$\mathbf{S}_t = \frac{4\lambda\delta}{N} \sum_{\ell=1}^t \text{sym} \left(\mathbf{V}^\top \Pi_{\mathbf{X}_{\ell-1}}^{\text{St}} \left(\mathbf{W}^\ell \mathbf{X}_{\ell-1} \right) \mathbf{M}_{\ell-1}^\top \right),$$

it follows that

$$\|\text{sym} \left(\mathbf{V}^\top \Pi_{\mathbf{X}_0}^{\text{St}} \left(\mathbf{W}^1 \mathbf{X}_0 \right) \mathbf{M}_0^\top \right)\| \leq 2rK_N \|\mathbf{G}_0\|^{1/2} \leq 2rK_N \frac{\sqrt{\gamma_1}}{\sqrt{N}},$$

where we used that $\|\mathbf{V}^\top \Pi_{\mathbf{X}_0}^{\text{St}} \left(\mathbf{W}^1 \mathbf{X}_0 \right)\| \leq 2rK_N$ from the proof of Lemma 4.7. Using the prescribed values for δ and K_N , we obtain that

$$\|\mathbf{S}_1\| = \frac{4\lambda\delta}{N} \|\text{sym} \left(\mathbf{V}^\top \Pi_{\mathbf{X}_0}^{\text{St}} \left(\mathbf{W}^1 \mathbf{X}_0 \right) \mathbf{M}_0^\top \right)\| \leq C_3 \frac{d_0}{N \log(N)^{3/2}} \ll c_0 \frac{\gamma_2}{N},$$

where $C_3 = C_3(\lambda, \sigma^2, r, \gamma_1, \gamma_2)$. This shows that $\mathcal{T}_{\text{bad}} > 1$.

We now introduce the following hitting times. For every $\varepsilon \in (0, 1)$, we define

$$\mathcal{T}_\varepsilon^{(\text{max},+)} = \min\{t \geq 0: \theta_{\max}(t) \geq \varepsilon\},$$

$$\mathcal{T}_\varepsilon^{(\text{min},+)} = \min\{t \geq 0: \theta_{\min}(t) \geq \varepsilon\},$$

$$\mathcal{T}_\varepsilon^{(\text{min},-)} = \min\{t \geq 0: \theta_{\min}(t) \leq \varepsilon\}.$$

Set the threshold $\varepsilon_N = \frac{\gamma_2}{2N}$ and consider $\varepsilon \in (0, \frac{1}{2})$. We want to establish a comparison inequality for $\theta_i(t)$ for every $t \leq \mathcal{T}_{\varepsilon_N}^{(\text{min},-)} \wedge \mathcal{T}_\varepsilon^{(\text{min},+)} \wedge \mathcal{T}_{1/2}^{(\text{max},+)} \wedge \mathcal{T}_\xi \wedge \mathcal{T}_{\text{bad}}$. We first note that, on the considered time interval,

$$\|\mathbf{S}_t\| \leq c_0 \lambda_{\min} \left(\mathbf{G}_0 + \frac{4\lambda^2 \delta}{\sqrt{N}} \sum_{\ell=1}^t \mathbf{G}_{\ell-1} (\mathbf{I}_r - \mathbf{G}_{\ell-1}) \right),$$

yielding

$$\|\mathbf{S}_t\|_{\mathbf{I}_r} \preceq c_0 \left(\mathbf{G}_0 + \frac{4\lambda^2\delta}{\sqrt{N}} \sum_{\ell=1}^t \mathbf{G}_{\ell-1} (\mathbf{I}_r - \mathbf{G}_{\ell-1}) \right).$$

Next, we provide a bound on $\|\mathbf{A}_t - \mathbf{A}_{t-1}\|$ since $\mathbf{A}_t = \sum_{\ell=1}^t \|\mathbf{A}_\ell - \mathbf{A}_{\ell-1}\|$. We first note that, on the considered time interval,

$$c_0 \lambda_{\min} \left(\frac{4\lambda\delta}{\sqrt{N}} \mathbf{G}_{t-1} (\mathbf{I}_r - \mathbf{G}_{t-1}) \right) = c_0 \frac{4\lambda\delta}{\sqrt{N}} \min_{1 \leq i \leq r} \{\theta_i(t-1)(1-\theta_i(t-1))\} \geq 2c_0 \lambda \frac{\delta}{\sqrt{N}} \theta_{\min}(t-1), \quad (6.2)$$

where we used the fact that $t \leq \mathcal{T}_{1/2}^{(\max,+)}$. According to the expression of \mathbf{A}_t from Lemma 4.7, the operator norm of the difference $\mathbf{A}_t - \mathbf{A}_{t-1}$ can be bounded using the definition of the stopping time \mathcal{T}_ξ as

$$\|\mathbf{A}_t - \mathbf{A}_{t-1}\| \leq C_5 \left(\frac{K_N^2 \delta^2}{N^2} + \frac{2\xi\delta^2}{N} \theta_{\min}(t-1) + \frac{K_N \sqrt{\xi} \delta^3}{N^2} \theta_{\min}(t-1)^{1/2} + \frac{\xi\delta^3}{N^{3/2}} \theta_{\min}(t-1) \right),$$

where C_5 is a constant which depends on $\lambda, \sigma^2, r, \gamma_1, \gamma_2$. Using the prescribed values for δ and K_N , we obtain that

$$\|\mathbf{A}_t - \mathbf{A}_{t-1}\| \leq C_6 \frac{d_0^2}{\log(N)^4} \theta_{\min}(t-1) \leq 2c_0 \lambda \frac{\delta}{\sqrt{N}} \theta_{\min}(t-1),$$

provided d_0 is a sufficiently small constant of order one and N is sufficiently large (with a benign condition). This estimate, combined with (6.2), leads to the inequality:

$$\|\mathbf{A}_t - \mathbf{A}_{t-1}\| \leq c_0 \lambda_{\min} \left(\frac{4\lambda\delta}{\sqrt{N}} \mathbf{G}_{t-1} (\mathbf{I}_r - \mathbf{G}_{t-1}) \right),$$

which, in turn, implies that

$$\mathbf{A}_t - \mathbf{A}_{t-1} \preceq c_0 \left(\frac{4\lambda\delta}{\sqrt{N}} \mathbf{G}_\ell (\mathbf{I}_r - \mathbf{G}_\ell) \right).$$

From (6.1) we obtain the following ordering for every $t \leq \mathcal{T}_{\varepsilon_N}^{(\min,-)} \wedge \mathcal{T}_\varepsilon^{(\min,+)} \wedge \mathcal{T}_{1/2}^{(\max,+)} \wedge \mathcal{T}_\xi \wedge \mathcal{T}_{\text{bad}}$:

$$\begin{aligned} \mathbf{G}_t &\succeq (1-c_0)\mathbf{G}_0 + 4(1-2c_0) \frac{\delta\lambda^2}{\sqrt{N}} \sum_{\ell=1}^t \mathbf{G}_{\ell-1} (\mathbf{I}_r - \mathbf{G}_{\ell-1}), \\ \mathbf{G}_t &\preceq (1+c_0)\mathbf{G}_0 + 4(1+2c_0) \frac{\delta\lambda^2}{\sqrt{N}} \sum_{\ell=1}^t \mathbf{G}_{\ell-1} (\mathbf{I}_r - \mathbf{G}_{\ell-1}), \end{aligned} \quad (6.3)$$

yielding

$$\begin{aligned} \theta_{\min}(t) &\geq (1-c_0) \frac{\gamma_2}{N} + 4(1-2c_0) \frac{\lambda^2\delta}{\sqrt{N}} \sum_{\ell=1}^t \theta_{\min}(\ell-1)(1-\theta_{\min}(\ell-1)) \\ \theta_{\max}(t) &\leq (1+c_0) \frac{\gamma_1}{N} + 4(1+2c_0) \frac{\lambda^2\delta}{\sqrt{N}} \sum_{\ell=1}^t \theta_{\max}(\ell-1)(1-\theta_{\max}(\ell-1)). \end{aligned} \quad (6.4)$$

According to Lemma A.2 and since $t \leq \mathcal{T}_\varepsilon^{(\min,+)} \wedge \mathcal{T}_{1/2}^{(\max,+)}$, the eigenvalues $\theta_1(t), \dots, \theta_r(t)$ are increasing in t . Consequently, we have that $\mathcal{T}_{\varepsilon_N}^{(\min,-)} > \mathcal{T}_\varepsilon^{(\min,+)} \wedge \mathcal{T}_{1/2}^{(\max,+)}$. Moreover, applying Lemma A.2 to (6.4) and performing some straightforward manipulations, we have that

$$\frac{(1-c_0)\gamma_2}{4N} \exp\left(\frac{4(1-2c_0) \frac{\delta}{\sqrt{N}} \lambda^2 t}{1 + 4(1-2c_0) \frac{\delta}{\sqrt{N}} \lambda^2} \right) \leq \theta_{\min}(t) \leq \theta_{\max}(t) \leq 2(1+c_0) \frac{\gamma_1}{N} \exp\left(4(1+2c_0) \frac{\delta}{\sqrt{N}} \lambda^2 t \right), \quad (6.5)$$

for every $t \leq \frac{K \log\left(\frac{N}{\gamma_1}\right)}{4(1+2c_0) \frac{\delta}{\sqrt{N}} \lambda^2} \wedge \mathcal{T}_\varepsilon^{(\min,+)} \wedge \mathcal{T}_{1/2}^{(\max,+)} \wedge \mathcal{T}_\xi \wedge \mathcal{T}_{\text{bad}}$, provided N is sufficiently large and K is an order 1 constant.

To complete the proof, we need to show that $\mathcal{T}_\varepsilon^{(\min,+)} \wedge \mathcal{T}_{1/2}^{(\max,+)} \leq \frac{K \log(\frac{N}{\gamma_1})}{4(1+c_0)\frac{\delta}{\sqrt{N}}\lambda^2} \wedge \mathcal{T}_{\text{bad}} \wedge \mathcal{T}_\xi$. We first focus on \mathcal{T}_{bad} . The comparison inequality (6.4) is valid on the considered time interval and implies that

$$\frac{4\lambda^2\delta}{\sqrt{N}} \sum_{\ell=1}^t \theta_{\min}(\ell-1) \leq \frac{1}{(1-2c_0)(1-\gamma_2/N)} \theta_{\min}(t) \leq 4\|\mathbf{G}_t\|,$$

for any $c_0 < 1/4$. Then, since $t \leq \mathcal{T}_\xi$, it holds that

$$\frac{4\lambda^2\delta}{\sqrt{N}} \sum_{\ell=1}^t \|\mathbf{G}_{\ell-1}\| = \frac{4\lambda^2\delta}{\sqrt{N}} \sum_{\ell=1}^t \theta_{\max}(\ell-1) \leq \frac{4\lambda^2\delta}{\sqrt{N}} \xi \sum_{\ell=1}^t \theta_{\min}(\ell-1) \leq 4\xi\|\mathbf{G}_t\|.$$

Combining this estimate with (6.5), we obtain that

$$\left(\frac{4\lambda\delta}{N}\right)^2 \sum_{\ell=1}^t \|\mathbf{G}_{\ell-1}\| \leq \frac{16\xi\delta}{N^{3/2}} \|\mathbf{G}_t\| \leq \frac{32\xi\gamma_1\delta}{N^{5/2}} (1+c_0) \exp\left(4(1+2c_0)\frac{\delta}{\sqrt{N}}\lambda^2 t\right). \quad (6.6)$$

We let $\hat{w}(t)$ denote the right-hand side of (6.6). Let us now find a lower bound on the other quantity appearing in the definition of the hitting time \mathcal{T}_{bad} on the same time interval. From the partial ordering (6.3), it follows that

$$\mathbf{G}_0 + \sum_{\ell=1}^t \frac{4\lambda^2\delta}{\sqrt{N}} \mathbf{G}_{\ell-1}(\mathbf{I}_r - \mathbf{G}_{\ell-1}) \succeq \frac{1}{1+2c_0} \mathbf{G}_t,$$

which leads to

$$\begin{aligned} \lambda_{\min}\left(\mathbf{G}_0 + \sum_{\ell=1}^t \frac{4\lambda^2\delta}{\sqrt{N}} \mathbf{G}_{\ell-1}(\mathbf{I}_r - \mathbf{G}_{\ell-1})\right) &\geq \frac{1}{1+2c_0} \theta_{\min}(t) \\ &\geq \frac{1}{1+2c_0} \frac{(1-c_0)\gamma_2}{4N} \exp\left(\frac{4(1-2c_0)\frac{\delta}{\sqrt{N}}\lambda^2 t}{1+4(1-2c_0)\frac{\delta}{\sqrt{N}}\lambda^2}\right). \end{aligned} \quad (6.7)$$

We let $\hat{v}(t)$ denote the function on the last line of (6.7). We then note that \mathcal{T}_{bad} can be bounded below by the hitting time

$$\tilde{\mathcal{T}}_{\text{bad}} = \min\{t \geq 0: c_0\hat{v}(t) \leq \|\mathbf{S}_t\|\}.$$

To control this hitting time, we first compare $\sqrt{\hat{w}(t)}$ with $c_0\hat{v}(t)$ using (6.6) and (6.7), i.e.,

$$\frac{\sqrt{\hat{w}(t)}}{c_0\hat{v}(t)} \leq \frac{\sqrt{32\xi\gamma_1}}{\gamma_2} \frac{\sqrt{\delta}}{c_0(1-c_0)N^{1/4}} \exp\left(\frac{2\delta}{\sqrt{N}}\lambda^2 t \left(1+2c_0 - \frac{2(1-2c_0)}{1+4(1-2c_0)\frac{\delta}{\sqrt{N}}\lambda^2}\right)\right).$$

Using the prescribed values for c_0 and δ , we see that the argument of the exponential in the previous display is negative for every $t \geq 0$, thus reaching the following bound:

$$\frac{\sqrt{\hat{w}(t)}}{c_0\hat{v}(t)} \leq \frac{\sqrt{32\xi\gamma_1}}{\gamma_2} \frac{\sqrt{\delta}}{c_0(1-c_0)N^{1/4}}.$$

This implies that we can bound $\tilde{\mathcal{T}}_{\text{bad}}$ using the hitting time

$$\tilde{\mathcal{T}}'_{\text{bad}} = \inf\left\{t \geq 0: \|\mathbf{S}_t\| \geq \frac{\gamma_2 c_0 (1-c_0) N^{1/4}}{\sqrt{32\xi\gamma_1\delta}} \sqrt{\hat{w}(t)}\right\}.$$

If we let $a(t) = \frac{\gamma_2 c_0 (1-c_0) N^{1/4}}{\sqrt{32\xi\gamma_1\delta}} \sqrt{\hat{w}(t)}$ and $b(t) = \frac{32\xi\gamma_1}{\gamma_2^2 c_0^2 (1-c_0^2)} \frac{\delta}{\sqrt{N}} a^2(t)$, Lemma 4.8 then implies that

$$\mathbb{P}(\exists t \in [0, T]: \|\mathbf{S}_t\| \geq a(t)) \leq 2r \exp\left(-K \frac{\gamma_2^2 c_0^2 (1-c_0^2) \sqrt{N}}{64\xi\gamma_1}\right) \leq 2r \exp\left(-\frac{c_1}{d_0}\right),$$

where we used the prescribed values for c_0 and δ , and the fact that $\lambda_{\max}(\mathbf{Q}_t) = \left(\frac{4\lambda\delta}{N}\right)^2 \sum_{\ell=1}^t \|\mathbf{G}_{\ell-1}\| \leq \hat{w}(t)$ with probability at least $1 - T \exp(-c_1 N) - 4r^2 T \exp\left(-c_2 \frac{K^2 N}{\sigma^2}\right)$ owing to Lemma 4.7. Choosing $d_0 = d_0(N)$ going to zero arbitrarily slowly, we have with high probability that $\mathcal{T}_\varepsilon^{(\min,+)} \wedge \mathcal{T}_{1/2}^{(\max,+)} \wedge \mathcal{T}_\xi \wedge \frac{K \log(N/\gamma_1)}{4(1+2c_0)\frac{\delta}{\sqrt{N}}\lambda^2} \leq \mathcal{T}_{\text{bad}}$.

We now need to control the hitting time \mathcal{T}_ξ . To this end, we study the maximal deviation between $\theta_{\min}(t)$ and $\theta_{\max}(t)$ using (6.5). Once $\theta_{\max}(t)$ exceeds ε , it remains of order 1, while the remaining eigenvalues continue to rise. We can therefore upper bound the maximal deviation between the largest and smallest eigenvalue of $\mathbf{G}(t)$ by computing an upper bound on the hitting time $\mathcal{T}_\varepsilon^{(\min,+)}$ (which is also an upper bound on the hitting time $\mathcal{T}_\varepsilon^{(\max,+)}$ since $\mathcal{T}_\varepsilon^{(\max,+)} < \mathcal{T}_\varepsilon^{(\min,+)}$) by solving

$$\frac{(1-c_0)\gamma_2}{4N} \exp\left(\frac{4(1-2c_0)\frac{\delta}{\sqrt{N}}\lambda^2 t}{1+4(1-2c_0)\frac{\delta}{\sqrt{N}}\lambda^2}\right) = \varepsilon,$$

which gives that

$$\mathcal{T}_\varepsilon^{(\min,+)} \leq T_{u,\varepsilon}^{\min} = \frac{\log\left(\frac{4\varepsilon N}{(1-c_0)\gamma_2}\right)}{\frac{4(1-2c_0)\frac{\delta}{\sqrt{N}}\lambda^2}{1+4(1-2c_0)\frac{\delta}{\sqrt{N}}\lambda^2}}.$$

Using this result along with the comparison inequality (6.5), we obtain that

$$\begin{aligned} \frac{\theta_{\max}(T_{u,\varepsilon}^{\min})}{\theta_{\min}(T_{u,\varepsilon}^{\min})} &\leq 8 \frac{(1+c_0)\gamma_1}{(1-c_0)\gamma_2} \exp\left(\frac{4c_0 + 4(1-2c_0)(1+2c_0)\lambda^2 \frac{\delta}{\sqrt{N}}}{1-2c_0} \log\left(\frac{4\varepsilon N}{(1-c_0)\gamma_2}\right)\right) \\ &\leq 8 \frac{(1+c_0)\gamma_1}{(1-c_0)\gamma_2} \left(\frac{4\varepsilon N}{(1-c_0)\gamma_2}\right)^{4c_0}, \end{aligned} \quad (6.8)$$

where we used the prescribed value for δ . Since $c_0 = \mathcal{O}\left(\frac{1}{\log(N)}\right)$, we conclude that the maximal deviation between θ_{\min} and θ_{\max} is at most a constant given by the last line of (6.8). Choosing ξ to be this constant, we deduce that $\mathcal{T}_\xi > \mathcal{T}_\varepsilon^{(\min,+)} \wedge \mathcal{T}_{1/2}^{(\max,+)}$. Moreover, since $\theta_{\min}(T_{u,\varepsilon}^{\min}) \leq \varepsilon$, from (6.8) we have that

$$\theta_{\max}(T_{u,\varepsilon}^{\min}) \leq \varepsilon^{1+4c_0} 8 \frac{(1+c_0)\gamma_1}{(1-c_0)\gamma_2} \left(\frac{4}{(1-c_0)\gamma_2}\right)^{4c_0} N^{4c_0}.$$

We can choose ε sufficiently small so that $\theta_{\max}(\mathcal{T}_\varepsilon^{(\min,+)}) \leq 1/2$. This shows that $\mathcal{T}_\varepsilon^{(\min,+)} \leq \mathcal{T}_{1/2}^{(\max,+)}$. We conclude the proof by choosing a time horizon T equal to $T_{u,\varepsilon}^{\min}$. \square

The following lemma shows that weak implies strong recovery.

Lemma 6.2. *Let $p = 2$ and assume that $\lambda_1 = \dots = \lambda_r \equiv \lambda$. Assume that $\delta = C_\delta \frac{d_0 \sqrt{N}}{\log(N)^2}$, where $C_\delta > 0$ is a constant and $d_0 = d_0(N)$ is a control parameter such that $d_0 \rightarrow 0$ as $N \rightarrow \infty$. Then, for every $\varepsilon, \varepsilon' \in (0, 1)$,*

$$\inf_{\mathbf{x}_0 \in \{\theta_{\min}(0) \geq \varepsilon\}} \mathbb{P}_{\mathbf{X}_0}(\theta_{\min}(T) \geq 1 - \varepsilon') \geq 1 - K_1 \frac{\log(N)^2}{d_0} e^{-c_1 N} - K_2 \frac{\log(N)^2}{d_0 N},$$

where the time horizon T is given by

$$T = \frac{\log(N)^2 \log\left(\frac{1-\varepsilon'}{\varepsilon}\right)}{C_\delta d_0 \lambda^2 \varepsilon}.$$

Theorem 3.7 then follows straightforwardly by combining Lemmas 6.1 and 6.2 using the strong Markov property.

Proof. In the following, we let $\varepsilon' > \varepsilon > 0$ and let $c_0 \in (0, \frac{1}{2})$ independent on N . According to Lemma 4.7, for a fixed time horizon T to be set later and for every $t \leq T$, we have the partial ordering

$$\begin{aligned} \mathbf{G}_t &\succeq \mathbf{G}_0 + \frac{4\lambda^2 \delta}{\sqrt{N}} \sum_{\ell=1}^t \mathbf{G}_{\ell-1} (\mathbf{I}_r - \mathbf{G}_{\ell-1}) - \mathbf{S}_t - \mathbf{A}_t \\ \mathbf{G}_t &\preceq \mathbf{G}_0 + \frac{4\lambda^2 \delta}{\sqrt{N}} \sum_{\ell=1}^t \mathbf{G}_{\ell-1} (\mathbf{I}_r - \mathbf{G}_{\ell-1}) + \mathbf{S}_t + \mathbf{A}_t, \end{aligned}$$

with probability at least $1 - T \exp(-cN) - 4r^2 T \exp\left(-\frac{K_N^2}{C\sigma^2}\right)$. Set $K_N = C_1 \sqrt{\log(N)}$. From the expression of \mathbf{A}_t in Lemma 4.7, we have the following bound on the increments: for every $t \geq 0$,

$$\|\mathbf{A}_t - \mathbf{A}_{t-1}\| \leq \alpha_2(r, \lambda, \sigma^2) \left(K_N^2 \frac{\delta^2}{N^2} + \frac{\delta^2}{N^2} + K_N \frac{\delta^3}{N^2} + \frac{\delta^3}{N^{3/2}} \right) \leq C_3 \frac{d_0^2}{\log(N)^4},$$

where we used the fact that $\|\mathbf{G}_t\| \leq 1$. Then, since $\mathbf{X}_0 \in \{\theta_{\min}(0) \geq \varepsilon\}$, for any $t \leq \mathcal{T}_{\varepsilon/2}^{(\min,-)} \wedge \mathcal{T}_{1-\varepsilon'}^{(\min,+)} \wedge \mathcal{T}_{1-\varepsilon}^{(\max,+)}$ it holds that

$$\frac{4\lambda^2\delta}{\sqrt{N}} \lambda_{\min}(\mathbf{G}_{t-1}(\mathbf{I}_r - \mathbf{G}_{t-1})) = \frac{4\lambda^2\delta}{\sqrt{N}} \min_{1 \leq i \leq r} \{\theta_i(t-1)(1 - \theta_i(t-1))\} \geq C_4(r, \lambda, \sigma^2) \frac{d_0^2}{\log(N)^2}.$$

Thus, up to a benign condition on N , it holds that, for every $t \leq \mathcal{T}_{\varepsilon/2}^{(\min,-)} \wedge \mathcal{T}_{1-\varepsilon'}^{(\min,+)} \wedge \mathcal{T}_{1-\varepsilon}^{(\max,+)}$,

$$\frac{c_0}{2} \frac{4\lambda^2\delta}{\sqrt{N}} \mathbf{G}_{t-1}(\mathbf{I}_r - \mathbf{G}_{t-1}) \succeq \mathbf{A}_t - \mathbf{A}_{t-1}.$$

Furthermore, for every $t \leq T$ and , we note that the martingale term \mathbf{S}_t satisfies

$$\begin{aligned} \|\mathbf{S}_t - \mathbf{S}_{t-1}\| &= \frac{4\lambda\delta}{N} \|\text{sym}\left(\mathbf{V}^\top \Pi_{\mathbf{X}_{t-1}}^{\mathbf{S}_t} (\mathbf{W}^t \mathbf{X}_{t-1}) \mathbf{M}_{t-1}^\top\right)\| \\ &\leq \frac{4\lambda\delta}{N} \|\mathbf{V}^\top \Pi_{\mathbf{X}_{t-1}}^{\mathbf{S}_t} (\mathbf{W}^t \mathbf{X}_{t-1})\| \|\mathbf{M}_{t-1}\| \\ &\leq \frac{4\lambda\delta}{N} \left(\|\mathbf{V}^\top \mathbf{W}^t \mathbf{X}_{t-1}\| + \|\mathbf{X}_{t-1}^\top \mathbf{W}^t \mathbf{X}_{t-1}\|\right) \\ &\leq 4r\lambda \frac{\delta}{N} K_N \\ &\leq C_5 \frac{d_0}{\sqrt{N} \log(N)^{3/2}}, \end{aligned}$$

where the second inequality follows by (4.17) and the third inequality holds with probability at least $1 - 4r^2T \exp\left(-\frac{K_N^2}{C\sigma^2}\right)$ according to (4.20). We thus obtain that, on the considered time interval,

$$\mathbf{G}_t \succeq \mathbf{G}_0 + \frac{4\lambda^2\delta}{\sqrt{N}} (1 - c_0) \sum_{\ell=1}^t \mathbf{G}_{\ell-1}(\mathbf{I}_r - \mathbf{G}_{\ell-1}),$$

with probability at least $1 - T \exp(-cN) - 4r^2T \exp\left(-\frac{K_N^2}{C\sigma^2}\right)$, ensuring that

$$\theta_{\min}(t) \geq \theta_{\min}(0) + \frac{4\lambda^2\delta}{\sqrt{N}} (1 - c_0) \varepsilon \sum_{\ell=1}^t \theta_{\min}(\ell - 1).$$

According to Lemma A.1 we then find that

$$\theta_{\min}(t) \geq \theta_{\min}(0) \left(1 + \frac{4\lambda^2\varepsilon\delta}{\sqrt{N}} (1 - c_0)\right)^t.$$

We then have the following upper bound on $\mathcal{T}_{1-\varepsilon'}^{(\min,+)}$:

$$\mathcal{T}_{1-\varepsilon'}^{(\min,+)} \geq \frac{\log\left(\frac{1-\varepsilon'}{\varepsilon}\right)}{\log\left(1 + \frac{4\lambda^2\delta}{\sqrt{N}} (1 - c_0)\varepsilon\right)}.$$

Since $x - \frac{x^2}{2} \leq \log(x)$ for every $x \in (0, 1)$, as long as $\frac{\delta}{\sqrt{N}}\lambda^2 \leq 1$, we have that

$$\mathcal{T}_{1-\varepsilon'}^{(\min,+)} \leq \frac{\sqrt{N} \log\left(\frac{1-\varepsilon'}{\varepsilon}\right)}{\delta\lambda^2\varepsilon}.$$

It suffices to take $c_0 = \frac{1}{2}$ and set the time horizon T to be the right-hand side of the previous equation to complete the proof. \square

APPENDIX A. COMPARISON INEQUALITIES FOR DISCRETE SEQUENCES

In this appendix, we present comparison inequalities for discrete time dynamical systems.

Lemma A.1 ([10, 2]). *Let $p \geq 2$ be an integer and $a_1, a_2, b_1, b_2 > 0$ be positive constants such that $a_1 \leq a_2$ and $b_1 \leq b_2$. Consider a sequence $(u_t)_{t \in \mathbb{N}}$ satisfying for all $t \in \mathbb{N}$,*

$$a_1 + b_1 \sum_{s=0}^{t-1} u_s^{p-1} \leq u_t \leq a_2 + b_2 \sum_{s=0}^{t-1} u_s^{p-1}.$$

Then,

(a) (Discrete Gronwall inequality) if $p = 2$, for all $t \in \mathbb{N}$ it holds that

$$a_1 (1 + b_1)^t \leq u_t \leq a_2 (1 + b_2)^t;$$

(b) (Discrete Bihari-LaSalle inequality) if $p \geq 3$, for all $t \in \mathbb{N}$ it holds that,

$$a_1 \left(1 - \frac{b_1}{(1 + b_1 u_{t-1}^{p-2})^{p-1}} a_1^{p-2} t \right)^{-\frac{1}{p-2}} \leq u_t \leq a_2 \left(1 - b_2 a_2^{p-2} t \right)^{-\frac{1}{p-2}}.$$

In particular, if $u_t \geq 0$ for all $t \in \mathbb{N}$, it holds that

$$a_1 \left(1 - b_1 a_1^{p-2} t \right)^{-\frac{1}{p-2}} \leq u_t \leq a_2 \left(1 - b_2 a_2^{p-2} t \right)^{-\frac{1}{p-2}}.$$

Lemma A.1 is proved in [10, 2]. The following lemma can be derived using analogous techniques. Since the proof mimics the arguments used for Lemma A.1, we omit the details here.

Lemma A.2. Let $a_1, a_2, b_1, b_2 > 0$ be positive constants such that $a_1 \leq a_2$ and $b_1 \leq b_2$. Consider a sequence $(u_t)_{t \in \mathbb{N}}$ taking values in $[0, 1]$ and satisfying for all $t \in \mathbb{N}$,

$$a_1 + b_1 \sum_{s=0}^{t-1} u_s (1 - u_s) \leq u_t \leq a_2 + b_2 \sum_{s=0}^{t-1} u_s (1 - u_s).$$

Then, for every $0 \leq t \leq \min \{t: u_t \geq \frac{1}{2}\}$,

$$\frac{\frac{a_1}{1-a_1} e^{\frac{b_1}{1+b_1} t}}{1 + e^{\frac{b_1}{1+b_1} t} \frac{a_1}{1-a_1}} \leq u_t \leq \frac{\frac{a_2}{1-a_2} e^{b_2 t}}{1 + e^{b_2 t} \frac{a_2}{1-a_2}}.$$

APPENDIX B. MATRIX CONCENTRATION INEQUALITIES

We provide a deviation inequality to control the operator norm of matrix martingale sequences. The next result is a straightforward modification of Theorem 2.3 from [61], which provides an extension of Freedman's inequality to matrix martingales.

Lemma B.1. Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in \mathbb{N}}, \mathbb{P})$ be a filtered probability space. Consider an adapted sequence $\{\mathbf{X}_t\}_{t \in \mathbb{N}}$ and a predictable sequence $\{\mathbf{V}_t\}_{t \in \mathbb{N}}$ of $d \times d$ self-adjoint matrices such that

$$\log(\mathbb{E}[\exp(\theta \mathbf{X}_t) | \mathcal{F}_{t-1}]) \preceq g(\theta) \mathbf{V}_t \quad \text{almost surely for each } \theta > 0,$$

for some function $g: (0, \infty) \rightarrow [0, \infty]$. Define the partial sum processes

$$\mathbf{Y}_t = \sum_{\ell=1}^t \mathbf{X}_\ell \quad \text{and} \quad \mathbf{W}_t = \sum_{\ell=1}^t \mathbf{V}_\ell.$$

Let T be a fixed time horizon. Then, for every scalar-valued continuous functions a and b we have that

$$\mathbb{P}(\exists t \in [0, T]: \lambda_{\max}(\mathbf{Y}_t) \geq a(t) \text{ and } \lambda_{\max}(\mathbf{W}_t) \leq b(t)) \leq d \inf_{\theta > 0} \sup_{t \in [0, T]} e^{-\theta a(t) + g(\theta) b(t)}.$$

Proof. Let \mathcal{E} denote the event

$$\mathcal{E} = \{\exists t \in [0, T]: \lambda_{\max}(\mathbf{Y}_t) \geq a(t) \text{ and } \lambda_{\max}(\mathbf{W}_t) \leq b(t)\}.$$

Let κ denote the stopping time given by

$$\kappa = \inf \{t \geq 0: \lambda_{\max}(\mathbf{Y}_t) \geq a(t) \text{ and } \lambda_{\max}(\mathbf{W}_t) \leq b(t)\}.$$

We observe that $\kappa \leq T < \infty$ on the event \mathcal{E} . Then, for every $t \in [0, T]$ and every $k \in [0, t]$, if \mathcal{E}_k^t denotes the event

$$\mathcal{E}_k^t = \{\lambda_{\max}(\mathbf{Y}_k) \geq a(t) \text{ and } \lambda_{\max}(\mathbf{W}_k) \leq b(t)\},$$

we have that

$$\mathcal{E} = \bigcup_{t=0}^T \left(\bigcup_{k=0}^t \mathcal{E}_k^t \right).$$

Now, for each fixed $\theta > 0$, consider the real-valued process $\{S_t = S_t(\theta): t \in \mathbb{N}\}$ given by

$$S_t(\theta) = \text{tr} \exp(\theta \mathbf{Y}_t - g(\theta) \mathbf{W}_t),$$

which is a real-valued positive supermartingale according to [61, Lemma 2.1]. For every $\theta > 0$, according to Lemma 2.2 from [61], we have that on each event $\bigcup_{k=0}^t \mathcal{E}_k^t$,

$$S_t \geq e^{\theta a(t) - g(\theta)b(t)},$$

so that on the event \mathcal{E} , it holds that

$$S_t \geq \inf_{t \in [0, T]} e^{\theta a(t) - g(\theta)b(t)}.$$

We then conclude the proof as done for [61, Theorem 2.3] upon noting that $\inf_{t \in [0, T]} e^{\theta a(t) - g(\theta)b(t)} = \sup_{t \in [0, T]} e^{-\theta a(t) + g(\theta)b(t)}$. \square

REFERENCES

- [1] E. Abbe, E. Boix-Adserà, and T. Misiakiewicz, “The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks”, [Proceedings of thirty fifth conference on learning theory](#), Vol. 178, edited by P.-L. Loh and M. Raginsky, Proceedings of Machine Learning Research (2022), pp. 4782–4887.
- [2] E. Abbe, E. Boix-Adserà, and T. Misiakiewicz, “Sgd learning on neural networks: leap complexity and saddle-to-saddle dynamics”, [Proceedings of thirty sixth conference on learning theory](#), Vol. 195, edited by G. Neu and L. Rosasco, Proceedings of Machine Learning Research (2023), pp. 2552–2623.
- [3] J. Ba, M. A. Erdogdu, T. Suzuki, Z. Wang, D. Wu, and G. Yang, “High-dimensional asymptotics of feature learning: how one gradient step improves the representation”, [Advances in neural information processing systems](#), Vol. 35, edited by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (2022), pp. 37932–37946.
- [4] F. Bach, *Breaking the curse of dimensionality with convex neural networks*, [J. Mach. Learn. Res.](#) **18**, 1–53 (2017), [MR3634886](#).
- [5] A. S. Bandeira, G. Cipolloni, D. Schröder, and R. van Handel, *Matrix concentration inequalities and free probability ii. two-sided bounds and applications*, (2024), [arXiv:2406.11453](#).
- [6] G. Ben Arous, C. Gerbelot, and V. Piccolo, *High-dimensional optimization for multi-spiked tensor pca*, (2024), [arXiv:2408.06401](#).
- [7] G. Ben Arous, R. Gheissari, J. Huang, and A. Jagannath, *High-dimensional SGD aligns with emerging outlier eigenspaces*, (2023), [arXiv:2310.03010](#).
- [8] G. Ben Arous, R. Gheissari, and A. Jagannath, *Bounding flows for spherical spin glass dynamics*, [Comm. Math. Phys.](#) **373**, 1011–1048 (2020), [MR4061404](#).
- [9] G. Ben Arous, R. Gheissari, and A. Jagannath, *Algorithmic thresholds for tensor PCA*, [Ann. Probab.](#) **48**, 2052–2087 (2020), [MR4124533](#).
- [10] G. Ben Arous, R. Gheissari, and A. Jagannath, *Online stochastic gradient descent on non-convex losses from high-dimensional inference*, [J. Mach. Learn. Res.](#) **22**, 1–51 (2021), [MR4279757](#).
- [11] G. Ben Arous, R. Gheissari, and A. Jagannath, “High-dimensional limit theorems for SGD: Effective dynamics and critical scaling”, [Advances in neural information processing systems](#), Vol. 35, edited by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (2022), pp. 25349–25362.
- [12] G. Ben Arous, R. Gheissari, and A. Jagannath, *High-dimensional limit theorems for SGD: effective dynamics and critical scaling*, [Comm. Pure Appl. Math.](#) **77**, 2030–2080 (2024), [MR4692886](#).
- [13] G. Ben Arous, D. Z. Huang, and J. Huang, *Long random matrices and tensor unfolding*, [Ann. Appl. Probab.](#) **33**, 5753–5780 (2023), [MR4677744](#).
- [14] G. Ben Arous, S. Mei, A. Montanari, and M. Nica, *The landscape of the spiked tensor model*, [Comm. Pure Appl. Math.](#) **72**, 2282–2330 (2019), [MR4011861](#).
- [15] R. Berthier, A. Montanari, and K. Zhou, *Learning time-scales in two-layers neural networks*, (2023), [arXiv:2303.00055](#).
- [16] A. Bietti, J. Bruna, and L. Pillaud-Vivien, *On learning gaussian multi-index models with gradient flow*, (2023), [arXiv:2310.19793](#).
- [17] A. Bietti, J. Bruna, C. Sanford, and M. J. Song, “Learning single-index models with shallow neural networks”, [Advances in neural information processing systems](#), Vol. 35, edited by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (2022), pp. 9768–9783.
- [18] L. Bottou, “Large-scale machine learning with stochastic gradient descent”, Proceedings of COMP-STAT’2010 (2010), pp. 177–186, [MR3362066](#).

- [19] L. Bottou and Y. Le Cun, “Large scale online learning”, *Advances in neural information processing systems*, Vol. 16, edited by S. Thrun, L. Saul, and B. Schölkopf (2003).
- [20] N. Boumal, *An introduction to optimization on smooth manifolds* (Cambridge University Press, Cambridge, 2023), pp. xviii+338, [MR4533407](#).
- [21] M. Celentano, C. Cheng, and A. Montanari, *The high-dimensional asymptotics of first order methods with random data*, (2021), [arXiv:2112.07572](#).
- [22] W.-K. Chen, *Phase transition in the spiked random tensor with Rademacher prior*, *Ann. Statist.* **47**, 2734–2756 (2019), [MR3988771](#).
- [23] W.-K. Chen, M. Handschy, and G. Lerman, *Phase transition in random tensors with multiple independent spikes*, *Ann. Appl. Probab.* **31**, 1868–1913 (2021), [MR4312849](#).
- [24] Y. Chikuse, *Statistics on special manifolds*, Vol. 174, Lecture Notes in Statistics (Springer-Verlag, New York, 2003), pp. xxvi+399, [MR1960435](#).
- [25] A. Damian, J. D. Lee, and M. Soltanolkotabi, “Neural networks can learn representations with gradient descent”, *Proceedings of thirty fifth conference on learning theory*, Vol. 178, edited by P.-L. Loh and M. Raginsky, Proceedings of Machine Learning Research (2022), pp. 5413–5452.
- [26] Y. Dandi, F. Krzakala, B. Loureiro, L. Pesce, and L. Stephan, *How two-layer neural networks learn, one (giant) step at a time*, (2023), [arXiv:2305.18270](#).
- [27] A. Dieuleveut, A. Durmus, and F. Bach, *Bridging the gap between constant step size stochastic gradient descent and Markov chains*, *Ann. Statist.* **48**, 1348–1382 (2020), [MR4124326](#).
- [28] R. Dudeja and D. Hsu, “Learning single-index models in gaussian space”, *Proceedings of the 31st conference on learning theory*, Vol. 75, edited by S. Bubeck, V. Perchet, and P. Rigollet, Proceedings of Machine Learning Research (2018), pp. 1887–1930.
- [29] A. Engel and C. Van den Broeck, *Statistical mechanics of learning* (Cambridge University Press, Cambridge, 2001), pp. xii+329, [MR1827275](#).
- [30] B. Follain and F. Bach, *Enhanced feature learning via regularisation: integrating neural networks and kernel methods*, (2024), [arXiv:2407.17280](#).
- [31] B. Follain, U. Simsekli, and F. Bach, *Nonparametric linear feature learning in regression through regularisation*, (2023), [arXiv:2307.12754](#).
- [32] S. Frei, Y. Cao, and Q. Gu, “Agnostic learning of a single neuron with gradient descent”, *Advances in neural information processing systems*, Vol. 33, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin (2020), pp. 5417–5428.
- [33] C. Gerbelot, E. Troiani, F. Mignacco, F. Krzakala, and L. Zdeborová, *Rigorous Dynamical Mean-Field Theory for Stochastic Gradient Descent Methods*, *SIAM J. Math. Data Sci.* **6**, 400–427 (2024), [MR4741502](#).
- [34] S. Goldt, M. S. Advani, A. M. Saxe, F. Krzakala, and L. Zdeborová, *Dynamics of stochastic gradient descent for two-layer neural networks in the teacher-student setup*, *J. Stat. Mech. Theory Exp.*, **124010**, 15 (2020), [MR4241363](#).
- [35] S. B. Hopkins, T. Schramm, J. Shi, and D. Steurer, “Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors”, *Proceedings of the forty-eighth annual acm symposium on theory of computing*, STOC ’16 (2016), pp. 178–191.
- [36] S. B. Hopkins, J. Shi, and D. Steurer, “Tensor principal component analysis via sum-of-square proofs”, *Proceedings of the 28th conference on learning theory*, Vol. 40, edited by P. Grünwald, E. Hazan, and S. Kale, Proceedings of Machine Learning Research (2015), pp. 956–1006.
- [37] J. Huang, D. Z. Huang, Q. Yang, and G. Cheng, *Power iteration for tensor PCA*, *J. Mach. Learn. Res.* **23**, 1–47 (2022), [MR4577080](#).
- [38] A. Jagannath, P. Lopatto, and L. Miolane, *Statistical thresholds for tensor PCA*, *Ann. Appl. Probab.* **30**, 1910–1933 (2020), [MR4132641](#).
- [39] I. M. Johnstone, *On the distribution of the largest eigenvalue in principal components analysis*, *Ann. Statist.* **29**, 295–327 (2001), [MR1863961](#).
- [40] C. Kim, A. S. Bandeira, and M. X. Goemans, “Community detection in hypergraphs, spiked tensor models, and sum-of-squares”, *2017 international conference on sampling theory and applications (sampta)* (2017), pp. 124–128.
- [41] Y. Le Cun, L. Bottou, Y. Bengio, and P. Haffner, *Gradient-based learning applied to document recognition*, *Proceedings of the IEEE* **86**, 2278–2324 (1998).
- [42] M. Lelarge and L. Miolane, *Fundamental limits of symmetric low-rank matrix estimation*, *Probab. Theory Related Fields* **173**, 859–929 (2019), [MR3936148](#).

- [43] T. Lesieur, L. Miolane, M. Lelarge, F. Krzakala, and L. Zdeborová, “Statistical and computational phase transitions in spiked tensor estimation”, [2017 IEEE International Symposium on Information Theory \(ISIT\)](#) (2017), pp. 511–515.
- [44] C. McDiarmid, “Concentration”, *Probabilistic methods for algorithmic discrete mathematics*, edited by M. Habib, C. McDiarmid, J. Ramirez-Alfonsin, and B. Reed (Springer Berlin Heidelberg, Berlin, Heidelberg, 1998), pp. 195–248.
- [45] F. Mignacco, F. Krzakala, P. Urbani, and L. Zdeborová, *Dynamical mean-field theory for stochastic gradient descent in Gaussian mixture classification*, *J. Stat. Mech. Theory Exp.*, Paper No. 124008, **23** (2021), [MR4412836](#).
- [46] A. Mousavi-Hosseini, S. Park, M. Girotti, I. Mitliagkas, and M. A. Erdogdu, *Neural networks efficiently learn low-dimensional representations with SGD*, (2023), [arXiv:2209.14863](#).
- [47] D. Needell, N. Srebro, and R. Ward, *Stochastic gradient descent, weighted sampling, and the randomized Kaczmarz algorithm*, *Mathematical Programming* **155**, 549–573 (2016).
- [48] Y. Nesterov, *Lectures on convex optimization*, Vol. 137, Springer Optimization and Its Applications (Springer, Cham, 2018), pp. xxiii+589, [MR3839649](#).
- [49] V. Pappas, *Traces of class/cross-class structure pervade deep learning spectra*, *J. Mach. Learn. Res.* **21**, 1–64 (2020), [MR4213433](#).
- [50] V. Pappas, X. Y. Han, and D. L. Donoho, *Prevalence of neural collapse during the terminal phase of deep learning training*, *Proc. Natl. Acad. Sci. USA* **117**, 24652–24663 (2020), [MR4250189](#).
- [51] A. Perry, A. S. Wein, and A. S. Bandeira, *Statistical limits of spiked tensor models*, *Ann. Inst. Henri Poincaré Probab. Stat.* **56**, 230–264 (2020), [MR4058987](#).
- [52] A. Perry, A. S. Wein, A. S. Bandeira, and A. Moitra, *Optimality and sub-optimality of PCA I: Spiked random matrix models*, *Ann. Statist.* **46**, 2416–2451 (2018), [MR3845022](#).
- [53] E. Richard and A. Montanari, “A statistical model for tensor PCA”, *Advances in neural information processing systems*, Vol. 27, edited by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger (2014), pp. 2897–2905.
- [54] H. Robbins and S. Monro, *A stochastic approximation method*, *Ann. Math. Statistics* **22**, 400–407 (1951), [MR42668](#).
- [55] V. Ros, G. Ben Arous, G. Biroli, and C. Cammarota, *Complex energy landscapes in spiked-tensor and simple glassy models: ruggedness, arrangements of local minima, and phase transitions*, *Phys. Rev. X* **9**, 011003 (2019).
- [56] D. Saad and S. A. Solla, *On-line learning in soft committee machines*, *Phys. Rev. E* **52**, 4225–4243 (1995).
- [57] S. Sarao Mannelli, G. Biroli, C. Cammarota, F. Krzakala, and L. Zdeborová, “Who is afraid of big bad minima? analysis of gradient-flow in spiked matrix-tensor models”, *Advances in neural information processing systems*, Vol. 32, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (2019).
- [58] S. Sarao Mannelli, F. Krzakala, P. Urbani, and L. Zdeborová, “Passed & spurious: descent algorithms and local minima in spiked matrix-tensor models”, *Proceedings of the 36th international conference on machine learning*, Vol. 97, edited by K. Chaudhuri and R. Salakhutdinov, Proceedings of Machine Learning Research (2019), pp. 4333–4342.
- [59] M. Soltanolkotabi, “Learning ReLU via gradient descent”, *Advances in neural information processing systems*, Vol. 30, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (2017).
- [60] Y. S. Tan and R. Vershynin, *Online stochastic gradient descent with arbitrary initialization solves non-smooth, non-convex phase retrieval*, *J. Mach. Learn. Res.* **24**, 1–47 (2023), [MR4582480](#).
- [61] J. A. Tropp, *Freedman’s inequality for matrix martingales*, *Electron. Commun. Probab.* **16**, 262–270 (2011), [MR2802042](#).
- [62] J. A. Tropp, *User-friendly tail bounds for sums of random matrices*, *Found. Comput. Math.* **12**, 389–434 (2012), [MR2946459](#).
- [63] R. Veiga, L. Stephan, B. Loureiro, F. Krzakala, and L. Zdeborová, *Phase diagram of stochastic gradient descent in high-dimensional two-layer neural networks*, *J. Stat. Mech. Theory Exp.*, Paper No. 114008, **28** (2023), [MR4725421](#).
- [64] R. Vershynin, *High-dimensional probability*, Vol. 47, Cambridge Series in Statistical and Probabilistic Mathematics, An introduction with applications in data science, With a foreword by Sara van de Geer (Cambridge University Press, Cambridge, 2018), pp. xiv+284, [MR3837109](#).

- [65] M. J. Wainwright, *High-dimensional statistics*, Vol. 48, Cambridge Series in Statistical and Probabilistic Mathematics, A non-asymptotic viewpoint (Cambridge University Press, Cambridge, 2019), pp. xvii+552, [MR3967104](#).
- [66] A. S. Wein, A. El Alaoui, and C. Moore, “The Kikuchi hierarchy and tensor PCA”, *2019 IEEE 60th Annual Symposium on Foundations of Computer Science* (IEEE Comput. Soc. Press, Los Alamitos, CA, 2019), pp. 1446–1468, [MR4228236](#).
- [67] Y. Wu and K. Zhou, “Lower bounds for the convergence of tensor power iteration on random overcomplete models”, [Proceedings of thirty sixth conference on learning theory](#), Vol. 195, edited by G. Neu and L. Rosasco, Proceedings of Machine Learning Research (2023), pp. 3783–3820.
- [68] Y. Xia, *A multiple-index model and dimension reduction*, *J. Amer. Statist. Assoc.* **103**, 1631–1640 (2008), [MR2504209](#).
- [69] Y. Xia, H. Tong, W. K. Li, and L.-X. Zhu, *An adaptive estimation of dimension reduction space*, *J. R. Stat. Soc. Ser. B Stat. Methodol.* **64**, 363–410 (2002), [MR1924297](#).
- [70] G. Yehudai and S. Ohad, “Learning a single neuron with gradient methods”, [Proceedings of thirty third conference on learning theory](#), Vol. 125, edited by J. Abernethy and S. Agarwal, Proceedings of Machine Learning Research (2020), pp. 3756–3786.
- [71] L. Zdeborová and F. Krzakala, *Statistical physics of inference: thresholds and algorithms*, *Advances in Physics* **65**, 453–552 (2016).