One-Bit Clustering for Two Component Sub-Gaussian Mixture Models

Junren Chen
Department of Statistics
Columbia University
New York, NY 10027
jc6315@columbia.edu

Yun Yang
Department of Mathematics
University of Maryland
College Park, MD 20740
yy84@umd.edu

Abstract

Clustering is a fundamental problem in statistics and machine learning. We propose the first one-bit clustering method for two-component sub-Gaussian mixture models. The method uses only one bit per entry of each sample obtained via a dithered quantizer. Under a mild non-spikiness condition on the cluster centers, we show that a variant of Lloyd's algorithm achieves a misclassification rate that decays exponentially with a signal-to-noise ratio comparable to that in the unquantized setting. This result further implies exact recovery under an explicit separation condition, which exceeds the optimal threshold for unquantized data by only a logarithmic factor. When the dimension p is sufficiently large, the non-spikiness condition can be enforced by applying a random rotation using a Haar distributed matrix prior to quantization. In particular, it holds with high probability when $p \gtrsim 1$ for partial recovery and $p \gtrsim \log n \log\log n$ for exact recovery, where n is the sample size. We also establish a minimax lower bound, showing that the misclassification rate and separation condition are in general optimal with sharp constants. Numerical results are provided to corroborate the theory and demonstrate the efficacy of the proposed method.

1 Introduction

Clustering is arguably a fundamental problem in statistics, signal processing, and machine learning. Suppose that the samples $\{X_i\}_{i=1}^n \in \mathbb{R}^p$ come from a few (unknown) centers $\{\theta_i\}_{i=1}^K \subset \mathbb{R}^p$. The goal of clustering is to learn the labels and group the samples from the same center together. The classical setting requires access to the samples $X_i \in \mathbb{R}^p$, which may be unrealistic in some modern applications such as distributed learning where the communication cost can be prohibitive and low-bit data may be used instead. This consideration raises a fundamental question:

Is accurate or even exact clustering possible under one-bit quantization?

Specifically, in this paper we ask whether accurate clustering is possible using one bit per entry. While similar questions have been studied in compressed sensing [19, 26, 13], matrix completion [10, 6], mean estimation [20, 4], covariance estimation [12, 5], and phase retrieval [14, 7], among others, to the best of our knowledge, this question remains unexplored in the context of clustering.

Our paper provides an affirmative answer to this question in the setting of a two-component symmetric sub-Gaussian mixture model centered at $\pm\theta$ for some unknown $\theta \in \mathbb{R}^p$:

$$X_i = \eta_i\theta + \varepsilon_i, \quad i = 1, 2, \cdots, n, \tag{1}$$

where $\eta = (\eta_1, \cdots, \eta_n)^T \in \{-1, 1\}^n$ are the labels to be estimated, and ε_i are sub-Gaussian noise vectors (see Assumption 1). We develop an estimator $\hat{\eta}$ that uses only one bit per entry from $\{X_i\}_{i=1}^n$

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

and achieves low, or even zero, misclassification rate

$$\ell(\hat{\eta}, \eta) = \frac{1}{n} \min \Big\{ \sum_{i=1}^{n} \mathbf{1}(\hat{\eta}_i \neq \eta_i),\ \sum_{i=1}^{n} \mathbf{1}(\hat{\eta}_i \neq -\eta_i) \Big\},$$

referred to as *partial recovery* and *exact recovery* [17, 15, 8, 23], respectively.

Our quantization scheme uses uniform dithering, that is, we add uniform "noise" distributed as $\mathrm{Unif}[-\lambda, \lambda]$ to the samples before applying the one-bit sign quantizer. We refer to λ as the dithering level. The benefits of dithering were observed early [27] and have recently received significant attention in the statistical estimation literature [12, 6, 32]. We will further elucidate the intuition and trade-offs underlying this quantization scheme. We also argue that direct quantization without dithering does not retain sufficient information for clustering in sub-Gaussian mixtures (cf. Remark 3).

The work most relevant to our paper is [23], which establishes the minimax-optimal misclassification rate for two symmetric Gaussian mixtures, that is, model (1) with $\varepsilon_i \sim N(0, \sigma^2 I_p)$:

$$\exp\Big(-\frac{1+o(1)}{2} r_{n,\sigma}^2\Big) \leq \inf_{\hat{\eta}} \sup_{\eta \in \{-1,1\}^n} \mathbb{E}[\ell(\hat{\eta}, \eta)] \leq \exp\Big(-\frac{1-o(1)}{2} r_{n,\sigma}^2\Big) \tag{2}$$

under $r_{n,\sigma} \gtrsim 1$, where $r_{n,\sigma} := \frac{\|\theta\|_2^2/\sigma^2}{\sqrt{\|\theta\|_2^2/\sigma^2 + p/n}}$ is the signal-to-noise ratio (SNR). Note that the sharp constant in the exponent is characterized, and the upper bound is achieved by a variant of Lloyd's algorithm [25, 22]. The bound (2) further implies exact recovery under the sharp separation condition

$$\|\theta\|_2^2 \geq (1+\epsilon)\sigma^2 \bigg(1 + \sqrt{1 + \frac{2p}{n \log n}}\bigg) \log n. \tag{3}$$

Some of these sharp results have been extended to mixtures with more than two components using semidefinite programming [8] and spectral clustering [21, 1].

We establish similar results under dithered one-bit quantization. In a nutshell, our estimator $\hat{\eta}_{1b}$, computed via a variant of Lloyd's algorithm with a small number of iterations, uses one bit per entry of $\{X_i\}_{i=1}^n$ and attains the misclassification rate $\ell(\hat{\eta}_{1b}, \eta) \leq \exp\Big(-\frac{(1-o(1))r_{n,\lambda}^2}{2}\Big)$ under $r_{n,\lambda} \gtrsim 1$, where $r_{n,\lambda} := \frac{\|\theta\|_2^2/\lambda^2}{\sqrt{\|\theta\|_2^2/\lambda^2 + p/n}}$ is our new signal-to-noise ratio with σ in $r_{n,\sigma}$ replaced by the dithering level λ. Under $\varepsilon_i \sim N(0, \sigma^2)$, λ can be chosen on the order of $\sigma\sqrt{\log(np)}$, only a logarithmic factor larger than σ. Hence, our rate is only slightly worse than (2), indicating that the quantization incurs little information loss. The misclassification bound immediately implies exact recovery under an appropriate separation condition, in contrast to most existing one-bit estimation results that only guarantee approximate recovery (cf. Remark 5). We also establish a lower bound showing that our misclassification rate exhibits a sharp constant in general. While the developments appear parallel to [23], substantially different and new techniques are required to establish these results. Moreover, there is an additional subtle interaction among the quantization, dimension, and the spikiness of the center (cf. Remarks 2, 4), which appears to be a novel phenomenon in the area of one-bit learning and motivates a Haar matrix preprocessing step.

Notation. Let $[m] = \{1, 2, \cdots, m\}$ for an integer m. Let $\mathrm{sign}(a) = 1$ for $a \geq 0$ and $\mathrm{sign}(a) = -1$ otherwise; this is applied elementwise to vectors. Let $\|v\|_\infty = \max_i |v_i|$ and $\|v\|_q = (\sum_i |v_i|^q)^{1/q}$ denote the max norm and ℓ_q norm of a vector $v \in \mathbb{R}^n$, respectively, and let $\|M\|_{op}$ denote the operator norm of a matrix M. We use $C, c, C_0, C_1, C', \cdots$ to denote universal constants whose values may vary from line to line. We write $T_1 = O(T_2)$ (or $T_1 \lesssim T_2$) to denote $T_1 \leq CT_2$, and write $T_1 = \Omega(T_2)$ (or $T_1 \gtrsim T_2$) to denote $T_1 \geq CT_2$. We also write $T_1 = \Theta(T_2)$ (or $T_1 \asymp T_2$) if both $T_1 = O(T_2)$ and $T_1 = \Omega(T_2)$ hold. The sub-Gaussian norm of a random variable X is defined as $\|X\|_{\psi_2} = \inf\{K > 0 : \mathbb{E}[\exp(X^2/K^2)] \leq 2\}$. More notation will be introduced as needed.

Overview. In Section 2, we first introduce our quantization procedure and clustering algorithm, and then present our one-bit clustering recovery guarantees, along with an overview of the technical proofs. In Section 3, we establish a minimax lower bound to demonstrate the sharpness of our results. Section 4 provides numerical examples to corroborate our theory, and Section 5 concludes the paper with several remarks. The complete proofs and most technical lemmas are deferred to the appendix.

Algorithm 1 One-bit Lloyd's Algorithm for 2-Component Symmetric Sub-Gaussian Mixtures

Require: Quantized samples $\dot{Y} = [\dot{X}_1, \cdots, \dot{X}_n]$, iteration number T_0

1: **Step 1: Initialization**

2: Let $\hat{v}$ be the leading eigenvector of $H(\dot{Y}^T\dot{Y})$ and let $\hat{\eta}^0 = \mathrm{sign}(\hat{v})$

3: **Step 2: Refinement**

4: **for** $k = 0, 1, 2, \ldots, T_0 - 1$ **do**

$$\hat{\eta}^{k+1} = \mathrm{sign}(H(\dot{Y}^T\dot{Y})\hat{\eta}^k) \tag{5}$$

5: **end for**

Output: $\hat{\eta}^{T_0}$

2 Main Results

Throughout the paper, we adopt the sub-Gaussian mixture model (1) and impose the following assumption.

Assumption 1. *$\{\varepsilon_i\}_{i=1}^n$ are i.i.d. noise vectors with independent, symmetric, zero-mean entries satisfying the σ-sub-Gaussian condition $\mathbb{E}[\exp(t\varepsilon_{ij})] \le \exp\big(\frac{t^2\sigma^2}{2}\big)$ for all $t \ge 0$ and some $\sigma > 0$.*

We collect the unquantized samples in $Y := [X_1, \cdots, X_n]$. Independent of $\{\varepsilon_i\}_{i=1}^n$, we draw dithers $\{\tau_i\}_{i=1}^n$ with entries i.i.d. uniformly distributed over $[-\lambda, \lambda]$, that is, $\tau_i \sim \mathrm{Unif}[-\lambda, \lambda]^p$, and quantize X_i to $\dot{X}_i = \mathrm{sign}(X_i + \tau_i) \in \{-1, 1\}^p$. This dithered quantization scheme has been used to address mean and covariance estimation under coarse quantization [2, 12, 5]. The intuition is that the expectation

$$\mathbb{E}_{\tau\sim\mathrm{Unif}[-\lambda,\lambda]}[\lambda\,\mathrm{sign}(a+\tau)] = T_\lambda(a) := \begin{cases} a\,, & \text{if } |a| \le \lambda, \\ \lambda\,\mathrm{sign}(a)\,, & \text{if } |a| > \lambda, \end{cases} \tag{4}$$

is a truncation of the input a (cf. Lemma 8). Therefore, for sufficiently large λ, $\lambda\dot{X}_i$ can serve as a good surrogate for X_i and retain sufficient information for clustering. We collect all the one-bit samples in a matrix $\dot{Y} = [\dot{X}_1, \cdots, \dot{X}_n]$.

To motivate our algorithm, we note that the minimax optimal procedure in [23] consists of two steps: (i) compute the top eigenvector $\hat{v}$ of $H(Y^TY)$ and set $\eta^0 = \mathrm{sign}(\hat{v})$, where the hollowing operator $H : \mathbb{R}^{n\times n} \to \mathbb{R}^{n\times n}$ removes the diagonal of a square matrix, that is, $H(M) = M - \mathrm{diag}(M)$; (ii) generate a sequence of estimates via a projected power iteration $\eta^{k+1} = \mathrm{sign}(H(Y^TY)\eta^k)$ for $k \ge 0$.[1] Our observation is that the hollowed Gram matrix $H(Y^TY)$ serves as a sufficient statistic for both steps. Combining this with the intuition that $\lambda\dot{X}_i$ acts as a surrogate for X_i, we propose to use $H(\lambda^2\dot{Y}^T\dot{Y})$ to replace the unavailable $H(Y^TY)$. Since λ^2 is absorbed into the retraction $\mathrm{sign}(\cdot)$, we arrive at Algorithm 1.

Remark 1. *While Algorithm 1 does not require the knowledge of λ, it is a key parameter that affects performance. To reduce bias, λ should be large enough to dominate most entries of $\{X_i\}_{i=1}^n$ (cf. Equation* (4)*). However, a larger λ also induces higher variance in $\lambda\dot{Y}$ and thus slower concentration. From another perspective, as $\lambda \to \infty$, $\dot{Y}$ converges to the non-informative* $\mathrm{Unif}\{-1, 1\}$ *in distribution. Therefore, choosing λ involves a bias–variance trade-off.*

2.1 Recovery Guarantees

Theorem 1 below shows that Algorithm 1 attains a misclassification rate that decays exponentially with the signal-to-noise ratio $r_{n,\lambda} := \frac{\|\theta\|_2^2/\lambda^2}{\sqrt{\|\theta\|_2^2/\lambda^2 + p/n}}$. An interesting feature is the appearance of the *spikiness* of the center, defined by $\mu(\theta) = \frac{\|\theta\|_\infty}{\|\theta\|_2}$. This quantity has previously appeared in the matrix completion literature [24, 10, 11] in connection with incoherence conditions.

[1] The sign here is a retraction that maps an estimate back to the parameter space $\{-1, 1\}^n$.

Theorem 1 (Partial recovery). *For every small $\nu > 0$, we assume $\mu(\theta) = \frac{\|\theta\|_\infty}{\|\theta\|_2} \lesssim \frac{1}{1+(p/n)^{1/4}}$. If*

$$\lambda \geq \|\theta\|_\infty + \sigma\sqrt{2(1+\nu)\log(np)} \tag{6}$$

and $1 \lesssim r_{n,\lambda} \lesssim \min\{(np)^{2\nu/5}, n^{2/5}\}$, then for $t \geq \lceil 3\log_4 n \rceil$, we have

$$\mathbb{P}\bigg(\ell(\hat{\eta}^t, \eta) \leq \exp\Big(-\frac{(1-\epsilon_{n,\lambda})r_{n,\lambda}^2}{2}\Big)\bigg) \geq 1 - r_{n,\lambda}^{-1/2} - 2(np)^{-\nu}\,, \tag{7}$$

for some small enough $\epsilon_{n,\lambda} = O\bigg(\frac{1}{r_{n,\lambda}} + \Big(\frac{\log n}{n}\Big)^{1/4}\bigg)$.

Our result resembles the sharp misclassification rate in [23], with the noise level σ in (2) replaced by λ. In addition, we require $\mu(\theta) \lesssim \frac{1}{1+(p/n)^{1/4}}$ and $r_{n,\lambda} \lesssim \min\{(np)^{2\nu/5}, n^{2/5}\}$. The latter appears to be a removable technical condition and is quite mild, since $r_{n,\lambda} \asymp \sqrt{\log n}$ already leads to exact recovery (see Corollary 1 below). The following remark further discusses the spikiness condition.

Remark 2. *To see why the spikiness bound is necessary, observe that*

$$r_{n,\lambda} \gtrsim 1 \text{ is equivalent to } \lambda \lesssim \frac{\|\theta\|_2}{1+(p/n)^{1/4}}\,, \tag{8}$$

thus $\mu(\theta) \lesssim \frac{1}{1+(p/n)^{1/4}}$ is required to ensure the existence of λ satisfying (6)*. In light of $\mu(\theta) \geq \frac{1}{\sqrt{p}}$, a necessary (but not sufficient) condition for our result to hold is $p \gtrsim 1$. We develop a Haar matrix preprocessing step in Section 2.2 to ensure that $\mu(\theta) \lesssim \frac{1}{1+(p/n)^{1/4}}$ holds after a random rotation of the data when $p \gtrsim 1$.*

The next remark shows that dithering is necessary to achieve Theorem 1.

Remark 3. *We show that direct quantization without dithering does not yield guarantees comparable to those in existing works [23, 15, 22, 1]. Consider $\theta = \frac{\Delta}{\sqrt{p}}\mathbf{1}_p$ (so that $\|\theta\|_2 = \Delta$) and let ε_i have i.i.d. Rademacher entries (i.e., $\mathbb{P}(\varepsilon_{ij} = 1) = \mathbb{P}(\varepsilon_{ij} = -1) = \frac{1}{2}$, which is encompassed by our setting with $\sigma \asymp 1$). Then, when $\Delta \leq \sqrt{p}/2$, we have $\dot{X}_i = \mathrm{sign}(X_i) = \mathrm{sign}(\eta_i\theta + \varepsilon_i) = \mathrm{sign}(\varepsilon_i) = \varepsilon_i$, which are pure noise and contain no information. Thus, in this example, accurate clustering can only occur when $\|\theta\|_2 = \Delta \geq \sqrt{p}/2$. In contrast, Theorem 1 only requires $\|\theta\|_2 \gtrsim (1+(p/n)^{1/4})\lambda$ which is substantially weaker in high dimensions. We will provide numerical result to confirm the importance of dithering in this setting (see Figure 2b).*

The proof of Theorem 1 builds on [23] but requires additional work to handle the quantization and introduces several new ideas to obtain a sharp constant in the exponent. The main difficulty is that the quantization breaks the rotational invariance of the Gaussian noise exploited in [23]. In Section 2.3, we provide an overview of the proof with an emphasis on these additional technical aspects. The complete proof is given in Appendix A.

Since $\ell(\hat{\eta}, \eta) \in \{\frac{i}{n} : i \in [n]\}$, $\ell(\hat{\eta}, \eta) < \frac{1}{n}$ implies $\hat{\eta} = \pm\eta$, namely exact recovery of the labels. Therefore, the partial recovery rate in Theorem 1 yields the following result. Its proof is given in Appendix B.

Corollary 1 (Exact recovery). *In the setting of Theorem 1 with a stronger spikiness condition $\mu^2(\theta) \leq \frac{0.99}{\log n + \sqrt{\log^2 n + \frac{2p\log n}{n}}}$, assume that the separation condition*

$$\|\theta\|_2^2 \geq (1+\epsilon)\lambda^2\bigg(1 + \sqrt{1 + \frac{2p}{n\log n}}\bigg)\log n \tag{9}$$

holds for some $\epsilon = O(\frac{1}{\sqrt{\log n}})$, then for any $t \geq \lceil 3\log_4 n \rceil$, $\hat{\eta}^t = \pm\eta$ holds with probability at least $1 - O(\frac{1}{\sqrt{\log n}}) - 2(np)^{-\nu}$.

In two-component Gaussian mixture models without quantization, it was established that Lloyd's algorithm [23] and semidefinite programming [8] attain exact recovery under the separation condition (3). The separation condition in Corollary 1 takes a similar form, with σ replaced by a slightly larger λ (see (6)).

Remark 4. *We impose the spikiness bound* $\mu^2(\theta) \le \frac{0.99}{\log n + \sqrt{\log^2 n + \frac{2p\log n}{n}}}$ *to ensure that* (6) *and* (9) *hold simultaneously under* $\lambda = \|\theta\|_\infty + \sigma\sqrt{2(1+\nu)\log(np)}$ *and* $\|\theta\|_2^2 \gtrsim \sigma^2 \log(np)(\log n + \sqrt{\frac{p\log n}{n}})$. *Since* $\mu^2(\theta) \ge \frac{1}{p}$, *the condition* $p \gtrsim \log n$ *is necessary, though not sufficient, for the spikiness condition to hold. We will see that an additional Haar matrix preprocessing step can help bypass the spikiness condition whenever* $p \gtrsim \log n \log\log n$.

Remark 5. *Estimation under one-bit quantization has been an active research area, but existing results typically provide only approximate recovery guarantees [19, 10, 12, 4]. In contrast, exact recovery is achievable in clustering under one-bit quantization, due to the fact that the desired label vector* η *takes discrete values in* $\{-1,1\}^n$. *Technically, while existing works control the impact of quantization on responses [13, 29] and covariance matrices [13, 5], our analysis examines the interplay between the quantization and the hollowed Gram matrix.*

2.2 Recovery Guarantees under Haar Matrix Preprocessing

A limitation of Theorem 1 and Corollary 1 is the spikiness condition on θ, which holds only when $p \gtrsim 1$ for partial recovery and $p \gtrsim \log n$ for exact recovery. The issue is that, in clustering, the center θ is unknown (indeed, estimating θ is itself an important problem [31]), so one cannot verify the spikiness condition before applying our one-bit clustering method. Note that even when p is sufficiently large, $\mu(\theta)$ can take any value in $[\frac{1}{\sqrt{p}}, 1]$.

We address this issue via a simple Haar matrix preprocessing step prior to one-bit quantization: we draw a Haar matrix $R \sim \mathrm{Haar}(\mathbb{O}(p))$[2] and transform the original samples $\{X_i\}_{i=1}^n$ to $\{RX_i\}_{i=1}^n$. The subsequent quantization and algorithm remain unchanged. In practice, the generation and communication of R can be controlled by a random seed and incur negligible cost (e.g., [28]). Moreover, to further reduce the storage and computational cost, the dense Haar matrix here can be further replaced by a suitable structured orthogonal matrix without changing the subsequent recovery guarantees.[3]

The Haar matrix resolves this issue by *whitening the center*: in light of

$$RX_i = \eta_i(R\theta) + R\varepsilon_i\,,$$

the original deterministic center θ is transformed into $R\theta$, which is uniformly distributed over $\|\theta\|_2 S^{p-1}$. A standard argument shows that $\mu(R\theta) \le \sqrt{3\log(p)/p}$ holds with high probability (cf. Lemma 17), and therefore the spikiness conditions in Theorem 1 and Corollary 1 are automatically satisfied in sufficiently high dimension:

- Partial recovery is attained when $p \gtrsim 1$, since this guarantees $\sqrt{\frac{\log p}{p}} \lesssim \frac{1}{1+(p/n)^{1/4}}$ (see Theorem 2);
- Exact recovery is attained when $p \gtrsim \log n \log\log n$, since this guarantees $\sqrt{\frac{\log p}{p}} \lesssim \frac{1}{\sqrt{\log n + \frac{p\log n}{n}}}$ (see Corollary 2).

The following theorem concerns the misclassification rate under the Haar matrix preprocessing.

Theorem 2 (Partial recovery without spikiness condition)**.** *Suppose that* $R \sim \mathrm{Haar}(\mathbb{O}(p))$ *is independent of everything else, and in our model we quantize* X_i *to* $\dot{X}_i = \mathrm{sign}(RX_i + \tau_i)$ *with* $\tau_i \sim \mathrm{Unif}[-\lambda,\lambda]^p$, *then we compute* $\{\hat{\eta}^k\}_{k\ge 0}$ *by Algorithm 1. Given small* $\nu > 0$, *assume* $\min\{n, p, r_{n,\lambda}\} \gtrsim 1$ *and* $r_{n,\lambda} \lesssim \min\{(np)^{2\nu/5}, n^{2/5}\}$, *where recall* $r_{n,\lambda} = \frac{\|\theta\|_2^2/\lambda^2}{\sqrt{\|\theta\|_2^2/\lambda^2 + p/n}}$. *If*

$$\lambda \ge \sqrt{\frac{3\log p}{p}}\|\theta\|_2 + \sigma\sqrt{2(1+\nu)\log(np)}\,, \tag{10}$$

[2]This means that R is uniformly distributed over the group of $p \times p$ orthonormal matrices.

[3]An example is $R = HD$ where H is a normalized Hadamard matrix and D is a diagonal matrix with independent random signs. Indeed, the goal of the Haar matrix is to flatten the center as in Lemma 17, and it is easy to see that $R = HD$ achieves the same goal.

then for any $t \geq \lceil 3\log_4 n \rceil$ we have

$$\mathbb{P}\bigg(\ell(\hat{\eta}^t, \eta) \leq \exp\bigg(-\frac{(1-\epsilon_{n,\lambda})r_{n,\lambda}^2}{2}\bigg)\bigg) \geq 1 - r_{n,\lambda}^{-1/2} - 2(np)^{-\nu} - 3p^{-1/4}$$

for some $\epsilon_{n,\lambda} = O(r_{n,\lambda}^{-1} + (\frac{\log n}{n})^{1/4} + \log^{-1/4}(np))$.

Overall, the result is proved by revisiting the arguments for Theorem 1. One aspect that requires adaptation is that, when extending to general sub-Gaussian noise, $R\varepsilon_i$ may have correlated entries. We address this issue using a noise decomposition and conditioning argument (cf. Remark 10 in the appendix). The complete proof appears in Appendix C.

Enforcing $\ell(\hat{\eta}^t, \eta) < \frac{1}{n}$ yields the following exact recovery guarantee, where we choose the minimal λ in (10) for simplicity. See Appendix D for the proof.

Corollary 2 (Exact recovery without spikiness condition)**.** *In the setting of Theorem 2 with*

$$\lambda = \sqrt{\frac{3\log p}{p}}\|\theta\|_2 + \sigma\sqrt{2(1+\nu)\log(np)}, \tag{11}$$

if for some $\epsilon = O(\frac{1}{\log^{1/4} n})$ the separation condition

$$\|\theta\|_2^2 \geq (1+\epsilon)\bigg(\sqrt{\frac{3\log p}{p}}\|\theta\|_2 + \sigma\sqrt{2(1+\nu)\log(np)}\bigg)^2\bigg(1+\sqrt{1+\frac{2p}{n\log n}}\bigg)\log n \tag{12}$$

holds, then for any $t \geq \lceil 3\log_4 n \rceil$, $\hat{\eta}^t = \pm\eta$ with probability at least $1 - O(\frac{1}{\log^{1/4} n} + (np)^{-\nu} + p^{-1/4})$. More specifically, if $\min\{n, p, r_{n,\lambda}\} \geq C_1(\nu)$ and $p \geq C_2(\nu)\cdot \log n \log\log n$ for sufficiently large constants $C_1(\nu), C_2(\nu)$ depending only on ν, then the separation condition in (12) *can be ensured by the more explicit condition*

$$\|\theta\|_2^2 \geq 2(1+2\nu)\sigma^2\bigg(1+\sqrt{1+\frac{2p}{n\log n}}\bigg)\log(np)\log n. \tag{13}$$

The above result states that exact recovery is achieved under (13) provided that $p \gtrsim_\nu \log n \log\log n$. The separation condition (13) is explicit and differs from the sharp separation condition (3) for two-component Gaussian mixture models without quantization by only an additional factor of $2\log(np)$.

Remark 6. *The tuning of λ appears to require knowledge of $(\|\theta\|_2, \sigma)$ in view of* (11)*, but we note that in many regimes it suffices to have an estimate on the noise level σ. In particular, if we assume $\|\theta\|_2 \leq c\sigma\sqrt{\frac{p\log(np)}{\log p}}$ for sufficiently small c,*[4] *then the term $\sqrt{\frac{3\log p}{p}}\|\theta\|_2$ is negligible compared to $\sigma\sqrt{2\log(np)}$, and hence the theory suggests $\lambda \approx \sigma\sqrt{2\log(np)}$. In practice, we recommend $\lambda = s\sigma\sqrt{2\log(np)}$ for some shrinkage parameter $s \in (0,1)$; see, e.g., the second experiment in Section 4. We leave a more thorough investigation of the tuning of λ for future work.*

Remark 7. *To extend our method to $p \lesssim 1$, a naive approach is to append additional zeros to the samples. That said, numerical results suggest that our method already performs well and achieves exact recovery in a dimension as low as $p = 5$; see Figure 1a in Section 4 for instance.*

2.3 Technical Overview (Theorem 1)

(A) Handing the quantization. Since in our one-bit setting $\lambda\dot{X}_i$ serves as a surrogate for X_i, a useful perspective is to view it as a clustering problem for $\{\lambda\dot{X}_i\}_{i=1}^n$, with the overall noise on the i-th sample being

$$u_i = \lambda\dot{X}_i - \eta_i\theta = \lambda\,\mathrm{sign}(X_i + \tau_i) - \eta_i\theta = \lambda\,\mathrm{sign}(\eta_i\theta + \varepsilon_i + \tau_i) - \eta_i\theta.$$

One difficulty here is that u_i is not mean-zero (a property that existing analyses [22, 23, 17] heavily rely on). We bypass this hurdle via a proof scheme [5] that draws connection between the one-bit quantizer and the uniform quantizer, consisting of three steps:

[4] This is mild because under $p \gtrsim \log n \log\log n$, the upper bound $\sigma\sqrt{\frac{p\log(np)}{\log p}}$ is much larger than $\|\theta\|_2 \gtrsim \sigma(\log^2(np)\log^2 n + \frac{p\log^2(np)\log n}{n})^{1/4}$ required in (13).

(i) Define the uniform quantizer with resolution $\delta > 0$ as $Q_\delta(a) = \delta(\lfloor \frac{a}{\delta} \rfloor + \frac{1}{2})$, and note that the (rescaled) one-bit quantizer and $Q_{2\lambda}$ are connected by (cf. Figure 3 in the appendix)

$$\begin{aligned} & Q_{2\lambda}(a) = \lambda \operatorname{sign}(a), \quad \forall |a| \le 2\lambda \\ \Longrightarrow\ & Q_{2\lambda}(a+\tau) = \lambda \operatorname{sign}(a+\tau) \text{ under } \tau \sim \mathrm{Unif}[-\lambda, \lambda], \quad \forall |a| \le \lambda \\ \Longrightarrow\ & \lambda \dot{X}_i = \underbrace{Q_{2\lambda}(X_i + \tau_i)}_{\tilde{X}_i}, \quad \forall i \in [n], \quad \text{if } \|Y\|_{\max} = \max_{i \in [n]} \|X_i\|_\infty \le \lambda. \end{aligned} \tag{14}$$

(ii) We then analyze the clustering of $\{\tilde{X}_i := Q_{2\lambda}(X_i + \tau_i)\}_{i=1}^n$ with overall noise $\tilde{u}_i = \tilde{X}_i - \eta_i \theta$ having independent, *zero-mean*, sub-Gaussian entries (cf. Lemma 12);

(iii) By the sub-Gaussian tail bound $\lambda \ge \|\theta\|_\infty + L\sqrt{2(1+\nu)\log(np)}$ ensures that $\|Y\|_\infty \le \lambda$ holds with high probability (w.h.p.), which together with (14) allows us to carry over the conclusion for $\{\tilde{X}_i\}_{i=1}^n$ to $\{\lambda \dot{X}_i\}_{i=1}^n$.

(B) Extending the argument in [23]. We focus on Step (ii). We analyze the performance of Algorithm 1 with $\tilde{Y} = [\tilde{X}_1, \cdots, \tilde{X}_n]$ rather than $\dot{Y}$. Note that the zero-mean $\tilde{u}_i = \tilde{X}_i - \eta_i \theta$ allows us to closely follow the analysis in [23] which (a) first shows a crude bound on the iterates $\ell(\hat{\eta}^k, \eta) = O(n^{-2} + r_{n,\lambda}^{-2})$ and (b) then tightens it to

$$\ell(\hat{\eta}^k, \eta) \lesssim \frac{1}{n^{3/2}} + r_{n,\lambda}^{5/2} P^* \ \text{ w.h.p.}, \quad (\forall k \ge \lceil 3 \log_4 n \rceil),$$

where, with $\delta_j = Q_{2\lambda}(\theta + \varepsilon_j + \tau_j) - \theta$ for $j \in [n]$, P^* is defined as

$$P^* := \mathbb{P}\bigg(\Big\langle \theta + \delta_1, \theta + \frac{1}{n-1} \sum_{j=2}^n \delta_j \Big\rangle < \frac{C\|\theta\|_2^2}{r_{n,\lambda}} \bigg). \tag{15}$$

This part of the argument is presented in Step 2 in Appendix A and only involves some straightforward extensions of the concentration bounds in [23] (cf. Lemmas 13–14).

(C) Establishing sharp bound on P^*. It remains to establish a sharp bound on P^* (accomplished in Step 4 of Appendix A). While the corresponding analysis in [23, Thm. 5] deals with standard Gaussian vector and hence rotational invariance readily yields sharp tail bound for the marginals, the δ_j's in Equation (15) are high-dimensional *non-Gaussian* random vectors and in turn the analysis here is much more entangled. We begin with

$$\Big\langle \theta + \delta_1, \theta + \frac{1}{n-1} \sum_{j=2}^n \delta_j \Big\rangle = \|\theta\|_2^2 + \Big\langle \theta, \frac{1}{n-1} \sum_{j=2}^n \delta_j \Big\rangle + \Big\langle \delta_1, \theta + \frac{1}{n-1} \sum_{j=2}^n \delta_j \Big\rangle. \tag{16}$$

In view of $\|\theta\|_2^2$, the term $\frac{C\|\theta\|_2^2}{r_{n,\lambda}}$ in (15) has minimal impact due to $r_{n,\lambda} \gtrsim 1$. By $\|\delta_j\|_{\psi_2} \lesssim \lambda$, the standard sub-Gaussian tail bound shows that $\langle \theta, \frac{1}{n-1} \sum_{j=2}^n \delta_j \rangle$ is negligible. The main bulk of techniques lies in treating $\langle \delta_1, \theta + \frac{1}{n-1} \sum_{j=2}^n \delta_j \rangle$:

(C.1) Establishing sharp bound on $\|\theta + \frac{1}{n-1} \sum_{j=2}^n \delta_j\|_2$. We start with $\big\|\theta + \frac{1}{n-1} \sum_{j=2}^n \delta_j\big\|_2^2 = \|\theta\|_2^2 + \big\|\frac{1}{n-1} \sum_{j=2}^n \delta_j\big\|_2^2 + 2\big\langle \theta, \frac{1}{n-1} \sum_{j=2}^n \delta_j \big\rangle$. A sub-Gaussian tail bound renders $\langle \theta, \frac{1}{n-1} \sum_{j=2}^n \delta_j \rangle$ negligible. To bound $\|\frac{1}{n-1} \sum_{j=2}^n \delta_j\|_2^2$, we first use a Bernstein's inequality to show that it sharply concentrates about $\mathbb{E}\|\frac{1}{n-1} \sum_{j=2}^n \delta_j\|_2^2 = \frac{1}{n-1}\mathbb{E}\|\delta_2\|_2^2$, and then leverage $Q_{2\lambda}(\theta + \varepsilon_i + \tau_i) \approx \lambda \operatorname{sign}(\theta + \varepsilon_i + \tau_i)$ to compute the expectation with sharp constant; see Lemma 15. We then reach $\|\frac{1}{n-1} \sum_{j=2}^n \delta_j\|_2^2 \le (1 + o(1))(\|\theta\|_2^2 + \frac{p\lambda^2}{n})$.

(C.2) Passing to the margins of δ_1. The sharp bound on $\|\theta + \frac{1}{n-1} \sum_{j=2}^n \delta_j\|_2$ readily yields $P^* \le P^{**} +$ negligible terms, where

$$P^{**} = \mathbb{P}\bigg(\langle \delta_1, \hat{\xi} \rangle < -(1 - o(1)) \frac{\|\theta\|_2^2}{\sqrt{\|\theta\|_2^2 + p\lambda^2/n}} \bigg), \text{ for } \hat{\xi} := \frac{\theta + (n-1)^{-1} \sum_{j=2}^n \delta_j}{\|\theta + (n-1)^{-1} \sum_{j=2}^n \delta_j\|_2} \in \mathbb{S}^{p-1}.$$

(C.3) Bounding P^{**} *sharply.* We reduce the random vector $\delta_1 = Q_{2\lambda}(\theta + \varepsilon_j + \tau_j) - \theta$ to the bounded $\tilde{\delta}_1 := \lambda \operatorname{sign}(\theta + \varepsilon_1 + \tau_1) - \mathbb{E}[\lambda \operatorname{sign}(\theta + \varepsilon_1 + \tau_1)]$ in Lemma 16 and then invoke Hoeffding's inequality, which happens to yield sharp bound in this regime.

3 Minimax Lower Bound

In this section, we establish a minimax lower bound for clustering based on the one-bit samples $(\dot{X}_i)_{i=1}^n$, showing that our results are sharp in general. Note that existing lower bounds (see [23, 8, 22]) are derived under isotropic Gaussian noise $\varepsilon_i \sim N(0, \sigma^2 I_p)$, and since the quantization breaks rotational invariance, our proof requires several different ideas. In particular, the main lemma is the anti-concentration bound for a binomial variable (cf. Lemma 1).

Theorem 3 (Minimax lower bound)**.** *Given $\|\theta\|_2 = \Delta$, assume that $(\varepsilon_i)_{i=1}^n$ are i.i.d. $N(0, \sigma^2 I_p)$ vectors, $\lambda \gtrsim \frac{\Delta}{\sqrt{p}} + \sigma$, and that a mild scaling condition $\log(\frac{2\lambda\sqrt{p}}{\Delta}) \le \frac{\lambda^2}{8\sigma^2}$ holds. Then let $\hat{\eta} = \hat{\eta}(\{\dot{X}_i\}_{i=1}^n)$ be a measurable function of $\{\dot{X}_i = \mathrm{sign}(X_i + \tau_i)\}_{i=1}^n$ (where $\tau_i \overset{iid}{\sim} \mathrm{Unif}[-\lambda, \lambda]^p$), we have*

$$\inf_{\hat{\eta}} \sup_{\theta, \eta} \mathbb{E}\big[\ell(\hat{\eta}, \eta)\big] \ge \frac{1}{\sqrt{2p}} \exp\Big(- \frac{(1+\epsilon)\Delta^2}{2\lambda^2}\Big)$$

where the constant $\epsilon \to 0$ when $\frac{\lambda}{p^{-1/2}\|\theta\|_2 + \sigma} \to \infty$.

We now discuss in what regimes the above lower bound implies the sharpness of our recovery guarantees in Section 2.

Remark 8. *By absorbing the leading factor $\frac{1}{\sqrt{2p}}$ into the exponent, the lower bound reads $\frac{1}{\sqrt{2p}} \exp\big(-\frac{(1+\epsilon)\|\theta\|_2^2}{2\lambda^2}\big) = \exp\big(-(1+\epsilon+\frac{\log(2p)\lambda^2}{\|\theta\|_2^2})\frac{\|\theta\|_2^2}{2\lambda^2}\big)$ and can be written as $\exp\big(-(1+\epsilon')\frac{\|\theta\|_2^2}{2\lambda^2}\big)$ for small enough ϵ' under $\|\theta\|_2^2 \gtrsim \lambda^2 \log(2p)$. Recall that our misclassification rates in Theorems 1, 2 read $\exp\big(-\frac{1-o(1)}{2} r_{n,\lambda^2}\big) = \exp\big(-\frac{1-o(1)}{2}\frac{(\|\theta\|_2/\lambda)^4}{(\|\theta\|_2/\lambda)^2 + p/n}\big)$. Hence, it reduces to $\exp\big(-\frac{1\pm o(1)}{2}\frac{\|\theta\|_2^2}{\lambda^2}\big)$ under $p \ll \frac{n\|\theta\|_2^2}{\lambda^2}$.*[5] *In conclusion, under the two scaling conditions of*

$$\|\theta\|_2^2 \gtrsim \lambda^2 \log(2p) \quad \textit{and} \quad p \ll \frac{n\|\theta\|_2^2}{\lambda^2}, \tag{17}$$

the lower bound matches the upper bounds in Theorems 1 and 2, indicating that $\exp\big(-\frac{1\pm o(1)}{2}\frac{\|\theta\|_2^2}{\lambda^2}\big)$ is the sharp misclassification rate. Since exact recovery is achieved if and only if $\ell(\hat{\eta}_t, \eta) < \frac{1}{n}$, the lower bound also indicates the sharpness of the separation conditions (1) and (13) in the regime of (17).

Remark 9. *However, we do not have a matching lower bound in some regimes of interest, such as the high-dimensional regime $p \gg \frac{n\|\theta\|_2^2}{\lambda^2}$, where our upper bound reads $\exp\big(-\frac{1-o(1)}{2}\frac{n\|\theta\|_2^4}{p\lambda^4}\big)$. We leave the development of tighter lower and upper bounds in these regimes to future work.*

We provide an overview of the proof here and relegate the complete proof to Appendix E.

Overview of the Proof for Theorem 3. Unlike in [22, 23], under quantization, the Bayes classifier does not have a closed form under general θ. Our remedy is to fix the specific center as $\theta = \frac{\Delta}{\sqrt{p}}\mathbf{1}_p$. We then place an independent Rademacher prior on the label η and leverage a standard argument (e.g., [16]) to get $\inf_{\hat{\eta}} \sup_{\theta,\eta} \mathbb{E}\big[\ell(\hat{\eta}, \eta)\big] \ge c \cdot \inf_{\bar{\eta}_j} \mathbb{E}_\pi \mathbb{E}\,|\bar{\eta}_j(\dot{X}_j) - \eta_j|$, where η_j is a Rademacher variable, and $\bar{\eta}_j$ is a measurable function of $\dot{X}_j$ for estimating η_i, $c > 0$ is a universal constant. Since the entries of θ are equal and the noise and dithering are symmetric, we find that $\eta_j = 1$ renders higher likelihood if and only if the number of 1's in $\dot{X}_j$ is higher than the number of -1's. Thus, the Bayes optimal selector attaining the infimum $\inf_{\bar{\eta}_j} \mathbb{E}_\pi \mathbb{E}\,|\bar{\eta}_j(\dot{X}_j) - \eta_j|$ is given by $\eta_j^* = \mathrm{sign}(\langle \dot{X}_j, \mathbf{1}\rangle)$. Note that it makes an error if and only if *at least half of the entries of $\dot{X}_j$ have different signs from η_j*. Observe that the entries of $\dot{X}_j$ are independent and have sign different from η_j with the same probability $P_- := \mathbb{P}(\frac{\|\theta\|_2}{\sqrt{p}} + \varepsilon_{ij} + \tau_{ij} < 0)$. Therefore, the number of $-\eta_j$ in $\dot{X}_j$ follows $\mathrm{Binomial}(p, P_-)$, and thus by Lemma 1 $\mathbb{E}_\pi \mathbb{E}\,|\bar{\eta}_j(\dot{X}_j) - \eta_j| = \mathbb{E}|\eta_j^* - \eta_j| = 2\mathbb{P}\Big(\mathrm{Binomial}(p, P_-) > \frac{p}{2}\Big) \ge \frac{1}{\sqrt{2p}} \exp\Big(-pD_{\mathrm{KL}}\big(\frac{1}{2}\,\|P_-\big)\Big)$. By further computing P_- and performing some algebra, we obtain the claimed lower bound under the scaling assumptions of the theorem. □

[5] We write this to denote $p \le \frac{cn\|\theta\|_2^2}{\lambda^2}$ for some small enough c.

4 Experiments

We provide experimental results to corroborate our theory and demonstrate the effectiveness of the proposed method. We consider two Gaussian mixtures with θ uniformly distributed over $\Delta\mathbb{S}^{p-1}$ and $\varepsilon_i \sim N(0, I_p)$,[6] otherwise stated. We run 200 iterations in Algorithm 1 and the reported results are averaged over 300 independent trials. All experiments were implemented using Matlab R2022a on a laptop with an Intel CPU up to 2.5 GHz and 32 GB RAM. We defer some details to Appendix G due to page limit.

Phase Transitions. The aim of our first two experiments is to illustrate the separation condition (9) for achieving exact recovery. Note that the minimal λ in (6) is approximately $\sqrt{2\log(np)}$ (cf. Equation (11), Remark 6). Further illustrations of the experimental designs can be found in Appendix G.1.

We start with a low-dimensional setting with $p = 5$, in which (9) is provably sharp (cf. Remark 8). We set the dithering level $\lambda = \sqrt{2\log n}$ in light of $\sqrt{2\log(np)} \approx \sqrt{2\log n}$ under $p = 5$. In this setting, the separation condition (9) approximately reduces to $\Delta \geq 2\log n$. We set $a = \log n$, $b = \Delta$ and test $a = 3.4 : 0.25 : 6.9$, $b = 4 : 0.25 : 16$. We then report the empirical rates of exact recovery in Figure 1a. Consistent with our theory, the results suggest that the phase transitions from partial recovery to exact recovery occur roughly in $b = 2a$.

Fix $n = 100$, we also follow [23] to provide a high-dimensional setting with $p = bn\log n$, $\Delta^2 = \lambda^2(1+\sqrt{a})\log n$, and $\lambda = \sqrt{2\log(np)}$. Hence, our (9) suggests that $a \geq 1 + 2b$ ensures exact recovery (w.h.p.). We test $b = 2.4 : 0.05 : 4.6$ and $a = 1.5 : 0.1 : 13$ and similarly report the empirical exact recovery rates in Figure 1b. While we do not have a matching lower bound in this regime (cf. Remark 9), it appears that $a \geq 1 + 2b$ remains close to the locations of the phase transitions.

One-bit v.s. Classical. Our third experiment confirms the efficacy of the proposed method by comparing the performance of [23] (using $\{X_i\}_{i=1}^n$) and Algorithm 1 (using $\{\dot{X}_i\}_{i=1}^n$). We set $\lambda = 0.3 \cdot \sqrt{2\log(np)}$ for our method with a shrinkage factor of 0.3; see further justification in Appendix G.2. We test $n = 100$, $p = 300$ and $\Delta = 0.1 : 0.1 : 6$ and report the misclassification rates in Figure 1c. Empirically, our one-bit clustering method requires *no more than twice the separation* Δ to achieve the same misclassification rate as in the classical clustering setting.

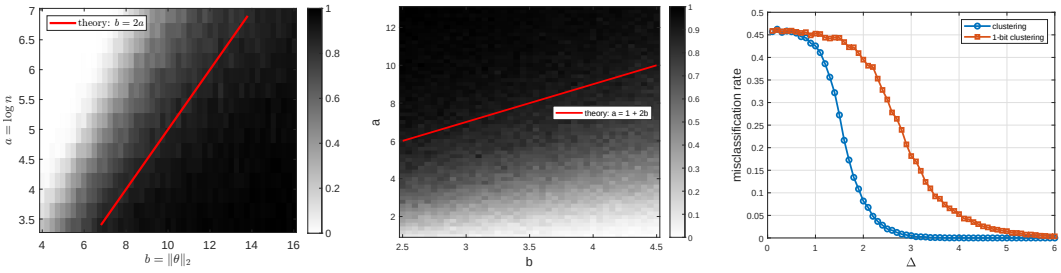

(a) Phase transition ($p = 5$) (b) Phase transition ($p = bn\log n$) (c) 1-bit clustering v.s. clustering

Figure 1: Phase transitions and comparison with classical clustering.

Importance of Haar Preprocessing. We demonstrate the importance of the Haar matrix preprocessing step. To this end, we adopt the one-bit clustering setting in Figure 1c and compare the results under a fixed $\theta = \Delta e_1$ (the spikiest center) with or without the Haar matrix preprocessing step. The results in Figure 2a show that the misclassification rates stay above 0.4 without out Haar preprocessing, thus confirming the importance of this step.

Importance of Random Dithering. To illustrate the essential role of random dithering in our result (as already discussed in Remark 3), we test $(n, p) = (100, 300)$ and the flattest center $\theta = \frac{\Delta}{\sqrt{p}}\mathbf{1}_p$ without further Haar preprocessing step under $\lambda = 0, 0.8, 1.2, 1.6, 2.0$. Note that $\lambda = 0$ and 0.8 correspond to the settings without random dithering and with overly small dithering level, respectively, and the results in Figure 2b show that the misclassification errors in these two settings almost do not

[6] In this case, due to $\theta \sim \mathrm{Unif}(\Delta\mathbb{S}^{p-1})$ we do not need a separate Haar matrix rotation step.

decrease over $\Delta \leq 5$. On the other hand, once the dithering level is above some threshold, the error curves decay as predicted in our theory. As this is the case for $\lambda = 1.2,\ 1.6,\ 2.0$, the results also suggest that our method is not very sensitive to λ.

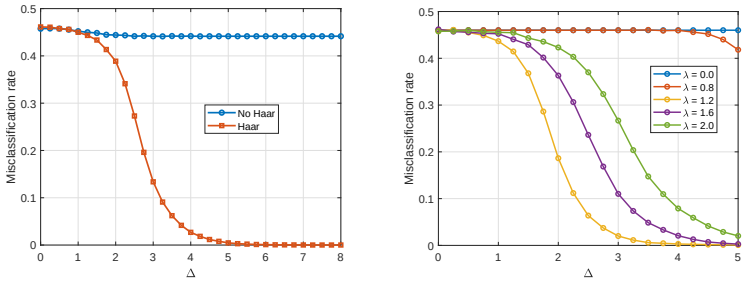

(a) Haar preprocessing (b) Random dithering

Figure 2: Importance of the Haar preprocessing and random dithering.

5 Conclusion

This paper provides the first one-bit clustering method for two symmetric sub-Gaussian mixtures. We adopt a dithered one-bit quantizer and establish partial and exact recovery guarantees that are only slightly worse than those in the classical unquantized setting. A lower bound is also provided to show that the recovery guarantees exhibit sharp constants in general. Compared to previous work on one-bit estimation, several new phenomena arise, including the achievability of exact recovery and a subtle interaction with the spikiness of the center. The latter can be addressed via a Haar matrix rotation step, which enables partial recovery under $p \gtrsim 1$ and exact recovery under $p \gtrsim \log n \log\log n$. There remain many interesting directions for future work, including tuning-free one-bit clustering (cf. Remark 6), sharp bounds in high dimensions (cf. Remark 9), the performance of other clustering algorithms under quantization, and extension to more than two mixtures that are not in symmetric positions.

Acknowledgment

The work was done while J.C. was affiliated with the Department of Mathematics, University of Maryland, supported by a Novikov Postdoctoral Fellowship.

References

[1] Emmanuel Abbe, Jianqing Fan, and Kaizheng Wang. An lp theory of pca and spectral clustering. *The Annals of Statistics*, 50(4):2359–2385, 2022.

[2] Pedro Abdalla and Junren Chen. Robust mean estimation under quantization. *arXiv preprint arXiv:2601.07074*, 2026.

[3] Robert B Ash. *Information theory*. Courier Corporation, 2012.

[4] T Tony Cai and Hongji Wei. Distributed adaptive gaussian mean estimation with unknown variance: Interactive protocol helps adaptation. *The Annals of Statistics*, 50(4):1992–2020, 2022.

[5] Junren Chen and Michael K Ng. A parameter-free two-bit covariance estimator with improved operator norm error rate. *Applied and Computational Harmonic Analysis*, page 101774, 2025.

[6] Junren Chen, Cheng-Long Wang, Michael K. Ng, and Di Wang. High dimensional statistical estimation under uniformly dithered one-bit quantization. *IEEE Transactions on Information Theory*, 69(8):5151–5187, 2023.

[7] Junren Chen and Ming Yuan. One-bit phase retrieval: Optimal rates and efficient algorithms. *IEEE Transactions on Information Theory*, 2026.

[8] Xiaohui Chen and Yun Yang. Cutoff for exact recovery of gaussian mixture models. *IEEE Transactions on Information Theory*, 67(6):4223–4238, 2021.

[9] Yuejie Chi, Yue M Lu, and Yuxin Chen. Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing*, 67(20):5239–5269, 2019.

[10] Mark A Davenport, Yaniv Plan, Ewout Van Den Berg, and Mary Wootters. 1-bit matrix completion. *Information and Inference: A Journal of the IMA*, 3(3):189–223, 2014.

[11] Mark A Davenport and Justin Romberg. An overview of low-rank matrix recovery from incomplete observations. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):608–622, 2016.

[12] Sjoerd Dirksen, Johannes Maly, and Holger Rauhut. Covariance estimation under one-bit quantization. *The Annals of Statistics*, 50(6):3538–3562, 2022.

[13] Sjoerd Dirksen and Shahar Mendelson. Non-gaussian hyperplane tessellations and robust one-bit compressed sensing. *Journal of the European Mathematical Society*, 23(9):2913–2947, 2021.

[14] Dylan Domel-White and Bernhard G Bodmann. Phase retrieval by binary questions: Which complementary subspace is closer? *Constructive Approximation*, 56(1):1–33, 2022.

[15] Yingjie Fei and Yudong Chen. Hidden integrality of sdp relaxations for sub-gaussian mixture models. In *Conference On Learning Theory*, pages 1931–1965. PMLR, 2018.

[16] Chao Gao, Zongming Ma, Anderson Y. Zhang, and Harrison H. Zhou. Community detection in degree-corrected block models. *Annals of Statistics*, 46(5):2153–2185, 2018.

[17] Christophe Giraud and Nicolas Verzelen. Partial recovery bounds for clustering with the relaxed k-means. *Mathematical Statistics and Learning*, 1(3):317–374, 2019.

[18] Robert M Gray and Thomas G Stockham. Dithered quantizers. *IEEE Transactions on Information Theory*, 39(3):805–812, 1993.

[19] Laurent Jacques, Jason N Laska, Petros T Boufounos, and Richard G Baraniuk. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Transactions on Information Theory*, 59(4):2082–2102, 2013.

[20] Alon Kipnis and John C Duchi. Mean estimation from one-bit measurements. *IEEE Transactions on Information Theory*, 68(9):6276–6296, 2022.

[21] Matthias Löffler, Anderson Y Zhang, and Harrison H Zhou. Optimality of spectral clustering in the gaussian mixture model. *The Annals of Statistics*, 49(5):2506–2530, 2021.

[22] Yu Lu and Harrison H Zhou. Statistical and computational guarantees of lloyd's algorithm and its variants. *arXiv preprint arXiv:1612.02099*, 2016.

[23] Mohamed Ndaoud. Sharp optimal recovery in the two component gaussian mixture model. *The Annals of Statistics*, 50(4):2096–2126, 2022.

[24] Sahand Negahban and Martin J Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *Journal of Machine Learning Research*, 13(1):1665–1697, 2012.

[25] Jiming Peng and Yu Wei. Approximating k-means-type clustering via semidefinite programming. *SIAM journal on optimization*, 18(1):186–205, 2007.

[26] Yaniv Plan and Roman Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *IEEE Transactions on Information Theory*, 59(1):482–494, 2012.

[27] Lawrence Roberts. Picture coding using pseudo-random noise. *IRE Transactions on Information Theory*, 8(2):145–154, 1962.

[28] Mayank Shrivastava, Berivan Isik, Qiaobo Li, Sanmi Koyejo, and Arindam Banerjee. Sketching for distributed deep learning: A sharper analysis. *Advances in Neural Information Processing Systems*, 37:6417–6447, 2024.

[29] Christos Thrampoulidis and Ankit Singh Rawat. The generalized lasso for sub-gaussian measurements with dithered quantization. *IEEE Transactions on Information Theory*, 66(4):2487–2500, 2020.

[30] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge University Press, 2018.

[31] Yihong Wu and Harrison H Zhou. Randomly initialized em algorithm for two-component gaussian mixture achieves near optimality in $o(\sqrt{n})$ iterations. *Mathematical Statistics & Learning*, 4, 2021.

[32] Chunlei Xu and Laurent Jacques. Quantized compressive sensing with rip matrices: The benefit of dithering. *Information and Inference: A Journal of the IMA*, 9(3):543–586, 2020.

A Proof of Theorem 1 (Partial recovery)

Proof. As $\lambda\dot{X}_i$ is the surrogate of X_i, we define the overall noise on the sample as

$$u_i = \lambda\dot{X}_i - \eta_i\theta = \lambda\,\mathrm{sign}(\eta_i\theta + \varepsilon_i + \tau_i) - \eta_i\theta$$

and can formulate the problem in matrix form:

$$\lambda\dot{Y} = \theta\eta^T + U$$

where $U = [u_1, \cdots, u_n]$. Therefore, the one-bit clustering problem can be treated as a clustering problem with noise U. While u_i is bounded and therefore sub-Gaussian, the difficulty is that u_i is not zero-mean. We bypass the difficulty by a technical trick developed in [5]: first, we analyze the uniform quantizer that leads to zero-mean noise; second, we show that, when $\lambda \geq \|\theta\|_\infty + L\sqrt{2\log(np)}$, then this uniform quantizer is identical to our one-bit quantizer with high probability.

Step 1: Introducing the Uniform Quantizer

The uniform quantizer is given by

$$Q_{2\lambda}(a) = 2\lambda\Big(\lfloor\frac{a}{2\lambda}\rfloor + \frac{1}{2}\Big),$$

and notice that

$$Q_{2\lambda}(a) = \lambda\,\mathrm{sign}(a)\,, \quad \forall |a| \leq 2\lambda. \tag{18}$$

See the following Figure 3.

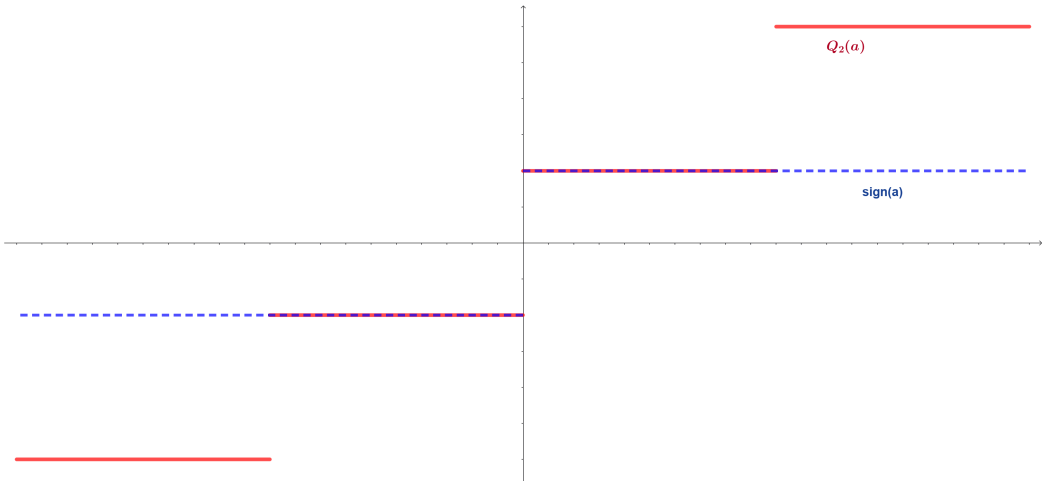

Figure 3: Graphical illustration of Equation (18): the red curve $Q_2(a)$ and the blue curve $\mathrm{sign}(a)$ are identical when $|a| < 2$.

We now proceed the analysis with the quantized samples

$$\tilde{X}_i = Q_{2\lambda}(X_i + \tau_i), \quad i = 1, \cdots, n,$$

which can be arranged in the matrix $\tilde{Y} = [\tilde{X}_1, \cdots, \tilde{X}_n]$. We then proceed to analyze the following two-stage procedure, replacing $\dot{Y}$ in Algorithm 1 by $\tilde{Y}$:

1. Let $\tilde{v}$ be the leading eigenvector of $H(\tilde{Y}^T\tilde{Y})$ and let $\tilde{\eta}^0 = \mathrm{sign}(\tilde{v})$;
2. Run $\tilde{\eta}^{k+1} = \mathrm{sign}(H(\tilde{Y}^T\tilde{Y})\tilde{\eta}^k)$ to obtain the sequence $\{\tilde{\eta}^i\}_{i\geq 0}$.

We let $e_i = \tilde{X}_i - X_i$ for $i \in [n]$, then we have that e_i is zero-mean and $O(\lambda)$ sub-Gaussian (cf. Lemma 12). We further let $u_i = e_i + \varepsilon_i$, then the observations $\tilde{X}_i$ can be expressed as

$$\tilde{X}_i = X_i + e_i = \eta_i\theta + \varepsilon_i + e_i = X_i + u_i,$$

or as the matrix form

$$\tilde{Y} = \theta\eta^T + W + E = \theta\eta^T + U$$

by letting

$$W = [\varepsilon_1, ..., \varepsilon_n], \quad E = [e_1, ..., e_n] \quad \text{and} \quad U = [u_1, ..., u_n].$$

Given that the columns of E and U are both zero-mean and sub-Gaussian, we can largely follow the argument in [23] to analyze the algorithm along with a transition from Gaussian noise to sub-Gaussian noise.

Step 2. Extending the Analysis of [23, Theorems 3 & 4]

Step 2.1. Crude Bound on Spectral Initialization

We use $\tilde{Y} = \theta\eta^T + U$ and start with

$$\frac{1}{n}\tilde{Y}^T\tilde{Y} = \frac{1}{n}(\theta\eta^T + U)^T(\theta\eta^T + U) = \frac{\|\theta\|_2^2}{n}\eta\eta^T + \underbrace{\frac{\eta\theta^T U + U^T\theta\eta^T}{n} + \frac{U^T U}{n}}_{:=Z_1},$$

which gives

$$H\Big(\frac{1}{n}\tilde{Y}^T\tilde{Y}\Big) = \frac{\|\theta\|_2^2}{n}\eta\eta^T - \frac{\|\theta\|_2^2 I_n}{n} + H(Z_1). \tag{19}$$

Therefore, by letting $Z_2 = H(Z_1) - \frac{\|\theta\|_2^2}{n} I_n$, we reach

$$\Big\| H\Big(\frac{1}{n}\tilde{Y}^T\tilde{Y}\Big) - \frac{\|\theta\|_2^2}{n}\eta\eta^T \Big\|_{op} \le \|Z_2\|_{op}. \tag{20}$$

Further noticing that the columns of U are independent and zero-mean, triangle inequality yields

$$\begin{aligned}
\|Z_2\|_{op} &\le \frac{\|\theta\|_2^2}{n} + 2\Big\| H\Big(\frac{\eta\theta^T U}{n}\Big)\Big\| + \Big\| H\Big(\frac{U^T U}{n}\Big)\Big\|_{op} \\
&\quad \blacktriangleright \text{ by } Z_2 = H\Big(\frac{\eta\theta^T U}{n}\Big) + H\Big(\frac{U^T\theta\eta^T}{n}\Big) + H\Big(\frac{U^T U}{n}\Big) - \frac{\|\theta\|_2^2}{n} I_n \\
&\le \frac{\|\theta\|_2^2}{n} + 4\Big\|\frac{\eta\theta^T U}{n}\Big\|_{op} + \Big\| H\Big(\frac{U^T U}{n} - \mathbb{E}\frac{U^T U}{n}\Big)\Big\|_{op} \\
&\quad \blacktriangleright \text{ by Lemmas 9 and 10} \\
&\le \frac{\|\theta\|_2^2}{n} + \frac{4\|\eta\|_2\|U^T\theta\|_2}{n} + \frac{2\|U^T U - \mathbb{E}(U^T U)\|_{op}}{n}. \qquad (21) \\
&\quad \blacktriangleright \text{ by } \|\eta\theta^T U\|_{op} \le \|\eta\|_2\|U^T\theta\|_2
\end{aligned}$$

Since the columns of U are independent, zero-mean and have sub-Gaussian norm bounded by

$$\|u_i\|_{\psi_2} \le \|\varepsilon_i\|_{\psi_2} + \|e_i\|_{\psi_2} \lesssim \sigma + \lambda \lesssim \lambda, \tag{22}$$

▶ by Assumption 1, Lemma 12, Equation (6)

Therefore, $U^T\theta \in \mathbb{R}^n$ is a vector with independent entries whose sub-Gaussian norms are bounded by $O(\lambda\|\theta\|_2)$. Thus, we invoke Lemma 13 to obtain

$$\mathbb{P}\Big(\|U^T\theta\|_2 \le C\lambda\|\theta\|_2\sqrt{n}\Big) \ge 1 - 2\exp(-2n), \tag{23}$$

and on the high-probability event $\|U^T\theta\|_2 \le C\lambda\|\theta\|_2\sqrt{n}$ we have

$$\frac{4\|\eta\|_2\|U^T\theta\|_2}{n} = \frac{4\|U^T\theta\|_2}{\sqrt{n}} \le 4C\lambda\|\theta\|_2. \tag{24}$$

Moreover, since U has independent rows of $O(\lambda)$ sub-Gaussian norms,[7] by Lemma 14 we have

$$\mathbb{P}\Big(\|U^T U - \mathbb{E}(U^T U)\|_{op} \le C\cdot\lambda^2 n\max\{1, \sqrt{\frac{p}{n}}\}\Big) \ge 1 - 2\exp(-2n). \tag{25}$$

[7] To see this, notice that $U = W + E$ and both W and E have independent entries of $O(\lambda)$ sub-Gaussian norms.

Therefore,

$$\begin{aligned}
&\left\| H\Big(\frac{1}{n}\tilde{Y}^T\tilde{Y}\Big) - \frac{\|\theta\|_2^2}{n}\eta\eta^T \right\|_{op} \le \|Z_2\|_{op} \\
&\le \frac{\|\theta\|_2^2}{n} + C\Big(\lambda\|\theta\|_2 + \lambda^2 \max\{1, \sqrt{p/n}\}\Big) \\
&\quad \blacktriangleright \text{ by substituting (24), (25) into (21)} \\
&\le \frac{\|\theta\|_2^2}{n} + C'\Big(\lambda\|\theta\|_2 + \lambda^2\sqrt{\frac{p}{n}}\Big), \qquad (26) \\
&\quad \blacktriangleright \text{ by } \|\theta\|_2 \ge \lambda \text{; see Equation (8)} \\
&\le \frac{\|\theta\|_2^2}{2} \\
&\quad \blacktriangleright \text{ by } \lambda \le \frac{c\|\theta\|_2}{1+(p/n)^{1/4}} \text{ for small enough } c
\end{aligned}$$

Note that the leading eigenvalue of $\frac{\|\theta\|_2^2\eta\eta^T}{n}$ is

$$\lambda_1\Big(\frac{\|\theta\|_2^2\eta\eta^T}{n}\Big) = \|\theta\|_2^2,$$

with the leading eigenvector being $\frac{\eta}{\sqrt{n}}$. Therefore, by Davis-Kahan's theorem (e.g., [9, Theorem 20]), the leading eigenvector of $H(\tilde{Y}^T\tilde{Y})$ denoted by $\tilde{v}$ satisfies

$$\min\left\{\Big\|\hat{v} - \frac{\eta}{\sqrt{n}}\Big\|_2, \Big\|\hat{v} + \frac{\eta}{\sqrt{n}}\Big\|_2\right\} \lesssim \frac{\|Z_2\|_{op}}{\|\theta\|_2^2} \lesssim \frac{1}{n} + \frac{\lambda}{\|\theta\|_2} + \frac{\lambda^2}{\|\theta\|_2^2}\sqrt{\frac{p}{n}}\,. \qquad (27)$$

Recall that $\tilde{\eta}^0 = \mathrm{sign}(\tilde{v})$, and therefore

$$\begin{aligned}
\frac{\min\{\|\tilde{\eta}^0 - \eta\|_1, \|\tilde{\eta}^0 + \eta\|_1\}}{n} &\le 2\min\left\{\Big\|\hat{v} - \frac{\eta}{\sqrt{n}}\Big\|_2^2, \Big\|\hat{v} + \frac{\eta}{\sqrt{n}}\Big\|_2^2\right\} \\
&\quad \blacktriangleright \text{ by Lemma 11} \\
&\lesssim \frac{1}{n^2} + \frac{\lambda^2}{\|\theta\|_2^2} + \frac{\lambda^4}{\|\theta\|_2^4}\frac{p}{n}. \\
&\quad \blacktriangleright \text{ by Equation (27)}
\end{aligned}$$

By using $r_{n,\lambda} = \frac{\|\theta\|_2^2/\lambda^2}{\sqrt{\|\theta\|_2^2/\lambda^2 + p/n}}$, it follows that

$$\mathbb{P}\left(\ell(\tilde{\eta}^0, \eta) \le C\Big[\frac{1}{n^2} + \frac{1}{r_{n,\lambda}^2}\Big]\right) \ge 1 - 4\exp(-2n) \qquad (28)$$

for some absolute constant C.

Step 2.2. Analyzing the Local Refinement

We now analyze the iterates of Lloyd's algorithm. We define the following events

$$A_i := \left\{\Big(\frac{H(\tilde{Y}^T\tilde{Y})_i^T}{n}\eta\Big)\eta_i \ge \frac{C^*\|\theta\|_2^2}{r_{n,\lambda}}\right\}, \quad i = 1, ..., n. \qquad (29)$$

where $H(\tilde{Y}^T\tilde{Y})_i^T$ denotes the i-th row of $H(\tilde{Y}^T\tilde{Y})$, C^* can be chosen to be large enough. We then define

$$B := \left\{\|Z_2\|_{op} \le C_1\left(\frac{\|\theta\|_2^2}{n} + \lambda\|\theta\|_2 + \lambda^2\sqrt{\frac{p}{n}}\right)\right\} \qquad (30)$$

that holds for some absolute constant C_1 with probability at least $1 - 4\exp(-2n)$ — this is a consequence of the analysis in Step 2.1, see Equation (26). Furthermore, we define the event

$$C := \left\{\frac{1}{n}\sum_{i=1}^n \mathbf{1}(A_i^c) \le \frac{C_2}{r_{n,\lambda}^2}\right\}, \qquad (31)$$

Moreover, we define

$$B' := \left\{ \ell(\tilde{\eta}^0, \eta) \le C_3 \Big[\frac{1}{n^2} + \frac{1}{r_{n,\lambda}^2} \Big] \right\} \tag{32}$$

that holds for some absolute constant C_3 with probability at least $1 - 4\exp(-2n)$; see Equation (28). Without loss of generality, we assume $\ell(\tilde{\eta}^0, \eta) = \frac{1}{n}\|\tilde{\eta}^0 - \eta\|_1$ and can write

$$B' = \left\{ \frac{\|\tilde{\eta}^0 - \eta\|_1}{n} \le C_3 \Big[\frac{1}{n^2} + \frac{1}{r_{n,\lambda}^2} \Big] \right\}. \tag{33}$$

By the above discussion,

$$\mathbb{P}(B) \ge 1 - 4\exp(-2n), \quad \text{and} \quad \mathbb{P}(B') \ge 1 - 4\exp(2n). \tag{34}$$

Iterates Stay Near η on $S := B \cap B' \cap C$. We first prove

$$\frac{\|\tilde{\eta}^k - \eta\|_1}{n} \mathbf{1}(S) \le \frac{C_4}{r_{n,\lambda}^2} + \frac{C_4}{n^2}, \quad \forall k \ge 0 \tag{35}$$

for some large enough constant C_4 by induction. By B' in Equation (33), (35) is trivial for $k = 0$. Now, we suppose

$$\frac{\|\tilde{\eta}^k - \eta\|_1}{n} \mathbf{1}(S) \le \frac{C_4}{r_{n,\lambda}^2} + \frac{C_4}{n^2} \tag{36}$$

and seek to prove

$$\frac{\|\tilde{\eta}^{k+1} - \eta\|_1}{n} \mathbf{1}(S) \le \frac{C_4}{r_{n,\lambda}^2} + \frac{C_4}{n^2}\,.$$

By $\tilde{\eta}^{k+1} = \mathrm{sign}(H(\tilde{Y}^T\tilde{Y})\tilde{\eta}^k)$, the i-th entry of $\tilde{\eta}^{k+1}$ is given by

$$(\tilde{\eta}^{k+1})_i = \mathrm{sign}\left(H(\tilde{Y}^T\tilde{Y})_i^T \tilde{\eta}^k \right) = \mathrm{sign}\left(H\Big(\frac{1}{n}\tilde{Y}^T\tilde{Y}\Big)_i^T \tilde{\eta}^k \right). \tag{37}$$

We now compute $H(\frac{1}{n}\tilde{Y}^T\tilde{Y})_i^T\tilde{\eta}^k$ in the following:

$$\begin{aligned}
&H(\frac{1}{n}\tilde{Y}^T\tilde{Y})_i^T\tilde{\eta}^k = H(\frac{1}{n}\tilde{Y}^T\tilde{Y})_i^T(\tilde{\eta}^k - \eta) + H(\frac{1}{n}\tilde{Y}^T\tilde{Y})_i^T\eta \\
&= (Z_2)_i^T(\tilde{\eta}^k - \eta) + \frac{\|\theta\|_2^2}{n}\eta_i\eta^T(\tilde{\eta}^k - \eta) + H(\frac{1}{n}\tilde{Y}^T\tilde{Y})_i^T\eta \\
&\quad \blacktriangleright \text{ by } H(\frac{1}{n}\tilde{Y}^T\tilde{Y}) = \frac{\|\theta\|_2^2}{n}\eta\eta^T + Z_2 \text{ from Equation (19)} \\
&= (Z_2)_i^T(\tilde{\eta}^k - \eta) - \frac{\eta_i\|\theta\|_2^2}{n}\big(n - \eta^T\tilde{\eta}^k\big) + H(\frac{1}{n}\tilde{Y}^T\tilde{Y})_i^T\eta \\
&\quad \blacktriangleright \text{ by } \eta^T(\tilde{\eta}^k - \eta) = \eta^T\tilde{\eta}^k - n \\
&= (Z_2)_i^T(\tilde{\eta}^k - \eta) - \eta_i\|\theta\|_2^2\frac{\|\tilde{\eta}^k - \eta\|_1}{n} + H(\frac{1}{n}\tilde{Y}^T\tilde{Y})_i^T\eta. \qquad (38) \\
&\quad \blacktriangleright \text{ by } \|\tilde{\eta}^k - \eta\|_1 = \frac{1}{2}\|\tilde{\eta}^k - \eta\|_2^2 = n - \eta^T\tilde{\eta}^k
\end{aligned}$$

Note that the error of $\tilde{\eta}$ can be written as

$$\frac{\|\tilde{\eta}^{k+1} - \eta\|_1}{n} = \frac{2}{n}\Bigg(\sum_{i:\eta_i=1} \mathbf{1}(H(\tilde{Y}^T\tilde{Y})_i^T\tilde{\eta}^k < 0) + \sum_{i:\eta_i=-1} \mathbf{1}(H(\tilde{Y}^T\tilde{Y})_i^T\tilde{\eta}^k > 0) \Bigg)$$

and therefore

$$\frac{\|\tilde{\eta}^{k+1} - \eta\|_1}{n}\mathbf{1}(S) = \frac{2}{n}\Bigg(\sum_{i:\eta_i=1} \mathbf{1}(H(\tilde{Y}^T\tilde{Y})_i^T\tilde{\eta}^k < 0)\mathbf{1}(S) + \sum_{i:\eta_i=-1} \mathbf{1}(H(\tilde{Y}^T\tilde{Y})_i^T\tilde{\eta}^k > 0)\mathbf{1}(S) \Bigg). \tag{39}$$

We separately discuss the cases of $\eta_i = 1$ and $\eta_i = -1$:

- If $\eta_i = 1$, then by (38)

$$H(\frac{1}{n}\tilde{Y}^T\tilde{Y})_i^T\tilde{\eta}^k = (Z_2)_i^T(\tilde{\eta}^k - \eta) - \|\theta\|_2^2\frac{\|\tilde{\eta}^k - \eta\|_1}{n} + H(\frac{1}{n}\tilde{Y}^T\tilde{Y})_i^T\eta$$

and the event A_i defined in Equation (29) reads

$$A_i = \left\{\frac{1}{n}H(\tilde{Y}^T\tilde{Y})_i^T\eta \geq \frac{C^*\|\theta\|_2^2}{r_{n,\lambda}}\right\}. \tag{40}$$

Therefore,

$$\begin{aligned}
&\mathbf{1}(H(\tilde{Y}^T\tilde{Y})_i^T\tilde{\eta}^k < 0) \leq \mathbf{1}(H(\frac{1}{n}\tilde{Y}^T\tilde{Y})_i^T\tilde{\eta}^k < 0, A_i) + \mathbf{1}(A_i^c)\\
&\leq \mathbf{1}\left((Z_2)_i^T(\tilde{\eta}^k - \eta) \leq -\frac{C^*\|\theta\|_2^2}{r_{n,\lambda}} + \frac{\|\theta\|_2^2\|\tilde{\eta}^k - \eta\|_1}{n}\right) + \mathbf{1}(A_i^c)\\
&\quad \blacktriangleright \text{ by Equations (38) and (40)}
\end{aligned}$$

which further leads to

$$\begin{aligned}
&\mathbf{1}(H(\tilde{Y}^T\tilde{Y})_i^T\tilde{\eta}^k < 0)\mathbf{1}(S)\\
&= \mathbf{1}\left((Z_2)_i^T(\tilde{\eta}^k - \eta) \leq -\frac{C^*\|\theta\|_2^2}{r_{n,\lambda}} + \frac{\|\theta\|_2^2\|\tilde{\eta}^k - \eta\|_1}{n},\ S\right) + \mathbf{1}(A_i^c, S)\\
&\leq \mathbf{1}\left((Z_2)_i^T(\tilde{\eta}^k - \eta) \leq -\frac{C^*\|\theta\|_2^2}{2r_{n,\lambda}} + \frac{C_4\|\theta\|_2^2}{n^2},\ S\right) + \mathbf{1}(A_i^c, S)\\
&\quad \blacktriangleright \text{ by Equation (36)}\\
&\leq \mathbf{1}\left((Z_2)_i^T(\tilde{\eta}^k - \eta) \leq -\frac{C^*\|\theta\|_2^2}{3r_{n,\lambda}},\ S\right) + \mathbf{1}(A_i^c, S)\\
&\quad \blacktriangleright \text{ by the assumption } \underbrace{\|\theta\|_2 \leq \lambda n\Big(1 + \frac{p^{1/4}}{n^{3/4}}\Big)}_{\text{due to } r_{n,\lambda} \lesssim n} \text{ and large enough } C^*
\end{aligned}$$

Further using $\mathbf{1}(T \leq -a) \leq \frac{T^2}{a^2}$ for any $a > 0$, we arrive at

$$\mathbf{1}(H(\tilde{Y}^T\tilde{Y})_i^T\tilde{\eta}^k < 0)\mathbf{1}(S) \leq \frac{[3r_{n,\lambda}(Z_2)_i^T(\tilde{\eta}^k - \eta)]^2}{(C^*)^2\|\theta\|_2^4}\mathbf{1}(S) + \mathbf{1}(A_i^c, S) \tag{41}$$

- If $\eta_i = -1$, then by argument parallel to the last dot point, we obtain

$$\mathbf{1}(H(\tilde{Y}^T\tilde{Y})_i^T\tilde{\eta}^k > 0)\mathbf{1}(S) \leq \frac{[3r_{n,\lambda}(Z_2)_i^T(\tilde{\eta}^k - \eta)]^2}{(C^*)^2\|\theta\|_2^4}\mathbf{1}(S) + \mathbf{1}(A_i^c, S) \tag{42}$$

The above disucssions then yield

$$\frac{\|\tilde{\eta}^{k+1}-\eta\|_1}{n}\mathbf{1}(S) \le \frac{2}{n}\sum_{i=1}^{n}\left(\frac{[3r_{n,\lambda}(Z_2)_i^T(\tilde{\eta}^k-\eta)]^2}{(C^*)^2\|\theta\|_2^4}\mathbf{1}(S)+\mathbf{1}(A_i^c,S)\right)$$

► by substituting (41) and (42) into (39)

$$\le \frac{18}{n}\frac{r_{n,\lambda}^2\sum_{i=1}^{n}[(Z_2)_i^T(\tilde{\eta}^k-\eta)]^2}{(C^*)^2\|\theta\|_2^4}\mathbf{1}(S)+\frac{2}{n}\sum_{i=1}^{n}\mathbf{1}(A_i^c,S)$$

► by simple algebra

$$\le \frac{18}{n}\frac{r_{n,\lambda}^2\|Z_2\|_{op}^2\|\tilde{\eta}^k-\eta\|_2^2}{(C^*)^2\|\theta\|_2^4}\mathbf{1}(S)+\frac{2}{n}\sum_{i=1}^{n}\mathbf{1}(A_i^c)\mathbf{1}(S) \tag{43}$$

► by $\sum_{i=1}^{n}[(Z_2)_i^T(\tilde{\eta}^k-\eta)]^2 = \|Z_2(\tilde{\eta}^k-\eta)\|_2^2 \le \|Z_2\|_{op}^2\|\tilde{\eta}^k-\eta\|_2^2$

$$\le \frac{\|\tilde{\eta}^k-\eta\|_1\mathbf{1}(S)}{n}O\left(\frac{r_{n,\lambda}^2}{(C^*)^2\|\theta\|_2^4}\left(\frac{\|\theta\|_2^4}{n^2}+\lambda^2\|\theta\|_2^2+\lambda^4\frac{p}{n}\right)\right)+O(r_{n,\lambda}^{-2})$$

► by $\|\tilde{\eta}^k-\eta\|_1 = \frac{1}{2}\|\tilde{\eta}^k-\eta\|_2^2$ and the events in Equations (30) and (31)

$$\le\left(\frac{1}{r_{n,\lambda}^2}+\frac{1}{n^2}\right)\underbrace{O\left(\frac{r_{n,\lambda}^2}{(C^*)^2}\left(\frac{1}{n^2}+\frac{\lambda^2}{\|\theta\|_2^2}+\frac{\lambda^4}{\|\theta\|_2^4}\frac{p}{n}\right)\right)}_{:=F}+O(r_{n,\lambda}^{-2})$$

► by the hypothesis in (36)

We now examine the scaling of F: since $r_{n,\lambda}^2\left(\frac{\lambda^2}{\|\theta\|_2^2}+\frac{\lambda^4}{\|\theta\|_2^4}\frac{p}{n}\right)=\Theta(1)$, we have $F \lesssim \frac{1}{(C^*)^2}\left(\frac{r_{n,\lambda}^2}{n^2}+\Theta(1)\right)$; moreover, $\|\theta\|_2 \le \lambda(n+(np)^{1/4})$ implies $r_{n,\lambda}=O(n)$ and therefore $F=O(\frac{1}{(C^*)^2})$; therefore, we can set C^* large enough to render a small enough F, yielding

$$\frac{\|\tilde{\eta}^{k+1}-\eta\|_1}{n}\mathbf{1}(S) \le \frac{C_4}{r_{n,\lambda}^2}+\frac{C_4}{n^2}$$

as long as C_4 is suitably large. The induction is complete and we have established

$$\frac{\|\tilde{\eta}^k-\eta\|_1}{n}\mathbf{1}(S) \le \frac{C_4}{r_{n,\lambda}^2}+\frac{C_4}{n^2}, \quad k\ge 0. \tag{44}$$

Yielding Sharp Bound. We shall proceed with an intermediate result in the above induction argument, which also holds for all $k \ge 0$: by Equation (43), we have

$$\frac{\|\tilde{\eta}^{k+1}-\eta\|_1}{n}\mathbf{1}(S) \le \frac{36r_{n,\lambda}^2\|Z_2\|_{op}^2}{(C^*)^2\|\theta\|_2^4}\cdot\frac{\|\tilde{\eta}^k-\eta\|_1}{n}\mathbf{1}(S)+\frac{2}{n}\sum_{i=1}^{n}\mathbf{1}(A_i^c)\mathbf{1}(S), \quad \forall k\ge 0.$$

As before, we substitute the bound on $\|Z_2\|_{op}$ from Equation (30) and take C^* large enough to obtain

$$\frac{\|\tilde{\eta}^{k+1}-\eta\|_1}{n}\mathbf{1}(S) \le \frac{1}{4}\cdot\frac{\|\tilde{\eta}^k-\eta\|_1}{n}\mathbf{1}(S)+\frac{2}{n}\sum_{i=1}^{n}\mathbf{1}(A_i^c)\mathbf{1}(S), \quad \forall k\ge 0.$$

We then iterate this inequality, along with $\frac{\|\tilde{\eta}_k-\eta\|_1}{n}\le 2$, to yield

$$\frac{\|\tilde{\eta}^k-\eta\|_1}{n}\mathbf{1}(S) \le \frac{2}{4^k}+\frac{C_5}{n}\sum_{i=1}^{n}\mathbf{1}(A_i^c)\mathbf{1}(S), \quad \forall k\ge 0.$$

Thus, it holds that

$$\frac{\|\tilde{\eta}^k-\eta\|_1}{n}\mathbf{1}(S) \le \frac{1}{n^2}+\frac{C_5}{n}\sum_{i=1}^{n}\mathbf{1}(A_i^c)\mathbf{1}(S), \qquad \forall k\ge\lceil 3\log_4 n\rceil. \tag{45}$$

We now take expectation,

$$\begin{aligned}
\mathbb{E}\frac{\|\tilde{\eta}^k-\eta\|_1}{n} &\le \mathbb{E}\Big(\frac{\|\tilde{\eta}^k-\eta\|_1}{n}\mathbf{1}(S)\Big)+\mathbb{E}\Big(\frac{\|\tilde{\eta}^k-\eta\|_1}{n}\Big[\mathbf{1}(B^c)+\mathbf{1}\big((B')^c\big)+\mathbf{1}(C^c)\Big]\Big)\\
&\qquad \blacktriangleright \text{ by } S=B\cap B'\cap C \text{ and hence } S^c=B^c\cup(B')^c\cup C^c\\
&\le \frac{1}{n^2}+\mathbb{E}\Big(\frac{C_5}{n}\sum_{i=1}^n\mathbf{1}(A_i^c)\Big)+2\big(\mathbb{P}(B^c)+\mathbb{P}\big((B')^c\big)+\mathbb{P}(C^c)\big)\\
&\qquad \blacktriangleright \text{ by Equation (45) and } \frac{\|\tilde{\eta}^k-\eta\|_1}{n}\le 2\\
&\le \frac{1}{n^2}+C_5\mathbb{P}(A_i^c)+16\exp(-2n)+2\mathbb{P}(C^c),\quad \forall k\ge\lceil 3\log_4 n\rceil. \qquad (46)\\
&\qquad \blacktriangleright \text{ by } \mathbb{P}(B^c)+\mathbb{P}((B')^c)\le 8\exp(-2n) \text{ from (34)}
\end{aligned}$$

Recall from Equation (31) that $C^c=\big\{\frac{1}{n}\sum_{i=1}^n\mathbf{1}(A_i^c)>\frac{C_2}{r_{n,\lambda}^2}\big\}$. Thus, by Markov's inequality we have

$$\begin{aligned}
\frac{C_2\mathbb{P}(C^c)}{r_{n,\lambda}^2}&\le\mathbb{E}\Big(\frac{1}{n}\sum_{i=1}^n\mathbf{1}(A_i^c)\Big)=\mathbb{P}(A_i^c)\\
\Longrightarrow \mathbb{P}(C^c)&=O(r_{n,\lambda}^2\mathbb{P}(A_i^c)).
\end{aligned}$$

Substituting this into (46), along with $r_{n,\lambda}\gtrsim 1$, yields

$$\mathbb{E}\frac{\|\tilde{\eta}^k-\eta\|_1}{n}\le\frac{2}{n^2}+C'r_{n,\lambda}^2\mathbb{P}(A_i^c),\qquad \forall k\ge\lceil 3\log_4 n\rceil \qquad (47)$$

where the sufficiently large n ensures $16\exp(-2n)\le\frac{1}{n^2}$. Recall from Equation (29) that

$$A_i^c=\Big\{\Big(\frac{H(\tilde{Y}^T\tilde{Y})_i^T}{n}\eta\Big)\eta_i<\frac{C^*\|\theta\|_2^2}{r_{n,\lambda}}\Big\}. \qquad (48)$$

By noticing

$$H(\tilde{Y}^T\tilde{Y})_i^T=\big(\langle\tilde{X}_i,\tilde{X}_1\rangle,\cdots,\langle\tilde{X}_i,\tilde{X}_{i-1}\rangle,0,\langle\tilde{X}_i,\tilde{X}_{i+1}\rangle,\cdots,\langle\tilde{X}_i,\tilde{X}_n\rangle\big),$$

we have that

$$\Big(\frac{H(\tilde{Y}^T\tilde{Y})_i^T}{n}\eta\Big)\eta_i=\frac{1}{n}\sum_{j\ne i}\eta_i\eta_j\langle\tilde{X}_i,\tilde{X}_j\rangle=\Big\langle\eta_i\tilde{X}_i,\frac{1}{n}\sum_{j\ne i}\eta_j\tilde{X}_j\Big\rangle.$$

By $\varepsilon_i\overset{d}{=}-\varepsilon_i$[8] and $\tilde{X}_i=\eta_i\theta+u_i$ where $u_i=Q_{2\lambda}(\eta_i\theta+\varepsilon_i+\tau_i)-\eta_i\theta$, we have

$$\begin{aligned}
\eta_i\tilde{X}_i&=\theta+\eta_iu_i=\theta+(Q_{2\lambda}(\theta+\eta_i\varepsilon_i+\eta_i\tau_i)-\theta)\\
&\overset{d}{=}\theta+(Q_{2\lambda}(\theta+\varepsilon_i+\tau_i)-\theta).
\end{aligned}$$

As a result, we can let $\delta_1,\cdots,\delta_n$ be i.i.d. and follow the distribution $Q_{2\lambda}(\theta+\varepsilon_i+\tau_i)-\theta$, then we have

$$\Big(\frac{H(\tilde{Y}^T\tilde{Y})_i^T}{n}\eta\Big)\eta_i\overset{\mathrm{d}}{=}\Big\langle\theta+\delta_1,\frac{1}{n}\sum_{j=2}^n(\theta+\delta_j)\Big\rangle \qquad (49)$$

and thus

$$\begin{aligned}
\mathbb{P}(A_i^c)&=\mathbb{P}\Big(\Big\langle\theta+\delta_1,\frac{1}{n}\sum_{j=2}^n(\theta+\delta_j)\Big\rangle<\frac{C^*\|\theta\|_2^2}{r_{n,\lambda}}\Big)\\
&\qquad\blacktriangleright \text{ by substituting (49) into (48)}\\
&\le\mathbb{P}\Big(\Big\langle\theta+\delta_1,\theta+\frac{1}{n-1}\sum_{j=2}^n\delta_j\Big\rangle<\frac{2C^*\|\theta\|_2^2}{r_{n,\lambda}}\Big):=P^*, \qquad (50)\\
&\qquad\blacktriangleright \text{ by } \frac{1}{n-1}<\frac{2}{n}
\end{aligned}$$

[8]We use $A\overset{d}{=}B$ to denote that A and B have the same distribution.

where $2C^*$ is some large enough absolute constant. Substituting this into (47) yields

$$\mathbb{E}\Big(\frac{1}{n}\|\tilde{\eta}^k-\eta\|_1\Big)\le\frac{2}{n^2}+C'r_{n,\lambda}^2P^*\quad\text{for any }\ k\ge\lceil 3\log_4 n\rceil,$$

and Markov's inequality further yields

$$\mathbb{P}\Big(\frac{\|\tilde{\eta}^k-\eta\|_1}{n}\ge\frac{2\sqrt{r_{n,\lambda}}}{n^2}+C'r_{n,\lambda}^{5/2}P^*\Big)\le r_{n,\lambda}^{-1/2}.$$

Combining with $r_{n,\lambda}=O(n)$, we reach

$$\mathbb{P}\bigg(\frac{\|\tilde{\eta}^k-\eta\|_1}{n}\le\frac{C}{n^{3/2}}+C'r_{n,\lambda}^{5/2}P^*\bigg)\ge 1-r_{n,\lambda}^{-1/2}. \tag{51}$$

We have now completed the analysis of Algorithm 1 with $\tilde{Y}$.

Step 3. Reduction to One-Bit Quantization

All that remains is to reduce to the one-bit quantization regime by transitioning from $\tilde{\eta}^k$ (obtained by running Algorithm 1 with $\dot{Y}$ being replaced by $\tilde{Y}$) to $\hat{\eta}^k$ (which is obtained by running Algorithm 1). To this end, we only need to show $\lambda\dot{Y}=\tilde{Y}$, namely

$$\lambda\,\mathrm{sign}(\eta_i\theta+\varepsilon_i+\tau_i)=\lambda\dot{X}_i=\tilde{X}_i=Q_{2\lambda}(\eta_i\theta+\varepsilon_i+\tau_i),\quad\forall i\in[n]. \tag{52}$$

The reason is that, on the event of (52), $\hat{\eta}^k=\tilde{\eta}^k$ holds for all $k\ge 0$, and therefore the bound for $\tilde{\eta}^k$ in Equation (51) transfers to $\hat{\eta}^k$ of interest. By the observation

$$Q_{2\lambda}(a)=\lambda\,\mathrm{sign}(a)\,,\quad\forall|a|\le 2\lambda,$$

it suffices to show $\|\eta_i\theta+\varepsilon_i+\tau_i\|_\infty\le 2\lambda$ for all $i\in[n]$. Using triangle inequality and combining with $\|\tau_i\|_\infty\le\lambda$, it suffices to show

$$\|W\|_{\max}:=\sup_{i\in[n]}\|\varepsilon_i\|_\infty\le\lambda-\|\theta\|_\infty,\qquad\forall i\in[n]. \tag{53}$$

By the tail bound of the entries of W (Lemma 2) and a union bound, it follows that

$$\mathbb{P}\Big(\|W\|_{\max}\ge\sigma\sqrt{2(1+\nu)\log(np)}\Big)\le\frac{2}{(np)^\nu}.$$

As we assume $\lambda-\|\theta\|_\infty\ge\sigma\sqrt{2(1+\nu)\log(np)}$, the desired result follows.

Step 4. Sharp Bound on P^*

In light of

$$\Big\langle\theta+\delta_1,\theta+\frac{1}{n-1}\sum_{j=2}^n\delta_j\Big\rangle=\|\theta\|_2^2+\Big\langle\theta,\frac{1}{n-1}\sum_{j=2}^n\delta_j\Big\rangle+\Big\langle\delta_1,\theta+\frac{1}{n-1}\sum_{j=2}^n\delta_j\Big\rangle$$

and the definition of P^* in Equation (50), we have

$$P^*=\mathbb{P}\Big(\underbrace{\Big\langle\theta,\frac{1}{n-1}\sum_{j=2}^n\delta_j\Big\rangle+\Big\langle\delta_1,\theta+\frac{1}{n-1}\sum_{j=2}^n\delta_j\Big\rangle<\Big(-1+\frac{C_1}{r_{n,\lambda}}\Big)\|\theta\|_2^2}_{\text{denoted by an event }\Gamma}\Big). \tag{54}$$

We first show that $\langle\theta,\frac{1}{n-1}\sum_{j=2}^n\delta_j\rangle$ is a higher order term. Since $\delta_j\overset{\mathrm{d}}{=}Q_{2\lambda}(\theta+\varepsilon_j+\tau_j)-\theta$ has independent, zero mean entries with $O(\lambda)$ sub-Gaussian norms (e.g., by argument parallel to Equation (22)), we can use Lemma 3 to obtain

$$\Big\|\frac{1}{n-1}\sum_{j=2}^n\delta_j\Big\|_{\psi_2}\lesssim\frac{\lambda}{\sqrt{n}}\implies\Big\|\Big\langle\theta,\frac{1}{n-1}\sum_{j=2}^n\delta_j\Big\rangle\Big\|_{\psi_2}\lesssim\frac{\lambda\|\theta\|_2}{\sqrt{n}}. \tag{55}$$

Therefore, let

$$\hat{B} := \left\{ \left| \left\langle \theta, \frac{1}{n-1} \sum_{j=2}^{n} \delta_j \right\rangle \right| < C\lambda \|\theta\|_2 \sqrt{\frac{\log n}{n}} \right\} \text{ with large enough } C, \tag{56}$$

then by (55) and a sub-Gaussian tail bound,

$$\mathbb{P}(\hat{B}^c) = \mathbb{P}\left(\left| \left\langle \theta, \frac{1}{n-1} \sum_{j=2}^{n} \delta_j \right\rangle \right| \geq C\lambda \|\theta\|_2 \sqrt{\frac{\log n}{n}} \right) \leq \frac{1}{n^2}. \tag{57}$$

Moreover, we let

$$\hat{B}_1 := \left\{ \left\| \frac{1}{n-1} \sum_{j=2}^{n} \delta_j \right\|_2^2 \leq \big(1 + O\Big(\frac{1}{n}\Big)\big) \frac{p\lambda^2}{n} + O\Big(\frac{1}{\sqrt{n}}\Big) \|\theta\|_2^2 \right\}, \tag{58}$$

then Lemma 15 gives

$$\mathbb{P}(\hat{B}_1^c) \leq 2\exp(-10 r_{n,\lambda}^2). \tag{59}$$

On the event $\hat{B} \cap \hat{B}_1$, along with $\|\theta\|_2 \gtrsim \lambda$, we have

$$\begin{aligned}
\left\| \theta + \frac{1}{n-1} \sum_{j=2}^{n} \delta_j \right\|_2^2 &= \|\theta\|_2^2 + \left\| \frac{1}{n-1} \sum_{j=2}^{n} \delta_j \right\|_2^2 + 2 \left\langle \theta, \frac{1}{n-1} \sum_{j=2}^{n} \delta_j \right\rangle \\
&\quad \blacktriangleright \text{ by simple algebra} \\
&\leq \|\theta\|_2^2 + \Big(1 + O(\frac{1}{n})\Big) \frac{p\lambda^2}{n} + O(\frac{1}{\sqrt{n}}) \|\theta\|_2^2 + 2C\lambda \|\theta\|_2 \sqrt{\frac{\log n}{n}} \\
&\quad \blacktriangleright \text{ by } \hat{B} \text{ in (56) and } \hat{B}_1 \text{ in (58)} \\
&= \left(1 + O\Big(\sqrt{\frac{\log n}{n}} \Big) \right) \|\theta\|_2^2 + \left(1 + O\Big(\frac{1}{\sqrt{n}} \Big) \right) \frac{p\lambda^2}{n}. \qquad (60) \\
&\quad \blacktriangleright \text{ by } \lambda \leq \|\theta\|_2 \text{; see Equation (8)}
\end{aligned}$$

Moreover, on the event of $\hat{B} \cap \hat{B}_1$, we have

$$\begin{aligned}
&\text{The event } \Gamma \text{ in Equation (54)} \\
\implies &\left\langle \delta_1, \theta + \frac{1}{n-1} \sum_{j=2}^{n} \delta_j \right\rangle < \Big(-1 + \frac{C_1}{r_{n,\lambda}} \Big) \|\theta\|_2^2 + \left| \left\langle \theta, \frac{1}{n-1} \sum_{j=2}^{n} \delta_j \right\rangle \right| \\
\implies &\left\langle \delta_1, \theta + \frac{1}{n-1} \sum_{j=2}^{n} \delta_j \right\rangle < \Big(-1 + \frac{C_1'}{r_{n,\lambda}} \Big) \|\theta\|_2^2 \\
&\quad \blacktriangleright \text{ by } \hat{B} \text{ in Equation (56)} \\
\implies &\left\langle \delta_1, \frac{\theta + (n-1)^{-1} \sum_{j=2}^{n} \delta_j}{\|\theta + (n-1)^{-1} \sum_{j=2}^{n} \delta_j\|_2} \right\rangle < - \frac{\Big(1 - O(r_{n,\lambda}^{-1}) - O((\frac{\log n}{n})^{1/4})\Big) \|\theta\|_2^2}{\sqrt{\|\theta\|_2^2 + \frac{p\lambda^2}{n}}}. \qquad (61) \\
&\quad \blacktriangleright \text{ by Equation (60) that holds under } \hat{B} \cap \hat{B}_1
\end{aligned}$$

Therefore, by letting $\hat{\xi} := \frac{\theta + (n-1)^{-1} \sum_{j=2}^{n} \delta_j}{\|\theta + (n-1)^{-1} \sum_{j=2}^{n} \delta_j\|_2} = (\hat{\xi}_1, \cdots, \hat{\xi}_p)^T$,

$$\begin{aligned}
P^* &= \mathbb{P}(\Gamma) \leq \mathbb{P}(\hat{B}^c) + \mathbb{P}(\hat{B}_1^c) + \mathbb{P}(\Gamma \cap \hat{B} \cap \hat{B}_1) \\
&\leq \frac{1}{n^2} + 2\exp(-10 r_{n,\lambda}^2) + \underbrace{\mathbb{P}\left(\langle \delta_1, \hat{\xi} \rangle < - \Big(1 - O\Big(r_{n,\lambda}^{-1} + (\frac{\log n}{n})^{1/4} \Big)\Big) \frac{\|\theta\|_2^2}{\sqrt{\|\theta\|_2^2 + \frac{p\lambda^2}{n}}} \right)}_{:= P^{**}} \qquad (62) \\
&\quad \blacktriangleright \text{ by Equations (57), (59) and (61)}
\end{aligned}$$

To establish sharp bound on P^{**}, we need a reduction to the one-bit case. Recall that

$$\delta_1 \stackrel{\mathrm{d}}{=} Q_{2\lambda}(\theta + \varepsilon_1 + \tau_1) - \theta,$$

we shall introduce a surrogate of δ_1 as

$$\tilde{\delta}_1 := \lambda \operatorname{sign}(\theta + \varepsilon_1 + \tau_1) - \mathbb{E}\big[\lambda \operatorname{sign}(\theta + \varepsilon_1 + \tau_1)\big] = (\tilde{\delta}_{11}, \cdots, \tilde{\delta}_{1p})^T.$$

By Lemma 16, for some $\epsilon_{n,\lambda} = O(r_{n,\lambda}^{-1} + (\frac{\log n}{n})^{1/4})$,

$$\begin{aligned} P^{**} &\le \mathbb{P}\bigg(\langle \tilde{\delta}_1, \hat{\xi} \rangle < -\frac{(1-\epsilon_{n,\lambda})\|\theta\|_2^2}{\sqrt{\|\theta\|_2^2 + \frac{p\lambda^2}{n}}} \bigg) + \frac{1}{n^{1+\nu}p^{\nu}} \\ &= \mathbb{P}\bigg(\sum_{j=1}^{p} \frac{\tilde{\delta}_{1j}\hat{\xi}_j}{\lambda} < -(1-\epsilon_{n,\lambda}) r_{n,\lambda} \bigg) + \frac{1}{n^{1+\nu}p^{\nu}}\,. \qquad (63) \\ &\qquad \blacktriangleright \text{ recall } r_{n,\lambda} = \frac{\|\theta\|_2^2}{\lambda^2} \Big/ \sqrt{\frac{\|\theta\|_2^2}{\lambda^2} + \frac{p}{n}} \end{aligned}$$

Since $\{\tilde{\delta}_{1j} : j \in [p]\}$ are independent, zero-mean, and $\|\lambda \operatorname{sign}(\theta + \varepsilon_1 + \tau_1)\|_\infty \le \lambda$, we shall use Hoeffding's inequality (see Lemma 6) to obtain

$$\mathbb{P}\bigg(\sum_{j=1}^{p} \frac{\tilde{\delta}_{1j}\hat{\xi}_j}{\lambda} < -(1-\epsilon_{n,\lambda}) r_{n,\lambda} \bigg) \le 2\exp\bigg(-\frac{(1-\epsilon_{n,\lambda})^2 r_{n,\lambda}^2}{2} \bigg). \qquad (64)$$

Step 5. Putting Pieces Together

We substitute (64) and (63) into (62) to obtain

$$P^{*} \le 2\exp\bigg(-\frac{(1-\epsilon_{n,\lambda})^2 r_{n,\lambda}^2}{2} \bigg) + \frac{1}{n^2} + 2\exp(-10 r_{n,\lambda}^2) + \frac{1}{n^{1+\nu}p^{\nu}}$$

Substituting this into (51) it follows that

$$\begin{aligned} \ell(\hat{\eta}^t, \eta) &\le C r_{n,\lambda}^{5/2} \exp\bigg(-\frac{(1-\epsilon_{n,\lambda}) r_{n,\lambda}^2}{2} \bigg) + O\bigg(\frac{1}{n^{3/2}} + \frac{r_{n,\lambda}^{5/2}}{n^2} + \frac{r_{n,\lambda}^{5/2}}{n^{1+\nu}p^{\nu}} \bigg) \\ &\qquad \blacktriangleright \text{ by } \exp(-10 r_{n,\lambda}^2) \le \exp(-\frac{1}{2} r_{n,\lambda^2}) \\ &\le \frac{1}{2}\exp\bigg(-\Big(1 - \epsilon_{n,\lambda} - \frac{2\log(2C r_{n,\lambda}^{5/2})}{r_{n,\lambda}^2}\Big)\frac{r_{n,\lambda}^2}{2} \bigg) + \frac{1}{5n} \\ &\qquad \blacktriangleright \text{ by } n \gtrsim 1,\ r_{n,\lambda} \lesssim \min\{(np)^{2\nu/5}, n^{2/5}\} \\ &= \frac{1}{2}\exp\bigg(-\frac{(1-\epsilon'_{n,\lambda}) r_{n,\lambda}^2}{2} \bigg) + \frac{1}{5n} \end{aligned}$$

for some $\epsilon_{n,\lambda}, \epsilon'_{n,\lambda} = O(\frac{1}{r_{n,\lambda}} + (\frac{\log n}{n})^{1/4})$. If $\frac{1}{2}\exp(-\frac{(1-\epsilon'_{n,\lambda}) r_{n,\lambda}^2}{2}) \le \frac{1}{5n}$, then we have $\ell(\hat{\eta}^t, \eta) \le \frac{2}{5n}$ and thus necessarily $\ell(\hat{\eta}^t, \eta) = 0$. If $\frac{1}{2}\exp(-\frac{(1-\epsilon'_{n,\lambda}) r_{n,\lambda}^2}{2}) > \frac{1}{5n}$, then we have

$$\ell(\hat{\eta}^t, \eta) \le \exp\bigg(-\frac{(1-\epsilon'_{n,\lambda}) r_{n,\lambda}^2}{2} \bigg),$$

as desired. □

B Proof of Corollary 1 (Exact Recovery)

Proof. Note that the separation condition in Equation (9) implies

$$\|\theta\|_2^2 \gtrsim \lambda^2 \Big(\log n + \sqrt{\frac{p \log n}{n}} \Big).$$

In light of $r_{n,\lambda} = \frac{\|\theta\|_2^2/\lambda^2}{\sqrt{\|\theta\|_2^2/\lambda^2 + p/n}} \asymp \min\{\frac{\|\theta\|_2}{\lambda}, \frac{\|\theta\|^2}{\lambda^2}\sqrt{\frac{n}{p}}\}$, we then have

$$\frac{1}{r_{n,\lambda}} \lesssim \max\{\frac{\lambda}{\|\theta\|_2}, \frac{\lambda^2}{\|\theta\|_2^2}\sqrt{\frac{p}{n}}\} = O(\frac{1}{\sqrt{\log n}}).$$

Under $n \gtrsim 1$, Theorem 1 guarantees that

$$\ell(\hat{\eta}^t, \eta) \le \exp(-\frac{(1-\epsilon_0)}{2} r_{n,\lambda}^2), \quad \text{for some } \epsilon_0 = O(\frac{1}{\sqrt{\log n}})$$

holds with probability at least $1 - O(\frac{1}{\sqrt{\log n}}) - 2(np)^{-\nu}$. Note that $\ell(\hat{\eta}^t, \eta) < n^{-1}$ implies $\ell(\hat{\eta}^t, \eta) = 0$. Therefore, to ensure $\hat{\eta}^t = \pm\eta$ with the promised probability, it suffices to have

$$\begin{aligned}
& \exp\Big(-\frac{(1-\epsilon_0)}{2} r_{n,\lambda}^2\Big) < \frac{1}{n} \\
\iff\ & \exp\Bigg(-\frac{1-\epsilon_0}{2}\frac{(\|\theta\|_2^2/\lambda^2)^2}{\frac{\|\theta\|_2^2}{\lambda^2} + \frac{p}{n}}\Bigg) < \frac{1}{n} \\
\iff\ & \frac{1-\epsilon_0}{2}\frac{(\|\theta\|_2^2/\lambda^2)^2}{\frac{\|\theta\|_2^2}{\lambda^2} + \frac{p}{n}} > \log n \\
\iff\ & (1-\epsilon_0)\Big(\frac{\|\theta\|_2^2}{\lambda^2}\Big)^2 - 2\log n \frac{\|\theta\|_2^2}{\lambda^2} - \frac{2p\log n}{n} > 0 \\
\iff\ & \|\theta\|_2^2 > \frac{1}{1-\epsilon_0}\lambda^2\Bigg(1 + \sqrt{1 + \frac{(1-\epsilon_0)2p}{n\log n}}\Bigg)\log n
\end{aligned} \tag{65}$$

This condition can be implied by (9) by noticing $\frac{1}{1-\epsilon_0} = 1 + \frac{\epsilon_0}{1-\epsilon_0} = 1 + O(\epsilon_0) = 1 + O(\frac{1}{\sqrt{\log n}})$ and $1 - \epsilon_0 \le 1$. □

C Proof of Theorem 2 (Partial Recovery without Spikiness Condition)

We start with a remark that summarizes the additional technicalities beyond the proof of Theorem 1.

Remark 10. *The result is proved by revisiting the proof of Theorem 1 and similarly relating $\lambda\dot{X}_i = \lambda\,\mathrm{sign}(\eta_i R\theta + R\varepsilon_i + \tau_i)$ to $\tilde{X}_i := Q_{2\lambda}(\eta_i R\theta + R\varepsilon_i + \tau_i)$. Yet the overall noise $u_i = Q_{2\lambda}(\eta_i R\theta + R\varepsilon_i + \tau_i) - \eta_i R\theta$ has possibly correlated entries, making some arguments no longer valid. Our general remedy is to further decompose u_i into $u_i = \big[Q_{2\lambda}(\eta_i R\theta + R\varepsilon_i + \tau_i) - (\eta_i R\theta + R\varepsilon_i)\big] + R\varepsilon_i := u_{i1} + R\varepsilon_i$ and then separately treat u_{i1} and $R\varepsilon_i$: note that u_{i1} has independent zero-mean entries (when conditioning on ε_i) and $R\varepsilon_i$ is simply a rotation of ε_i with independent zero-mean entries. Following this idea, most arguments can be readily extended by triangle inequality.*

Proof. By letting $\dot{Y} = [\dot{X}_1, ..., \dot{X}_n]$, $Y = [X_1, ..., X_n]$, $D = [\tau_1, ..., \tau_n]$, $W = [\varepsilon_1, ..., \varepsilon_n]$, we can write

$$\dot{Y} = \mathrm{sign}(RY + D) = \mathrm{sign}(R\theta\eta^T + RW + D).$$

The differences from Theorem 1 are that the center becomes $\theta_R := R\theta$, and that the noise matrix becomes $W_R := RW = [R\varepsilon_1, \cdots, R\varepsilon_n]$, which has independent columns but may not have independent rows. By Lemma 17, we assume that $\mu(\theta_R) \le \sqrt{\frac{3\log p}{p}}$ holds with the promised probability. This proof will follow that of Theorem 1 with necessary adaptations.

Step 1. Checking Steps 1–3 in Appendix A. We now re-iterate the proof of Theorem 1 to reach the following: with probability at least $1 - r_{n,\lambda}^{-1/2} - 2(np)^{-\nu} - 3p^{-1/4}$, we have that

$$\ell(\hat{\eta}^t, \eta) \le \frac{C'}{n^{3/2}} + C'' r_{n,\lambda}^{5/2} P^*, \tag{66}$$

where P^* is defined as

$$P^* := \mathbb{P}\left(\left\langle \theta+\delta_1, \theta+\frac{1}{n-1}\sum_{j=2}^{n}\delta_j\right\rangle < \frac{C_1\|\theta\|_2^2}{r_{n,\lambda}}\right) \tag{67}$$

where $\delta_1, \cdots, \delta_n$ are i.i.d. random vectors in $\mathbb{R}^p$ having the same distribution as $Q_{2\lambda}(\theta_R + R\varepsilon_i + \tau_i) - \theta_R$. In fact, we only need to modify the concentration bounds, as detailed in the following:

- Note that we now have $\tilde{Y} = Q_{2\lambda}(R\theta\eta^T + RW + D) = R\theta\eta^T + U$ and thus $U = Q_{2\lambda}(R\theta\eta^T + RW + D) - R\theta\eta^T$.
- Since U has independent columns of $O(\lambda)$ sub-Gaussian norms, the argument in Equations (23)–(24) for bounding $\|U^T\theta\|_2$ remains valid.
- Since U has possibly correlated rows, Equation (25) via Lemma 14 is no longer valid, and we should use Lemma 18 to achieve the same bound, by the decomposition

$$U = \underbrace{Q_{2\lambda}(R\theta\eta^T + RW + D) - (R\theta\eta^T + RW)}_{:=U_1} + RW := U_1 + RW \tag{68}$$

  where, given $R \in \mathbb{O}(p)$, W has i.i.d. entries of $O(\sigma)$ sub-Gaussian norms, and $U_1|W$ also has i.i.d. entries of $O(\lambda)$ sub-Gaussian norms.
- It is evident that (53) remains true for $\|RW\|_{\max}$.

Step 2. Checking Steps 4–5 in Appendix A. The additional challenge is that $\delta_1, ..., \delta_n$ may not have independent entries. Analogously to (68), we decompose $\{\delta_j\}_{j=1}^n$ as

$$\delta_j = \underbrace{Q_{2\lambda}(\theta_R + R\varepsilon_j + \tau_j) - (\theta_R + R\varepsilon_j)}_{:=\delta_j'} + R\varepsilon_i \tag{69}$$

and separately treat δ_j' and $R\varepsilon_i$. It turns out that the adaptation of Equations (55)–(57) is simple, and we still have

$$\hat{B} := \left\{\left|\langle\theta, \frac{1}{n-1}\sum_{j=2}^{n}\delta_j\rangle\right| \le C\lambda\|\theta\|_2\sqrt{\frac{\log n}{n}}\right\}$$

holds with probability at least $1 - \frac{1}{n^2}$. By Lemma 19, we also have that

$$\hat{B}_1 := \left\{\left\|\frac{1}{n-1}\sum_{j=2}^{n}\delta_j\right\|_2^2 \le (1+O(\frac{1}{n}))\frac{p\lambda^2}{n} + O(\frac{1}{\sqrt{n}})\|\theta\|_2^2\right\}$$

holds with probability at least $1 - 6\exp(-10r_{n,\lambda}^2)$. On the event $\hat{B} \cap \hat{B}_1$, following the derivations of Equations (60) and (61), we have

$$\left\|\theta + \frac{1}{n-1}\sum_{j=2}^{n}\delta_j\right\|_2^2 \le \left(1 + O\Big(\sqrt{\frac{\log n}{n}}\Big)\right)\|\theta\|_2^2 + \left(1 + O\left(\frac{1}{\sqrt{n}}\right)\right)\frac{p\lambda^2}{n},$$

and thus the event defining P^* in Equation (67) implies

$$\lambda^{-1}\langle\delta_1, \hat{\xi}\rangle \le -(1 - O(r_{n,\lambda}^{-1}) - O\Big((\frac{\log n}{n})^{1/4}\Big)r_{n,\lambda}$$

for a vector $\hat{\xi} \in \mathbb{S}^{p-1}$ independent of δ_1; therefore,

$$P^* \le \mathbb{P}\left(\frac{\langle\delta_1, \hat{\xi}\rangle}{\lambda} \le -\Big(1 - O(r_{n,\lambda}^{-1}) - O((\frac{\log n}{n})^{1/4})\Big)r_{n,\lambda}\right),$$

and further using $\delta_1 = \delta_1' + R\varepsilon_1$ in Equation (69) yields

$$\begin{aligned} P^* &\le \mathbb{P}\left(\frac{\langle\delta_1', \hat{\xi}\rangle}{\lambda} \le -\left(1 - O(r_{n,\lambda}^{-1}) - O((\frac{\log n}{n})^{1/4}) - \frac{1}{\log^{1/4}(np)}\right)r_{n,\lambda}\right) \\ &\quad + \mathbb{P}\left(\frac{\langle\varepsilon_1, R^T\hat{\xi}\rangle}{\lambda} \le -\frac{r_{n,\lambda}}{\log^{1/4}(np)}\right) + \frac{1}{n^2} + 6\exp(-10r_{n,\lambda}^2) \\ &:= P^{**} + P_2^* + \frac{1}{n^2} + 6\exp(-10r_{n,\lambda}^2). \end{aligned} \tag{70}$$

By the sub-Gaussian tail bound, $P_2^* \le \exp(-\sqrt{\log(np)} \cdot r_{n,\lambda}^2)$, which is negligible. To obtain sharp bound on P^{**}, we need a reduction to the one-bit case. As a surrogate of

$$\delta_1' = Q_{2\lambda}(\theta_R + R\varepsilon_1 + \tau_1) - (\theta_R + R\varepsilon_1),$$

we introduce $\tilde{\delta}_1' = \lambda \operatorname{sign}(\theta_R + R\varepsilon_1 + \tau_1) - \mathbb{E}_{\tau_1}[\lambda \operatorname{sign}(\theta_R + R\varepsilon_1 + \tau_1)]$. We invoke Lemma 20 to reach

$$P^{**} \le \mathbb{P}\Big(\frac{\langle \tilde{\delta}_1', \hat{\xi}\rangle}{\lambda} \le -(1-\epsilon_{n,\lambda})r_{n,\lambda}\Big) + \frac{1}{n^{1+\nu}p^{\nu}}.$$

Furthermore, by conditioning on ε_1 and using the randomness of τ_1 (so that the entries of $\tilde{\delta}_1'$ are independent, zero-mean), we can invoke Hoeffding's inequality as in Equation (64) to reach

$$\mathbb{P}\bigg(\frac{\langle \tilde{\delta}_1', \hat{\xi}\rangle}{\lambda} \le -(1-\epsilon_{n,\lambda})r_{n,\lambda}\bigg) \le 2\exp\Big(-\frac{(1-\epsilon_{n,\lambda})^2 r_{n,\lambda}^2}{2}\Big) = 2\exp\Big(-\frac{1}{2}(1-\epsilon_{n,\lambda}')r_{n,\lambda}^2\Big)$$

for some $\epsilon_{n,\lambda}' = O(r_{n,\lambda}^{-1} + (\frac{\log n}{n})^{1/4} + \frac{1}{\log^{1/4}(np)})$. Substituting this and $P_2^* \le \exp(-\sqrt{\log(np)} \cdot r_{n,\lambda}^2)$ into (70) yields

$$P^* \le 3\exp\Big(-\frac{1}{2}(1-\epsilon_{n,\lambda}')r_{n,\lambda}^2\Big) + \frac{1}{n^2} + 6\exp(-10r_{n,\lambda}^2) + \frac{1}{n^{1+\nu}p^{\nu}}.$$

Further combining with (66), along with $r_{n,\lambda} \lesssim \min\{n^{2/5}, (np)^{2\nu/5}\}$ and some simple reasoning used in the proof of Theorem 1, yields the claim. □

D Proof of Corollary 2 (Exact Recovery without Spikiness Condition)

Proof. In view of (11) and (12), we have $\|\theta\|_2^2 \gtrsim \lambda^2(\log n + \sqrt{\frac{p\log n}{n}})$. Combining with $r_{n,\lambda} = \frac{\|\theta\|_2^2/\lambda^2}{\sqrt{\|\theta\|_2^2/\lambda^2 + p/n}} \asymp \min\{\frac{\|\theta\|_2}{\lambda}, \frac{\|\theta\|^2}{\lambda^2}\sqrt{\frac{n}{p}}\}$, we then have

$$\frac{1}{r_{n,\lambda}} \lesssim \max\{\frac{\lambda}{\|\theta\|_2}, \frac{\lambda^2}{\|\theta\|_2^2}\sqrt{\frac{p}{n}}\} = O(\frac{1}{\sqrt{\log n}}).$$

As a result, Theorem 2 yields that for some $\epsilon_{n,\lambda} = O(\log^{-1/4} n)$ and for any $t \ge \lceil 3\log_4 n \rceil$, it holds with the promised probability that

$$\ell(\hat{\eta}^t, \eta) \le \exp\Big(-\frac{(1-\epsilon_{n,\lambda})r_{n,\lambda}^2}{2}\Big).$$

This leads to the desired $\hat{\eta}^t = \pm\eta$ if

$$\exp\Big(-\frac{(1-\epsilon_{n,\lambda})r_{n,\lambda}^2}{2}\Big) < \frac{1}{n}$$

$$\iff \|\theta\|_2^2 > \frac{1}{1-\epsilon_{n,\lambda}}\lambda^2\bigg(1 + \sqrt{1 + \frac{(1-\epsilon_0)2p}{n\log n}}\bigg)\log n.$$

► by calculations parallel to Equation (65)

This can be implied by

$$\|\theta\|_2^2 \ge (1+\epsilon)\lambda^2\bigg(1 + \sqrt{1 + \frac{2p}{n\log n}}\bigg)\log n, \quad \text{for some } \epsilon = O(\log^{-1/4} n),$$

which is exactly what we assume in Equations (11) and (12).

It remains to show that the condition (12) can be implied by the more explicit condition (13) under $\min\{n, p, r_{n,\lambda}\} \ge C_1(\nu)$ and $p \ge C_2(\nu)\log n \log\log n$ for some sufficiently large constants $C_1(\nu), C_2(\nu)$ depending on ν only. For some sufficiently large $C_0(\nu)$, we use the inequality

$$(A_1 + A_2)^2 = A_1^2 + A_2^2 + 2A_1A_2 \le (1 + C_0(\nu))A_1^2 + (1 + \frac{1}{C_0(\nu)})A_2^2$$

to achieve

$$\begin{aligned}
&\left(\sqrt{\frac{3\log p}{p}}\|\theta\|_2+\sigma\sqrt{2(1+\nu)\log(np)}\right)^2\\
&\le \frac{3[1+C_0(\nu)]\log p\|\theta\|_2^2}{p}+\Big[1+\frac{1}{C_0(\nu)}\Big](1+\nu)2\sigma^2\log(np)\\
&\le \frac{6C_0(\nu)\log p\|\theta\|_2^2}{p}+(1+1.5\nu)2\sigma^2\log(np).
\end{aligned}$$

► by choosing $C_0(\nu)$ large enough to ensure this

In turn a sufficient condition for Equation (12) is

$$\|\theta\|_2^2\ge(1+O(\log^{-1/4}n))\left(\frac{6C_0(\nu)\log p\|\theta\|_2^2}{p}+(1+1.5\nu)2\sigma^2\log(np)\right)\left(1+\sqrt{1+\frac{2p}{n\log n}}\right)\log n,$$

and by rearranging, this is implied by

$$\begin{aligned}
&\left(1-\tilde{C}_0(\nu)\frac{\log n\log p}{p}\big[1+\sqrt{\frac{p}{n\log n}}\big]\right)\|\theta\|_2^2\\
&\qquad\ge\big(1+O(\log^{-1/4}n)\big)(1+1.5\nu)2\sigma^2\log(np)\left(1+\sqrt{1+\frac{2p}{n\log n}}\right)\log n
\end{aligned}\tag{71}$$

for some constant $\tilde{C}_0(\nu)$ depending on ν only. Note that (71) can be written as

$$\|\theta\|_2^2\ge\underbrace{\frac{\big(1+O(\log^{-1/4}n)\big)(1+1.5\nu)}{\big(1-\tilde{C}_0(\nu)\frac{\log n\log p}{p}\big[1+\sqrt{\frac{p}{n\log n}}\big]\big)}}_{F_1}2\sigma^2\left(1+\sqrt{1+\frac{2p}{n\log n}}\right)\log(np)\log n$$

All that remains is to show that the leading factor F_1 is bounded by $1+2\nu$. This is ensured by the following observations: (i) $n\ge C_1(\nu)$ with large enough $C_1(\nu)$ ensures that $O(\log^{-1/4}n)\le c_1(\nu)$ for some small enough $c_1(\nu)$; (ii) $p\ge C_2(\nu)\log n\log\log n$ with large enough $C_2(\nu)$ implies that $\tilde{C}_0(\nu)\frac{\log n\log p}{p}\le c_2(\nu)$ for some small enough $c_2(\nu)$; (iii) $n,p\ge C_1(\nu)$ for large enough $C_1(\nu)$ ensures that

$$\frac{\tilde{C}_0(\nu)\log n\log p}{p}\sqrt{\frac{p}{n\log n}}=\tilde{C}_0(\nu)\sqrt{\frac{\log^2 p\log n}{np}}\le c_3(\nu)\quad\text{for some small enough } c_3(\nu).$$

The proof is complete. □

E Proof of Theorem 3 (Minimax Lower Bound)

Proof. We fix $\theta=\frac{\|\theta\|_2}{\sqrt{p}}\mathbf{1}_p$ and place an independent Rademacher prior on the label η, meaning that we let $\eta_1,\cdots,\eta_n$ be i.i.d. uniformly distributed over $\{-1,1\}$. Let $\mathbb{E}_\pi$ be the expectation on the prior, then we have

$$\inf_{\hat{\eta}}\sup_{\theta,\eta}\mathbb{E}\big[\ell(\hat{\eta},\eta)\big]\ge\inf_{\hat{\eta}}\mathbb{E}_\pi\mathbb{E}\big[\ell(\hat{\eta},\eta)\big].$$

We let $T\subset[n]$ be of size $\lfloor n/2\rfloor+1$, then we can handle the sign flipping associated with $\ell(\hat{\eta},\eta)$ by argument in [16], reaching

$$\begin{aligned}
\inf_{\hat{\eta}}\mathbb{E}_\pi\mathbb{E}\big[\ell(\hat{\eta},\eta)\big]&\ge\frac{c}{|T^c|}\sum_{i\in T^c}\inf_{\tilde{\eta}^i}\mathbb{E}_\pi\mathbb{E}|\tilde{\eta}^i(\dot{Y},(\eta_j)_{j\in T})-\eta_i|\\
&\ge\frac{c}{|T^c|}\sum_{i\in T^c}\inf_{\tilde{\eta}^i}\mathbb{E}_\pi\mathbb{E}|\tilde{\eta}^i(\dot{Y},(\eta_j)_{j\ne i})-\eta_i|\\
&\qquad\text{► since }(\eta_j)_{j\ne i}\text{ contains more information than }(\eta_j)_{j\in T}
\end{aligned}$$

for some constant c. Moreover, since the pairs $(\dot{Y}_j, \eta_j)$ are independently distributed, we have

$$\inf_{\hat{\eta}} \mathbb{E}_\pi \mathbb{E}\big[\ell(\hat{\eta}, \eta)\big] \geq c \cdot \inf_{\bar{\eta}_j} \mathbb{E}_\pi \mathbb{E}|\bar{\eta}_j(\dot{X}_j) - \eta_j|, \tag{72}$$

where η_i is a Rademacher variable, and $\bar{\eta}_i$ is a measurable function of $\dot{X}_i$ for estimating η_i. We now let

$$P_1 = \mathbb{P}\Big(\operatorname{sign}\Big(\frac{\|\theta\|_2}{\sqrt{p}} + \varepsilon_{ij} + \tau_{ij} \Big) = 1 \Big),$$
$$P_{-1} = \mathbb{P}\Big(\operatorname{sign}\Big(\frac{\|\theta\|_2}{\sqrt{p}} + \varepsilon_{ij} + \tau_{ij} \Big) = -1 \Big),$$

which satisfy $P_1 > \frac{1}{2} > P_{-1}$ by the symmetry of ε_{ij} and τ_{ij}. Then, it is easy to write the likelihood of $\dot{X}_i$ under $\eta_i = 1$ and $\eta_i = -1$ as

$$f_1(\dot{Y}_j) = P_1^{|\{i \in [p] : \dot{Y}_{ij} = 1\}|} P_{-1}^{p - |\{i \in [p] : \dot{Y}_{ij} = 1\}|},$$
$$f_{-1}(\dot{Y}_j) = P_1^{p - |\{i \in [p] : \dot{Y}_{ij} = 1\}|} P_{-1}^{|\{i \in [p] : \dot{Y}_{ij} = 1\}|}.$$

Therefore, by Bayesian decision rule, the selector

$$\eta_k^* = \mathbf{1}(|\{i \in [p] : \dot{Y}_{ij} = 1\}| \geq \frac{p}{2}) - \mathbf{1}(|\{i \in [p] : \dot{Y}_{ij} = 1\}| < \frac{p}{2}) = \operatorname{sign}(\dot{Y}_j^T \mathbf{1})$$

attains the minimum of the right-hand side of (72). Then, it is not hard to see that

$$\mathbb{E}_\pi \mathbb{E}|\bar{\eta}_j(\dot{X}_j) - \eta_j| = 2\mathbb{P}\Big(\text{Binomial}(p, P_{-1}) > \frac{p}{2} \Big) \geq \frac{1}{\sqrt{2p}} \exp\Big(- p D_{KL}\Big(\frac{1}{2} \| P_{-1} \Big) \Big) \tag{73}$$

where the last inequality is due to the anti-concentration bound for binomial variable (cf. Lemma 1). We now lower bound P_{-1} as follows:

$$\begin{aligned}
P_{-1} &= \mathbb{P}\Big(\frac{\|\theta\|_2}{\sqrt{p}} + \varepsilon_{ij} + \tau_{ij} < 0 \Big) \\
&\geq \mathbb{E}\left(\mathbf{1}\Big(\Big| \frac{\|\theta\|_2}{\sqrt{p}} + \varepsilon_{ij} \Big| \leq \lambda \Big) \mathbf{1}\Big(\tau_{ij} < -\varepsilon_{ij} - \frac{\|\theta\|_2}{\sqrt{p}} \Big) \right) \\
&= \mathbb{E}\left(\mathbf{1}\Big(\Big| \frac{\|\theta\|_2}{\sqrt{p}} + \varepsilon_{ij} \Big| \leq \lambda \Big) \cdot \frac{\lambda - \varepsilon_{ij} - p^{-1/2}\|\theta\|_2}{2\lambda} \right) \\
&\quad \blacktriangleright \text{ by the randomness of } \tau_{ij} \sim \mathrm{Unif}[-\lambda, \lambda] \\
&= \mathbb{E}\left(\frac{\lambda - \varepsilon_{ij} - p^{-1/2}\|\theta\|_2}{2\lambda} \right) - \mathbb{E}\left(\frac{\lambda - \varepsilon_{ij} - p^{-1/2}\|\theta\|_2}{2\lambda} \mathbf{1}\Big(\Big| \frac{\|\theta\|_2}{\sqrt{p}} + \varepsilon_{ij} \Big| > \lambda \Big) \right) \\
&\geq \frac{1}{2} - \frac{\|\theta\|_2}{2\lambda\sqrt{p}} - E_1 - E_2,
\end{aligned} \tag{74}$$

where we let

$$E_1 := \mathbb{E}\left(\left| \frac{\lambda - p^{-1/2}\|\theta\|_2}{2\lambda} \right| \mathbf{1}\Big(\Big| \frac{\|\theta\|_2}{\sqrt{p}} + \varepsilon_{ij} \Big| > \lambda \Big) \right),$$
$$E_2 := \mathbb{E}\left(\frac{|\varepsilon_{ij}|}{2\lambda} \mathbf{1}\Big(\Big| \frac{\|\theta\|_2}{\sqrt{p}} + \varepsilon_{ij} \Big| > \lambda \Big) \right).$$

By standard Gaussian tail bound,

$$E_1 \leq \frac{1}{2}\mathbb{P}\Big(|\varepsilon_{ij}| > \lambda - \frac{\|\theta\|_2}{\sqrt{p}} \Big) \leq \exp\Big(- \frac{(\lambda - p^{-1/2}\|\theta\|_2)^2}{2\sigma^2} \Big).$$

Next, we have

$$\begin{aligned}
E_2 &\leq \frac{1}{2\lambda}\mathbb{E}\big(|\varepsilon_{ij}| \mathbf{1}(|\varepsilon_{ij}| > \lambda - \|\theta\|_2/\sqrt{p}) \big) \\
&= \frac{\sigma}{2\lambda} \int_{\sigma^{-1}(\lambda - \frac{\|\theta\|_2}{\sqrt{p}})}^{\infty} \sqrt{\frac{2}{\pi}} s \exp(-\frac{s^2}{2}) \, \mathrm{d}s \\
&= \frac{\sigma}{\lambda\sqrt{2\pi}} \exp\Big(- \frac{(\lambda - p^{-1/2}\|\theta\|_2)^2}{2\sigma^2} \Big).
\end{aligned}$$

Substituting these bounds into (74) yields

$$P_{-1} \geq \frac{1}{2} - \Big(\frac{\|\theta\|_2}{2\lambda\sqrt{p}} + \Big(1 + \frac{\sigma}{\lambda\sqrt{2\pi}}\Big) \exp\Big(- \frac{(\lambda - p^{-1/2}\|\theta\|_2)^2}{2\sigma^2}\Big)\Big) := \frac{1}{2} - \xi_0,$$

where we let

$$\xi_0 := \frac{\|\theta\|_2}{2\lambda\sqrt{p}} + \Big(1 + \frac{\sigma}{\lambda\sqrt{2\pi}}\Big) \exp\Big(- \frac{(\lambda - p^{-1/2}\|\theta\|_2)^2}{2\sigma^2}\Big),$$

and under $\lambda \gtrsim \sigma + p^{-1/2}\|\theta\|_2$ and $\log(\frac{2\lambda\sqrt{p}}{\|\theta\|_2}) \leq \frac{\lambda^2}{8\sigma^2}$, we have

$$\xi_0 \leq \frac{\|\theta\|_2}{2\lambda\sqrt{p}} + 2\exp\Big(- \frac{\lambda^2}{4\sigma^2}\Big) = \frac{\|\theta\|_2}{2\lambda\sqrt{p}}\Big(1 + 2\exp\Big(- \frac{\lambda^2}{8\sigma^2}\Big)\Big). \tag{75}$$

Note that ξ_0 is small enough under the assumption $\lambda \geq C(\frac{\|\theta\|_2}{\sqrt{p}} + \sigma)$. Continuing from (73),

$$\begin{aligned}
\mathbb{E}_\pi \mathbb{E}\,|\bar{\eta}_j(\dot{X}_j) - \eta_j| &\geq \frac{1}{\sqrt{2p}} \exp\Big(- pD_{KL}\Big(\frac{1}{2}\Big\|\frac{1}{2} - \xi_0\Big)\Big) \\
&= \frac{1}{\sqrt{2p}} \exp\Big(\frac{p}{2}\log(1 - 4\xi_0^2)\Big) \\
&\geq \frac{1}{\sqrt{2p}} \exp\big(- 2(1 + o(1))p\xi_0^2\big),
\end{aligned}$$

where in the last inequality we use $\log(1 - 4\xi_0^2) \geq -4(1 + o(1))\xi_0^2$ for some $o(1) \to 0$ as $\xi_0 \to 0$, which is ensured by $\frac{\lambda}{p^{-1/2}\|\theta\|_2 + \sigma} \to \infty$. Further using $\xi_0 \leq (1 + o(1))\frac{\|\theta\|_2}{2\lambda\sqrt{p}}$ from (75), we arrive at the desired claim. □

F Technical Lemmas

Lemma 1 (Anti-concentration of Binomial variable (e.g., [3]))**.** *Let $X \sim \mathrm{Binomial}(N, q)$, then for any $k = 0, 1, \cdots, N$ we have $\mathbb{P}\big(X = k\big) \geq \frac{1}{\sqrt{2N}} \exp\big(- N \cdot D_{\mathrm{KL}}(\frac{k}{N}\big\|q)\big)$ where $D_{\mathrm{KL}}(q\|q') := q\log(\frac{q}{q'}) + (1 - q)\log(\frac{1-q}{1-q'})$.*

Lemma 2 (Sub-Gaussian tail bound)**.** *Under Assumption 1, for any $v \in \mathbb{R}^p$ with $\|v\|_2 = 1$, we have*

$$\max\Big\{\mathbb{P}(v^T\varepsilon_i \geq t), \mathbb{P}(v^T\varepsilon_i \leq -t)\Big\} \leq \exp\Big(- \frac{\sigma^2 t^2}{2}\Big), \quad \forall t \geq 0.$$

Proof. We only bound $\mathbb{P}(v^T\varepsilon_i \geq t)$ since $\mathbb{P}(v^T\varepsilon_i \leq -t) = \mathbb{P}((-v)^T\varepsilon_i \geq t)$. For $t = 0$ the bound is trivial. We suppose $t > 0$. Then for any $s > 0$,

$$\begin{aligned}
\mathbb{P}(v^T\varepsilon_i \geq t) &= \mathbb{P}\Big(\exp(sv^T\varepsilon_i) \geq \exp(st)\Big) \\
&\leq \exp(-st)\mathbb{E}\Big[\exp(sv^T\varepsilon_i)\Big] \\
&\quad \blacktriangleright \text{ by Markov's inequality} \\
&= \exp(-st)\mathbb{E}\Big[\prod_{j=1}^p \exp(sv_j\varepsilon_{ij})\Big] \\
&= \exp(-st)\prod_{j=1}^p \mathbb{E}\big[\exp(sv_j\varepsilon_{ij})\big] \\
&\quad \blacktriangleright \text{ by independence of entries of } \varepsilon_i \\
&\leq \exp(-st)\prod_{j=1}^p \exp\Big(\frac{\sigma^2 s^2 v_j^2}{2}\Big) = \exp\Big(- st + \frac{\sigma^2 s^2}{2}\Big).
\end{aligned}$$

Setting $s = \frac{t}{\sigma^2}$ yields the claim. □

Lemma 3. *[30, Proposition 2.6.1] Let $X_1, \cdots, X_N$ be independent, zero-mean variables such that $\|X_i\|_{\psi_2} < \infty$. Then for some universal constant C,*

$$\Big\| \sum_{i=1}^N X_i \Big\|_{\psi_2}^2 \leq C \sum_{i=1}^N \|X_i\|_{\psi_2}^2$$

Lemma 4. *[30, Lemma 2.7.7] Define the sub-exponential norm (or ψ_1-norm) of X as $\|X\|_{\psi_1} = \inf\{K > 0 : \mathbb{E}\exp(|X|/K) \leq 2\}$. Let X, Y be sub-Gaussian, then XY is sub-exponential: $\|XY\|_{\psi_1} \leq \|X\|_{\psi_2}\|Y\|_{\psi_2}$.*

Lemma 5. *[30, Exercise 2.7.10] For some absolute constant C, $\|X - \mathbb{E}X\|_{\psi_1} \leq C\|X\|_{\psi_1}$.*

Lemma 6. *[30, Theorem 2.2.6] Let $X_1, ..., X_N$ be independent random variables. Assume that $X_i \in [m_i, M_i]$ for every i. Then, for any $t > 0$ we have*

$$\mathbb{P}\Big(\sum_{i=1}^N (X_i - \mathbb{E}X_i) \geq t \Big) \leq \exp\Big(- \frac{2t^2}{\sum_{i=1}^N (M_i - m_i)^2} \Big)$$

Lemma 7. *[30, Theorem 2.8.1] Let $X_1, ..., X_N$ be independent, zero-mean, sub-exponential random variables. Then for every $t \geq 0$, for some absolute constant c we have*

$$\mathbb{P}\left(\Big| \sum_{i=1}^N X_i \Big| \geq t \right) \leq 2\exp\left(-c \min\Big\{ \frac{t^2}{\sum_{i=1}^N \|X_i\|_{\psi_1}^2}, \frac{t}{\max_{1\leq i\leq N} \|X_i\|_{\psi_1}} \Big\} \right)$$

Lemma 8. *[2, Section 2.1] Let $a \in \mathbb{R}$ and $\tau \sim \mathrm{Unif}[-\lambda, \lambda]$, we have*

$$\mathbb{E}[\lambda \operatorname{sign}(a + \tau)] = T_\lambda(a) := \begin{cases} a\,, & \text{if } |a| \leq \lambda \\ \lambda \operatorname{sign}(a)\,, & \text{if } |a| > \lambda \end{cases}.$$

Lemma 9. *Let $H(Z) = Z - \operatorname{diag}(Z)$. For any $Z \in \mathbb{R}^{n\times n}$ we have $\|H(Z)\|_{op} \leq 2\|Z\|_{op}$.*

Proof. By triangle inequality,

$$\|H(Z)\|_{op} = \|Z - \operatorname{diag}(Z)\|_{op} \leq \|Z\|_{op} + \|\operatorname{diag}(Z)\|_{op} \leq 2\|Z\|_{op}.$$

□

Lemma 10. *[23, Lemma 2] For a random matrix W with independent, zero-mean columns, we have*

$$\|H(W^TW)\|_{op} \leq 2\|W^TW - \mathbb{E}(W^TW)\|_{op}$$

Lemma 11. *[23, Lemma 5] For any $x \in \{-1, 1\}^n$ and $y \in \mathbb{R}^n$, we have*

$$\frac{\|\operatorname{sign}(y) - x\|_1}{n} \leq 2\Big\| y - \frac{x}{\sqrt{n}} \Big\|_2^2.$$

Lemma 12. *Let $Q_{2\lambda}(a) = 2\lambda(\lfloor \frac{a}{2\lambda} \rfloor + \frac{1}{2})$ for $a \in \mathbb{R}$ and $\tau \sim \mathrm{Unif}[-\lambda, \lambda]$. Then for any $a \in \mathbb{R}$, we have*

$$\mathbb{E}\big[Q_{2\lambda}(a + \tau)\big] = a, \quad |Q_{2\lambda}(a + \tau) - a| \leq 2\lambda \quad \textit{and} \quad \|Q_{2\lambda}(a + \tau) - a\|_{\psi_2} = O(\lambda)$$

Proof. $\mathbb{E}\big[Q_{2\lambda}(a + \tau)\big] = a$ can be found in [18]. For the second result, the definition of $Q_{2\lambda}(a)$ gives $\sup_{a\in\mathbb{R}} |Q_{2\lambda}(a) - a| \leq \lambda$, hence

$$|Q_{2\lambda}(a + \tau) - a| \leq |Q_{2\lambda}(a + \tau) - (a + \tau)| + |\tau| \leq 2\lambda.$$

Since $Q_{2\lambda}(a + \tau) - a$ is zero-mean and bounded by 2λ, $\|Q_{2\lambda}(a + \tau) - a\|_{\psi_2} = O(\lambda)$ immediately follows (e.g., [30, Section 2]). □

Lemma 13. *Let $x \in \mathbb{R}^n$ be a random vector with independent entries with sub-Gaussian norms bounded by A. Then for some absolute constant C, we have*

$$\mathbb{P}\big(\|x\|_2 \leq C \cdot A\sqrt{n}\big) \geq 1 - \exp(-2n).$$

Proof. Let x_i be the i-th entry of x, then $\|x\|_2^2 = \sum_{j=1}^n x_j^2$, where $x_1^2, ..., x_n^2$ are sub-exponential variables with sub-exponent norm bounded by (see Lemma 4)

$$\|x_i^2\|_{\psi_1} \le \|x_i\|_{\psi_2}^2 \le A^2.$$

By centering (c.f., Lemma 5), we have $\|x_i^2 - \mathbb{E}x_i^2\|_{\psi_1} = O(A^2)$. Therefore, Bernstein's inequality (cf. Lemma 7) gives

$$\mathbb{P}\Big(\big|\|x\|_2^2 - \mathbb{E}\|x\|_2^2\big| \ge t\Big) \le 2\exp\left(-c\min\left\{\frac{t^2}{nA^4}, \frac{t}{A^2}\right\}\right).$$

Setting $t = C'nA^2$ with large enough C' yields

$$\mathbb{P}\Big(\|x\|_2^2 \ge C'nA^2 + \mathbb{E}\|x\|_2^2\Big) \le 2\exp(-2n).$$

Moreover, by the sub-Gaussianity of x_i we have $\mathbb{E}x_i^2 = O(A^2)$, and therefore $\mathbb{E}\|x\|_2^2 = O(nA^2)$, which completes the proof. □

Lemma 14. *Suppose that $U \in \mathbb{R}^{p\times n}$ have independent rows with sub-Gaussian norms bounded by A, then for some absolute constant C we have*

$$\mathbb{P}\left(\|U^TU - \mathbb{E}(U^TU)\|_{op} \ge CA^2n\max\big\{1, \sqrt{\frac{p}{n}}\big\}\right) \le 2\exp(-2n).$$

Proof. To control $\|U^TU - \mathbb{E}(U^TU)\|_{op} = \sup_{v\in\mathbb{S}^{n-1}}\big[\|Uv\|_2^2 - \mathbb{E}\|Uv\|_2^2\big]$, we start with a discretization (see, e.g., [30, Exercise 4.4.3]): let $N_{1/3}$ be a $1/3$-net of $\mathbb{S}^{n-1}$ of cardinality bounded by 9^n, we have

$$\|U^TU - \mathbb{E}(U^TU)\|_{op} \le 3\sup_{v\in N_{1/3}}\big[\|Uv\|_2^2 - \mathbb{E}\|Uv\|_2^2\big]. \tag{76}$$

For any $v \in N_{1/3}$, we note that $Uv = (w_1, ..., w_p)^T \in \mathbb{R}^p$ has independent entries with sub-Gaussian norms bounded by A, and therefore $\|w_i^2\|_{\psi_1} \le \|w_i\|_{\psi_2}^2 = O(A^2)$ by Lemma 4. We now invoke Bernstein's inequality (see Lemma 7) to arrive at

$$\begin{aligned}\mathbb{P}\Big(\Big|\|Uv\|_2^2 - \mathbb{E}\|Uv\|_2^2\Big| \ge t\Big) &= \mathbb{P}\Big(\Big|\sum_{i=1}^p[w_i^2 - \mathbb{E}w_i^2]\Big| \ge t\Big)\\ &\le 2\exp\Big(-c\min\Big\{\frac{t^2}{pA^4}, \frac{t}{A^2}\Big\}\Big), \qquad \forall t \ge 0.\end{aligned}$$

Taking a union bound, along with Equation (76), yields

$$\mathbb{P}\Big(\|U^TU - \mathbb{E}(U^TU)\|_{op} \ge t\Big) \le 2\exp\Big(n\log 9 - c'\min\Big\{\frac{t^2}{pA^4}, \frac{t}{A^2}\Big\}\Big), \quad \forall t \ge 0.$$

Setting $t = C'A^2n\max\{1, \sqrt{\frac{p}{n}}\}$ with large enough C', we obtain that

$$\|U^TU - \mathbb{E}(U^TU)\|_{op} \le C'A^2n\max\{1, \sqrt{\frac{p}{n}}\}$$

with probability at least $1 - 2\exp(-2n)$, as claimed. □

Lemma 15. *In the proof of Theorem 1 in Appendix A, we have*

$$\mathbb{P}\left(\Big\|\frac{1}{n-1}\sum_{j=2}^n \delta_j\Big\|_2^2 < \frac{p\lambda^2}{n}\Big(1 + \frac{C}{n}\Big) + \frac{C'\|\theta\|_2^2}{\sqrt{n}}\right) \ge 1 - 2\exp(-10r_{n,\lambda}^2)$$

Proof. Recall that $\delta_j \overset{d}{=} Q_{2\lambda}(\theta + \varepsilon_j + \tau_j) - \theta$ for $j \in [n]$, and notice that the entries of $\delta_2, \cdots, \delta_n$ are independent, and entries of $\frac{1}{n-1}\sum_{j=2}^n \delta_j$ have $O(\frac{\lambda}{\sqrt{n}})$ sub-Gaussian norms:

$$\max_{i\in[p]}\Big\|\frac{1}{n-1}\sum_{j=2}^n \delta_{ji}\Big\|_{\psi_2} \lesssim \frac{\lambda}{\sqrt{n}}.$$

Therefore, by Lemma 4 we have

$$\max_{i\in[p]}\left\|\Big(\frac{1}{n-1}\sum_{j=2}^{n}\delta_{ji}\Big)^2\right\|_{\psi_1}\lesssim\frac{\lambda^2}{n},$$

and thus Bernstein's inequality (Lemma 7) gives, for any $t\geq 0$, that

$$\mathbb{P}\bigg(\bigg\|\frac{1}{n-1}\sum_{j=2}^{n}\delta_j\bigg\|_2^2\geq\mathbb{E}\bigg\|\frac{1}{n-1}\sum_{j=2}^{n}\delta_j\bigg\|_2^2+t\bigg)\leq 2\exp\left(-c\min\left\{\frac{n^2t^2}{p\lambda^4},\frac{nt}{\lambda^2}\right\}\right). \tag{77}$$

We set $t=C\max\{\frac{\lambda^2\sqrt{p}\cdot r_{n,\lambda}}{n},\frac{\lambda^2 r_{n,\lambda}^2}{n}\}$ with large enough C to establish that

$$\begin{aligned}&\mathbb{P}\bigg(\bigg\|\frac{1}{n-1}\sum_{j=2}^{n}\delta_j\bigg\|_2^2\geq\mathbb{E}\bigg\|\frac{1}{n-1}\sum_{j=2}^{n}\delta_j\bigg\|_2^2+C\max\left\{\frac{\lambda^2\sqrt{p}\cdot r_{n,\lambda}}{n},\frac{\lambda^2 r_{n,\lambda}^2}{n}\right\}\bigg)\\&\leq 2\exp(-10r_{n,\lambda}^2).\end{aligned} \tag{78}$$

It remains to compute the expectation. In light of the i.i.d. zero-mean δ_j's,

$$\mathbb{E}\bigg\|\frac{1}{n-1}\sum_{j=2}^{n}\delta_j\bigg\|_2^2=\sum_{i=1}^{p}\mathbb{E}\bigg|\frac{1}{n-1}\sum_{j=2}^{n}\delta_{ji}\bigg|^2=\frac{1}{n-1}\sum_{i=1}^{p}\mathbb{E}[\delta_{ji}^2]. \tag{79}$$

Furthermore, we have

$$\begin{aligned}\mathbb{E}[\delta_{ji}^2]&=\mathbb{E}\Big|Q_{2\lambda}(\theta_i+\varepsilon_{ji}+\tau_{ji})-\theta_i\Big|^2=\mathbb{E}|Q_{2\lambda}(\theta_i+\varepsilon_{ji}+\tau_{ji})|^2-\theta_i^2\\&\quad\blacktriangleright\text{ by }\mathbb{E}[Q_{2\lambda}(\theta_i+\varepsilon_{ji}+\tau_{ji})]=\theta_i\text{ (see Lemma 12)}\\&\leq\lambda^2+\underbrace{\mathbb{E}\Big(|Q_{2\lambda}(\theta_i+\varepsilon_{ji}+\tau_{ji})|^2\mathbf{1}(|\varepsilon_{ji}|\geq\sigma\sqrt{2(1+\nu)\log(np)})\Big)}_{:=T_1}-\theta_j^2.\end{aligned} \tag{80}$$

$$\blacktriangleright\text{ by }\mathbb{E}\Big(|Q_{2\lambda}(\theta_i+\varepsilon_{ji}+\tau_{ji})|^2\mathbf{1}(|\varepsilon_{ji}|<\sigma\sqrt{2(1+\nu)\log(np)})\Big)\leq\lambda^2$$

which holds because under (6), $|\theta_i+\varepsilon_{ji}+\tau_{ji}|\leq 2\lambda$, and $Q_{2\lambda}(a)=\pm\lambda$ when $|a|\leq 2\lambda$

Moreover,

$$\begin{aligned}T_1&\leq\mathbb{E}\Big((3\lambda+|\varepsilon_{ji}|)^2\mathbf{1}(|\varepsilon_{ji}|\geq\sigma\sqrt{2(1+\nu)\log(np)})\Big)\\&\quad\blacktriangleright\text{ by }|Q_{2\lambda}(\theta_i+\varepsilon_{ji}+\tau_{ji})|\leq|\theta_i+\varepsilon_{ji}+2\lambda|\leq 3\lambda+|\varepsilon_{ji}|\\&\leq 18\lambda^2\mathbb{P}\big(|\varepsilon_{ji}|\geq\sigma\sqrt{2(1+\nu)\log(np)}\big)+2\mathbb{E}\big[\varepsilon_{ji}^2\mathbf{1}(|\varepsilon_{ji}|\geq\sigma\sqrt{2(1+\nu)\log(np)})\big]\\&\quad\blacktriangleright\text{ by }(A_1+A_2)^2\leq 2A_1^2+2A_2^2\\&\leq\frac{C'\lambda^2}{(np)^{1+\nu}}.\end{aligned} \tag{81}$$

$$\blacktriangleright\text{ by Lemma 2 and }\mathbb{E}\big[\varepsilon_{ji}^2\mathbf{1}(|\varepsilon_{ji}|\geq\sigma\sqrt{2(1+\nu)\log(np)})\big]\lesssim\frac{\lambda^2}{(np)^{1+\nu}}\text{ (see the following)}$$

We now explain $\mathbb{E}\big[\varepsilon_{ji}^2\mathbf{1}(|\varepsilon_{ji}|\geq\sigma\sqrt{2(1+\nu)\log(np)})\big]\lesssim\frac{\lambda^2}{(np)^{1+\nu}}$:

$$\begin{aligned}&\mathbb{E}\left[\varepsilon_{ji}^2\mathbf{1}(|\varepsilon_{ji}|\geq\sigma\sqrt{2(1+\nu)\log(np)})\right]\\&=\int_0^\infty\mathbb{P}\left(\varepsilon_{ji}^2\mathbf{1}(|\varepsilon_{ji}|\geq\sigma\sqrt{2(1+\nu)\log(np)})\geq t\right)dt\\&\quad\blacktriangleright\text{ by tail integral representation of expectation}\\&=\int_{2\sigma^2(1+\nu)\log(np)}^\infty\mathbb{P}\left(|\varepsilon_{ji}|^2\geq t\right)dt\\&\leq\int_{2\sigma^2(1+\nu)\log(np)}^\infty 2\exp\Big(-\frac{t}{2\sigma^2}\Big)\,dt\\&\quad\blacktriangleright\text{ by Lemma 2}\\&=\frac{4\sigma^2}{(np)^{1+\nu}}\leq\frac{4\lambda^2}{(np)^{1+\nu}}.\end{aligned}$$

Substituting (81) into (80) yields

$$\mathbb{E}[\delta_{ji}^2] \le \big(1+\frac{C'}{np}\big)\lambda^2 - \theta_j^2, \quad \forall i \in [p]$$

and substituting this into (79) yields

$$\mathbb{E}\bigg\|\frac{1}{n-1}\sum_{j=2}^{n}\delta_j\bigg\|_2^2 \le \frac{1}{n-1}\left[p\lambda^2\Big(1+\frac{C'}{np}\Big) - \|\theta\|_2^2\right]. \tag{82}$$

In light of $r_{n,\lambda} = \frac{\|\theta\|_2^2/\lambda^2}{\sqrt{\|\theta\|_2^2/\lambda^2 + p/n}} \asymp \min\{\frac{\|\theta\|_2}{\lambda}, \frac{\|\theta\|_2^2}{\lambda^2}\sqrt{\frac{n}{p}}\}$ and therefore

$$\frac{\lambda^2\sqrt{p}\cdot r_{n,\lambda}}{n} = O\Big(\frac{\|\theta\|_2^2}{\sqrt{n}}\Big) \qquad \text{and} \qquad \frac{\lambda^2 r_{n,\lambda}^2}{n} = O\Big(\frac{\|\theta\|_2^2}{n}\Big). \tag{83}$$

Substituting (82) and (83) into (78) completes the proof. □

Lemma 16. *In the proof of Theorem 1 in Appendix A, we have*

$$P^{**} \le \mathbb{P}\bigg(\langle\tilde{\delta}_1, \hat{\xi}\rangle < -\frac{(1-\epsilon_{n,\lambda})\|\theta\|_2^2}{\sqrt{\|\theta\|_2^2 + \frac{p\lambda^2}{n}}}\bigg) + \frac{1}{n^{1+\nu}p^{\nu}}$$

for some $\epsilon_{n,\lambda} = O(r_{n,\lambda}^{-1} + (\frac{\log n}{n})^{1/4})$.

Proof. Recall that $\delta_1 = Q_{2\lambda}(\theta+\varepsilon_1+\tau_1) - \theta$ and $\tilde{\delta}_1 = \lambda\,\mathrm{sign}(\theta+\varepsilon_1+\tau_1) - \mathbb{E}[\lambda\,\mathrm{sign}(\theta+\varepsilon_1+\tau_1)]$. We define the event

$$\hat{B}_2 := \Big\{\max_{i\in[p]}|\theta_i+\varepsilon_{1i}| \le \lambda\Big\}. \tag{84}$$

Under $\lambda \ge \|\theta\|_\infty + \sigma\sqrt{2(1+\nu)\log(np)}$, Lemma 2 gives

$$\mathbb{P}(\hat{B}_2) \ge \mathbb{P}\left(\max_{i\in[p]}|\varepsilon_{1i}| \le \sigma\sqrt{2(1+\nu)\log(np)}\right) \ge 1 - \frac{p}{(np)^{1+\nu}} = 1 - \frac{1}{n^{1+\nu}p^{\nu}}. \tag{85}$$

Moreover, for any $i \in [p]$,

$$\begin{aligned}
&|\mathbb{E}[\lambda\,\mathrm{sign}(\theta_i+\varepsilon_{1i}+\tau_{1i})] - \theta_i| \\
&= |\mathbb{E}[\lambda\,\mathrm{sign}(\theta_i+\varepsilon_{1i}+\tau_{1i}) - (\theta_i+\varepsilon_{1i})]| \\
&= |\mathbb{E}[\lambda\,\mathrm{sign}(\theta_i+\varepsilon_{1i}+\tau_{1i}) - (\theta_i+\varepsilon_{1i})]\mathbf{1}(|\theta_i+\varepsilon_{1i}| \ge \lambda)| \\
&\quad \blacktriangleright \text{ by } \mathbb{E}_{\tau_{1i}}[\lambda\,\mathrm{sign}(\theta_i+\varepsilon_{1i}+\tau_{1i})\mathbf{1}(|\theta_i+\varepsilon_{1i}| < \lambda) \\
&\qquad = \mathbb{E}_{\tau_{1i}}[Q_{2\lambda}(\theta_i+\varepsilon_{1i}+\tau_{1i})\mathbf{1}(|\theta_i+\varepsilon_{1i}| < \lambda)] \\
&\qquad = (\theta_i+\varepsilon_{1i})\mathbf{1}(|\theta_i+\varepsilon_{1i}| < \lambda); \text{ see Lemma 12} \\
&\le \mathbb{E}\Big[(\lambda + |\theta_i| + |\varepsilon_{1i}|)\mathbf{1}\Big(|\varepsilon_{1i}| \ge \sigma\sqrt{2\log(np)}\Big)\Big] \\
&\quad \blacktriangleright \text{ by triangle inequality and (6)} \\
&\le \frac{2\lambda}{np} + \int_{\sigma\sqrt{2\log(np)}}^{\infty}\mathbb{P}(|\varepsilon_{1i}| \ge t)\,dt \\
&\le \frac{2\lambda}{np} + \int_{\sigma\sqrt{2\log(np)}}^{\infty} 2\exp\Big(-\frac{t^2}{2\sigma^2}\Big)\,dt \le \frac{3\lambda}{np}. \\
&\quad \blacktriangleright \text{ by Lemma 2}
\end{aligned}$$

Combining with $\|\theta\|_2 \ge \lambda$ leads to

$$\|\mathbb{E}[\lambda\,\mathrm{sign}(\theta+\varepsilon_1+\tau_1)] - \theta\|_2 \le \frac{3\lambda}{n\sqrt{p}} \le \frac{6}{n}\frac{\|\theta\|_2^2}{\sqrt{\|\theta\|_2^2 + \frac{p\lambda^2}{n}}} \tag{86}$$

In view of Equation (62), for some positive $\epsilon_{n,\lambda}, \epsilon'_{n,\lambda} = O(r_{n,\lambda}^{-1} + (\frac{\log n}{n})^{1/4})$, we have that

$$P^{**} = \mathbb{P}\left(\langle Q_{2\lambda}(\theta+\varepsilon_1+\tau_1)-\theta, \hat{\xi}\rangle < -(1-\epsilon_{n,\lambda})\frac{\|\theta\|_2^2}{\sqrt{\|\theta\|_2^2+\frac{p\lambda^2}{n}}}\right)$$

$$\le \mathbb{P}\left(\langle Q_{2\lambda}(\theta+\varepsilon_1+\tau_1)-\theta, \hat{\xi}\rangle < -(1-\epsilon_{n,\lambda})\frac{\|\theta\|_2^2}{\sqrt{\|\theta\|_2^2+\frac{p\lambda^2}{n}}},\ \hat{B}_2\right) + \frac{1}{n^{1+\nu}p^{\nu}}$$

▶ by Equation (85)

$$\le \mathbb{P}\Bigg(\langle \lambda\,\mathrm{sign}(\theta+\varepsilon_1+\tau_1) - \mathbb{E}[\lambda\,\mathrm{sign}(\theta+\varepsilon_1+\tau_1)], \hat{\xi}\rangle$$
$$< -(1-\epsilon_{n,\lambda})\frac{\|\theta\|_2^2}{\sqrt{\|\theta\|_2^2+\frac{p\lambda^2}{n}}} + \|\mathbb{E}[\lambda\,\mathrm{sign}(\theta+\varepsilon_1+\tau_1)]-\theta\|_2\Bigg) + \frac{1}{n^{1+\nu}p^{\nu}}$$

▶ by $Q_{2\lambda}(\theta+\varepsilon_1+\tau_1) = \lambda\,\mathrm{sign}(\theta+\varepsilon_1+\tau_1)$ holds on the event $\hat{B}_2$ in (84)

$$\le \mathbb{P}\left(\langle \delta_1, \hat{\xi}\rangle < -(1-\epsilon'_{n,\lambda})\frac{\|\theta\|_2^2}{\sqrt{\|\theta\|_2^2+\frac{p\lambda^2}{n}}}\right) + \frac{1}{n^{1+\nu}p^{\nu}}\,.$$

▶ by Equation (86)

This completes the proof. □

Lemma 17. *Let $\theta \sim \mathrm{Unif}(\Delta S^{p-1})$, then we have*

$$\mathbb{P}\left(\mu(\theta) \ge \sqrt{\frac{3\log p}{p}}\right) \le 1 - 2p^{-1/4} - 2\exp(-cp). \tag{87}$$

Proof. Since $\mu(\cdot)$ is invariant under rescaling, we can treat $\theta \sim \mathrm{Unif}(S^{p-1})$ only. It is well known that $\theta \overset{\mathrm{d}}{=} \frac{g}{\|g\|_2}$ for $g \sim \mathcal{N}(0, I_p)$. By standard Gaussian tail bound, along with a union bound, it is not hard to show $\|g\|_\infty \le \sqrt{2.5\log p}$ with probability at least $1-2p^{-1/4}$. Furthermore, it is standard to show that $\|g\|_2 \ge 1.01\sqrt{p}$ with probability at least $1-2\exp(-cp)$ for some absolute constant c (e.g., [30, Theorem 3.1.1]). Combining both bounds completes the proof. □

Lemma 18 (Generalization of Lemma 14 for Appendix C)**.** *Given $R \in \mathbb{O}(p)$. Suppose that $U_1 \in R^{p\times n}$ has independent zero-mean entries of sub-Gaussian norms bounded by A_1; when conditioning on U_1, $U_2 \in R^{p\times n}$ has independent zero-mean rows of sub-Gaussian norms bounded by A_2. Then $U = RU_1 + U_2$ satisfies*

$$\mathbb{P}\left(\|U^TU - \mathbb{E}(U^TU)\|_{op} \le C(A_1^2+A_2^2)\cdot n\max\{1, \sqrt{\frac{p}{n}}\}\right) \le 8\exp(-2n).$$

Proof. By $U^TU = U_1^TU_1 + U_2^TU_2 + U_2^TRU_1 + U_1^TR^TU_2$ and $\mathbb{E}(U_2^TRU_1) = 0$ (conditioning on U_1 and using randomness of U_2), triangle inequality yields

$$\|U^TU - \mathbb{E}(U^TU)\|_{op} \le \|U_1^TU_1 - \mathbb{E}(U_1^TU_1)\|_{op} + \|U_2^TU_2 - \mathbb{E}(U_2^TU_2)\|_{op} + 2\|U_2^TRU_1\|_{op}.$$

By Lemma 14, with probability at least $1-4\exp(-2n)$, we have

$$\|U_1^TU_1 - \mathbb{E}(U_1^TU_1)\|_{op} \lesssim A_1^2 n\max\{1, \sqrt{\frac{p}{n}}\}, \quad \|U_2^TU_2 - \mathbb{E}(U_2^TU_2)\|_{op} \lesssim A_2^2 n\max\{1, \sqrt{\frac{p}{n}}\}.$$

It remains to bound $\|U_2^TRU_1\|_{op}$. By a standard bound on random matrix with independent, zero-mean sub-Gaussian entries (e.g., [30, Theorem 4.4.5]), we have that

$$\mathbb{P}\Big(\|U_1\|_{op} \le C'A_1(\sqrt{n}+\sqrt{p})\Big) \ge 1 - 2\exp(-2n).$$

To bound $\|U_2^T R U_1\|_{op}$, we assume that $\|U_1\|_{op} \le C' A_1(\sqrt{n}+\sqrt{p})$ holds, and we utilize the randomness of U_2 by conditioning on U_1. Let $N_{1/4}$ be a $1/4$-net of S^{n-1} of cardinality smaller than 9^n, then a simple discretization of the operator norm (e.g., [30, Exercise 4.4.3]) yields

$$\|U_2^T R U_1\|_{op} \le 2 \sup_{v_1, v_2 \in N_{1/4}} (U_2 v_2)^T R U_1 v_1. \tag{88}$$

For any pair $(v_1, v_2) \in N_{1/4} \times N_{1/4}$, we have $\|R U_1 v_1\|_2 \le \|U_1\|_{op} \le C' A_1(\sqrt{n}+\sqrt{p})$, and by the assumption on U_2 (conditioning on U_1) we also have that the entries of $U_2 v_2$ are independent, zero-mean, and of sub-Gaussian norms bounded by A_2. Therefore, $\|U_2 v_2\|_{\psi_2} = O(A_2)$ (see Lemma 3). Taken collectively, we reach $\|(U_2 v_2)^T R U_1 v_1\|_{\psi_2} = O(A_1 A_2(\sqrt{n}+\sqrt{p}))$. Thus, the standard sub-Gaussian tail bound, along with a union bound over $(v_1, v_2) \in N_{1/4} \times N_{1/4}$, establishes

$$\mathbb{P}\left(\sup_{v_1, v_2 \in N_{1/4}} (U_2 v_2)^T R U_1 v_1 \ge C'' A_1 A_2 (\sqrt{n}+\sqrt{p}) t \right) \le 2 \exp\Big(2n \log 9 - c t^2\Big), \quad \forall t \ge 0.$$

Setting $t \asymp \sqrt{n}$ and recalling (88), we obtain that $\|U_2^T R U_1\|_{op} \lesssim C'' A_1 A_2 n \max\{1, \sqrt{\frac{p}{n}}\}$ holds with probability at least $1 - 2\exp(-2n)$. Putting all the pieces together concludes the proof. □

Lemma 19 (Adaptation of Lemma 15 for Appendix C)**.** *In the proof of Theorem 2 in Appendix C with some given $R \in \mathbb{O}(p)$, where $\delta_j = Q_{2\lambda}(\theta_R + R\varepsilon_j + \tau_j) - \theta_R$ for $j \in [n]$, we have*

$$\mathbb{P}\bigg(\Big\| \frac{1}{n-1} \sum_{j=2}^n \delta_j \Big\|_2^2 < \frac{p\lambda^2}{n}\Big(1 + \frac{C}{n}\Big) + \frac{C' \|\theta\|_2^2}{\sqrt{n}} \bigg) \ge 1 - 6\exp(-10 r_{n,\lambda}^2).$$

Proof. We let $\delta_j' := Q_{2\lambda}(\theta_R + R\varepsilon_j + \tau_j) - (\theta_R + R\varepsilon_j)$ and thus $\delta_j = \delta_j' + R\varepsilon_j$, yielding

$$\Big\| \frac{1}{n-1} \sum_{j=2}^n \delta_j \Big\|_2^2 = \Big\| \frac{1}{n-1} \sum_{j=2}^n \delta_j' \Big\|_2^2 + \Big\| \frac{1}{n-1} \sum_{j=2}^n \varepsilon_j \Big\|_2^2 + 2 \Big\langle \frac{1}{n-1} \sum_{j=2}^n \delta_j', \frac{1}{n-1} \sum_{j=2}^n R\varepsilon_j \Big\rangle. \tag{89}$$

Note that $\frac{1}{n-1}\sum_{j=2}^n \delta_j'$ and $\frac{1}{n-1}\sum_{j=2}^n \varepsilon_j$ have independent, zero-mean entries of $O(\frac{\lambda}{\sqrt{n}})$ sand $O(\frac{\sigma}{\sqrt{n}})$ sub-Gaussian norms, respectively. By Bernstein's inequality as in Equations (77)–(78), with probability at least $1 - 4\exp(-10 r_{n,\lambda}^2)$ we have

$$\Big\| \frac{1}{n-1} \sum_{j=2}^n \delta_j' \Big\|_2^2 \le \mathbb{E} \Big\| \frac{1}{n-1} \sum_{j=2}^n \delta_j' \Big\|_2^2 + C \max\bigg\{ \frac{\lambda^2 \sqrt{p} \cdot r_{n,\lambda}}{n}, \frac{\lambda^2 r_{n,\lambda}^2}{n} \bigg\}, \tag{90}$$

$$\Big\| \frac{1}{n-1} \sum_{j=2}^n \varepsilon_j \Big\|_2^2 \le \mathbb{E} \Big\| \frac{1}{n-1} \sum_{j=2}^n \varepsilon_j \Big\|_2^2 + C \max\bigg\{ \frac{\sigma^2 \sqrt{p} \cdot r_{n,\lambda}}{n}, \frac{\sigma^2 r_{n,\lambda}^2}{n} \bigg\}. \tag{91}$$

We assume these hold. Combining with $\mathbb{E}\|\frac{1}{n-1}\sum_{j=2}^n \varepsilon_j\|_2^2 = O(\frac{\sigma^2 p}{n})$, we have

$$\Big\| \frac{1}{n-1} \sum_{j=2}^n \varepsilon_j \Big\|_2^2 \lesssim \frac{\sigma^2 (p + r_{n,\lambda}^2)}{n}.$$

Recalling that $\frac{1}{n-1}\sum_{j=2}^n \delta_j'$ has independent zero-mean entries of $O(\frac{\lambda}{\sqrt{n}})$ sub-Gaussian norms, we have

$$\Big\| \Big\langle \frac{1}{n-1} \sum_{j=2}^n \delta_j', \frac{1}{n-1} \sum_{j=2}^n R\varepsilon_j \Big\rangle \Big\|_{\psi_2} \lesssim \frac{\lambda}{\sqrt{n}} \Big\| \frac{1}{n-1} \sum_{j=2}^n \varepsilon_j \Big\|_2 \lesssim \frac{\sigma\lambda(\sqrt{p} + r_{n,\lambda})}{n},$$

which implies that

$$\Big| \Big\langle \frac{1}{n-1} \sum_{j=2}^n \delta_j', \frac{1}{n-1} \sum_{j=2}^n R\varepsilon_j \Big\rangle \Big| \lesssim \frac{\sigma\lambda(\sqrt{p} r_{n,\lambda} + r_{n,\lambda}^2)}{n} \tag{92}$$

holds with probability at least $1 - 2\exp(-10r_{n,\lambda}^2)$.

We now bound

$$\mathbb{E}\left\|\frac{1}{n-1}\sum_{j=2}^{n}\delta_j'\right\|_2^2 = \sum_{i=1}^{p}\mathbb{E}\left|\frac{1}{n-1}\sum_{j=2}^{n}\delta_{ji}'\right|^2 = \sum_{i=1}^{p}\sum_{j=2}^{n}\frac{\mathbb{E}[\delta_{ji}']^2}{(n-1)^2} = \sum_{i=1}^{p}\frac{\mathbb{E}[\delta_{ji}']^2}{n-1}, \tag{93}$$

where the last two inequalities hold because δ_j's are independent and zero-mean. Letting R_i^T be the i-th row of R, we expect over τ_{ji} and ε_j in sequence to reach

$$\begin{aligned}
&\mathbb{E}[\delta_{ji}']^2 = \mathbb{E}\Big[Q_{2\lambda}(\theta_{R,i} + R_i^T\varepsilon_j + \tau_{ji}) - (\theta_{R,i} + R_i^T\varepsilon_j)\Big]^2 \\
&= \mathbb{E}_{\varepsilon_j}\mathbb{E}_{\tau_{ji}}\left[Q_{2\lambda}(\theta_{R,i} + R_i^T\varepsilon_j + \tau_{ji})^2 + (\theta_{R,i} + R_i^T\varepsilon_j)^2 - 2(\theta_{R,i} + R_i^T\varepsilon_j)Q_{2\lambda}(\theta_{R,i} + R_i^T\varepsilon_j + \tau_{ji})\right] \\
&= \mathbb{E}_{\varepsilon_j}\left[\mathbb{E}_{\tau_{ji}}\big[Q_{2\lambda}(\theta_{R,i} + R_i^T\varepsilon_j + \tau_{ji})^2\big] - (\theta_{R,i} + R_i^T\varepsilon_j)^2\right] \\
&\quad \blacktriangleright \text{ by Lemma 12} \\
&= \mathbb{E}\big[Q_{2\lambda}(\theta_{R,i} + R_i^T\varepsilon_j + \tau_{ji})^2\big] - \theta_{R,i}^2 - \mathbb{E}_{\varepsilon_j}(R_i^T\varepsilon_j)^2
\end{aligned} \tag{94}$$

Moreover, by $|Q_{2\lambda}(a) - a| \le \lambda$ and , we have

$$\begin{aligned}
&\mathbb{E}[Q_{2\lambda}(\theta_{R,i} + R_i^T\varepsilon_j + \tau_{ji})^2] \le \lambda^2 + \mathbb{E}\big[Q_{2\lambda}(\theta_{R,i} + R_i^T\varepsilon_j + \tau_{ji})^2\mathbf{1}(|\theta_{R,i} + R_i^T\varepsilon_j| > \lambda)\big] \\
&\quad \blacktriangleright \text{ by } |Q_{2\lambda}(a + \tau_{ji})\mathbf{1}(|a| \le \lambda)| = \lambda \text{ for any } a \in \mathbb{R} \\
&\le \lambda^2 + \mathbb{E}\left[\Big(2\lambda + |\theta_{R,i}| + |R_i^T\varepsilon_j|\Big)^2 \cdot \mathbf{1}\Big(|R_i^T\varepsilon_j| \ge \sigma\sqrt{2(1+\nu)\log(np)}\Big)\right] \\
&\quad \blacktriangleright \text{ by Equation (10)} \\
&\le \lambda^2 + 18\lambda^2\mathbb{P}\Big(|R_i^T\varepsilon_j| \ge \sigma\sqrt{2(1+\nu)\log(np)}\Big) \\
&\qquad\qquad + 2\mathbb{E}\left[|R_i^T\varepsilon_j|^2\mathbf{1}\Big(|R_i^T\varepsilon_j| \ge \sigma\sqrt{2(1+\nu)\log(np)}\Big)\right] \\
&\quad \blacktriangleright \text{ by } |\theta_{R,i}| \le \lambda \text{ and } (A+B)^2 \le 2A^2 + 2B^2 \\
&\le \left(1 + \frac{18}{(np)^{1+\nu}}\right)\lambda^2 + 2\mathbb{E}\left[|R_i^T\varepsilon_j|^2\mathbf{1}\Big(|R_i^T\varepsilon_j| \ge \sigma\sqrt{2(1+\nu)\log(np)}\Big)\right]. \\
&\quad \blacktriangleright \text{ by Lemma 2}
\end{aligned}$$

By the tail bound of $|R_i^T\varepsilon_j|$ (cf. Lemma 2), we have

$$\begin{aligned}
&\mathbb{E}\left[|R_i^T\varepsilon_j|^2\mathbf{1}\Big(|R_i^T\varepsilon_j| \ge \sigma\sqrt{2(1+\nu)\log(np)}\Big)\right] \\
&= \int_0^\infty \mathbb{P}\left(|R_i^T\varepsilon_j|^2\mathbf{1}\Big(|R_i^T\varepsilon_j| \ge \sigma\sqrt{2(1+\nu)\log(np)}\Big) \ge t\right)dt \\
&\quad \blacktriangleright \text{ by tail integral representation of expectation} \\
&= \int_{2\sigma^2(1+\nu)\log(np)}^\infty \mathbb{P}\left(|R_i^T\varepsilon_j|^2 \ge t\right)dt \\
&\le 2\int_{2\sigma^2(1+\nu)\log(np)}^\infty \exp\Big(-\frac{t}{2\sigma^2}\Big)\,dt = \frac{4\sigma^2}{(np)^{1+\nu}}. \\
&\quad \blacktriangleright \text{ by Lemma 2}
\end{aligned}$$

Combining the preceding two displays establishes $\mathbb{E}\left[Q_{2\lambda}(\theta_{R,i} + R_i^T\varepsilon_j + \tau_{ji})^2\right] \le \left(1 + \frac{20}{(np)^{1+\nu}}\right)\lambda^2$, which together with (94) and (93) yields

$$\mathbb{E}\left\|\frac{1}{n-1}\sum_{j=2}^{n}\delta_j'\right\|_2^2 \le \frac{1}{n-1}\left[(1+\frac{1}{n})p\lambda^2 - \|\theta\|_2^2 - \mathbb{E}\|\varepsilon_j\|_2^2\right]. \tag{95}$$

We now substitute Equations (95), (92), (90)–(91) into Equation (89), along with noticing $\mathbb{E}\|\frac{1}{n-1}\sum_{j=2}^{n}\varepsilon_j\|_2^2 = \frac{1}{n-1}\mathbb{E}\|\varepsilon_j\|_2^2$ in Equation (91) and $r_{n,\lambda} \lesssim \min\{\frac{\|\theta\|_2}{\lambda}, \frac{\|\theta\|_2^2}{\lambda^2}\sqrt{\frac{n}{p}}\}$, to establish

$$\left\|\frac{1}{n-1}\sum_{j=2}^{n}\delta_j\right\|_2^2 \le \frac{p\lambda^2}{n}\Big(1 + \frac{C}{n}\Big) + \frac{C'\|\theta\|_2^2}{\sqrt{n}}$$

with the promised probability. □

Lemma 20 (Adaptation of Lemma 16 for Appendix C)**.** *In the proof of Theorem 2 in Appendix C, we have*

$$P^{**} \le \mathbb{P}\left(\frac{\langle \tilde{\delta}_1', \hat{\xi}\rangle}{\lambda} \le -(1-\epsilon_{n,\lambda})r_{n,\lambda}\right) + \frac{1}{n^{1+\nu}p^{\nu}}$$

for some $\epsilon_{n,\lambda} = O(r_{n,\lambda}^{-1} + (\frac{\log n}{n})^{1/4}) + \frac{1}{\log^{1/4}(np)}$.

Proof. Recall that P^{**} is defined in (70) with $\delta_1' := Q_{2\lambda}(\theta_R + R\varepsilon_1 + \tau_1) - (\theta_R + R\varepsilon_1)$, and as a surrogate we introduce $\tilde{\delta}_1' = \lambda\,\mathrm{sign}(\theta_R + R\varepsilon_1 + \tau_1) - \mathbb{E}_{\tau_1}[\lambda\,\mathrm{sign}(\theta_R + R\varepsilon_1 + \tau_1)]$. Notice that on the event $\hat{B}_2 = \{\|\theta_R + R\varepsilon_1\|_\infty \le \lambda\}$, we have $\delta_1' = \tilde{\delta}_1'$. Also, by $\lambda \ge (\frac{3\log p}{p})^{1/2} + \sigma\sqrt{2(1+\nu)\log(np)} \ge \|\theta_R\|_\infty + \sigma\sqrt{2(1+\nu)\log(np)}$, Lemma 2 along with a union bound gives

$$\mathbb{P}(\hat{B}_2) \ge \mathbb{P}\Big(\|R\varepsilon_1\|_\infty \le \sigma\sqrt{2(1+\nu)\log(np)}\Big) \ge 1 - \frac{1}{n^{1+\nu}p^{\nu}}.$$

Therefore, for some $\epsilon_{n,\lambda} = O(r_{n,\lambda}^{-1} + (\frac{\log n}{n})^{1/4}) + \frac{1}{\log^{1/4}(np)}$, we have

$$P^{**} = \mathbb{P}\left(\frac{\langle \tilde{\delta}_1', \hat{\xi}\rangle}{\lambda} \le -(1-\epsilon_{n,\lambda})r_{n,\lambda}\right) \le \mathbb{P}\left(\frac{\langle \delta_1', \hat{\xi}\rangle}{\lambda} \le -(1-\epsilon_{n,\lambda}r_{n,\lambda}), \hat{B}_2\right) + \mathbb{P}(\hat{B}_2^c),$$

which then yields the claim immediately. □

G Supplemental for Section 4

G.1 Further Details (Figures 1a, 1b)

The main aim of Figures 1a, 1b is to corroborate Corollary 1. Note that separation condition (9), when ignoring the factor of $1+\epsilon$, reads

$$\Delta^2 \ge \lambda^2\left(1 + \sqrt{1 + \frac{2p}{n\log n}}\right)\log n. \tag{96}$$

Also, the minimum of λ in Equation (6) is approximately $\sigma\sqrt{2\log(np)} = \sqrt{2\log(np)}$, by ignoring the factor of $1+\nu$ and observing that $\|\theta\|_\infty$ is typically dominated by $\sqrt{\log(np)}$. We therefore set λ approximately as $\sqrt{2\log(np)}$.

G.1.1 Low-Dimensional Setting (Figure 1a)

Recall that we set $a = \log n$, $b = \Delta$, and fix

$$p = 5.$$

In our experiment, given a we set $n = \lceil \exp(a) \rceil$.

Since we test $a \in [3.4, 6.9]$, the range of n is roughly $[30, 1000]$. Hence, under $p = 5$, we make further simplifications:

- Since $\sqrt{2\log(np)} \approx \sqrt{2\log n}$, in this experiment we simply set $\lambda = \sqrt{2\log n}$;
- Since $\frac{p}{n\log n} \ll 1$, approximately, the separation condition (96) further reduces to

  $$\Delta^2 \ge 2\lambda^2 \log n = (2\log n)^2,$$

  that is,

  $$\Delta = 2\log n.$$

  As such, the separation condition is roughly $b = 2a$.

G.1.2 High-Dimensional Setting (Figure 1b)

Recall that we set $n = 100$, $p = bn\log n$, $\Delta^2 = \lambda^2(1+\sqrt{a})\log n$, and $\lambda = \sqrt{2\log(np)}$. For this experiment, the setting provided in the main text is already clear. Hence, we only derive the separation condition (on (a, b)) here. In fact, substituting $\Delta^2 = \lambda^2(1+\sqrt{a})\log n$ and $p = bn\log n$ into (96) gives

$$\lambda^2(1+\sqrt{a}) \geq \lambda^2\left(1+\sqrt{1+2b}\right)\log n,$$

which then reduces to

$$a = 1 + 2b.$$

We note that [23] adopted the same experimental design.

G.2 Shrinking the Dithering Level (Figure 1c)

Although (6) is required in our theoretical results, in practice we suggest $\lambda = s\sqrt{2\log(np)}$ for some shrinkage parameter $0 < s < 1$ to draw a better bias-and-variance tradeoff (cf. Remark 1), in a spirit similar to [5, Remark 1] for two-bit covariance estimation. Indeed, in light of

$$r_{n,\lambda} = \frac{\|\theta\|_2^2/\lambda^2}{\sqrt{\frac{\|\theta\|_2^2}{\lambda^2} + \frac{p}{n}}},$$

an larger λ leads to smaller $c_{n,\lambda}$. Hence, using overly large λ leads to performance degradation. While (6) in Theorem 1 reads

$$\lambda \geq \|\theta\|_\infty + \sigma\sqrt{2(1+\nu)\log(np)} \approx \sigma\sqrt{2\log(np)},$$

and ensures that $\lambda \geq \max_{i\in[n]}\|X_i\|_\infty$ with high probability, in practice, a smaller λ could be sufficient for rendering small bias, and in this case using a smaller λ typically leads to better estimation performance.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: We write the abstract and introduction based on the contributions and scope of this paper.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: We discuss the limitations of our work in Remark 9 and Section 5, which are then left for future work.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We provide the full set of assumptions in Assumption 1 and the complete and correct proofs of our theorems are provided in Appendix.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We disclose all the information needed to reproduce the experimental results of the paper in Section 4. The code will also be open sourced after acceptance.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We use synthetic data produced by certain random seed and the code will be open sourced after acceptance.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: We disclose all the relevant information, such as choices of parameters, in Section 4 and Appendix G.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: We have reported the empirical success rates and misclassification rates in Figure 1 that are averaged over 300 independent trials.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: We have provided sufficient information on the computer resources needed to reproduce the experiments in Section 4.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: We checked the NeurIPS Code of Ethics and confirm that research conducted in the paper conform it.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

    Justification: This is primarily a theoretical work and we are not aware of potential positive societal impacts and negative societal impacts of this work.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.

- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: As far as we can tell, this paper poses no such risks.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [N/A]

    Justification: This paper does not use existing assets.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: This paper does not release new assets.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: This paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: This paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: LLM is used only for writing and editing.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.