

Partial recovery bounds for clustering with the relaxed K -means

Christophe Giraud and Nicolas Verzelen

Abstract. We investigate the clustering performances of the relaxed K -means in the setting of sub-Gaussian Mixture Model (sGMM) and Stochastic Block Model (SBM). After identifying the appropriate signal-to-noise ratio (SNR), we prove that the misclassification error decays exponentially fast with respect to this SNR. These partial recovery bounds for the relaxed K -means improve upon results currently known in the sGMM setting. In the SBM setting, applying the relaxed K -means SDP allows us to handle general connection probabilities whereas other SDPs investigated in the literature are restricted to the (dis-)assortative case (where within group probabilities are larger than between group probabilities). Again, this partial recovery bound complements the state-of-the-art results. All together, these results put forward the versatility of the relaxed K -means.

Mathematics Subject Classification (2010). 62H30, 68T10.

Keywords. Relaxed K -means, clustering, partial recovery bound, mixture of sub-Gaussian, stochastic block model.

1. Introduction

The problem of clustering is that of grouping similar “objects” in a data set. It encompasses many different instances such as partitioning points in a metric space, or partitioning the nodes of a graph.

1.1. K -means and a convex relaxation. When these objects can be represented as vectors in a Euclidean space, some of the most standard clustering approaches are based on the minimization of the K -means criterion [35]. Observing n objects and writing $X_a \in \mathbb{R}^p$ for the object $a \in \{1, \dots, n\}$, the K -means criterion of a partition $G = (G_1, \dots, G_k)$ of $\{1, \dots, n\}$ is defined as

$$\text{Crit}(G) = \sum_{k=1}^K \sum_{a \in G_k} \left\| X_a - \frac{1}{|G_k|} \sum_{b \in G_k} X_b \right\|^2, \quad (1)$$

where $\|\cdot\|$ stands for the Euclidean norm. This criterion quantifies the dispersion of each group around its centroid in order to favor homogeneous partitions. A K -means

procedure then aims at finding a partition \hat{G} that minimizes, at least locally, the K -means criterion. However, solving this problem is NP-hard and it is even hard to approximate [7].

In general, iterative procedures such as Lloyd’s algorithm [35] and its variants [5] are only shown to converge to a local minimum of the K -means criterion. Alternatively, Peng and Wei [41] have suggested to relax the K -means criterion to a Semi-Definite Program (SDP) followed by a rounding step. See the next section for a definition. The resulting program is provably solvable in polynomial time. This work is dedicated to promoting Peng and Wei’s procedure and some of its variants by (i) putting forward its versatility by handling both vector and general graph clustering problems and (ii) assessing its near-optimal performances.

1.2. Sub-Gaussian Mixture Models (sGMM) and Stochastic Block Models (SBM).

In the computer-science and statistical literature, the most popular approach to assess the performances of a procedure is the “model-based” strategy. It assumes there exists a true unknown partition G of the “objects” and that the data have been randomly generated from a probability distribution rendering this partition. Then, one can assess the performances of a clustering procedure by comparing the partition estimated from the data to G .

For vector clustering, it is classical to assume that the vectors X_a are distributed according to a sub-Gaussian Mixture Model (sGMM). In a sGMM with partition G , the random variables X_a are assumed to be independent and for $a \in G_k$, the random variable X_a is assumed to follow a sub-Gaussian distribution centered at $\mu_k \in \mathbb{R}^p$ and with covariance matrix Σ_k . In other words, variables X_a whose indices a belongs to the same group are identically distributed and variables X_a and X_b whose indices belong two different groups have different means. See Section 3.1 for a definition.

Node clustering in a network has been widely investigated within the framework of Stochastic Block Models (SBM) [27] and its variants. According to a SBM with partition G , the network edges are sampled independently and the probability of presence of an edge between any two nodes $a \in G_k$ and $b \in G_l$ is equal to some quantity $P_{kl} \in [0, 1]$ only depending on the groups. In other words, two nodes a and b belonging to the same group in G share the same probability of being connected to any other node c .

These two random models have attracted a lot of attention in the last decade. See e.g. [1, 39] for two recent reviews on SBM and [17, 19, 28, 31, 34, 36, 37, 43, 44, 46] for recent contributions on sGMM. A large body of the literature on these two models focuses on pinpointing the right scaling between the model parameters allowing to recover the partition G from the data. For sGMM, this translates into identifying the minimal distance $\min_{k \neq l} \|\mu_k - \mu_l\|$ within the mixtures means, such that, there exists a clustering procedure, if possible running in polynomial time, that recovers G with high probability. Most of the works concentrate on two types of recovery: perfect recovery, where one wants to recover exactly the partition G with high probability and

weak-recovery where the estimated partition \hat{G} is only required to be more accurate than random guessing. The goal is then to identify the precise threshold at which perfect or weak recovery can occur. We refer to [1] for a review of these questions in SBM. Between these two extreme regimes, when the best possible classification is neither perfect nor trivial, the objective is to maximize the proportion of well-classified data. Given two partitions $\hat{G} = (\hat{G}_1, \dots, \hat{G}_K)$ and $G = (G_1, \dots, G_K)$ of $\{1, \dots, n\}$ into K non-void groups, we define the proportion of non-matching points

$$err(\hat{G}, G) = \min_{\pi \in \mathcal{S}_K} \frac{1}{2n} \sum_{k=1}^K |G_k \Delta \hat{G}_{\pi(k)}|, \tag{2}$$

where $A \Delta B$ represents the symmetric difference between the two sets A and B and \mathcal{S}_K represents the set of permutations on $\{1, \dots, K\}$. When \hat{G} is a partition estimating G , we refer to $err(\hat{G}, G)$ as the misclassification proportion (or error) of the clustering. The problem of minimizing this error has attracted less attention but see [2, 8, 16–18, 20, 22, 24, 47] for some related contributions.

Among the polynomial-time clustering procedures, Semi-Definite-Programs (SDP) have proved to be versatile and they have been investigated in a large range of clustering problems, including clustering in SBM [15, 20, 24, 25, 30, 42], sGMM [17, 34, 37, 44] or in block covariance models [9, 11]. While not always reaching the exact threshold for weak/perfect clustering in several cases [30, 42], SDP algorithms are versatile enough in order to enjoy some robustness properties [20, 38, 42], which are not met by more specialized algorithms (see [38] for more details). However, most SDPs require the partition to be balanced or that, at least, the size of each group is known in advance. Besides, all SDPs studied for SBM clustering arise as convex relaxations of min-cut optimization problems [15, 20, 24, 25, 30, 33, 42] and therefore only fall within the framework of assortative SBM where within group probabilities of connection are larger than between group probabilities of connection. In other words, the diagonal entries of P have to be larger than its off-diagonal entries.

1.3. Our contribution. In this work, we provide misclassification error bounds for the relaxed K -means of Peng and Wei [41] combined with a rounding step, both in the sGMM and the SBM frameworks. Compared to other SDPs, this convex relaxation of K -means has the nice feature to only require the knowledge of the number of groups (which can sometimes be estimated, as in [11]). Hence, there is no need to know the size of the clusters, nor the parameters of the model. The details about this SDP and the subsequent rounding step are given in Section 2.

Some of the first partial recovery results for SDPs have been derived using the Grothendieck inequality [17, 24]. The corresponding misclassification error bounds scale with a square-root decay with respect to an appropriate signal-to-noise ratio. More recently, Fei and Chen [20] have dramatically improved such bounds in the

context of assortative SBM, by proving misclassification error bounds for their SDP that actually exponentially decays with respect to this signal-to-noise ratio. Our results and our proof techniques are inspired by this work.

Let us first give a glimpse of our results on sGMM, by specifying it to the special case of Gaussian mixture models, with K groups of equal size $m = n/K$ and equal covariance Σ . The general statement of the results for possibly unbalanced groups and unequal covariances in sGMM is postponed to Section 3. Write $\Delta = \min_{k \neq l} \|\mu_k - \mu_l\|$ for the minimal Euclidean distance between the means of the components and write $R_\Sigma = |\Sigma|_F^2 / |\Sigma|_{\text{op}}^2$ for the ratio between the square Frobenius norm of Σ and the square operator norm of Σ . This ratio can be interpreted as an effective rank of Σ and is always smaller than the ambient dimension p . In the sequel, c stands for a positive numerical constant. Then, Theorem 1 in Section 3 entails that, with high probability, the proportion of misclassified observations decreases exponentially fast with the signal to noise ratio

$$s^2 = \frac{\Delta^2}{|\Sigma|_{\text{op}}} \wedge \frac{n\Delta^4}{K|\Sigma|_F^2}, \quad (3)$$

at least, as long as the condition $s^2 \geq cK$, or equivalently

$$\Delta^2 \geq c|\Sigma|_{\text{op}} \left(1 \vee \sqrt{\frac{R_\Sigma}{n}} \right) K \quad (4)$$

is met. The shape (3) of the signal-to-noise ratio is new and it differs from the classical signal-to-noise ratio $\tilde{s}^2 = \Delta^2 / |\Sigma|_{\text{op}}$ considered e.g. in [36]. We explain in Section 3.3 why the exponential decay should be with respect to s^2 , at least in the isotropic case. Since $\text{err}(\hat{G}, G) \leq 1/n$ implies that the partition \hat{G} is equal to G , the exponential decay with respect to s^2 ensures perfect recovery of the clustering with high-probability when $s^2 \geq c(K \vee \log(n))$, recovering the results of [44]. It also ensures a better than random guess clustering when (4) is met, which improves, in high-dimensional setting, upon state-of-the art results in [36, 37].

On the SBM side, we explain how the relaxed K -means procedure can be applied to general SBM to cluster nodes presenting similar connectivity profiles. Instead of the previously discussed SDPs that look for a partition with maximal within-group connectivity, this allows us to handle general unknown connection matrices P and thereby going far beyond the assortative case. Denoting by m the size of the smallest group in G , we prove that, with high probability, the misclassification proportion decreases exponentially fast with the signal-to-noise ratio

$$s^2 = m \cdot \min_{j \neq k} \frac{\|P_{j\cdot} - P_{k\cdot}\|^2}{|P|_\infty}, \quad (5)$$

at least as long as the condition $s^2 \geq cn/m$ is met. Here, $P_{j\cdot}$ stands for the j -th row of P and $|P|_\infty$ denotes the supremum norm. Note that this result encompasses

sparse graph, where the connection probability may scale as a constant divided by n . When specified to the classical case with all within-group probabilities equal to p and all between-group probabilities equal to q , with $q < p$, and all groups of the same size, we recover the results obtained by [20] for a relaxed version of the MLE, but without knowing that we are in the assortative case, nor knowing the group sizes.

1.4. Connection to the literature. Only a few papers have previously proved theoretical properties on the relaxed K -means of [41]. [6, 29, 34] obtain perfect recovery results for the so-called stochastic ball models and Gaussian mixture models, and [37] provides bounds on the estimation of the centers of the means in the sGMM, under a condition stronger than (4). Closer to the present paper, [11, 44] (see also [10]) prove perfect recovery results in the setting of block covariance models and sGMM. To the best of our knowledge, the main result of [44] provides the weaker condition in high-dimension ($p \geq n$) ensuring perfect recovery with polynomial-time algorithm in the sGMM. This condition is $s^2 \geq c(K \vee \log(n))$, with s^2 defined by (3). Theorem 1 below extends this result to the partial recovery regime, in the sense that the main result in [44] can be recovered from Theorem 1.

In the sGMM setting, the paper [36] derives partial recovery results for Lloyd algorithm (with a suitable initialization). To the best of our knowledge, these results are the strongest ones in the literature. In the setting discussed above, they prove a decay of the misclassification proportion exponentially fast relative to $\tilde{s}^2 = \Delta^2/|\Sigma|_{\text{op}}$, with a minimal signal-to-noise requirement

$$\Delta^2 \geq c|\Sigma|_{\text{op}}K^2\left(1 \vee \frac{pK}{n}\right).$$

Here, the requirement is proportional to $K^2|\Sigma|_{\text{op}}$ in low-dimension, which is larger than our $K|\Sigma|_{\text{op}}$ by a factor K . In high-dimension, since we always have $R_\Sigma \leq p$, the factor pK/n is also much larger than our $\sqrt{R_\Sigma/n}$ in condition (4). This more limited range of validity is partially due to the fact that [36] investigates exponential decay with respect to $\tilde{s}^2 = \Delta^2/|\Sigma|_{\text{op}}$, rather than the suitable signal-to-noise ratio s^2 given by (3). We refer to Section 3.3 (and Appendix B) for a short explanation of this point. Yet, compared to us, [36] have a tight constant in the exponential rate.

During the wrap-up of this paper, we became aware of an independent and simultaneous work of Fei and Chen [21], which also investigate partial recovery in sGMM with another SDP. They show interesting connections between their SDP and the error of the supervised classification problem with known centers, and derive some partial recovery bounds based on it. Their results also have a more limited range of validity than ours, as they require groups of the same size and a minimal signal-to-noise condition of the form

$$\Delta^2 \geq c|\Sigma|_{\text{op}}\left(K\left(1 \vee \frac{p}{n}\right) + \sqrt{\frac{Kp \log(n)}{n}}\right),$$

instead of (4). As before, this more limited condition is partially due to the fact that they investigate exponential decay with respect to $\tilde{s}^2 = \Delta^2/|\Sigma|_{\text{op}}$, instead of s^2 given by (3).

In the SBM setting, most results on partial recovery [2, 16, 20, 22, 47] cover the assortative setting. The papers [16, 22, 47] investigate some two-steps procedures based on a spectral algorithm. The papers [22, 47] derive tight misclassification bounds for their algorithm, showing sharp exponential decay with respect to the signal-to-noise ratio. Closer to us, [20] proves similar results for an SDP, with less tight constants than [22, 47], but a wider range of validity. Compared to these results, our results does not provide sharp constants as in [22, 47]. Yet, they provide some new results for partial recovery in non-assortative cases and they only require a mild condition on the size of the smallest cluster. To the best of our knowledge, (i) our results are the first results about clustering with an SDP in non-assortative cases and (ii) the only known exponential bounds for partial recovery in general SBM are those of [2] which handle the sparse setting where the matrix P scales as $P = P_0/n$, with P_0 fixed and $n \rightarrow \infty$. Their results are optimal in the vicinity of the weak recovery threshold. Our results cover a setting with slightly more signal (the misclassification error has to be smaller than e^{-cK}), and the results do not overlap. In particular, as discussed in Section 4, our exponential rate (5) involved in Theorem 2 is faster by (at least) a factor K than the exponential rate involved in [2], though the rate of [2] cannot be improved in the vicinity of the weak recovery threshold. Hence both results are more complementary than comparable.

Since our work has been made available, two follow-up papers have extended and complemented our results. [13] have proved an exponential clustering error with respect to the SNR s^2 (3) for sGMM in infinite dimensional Hilbert spaces with common covariance matrix. For spherical mixtures with $K = 2$, [40] has proved a minimax lower bound showing that the clustering error $\exp(-cs^2)$ is optimal.

1.5. Organization and notation. The paper is organized as follows. In Section 2, we recall the relaxed K -means SDP derived by [41] and we explain how the partition is derived from the solution of the SDP. Section 3 covers the sGMM and Section 4 covers the SBM. We explain the main lines of the proofs and discuss the main arguments in Section 5. Finally, the full proof of the two main theorems can be found in Section 6 and Section 7.

Notation. To any matrix M we denote by $\text{Row}(M)$ the set of its rows, by $|M|_1$ the ℓ^1 norm of its entries, by $|M|_{\text{op}}$ its operator norm with respect to the ℓ^2 norm, by $|M|_F$ its Frobenius norm, by $|M|_*$ its nuclear norm and by $\text{Tr}(M)$ its trace. We also associate to a diagonal matrix D , the pseudo-norm $|D|_V = \max_a D_{aa} - \min_a D_{aa}$. Besides, for any A, B with the same dimensions, we write $\langle A, B \rangle = \sum_{ab} A_{ab}B_{ab}$ for its canonical inner product. For a vector x we write $\|x\|_2$ for its Euclidean norm

and $\langle \cdot, \cdot \rangle$ for the corresponding inner product. In the sequel, $\mathbf{1}$ stands for the indicator function.

For two sequences u_n and v_n (possibly depending on other parameters), we write $u_n \lesssim v_n$ (resp. $u_n \gtrsim v_n$) when there exists some numerical constant $c > 0$ such that $u_n \leq cv_n$ (resp. $u_n \geq cv_n$). Given $x, y \in \mathbb{R}$, $x \vee y$ (resp. $x \wedge y$) stands for the maximum (resp. the minimum) of x and y .

2. Relaxed K -means

We have n “objects” that we want to cluster. In the case of sGMM, these objects are p -dimensional vectors, and in the case of SBM they corresponds to the nodes of a graph. For each object a , we have a p -dimensional vector of observations: In the sGMM setting, the observation related to a is the vector $X_a \in \mathbb{R}^p$ and in the SBM the observation is the vector $X_a \in \{0, 1\}^n$ recording presence/absence of edges between a and the other nodes (hence $p = n$ is this case). We denote by $\mathbf{X} \in \mathbb{R}^{n \times p}$, the $n \times p$ matrix whose a -th row is given by X_a . In particular, in the SBM case, \mathbf{X} is simply the adjacency matrix of the graph.

Peng and Wei [41] have observed that any partition G of $\{1, \dots, n\}$ can be uniquely represented by a $n \times n$ matrix $B \in \mathbb{R}^{n \times n}$ such that $B_{ab} = 0$ if and only if a and b are in different group and $B_{ab} = 1/|G_k|$ if a and b are in the same group G_k . The collection of such matrices when G spans the collection of all partitions with K groups may be described as

$$\mathcal{P} = \{B \in \mathbb{R}^{n \times n} : \text{symmetric}, B^2 = B, \text{Tr}(B) = K, B\mathbf{1} = \mathbf{1}, B \geq 0\}.$$

Here, $B \geq 0$ means that all entries of B are nonnegative and B^2 refers to the matrix product of B with itself. Peng and Wei [41] have shown that minimizing the classical K -means criterion (1) is equivalent to maximizing $\langle \mathbf{X}\mathbf{X}^T, B \rangle$ over the space \mathcal{P} . Writing \tilde{B} for such a maximizer, the K -means clustering is obtained from \tilde{B} by grouping together indices a, b which have a non-zero entry \tilde{B}_{ab} .

The constraint set \mathcal{P} is non-convex and solving the K -means problem is NP-hard [7]. Peng and Wei [41] then propose to relax the constraint set \mathcal{P} by dropping the condition $B^2 = B$ to consider

$$\mathcal{C} = \{B \in \mathbb{R}^{n \times n} : \text{Positive Semi Definite}, \text{Tr}(B) = K, B\mathbf{1} = \mathbf{1}, B \geq 0\},$$

and hence solve the relaxed K -means SDP

$$\hat{B} \in \operatorname{argmax}_{B \in \mathcal{C}} \langle \mathbf{X}\mathbf{X}^T, B \rangle. \tag{6}$$

Obviously, the solution \hat{B} does not necessarily belong to \mathcal{P} and then does not provide a clustering. One has therefore to rely on a rounding step to obtain a proper partition. If \hat{B} is close to the true matrix B , one should expect that rows of \hat{B} belonging to the

same group are similar. This is why the final step is obtained by applying a clustering algorithm on the rows of \hat{B} . As in [20], we apply here an approximate K -medoids on the rows of \hat{B} . Let us detail this final step.

To any partition $G = (G_1, \dots, G_K)$ of $\{1, \dots, n\}$ into K (non empty) groups, we can associate a partition matrix $A \in \mathbb{R}^{n \times K}$ defined by $A_{ik} = \mathbf{1}_{i \in G_k}$. Let us denote by \mathcal{A}_K the set of such matrices and by $\text{Rows}(\hat{B})$ the set of the rows of \hat{B} . Then a ρ -approximate K -medoids on the rows of \hat{B} is a pair (\hat{A}, \hat{M}) with $\hat{A} \in \mathcal{A}_K$, $\hat{M} \in \mathbb{R}^{k \times n}$, $\text{Rows}(\hat{M}) \subset \text{Rows}(\hat{B})$ and fulfilling

$$|\hat{A}\hat{M} - \hat{B}|_1 \leq \rho \min_{\substack{A \in \mathcal{A}_K, \\ \text{Rows}(M) \subset \text{Rows}(\hat{B})}} |AM - \hat{B}|_1. \quad (7)$$

We refer to [12] for a polynomial-time algorithm producing such an output (\hat{A}, \hat{M}) with $\rho = 7$. Then a partition is obtained from \hat{A} by setting $\hat{G}_k = \{i : \hat{A}_{ik} = 1\}$.

In the sequel, \hat{G} is said to be a relaxed K -means solution if \hat{G} is derived from any 7-approximate K -medoids on the rows of \hat{B} obtained in (6). All our partial recovery bounds are for this partition \hat{G} .

As a final remark, K -means is known to suffer from a bias which tends to produce groups of similar width, see e.g. [44]. As explained in [11, 44], K -means and its relaxed version can be debiased when useful (e.g. for high-dimensional mixtures with unequal traces $\text{Tr}(\Sigma_k)$). We refer to Section 3.4 for details.

3. Clustering sub-Gaussian mixtures

3.1. Model. We observe n independent random vectors $X_1, \dots, X_n \in \mathbb{R}^p$. We assume that there exists an unknown partition (G_1, \dots, G_K) of $\{1, \dots, n\}$ and K unknown p -dimensional vectors $\mu_1, \dots, \mu_K \in \mathbb{R}^p$, such that

$$X_a = \mu_k + E_a \quad \text{for any } a \in G_k,$$

with E_a centered, independent with covariance Σ_k . We recall that $m = \min_k |G_k|$ stands for the size of the smallest group.

The larger the Euclidean distance between two centers $\Delta_{jk} = \|\mu_k - \mu_j\|$, the more easily we can recover the unknown partition from the observations X_1, \dots, X_n . Hence, we denote by $\Delta = \min_{j \neq k} \Delta_{jk}$ the minimal distance between two distinct centers, which will represent the signal part in the signal-to-noise ratio.

The hardness of the clustering problem also depends on the concentration of the random vectors E_a around zero. A common distributional assumption when analyzing clustering is the sub-Gaussianity assumption. For a centered random vector $Z \in \mathbb{R}^p$ and $L > 0$, we say that Z is $\text{SubG}(LI_p)$ if the random variables Z_1, \dots, Z_p are independent and $\mathbb{E}[\exp(tZ_i)] \leq \exp(t^2L^2/2)$ for all $t \in \mathbb{R}$

and $i = 1, \dots, p$. For sub-Gaussian mixture, we make the following distributional assumption.

Assumption A1. There exists $L > 0$ such that $\Sigma_k^{-1/2} E_a$ is $\text{SubG}(L^2 I_p)$.

Under this assumption, two quantities mainly drive the noise width in the signal-to-noise ratio: the maximum scaled operator norm of the covariances

$$\sigma^2 = L^2 \max_k |\Sigma_k|_{\text{op}},$$

and the maximum scaled Frobenius norm of the covariances

$$\nu^2 = L^2 \max_k |\Sigma_k|_F.$$

Actually, as shown in Theorem 1 below, the misclassification error of the relaxed K -means decreases exponentially fast with the signal-to-noise ratio

$$s^2 = \frac{\Delta^2}{\sigma^2} \wedge \frac{m \Delta^4}{\nu^4}, \tag{8}$$

where m denotes the size of the smallest cluster. This particular definition of the signal-to-noise ratio is new and is further discussed below.

Remarks. (1) When the random variable E_a is normally distributed with covariance Σ_k , it fulfills Assumption A1 with $L = 1$.

(2) We observe that the random variable E_1, \dots, E_n are all sub-Gaussian $\text{SubG}(\sigma^2 I_p)$, and we can always upper-bound ν^2 by $\nu^2 \leq p \sigma^2$, with equality in the spherical case where the covariances Σ_k are proportional to the identity matrix. Yet, this upper bound is crude when the covariances are far from being proportional to the identity matrix.

3.2. Partial recovery bound. Let us denote by $\Gamma \in \mathbb{R}^{n \times n}$ the diagonal matrix with entries $\Gamma_{aa} = \text{Tr}(\Sigma_k)$ for $a \in G_k$. As shown in [44] (see also [11]), when the trace of the covariances $\Sigma_1, \dots, \Sigma_K$ are unequal, i.e. when Γ is not proportional to the identity, it is useful to de-bias the relaxed K -means (6) by removing from $\mathbf{X}\mathbf{X}^T$ a preliminary estimator $\hat{\Gamma}$ of Γ . This estimator can be $\hat{\Gamma} = 0$ (no correction) when the covariances have equal traces or be equal to (17) as defined in Section 3.4 when the trace of the covariances are unequal. This leads to the so-called Pecok estimator [11, 44]

$$\hat{B} \in \text{argmax}_{B \in \mathcal{C}} \langle \mathbf{X}\mathbf{X}^T - \hat{\Gamma}, B \rangle. \tag{9}$$

As explained in Section 2, \hat{G} is then computed as any 7-approximate K -medoid solution of \hat{B} . We recall the notation $|D|_V = \max_a D_{aa} - \min_a D_{aa}$.

Theorem 1. *There exist three positive constants c, c', c'' such that the following holds. Assume that Assumption A1 holds,*

$$\Delta^2 \geq \frac{64|\Gamma - \hat{\Gamma}|_V}{m}, \quad (10)$$

and that s^2 (as defined in (8)) satisfies $s^2 \geq c''n/m$, then the proportion of misclassified points is upper bounded by

$$\text{err}(\hat{G}, G) \leq e^{-c's^2},$$

with probability at least $1 - c/n^2$,

Leaving aside Condition (10), which will be discussed below, one observes that the misclassification error decreases exponentially fast with respect to the signal-to-noise ratio s^2 as soon as s^2 is large enough. We further discuss this result in the next two paragraphs, first discussing the cases of equal covariance traces and then turning to unequal trace case.

3.3. Equal traces case. We assume here that all the covariance matrices Σ_k have equal trace. Hence, Γ is proportional to the identity and $|\Gamma|_V = 0$. So, when choosing $\hat{\Gamma} = 0$ in (9), Condition (10) becomes $\Delta^2 \geq 0$ which always holds.

Non-trivial recovery. Theorem 1 ensures non-trivial recovery as soon as $s^2 \gtrsim n/m$. Introduce R_Σ as the ratio

$$R_\Sigma = \frac{v^4}{\sigma^4} = \frac{\max_{k=1,\dots,K} |\Sigma_k|_F^2}{\max_{k=1,\dots,K} |\Sigma_k|_{\text{op}}^2} \leq \max_{k=1,\dots,K} \frac{|\Sigma_k|_F^2}{|\Sigma_k|_{\text{op}}^2} \leq p, \quad (11)$$

which can be interpreted as an effective rank of the mixture model. In order to compare this result with those in the literature, let us discuss this condition in the special case of balanced partition with K groups of equal size $m = n/K$. Theorem 1 guaranties a non-trivial recovery as soon as $s^2 \gtrsim K$, or equivalently

$$\frac{\Delta^2}{\sigma^2} \gtrsim \left(1 \vee \sqrt{\frac{R_\Sigma}{n}}\right) K, \quad (12)$$

with a misclassification error upper-bounded by $e^{-c'K}$ with high-probability. Taking for granted that, as we advocate below, the exponential decay $e^{-c's^2}$ is optimal in some cases, we cannot hope for a weaker condition than (12) to ensure a misclassification error of at most $e^{-c'K}$ when $K \lesssim \log(n)$ (for larger K a misclassification proportion of $e^{-c'K} \leq 1/n$ ensures perfect recovery). Yet, it is possible to get a weaker dependence in K when $K \gtrsim \log(n)$. Actually, in a large sample setting where $n \gtrsim p^3 K^2 \log(pK)$, [46] derives clustering guaranties for an iterate spectral

clustering under the weaker condition $\Delta^2 \gtrsim \sigma^2(\sqrt{K \log(n)} + \log(n))$. Nevertheless, we emphasize that our result, contrary to theirs, holds in the high-dimensional regime $p \geq n^{1/3}$. For the different goal of learning the means in a large sample size setting, some recent papers [19,28,31] have shown that a separation $\Delta \geq K^\epsilon$ in the number K of cluster is enough for learning the means in polynomial time, when the sample size is larger than $n \geq \text{poly}(p^{1/\epsilon}, k)$. For the same question, [37] has shown that, in a large sample size setting, the relaxed K -means succeeds to learn the means when $\Delta^2/\sigma^2 \gtrsim K^2$, which is a stronger requirement than (12). Turning back to our problem of deriving exponential bound for the misclassification proportion, [36] provides such bounds for the Lloyd algorithm under the minimal requirement

$$\frac{\Delta^2}{\sigma^2} \gtrsim K^2 \left(1 \vee \frac{pK}{n}\right), \tag{13}$$

which, again, is stronger than (12). To the best of our knowledge, our result is the first result of this kind for an SDP in this setting. We mention yet, that in an independent and simultaneous work, Fei and Chen [21] have derived a similar in spirit result in the very precise setting where the groups are of equal size. Actually, for an SDP taking as input that all groups have the same size n/K , [21] shows non-trivial recovery when

$$\frac{\Delta^2}{\sigma^2} \gtrsim \left(1 \vee \frac{p}{n}\right)K + \sqrt{\frac{Kp \log(n)}{n}}. \tag{14}$$

Since we always have $R_\Sigma \leq p$, the requirement (14) of [21] or (13) of [36] are stronger than (12), especially in the practical case where p is larger than n , but the effective rank R_Σ is small compared to p .

Intermediate regime. In the intermediate regime where $n/m \lesssim s^2 \lesssim (n/m) \vee \log(n)$, the misclassification rate of our procedure decays at the exponential rate $s^2 = \frac{\Delta^2}{\sigma^2} \wedge \frac{m\Delta^4}{v^4}$. To simplify the discussion, we assume again that the clustering is made of K groups of equal size $m = n/K$. For spherical mixtures, the misclassification rate of the Bayes classifier decays at the exponential rate $\tilde{s}^2 = \frac{\Delta^2}{\sigma^2} \geq s^2$. In view of the definition of s^2 and R_Σ , Theorem 2 ensures that relaxed K -means achieves this optimal rate as soon as

$$\frac{\Delta^2}{\sigma^2} \gtrsim \left(1 \vee \frac{R_\Sigma}{n}\right)K. \tag{15}$$

Such exponential rates have already been obtained in [21, 36] but under the corresponding stronger separation conditions (13) and (14). Nevertheless, the numerical constants in the exponential rate of [36] are tighter than ours.

In the high-dimensional setting $R_\Sigma > n$, the misclassification rate of relaxed K -means decays at the slower exponential rate $\frac{\Delta^4 n}{\sigma^4 K R_\Sigma} < \frac{\Delta^2}{\sigma^2}$ when the distances between the means satisfy

$$\sqrt{\frac{R_\Sigma}{n}} K \lesssim \frac{\Delta^2}{\sigma^2} \lesssim \frac{R_\Sigma}{n} K. \quad (16)$$

To the best of our knowledge, this moderate signal regime was not previously covered in the literature. The discrepancy between the rates $\frac{\Delta^4 n}{\sigma^4 K R_\Sigma}$ and $\frac{\Delta^2}{\sigma^2}$ may seem suboptimal. Yet, as we explain below, this discrepancy is inherent to the lack of knowledge of the location of the means of the clusters, and it seems unavoidable in our clustering setting.

Actually, let us consider the arguably simpler problem of Gaussian supervised classification with a two-class balanced partition, a common spherical covariance $\Sigma_k = \sigma I_p$ and opposite means $\mu_{-1} = -\mu_1$ uniformly distributed on the Euclidean sphere $\partial B(0, \Delta/2)$. More precisely, assume that we have n labeled observations $(X_a, Z_a) \in \mathbb{R}^p \times \{-1, 1\}$, for $a = 1, \dots, n$ distributed as follows. The labels Z_1, \dots, Z_n are i.i.d. with uniform distribution on $\{-1, 1\}$, a random vector $\mu \in \mathbb{R}^p$ is sampled uniformly over the sphere $\partial B(0, \Delta/2)$ independently of Z_1, \dots, Z_n , and, conditionally on Z_1, \dots, Z_n, μ , the X_a are independent Gaussian random variables with mean $Z_a \mu$ and covariance $\sigma^2 I_p$. Then, direct computations (see Appendix B for details) show that the classifier minimizing the probability of misclassification of a new observation is

$$\hat{h}(x) = \text{sign}\left(\left\langle \frac{1}{n} \sum_{a=1}^n Z_a X_a, x \right\rangle\right).$$

According to the invariance of the distribution by rotation, the probability of misclassification is given by

$$\mathbb{P}[Z_{\text{new}} \neq \hat{h}(X_{\text{new}})] = \mathbb{P}\left[\left\langle \frac{\mu}{\sigma} + \frac{\epsilon}{\sqrt{n}}, \frac{\mu}{\sigma} + \epsilon' \right\rangle < 0\right],$$

where $\mu = [\Delta/2, 0, \dots, 0] \in \mathbb{R}^p$, and ϵ, ϵ' are two independent standard Gaussian vectors in \mathbb{R}^p . For

$$\frac{\Delta^2}{\sigma^2} \gtrsim \left(1 \vee \frac{p}{n}\right),$$

this error is of exponential order Δ^2/σ^2 whereas, for

$$\left(1 \vee \sqrt{\frac{p}{n}}\right) \lesssim \frac{\Delta^2}{\sigma^2} \lesssim \left(1 \vee \frac{p}{n}\right),$$

the error is of exponential order $n\Delta^4/p\sigma^4$, see again Appendix B for details. Hence, even in this simpler toy model, the exponential rate with respect to $s^2 = \frac{\Delta^2}{\sigma^2} \wedge \frac{n\Delta^4}{2p\sigma^4}$ is

intrinsic. To formalize this argument, one should prove rigorously a minimax lower bound for our clustering problem, but this is beyond the scope of this paper.

When the common covariance Σ is not spherical, we cannot argue anymore that the exponential rate s^2 is optimal. Actually, in the Gaussian supervised classification setting with known means μ_{-1}, μ_1 and (common) covariance Σ , the probability of misclassification is known to decay exponentially fast with the Mahalanobis distance $d_{\Sigma}^2(\mu_{-1}, \mu_1) = (\mu_1 - \mu_{-1})^T \Sigma^{-1} (\mu_1 - \mu_{-1})$ rather than $\Delta^2/\sigma^2 = \|\mu_1 - \mu_{-1}\|^2/|\Sigma|_{\text{op}}$, see e.g. Section 9.5.1 in [23]. Hence, we expect that the optimal rate of decay should involve d_{Σ}^2 instead of Δ^2/σ^2 , at least for Gaussian mixtures when the sample size is large. Since K -means criterion is tightly linked to isotropic (sub-)Gaussian mixtures, it is unlikely that the relaxed K -means procedure enjoys an exponential decay with respect to d_{Σ}^2 . In addition, we explain in the discussion page 340 at the end of Section 5.1, that the term $m\Delta^4/v^4$ showing up in (8) is a variance term which seems hard to avoid. We emphasize yet that, despite these drawbacks, the exponential decay $e^{-c's^2}$ remains interesting in this setting, as it is “dimensionless,” in the sense that it depends only on the effective rank R_{Σ} of Σ and not on the ambient dimension p .

Perfect recovery. Theorem 1 ensures perfect recovery as soon as $s^2 \gtrsim \log(n) \vee (n/m)$. This requirement exactly matches the requirement derived in Theorem 1 of [44] and it is, to the best of our knowledge, the sharpest known result for polynomial-time algorithms in high-dimension ($p \geq n$). In the large sample size setting where $n \gtrsim p^3 K^2 \log(pK)$, [46] ensures perfect recovery for an iterative spectral clustering under the condition $\Delta^2 \gtrsim \sigma^2(\sqrt{K \log(n)} + \log(n))$ which is weaker than our $s^2 \gtrsim \log(n) \vee (n/m)$ for $K \gg \log(n)$. Again, this lack of optimality is likely to be an artifact of the proof which is only valid when $s^2 \gtrsim K$. Actually, when $K \gg \log(n)$, the condition $s^2 \gtrsim K$ enforces $e^{-c's^2} \ll 1/n$.

3.4. Unequal trace case. As long as Condition (10) of Theorem 1 is satisfied, the misclassification rates $\text{err}(\hat{G}, G)$ is less than $e^{-c's^2}$ as for the equal trace case. Let us first discuss some regimes and choice of $\hat{\Gamma}$ under which (10) is valid, and then discuss the upper-bound.

3.4.1. Choice of $\hat{\Gamma}$. First, observe that, when the covariance matrices Σ_k have unequal traces, uncorrected convex K -means may not satisfy (10) if $\|\Gamma\|_V \geq m\Delta^2/64$. Such a behavior is not an artifact of our proof techniques but is intrinsic for K -means as argued in Proposition 3 of [44].

When the covariance matrices Σ_k have unequal traces, we suggest to use in (9) the estimator $\hat{\Gamma}$ introduced in [11, 44] and defined as follows. For any $a, b \in \{1, \dots, n\}$, let

$$V(a, b) := \max_{c, d \in \{1, \dots, n\} \setminus \{a, b\}} \left| \left\langle X_a - X_b, \frac{X_c - X_d}{\|X_c - X_d\|_2} \right\rangle \right|,$$

denote a measure of dissimilarity between a and b . Then, let

$$\hat{b}_1 = \operatorname{argmin}_{b \in \{1, \dots, n\} \setminus \{a\}} V(a, b) \quad \text{and} \quad \hat{b}_2 := \operatorname{argmin}_{b \in \{1, \dots, n\} \setminus \{a, \hat{b}_1\}} V(a, b)$$

denote the two indices most similar to a with respect to this dissimilarity. Our estimator $\hat{\Gamma}$ is defined by

$$\hat{\Gamma} := \operatorname{Diag}((X_a - X_{\hat{b}_1}, X_a - X_{\hat{b}_2})_{a=1, \dots, n}). \quad (17)$$

When $m > 2$, denoting by $\gamma^2 = L^2 \max_k \operatorname{Tr}(\Sigma_k)$ the maximum scaled trace of the covariances (with L defined in Assumption A1), Proposition 4 of [44] ensures, that with probability higher than $1 - c/n^2$

$$|\hat{\Gamma} - \Gamma|_\infty \lesssim (\sigma^2 \log n + \sigma \gamma \sqrt{\log n}). \quad (18)$$

Hence, Condition (10) holds with probability larger than $1 - c/n^2$ as soon as the condition

$$\Delta^2 \gtrsim \frac{\sigma^2 \log(n) + \gamma \sigma \sqrt{\log(n)}}{m} \quad (19)$$

is met. This condition can be compared to the condition $s^2 \gtrsim n/m$ arising in Theorem 1 by reformulating it as

$$\frac{\Delta^2}{\sigma^2} \wedge \frac{m \Delta^4}{\gamma^2 \sigma^2} \gtrsim \frac{\log(n)}{m}.$$

In particular, we observe that this requirement is weaker than the condition $s^2 \gtrsim n/m$ when

$$|\Sigma_k|_{\text{op}} \operatorname{Tr}(\Sigma_k) \lesssim \frac{n}{\log(n)} |\Sigma_k|_F^2, \quad \text{for all } k = 1, \dots, K. \quad (20)$$

This last condition is mild. For instance, it is met when the ratio between the singular values σ_ℓ/σ_1 of each Σ_k decays faster than $1/\ell^{1+\epsilon}$ with $\epsilon > 0$. It also holds when the condition number of each Σ_k is upper bounded by $n/\log(n)$.

3.4.2. Discussion of the exponential rate. Contrary to the equal trace case, even for spherical covariances $\Sigma_k = \sigma_k^2 I_p$, we cannot argue anymore that our exponential rate s^2 is optimal. Actually, in this case the signal-to-noise ratio s^2 should involve

$$\bar{s}^2 = \min_{j \neq k} \frac{\|\mu_j - \mu_k\|^2}{|\Sigma_k|_{\text{op}} \vee |\Sigma_j|_{\text{op}}} \quad \text{rather than} \quad \frac{\min_{j \neq k} \|\mu_j - \mu_k\|^2}{\max_k |\Sigma_k|_{\text{op}}},$$

at least in the Gaussian case with large sample sizes. Some results on perfect recovery with respect to this signal-to-noise ratio \bar{s}^2 have been derived in [3, 46] in the asymptotic setting where $n \rightarrow \infty$. We do not know whether exponential decay with respect to \bar{s}^2 can be proved for the corrected relaxed K -means.

We emphasize that previously mentioned works [3,37,46] are restricted to perfect recovery in the large sample size setting where (at least) $n \gg p$, while Theorem 1 provides a partial recovery bound that holds under conditions (19) or (20) which are non-asymptotic and can hold in any dimension p . As for [21], they prove exponential decay with respect to $\tilde{s}^2 = \min_{j \neq k} \|\mu_j - \mu_k\|^2 / \max_k |\Sigma_k|_{\text{op}}$ in a higher signal regime and only for clusters of equal sizes.

4. Clustering in Stochastic Block Models (SBM)

4.1. Stochastic Block Model (SBM). We observe an undirected graph with n nodes labeled by $a = 1, \dots, n$. We assume that the edges are independent and that there exists an unknown partition (G_1, \dots, G_K) of the nodes and a symmetric matrix $P \in [0, 1]^{K \times K}$ such the probability to have an edge between $a \in G_k$ and $b \in G_j$ is equal to P_{jk} when $a \neq b$ and 0 when $a = b$. In other words, the adjacency matrix $\mathbf{X} \in \mathbb{R}^{n \times n}$ is symmetric, with zero diagonal and independent lower-diagonal entries fulfilling $\mathbb{E}[\mathbf{X}_{ab}] = P_{jk}$ for any $a \in G_k$ and $b \in G_j$, with $a \neq b$.

Let us denote by m_k the cardinality of G_k . There is a strong interest in the analysis of recovery properties of SDP like

$$\max_{B \in \mathcal{C}'} \langle \mathbf{X}, B \rangle, \quad \text{with } \mathcal{C}' = \left\{ B : \text{PSD}, B_{ab} \geq 0, B_{aa} = 1, |B|_1 = \sum_k m_k^2 \right\}, \quad (21)$$

or variants of it. Such SDP are derived as a convex relaxation of the MLE optimization in the case where within group probability of connection P_{kk} are all equal to p and between group probability of connection P_{jk} are all equal to q with $q < p$, see e.g. [15, 42]. In the assortative setting, where the within group probabilities of connection are larger than the between group probabilities of connection, such SDP enjoy some very nice properties at all regimes, see [15, 18, 20, 24, 25, 42].

Yet, we observe that SDP like (21) seek for partitions maximizing within group connectivity. So they are tied to the assortative setting and we cannot expect some good performances far away from this setting.

Instead of these SDP, we propose to solve (6) or a variation of it, which is a relaxed-version of K -means applied to the adjacency matrix. The heuristic is as follows: two nodes a and b belonging to the same group share the same connectivity profile, that is, up to diagonal terms, the expectation of the columns X_a and X_b are equal. Therefore, it is tempting to recover the groups of the SBM by using distance clustering on the columns of the adjacency matrix \mathbf{X} . In particular, (6) seeks for groups of nodes sharing similar connectivity profiles, instead of groups with maximal within group connectivity. The main difference between (6) and the SDP like (21), is that the maximization is applied to $\mathbf{X}\mathbf{X}^T = \mathbf{X}^2$ instead of \mathbf{X} . So, compared to SDP like (21), the SDP (6) seeks for a partition where the groups have a high density of common neighbors rather than a high density of connections. Hence, the SDP (6) is

not tied to the assortative case and can handle arbitrary matrices P . We also point out, that, contrary to previous SDP investigated in the literature for SBM clustering, the relaxed K -means (6) does not require the knowledge of the size of the groups, nor the knowledge of some parameters of P . Yet, in order to handle appropriately the sparse setting where $|P|_\infty = o(\log(n)/n)$, we need to add a constraint on $|B|_\infty$ in the program (6) to prevent the solution of K -means to produce too unbalanced partitions. As pointed in the previous section, the norm $|B|_\infty$ of a matrix $B \in \mathcal{P}$ corresponding to a true partition is the inverse of the size of the smallest cluster. Thus, adding a constraint on $|B|_\infty$ will avoid the formation of too unbalanced partitions.

Hence, we propose to solve the following constrained version of the relaxed K -means (6)

$$\hat{B} \in \operatorname{argmax}_{B \in \mathcal{C}} \langle \mathbf{X}\mathbf{X}^T, B \rangle, \quad (22)$$

where, for $K/n \leq \alpha \leq 1$,

$$\mathcal{C}_\alpha = \{B \in \mathbb{R}^{n \times n} : \text{Positive Semi Definite, } \operatorname{Tr}(B) = K, B1 = 1, 0 \leq B \leq \alpha\}.$$

As explained in Theorem 2 below, the parameter α can be chosen equal to 1 when $|P|_\infty \geq \log(n)/n$, but its choice is more constrained when $|P|_\infty \leq \log(n)/n$. We explain below Theorem 2 how α can be chosen in a data-driven way.

As for the mixture of sub-Gaussian, the signal strength driving the exponential decay will be related to the Euclidean distance between two rows of the expected adjacency matrix. A row of \mathbf{X} can be written as $X_a = \mu_k - P_{kk}e_a + E_a$, with $[\mu_k]_b = P_{jk}$ and $E_a = X_a - \mathbb{E}[X_a]$ for $a \in G_k$ and $b \in G_j$. The square Euclidean distance between μ_k and μ_j is

$$\|\mu_k - \mu_j\|^2 = \sum_{\ell} m_{\ell} (P_{k\ell} - P_{j\ell})^2 \geq m \|P_{k\cdot} - P_{j\cdot}\|^2.$$

Similarly to the sub-Gaussian setting, we define

$$\Delta^2 := \min_{j \neq k} \Delta_{jk}^2; \quad \Delta_{jk}^2 := \sum_{\ell} m_{\ell} (P_{k\ell} - P_{j\ell})^2 \geq m \|P_{k\cdot} - P_{j\cdot}\|^2, \quad (23)$$

which represents the signal strength in our analysis.

We point out that $\|P_{k\cdot} - P_{j\cdot}\| \geq \sqrt{2} \lambda_{\min}(P)$, with $\lambda_{\min}(P)$ the smallest eigenvalue of P , so that $\Delta \geq \sqrt{2} m \lambda_{\min}(P)$.

Since the variance of a Bernoulli variable with small probability of success is roughly equal to this probability, we control the variance of E_a with the following assumption.

Assumption A1'. We have $|P|_\infty \leq L$.

Under this assumption, we will prove that the misclassification error of the relaxed K -means decreases exponentially fast with the signal-to-noise ratio $s^2 = \Delta^2/L$.

4.2. Main result. We consider a partition \hat{G} obtained from 7-approximate K -medoid solution of \hat{B} as obtained in (22). Next theorem provides an upper bound on the misclassified nodes decreasing exponentially fast with s^2 .

Theorem 2. Assume that Assumption $A1'$ holds and set $s^2 = \Delta^2/L$, with Δ defined by (23). Then, there exist three positive constants c, c', c'' , such that for any $1/m \leq L \leq 1/\log(n)$,

$$\frac{1}{m} \leq \alpha \leq \alpha(L) := \frac{K^3}{n} e^{4nL} \tag{24}$$

and

$$s^2 \geq c''n/m, \tag{25}$$

with probability at least $1 - c/n^2$, the proportion of misclassified nodes is upper bounded by

$$\text{err}(\hat{G}, G) \leq e^{-c's^2}.$$

We observe that we always have

$$\alpha(L) = \frac{K^3}{n} e^{4nL} \geq \frac{K^3}{m} \times \frac{m}{n} e^{4n/m} \geq \frac{10^4}{m},$$

so Condition (24) is non-void. We discuss this condition into more details in the next section.

In practice, we can set α to the value $\hat{\alpha} = \frac{K^3}{n} e^{2nd_{\mathbf{x}}} \wedge 1$, where $d_{\mathbf{x}}$ denotes the density of the graph. Next corollary provides a partial recovery bound when plugin this data driven choice for α in (22).

Corollary 1. Assume that Assumption $A1'$ holds, that $1/m \leq L \leq 1/\log(n)$, that the density $d_{\mathbf{x}}$ fulfills $\mathbb{E}[d_{\mathbf{x}}] \geq n^{-1}$ and

$$\frac{1}{m} \leq \frac{K^3}{n} e^{n\mathbb{E}[d_{\mathbf{x}}]}. \tag{26}$$

Then, there exist three positive constants c, c', c'' , such that when $s^2 \geq c''n/m$, and when setting $\alpha = \hat{\alpha}$ in (22), then $\text{err}(\hat{G}, G) \leq e^{-c's^2}$ with probability at least $1 - c/n^2$.

Indeed, $n(n-1)d_{\mathbf{x}}/2$ is stochastically dominated by a binomial distribution with parameter $(n(n-1)/2, |P|_{\infty})$ so that $d_{\mathbf{x}} \leq 2|P|_{\infty}$ with probability higher than $1 - e^{-0.3n(n-1)|P|_{\infty}}$ and hence $\hat{\alpha} \leq \alpha(L) \wedge 1$. Conversely, the Bernstein inequality together with $\mathbb{E}[d_{\mathbf{x}}] \geq n^{-1}$ ensures that $d_{\mathbf{x}} \geq \mathbb{E}[d_{\mathbf{x}}]/2$ with probability larger than $1 - c/n^2$. Hence (26) ensures that $m^{-1} \leq \hat{\alpha}$ with probability larger than $1 - c/n^2$.

4.3. Discussion. An SDP not tied to the assortative case. Our results cover a wide range of settings going beyond the assortative case usually handled by SDP algorithms. This is due to the fact that usual SDP criteria considered SBM clustering are derived as convex relaxation of MLE in the assortative case; while (6) is derived as a convex relaxation of K -means.

On Condition (24). The constraint $|B|_\infty \leq \alpha(L)$ is needed in our proofs in order to avoid the concentration of \widehat{B} on the high-degree nodes. We observe first that $\alpha(L) \geq 1$ when

$$L \geq \frac{\log(n/K^3)}{4n},$$

in which case we can take $\alpha = 1$, which amounts to remove the constraint.

For smaller value of L the constraint $|B|_\infty \leq \alpha(L)$ becomes active, with $\alpha(L)$ decreasing when L decreases down to the extreme value $L = 1/m$. We emphasize again that the condition on α does not require the knowledge of the true size of the groups but only constrains the size of the smallest group. Besides, Condition (24) can be met as long as $m \geq \frac{n}{K^3} e^{-4nL}$. For L scaling as l_0/n , the size of smallest cluster can still be as small as $\frac{n}{K^3} e^{-4l_0}$ allowing for unbalanced partitions.

Assortative case. To start with, let us make explicit the value of $s^2 = \Delta^2/L$ in the assortative case, with within group probabilities of connection p , between group probabilities of connection q (with $q < p$) and balanced group sizes ($m \approx n/K$). In this case, $s^2 = 2m(p - q)^2/p$, and we obtain the same rate of exponential decay as in [2, 16, 20, 22, 47], but without the tight constants of [22, 47] in the exponential rate. We emphasize yet that we handle unknown group sizes, and with only a mild constraint on the group the sizes. Besides, Theorem 2 ensures perfect recovery for

$$\frac{(p - q)^2}{p} \gtrsim \frac{K(K \vee \log(n))}{n},$$

matching the best known results (up to constants) for polynomial-time algorithms [15].

Partial recovery for the general model. To the best of our knowledge, outside the assortative case, the only other exponentially decaying misclassification error is stated in Theorem 4 in [2] for a quite different procedure. Their results do not cover the same regime as ours, since they focus on the sparse regime where $P = P_0/n$ with P_0 a fixed matrix and $n \rightarrow \infty$. For simplicity, let us discuss again the case of balanced groups where $m \approx n/K$. With our notation, Theorem 4 in [2] shows (under some conditions) that, in the sparse regime, the misclassification error is upper-bounded by $e^{-c\tilde{s}^2}$ where

$$\tilde{s}^2 = \frac{m\lambda_{\min}(P)^2}{K\lambda_{\max}(P)}.$$

Since $\lambda_{\min}^2(P) \leq 2\Delta^2/m$ and $\lambda_{\max}(P) \geq |P|_\infty$, we observe that

$$\tilde{s}^2 \leq \frac{2\Delta^2}{K|P|_\infty} = \frac{2}{K}s^2,$$

so that their exponential decay with respect to \tilde{s}^2 is slower than the exponential decay with respect to s^2 . Yet, this discrepancy between the rates is partly due to the fact that the exponential decay $e^{-c\tilde{s}^2}$ in [2] is valid in regimes where our Condition (25) for Theorem 2 is not met. In other words, when the signal is low, [2] achieves the exponential decay $e^{-c\tilde{s}^2}$ which is not covered by our theory, while for stronger signals (where $s^2 \gtrsim K$ holds), our exponential decay e^{-cs^2} is faster at least by a K factor. We again emphasize that Theorem 2 is also valid in denser regime than that of [2].

Perfect recovery for the general model. From Theorem 2, we derive that relaxed K -means achieves exact recovery as long as

$$s^2 \gtrsim \frac{n}{m} \vee \log(n). \tag{27}$$

Again, the only other results we are aware of in the general model is from [2] where the authors consider the asymptotic regime $P = \bar{P} \log(n)/n$ with \bar{P} (and therefore also K) fixed and $n \rightarrow \infty$. In this setting, they proved that perfect recovery is possible in the balanced case ($m = K/n$) if and only if

$$\lim_{n \rightarrow \infty} \frac{m}{\log(n)} \min_{j \neq k} D_+(P_{j:} || P_{k:}) > 1, \tag{28}$$

where, for two vectors q and p ,

$$D_+(q || p) = \max_{t \in [0,1]} \sum_x p_x \left(1 - t + t \frac{q_x}{p_x} - \left(\frac{q_x}{p_x} \right)^t \right). \tag{29}$$

Since Δ^2 is based on the Euclidean distances between the columns of P instead of D_+ , our results cannot guaranty perfect recovery up to the exact threshold of (28). Yet, in the case where $\min_{j,k} P_{jk} / \max_{j,k} P_{jk}$ is bounded away from zero, we can compare Δ^2 to $\min_{j \neq k} D_+(P_{j:} || P_{k:})$. Actually, according to Lemma 18 in Appendix A, we have

$$D_+(q || p) \leq \frac{1}{4\rho} \sum_x \frac{(p_x - q_x)^2}{p_x}, \quad \text{when } \min_x \frac{q_x}{p_x} \geq \rho > 0.$$

Hence, when $\min_{j,k} P_{jk} \geq \rho \max_{j,k} P_{jk}$, we have

$$m \cdot \min_{j \neq k} D_+(P_{j:} || P_{k:}) \leq \frac{m}{4\rho^2} \min_{j \neq k} \frac{\|P_{j:} - P_{k:}\|^2}{|P|_\infty} = \frac{s^2}{4\rho^2},$$

so that s^2 and $m \cdot \min_{j \neq k} D_+(P_{j:} || P_{k:})$ differs from at most a factor $1/(4\rho^2)$. As a consequence, in the asymptotic setting of [2], Condition (27) achieves up to constants (depending on ρ) the optimal threshold (29).

When the ratio $\max_{j,k} P_{jk} / \min_{j,k} P_{jk}$ is unbounded, our SNR s^2 involving $|P|_\infty$ is no longer optimal. For example, in the weak assortative setting with clusters of

equal sizes, [4] derives some performance bounds which are better than ours when $\max_{jk} P_{jk} / \min_{jk} P_{jk}$ is unbounded.

Finally, we point out that we do not need to de-bias the relaxed K -means as in (9) for the sGMM. This is due to the fact that the size of the bias is small compared to the size of the fluctuations in this setting.

5. Outline of the proofs

We write $B^* \in \mathcal{C}$ for the matrix associated to the true partition G of the data set. Following the definition in Section 2, we have $B_{ab}^* = 0$ unless a and b belong to the same group, in which case $B_{ab}^* = 1/|G_k|$.

5.1. Outline of the proof of Theorem 1. In this section, we describe the main lines of the proof of Theorem 1. We refer to Section 6 for all the details. The proof relies on three main arguments detailed below:

- (1) First, similarly to [14, 20], the misclassification proportion of the final K -medoid clustering (7) can be directly controlled by the ℓ^1 -norm $|\hat{B} - B^*|_1$, see Section 6.3.
- (2) Second, by comparing $\langle \mathbf{X}\mathbf{X}^T - \hat{\Gamma}, B^* \rangle \leq \langle \mathbf{X}\mathbf{X}^T - \hat{\Gamma}, \hat{B} \rangle$, we can upper-bound $|B^* \hat{B} - B^*|_1$, (which is closely related to $|\hat{B} - B^*|_1$) by some “noise” terms.
- (3) Third, a careful analysis of the noise terms provides the claimed result. Following [20], we use the key inequality

$$\sum_{i=1}^n a_i b_i \leq \sum_{i=1}^{|b|_1} a_i$$

for any $a_1 \geq a_2 \geq \dots \geq a_n$ and $b_1, \dots, b_n \in [0, 1]$, combined with tight upper-bounds on sums of ordered statistics. This bound involving ordered statistics is tighter than the classical

$$\sum_{i=1}^n a_i b_i \leq |a|_\infty |b|_1$$

combined with upper-bound on ℓ^∞ -norms. We underline, that this classical reasoning based on ℓ^∞ -norms is not tight enough in order to handle partial recovery results in this setting.

For simplicity, we assume throughout this section that all the inter-cluster distances $\Delta_{jk} = \|\mu_k - \mu_j\|$ are equal to $\Delta = \min_{j \neq k} \Delta_{jk}$, for $j \neq k$. We refer to Section 6 for the general case. As claimed in the first point above, our main task is to prove the bound $|\hat{B} - B^*|_1 \leq n e^{-c's^2}$, with probability at least $1 - c/n^2$.

Let $A \in \mathbb{R}^{n \times K}$ denote the membership matrix $A_{ik} = \mathbf{1}_{i \in G_k}$, let $\boldsymbol{\mu} \in \mathbb{R}^{K \times p}$ be the matrix whose k -th rows is given by μ_k and let $E \in \mathbb{R}^{n \times p}$ be the matrix whose a -th rows is given by E_a . The observed matrix \mathbf{X} can then be written as $\mathbf{X} = A\boldsymbol{\mu} + E$. For any matrix $B \in \mathcal{C}$, expanding the product $\mathbf{X}\mathbf{X}^T$, we can decompose the scalar product

$$\begin{aligned} \langle \mathbf{X}\mathbf{X}^T - \hat{\Gamma}, B^* - B \rangle &= \langle A\boldsymbol{\mu}\boldsymbol{\mu}^T A^T, B^* - B \rangle + \langle EE^T - \Gamma, B^* - B \rangle \\ &\quad + \langle \Gamma - \hat{\Gamma}, B^* - B \rangle + \langle A\boldsymbol{\mu}E^T + E\boldsymbol{\mu}^T A^T, B^* - B \rangle. \end{aligned}$$

The first term can be interpreted as a signal term, which is minimized in \mathcal{C} at $B = B^*$. The three remaining terms involve the noise.

Signal term. Some basic algebra (see Lemma 4 in Section 6) shows that this writes as

$$\langle A\boldsymbol{\mu}\boldsymbol{\mu}^T A^T, B^* - \hat{B} \rangle = \frac{1}{4} \Delta^2 |B^* - B^* \hat{B}|_1. \quad (30)$$

Since $B^* \in \mathcal{C}$, by definition of \hat{B} we have $\langle \mathbf{X}\mathbf{X}^T - \hat{\Gamma}, B^* - \hat{B} \rangle \leq 0$ and hence

$$\frac{1}{4} \Delta^2 |B^* - B^* \hat{B}|_1 \leq \text{noise terms}.$$

With a suitable control of the three noise terms we can therefore hope to get a control on $|B^* - B^* \hat{B}|_1$. Relying on the following inequality

$$|B^* - \hat{B}|_1 \leq \frac{2n}{m} |B^* - B^* \hat{B}|_1 \quad (31)$$

proved in Lemma 1, Section 6, this will allows us in turn to control $|B^* - \hat{B}|_1$. Let us explain how we can control each of the three noise terms, and how they contribute to the final result.

Quadratic terms. The quadratic term $\langle EE^T - \Gamma, \hat{B} - B^* \rangle$ is the most delicate one. We observe first, that for any matrix M , the product $B^* M$ averages the rows of M within the groups

$$[B^* M]_{ab} = \frac{1}{|G_k|} \sum_{c \in G_k} M_{cb}, \quad \text{for all } a \in G_k.$$

As a consequence, the variance of the entries of $B^*(EE^T - \Gamma)$ is reduced by a factor at least \sqrt{m} compared to $EE^T - \Gamma$. So decomposing

$$\begin{aligned} EE^T - \Gamma &= B^*(EE^T - \Gamma) + (EE^T - \Gamma)B^* \\ &\quad - B^*(EE^T - \Gamma)B^* + (I - B^*)(EE^T - \Gamma)(I - B^*), \end{aligned}$$

we observe that we need to control three terms similar to $\langle B^*(EE^T - \Gamma), \hat{B} - B^* \rangle$, and a last term $\langle (I - B^*)(EE^T - \Gamma)(I - B^*), \hat{B} - B^* \rangle$. Since B^* is a projection matrix, this last term involves the projection of $EE^T - \Gamma$ onto the orthogonal of the range of B^* . The three first terms benefits from the averaging effect described above, but not the last one. Instead, we control the last term as in [11, 44]

$$\begin{aligned} & \langle (I - B^*)(EE^T - \Gamma)(I - B^*), \hat{B} - B^* \rangle \\ & \leq |EE^T - \Gamma|_{\text{op}} |(I - B^*)(\hat{B} - B^*)(I - B^*)|_* \\ & \leq \frac{1}{2m} |EE^T - \Gamma|_{\text{op}} |B^* - B^*\hat{B}|_1, \end{aligned} \quad (32)$$

see Lemma 1 for the last inequality. Since $|EE^T - \Gamma|_{\text{op}} \lesssim v^2\sqrt{n} + \sigma^2n$ with probability at least $1 - c/n^2$, we obtain that (32) is upper-bounded by $\Delta^2 |B^* - B^*\hat{B}|_1/16$ under the assumption $s^2 \gtrsim n/m$. So this term is smaller than half of the signal term, and we can remove it at the price of losing a factor 2 in the signal level. We emphasize that the condition $s^2 \gtrsim n/m$ is exactly tailored to get this control, and hence it is fully driven by the upper-bound $|EE^T - \Gamma|_{\text{op}} \lesssim v^2\sqrt{n} + \sigma^2n$.

Let us now turn to the three terms of the form

$$\langle B^*(EE^T - \Gamma), \hat{B} - B^* \rangle = \langle B^*(EE^T - \Gamma), B^*\hat{B} - B^* \rangle.$$

The simple inequality $\langle A, B \rangle \leq |A|_{\infty}|B|_1$ as in [11, 44] leads to the control

$$\langle B^*(EE^T - \Gamma), \hat{B} - B^* \rangle \leq \frac{v^2\sqrt{\log(n)} + \sigma^2\log(n)}{\sqrt{m}} |B^* - B^*\hat{B}|_1, \quad (33)$$

with high probability. This control is good enough to prove perfect recovery at the right scale, but it is too crude in order to exhibit partial recovery rates. Instead, we adapt the clever analysis of [20], which relies on the upper-bound,

$$\langle A, B \rangle \leq \sum_{j=1}^{|B|_1} A_{(j)}, \quad \text{for any } B \text{ with } 0 \leq B_{ab} \leq 1,$$

where $A_{(1)} \geq A_{(2)} \geq \dots$ are the entries of A ranked in decreasing order and

$$\sum_{j=1}^b a_j = a_1 + \dots + a_{[b]} + (b - [b])a_{[b]+1},$$

where $[b]$ is the integer part of b . Lemma 9 based on the Hanson–Wright inequality provides a control of the sum of the ordered statistics of $B^*(EE^T - \Gamma)$ ensuring that with probability at least $1 - c/n^2$, the following inequality holds

$$\begin{aligned} & \langle B^*(EE^T - \Gamma), B - B^* \rangle \\ & \lesssim \left(\frac{v^2}{\sqrt{m}} \sqrt{\log\left(\frac{nK^3}{|B^* - B^*B|_1}\right)} \vee \sigma^2 \log\left(\frac{nK^3}{|B^* - B^*B|_1}\right) \right) |B^* - B^*B|_1, \end{aligned} \quad (34)$$

simultaneously for all $B \in \mathcal{C}$ fulfilling $|B^* - B^*B|_1 \lesssim m$. The main difference compared to (33), is that the $\log(n)$ has been replaced by something of the form $\log(n/|B^* - B^*B|_1)$. As it will appear clearly in the last step, moving from $\log(n)$ to $\log(n/|B^* - B^*B|_1)$ is the key to obtain the control $ne^{-c's^2}$ on $|B^* - B^*\hat{B}|_1$. We point out that (34) is only guaranteed for $B \in \mathcal{C}$ fulfilling $|B^* - B^*B|_1 \lesssim m$. So we need to first get such a bound.

Contrary to [20], we do not use Grothendieck's inequality for a preliminary control, but instead we apply a first time our upper-bound in order to get the rough bound $|B^* - B^*B|_1 \leq nK^3e^{-c\sqrt{n/m}} \lesssim m$ when $s^2 \gtrsim m/n$ and then apply again our analysis by using $|B^* - B^*B|_1 \lesssim m$. We refer to Section 6.2 below Lemma 8 (page 345) for the details.

Gamma term. The term $\langle \Gamma - \hat{\Gamma}, \hat{B} - B^* \rangle$ can be directly controlled by Inequality (B12) of [44], which is recalled in Lemma 5 for convenience

$$\langle \Gamma - \hat{\Gamma}, \hat{B} - B^* \rangle \leq \frac{2}{m} |\Gamma - \hat{\Gamma}|_V |B^* - B^*\hat{B}|_1.$$

Condition (10) exactly ensures that the right hand side is upper-bounded by

$$\Delta^2 |B^* - B^*\hat{B}|_1 / 32.$$

So this term is smaller than a quarter of the signal term, and we can again remove it at the price of loosing another factor 2 in the signal level.

Cross-products term. It remains to upper-bound the cross-products term

$$\langle A\mu E^T + E\mu^T A^T, \hat{B} - B^* \rangle = \sum_{j \neq k} \sum_{a \in G_k, b \in G_j} \langle E_a - E_b, \mu_k - \mu_j \rangle B_{ab}.$$

We recall that we describe here the case where the Δ_{jk} are all equal to Δ for $j \neq k$. The general case is treated in Section 6. We observe first that $\langle E_a, \mu_k - \mu_j \rangle$ is sub-Gaussian $\text{SubG}(\sigma^2 \Delta^2)$ and

$$\sum_{j \neq k} \sum_{a \in G_k, b \in G_j} |B_{ab}| = \frac{1}{2} |B^* - B^*\hat{B}|_1,$$

see Lemma 1 for this last equality. Hence, building again on the inequality $\langle A, B \rangle \leq \sum_{j=1}^{|B|_1} A_{(j)}$ and deviation bounds for sub-Gaussian random variables, we obtain that

$$\langle A\mu E^T + E\mu^T A^T, \hat{B} - B^* \rangle \leq |B^* - B^*\hat{B}|_1 \sqrt{\sigma^2 \Delta^2 \log\left(\frac{nK^3}{|B^* - B^*\hat{B}|_1}\right)}, \quad (35)$$

with probability at least $1 - c/n^2$.

Conclusion. Take for granted that we have $|B^* - B^* \hat{B}|_1 \lesssim m$. Then combining the above bounds leads to

$$\begin{aligned} & \Delta^2 |B^* - B^* \hat{B}|_1 \\ & \lesssim \left(\frac{\nu^2}{\sqrt{m}} \sqrt{\log \left(\frac{nK^3}{|B^* - B^* \hat{B}|_1} \right)} \vee \sigma^2 \log \left(\frac{nK^3}{|B^* - B^* \hat{B}|_1} \right) \right) |B^* - B^* \hat{B}|_1, \end{aligned}$$

with probability at least $1 - c''/n^2$. This bound can be rewritten as

$$s^2 \lesssim \log \left(\frac{nK^3}{|B^* - B^* \hat{B}|_1} \right)$$

and hence $|B^* - B^* \hat{B}|_1 \leq nK^3 e^{-c's^2}$.

In light of the three last bounds, we can traceback the terms contributing to the two parts of $s^2 = (\Delta^2/\sigma^2) \wedge (m\Delta^4/\nu^4)$. The second term $m\Delta^4/\nu^4$ in s^2 comes from the first term in the upper-bound (34). The ratio ν^2/\sqrt{m} in (34) is driven by the variance term in the Hanson–Wright inequality, so there is very little room for possible improvement. The first term Δ^2/σ^2 in s^2 mainly comes from (35), and secondary from the second term in (34). As above, the term $\Delta\sigma$ arising in (35) is a variance term, so there is again very little room for possible improvement on this side.

Recalling Inequality (31), the bound $|B^* - B^* \hat{B}|_1 \leq nK^3 e^{-c's^2}$ obtained above ensures that $|B^* - \hat{B}|_1 \leq 2nK^3(n/m)e^{-c's^2}$. This last inequality does not seem to meet our expectations. Yet, when $s^2 \gtrsim n/m$, it enforces the targeted inequality $|B^* - \hat{B}|_1 \leq ne^{-c's^2}$ for some smaller constant c' .

5.2. Outline of the proof of Theorem 2. Let us define D the diagonal matrix with $D_{aa} = P_{kk}$ for $a \in G_k$ and $\mathbf{X}' = \mathbf{X} + D$. We observe that

$$|\mathbf{X}\mathbf{X}^T - \mathbf{X}'(\mathbf{X}')^T|_\infty \leq |D|_\infty (|\mathbf{X}|_\infty + |\mathbf{X}'|_\infty) \leq 2L,$$

so, when $\Delta^2 \gtrsim Ln/m$ we have

$$\begin{aligned} |\langle \mathbf{X}\mathbf{X}^T - \mathbf{X}'(\mathbf{X}')^T, \hat{B} - B^* \rangle| & \leq 2L |B^* - \hat{B}|_1 \\ & \leq 4L \frac{n}{m} |B^* - B^* \hat{B}|_1 \leq 0.05 \Delta^2 |B^* - B^* \hat{B}|_1. \end{aligned}$$

So we can replace \mathbf{X} by \mathbf{X}' in our analysis, since this term is smaller than a fraction of the signal term.

Since $\mathbf{X}'_a = \mu_k + E_a$, with $[\mu_k]_b = P_{jk}$ and $E_{ab} = X_{ab} - \mathbb{E}[X_{ab}]$ for $b \in G_j$, the proof of Theorem 2 follows similar lines as that of Theorem 1. The main differences lies in the symmetry of \mathbf{X} and the different stochastic control of the

Bernoulli variables. We briefly sketch the main lines below. Again, for simplicity, we assume in this section that $\Delta_{jk} = \Delta$ for all $j \neq k$.

As in the previous section, we have

$$\begin{aligned} \langle \mathbf{X}'(\mathbf{X}')^T, B^* - \hat{B} \rangle &= \frac{1}{4} \Delta^2 |B^* - B^* \hat{B}|_1 \\ &\quad + \langle EE^T, B^* - \hat{B} \rangle + \langle A\mu E^T + E\mu^T A^T, B^* - \hat{B} \rangle. \end{aligned}$$

Cross-product terms. The cross-product terms are handled similarly as before, the main difference is that we rely on Bernstein’s inequality instead of sub-Gaussian deviations, producing an additional term. Since $\text{var}(E_{ai}) \leq L$ and $|\mu|_\infty \leq L$, we get the same bound as before, with σ^2 replaced by L and the additional term equal to L times the sum of $|B^* - B^* \hat{B}|_1$ ordered exponential random variables

$$\begin{aligned} \langle A\mu E^T + E\mu^T A^T, \hat{B} - B^* \rangle &\leq |B^* - B^* \hat{B}|_1 \sqrt{L \Delta^2 \log \left(\frac{nK^3}{|B^* - B^* \hat{B}|_1} \right)} \\ &\quad + L |B^* - B^* \hat{B}|_1 \log \left(\frac{nK^3}{|B^* - B^* \hat{B}|_1} \right). \quad (36) \end{aligned}$$

Quadratic terms. We use the same decomposition as before, except that Γ is absent here. Actually the operator norm of Γ is upper bounded by nL which is smaller than the size of the fluctuations of EE^T around Γ .

First, we consider the expression involving $(I - B^*)EE^T(I - B^*)$. When $L \geq \log(n)/n$, we can directly use the bound (32) since, in that regime, $|EE^T|_{\text{op}} \leq cnL$ (see e.g. [32]). Unfortunately, for a smaller L , the operator norm $|EE^T|_{\text{op}}$ can only be upper-bounded by $nL + \log(n)$ (up to a multiplicative constant). Actually, even for L of order $1/n$, the nodes with the highest degree enforces an operator norm of size at least $\log(n)/\log \log(n)$. Being compelled to follow an alternative approach for bounding this expression, we use, classically for sparse graphs, a trimming argument which amounts to remove high degree nodes to the adjacency matrix. For the trimmed adjacency matrix we can now apply the bound (32) since the operator norm of this trimmed matrix is at most cnL . Then, it remains to upper bound the residual term by a ℓ^1/ℓ^∞ bound. Relying on the box-constraint $|B|_\infty \leq \alpha(L)$, we will then guaranty that this residual remains under control.

As for the term $\langle B^*EE^T, \hat{B} - B^* \rangle$, we need to control some quadratic forms of centered Bernoulli variables. We get a control of the right order by splitting them into pieces and considering apart different sub-cases. The symmetry of \mathbf{X} induces some interlaced dependencies that must be handled with care. It is the main hurdle of the proof.

At the end of the day, we obtain a bound of the form

$$\langle EE^T, B^* - \widehat{B} \rangle \leq 0.05\Delta^2 |B^* - B^* \widehat{B}|_1 + L |B^* - B^* \widehat{B}|_1 \log \left(\frac{nK^3}{|B^* - B^* \widehat{B}|_1} \right),$$

for the quadratic terms. Then, the end of the proof of Theorem 2 follows the same line as those of Theorem 1.

6. Proof of Theorem 1

We provide in this section the full proof of Theorem 1. We recall that $m = \min_k |G_k|$ stands for the size of the smallest group. In the sequel, B_{aG_k} stands for a -th row of B restricted to the column in G_k whereas $B_{G_k G_l}$ stands for the restriction of B to rows in G_k and columns in G_l . For two indices j and k , $(-1)^{j=k}$ equals -1 when $j = k$ and 1 when $j \neq k$.

6.1. A few useful formulas. We start by gathering some useful formulas, in particular relating $|B^* - B|_1$ to $|B^* - B^* B|_1$.

Lemma 1. *We have for any $B \in \mathcal{C}$*

$$\begin{aligned} (B^* - B^* B)_{ab} &= \frac{1}{m_k} - \frac{1}{m_k} \sum_{c \in G_k} B_{cb} && \text{if } a, b \in G_k, \\ &= -\frac{1}{m_k} \sum_{c \in G_k} B_{cb} && \text{if } a \in G_k, b \notin G_k \end{aligned}$$

and

$$|B^* - B^* B B^*|_1 = |B^* - B^* B|_1 = 2 \sum_{j \neq k} |B_{G_j G_k}|_1,$$

and

$$|(I - B^*) B (I - B^*)|_* \leq \frac{|B^* - B^* B|_1}{2m}$$

and

$$|B^* - B|_1 \leq \frac{2n}{m} |B^* - B^* B|_1.$$

Besides, for any $n \times n$ matrix B , we have

$$|B^* - B^* B|_1 \vee |B^* - B B^*|_1 \leq |B^* - B|_1.$$

Proof of Lemma 1. The two first displays follows from direct computations. The third display is given by (57) in [11]. For the next to last display, we observe that

$$|(I - B^*) B (I - B^*)|_1 \leq n |(I - B^*) B (I - B^*)|_* \leq \frac{n}{2m} |B^* - B^* B|_1.$$

The claim follows from

$$B = B^*B + (I - B^*)B(I - B^*) + BB^* - B^* + B^* - B^*BB^*$$

and hence

$$|B^* - B|_1 \leq |B^* - B^*B|_1 + |(I - B^*)B(I - B^*)|_1 + |BB^* - B^*|_1 + |B^* - B^*BB^*|_1.$$

For the last display, we use that for $a \in G_k$,

$$(B^*M)_{ab} = |G_k|^{-1} \sum_{c \in G_k} M_{cb}.$$

By the triangular inequality, this implies that

$$|B^*M|_1 = \sum_{k=1}^K \sum_{a \in G_k} \sum_{b=1}^n \left| |G_k|^{-1} \sum_{c \in G_k} M_{cb} \right| \leq |M|_1.$$

Similarly $|MB^*|_1 \leq |M|_1$ and the last display follows by taking $M = B^* - B$ and using $(B^*)^2 = B^*$. \square

The next two lemmas recall two useful probabilistic bounds.

Lemma 2 (Hanson–Wright inequality). *Let ε be the vector obtained by concatenation of $\Sigma_{k(1)}^{-1/2} E_1, \dots, \Sigma_{k(n)}^{-1/2} E_n$.*

Under Assumption A1, the random vector ε is sub-Gaussian $\text{SubG}(L^2 I_{np})$ and for all $t > 0$

$$\mathbb{P} [\varepsilon^T A \varepsilon - \mathbb{E}[\varepsilon^T A \varepsilon] \geq L^2 (|A|_F \sqrt{t} + |A|_{\text{op}} t)] \leq e^{-ct}.$$

We refer to [45] for a proof of this lemma. Next lemma rephrases Lemma A1 in [44].

Lemma 3. *Let E be the $n \times p$ matrix $E^T = [E_1, \dots, E_n]$ Under Assumption A1, we have for all $t > 0$*

$$\mathbb{P} [|EE^T - \mathbb{E}[EE^T]|_{\text{op}} \geq v^2 \sqrt{t} + \sigma^2 t] \leq 2 \times 9^n e^{-ct}.$$

6.2. Bounding $|B^* - B^* \hat{B}|_1$. We recall that $s^2 = \min(\Delta^2/\sigma^2, m\Delta^4/v^4)$. As explained in Section 5.1, the main step in the proof of Theorem 1 is to prove that for $s^2 \gtrsim \frac{n}{m}$, we have with probability at least $1 - c/n^2$

$$|B^* - B^* \hat{B}|_1 \leq nK^3 \exp(-cs^2). \tag{37}$$

For any $B \in \mathcal{C}$, we have the decomposition

$$\langle \mathbf{X}\mathbf{X}^T - \hat{\Gamma}, B^* - B \rangle = \langle S + \Gamma - \hat{\Gamma} + W + W', B^* - B \rangle$$

where for $a \in G_k, b \in G_j$,

$$\begin{aligned} S_{ab} &= -0.5\|\mu_k - \mu_j\|^2 \\ W'_{ab} &= \langle E_a - E_b, \mu_k - \mu_j \rangle \\ W_{ab} &= \langle E_a, E_b \rangle - \mathbb{E}[\langle E_a, E_b \rangle]. \end{aligned}$$

Since $B^* \in \mathcal{C}$, we have $\langle \mathbf{X}\mathbf{X}^T - \hat{\Gamma}, B^* - \hat{B} \rangle \leq 0$ by definition of \hat{B} , and hence

$$\langle S, B^* - \hat{B} \rangle \leq \langle \Gamma - \hat{\Gamma}, \hat{B} - B^* \rangle + \langle W, \hat{B} - B^* \rangle + \langle W', \hat{B} - B^* \rangle. \quad (38)$$

For the term in the left-hand side of (38), direct computations combined with Lemma 1 give the following evaluation of the signal.

Lemma 4. *We set $b_{jk} = |B_{G_k G_j}|_1$ and $\Delta_{jk} = \|\mu_k - \mu_j\|$. Then, for any $B \in \mathcal{C}$, we have*

$$\langle S, B^* - B \rangle = 0.5 \sum_{j \neq k} \Delta_{jk}^2 b_{jk} \geq \frac{\Delta^2}{4} |B^* - B^* B|_1.$$

For the first term in the right-hand side of (38), we lift from [44] the next lemma (see (B12) in [44]). Recall that $|D|_V = \max_a D_{aa} - \min_a D_{aa}$.

Lemma 5. *For any diagonal $\hat{\Gamma}$, and any $B \in \mathcal{C}$, we have*

$$\langle \Gamma - \hat{\Gamma}, B - B^* \rangle \leq \frac{2}{m} |\Gamma - \hat{\Gamma}|_V |B^* - B^* B|_1.$$

For the second term in the right-hand side of (38), as explained in the Section 5.1, we decompose W into $W = (I - B^*)W(I - B^*) + B^*W + WB^* - B^*WB^*$.

In order to control the scalar product involving $(I - B^*)W(I - B^*)$, we get by combining Lemma 3 and Lemma 1 the following bound.

Lemma 6. *Under Assumption A1, we have with probability at least $1 - c/n^2$*

$$\langle (I - B^*)W(I - B^*), B - B^* \rangle \leq \frac{v^2 \sqrt{n} + \sigma^2 n}{m} |B^* - B^* B|_1,$$

simultaneously for all $B \in \mathcal{C}$.

Hence, when $s^2 \gtrsim \frac{n}{m}$ and when (10) holds, we have

$$\langle (I - B^*)W(I - B^*), \hat{B} - B^* \rangle + \langle \Gamma - \hat{\Gamma}, \hat{B} - B^* \rangle \leq 0.75 \langle S, B^* - \hat{B} \rangle.$$

The remaining terms involving W are controlled by the next lemma proved in Section 6.2.1.

Lemma 7. *We set $\delta = |B^*(B^* - B)|_1$ and we assume that Assumption A1 holds. Then, with probability at least $1 - c/n^2$, we have for all $B \in \mathcal{C}$*

$$\langle B^*W, B - B^* \rangle \lesssim \frac{\delta}{\sqrt{m}} (v^2 \sqrt{\log(nK^3/\delta)} + \sigma^2 \sqrt{\delta} \vee 1 \log(nK^3/\delta)).$$

The same bound holds for $\langle B^*WB^*, B - B^* \rangle$.

It remains to control the last term in the right-hand side of (38) with the following lemma, proved in Section 6.2.2.

Lemma 8. *We set $\delta = |B^*(B^* - B)|_1$. Under Assumption A1, with probability at least $1 - c/n^2$, we have for all $B \in \mathcal{C}$*

$$\begin{aligned} \langle W', B - B^* \rangle &\lesssim \sigma \sum_{j \neq k} \Delta_{jk} b_{jk} \sqrt{\log(nK/b_{jk})} \\ &\lesssim \sigma \sqrt{\langle S, B^* - B \rangle} \sqrt{\delta \log(nK^3/\delta)}. \end{aligned}$$

Conclusion. Focusing on \hat{B} , we set $\delta = |B^*(B^* - \hat{B})|_1$. Combining (38) with

$$\begin{aligned} \langle W, \hat{B} - B^* \rangle &= \langle (I - B^*)W(I - B^*), \hat{B} - B^* \rangle \\ &\quad + 2\langle B^*W, \hat{B} - B^* \rangle - \langle B^*WB^*, \hat{B} - B^* \rangle, \end{aligned}$$

and the five previous lemmas, we obtain that when $s^2 \gtrsim \frac{n}{m}$ and when (10) holds, with probability at least $1 - c/n^2$

$$\begin{aligned} \langle S, B^* - \hat{B} \rangle &\lesssim \sigma \sqrt{\langle S, B^* - \hat{B} \rangle} \sqrt{\delta \log(nK^3/\delta)} \\ &\quad \vee \frac{\delta}{\sqrt{m}} (\nu^2 \sqrt{\log(nK^3/\delta)} + \sigma^2 \sqrt{\delta} \vee 1 \log(nK^3/\delta)). \end{aligned}$$

According to Lemma 4, we have $\langle S, B^* - \hat{B} \rangle \geq \Delta^2 \delta / 4$, and hence the previous bound ensures that

$$\begin{aligned} \Delta^2 &\lesssim \sigma^2 \log(nK^3/\delta) \vee \frac{1}{\sqrt{m}} (\nu^2 \sqrt{\log(nK^3/\delta)} \vee \sigma^2 \sqrt{\delta} \vee 1 \log(nK^3/\delta)) \\ &\lesssim \sigma^2 \log(nK^3/\delta) \vee \frac{1}{\sqrt{m}} (\nu^2 \sqrt{\log(nK^3/\delta)} \vee \sigma^2 \sqrt{\delta} \log(nK^3/\delta)). \end{aligned} \quad (39)$$

Since $|B^*|_1 = |\hat{B}|_1 = n$ and $\delta \leq |B^* - \hat{B}|_1$ (see Lemma 1), we always have $\delta \leq 2n$ by the triangular inequality. Let us prove that we actually have $\delta \lesssim m$. From (39) and $\delta \leq 2n$, we obtain

$$\Delta^2 \lesssim \sigma^2 \sqrt{\frac{n}{m}} \log(nK^3/\delta) \vee \frac{1}{\sqrt{m}} (\nu^2 \sqrt{\log(nK^3/\delta)})$$

and hence

$$\begin{aligned} \delta &\leq nK^3 \exp\left(-c\left(\sqrt{\frac{m}{n}} \frac{\Delta^2}{\sigma^2}\right) \wedge \frac{m\Delta^4}{\nu^4}\right) \\ &\leq nK^3 \exp\left(-c\sqrt{\frac{m}{n}} s^2\right) \leq nK^3 \exp\left(-c'\sqrt{\frac{n}{m}}\right), \end{aligned}$$

where the last bound comes from $s^2 \gtrsim m/n$. Hence,

$$\delta \lesssim m (n/m)^4 \exp(-c' \sqrt{n/m}) \lesssim m.$$

We can now conclude. Since $\delta \lesssim m$, the bound (39) gives

$$\Delta^2 \lesssim \sigma^2 \log(nK^3/\delta) \vee \frac{1}{\sqrt{m}} (v^2 \sqrt{\log(nK^3/\delta)})$$

from which follows

$$\delta \leq nK^3 \exp(-cs^2).$$

The proof of (37) is complete.

6.2.1. Proof of Lemma 7. In the following we use the notation $m_k = |G_k|$, and hence $m = \min_k m_k$. Since B^* is a projection matrix, we observe that

$$\langle B^*W, B - B^* \rangle = \langle B^*W, B^*B - B^* \rangle.$$

We have

$$\langle B^*W, B^*B - B^* \rangle = \sum_{k,j=1}^K \sum_{b \in G_j} z_b^{(j,k)} \beta_{kb},$$

where $\beta_{kb} = m_k |(B^* - B^*B)_{ab}|$ with $a \in G_k$ and

$$z_b^{(j,k)} = \frac{(-1)^{j=k}}{m_k} \sum_{a \in G_k} W_{ab}, \quad \text{for } b \in G_j.$$

We observe that $0 \leq \beta_{kb} \leq 1$ and $|\beta_{kG_j}|_1 = |(B^* - B^*B)_{G_k G_j}|_1 =: b_{jk}$. Hence, writing $z_{(1)}^{(j,k)} \geq z_{(2)}^{(j,k)} \geq \dots$ for the sequence $\{z_b^{(j,k)} : b \in G_j\}$ ranked in decreasing order, we have

$$\langle B^*W, B^*B - B^* \rangle \leq \sum_{j,k=1}^K \sum_{u=1}^{b_{jk}} z_{(u)}^{(j,k)},$$

with the convention that for $b = r + f$ with r integer and $0 \leq f < 1$,

$$\sum_{u=1}^b a_u = (a_1 + \dots + a_r) + f a_{r+1} \leq \left(\sum_{u=1}^r a_u \right) \vee \left(\sum_{u=1}^{r+1} a_u \right). \quad (40)$$

and for $0 \leq b < 1$, $\sum_{u=1}^b a_u \leq b a_{(1)}$.

We control the sum of ordered statistics by the next lemma proved at the end of this section.

Lemma 9. For any integer q in $[1, m_j]$ and $t \geq 0$, we have

$$\mathbb{P} \left[\sum_{u=1}^q z_{(u)}^{(j,k)} \gtrsim \sqrt{\frac{q}{m}} (v^2 \sqrt{t} + \sigma^2 t) \right] \leq C_{m_j}^q e^{-ct}.$$

Let us choose $t_q = c'' q \log(nK/q)$. Since $C_n^q \leq (en/q)^q$,

$$\sum_{q=1}^{m_j} C_{m_j}^q e^{-ct_q} \leq \sum_{q=1}^n e^{-c' q \log(nK/q)} \lesssim \frac{1}{(nK)^2},$$

where the last bound can be obtained e.g. by

$$\sum_{q=1}^n e^{-c' q \log(nK/q)} \leq \sum_{q=1}^{\sqrt{n}} e^{-0.5c' q \log(nK)} + \sum_{q=\sqrt{n}}^n e^{-c' q \log(K)} \lesssim \frac{1}{(nK)^2}.$$

Hence, with probability at least $1 - c/(nK)^2$, we have simultaneously for all integers $1 \leq q \leq m_j$,

$$\sum_{u=1}^q z_{(u)}^{(j,k)} \lesssim q \sqrt{\frac{1}{m}} (v^2 \sqrt{\log(nK/q)} + \sigma^2 \sqrt{q} \log(nK/q)).$$

From (40), we deduce that

$$\sum_{u=1}^{b_{jk}} z_{(u)}^{(j,k)} \lesssim b_{jk} \sqrt{\frac{1}{m}} (v^2 \sqrt{\log(nK/b_{jk})} + \sigma^2 \sqrt{b_{jk} \vee 1} \log(nK/(b_{jk} \vee 1))).$$

With a union bound over $j, k = 1, \dots, K$ we obtain that the inequality above holds simultaneously for all j, k with probability at least $1 - c/n^2$. As a consequence, with Jensen's inequality and $\sqrt{b_{jk}} \leq \sqrt{\delta}$, we get

$$\begin{aligned} \langle B^* W, B^* B - B^* \rangle &\lesssim \sum_{j,k=1}^K b_{jk} \sqrt{\frac{1}{m}} (v^2 \sqrt{\log(nK/b_{jk})} + \sigma^2 \sqrt{\delta \vee 1} \log(nK/b_{jk})) \\ &\lesssim \frac{\delta}{\sqrt{m}} (v^2 \sqrt{\log(nK^3/\delta)} + \sigma^2 \sqrt{\delta \vee 1} \log(nK^3/\delta)). \end{aligned}$$

The term $\langle B^* W B^*, B^* - B \rangle = \langle B^* W B^*, B^* (B^* - B) B^* \rangle$ can be handled in the same way as $\langle B^* W, B - B^* \rangle$, by noticing that $|B^* (B^* - B) B^*|_1 = |B^* (B^* - B)|_1$ according to Lemma 1.

Proof of Lemma 9. Let ε be the vector obtained by concatenation of

$$\Sigma_{k(1)}^{-1/2} E_1, \dots, \Sigma_{k(n)}^{-1/2} E_n,$$

which is sub-Gaussian $\text{SubG}(L^2 I_{np})$. Let \otimes refer to the Kronecker product. For a subset $Q \subset G_j$ with cardinality q , we have

$$\sum_{b \in Q} z_b^{(j,k)} = \frac{(-1)^{j=k}}{m_k} \sum_{a \in G_k, b \in Q} (E_a^T E_b - \mathbb{E}[E_a^T E_b]).$$

For a subset $Q \subset \{1, \dots, n\}$, define $1_Q \in \{0, 1\}^p$ such that $(1_Q)_i = 1$ if $i \in Q$. Define

$$A = \frac{(-1)^{j=k}}{m_k} 1_{G_k} 1_Q^T$$

fulfilling $|A|_F = |A|_{\text{op}} = \sqrt{q/m_k}$. Then, with $\Sigma^{(j,k)} = \Sigma_k^{1/2} \Sigma_j^{1/2}$, we have

$$\sum_{b \in Q} z_b^{(j,k)} = \varepsilon^T (\Sigma^{(j,k)} \otimes A) \varepsilon - \mathbb{E} [\varepsilon^T (\Sigma^{(j,k)} \otimes A) \varepsilon].$$

Hence, since $|\Sigma^{(j,k)}|_F \leq v^2/L^2$ and $|\Sigma^{(j,k)}|_{\text{op}} \leq \sigma^2/L^2$, we have

$$|\Sigma^{(j,k)} \otimes A|_F \leq v^2 \sqrt{q}/(L^2 \sqrt{m}) \quad \text{and} \quad |\Sigma^{(j,k)} \otimes A|_{\text{op}} \leq \sigma^2 \sqrt{q}/(L^2 \sqrt{m}).$$

The bound then follows from Lemma 2, and the $C_{m_j}^q$ possible choices of subset Q . \square

6.2.2. Proof of Lemma 8. To start with, we observe that

$$\langle W', B - B^* \rangle = \sum_{j \neq k} \sum_{a \in G_k, b \in G_j} \langle E_a - E_b, \mu_k - \mu_j \rangle B_{ab}.$$

By symmetry, it is enough to control $\sum_{j \neq k} \sum_{a \in G_k, b \in G_j} \langle E_a, \mu_k - \mu_j \rangle B_{ab}$. Define $b_{jk} = |B_{G_k G_j}|_1$. Since $|B_{a G_j}|_1 \leq 1$, we have

$$\begin{aligned} \sum_{j \neq k} \sum_{a \in G_k, b \in G_j} \langle E_a, \mu_k - \mu_j \rangle B_{ab} &= \sum_{j \neq k} \sum_{a \in G_k} \langle E_a, \mu_k - \mu_j \rangle |B_{a G_j}|_1 \\ &\leq \sum_{j \neq k} \sum_{a=1}^{b_{jk}} z_{(a)}^{(j,k)}, \end{aligned}$$

where $z_{(1)}^{(j,k)} \geq z_{(2)}^{(j,k)} \geq \dots$ corresponds to the values $(z_a^{(j,k)} = \langle E_a, \mu_k - \mu_j \rangle : a \in G_k)$ ranked in decreasing order. The next lemma, proved at the end of this section, provides a control on the sum of the ordered statistics.

Lemma 10. *Under Assumption A1, with probability at least $1 - c/n^2$, simultaneously for all $j \neq k$ and all integers $1 \leq q \leq m_k$*

$$\sum_{a=1}^q z_{(a)}^{(j,k)} \lesssim \sigma \Delta_{jk} q \sqrt{\log(nK/q)}.$$

As a consequence, the Cauchy–Schwarz and Jensen inequalities, and (40) ensure that

$$\begin{aligned} \langle W', B - B^* \rangle &\lesssim \sigma \sum_{j \neq k} \Delta_{jk} b_{jk} \sqrt{\log(nK/b_{jk})} \\ &\lesssim \sigma \sqrt{\sum_{j \neq k} \Delta_{jk}^2 b_{jk}} \sqrt{\delta \sum_{j \neq k} \delta^{-1} b_{jk} \log(nK/b_{jk})} \\ &\lesssim \sigma \sqrt{\langle S, B^* - B \rangle} \sqrt{\delta \log\left(\sum_{j \neq k} nK/\delta\right)} \\ &\leq \sigma \sqrt{\langle S, B^* - B \rangle} \sqrt{\delta \log(nK^3/\delta)}, \end{aligned}$$

which is the inequality claimed in Lemma 8.

It remains to prove Lemma 10. Let $q \leq m_k$ denote a positive integer. Since the $(z_a^{(k,j)} : a \in G_k)$ are independent and sub-Gaussian $\text{SubG}(\sigma^2 \Delta_{jk}^2)$, for any (a_1, \dots, a_q) we have $\sum_{i=1}^q z_{a_i}^{(j,k)}$ sub-Gaussian $\text{SubG}(q\sigma^2 \Delta_{jk}^2)$ under Assumption A1. Hence

$$\mathbb{P}\left[\sum_{a=1}^q z_{(a)}^{(j,k)} > t\right] \leq C_{m_k}^q e^{-t^2/(2\sigma^2 \Delta_{jk}^2 q_{jk})}.$$

For $t_{q_{jk}} = c' \sigma \Delta_{jk} q_{jk} \sqrt{\log(nK/q_{jk})}$, we then have

$$\begin{aligned} \mathbb{P}\left[\exists q_{jk} : \sum_{a=1}^{q_{jk}} z_{(a)}^{(j,k)} > c' \sigma \Delta_{jk} q_{jk} \sqrt{\log(nK/q_{jk})}\right] \\ \leq \sum_{j \neq k} \sum_{q_{j,k}=1}^{m_k} C_{m_k}^{q_{jk}} e^{-c q_{jk} \log(nK/q_{jk})} \\ \lesssim \frac{K^2}{(nK)^c} \lesssim \frac{1}{n^2} \end{aligned}$$

for $c \geq 2$.

6.3. Final clustering bound. We recall that, similarly to [20], the final clustering is provided by a 7-approximate K -medoids (7) on the rows of the matrix \hat{B} output by (6) or (9). Next lemma connects the misclassification error $\text{err}(\hat{G}, G)$ defined by (2) to the ℓ^1 -norm $|B^* \hat{B} - B^*|_1$.

Lemma 11. *The proportion $\text{err}(\hat{G}, G)$ of misclassified points is upper bounded by*

$$\text{err}(\hat{G}, G) \leq 60 \left(\frac{n}{m}\right)^2 \frac{|B^* \hat{B} - B^*|_1}{n}.$$

Combining Lemma 11 and (37), we get that, when $s^2 \gtrsim \frac{n}{m}$, with probability at least $1 - c/n^2$,

$$\text{err}(\hat{G}, G) \leq 60 \left(\frac{n}{m}\right)^5 e^{-c's^2} \leq e^{-c''s^2}.$$

Theorem 1 then follows.

6.3.1. Proof of Lemma 11. The proof of Lemma 11 is close to the proof of Proposition 3 in [20]. We sketch below the main lines of this proof, referring to [20] when the arguments are the same.

We define $\tilde{B} = \hat{A}\hat{M}$, with (\hat{A}, \hat{M}) obtained in (7) and we define $A^* \in \mathbb{R}^{n \times k}$ by $A_{ak}^* = \mathbf{1}_{a \in G_k}$. We also define the sets

$$T_k = \{a \in G_k : |B_{a:}^* - (\tilde{B}B^*)_{a:}|_1 < 1\}, \quad S_k = G_k \setminus T_k, \quad R_1 = \{k : T_k = \emptyset\}, \\ R_2 = \{k : T_k \neq \emptyset \text{ and } \hat{A}_{a:} = \hat{A}_{b:} \forall a, b \in T_k\}, \text{ and } R_3 = \{1, \dots, K\} \setminus (R_1 \cup R_2).$$

The proof is decomposed into 4 steps.

Step 1: For $a \in T_k$ and $b \in T_j$ with $k \neq j$, we have $|B_{a:}^* - B_{b:}^*|_1 \geq 2$ since $a \in G_k$ and $b \in G_j$, and $|B_{a:}^* - (\tilde{B}B^*)_{a:}|_1 + |B_{b:}^* - (\tilde{B}B^*)_{b:}|_1 < 2$ by definition of T_k and T_j . Hence,

$$|(\tilde{B}B^*)_{a:} - (\tilde{B}B^*)_{b:}|_1 \geq |B_{a:}^* - B_{b:}^*|_1 - |B_{a:}^* - (\tilde{B}B^*)_{a:}|_1 - |B_{b:}^* - (\tilde{B}B^*)_{b:}|_1 > 0,$$

and so $\tilde{B}_{a:} \neq \tilde{B}_{b:}$ from which follows that $\hat{A}_{a:} \neq \hat{A}_{b:}$.

Hence, for $j, k \in R_2$, $a \in T_k$ and $b \in T_j$, we have

$$\hat{A}_{a:} \neq \hat{A}_{b:} \quad \text{if and only if } j \neq k.$$

So all points in $\cup_{k \in R_2} T_k$ are well classified. Hence, there exists a permutation π such that

$$|\{a : \hat{A}_{\pi(a):} \neq A_{a:}^*\}| \leq n - \sum_{k \in R_2} |T_k| = S + \sum_{k \in R_3} |T_k| \leq S + |R_3|n,$$

where $S = \sum_{k=1}^K |S_k|$.

Step 2: The same arguments as in Claim 2 in [20] ensures that $|R_3| \leq |R_1|$.

Step 3: We prove now the inequalities $m|R_1| \leq S \leq |B^* - \tilde{B}B^*|_1$. Actually, we have

$$S \geq \sum_{k \in R_1} |S_k| = \sum_{k \in R_1} |G_k| \geq m|R_1|$$

by definition of R_1 and $m = \min_k |G_k|$. In addition, since $|B_{a:}^* - (\tilde{B}B^*)_{a:}|_1 \geq 1$ for $a \in \cup_k S_k$,

$$S \leq \sum_k \sum_{a \in S_k} |B_{a:}^* - (\tilde{B}B^*)_{a:}|_1 \leq |B^* - \tilde{B}B^*|_1.$$

Step 4: The same arguments as in Claim 4 in [20] ensure that

$$|B^* - \tilde{B}|_1 \leq 15|B^* - \hat{B}|_1.$$

We can now conclude. Combining the 3 first steps, we obtain that

$$|\{a : \hat{A}_{\pi(a)} \neq A_{a:}^*\}| \leq S(1 + n/m) \leq (1 + n/m)|B^* - \tilde{B}B^*|_1.$$

In addition, Step 4 and Lemma 1 ensure that

$$|B^* - \tilde{B}B^*|_1 \leq |B^* - \tilde{B}|_1 \leq 15|B^* - \hat{B}|_1 \leq 30\frac{n}{m}|B^* - B^*\hat{B}|_1.$$

The claim of Lemma 11 then holds by combining the last two displays.

7. Proof of Theorem 2

We provide in this section a full proof of Theorem 2. The lines are very closed to those of the proof of Theorem 1. In particular, all we need is to prove in (37) in our setting. As explained in Section 5.2, we have

$$\langle \mathbf{X}\mathbf{X}^T, B^* - \hat{B} \rangle - \langle \mathbf{X}'(\mathbf{X}')^T, B^* - \hat{B} \rangle \leq 0.05\Delta^2|B^* - B^*\hat{B}|_1,$$

so

$$\begin{aligned} \langle S, B^* - \hat{B} \rangle &= 0.5 \sum_{j \neq k} \Delta_{jk}^2 b_{jk} \\ &\leq \langle EE^T, \hat{B} - B^* \rangle + \langle W', \hat{B} - B^* \rangle + 0.05\Delta^2|B^* - B^*\hat{B}|_1, \end{aligned}$$

with $W'_{ab} = \langle E_a - E_b, \mu_k - \mu_j \rangle$ for $a \in G_k$ and $b \in G_j$.

The cross-product $\langle W', \hat{B} - B^* \rangle$ can be easily bounded with a variation of Lemma 8.

Lemma 12. *With probability at least $1 - c/n^2$, we have for all $B \in \mathcal{C}$*

$$\langle W', B - B^* \rangle \lesssim \sqrt{\langle S, B^* - B \rangle} \sqrt{\delta L \log(nK^3/\delta)} + L\delta \log\left(\frac{nK^3}{\delta}\right),$$

where $\delta = |B^*(B^* - B)|_1$.

The term $\langle EE^T, \hat{B} - B^* \rangle$ must be handled with more care. We first focus on

$$\langle (I - B^*)EE^T(I - B^*), \hat{B} - B^* \rangle.$$

Lemma 13. *When $\Delta^2 \gtrsim Ln/m$, with probability at least $1 - c/n^2$, we have*

$$\langle (I - B^*)EE^T(I - B^*), \hat{B} - B^* \rangle \leq 0.05\Delta^2|B^* - B^*\hat{B}|_1,$$

when $L \geq \log(n)/n$, and

$$\langle (I - B^*)EE^T(I - B^*), \hat{B} - B^* \rangle \leq 0.05\Delta^2 |B^* - B^* \hat{B}|_1 + 0.04n(nL)^2 e^{-cnL},$$

when $L \leq \log(n)/n$ and $\alpha(L) \leq K^3 n^{-1} e^{cnL}$. In particular, either

$$|B^* - B^* \hat{B}|_1 \leq nK^3 e^{-c\Delta^2/4L}$$

or

$$\langle (I - B^*)EE^T(I - B^*), \hat{B} - B^* \rangle \leq 0.06\Delta^2 |B^* - B^* \hat{B}|_1.$$

It remains to control the average terms of the form $\langle B^*EE^T, \hat{B} - B^* \rangle$.

Lemma 14. When $\Delta^2 \gtrsim Ln/m$, with probability at least $1 - c/n^2$, we have

$$\langle B^*EE^T, \hat{B} - B^* \rangle \leq 0.05\Delta^2\delta + c' L\delta \log\left(\frac{nK^3}{\delta}\right).$$

Putting the last three lemmas together, we conclude that, either $\delta = |B^* - B^* \hat{B}|_1$ fulfills $\delta \leq nK^3 e^{-c\Delta^2/4L}$, or with probability larger than $1 - c''/n^2$, we have

$$\delta\Delta^2 \lesssim \delta L \log(nK^3/\delta).$$

In any case, it follows that $\delta \leq nK^3 e^{-c'\Delta^2/L}$ with probability at least $1 - c/n^2$, which gives (37). We conclude the proof of Theorem 2 by following the same lines as for Theorem 1.

7.1. Proof of Lemma 12. As in the proof of Lemma 8, we denote by $z_{(1)}^{(j,k)} \geq z_{(2)}^{(j,k)} \geq \dots$ the values $(z_a^{(j,k)} = \langle E_a, \mu_k - \mu_j \rangle : a \in G_k)$ ranked in decreasing order and we have

$$\langle W', B - B^* \rangle \leq \sum_{j \neq k} \sum_{a=1}^{|B_{G_j G_k}|_1} z_{(a)}^{(j,k)},$$

where we use the same convention as in (40) when $|B_{G_j G_k}|_1$ is not an integer. For any $Q \subset G_k$ with cardinality q we have

$$\sum_{a \in Q} z_a^{(j,k)} = \sum_{a \in Q} \sum_{i=1}^n [(E_{ai} \mathbf{1}_{i < a} + E_{ia} \mathbf{1}_{i > a})(\mu_k - \mu_j)_i].$$

Since each variable E_{ab} for $a < b$ appears at most twice,

$$\sum_i \text{var}(E_{ai}(\mu_k - \mu_j)_i) \leq L\Delta_{jk}^2 \quad \text{and} \quad |E_{ai}(\mu_k - \mu_j)_i| \leq 2L,$$

Bernstein's inequality ensures that

$$\mathbb{P} \left[\sum_{a \in \mathcal{Q}} z_a^{(j,k)} \gtrsim \sqrt{qL\Delta_{jk}^2} + Lt \right] \leq e^{-t}.$$

The conclusion of Lemma 12 is then derived by following the same lines as in the proof of Lemma 8.

7.2. Proof of Lemma 13. We first consider the case $L \geq \log(n)/n$ and then turn to the sparse case $L \leq \log(n)/n$. When $L \geq \log(n)/n$, according e.g. to Theorem 5.2 in [32] we have $|EE^T|_{\text{op}} \lesssim nL$, so (32) ensures that when $\Delta^2 \gtrsim Ln/m$

$$\langle (I - B^*)EE^T(I - B^*), \hat{B} - B^* \rangle \lesssim \frac{nL}{m} |B^* - B^*\hat{B}|_1 \leq 0.05\Delta^2 |B^* - B^*\hat{B}|_1.$$

Let us now consider the case $L \leq \log(n)/n$. As mentioned in Section 5.2, we use a trimming argument. Let \mathbf{X}^{tr} be the matrix \mathbf{X} where we have removed the nodes with degrees larger than γnL , with $\gamma = 2^8 + 2$, and set $E^{\text{tr}} = \mathbf{X}^{\text{tr}} - \mathbb{E}[\mathbf{X}]$. It is known that removing the high-degree nodes drastically reduces the operator norm of the adjacency matrix. For instance, Lemma 5 in [20] ensures that $|E^{\text{tr}}(E^{\text{tr}})^T|_{\text{op}} \lesssim nL$ so when $\Delta^2 \gtrsim Ln/m$, we have

$$\langle (I - B^*)E^{\text{tr}}(E^{\text{tr}})^T(I - B^*), \hat{B} - B^* \rangle \leq 0.05\Delta^2 |B^* - B^*\hat{B}|_1. \quad (41)$$

Similarly as in [20], we bound now the residual terms with a ℓ^1/ℓ^∞ bound. Compared to [20], the main additional difficulty comes from the quadratic residuals, whereas the SDP in [20] was only linear in \mathbf{X} . The first step is to bound the ℓ^1 norm of the residual terms by the sum of the square degrees of the trimmed nodes. We start from

$$\begin{aligned} \langle EE^T - E^{\text{tr}}(E^{\text{tr}})^T, (I - B^*)\hat{B}(I - B^*) \rangle \\ \leq |(I - B^*)\hat{B}(I - B^*)|_\infty |EE^T - E^{\text{tr}}(E^{\text{tr}})^T|_1 \\ \leq 2\alpha(L) |EE^T - E^{\text{tr}}(E^{\text{tr}})^T|_1 \\ \leq 2\alpha(L) (|(E - E^{\text{tr}})^2|_1 + 2|(E - E^{\text{tr}})E^{\text{tr}}|_1) \end{aligned} \quad (42)$$

since $|\hat{B}|_\infty \leq \alpha(L)$.

Control of $|(E - E^{\text{tr}})^2|_1$. The matrix $E - E^{\text{tr}} = \mathbf{X} - \mathbf{X}^{\text{tr}}$ is the adjacency matrix of the graph where we have only kept the edges involving at least one node with degree larger than γnL and their neighbors. The ℓ^1 norm $|(E - E^{\text{tr}})^2|_1$ then counts the number of paths of size 2 in this graph. Write \mathcal{T} for the set of nodes with degree larger than γnL . To evaluate the number of paths (i_1, i_2, i_3) of size 2 in this graph, we consider apart the two cases $i_2 \in \mathcal{T}$ and $i_2 \notin \mathcal{T}$.

Consider first the case where the node i_2 belongs to \mathcal{T} . Since both i_1 and i_3 are neighbors of i_2 , we have at most $\sum_{i \in \mathcal{T}} d^2(i)$ such paths, where $d(i)$ is the degree of i . Consider now the case where $i_2 \notin \mathcal{T}$. In this case, i_1 and i_3 belong to \mathcal{T} . Since the degree of i_2 is less than $\gamma n L \leq d(i_1)$, the number of such paths is again smaller than $\sum_{i \in \mathcal{T}} d^2(i)$. So, we have the bound

$$|(E - E^{\text{tr}})^2|_1 \leq 2 \sum_{i \in \mathcal{T}} d^2(i).$$

Control of $|(E - E^{\text{tr}})E^{\text{tr}}|_1$. We have

$$|(E - E^{\text{tr}})E^{\text{tr}}|_1 \leq |(E - E^{\text{tr}})\mathbf{X}^{\text{tr}}|_1 + |(E - E^{\text{tr}})\mathbb{E}[\mathbf{X}]|_1,$$

and we bound separately the two terms in the right hand side of the above inequality. First, we notice that $|(E - E^{\text{tr}})\mathbf{X}^{\text{tr}}|_1$ corresponds to the number of size 2 paths (i_1, i_2, i_3) such that i_1 belongs to \mathcal{T} and both i_2 and i_3 do not belong to \mathcal{T} . Since $d(i_2) < d(i_1)$, we have again

$$|(E - E^{\text{tr}})\mathbf{X}^{\text{tr}}|_1 \leq \sum_{i \in \mathcal{T}} d^2(i).$$

As for $|(E - E^{\text{tr}})\mathbb{E}[\mathbf{X}]|_1$, this corresponds to the sum of the weights $(E - E^{\text{tr}})_{i_1 i_2} \mathbb{E}[\mathbf{X}_{i_2 i_3}]$ associated to paths (i_1, i_2, i_3) . Since the weight $\mathbb{E}[\mathbf{X}_{i_2 i_3}]$ of the edge (i_2, i_3) is less than L , the total weight of paths starting from $i_1 \in \mathcal{T}$ is upper-bounded by $nL \sum_{i \in \mathcal{T}} d(i)$ and the total weight of paths starting from $i_1 \notin \mathcal{T}$ is also upper-bounded by $nL \sum_{i \in \mathcal{T}} d(i)$, so

$$|(E - E^{\text{tr}})\mathbb{E}[\mathbf{X}]|_1 \leq 2nL \sum_{i \in \mathcal{T}} d(i) \leq \sum_{i \in \mathcal{T}} d^2(i)$$

since \mathcal{T} is made of high-degree nodes.

Coming back to (42), we conclude that

$$\langle EE^T - E^{\text{tr}}(E^{\text{tr}})^T, (I - B^*)\hat{B}(I - B^*) \rangle \leq 8\alpha(L) \sum_{i \in \mathcal{T}} d^2(i). \quad (43)$$

It remains to bound the sum of the squared highest degrees.

Control of $\sum_{i \in \mathcal{T}} d^2(i)$. We control the sum with a stratification argument. First, we shall get rid of the dependencies in \mathbf{X} that are due to the symmetry.

$$d^2(i) = \left(\sum_{j:j>i} \mathbf{X}_{ij} + \sum_{j:j<i} \mathbf{X}_{ij} \right)^2 \leq 2 \left(\sum_{j:j>i} \mathbf{X}_{ij} \right)^2 + 2 \left(\sum_{j:j<i} \mathbf{X}_{ij} \right)^2.$$

For a node i , we write $d_1(i) = \sum_{j:j>i} \mathbf{X}_{ij}$ and $d_2(i) = \sum_{j:j<i} \mathbf{X}_{ij}$. As a consequence,

$$\sum_{i \in \mathcal{I}} d^2(i) \leq 4 \sum_{i=1}^n d_1^2(i) \mathbf{1}_{d_1(i) \geq \gamma n L / 2} + 4 \sum_{i=1}^n d_2^2(i) \mathbf{1}_{d_2(i) \geq \gamma n L / 2}.$$

We focus on the first term, the second term can be bounded in the same way by symmetry. The following technical Lemma is stated in general form as it will be applied several times in the manuscript. Henceforth, \log_2 refers to the binary logarithm.

Lemma 15. *Consider any $\ell > 0$ such that $\ell L \geq 1$. Let $I \subset \{1, \dots, n\}$, $J_i \subset \{i + 1, \dots, n\}$, with $|J_i| \leq \ell$ and $S_i = \sum_{j \in J_i} E_{ij}$ for $i = 1, \dots, n$. For $r_0 \geq 2$ and any integer $r \geq 1$ we set $y_r = 2^{r_0+r} \ell L$ and $I_r = \{i \in I : y_r < S_i \leq y_{r+1}\}$. Then, for $1/14 \leq \tau \leq 1/2$, we have*

$$\mathbb{P} \left[\bigcap_{r \geq 1} \{|I_r| \leq 2n 2^{-\tau(r+r_0-2)y_r}\} \right] \geq 1 - \frac{1 + \log_2(\tau^{-1} \log(2n))}{(2n)^{(1-\tau)/4\tau}},$$

for $n \geq 2$. In addition, we always have $S_i \geq -\ell L$ and for $I_0 = \{i \in I : S_i \leq y_1\}$

$$\mathbb{P} \left[\sum_{i \in I_0} S_i^2 \lesssim n \ell L \right] \geq 1 - 1/n^3.$$

Since $\mathbf{X}_{ij} \in [0, 1]$, $d_1(i) \leq nL + \sum_{j \geq i} E_{ij}$. Take $r_0 = 6$, $I = \{1, \dots, n-1\}$ and $J_i = \{i + 1, \dots, n\}$ for $i \in I$. Since we restrict ourselves to indices i such that $d_1(i) \geq \gamma n L / 2$, our choice of γ implies that $S_i > y_1$. Taking $\ell = n$ and $\tau = 1/10$ in Lemma 15, we obtain with probability at least $1 - c/n^2$ that

$$\begin{aligned} \sum_{i=1}^n d_1^2(i) \mathbf{1}_{d_1(i) \geq \gamma n L / 2} &\leq \sum_{r=1}^{\infty} |I_r| (nL + 2^{r_0+r+1} nL)^2 \\ &\leq 2^8 n(nL)^2 \sum_{r=1}^{\infty} 2^{-(r_0+r-2)(nL 2^{r_0+r} / 10 - 2)} \\ &\leq 2^8 n(nL)^2 \sum_{r=1}^{\infty} 2^{-11.8(r_0+r-2)nL} \\ &\leq 2^9 n(nL)^2 e^{-40nL}. \end{aligned}$$

So putting the pieces together with (42), we obtain

$$\begin{aligned} \langle EE^T - E^{\text{tr}}(E^{\text{tr}})^T, (I - B^*)\widehat{B}(I - B^*) \rangle &\leq 2^{12}\alpha(L)n(Ln)^2e^{-40nL} \\ &\leq 0.04(Ln)^2K^3e^{-4nL}, \end{aligned}$$

where the last inequality holds when $L \geq 1/m \geq 2/n$ and $\alpha(L) \leq \frac{K^3}{n}e^{4nL}$. Then, with (41), we conclude that

$$\langle (I - B^*)EE^T(I - B^*), \widehat{B} - B^* \rangle \leq 0.05\Delta^2|B^* - B^*\widehat{B}|_1 + 0.04(Ln)^2K^3e^{-4nL}. \quad (44)$$

Let us prove the last statement of the lemma. Assume that $\delta = |B^* - B^*\widehat{B}|_1 \geq nK^3e^{-\Delta^2/L}$ and hence

$$\delta\Delta^2 \geq nLK^3(\Delta^2/L)e^{-\Delta^2/L}.$$

Since $n/m \lesssim \Delta^2/L \leq 4nL$ and xe^{-x} is decreasing for $x > 1$, then $(\Delta^2/L)e^{-\Delta^2/L} \geq 4nLe^{-4nL}$ and $\delta\Delta^2 \geq 4K^3(nL)^2e^{-4nL}$. Coming back to (44) concludes the proof.

7.3. Proof of Lemma 14. In order to properly handle the dependences between the symmetric entries of E we split E into two parts $E = U + U^T$ where the upper triangular matrix U is such that $U_{ab} = E_{ab}$ for $a < b$ and $U_{ab} = 0$ else. We have $E^2 = U^2 + (U^T)^2 + UU^T + U^TU$ and by symmetry we only need to control $\langle B^*U^TU, \widehat{B} - B^* \rangle$ and $\langle B^*(U^T)^2, \widehat{B} - B^* \rangle$.

7.3.1. Case U^TU . As in the proof of Lemma 7, all we need is to prove the following bound. For $1 \leq j, k \leq K$ and $b \in G_j$, define

$$z_b^{(j,k)} = \left\{ \frac{(-1)^{j=k}}{m_k} \sum_{i=1}^n \sum_{a \in G_k} U_{ia}U_{ib} : b \in G_j \right\}.$$

Let $z_{(1)}^{(j,k)} \geq z_{(2)}^{(j,k)} \geq \dots$ be the random variables $z_b^{(j,k)}$ ranked in decreasing order.

Lemma 16. *There exists an event Ω of probability at least $1 - c/n^2$, such that for all $1 \leq j, k \leq K$ and for any integer $q \in [1, m_j]$ and $t \geq 0$, we have*

$$\mathbb{P} \left[\Omega \cap \left\{ \sum_{u=1}^q z_{(u)}^{(j,k)} \gtrsim L \left(\frac{qn}{m_k} + t \right) \right\} \right] \leq 3C_{m_j}^q e^{-c't}. \quad (45)$$

To conclude from Lemma 16, we simply apply as in Lemma 7 a union bound

$$\begin{aligned}
 \mathbb{P} \left[\exists q : \sum_{u=1}^q z_{(u)}^{(j,k)} \gtrsim \frac{nLq}{m_k} + Lq \log \left(\frac{nK}{q} \right) \right] \\
 \leq \mathbb{P}[\Omega^c] + \mathbb{P} \left[\Omega \cap \left\{ \exists q : \sum_{u=1}^q z_{(u)}^{(j,k)} \gtrsim \frac{nLq}{m_k} + Lq \log \left(\frac{nK}{q} \right) \right\} \right] \\
 \leq \frac{c}{n^2} + 3 \sum_{j,k} \sum_{q_{j,k}=1}^n C_{m_k}^{q_{j,k}} e^{-cq_{j,k} \log(nK/q_{j,k})} \\
 \lesssim \frac{1}{n^2} + \frac{K^2}{(nK)^c} \lesssim \frac{1}{n^2},
 \end{aligned}$$

for $c \geq 2$. We denote $b_{jk} = |B_{G_k G_j}|_1$ so that $\delta = \sum_{j,k} b_{jk}$. Then, arguing as in Lemma 7, we use order variables with the convention (40) together with Jensen's inequality to conclude

$$\begin{aligned}
 \langle B^* U U^T, \hat{B} - B^* \rangle &\leq \sum_{k,j=1}^K \sum_{u=1}^{b_{jk}} z_{(u)}^{(j,k)} \\
 &\lesssim \sum_{k,j=1}^K b_{jk} L \left(\frac{n}{m_k} + \log \left(\frac{nK}{b_{jk}} \right) \right) \\
 &\lesssim \delta L \left(\frac{n}{m} + L \log \left(\frac{nK^3}{\delta} \right) \right).
 \end{aligned}$$

Since we assume that $\Delta \gtrsim Ln/m$, we have proved the desired bound.

Proof of Lemma 16. With the notation of Lemma 15, for $1 \leq k \leq K$ let us take $J_i^{(k)} = G_k \cap \{i+1, \dots, n\}$, $\ell^{(k)} = m_k$, $\tau = 1/14$, $r_0 = 2$, and $S_i^{(k)} = \sum_{b \in J^{(k)}} U_{ib} = \sum_{b \in J^{(k)}} E_{ib}$. Define accordingly, the sets $I_r^{(k)}$ and

$$\Omega_k = \left\{ \sum_{i \in I_0^{(k)}} (S_i^{(k)})^2 \lesssim nm_k L \right\} \cap \bigcap_{r \geq 1} \left\{ |I_r^{(k)}| \leq 2n2^{-\tau m_k L r 2^{r_0+r}} \right\}, \quad (46)$$

and $y_r^{(k)} = 2^{r_0+r} \ell^{(k)} L$. Then, according to Lemma 15, the event $\Omega = \bigcap_{k=1}^K \Omega_k$ holds with probability at least $1 - c/n^2$. Let us now prove (45). We consider apart the case $j \neq k$ and $j = k$. In the remainder of the proof, k is fixed and to alleviate the notation, we simply write I_r and S_i for $I_r^{(k)}$ and $S_i^{(k)}$.

Case $j \neq k$. Let Q be a subset of G_j with cardinality q , and set $T_{i,Q} = \sum_{b \in Q} U_{ib}$ for $i \in \{1, \dots, n\}$. Then, we have

$$\sum_{b \in Q} z_b^{(j,k)} = \frac{1}{m_k} \sum_{i=1}^n S_i T_{i,Q}.$$

Since all the entries of U are independent and $Q \cap J_i = \emptyset$, the $T_{i,Q}$'s are independent from the S_i 's. Let us first upper-bound the sum

$$\sum_{i \in I_0} S_i T_{i,Q} = \sum_{i=1}^n (S_i \mathbf{1}_{S_i^2 \leq y_1^2}) T_{i,Q}$$

on Ω . Working conditionally on the $(S_i : i \in I_0)$, we have with Bernstein's inequality

$$\mathbb{P} \left[\sum_{i \in I_0} S_i T_{i,Q} \gtrsim \sqrt{qL \sum_{i \in I_0} S_i^2 t} + m_k L t \mid S_i : i \in I_0 \right] \leq e^{-t},$$

Hence, since $\sum_{i \in I_0} S_i^2 \lesssim n m_k L$ on Ω , we have

$$\mathbb{P} \left[\Omega \cap \sum_{i \in I_0} S_i T_{i,Q} \gtrsim L \sqrt{q n m_k t} + m_k L t \right] \leq e^{-t}, \quad (47)$$

for any $t > 0$. Let us now upper-bound

$$\sum_{r \geq 1} \sum_{i \in I_r} S_i T_{i,Q}$$

on Ω . Since such S_i are positive, we have for $\lambda > 0$

$$\begin{aligned} & \mathbb{E} \left[\exp \left(\lambda \sum_{r \geq 1} \sum_{i \in I_r} S_i T_{i,Q} \right) \mathbf{1}_\Omega \right] \\ & \leq \mathbb{E} \left[\exp \left(Lq \sum_{r \geq 1} \sum_{i \in I_r} (e^{\lambda S_i} - 1 - \lambda S_i) \right) \mathbf{1}_\Omega \right] \\ & \leq \mathbb{E} \left[\exp \left(Lq \sum_{r \geq 1} |I_r| (e^{2\lambda m_k L 2^{r_0+r}} - 1 - 2\lambda m_k L 2^{r_0+r}) \right) \mathbf{1}_\Omega \right] \\ & \leq \exp \left(2nLq \sum_{r \geq 1} 2^{-\tau m_k L r 2^{r_0+r}} e^{2\lambda m_k L 2^{r_0+r}} \right). \end{aligned}$$

For $\lambda = \tau \log(2)/4$, we have

$$\begin{aligned} \mathbb{E} \left[\exp \left(\frac{\tau \log(2)}{4} \sum_{r \geq 1} \sum_{i \in I_r} S_i T_{i,Q} \right) \mathbf{1}_\Omega \right] & \leq \exp \left(2nLq \sum_{r \geq 1} 2^{-\tau m_k L r 2^{r_0+r-1}} \right) \\ & \leq e^{2nLq e^{-c m_k L}} \leq e^{nLq}, \end{aligned}$$

since $L \gtrsim m_k$. This gives

$$\mathbb{P} \left[\Omega \cap \sum_{r \geq 1} \sum_{i \in I_r} S_i T_{i,Q} > t \right] \leq e^{-\frac{\tau \log(2)t}{4} + nLq}.$$

Together with (47), we obtain that

$$\mathbb{P} \left[\Omega \cap \sum_{i=1}^n S_i T_{i,Q} \gtrsim nLq + m_k Lt \right] \leq 2e^{-t}$$

and so

$$\mathbb{P} \left[\Omega \cap \max_{Q \subset G_j, |Q|=q} \sum_{b \in Q} z_b^{(j,k)} \gtrsim \frac{nLq}{m_k} + Lt \right] \leq 2C_{m_j}^q e^{-t}.$$

Case $j = k$. We start from

$$\sum_{b \in Q} z_b^{(j,k)} = -m_k^{-1} \sum_{i=1}^n S_i T_{i,Q}.$$

Unfortunately, the sums S_i and $T_{i,Q}$ are no longer independent as $Q \subset G_k$ and $T_{i,Q}$ is therefore a subsum of S_i . We define $k(i)$ as the index in $\{1, \dots, K\}$ such that $i \in G_{k(i)}$ and we set

$$J_i = G_k \cap \{i + 1, \dots, n\}, \quad L_i = P_{k(i)k}, \quad \text{and} \quad N_i = \sum_{a \in J_i} X_{ia} = S_i + |J_i|L_i.$$

We observe that conditionally on N_i , the sum $H_i = \sum_{b \in Q \cap \{i+1, \dots, n\}} X_{ib}$ follows a hypergeometric distribution with parameter $(q_i, N_i, |J_i|)$, where

$$q_i = |Q \cap \{i + 1, \dots, n\}| \leq q.$$

Let H'_i be a random variable with binomial $(q_i, N_i/|J_i|)$ distribution conditionally on N_i . Since $x \rightarrow e^{-\lambda S_i x}$ is continuous and convex, according to Theorem 4 of [26], we have conditionally on S_i

$$\mathbb{E}[e^{-\lambda S_i H_i} | S_i] \leq \mathbb{E}[e^{-\lambda S_i H'_i} | S_i].$$

Hence, conditionally on S_i , we apply Chernoff bound to $\sum_{i \in I_0} S_i H_i$ together with the above control of the Laplace transform. This allows us to get a Bernstein-like inequality. Hence, with probability higher than $1 - e^{-t}$ we have

$$-\sum_{i \in I_0} S_i T_{i,Q} = -\sum_{i \in I_0} S_i (H_i - q_i L_i)$$

$$\begin{aligned}
&\leq -\sum_{i \in I_0} q_i S_i \left(\frac{N_i}{|J_i|} - L_i \right) + c \sqrt{\sum_{i \in I_0} q_i S_i^2 \frac{N_i}{|J_i|} t} + cm_k L t \\
&\leq -\sum_{i \in I_0} \frac{q_i S_i^2}{|J_i|} + c \sqrt{\max_{i \in I_0} N_i \sum_{i \in I_0} \frac{q_i S_i^2}{|J_i|} t} + cm_k L t \\
&\lesssim m_k L t,
\end{aligned}$$

where we used in the last line that $N_i = S_i + |J_i|L_i \leq (1 + 2^{r_0})m_k L$ for $i \in I_0$.

When $S_i > 0$, we have

$$-S_i T_{i,Q} \leq qL S_i,$$

since each entry E_{ab} is larger or equal to $-L$. For $i \in I_r$ with $r \geq 1$, we have $y_r < S_i \leq 2y_r$ and hence, on Ω

$$\begin{aligned}
-\sum_i \sum_{r \geq 1} S_i T_{i,Q} \mathbf{1}_{i \in I_r} &\leq 2qL \sum_{r \geq 1} |I_r| y_r \\
&\leq 4nqL \sum_{r \geq 1} 2^{r_0+r} m_k L 2^{-\tau m_k L r 2^{r_0+r}} \\
&\leq 4nqL 2^{-cm_k L},
\end{aligned}$$

since $m_k L \gtrsim 1$. Hence arguing as before we obtain that for $t \geq 0$

$$\mathbb{P} \left[\Omega \cap \max_{Q \subset G_k : |Q|=q} \sum_{b \in Q} z_b^{(j,k)} \gtrsim \frac{nLq}{m_k} + Lt \right] \leq C_m^q e^{-t}.$$

The proof of Lemma 16 is complete. \square

7.3.2. Case $(U^T)^2$. The case $(U^T)^2$ is somewhat more messy, due to interlaced rows/columns dependences. Recall, that for two indices j and k , $(-1)^{j=k}$ equals -1 when $i = k$ and 1 when $i \neq k$. For $1 \leq j \leq k$ and $b \in G_j$, define

$$z_b^{(j,k)} = \frac{(-1)^{j=k}}{m_k} \sum_{i=1}^n \sum_{a \in G_k} U_{ia} U_{bi},$$

and let $z_{(1)}^{(j,k)} \geq z_{(2)}^{(j,k)} \geq \dots$ be the random variables $z_b^{(j,k)}$ ranked in decreasing order. Compared to $U^T U$, the difficulty is that the same random variables U_{ia} can occur several times in the definition of $z_b^{(j,k)}$.

Lemma 17. *There exists an event $\underline{\Omega}$ of probability at least $1 - c/n^2$, such that for all $1 \leq j, k \leq K$ and for any integer $q \in [1, m_k]$ and $t \geq 0$, we have*

$$\mathbb{P} \left[\underline{\Omega} \cap \left\{ \sum_{u=1}^q z_{(u)}^{(j,k)} \gtrsim L \left(\frac{qn}{m_k} + t \right) \right\} \right] \leq c'' C_{m_k}^q e^{-c't}. \quad (48)$$

As for the previous case, we easily conclude from Lemma 17 by applying a union bound together with Jensen's inequality.

Proof of Lemma 17. The event $\underline{\Omega}$ is defined as the intersection

$$\underline{\Omega} = \Omega \cap \bigcap_{u=1}^7 \bigcap_{k=1}^K \underline{\Omega}_k^{(u)},$$

where the event $\Omega = \bigcap_k \Omega_k$ is introduced in (46) the previous proof and the events $\underline{\Omega}_k^{(u)}$ are defined along the proof. Let us split the sum

$$(-1)^{j=k} \sum_{i=1}^n \sum_{a \in G_k, b \in Q} U_{ia} U_{bi}$$

into two parts depending whether $i \in G_k$ or not.

Case $i \notin G_k$. For $Q \subset G_j$, consider the sum

$$(-1)^{j=k} \sum_{i \notin G_k} \sum_{a \in G_k, b \in Q} U_{ia} U_{bi} = (-1)^{j=k} \sum_{i \notin G_k} S_i T'_{i,Q}, \quad (49)$$

with $S_i = \sum_{a \in G_k} U_{ia}$ as in the proof of Lemma 16 (we dropped the exponent (k) to alleviate the notation) and $T'_{i,Q} = \sum_{b \in Q} U_{bi}$. In the collection of (S_i)'s and ($T'_{i,Q}$)'s, all the random variables are independent since the sums respectively run on the sets $G_k^c \times G_k$ and $Q \times G_k^c$ that do not intersect. As a consequence, (49) is handled exactly as the case $j \neq k$ in Lemma 16. We conclude that for all $1 \leq j, k \leq K$ and all $Q \subset G_j$ of size q , we have

$$\mathbb{P} \left[\Omega \cap \left\{ (-1)^{j=k} \sum_{i \notin G_k} \sum_{a \in G_k, b \in Q} U_{ia} U_{bi} \gtrsim \frac{nLq}{m_k} + Lt \right\} \right] \lesssim e^{-t}. \quad (50)$$

Case $i \in G_k$ and $j \neq k$. For $Q \subset G_j$, we consider the sum

$$(-1)^{j=k} \sum_{i \in G_k} \sum_{a \in G_k, b \in Q} U_{ia} U_{bi} = (-1)^{j=k} \sum_{i \in G_k} S_i T'_{i,Q}, \quad (51)$$

with $S_i = \sum_{a \in G_k} U_{ia}$ and $T'_{i,Q} = \sum_{b \in Q} U_{bi}$. As above, the indices run in $G_k \times G_k$ and $Q \times G_k$ which do not intersect. Again, (51) is handled exactly as the case $j \neq k$ in Lemma 16. We conclude that for all $1 \leq j, k \leq K$ with $j \neq k$ and all $Q \subset G_j$ of size q , we have

$$\mathbb{P} \left[\Omega \cap \left\{ \sum_{i \in G_k} \sum_{a \in G_k, b \in Q} U_{ia} U_{bi} \gtrsim \frac{nLq}{m_k} + Lt \right\} \right] \lesssim e^{-t}. \quad (52)$$

Case $i \in G_k$ and $j = k$. It remains to upper-bound the sum

$$- \sum_{i \in G_k} \sum_{a \in G_k} U_{ia} \sum_{b \in Q} U_{bi}, \quad (53)$$

for $Q \subset G_k$. It is the main hurdle of the proof. Indeed, we multiply row sums $\sum_{a \in G_k} U_{ia}$ of the matrix U (restricted to $G_k \times G_k$) to columns sums of the matrix U . As in the previous proof, we consider separately small and large row sum of U . Define the set of indices corresponding to small S_i 's

$$\underline{I}_0 = \{i \in G_k : S_i \leq 5m_k L\}, \quad (54)$$

where we recall that $S_i = \sum_{a \in G_k} U_{ia}$.

Define $\ell_b = |G_k \cap \{b+1, \dots, n\}| \leq m_k$ and $N_b = S_b + \ell_b P_{kk}$. Given the collection (S_i) , $i \in G_k$, the binary random variable $U_{bi} + P_{kk}$ is distributed as a sampling of size 1 in an urn of size ℓ_b containing N_b ones. Hence, we split the sum (53) into three pieces to center the random variables U_{bi} .

$$\begin{aligned} & - \sum_{i \in G_k} \sum_{a \in G_k, a > i} \sum_{b \in Q, b < i} U_{ia} U_{bi} \\ &= - \sum_{b \in Q} \sum_{i \in G_k, i > b} S_i U_{bi} \\ &= - \sum_{i \in G_k} S_i \sum_{b \in Q, b < i} \left(\frac{N_b}{\ell_b} - P_{kk} \right) - \sum_{i \in G_k} S_i \sum_{b \in Q, b < i} \left(U_{bi} + P_{kk} - \frac{N_b}{\ell_b} \right) \\ &\leq - \sum_{b \in Q} \left(\frac{N_b}{\ell_b} - P_{kk} \right) \sum_{i \in G_k, i > b} S_i - \sum_{i \in \underline{I}_0} S_i \sum_{b \in Q, b < i} \left(U_{bi} + P_{kk} - \frac{N_b}{\ell_b} \right) \\ &\quad + \sum_{i \in \underline{I}_0^c} S_i \sum_{b \in Q, b < i} \frac{N_b}{\ell_b} \\ &= E_1 + E_2 + E_3, \end{aligned}$$

where we used that $S_i > 0$ for $i \in \underline{I}_0^c$.

We shall prove that the three following bounds hold, for any $t \geq 1$,

$$\mathbf{1}_{\underline{\Omega}} |E_1| \lesssim qnL, \quad (55)$$

$$\mathbb{P}[\underline{\Omega} \cap \{E_2 \gtrsim qnL + m_k Lt\}] \leq 2e^{-t}, \quad (56)$$

$$\mathbb{P}[\underline{\Omega} \cap \{E_3 \gtrsim qnL + t\}] \leq e^{-t}, \quad (57)$$

and that $\mathbb{P}[\underline{\Omega}] \geq 1 - c/n^2$. Gathering these three bounds, we obtain for all $1 \leq k \leq K$ and all $Q \subset G_k$, we have

$$\mathbb{P} \left[-\mathbf{1}_{\underline{\Omega}} \sum_{i \notin G_k} \sum_{a \in G_k, b \in Q} U_{ia} U_{bi} \gtrsim \frac{nLq}{m_k} + Lt \right] \lesssim e^{-t}. \quad (58)$$

Together with (50) and (52), this concludes the proof. It remains to show (55)–(57).

Control of E_1 . The random variable $\sum_{i>b, i \in G_k} S_i$ is distributed as a Binomial random variable with parameters P_{kk} and $n_b \leq \ell_b^2/2$. By Bernstein's inequality together with a union bound, we derive that, on an event $\underline{\Omega}_k^{(1)}$ with probability higher than $1 - 1/n^3$, we have

$$\left| \sum_{i>b} S_i \right| \lesssim \ell_b \sqrt{L \log(n)} + \log(n) \tag{59}$$

$$|N_b| \lesssim \ell_b L + \sqrt{\ell_b L \log(n)} + \log(n) \lesssim \ell_b L + \log(n), \tag{60}$$

uniformly on $b \in G_k$. Consider any b such that $\ell_b^2 L \geq \log(n)$. Then,

$$\left| \left(\frac{N_b}{\ell_b} - P_{kk} \right) \sum_{i \in G_k, i>b} S_i \right| \lesssim L^{3/2} \ell_b \sqrt{\log(n)} + \log^{3/2}(n) \sqrt{L} \lesssim nL, \tag{61}$$

since $m_k L \geq 1$ and $L \leq 1/\log(n)$.

Now assume that $\ell_b^2 L \leq \log(n)$ and $L \geq \log(n)/n$. By definition, we have $N_b \leq \ell_b$ almost surely. Together with (59), this leads us to

$$\left| \left(\frac{N_b}{\ell_b} - P_{kk} \right) \sum_{i \in G_k, i>b} S_i \right| \lesssim \log(n) \leq nL. \tag{62}$$

Next, we consider the case where $\ell_b^2 L \leq \log(n)$ and $L \leq \log(n)/n$. If $\ell_b > n^{1/4}$, then (59) leads us to

$$\left| \left(\frac{N_b}{\ell_b} - P_{kk} \right) \sum_{i \in G_k, i>b} S_i \right| \lesssim L \log(n) + \frac{\log^2(n)}{n^{1/4}} \lesssim 1 \lesssim nL. \tag{63}$$

Let b_0 such that $\ell_{b_0} = \lfloor n^{1/4} \rfloor$ (if it exists). Then, $\sum_{i \in G_k, i>b_0} N_i$ follows a binomial distribution with parameters $P_{kk} \leq \log(n)/n$ and $r \leq \sqrt{n}$. On an event $\underline{\Omega}_k^{(2)}$ with probability higher than $1 - 1/n^3$, it is no higher than 5. Under this event, we have for any $b \geq b_0$,

$$\left| \left(\frac{N_b}{\ell_b} - P_{kk} \right) \sum_{i \in G_k, i>b} S_i \right| \lesssim 1 \leq nL. \tag{64}$$

Gathering (61)–(64) and summing over all $b \in Q$, we have proved (55).

Control of E_2 . We work conditionally to S_i . In such a case, for a fixed b , the random variables $((U_{bi} + P_{kk})_{i>b})$ are distributed as a sampling without replacement of N_b ones in an urn of size ℓ_b . Then, according to Theorem 4 of [26], the Laplace transform of $-\sum_{b \in Q} \sum_{i>b} S_i (U_{bi} + P_{kk} - N_b/\ell_b)$ conditional to the $(S_i)_i$ is upper bounded by that of $-\sum_{b \in Q} \sum_{i>b} S_i (\tilde{E}_{ib} - N_b/\ell_b)$ where the \tilde{E}_{ib} are independent and follow a Bernoulli distribution with parameters N_b/ℓ_b .

Hence, we can apply Bernstein's inequality conditionally to S_i to obtain that with probability at least $1 - e^{-t}$

$$E_2 \lesssim \sqrt{\sum_{i \in \underline{I}_0} S_i^2 \left(\sum_{b \in Q, b < i} \frac{N_b}{\ell_b} \right) t} + m_k L t, \quad (65)$$

since $\sup_{i \in \underline{I}_0} |S_i| \leq 5m_k L$. We define $s^* = \log^2(n)$, $Q_- = \{b \in Q : \ell_b \leq s^*\}$ and $Q_+ = \{b \in Q : \ell_b > s^*\}$. Then, we split the sum into two parts

$$\sum_{i \in \underline{I}_0} S_i^2 \sum_{b \in Q, b < i} \frac{N_b}{\ell_b} = \sum_{i \in \underline{I}_0} S_i^2 \sum_{b \in Q_+, b < i} \frac{N_b}{\ell_b} + \sum_{i \in \underline{I}_0} S_i^2 \sum_{b \in Q_-, b < i} \frac{N_b}{\ell_b}.$$

Sum over Q_+ . By Bernstein's inequality, we have, simultaneously for all $i \in G_k$,

$$|S_i| \mathbf{1}_{i \in \underline{I}_0} \lesssim \sqrt{m_k L \log(n)} + (\log(n) \wedge (m_k L)) \lesssim \sqrt{m_k L \log(n)}$$

on an event $\underline{\Omega}_k^{(3)}$ with probability higher than $1 - 1/n^3$. Since the random variables $S_i^2 \mathbf{1}_{i \in \underline{I}_0}$ are independent and their variance is less than $6\ell_b^2 L^2 + \ell_b L$, we derive from Bernstein's inequality that, on an event $\underline{\Omega}_k^{(3)} \cap \underline{\Omega}_k^{(4)}$ with probability higher than $1 - 2/n^3$, for all $b \in Q_+$, we have

$$\begin{aligned} \sum_{i > b} S_i^2 \mathbf{1}_{i \in \underline{I}_0} &\lesssim \ell_b^2 L + \sqrt{[\ell_b^3 L^2 + \ell_b^2 L] \log(n)} + m_k L \log^2(n) \\ &\lesssim \ell_b m_k L + m_k L \log^2(n) \lesssim \ell_b m_k L, \end{aligned} \quad (66)$$

since $\log^2(n) \leq \ell_b \leq m_k$ and $m_k L \geq 1$. Under this event, we obtain

$$\sum_{i \in \underline{I}_0} S_i^2 \sum_{b \in Q_+, b < i} \frac{N_b}{\ell_b} \lesssim m_k L \sum_{b \in Q_+} N_b.$$

Then, $\sum_{b \in Q_+} N_b$ is stochastically dominated by a Binomial distribution with parameters L and qm_k , hence we derive from Bernstein's inequality that

$$\sum_{b \in Q_+} N_b \lesssim qm_k L + t,$$

with probability higher than $1 - e^{-t}$. Gathering (65), (66), and the last bound, we obtain

$$\mathbb{P}(\underline{\Omega} \cap \{E_2^+ \gtrsim m_k L(\sqrt{qt} + t)\}) \leq 2e^{-t}, \quad (67)$$

where E_2^+ corresponds to the sum E_2 restricted to the indices $b \in Q_+$.

Sum over Q_- . We consider two subcases. First assume that $L \geq n^{-1/4}$. Since $\sum_{i:\ell_i \leq s^*} S_i^2 \leq \sum_{i=1}^{\log^2(n)-1} i^2 \leq \log^6(n)$ and, since $N_b \leq \ell_b$, we derive that

$$\sum_{i \in I_0} S_i^2 \sum_{b \in Q_-, b < i} \frac{N_b}{\ell_b} \leq q \log^6(n) \lesssim qL^2 nm_k. \tag{68}$$

Next, we assume that $L \leq n^{-1/4}$. Let $b_* = \min\{b \in G_k : \ell_b \leq s^*\}$. Since, by definition of s^* and b_* , we have

$$|\{(i, a) : i, a \in G_k, i \geq b_* + 1, a \geq i + 1\}| \leq (s^*)^2/2 = \log^4(n)/2,$$

the sums $(S_i)_{i > b_*}$ involves less than $\log^4(n)/2$ independent Bernoulli random variables with parameters less than L . Hence, on an event $\underline{\Omega}_k^{(5)}$ with probability larger than $1 - 1/n^3$, at most 10 of them are equal to one and

$$\sum_{i \in I_0} S_i^2 \sum_{b \in Q_-, b < i} \frac{N_b}{\ell_b} \lesssim q(L^2 \log^6(n) + 1) \lesssim qL^2 nm_k, \tag{69}$$

since $m_k L \geq 1$. Gathering (65)–(69), we have proved (56).

Control of E_3 . If $L \geq \log(n)/m_k$, then Bernstein’s inequality enforces that $I_0^c = \emptyset$ and therefore $E_3 = 0$ with probability higher than $1 - 1/n^3$. Let us call $\underline{\Omega}_k^{(6)}$ the corresponding event. Hence, we assume henceforth that $L \leq \log(n)/m_k$. We claim that, on an event $\underline{\Omega}_k^{(7)}$ with probability larger than $1 - 1/n^3$, we have

$$\sum_{i > b} S_i \mathbf{1}_{i \in I_0^c} \lesssim \ell_b + \log^3(n), \tag{70}$$

uniformly over all $b \in G_k$. The proof of this claim is a slight variation on the proof of Lemma 15. We provide it here for the sake of completeness. With probability higher than $1 - 1/n^3$, we have $\max_i S_i \lesssim \log(n)$. Write \log_2 for the binary logarithm. Fix any $b \in G_k$ and decompose

$$\sum_{i > b} S_i \mathbf{1}_{i \in I_0^c} \leq \sum_{r=1}^{\lfloor \log_2(c'' \log(n)/m_k L) \rfloor} \sum_{i > b} 52^r L m_k \mathbf{1}_{S_i \in [5 \cdot 2^{r-1} L m_k, 5 \cdot 2^r L m_k]}.$$

The random variables $\sum_{i > b} \mathbf{1}_{S_i \in [5 \cdot 2^{r-1} L m_k, 5 \cdot 2^r L m_k]}$ are stochastically dominated by binomial distributions with parameters ℓ_b and p_r , where $p_r \leq e^{-c' 2^r L m_k}$ is the probability that a Binomial distribution with parameters (m_k, L) is larger than $5 \cdot 2^{r-1} L m_k$. Applying Bernstein’s inequality together with a union bound, we conclude that, simultaneously for all b and all r ,

$$\sum_{i > b} \mathbf{1}_{S_i \in [5 \cdot 2^{r-1} L m_k, 5 \cdot 2^r L m_k]} \lesssim \ell_b e^{-c' 2^{r-1} L m_k} + \log(n),$$

with probability higher than $1 - 1/n^3$. This leads us to

$$\sum_{i>b} S_i \mathbf{1}_{i \in I_0^c} \leq \sum_{r=1}^{\lfloor \log_2(c'' \log(n)/m_k L) \rfloor} c 2^r L m_k [\ell_b e^{-c' 2^r L m_k} + \log(n)] \lesssim \ell_b + \log^3(n),$$

since $L m_k \geq 1$. We have proved the claim (70).

As for E_2 , we decompose $Q = Q_+ \cup Q_-$, with s^* now set to $s^* = \log^3(n)$.

Sum over Q_+ . Let us work on the event $\underline{\Omega}_k^{(7)}$ so that (70) holds. Hence, as for E_2 , according to Bernstein's inequality, we have with probability larger than $1 - e^{-t}$

$$\sum_{b \in Q_+} \frac{N_b}{\ell_b} \sum_{i \in I_0^c, i>b} S_i \lesssim \sum_{b \in Q_+} N_b \lesssim q m_k L + t. \quad (71)$$

Sum over Q_- . As for E_2 , we consider two subcases depending whether $L \leq n^{-1/4}$ or $L > n^{-1/4}$. If $L \leq n^{-1/4}$, we have, as argued in E_2 , that $\sum_{i>b_*} N_i$ is less than 10 on an event $\underline{\Omega}_k^{(5)}$ with probability larger than $1 - 1/n^2$. Under this event, we have

$$\sum_{b \in Q_-} \frac{N_b}{\ell_b} \sum_{i \in I_0^c, i>b} S_i \lesssim q \lesssim q m_k L. \quad (72)$$

If $L > n^{-1/4}$, we use the crude bound

$$\sum_{b \in Q_-} \frac{N_b}{\ell_b} \sum_{i \in I_0^c, i>b} S_i \leq q (s^*)^2 \lesssim q L n. \quad (73)$$

Putting (71)–(73) together, we conclude that on the event $\underline{\Omega}$ which has a probability larger than $1 - 7/n^2$ we have proved (57).

7.4. Proof of Lemma 15. With no loss of generality we only consider the case $I = \{1, \dots, n\}$. Let us first handle the upper-bound on the sum $\sum_{i=1}^n S_i^2 \mathbf{1}_{S_i^2 \leq y_1^2}$, with $y_1 = 2^{r_0+1} \ell L$. We observe that the variance of S_i^2 is upper bounded by $3\ell L$, hence, according to Bernstein's inequality,

$$\sum_{i=1}^n S_i^2 \mathbf{1}_{S_i^2 \leq y_1^2} \lesssim \ell L (n + \sqrt{\log(n)}) + y_1^2 \log(n),$$

with probability at least $1 - 1/n^3$. We observe that

$$y_1^2 \log(n) \lesssim (\ell L)^2 \log(n) \lesssim n \ell L$$

according to the assumptions on L , so with probability at least $1 - 1/n^3$

$$\sum_{i=1}^n S_i^2 \mathbf{1}_{S_i^2 \leq y_1^2} \lesssim n\ell L.$$

We now turn to the first part of the lemma. We shall first work around a bound of $\mathbb{P}(y_r < S_i \leq y_{r+1})$. We set $h(x) = (1 + x) \log(1 + x) - x \geq x \log(x/4)$ for $x \geq 0$ and $y_r = 2^{r_0+r} \ell L$. Denoting σ_i^2 the variance of S_i , we deduce from Bennett's inequality that, for $r \geq 1$,

$$\begin{aligned} \mathbb{P}(y_r < S_i \leq y_{r+1}) &\leq \mathbb{P}(S_i > y_r) \leq e^{-\sigma_i^2 h(y_r/\sigma_i^2)} \leq \exp\left[-y_r \log\left(\frac{y_r}{4\sigma_i^2}\right)\right] \\ &\leq 2^{-\ell L(r+r_0-2)2^{r_0+r}} =: p_r, \end{aligned}$$

since $\sigma_i^2 \leq \ell L$.

Next, we use again Bennett's inequality to ensure that $|I_r|$ is small. For $1/14 \leq \tau \leq 1/2$, and $r \geq 1$, Bennett's inequality again ensures that

$$\mathbb{P}[|I_r| \geq 2np_r^\tau] \leq \exp(-np_r h(p_r^{-(1-\tau)})) \leq \exp(-0.5np_r^\tau \log(p_r^{-(1-\tau)})),$$

since $h(x) \geq x \log(x)/2$ for $x \geq e^2$ and since $\tau \leq 1/2$ and $p_r \leq 2^{-16}$ is small enough. Since $r_0 \geq 2$, we have $2np_r^\tau < 1$ for $r \geq r^* = \log_2(\tau^{-1} \log(2n))$. Let Ω' be the event

$$\Omega' = \bigcap_{r \geq 1} \{|I_r| \leq 2np_r^\tau\} = \bigcap_{r=1}^{\lfloor r^* \rfloor} \{|I_r| \leq 2np_r^\tau\} \cap \{|\{i : S_i > y_{r^*}\}| = 0\}.$$

So we have

$$\mathbb{P}(\Omega'^c) \leq \sum_{r=1}^{\lfloor r^* \rfloor} \exp(-0.5np_r^\tau \log(p_r^{-(1-\tau)})) + \exp(-0.5np_{r^*}^\tau \log(p_{r^*}^{-(1-\tau)})).$$

For $x \in (0, 1)$, the function $\phi_\tau: x \mapsto x^\tau \log(x)$ is decreasing when $\tau \log(x) < -1$ and then increasing. Since $(2n)^{-1/\tau} \leq p_r \leq 2^{-16}$, $\phi_\tau(p_r)$ achieves its maximum at $p_r = (2n)^{-1/\tau}$ for $\tau \in [1/10, 1/2]$. For $\tau \in [1/14, 1/10)$, its maximum is achieved either at $p_r = 2^{-16}$ or $p_r = (2n)^{-1/\tau}$ and it can be checked to be reached at $p_r = (2n)^{-1/\tau}$ as long as $(2n)^{-1/\tau} \leq 2^{-16}$. Note that for smaller nT , the event Ω' reduces to $\{|\{i : S_i > y_{r^*}\}| = 0\}$. For $n \geq 2$, we conclude that

$$\mathbb{P}(\Omega'^c) \leq (\lfloor r^* \rfloor + 1) \exp\left(-\frac{1-\tau}{4\tau} \log(2n)\right) \leq \frac{\log_2(\tau^{-1} \log(2n)) + 1}{(2n)^{(1-\tau)/4\tau}}.$$

Acknowledgements. We thank the editors and anonymous reviewers for their helpful suggestions. Christophe Giraud is partially supported by the LabEx LMH, ANR-11-LABX-0056-LMH and the CNRS PICS grant HighClust.

A. About CH-divergence

Lemma 18. *The CH-divergence defined by (29) fulfills*

$$D_+(q||p) \leq \frac{1}{4\rho} \sum_x \frac{(p_x - q_x)^2}{p_x}, \quad \text{when } \min_x \frac{q_x}{p_x} \geq \rho > 0.$$

Proof. The function $f_t(y) = 1 - t + ty - y^t$ is convex and fulfills $f_t(1) = 0 = f_t'(1)$ and

$$f_t''(1 + u) \leq \frac{u^2}{4\rho}, \quad \text{for all } 1 + u \geq \rho \text{ and } t \in [0, 1].$$

Setting $u_x = (q_x - p_x)/p_x$, we get the claimed result. \square

B. Lower bound on misclassification probability in supervised Gaussian classification with unknown means

In this section, we derive a lower-bound on the misclassification probability of the Bayes classifier, in the Gaussian supervised classification problem with two balanced classes, with identical spherical covariances $\Sigma_k = \sigma I_p$ and opposite means $\mu_{-1} = -\mu_1$ uniformly distributed on the Euclidean sphere $\partial B(0, \Delta/2)$ in \mathbb{R}^p .

We denote by $\mathcal{L} = (X_a, Z_a)_{a=1, \dots, n}$ the learning sample distributed as follows. The labels Z_1, \dots, Z_n are i.i.d. with uniform distribution on $\{-1, 1\}$, a random vector $\mu \in \mathbb{R}^p$ is sampled uniformly over the sphere $\partial B(0, \Delta/2)$ independently of Z_1, \dots, Z_n , and, conditionally on Z_1, \dots, Z_n, μ , the X_a are independent Gaussian random variables with mean $Z_a \mu$ and covariance $\sigma^2 I_p$.

The classifier minimizing the misclassification probability $\mathbb{P}[Z_{\text{new}} \neq \hat{h}(X_{\text{new}})]$ over all the $\sigma(\mathcal{L})$ -measurable classifiers \hat{h} is the Bayes classifier given by

$$\hat{h}(x) = \text{sign}(\mathbb{P}[Z = 1|X = x, \mathcal{L}] - \mathbb{P}[Z = -1|X = x, \mathcal{L}]).$$

Let us compute the Bayes classifier in our setting. Indeed, the classification problem is scaling-invariant.

For $\delta \in \{-1, 1\}$ and $x \in \mathbb{R}^p$, we have

$$\mathbb{P}[Z = \delta|X = x, \mathcal{L}] = \int_{\partial B(0, \Delta/2)} \mathbb{P}[Z = \delta|X = x, \mathcal{L}, \mu] d \mathbb{P}[\mu|X = x, \mathcal{L}].$$

Direct computations give

$$\begin{aligned} \mathbb{P}[Z = \delta|X = x, \mathcal{L}, \mu] &= \mathbb{P}[Z = \delta|X = x, \mu] \\ &= \frac{e^{-0.5\|\delta x - \mu\|^2/\sigma^2}}{e^{-0.5\|x + \mu\|^2/\sigma^2} + e^{-0.5\|x - \mu\|^2/\sigma^2}} \end{aligned}$$

and

$$d \mathbb{P} [\mu | X = x, \mathcal{L}] \propto (e^{-0.5\|x+\mu\|^2/\sigma^2} + e^{-0.5\|x-\mu\|^2/\sigma^2}) e^{-0.5 \sum_a \|Z_a X_a - \mu\|^2/\sigma^2}.$$

Hence, by using that $\|\mu\| = \Delta/2$ on $\partial B(0, \Delta/2)$, and by denoting by γ the uniform distribution on $\partial B(0, \Delta/2)$, we obtain

$$\begin{aligned} & \mathbb{P}[Z = \delta | X = x, \mathcal{L}] \\ &= \frac{\int_{\partial B(0, \Delta/2)} e^{-0.5\|\delta x - \mu\|^2/\sigma^2} e^{-0.5 \sum_a \|Z_a X_a - \mu\|^2/\sigma^2} d\gamma(\mu)}{\int_{\partial B(0, \Delta/2)} (e^{-0.5\|x+\mu'\|^2/\sigma^2} + e^{-0.5\|x-\mu'\|^2/\sigma^2}) e^{-0.5 \sum_a \|Z_a X_a - \mu'\|^2/\sigma^2} d\gamma(\mu')} \\ &= \frac{\int_{\partial B(0, \Delta/2)} e^{-(\delta x + \sum_a Z_a X_a, \mu)/\sigma^2} d\gamma(\mu)}{\int_{\partial B(0, \Delta/2)} e^{-(x + \sum_a Z_a X_a, \mu')/\sigma^2} d\gamma(\mu') + \int_{\partial B(0, \Delta/2)} e^{-(-x + \sum_a Z_a X_a, \mu')/\sigma^2} d\gamma(\mu')}. \end{aligned}$$

Since $F(v) = \int_{\partial B(0, \Delta/2)} e^{(v, \mu)} d\gamma(\mu)$ depends only on $\|v\|$ and is monotone increasing with $\|v\|$, we obtain that

$$\begin{aligned} \mathbb{P}[Z = 1 | X = x, \mathcal{L}] &> \mathbb{P}[Z = -1 | X = x, \mathcal{L}] \\ &\iff \|x + \sum_a Z_a X_a\|^2 > \|-x + \sum_a Z_a X_a\|^2 \\ &\iff \langle x, \sum_a Z_a X_a \rangle > 0, \end{aligned}$$

and finally

$$\hat{h}(x) = \text{sign}\left(\left\langle \frac{1}{n} \sum_{a=1}^n Z_a X_a, x \right\rangle\right).$$

For any $\sigma > 0$, the probability of misclassification of the Bayes classifier is given by

$$\begin{aligned} \mathbb{P} [Z_{\text{new}} \neq \hat{h}(X_{\text{new}})] &= \int_{\partial B(0, \Delta/2)} \mathbb{P} [Z \hat{h}(X) < 0 | \mu] d\gamma(\mu) \\ &= \int_{\partial B(0, \Delta/2)} \mathbb{P} \left[\left\langle \mu + \frac{\sigma}{\sqrt{n}} \epsilon, \mu + \sigma \epsilon' \right\rangle < 0 | \mu \right] d\gamma(\mu), \end{aligned}$$

where ϵ and ϵ' are two independent standard Gaussian random variables in \mathbb{R}^p . The above conditional probability is invariant over $\partial B(0, \Delta/2)$ hence we only need to evaluate it for a fixed $\mu \in \partial B(0, \Delta/2)$, say $\mu_\Delta = [\Delta/2, 0, \dots, 0]$. Let us set

$$W = -2(\Delta \sqrt{1 + 1/n})^{-1} \left\langle \mu_\Delta, \frac{1}{\sqrt{n}} \epsilon + \epsilon' \right\rangle,$$

which follows a standard Gaussian distribution in \mathbb{R} and $S = -\langle \epsilon, \epsilon' \rangle$. Then, we have

$$\begin{aligned} \mathbb{P} [Z_{\text{new}} \neq \hat{h}(X_{\text{new}})] &= \mathbb{P} \left[\left\langle \mu_{\Delta} + \frac{\sigma}{\sqrt{n}}\epsilon, \mu_{\Delta} + \sigma\epsilon' \right\rangle < 0 \right] \\ &= \mathbb{P} \left[\frac{\Delta^2}{4\sigma^2} < \frac{\Delta}{2\sigma} \sqrt{1 + \frac{1}{n}} W + \frac{1}{\sqrt{n}} S \right]. \end{aligned}$$

We observe that $(-W, S)$ has the same distribution as (W, S) , hence by a union bound

$$\mathbb{P} \left[\frac{\Delta^2}{4\sigma^2} < \frac{\Delta}{2\sigma} \sqrt{1 + \frac{1}{n}} |W| + \frac{1}{\sqrt{n}} S \right] \leq 2 \mathbb{P} \left[\frac{\Delta^2}{4\sigma^2} < \frac{\Delta}{2\sigma} \sqrt{1 + \frac{1}{n}} W + \frac{1}{\sqrt{n}} S \right].$$

By using that the distributions of S and W are symmetric and that $a \vee (2b - 2a) \geq a \vee (b/2)$, we get

$$\begin{aligned} &\mathbb{P} \left[\frac{\Delta^2}{4\sigma^2} < \frac{\Delta}{2\sigma} \sqrt{1 + \frac{1}{n}} |W| + \frac{1}{\sqrt{n}} S \right] \\ &\geq \mathbb{P} \left[\frac{\Delta^2}{4\sigma^2} < \frac{1}{\sqrt{n}} S \right] + \mathbb{P} \left[\frac{\Delta^2}{2\sigma^2} < \frac{\Delta}{2\sigma} \sqrt{1 + \frac{1}{n}} |W|; \frac{1}{\sqrt{n}} |S| \leq \frac{\Delta^2}{4\sigma^2} \right] \\ &\geq \mathbb{P} \left[\frac{\Delta^2}{4\sigma^2} < \frac{1}{\sqrt{n}} S \right] \vee \left(\mathbb{P} \left[\frac{\Delta}{\sigma} < |W| \right] - \mathbb{P} \left[\frac{1}{\sqrt{n}} |S| > \frac{\Delta^2}{4\sigma^2} \right] \right) \\ &\geq \mathbb{P} \left[\frac{\Delta^2}{4\sigma^2} < \frac{1}{\sqrt{n}} S \right] \vee \left(2 \mathbb{P} \left[\frac{\Delta}{\sigma} < W \right] - 2 \mathbb{P} \left[\frac{1}{\sqrt{n}} S > \frac{\Delta^2}{4\sigma^2} \right] \right) \\ &\geq \mathbb{P} \left[\frac{\Delta^2}{4\sigma^2} < \frac{1}{\sqrt{n}} S \right] \vee \frac{1}{2} \mathbb{P} \left[\frac{\Delta}{\sigma} < W \right]. \end{aligned}$$

Putting the pieces together, we get

$$4 \mathbb{P} [Z_{\text{new}} \neq \hat{h}(X_{\text{new}})] \geq \mathbb{P} \left[\frac{\Delta}{\sigma} < W \right] \vee \mathbb{P} \left[\frac{\Delta^2}{4\sigma^2} < \frac{1}{\sqrt{n}} S \right].$$

We observe that S is distributed as the product of a standard Gaussian real random variable with the square-root of an independent χ^2 random variable with p degrees of freedom. Since a χ^2 random variable with p degrees of freedom is larger than $p/2$ with probability larger than $1/2$, we then have

$$\mathbb{P} \left[S > \frac{\sqrt{n} \Delta^2}{4\sigma^2} \right] \geq \frac{1}{2} \mathbb{P} \left[\sqrt{\frac{p}{2}} W > \frac{\sqrt{n} \Delta^2}{4\sigma^2} \right] \geq \frac{1}{2} \mathbb{P} \left[W > \frac{\sqrt{n} \Delta^2}{\sqrt{8p} \sigma^2} \right].$$

We then obtain the lower bound on the Bayes probability of misclassification

$$\mathbb{P} [Z_{\text{new}} \neq \hat{h}(X_{\text{new}})] \geq c \exp \left(-c' \left(\frac{\Delta^2}{\sigma^2} \wedge \frac{n\Delta^4}{p\sigma^4} \right) \right),$$

for some numerical constants $c, c' > 0$.

References

- [1] E. Abbe, Community detection and stochastic block models: recent developments, *J. Mach. Learn. Res.*, **18** (2017), Paper No. 177, 86pp (2018). [Zbl 1403.62110](#) [MR 3827065](#)
- [2] E. Abbe and C. Sandon, Community detection in general stochastic block models: fundamental limits and efficient algorithms for recovery, in *2015 IEEE 56th Annual Symposium on Foundations of Computer Science—FOCS 2015*, 670–688, IEEE Computer Soc., Los Alamitos, CA, 2015. [MR 3473334](#)
- [3] D. Achlioptas and F. McSherry, On spectral learning of mixtures of distributions, in *Learning theory*, 458–469, Lecture Notes in Comput. Sci., 3559, Lecture Notes in Artificial Intelligence, Springer, Berlin, 2005. [Zbl 1137.68512](#) [MR 2203280](#)
- [4] A. A. Amini and E. Levina, On semidefinite relaxations for the block model, 2014. [arXiv:1406.5647](#)
- [5] D. Arthur and S. Vassilvitskii, k -means++: the advantages of careful seeding, in *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1027–1035, ACM, New York, 2007. [Zbl 1302.68273](#) [MR 2485254](#)
- [6] P. Awasthi, A. S. Bandeira, M. Charikar, R. Krishnaswamy, S. Villar, and R. Ward, Relax, no need to round: integrality of clustering formulations, in *ITCS'15—Proceedings of the 6th Innovations in Theoretical Computer Science*, 191–200, ACM, New York, 2015.. [Zbl 1364.62144](#) [MR 3419009](#)
- [7] P. Awasthi, M. Charikar, R. Krishnaswamy, and A. K. Sinop, The hardness of approximations of Euclidean k -means, in *31st International Symposium on Computational Geometry*, 754–767, LIPIcs. Leibniz Int. Proc. Inform., 34, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2015. [Zbl 1378.68048](#) [MR 3392820](#)
- [8] M. Azizyan, A. Singh, and L. Wasserman, Minimax theory for high-dimensional gaussian mixtures with sparse mean separation, in *Advances in Neural Information Processing Systems 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (eds.), 2139–2147, Curran Associates, Inc., 2013.
- [9] Q. Berthet, P. Rigollet, and P. Srivastava, Exact recovery in the Ising blockmodel, to appear in *Ann. Statist.* [arXiv:1612.03880](#)
- [10] F. Bunea, C. Giraud, X. Luo, M. Royer, and N. Verzelen, Model assisted variable clustering: Minimax-optimal recovery and algorithms, 2015. [arXiv:1508.01939](#)
- [11] F. Bunea, C. Giraud, M. Royer, and N. Verzelen, PECOK: a convex optimization approach to variable clustering, 2016. [arXiv:1606.05100](#)
- [12] M. Charikar, S. Guha, É. Tardos, and D. B. Shmoys, A constant-factor approximation algorithm for the k -median problem, *J. Comput. System Sci.*, **65** (2002), no. 1, Special issue on STOC (Atlanta, GA, 1999), 129–149. [Zbl 1023.90037](#) [MR 1946291](#)
- [13] X. Chen and Y. Yang, Hanson–Wright inequality in Hilbert spaces with application to K -means clustering for non-Euclidean data, 2018. [arXiv:1810.11180](#)
- [14] Y. Chen, X. Li, and J. Xu, Convexified modularity maximization for degree-corrected stochastic block models, 2015. [arXiv:1512.08425](#)

- [15] Y. Chen and J. Xu, Statistical-computational tradeoffs in planted problems and submatrix localization with a growing number of clusters and submatrices, *J. Mach. Learn. Res.*, **17** (2016), Paper No. 27, 57pp. [Zbl 1360.62320](#) [MR 3491121](#)
- [16] P. Chin, A. Rao, and V. Vu, Stochastic Block Model and Community Detection in Sparse Graphs: A spectral algorithm with optimal rate of recovery, in *Proceedings of The 28th Conference on Learning Theory (Paris, France, 03–06 July 2015)*, P. Grünwald, E. Hazan, and S. Kale (eds.), 391–423, Proceedings of Machine Learning Research, 40, PMLR, 2015.
- [17] S. Chrétien, C. Dombry, and A. Faivre, A semi-definite programming approach to low dimensional embedding for unsupervised clustering, 2016. [arXiv:1606.09190](#)
- [18] Y. Deshpande, E. Abbe, and A. Montanari, Asymptotic mutual information for the binary stochastic block model, in *2016 IEEE International Symposium on Information Theory (ISIT, Barcelona, Spain, July 2016)*, 185–189, IEEE, 2016.
- [19] I. Diakonikolas, D. M. Kane, and A. Stewart, List-decodable robust mean estimation and learning mixtures of spherical Gaussians, CoRR, 2017. [arXiv:1711.07211](#)
- [20] Y. Fei and Y. Chen, Exponential error rates of SDP for block models: beyond Grothendieck’s inequality, *IEEE Trans. Inform. Theory*, **65** (2019), no. 1, 551–571. [Zbl 07003336](#) [MR 3901009](#)
- [21] Y. Fei and Y. Chen, Hidden integrality of SDP relaxation for sub-Gaussian mixture models, 2018. [arXiv:1803.06510](#)
- [22] C. Gao, Z. Ma, A. Y. Zhang, and H. H. Zhou, Achieving optimal misclassification proportion in stochastic block models, *J. Mach. Learn. Res.*, **18** (2017), Paper No. 60, 45pp. [Zbl 06781381](#) [MR 3687603](#)
- [23] C. Giraud, *Introduction to high-dimensional statistics*, Monographs on Statistics and Applied Probability, 139, CRC Press, Boca Raton, FL, 2015. [Zbl 1341.62011](#) [MR 3307991](#)
- [24] O. Guédon and R. Vershynin, Community detection in sparse networks via Grothendieck’s inequality, *Probab. Theory Related Fields*, **165** (2016), no. 3-4, 1025–1049. [Zbl 1357.90111](#) [MR 3520025](#)
- [25] B. Hajek, Y. Wu, and J. Xu, Semidefinite programs for exact recovery of a hidden community, 2016. [arXiv:1602.06410](#)
- [26] W. Hoeffding, Probability inequalities for sums of bounded random variables, *J. Amer. Statist. Assoc.*, **58** (1963), no. 301, 13–30. [Zbl 0127.10602](#) [MR 144363](#)
- [27] P. W. Holland, K. Blackmond Laskey, and S. Leinhardt, Stochastic blockmodels: first steps, *Social Networks*, **5** (1983), no. 2, 109–137. [MR 718088](#)
- [28] S. B. Hopkins and J. Li, Mixture models, robustness, and sum of squares proofs, in *STOC’18—Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, 1021–1034, ACM, New York, 2018. [MR 3826314](#)
- [29] T. Iguchi, D. G. Mixon, J. Peterson, and S. Villar, On the tightness of an SDP relaxation of k -means, 2015. [arXiv:1505.04778](#)
- [30] A. Javanmard, A. Montanari, and F. Ricci-Tersenghi, Phase transitions in semidefinite relaxations, *Proc. Natl. Acad. Sci. USA*, **113** (2016), no. 16, E2218–E2223. [Zbl 1359.62188](#) [MR 3494080](#)

- [31] P. K. Kothari and J. Steinhardt, Better agnostic clustering via relaxed tensor norms, 2017. [arXiv:1711.07465](#)
- [32] J. Lei and A. Rinaldo, Consistency of spectral clustering in stochastic block models, *Ann. Statist.*, **43** (2015), no. 1, 215–237. [Zbl 1308.62041](#) [MR 3285605](#)
- [33] X. Li, Y. Chen, and J. Xu, Convex relaxation methods for community detection, 2018. [arXiv:1810.00315](#)
- [34] X. Li, Y. Li, S. Ling, T. Strohmer, and K. Wei, When do birds of a feather flock together? k -means, proximity, and conic programming, *Math. Program.*, published online: 29 September 2018, 47pp. [doi:10.1007/s10107-018-1333-x](#)
- [35] S. Lloyd, Least squares quantization in PCM, *IEEE Trans. Inform. Theory*, **28** (1982), no. 2, 129–137. [Zbl 0504.94015](#) [MR 651807](#)
- [36] Y. Lu and H. H. Zhou, Statistical and computational guarantees of Lloyd’s algorithm and its variants, 2016. [arXiv:1612.02099](#)
- [37] D. G. Mixon, S. Villar, and R. Ward, Clustering sub-Gaussian mixtures by semidefinite programming, *Inf. Inference*, **6** (2017), no. 4, 389–415. [Zbl 1381.62189](#) [MR 3764529](#)
- [38] A. Moitra, W. Perry, and A. S. Wein, How robust are reconstruction thresholds for community detection? in *STOC’16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, 828–841, ACM, New York, 2016. [Zbl 1381.62190](#) [MR 3536617](#)
- [39] C. Moore, The computer science and physics of community detection: landscapes, phase transitions, and hardness, *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS*, (2017), no. 121, 26–61. [MR 3699594](#)
- [40] M. Ndaoud, Sharp optimal recovery in the two component Gaussian mixture model, 2018. [arXiv:1812.08078](#)
- [41] J. Peng and Y. Wei, Approximating k -means-type clustering via semidefinite programming, *SIAM J. Optim.*, **18** (2007), no. 1, 186–205. [Zbl 1146.90046](#) [MR 2299680](#)
- [42] A. Perry and A. S. Wein, A semidefinite program for unbalanced multisection in the stochastic block model, 2015. [arXiv:1507.05605](#)
- [43] O. Regev and A. Vijayaraghavan, On learning mixtures of well-separated Gaussians, in *58th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2017*, 85–96, IEEE Computer Soc., Los Alamitos, CA, 2017. [MR 3734220](#)
- [44] M. Royer, Adaptive clustering through semidefinite programming, *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), 1795–1803, Curran Associates, Inc., 2017.
- [45] M. Rudelson and R. Vershynin, Hanson–Wright inequality and sub-Gaussian concentration, *Electron. Commun. Probab.*, **18** (2013), no. 82, 9pp. [Zbl 1329.60056](#) [MR 3125258](#)
- [46] S. Vempala and G. Wang, A spectral algorithm for learning mixture models, *J. Comput. System Sci.*, **68** (2004), no. 4, 841–860. [Zbl 1074.68028](#) [MR 2059647](#)
- [47] S.-Y. Yun and A. Proutière, Accurate community detection in the stochastic block model via spectral algorithms, 2014. [arXiv:1412.7335](#)

Received 19 July, 2018; revised 04 March, 2019

C. Giraud, Laboratoire de Mathématiques d'Orsay, Université Paris-Sud,
CNRS, Université Paris-Saclay, 91405 Orsay, France

E-mail: christophe.giraud@math.u-psud.fr

N. Verzelen, INRA, Montpellier SupAgro, MISTEA,
Université Montpellier, 34060 Montpellier, France

E-mail: nicolas.verzelen@inra.fr