# Separating populations with wide data: A spectral analysis 

Avrim Blum*<br>Computer Sciene Departement<br>Carnegie Mellon University<br>5000 Forbes Ave<br>Pittsburgh, PA 15213, USA<br>e-mail: avrim@cs.cmu.edu<br>Amin Coja-Oghlan ${ }^{\dagger}$<br>University of Edinburgh<br>Informatics Forum<br>10 Crichton Street<br>Edinburgh EH8 9AB, UK<br>e-mail: acoghlan@inf.ed.ac.uk

Alan Frieze ${ }^{\ddagger}$<br>Departement of Mathmatics<br>Carnegie Mellon University<br>5000 Forbes Ave<br>Pittsburgh, PA 15213, USA<br>e-mail: alan@random.math.cmu.edu

## Shuheng Zhou ${ }^{8}$

Seminar für Statistik
ETH Zentrum, CH-8092 Zürich, Switzerland
e-mail: zhou@stat.math.ethz.ch


#### Abstract

In this paper, we consider the problem of partitioning a small data sample drawn from a mixture of $k$ product distributions. We are interested in the case that individual features are of low average quality $\gamma$, and we want to use as few of them as possible to correctly partition the sample. We analyze a spectral technique that is able to approximately optimize the total data size - the product of number of data points $n$ and the number of features $K$-needed to correctly perform this partitioning as a function of $1 / \gamma$ for $K>n$. Our goal is motivated by an application in clustering individuals according to their population of origin using markers, when the divergence between any two of the populations is small.


[^0]> AMS 2000 subject classifications: Primary $60 \mathrm{~K} 35,60 \mathrm{~K} 35$; secondary 60 K 35.
> Keywords and phrases: mixture of product distributions, clustering, small sample, spectral analysis.
> Received September 2008 .

## Contents

1 Introduction ..... 77
1.1 Related work ..... 80
2 A simple algorithm using singular vectors ..... 82
2.1 Generating a large gap in $s_{1}(\mathcal{X}), s_{2}(\mathcal{X})$ ..... 85
2.2 Detailed analysis for the simple algorithm ..... 86
2.3 Proof of Theorem 2.4 ..... 87
2.4 Correctness of classification for the simple algorithm ..... 88
3 The algorithm Partition ..... 91
3.1 Preliminaries ..... 91
3.2 Description of the algorithm ..... 93
3.3 Proof of Lemma 3.5 ..... 95
3.4 Proof of Lemma 3.6 ..... 98
4 Experiments ..... 99
Acknowledgements ..... 101
A More Proofs for the simple algorithm classify ..... 101
A. 1 Proof of Lemma 2.1 ..... 101
A. 2 Some Propositions regarding the static matrices ..... 104
A. 3 Proofs of Proposition 2.5 and 2.6 ..... 106
A. 4 Proof of Lemma 2.7 ..... 108
B Proof of Lemma 2.16 ..... 109
References ..... 112

## 1. Introduction

We explore a type of classification problem that arises in the context of computational biology. The problem is that we are given a small sample of size $n$, e.g., DNA of $n$ individuals (think of $n$ in the hundreds or thousands), each described by the values of $K$ features or markers, e.g., SNPs (Single Nucleotide Polymorphisms, think of $K$ as an order of magnitude larger than $n$ ). Our goal is to use these features to classify the individuals according to their population of origin. Features have slightly different probabilities depending on which population the individual belongs to, and are assumed to be independent of each other (i.e., our data is a small sample from a mixture of $k$ very similar product distributions). The objective we consider is to minimize the total data size $D=n K$ needed to correctly classify the individuals in the sample as a function of the "average quality" $\gamma$ of the features, under the assumption that $K>n$. Throughout the paper, we use $p_{i}^{j}$ and $\mu_{i}^{j}$ as shorthands for $p_{i}^{(j)}$ and $\mu_{i}^{(j)}$ respectively.

Statistical Model: We have $k$ probability spaces $\Omega_{1}, \ldots, \Omega_{k}$ over the set $\{0,1\}^{K}$. Further, the components (features) of $z \in \Omega_{t}$ are independent and $\operatorname{Pr}_{\Omega_{t}}\left[z_{i}=1\right]=p_{t}^{i}(1 \leq t \leq k, 1 \leq i \leq K)$. Hence, the probability spaces $\Omega_{1}, \ldots, \Omega_{k}$ comprise the distribution of the features for each of the $k$ populations. Moreover, the input of the algorithm consists of a collection (mixture) of $n=\sum_{t=1}^{k} N_{t}$ unlabeled samples, $N_{t}$ points from $\Omega_{t}$, and the algorithm is to determine for each data point from which of $\Omega_{1}, \ldots, \Omega_{k}$ it was chosen. In general we do not assume that $N_{1}, \ldots, N_{t}$ are revealed to the algorithm; but we do require some bounds on their relative sizes. An important parameter of the probability ensemble $\Omega_{1}, \ldots, \Omega_{k}$ is the measure of divergence

$$
\begin{equation*}
\gamma=\min _{1 \leq s<t \leq k} \frac{\sum_{i=1}^{K}\left(p_{s}^{i}-p_{t}^{i}\right)^{2}}{K} \tag{1.1}
\end{equation*}
$$

between any two distributions. Note that $\sqrt{K \gamma}$ measures the Euclidean distance between the means of any two distributions and thus represents their separation. Further, let $N=n / k$ (so if the populations were balanced we would have $N$ of each type) and assume from now on that $k N<K$. Let $D=n K$ denote the size of the data-set. In addition, let $\sigma^{2}=\max _{i, t} p_{t}^{i}\left(1-p_{t}^{i}\right)$ denote the maximum variance of any random bit.

The biological context for this problem is we are given DNA information from $n$ individuals from $k$ populations of origin and we wish to classify each individual into the correct category. DNA contains a series of markers called SNPs, each of which has two variants (alleles). Given the population of origin of an individual, the genotypes can be reasonably assumed to be generated by drawing alleles independently from the appropriate distribution. The following theorem gives a sufficient condition for a balanced ( $N_{1}=N_{2}$ ) input instance when $k=2$.

Theorem 1.1 (Zhou (2006)). Assume $N_{1}=N_{2}=N$. If $K=\Omega\left(\frac{\ln N}{\gamma}\right)$ and $K N=\Omega\left(\frac{\ln N \log \log N}{\gamma^{2}}\right)$ then with probability $1-1 / \operatorname{poly}(N)$, among all balanced cuts in the complete graph formed among $2 N$ sample individuals, the maximum weight cut corresponds to the partition of the $2 N$ individuals according to their population of origin. Here the weight of a cut is the sum of weights across all edges in the cut, and the edge weight equals the Hamming distance between the bit vectors of the two endpoints.

Variants of the above theorem, based on a model that allows two random draws from each SNP for an individual, are given in Chaudhuri et al. (2007); Zhou (2006). In particular, notice that edge weights based on the inner-product of two individuals' bit vectors correspond to the sample covariance, in which case the max-cut corresponds to the correct partition Zhou (2006) with high probability. Finding a max-cut is computationally intractable; hence in the same paper Chaudhuri et al. (2007), a hill-climbing algorithm is given to find the correct partition for balanced input instances but with a stronger requirement on the sizes of both $K$ and $n K$.

A Spectral Approach: In this paper, we construct two simpler algorithms using spectral techniques, attempting to reproduce conditions above. In particular, we study the requirements on the parameters of the model (namely, $\gamma, N$, $k$, and $K$ ) that allow us to classify every individual correctly and efficiently with high probability.

The two algorithms Classify and Partition compare as follows. Both algorithms are based on spectral methods originally developed in graph partitioning. More precisely, Theorem 1.2 is based on computing the singular vectors with the two largest singular values for each of the $n \times K$ input random matrix. The procedure is conceptually simple, easy to implement, and efficient in practice. For simplicity, Procedure Classify assumes the separation parameter $\gamma$ is known to decide which singular vector to examine; in practice, one can just try both singular vectors as we do in the simulations. Proof techniques for Theorem 1.2, however, are difficult to apply to cases of multiple populations, i.e., $k>2$. Procedure Partition is based on computing a rank- $k$ approximation of the input random matrix and can cope with a mixture of a constant number of populations. It is more intricate for both implementation and execution than Classify. It does not require $\gamma$ as an input, while only requires that the constant $k$ is given. We prove the following theorems.
Theorem 1.2. Let $\omega=\frac{\min \left(N_{1}, N_{2}\right)}{n}$ and $\omega_{\min }$ be a lower bound on $\omega$. Let $\gamma$ be given. Assume that $K>2 n \ln n$ and $k=2$. Procedure Classify allows us to separate two populations w.h.p., when $n \geq \Omega\left(\frac{\sigma^{2}}{\gamma \omega_{\min } \omega}\right)$, where $\sigma^{2}$ is the largest variance of any random bit, i.e. $\sigma^{2}=\max _{i, t} p_{t}^{i}\left(1-p_{t}^{i}\right)$. Thus if the populations are roughly balanced, then $n \geq \frac{c}{\gamma}$ suffices for some constant $c$.

This implies that the data required is $D=n K=O\left(\ln n \sigma^{4} / \gamma^{2} \omega^{2} \omega_{\min }^{2}\right)$. Let $P_{s}=\left(p_{s}^{i}\right)_{i=1, \ldots, K}$, we have

$$
\begin{equation*}
\left\|P_{1}-P_{2}\right\|_{2}=\sqrt{K \gamma}=\sqrt{\sum_{i=1}^{K}\left(p_{1}^{i}-p_{2}^{i}\right)^{2}} \geq \frac{\sigma}{\omega_{\min } \omega} \sqrt{\ln n} \tag{1.2}
\end{equation*}
$$

Theorem 1.3. Let $\omega=\frac{\min \left(N_{1}, \ldots, N_{k}\right)}{n}$. There is a polynomial time algorithm Partition that satisfies the following. Suppose that $K>n \log n, \gamma>K^{-2}$, $n>\frac{C_{k} \sigma^{2}}{\gamma \omega}$ for some large enough constant $C_{k}$, and $\omega=\Omega(1)$. Then given the empirical $n \times K$ matrix comprising the $K$ features for each of the $n$ individuals along with the parameter $k$, PARTITION separates the $k$ populations correctly w.h.p.

Summary and Future Direction: Note that unlike Theorem 1.1, both Theorem 1.2 and Theorem 1.3 require a lower bound on $n$, even when $k=2$ and the input instance is balanced. We illustrate through simulations to show that this seems not to be a fundamental constraint of the spectral techniques; our experimental results show that even when $n$ is small, by increasing $K$ so that $n K=\Omega\left(1 / \gamma^{2}\right)$, one can classify a mixture of two populations using ideas in Procedure Classify with success rate reaching an "oracle" curve, which is computed
assuming that distributions are known, where success rate means the ratio between correctly classified individuals and $N$. Exploring the tradeoffs of $n$ and $K$ that are sufficient for classification, when sample size $n$ is small, is both of theoretical interests and practical value.

Outline of the paper: The paper is organized as follows. In Section 1.1 we discuss related work. Then, in Section 2 we describe the algorithm Classify for Theorem 1.2 and outline its analysis. Some (very) technical details of the analysis are deferred to the appendix. Section 3 deals with the algorithm Partition for Theorem 1.3. Finally, in Section 4 we report some experimental results on Classify.

### 1.1. Related work

In their seminal paper Pritchard et al. (2000), Pritchard, Stephens, and Donnelly presented a model-based clustering method to separate populations using genotype data. They assume that observations from each cluster are random from some parametric model. Inference for the parameters corresponding to each population is done jointly with inference for the cluster membership of each individual, and $k$ in the mixture, using Bayesian methods.

The idea of exploiting the eigenvectors with the first two eigenvalues of the adjacency matrix to partition graphs goes back to the work of Fiedler Fiedler (1973), and has been used in the heuristics for various NP-hard graph partitioning problems (e.g., Fjallstrom (1998)). The main difference between graph partitioning problems and the classification problem that we study is that the matrices occurring in graph partitioning are symmetric and hence diagonalizable, while our input matrix is rectangular in general. Thus, the contribution of Theorem 1.2 is to show that a conceptually simple and efficient algorithm based on singular value decompositions performs well in the framework of a fairly general probabilistic model, where probabilities for each of the $K$ features for each of the $k$ populations are allowed to vary. Indeed, the analysis of Classify requires exploring new ideas such as the Separation Lemma and the normalization of the random matrix $X$, for generating a large gap between top two singular values of the expectation matrix $\mathcal{X}$ and for bounding the angle between random singular vectors and their static correspondents, details of which are included in Section 2 with analysis in full version.

Procedure Partition and its analysis build upon the spectral techniques of McSherry (2001) on graph partitioning, and an extension due to Coja-Oghlan (2006). McSherry provides a comprehensive probabilistic model and presents a spectral algorithm for solving the partitioning problem on random graphs, provided that a separation condition similar to (1.2) is satisfied. Indeed, McSherry (2001) encompasses a considerable portion of the prior work on Graph Coloring, Minimum Bisection, and finding Maximum Clique. Moreover, McSherry's approach easily yields an algorithm that solves the classification problem studied in the present paper under similar assumptions as in Theorem 1.3, provided
that the algorithm is given the parameter $\gamma$ as an additional input; this is actually pointed out in the conclusions of McSherry (2001). In the context of graph partitioning, an algorithm that does not need the separation parameter as an input was devised in Coja-Oghlan (2006). The main difference between Partition and the algorithm presented in Coja-Oghlan (2006) is that PartiTION deals with the asymmetric $n \times K$ matrix of individuals/features, whereas Coja-Oghlan (2006) deals with graph partitioning (i.e., a symmetric matrix).

There are two streams of related work in the learning community. The first stream is the recent progress in learning from the point of view of clustering: given samples drawn from a mixture of well-separated Gaussians (component distributions), one aims to classify each sample according to which component distribution it comes from, as studied in Dasgupta (1999), Dasgupta and Schulman (2000), Arora and Kannan (2001); Vempala and Wang (2002); Achlioptas and McSherry (2005); Kannan et al. (2005); Dasgupta et al. (2005). This framework has been extended to more general distributions such as logconcave distributions in Achlioptas and McSherry (2005); Kannan et al. (2005) and heavy-tailed distributions in Dasgupta et al. (2005), as well as to more than two populations. These results focus mainly on reducing the requirement on the separations between any two centers $P_{1}$ and $P_{2}$. In contrast, we focus on the sample size $D$. This is motivated by previous results Chaudhuri et al. (2007); Zhou (2006) stating that by acquiring enough attributes along the same set of dimensions from each component distribution, with high probability, we can correctly classify every individual.

While our aim is different from those results, where $n>K$ is almost universal and we focus on cases $K>n$, we do have one common axis for comparison, the $\ell_{2}$-distance between any two centers of the distributions. In earlier works Dasgupta and Schulman (2000); Arora and Kannan (2001), the separation requirement depended on the number of dimensions of each distribution; this has recently been reduced to be independent of $K$, the dimensionality of the distribution for certain classes of distributions Achlioptas and McSherry (2005); Kannan et al. (2005). This is comparable to our requirement in (1.2) for the discrete distributions. For example, according to Theorem 7 in Achlioptas and McSherry (2005), in order to separate the mixture of two Gaussians,

$$
\begin{equation*}
\left\|P_{1}-P_{2}\right\|_{2}=\Omega\left(\frac{\sigma}{\sqrt{\omega}}+\sigma \sqrt{\log n}\right) \tag{1.3}
\end{equation*}
$$

is required. Besides Gaussian and Logconcave, a general theorem: Theorem 6 in Achlioptas and McSherry (2005) is derived that in principle also applies to mixtures of discrete distributions. The key difficulty of applying their theorem directly to our scenario is that it relies on a concentration property of the distribution (Eq. (10) of Achlioptas and McSherry (2005)) that need not hold in our case. In addition, once the distance between any two centers is fixed (i.e., once $\gamma$ is fixed in the discrete distribution), the sample size $n$ in their algorithms is always larger than $\Omega\left(\frac{K}{\omega} \log ^{5} K\right)$ Achlioptas and McSherry (2005); Kannan et al. (2005) for log-concave distributions (in fact, in Theorem 3 of Kannan et al. (2005), they discard at least this many individuals in order to
correctly classify the rest in the sample), and larger than $\Omega\left(\frac{K}{\omega}\right)$ for Gaussians Achlioptas and McSherry (2005), whereas in our case, $n<K$ always holds. Hence, our analysis allows one to obtain a clean bound on $n$ in the discrete case.

The second stream of work is under the PAC-learning framework, where given a sample generated from some target distribution $Z$, the goal is to output a distribution $Z_{1}$ that is close to $Z$ in Kullback-Leibler divergence: $K L\left(Z \| Z_{1}\right)$, where $Z$ is a mixture of product distributions over discrete domains or Gaussians Kearns et al. (1994); Freund and Mansour (1999); Cryan (1999); Cryan et al. (2002); Mossel and Roch (2005); Feldman et al. (2005, 2006). They do not require a minimal distance between any two distributions, but they do not aim to classify every sample point correctly either, and in general require much more data.

Our work is also related to the use of principal component analysis (" PCA ") in genetics Patterson et al. (2006); Price et al. (2006). The basic approach in these papers is to use the eigenvectors of a covariance matrix between samples to analyze a mixture of populations. While Patterson et al. (2006); Price et al. (2006) study the use of spectral methods empirically, the crucial point of the present work is that we prove rigorously that spectral methods succeed on a certain (simple) probabilistic model. Hence, our work can be seen as a further theoretical justification of the practical use of PCA. A difference between the present paper and Patterson et al. (2006); Price et al. (2006) is that we actually aim to assign each individual to exactly one of the populations. By contrast, Patterson et al. (2006); Price et al. (2006) just assign each individual a real "weight" for each population: essentially the eigenvectors with the dominant eigenvalues corresponding to the populations, and each individual is assigned its projection on these dominant eigenvectors. The algorithm Classify is somewhat similar to PCA, but Partition is conceptually more involved. In addition, our experimental results show a phase transition phenomenon similar to what was observed in Patterson et al. (2006) in detecting population structure using simulated data.

## 2. A simple algorithm using singular vectors

As described in Theorem 1.2, we assume we have a mixture of two product distributions. Let $N_{1}, N_{2}$ be the number of individuals from each population class. Our goal is to correctly classify all individuals according to their distributions. Let $n=2 N=N_{1}+N_{2}$, and refer to the case when $N_{1}=N_{2}$ as the balanced input case. For convenience, let us redefine " $K$ " to assume we have $O(\log n)$ blocks of $K$ features each (so the total number of features is really $O(K \log n))$ and we assume that each set of $K$ features has divergence at least $\gamma$. (If we perform this partitioning of features into blocks randomly, then with high probability this divergence has changed by only a constant factor for most blocks.)

The high-level idea of the algorithm is now to repeat the following procedure for each block of $K$ features: use the $K$ features to create an $n \times K$ matrix $X$, such
that each row $X_{i}, i=1, \ldots, n$, corresponds to a feature vector for one sample point, across its $K$ dimensions. We then compute the top two left singular vectors $u_{1}, u_{2}$ of $X$ and use these to classify each sample. This classification induces some probability of error $f$ for each individual at each round, so we repeat the procedure for each of the $O(\log n)$ blocks and then take majority vote over different runs. Each round we require $K \geq n$ features, so we need $O(n \log n)$ features total in the end.

In more detail, we repeat the following procedure $O(\log n)$ times. Let $T=$ $\frac{15 N}{32} \sqrt{3 \omega_{\min } \gamma}$, where $\omega_{\min }$ is the lower bound on the minimum weight $\min \left\{\frac{N_{1}}{2 N}, \frac{N_{2}}{2 N}\right\}$, which is independent of an actual instance. Let $s_{1}(X), s_{2}(X)$ be the top two singular values of $X$.

Procedure Classify: Given $\gamma, N, \omega_{\min }$. Assume that $N \gg \frac{1}{\gamma}$,

- Normalization: use the $K$ features to form a random $n \times K$ matrix $X$; Each individual random variable $X_{i, j}$ is a normalized random variable based on the original Bernoulli r.v. $b_{i, j} \in\{0,1\}$ with $\operatorname{Pr}\left[b_{i, j}=1\right]=p_{1}^{j}$ for $X_{i} \in P_{1}$ and $\operatorname{Pr}\left[b_{i, j}=1\right]=p_{2}^{j}$ for $X_{i} \in P_{2}$, such that $X_{i, j}=\frac{b+1}{2}$.
- Take top two left singular vectors $u_{1}, u_{2}$ of $X$, where $u_{i}=\left[u_{i, 1}, \ldots, u_{i, n}\right], i=$ 1,2 .

1. If $s_{2}(X)>T=\frac{15 N}{32} \sqrt{3 \omega_{\min } \gamma}$, use $u_{2}$ to partition the individuals with 0 as the threshold, i.e., partition $j \in[n]$ according to $u_{2, j}<0$ or $u_{2, j} \geq 0$.
2. Otherwise, use $u_{1}$ to partition, with mixture mean $M=\sum_{i=1}^{n} u_{1, n}$ as the threshold.

Analysis of the Simple Algorithm: Our analysis is based on comparing entries in the top two singular vectors of the normalized random $n \times K$ matrix $X$, with those of a static matrix $\mathcal{X}$, where each entry $\mathcal{X}_{i, j}=\mathbf{E}\left[X_{i, j}\right]$ is the expected value of the corresponding entry in $X$. Hence $\forall i=1, \ldots, N_{1}$, $\mathcal{X}_{i}=\left[\mu_{1}^{1}, \mu_{1}^{2}, \ldots, \mu_{1}^{K}\right]$, where $\mu_{1}^{j}=\frac{1+p_{1}^{j}}{2}, \forall j$, and $\forall i=N_{1}+1, \ldots, n, \mathcal{X}_{i}=$ $\left[\mu_{2}^{1}, \mu_{2}^{2}, \ldots, \mu_{2}^{K}\right]$, where $\mu_{2}^{j}=\frac{1+p_{2}^{j}}{2}, \forall j$. We assume the divergence is exactly $\gamma$ among the $K$ features that we have chosen in all calculations.

The inspiration for this approach is based on Lemma 2.1, whose proof is built upon Theorem A. 2 that is presented in a lecture note by Spielman (2002). For an $n \times K$ matrix $A$, let $s_{1}(A) \geq s_{2}(A) \geq \cdots \geq s_{n}(A)$ be singular values of $A$. Let $u_{1}, \ldots, u_{n}, v_{1}, \ldots, v_{n}$, be the $n$ left and right singular vectors of $X$, corresponding to $s_{1}(X), \ldots, s_{n}(X)$ such that $\left\|u_{i}\right\|_{2}=1,\left\|v_{i}\right\|_{2}=1, \forall i$. We denote the set of $n$ left and right singular vectors of $\mathcal{X}$ with $\bar{u}_{1}, \ldots, \bar{u}_{n}, \bar{v}_{1}, \ldots, \bar{v}_{n}$.

Lemma 2.1. Let $X$ be the random $n \times K$ matrix and $\mathcal{X}$ its expected value matrix. Let $A=X-\mathcal{X}$ be the zero-mean random matrix. Let $\theta_{i}$ be the angle between two vectors: $\left[u_{i}, v_{i}\right],\left[\bar{u}_{i}, \bar{v}_{i}\right]$, where $\left\|\left[u_{i}, v_{i}\right]\right\|_{2}=\left\|\left[\bar{u}_{i}, \bar{v}_{i}\right]\right\|_{2}=2$ and $[u, v]$
represents a vector that is the concatenation of two vectors $u, v$.

$$
\begin{equation*}
\left\|u_{i}-\bar{u}_{i}\right\|_{2} \leq\left\|\left[u_{i}, v_{i}\right]-\left[\bar{u}_{i}, \bar{v}_{i}\right]\right\|_{2} \approx 2 \theta_{i} \approx 2 \sin \left(\theta_{i}\right) \leq \frac{4 s_{1}(A)}{\operatorname{gap}(i, \mathcal{X})} \tag{2.1}
\end{equation*}
$$

where $\operatorname{gap}(i, \mathcal{X})=\min _{j \neq i}\left|s_{i}(\mathcal{X})-s_{j}(\mathcal{X})\right|$.
We first bound the largest singular value $s_{1}(A)=s_{1}(X-\mathcal{X})$ of $\left(a_{i, j}\right)$ with independent zero-mean entries, which defines the Euclidean operator norm

$$
\begin{equation*}
\left\|\left(a_{i, j}\right)\right\|:=\sup \left\{\sum_{i, j} a_{i, j} x_{i} y_{j}: \sum x_{i}^{2} \leq 1, \sum y_{i}^{2} \leq 1\right\} \tag{2.2}
\end{equation*}
$$

The behavior of the largest singular value of an $n \times m$ random matrices $A$ with i.i.d. entries is well studied. Latala (2005) shows that the weakest assumption for its regular behavior is boundedness of the fourth moment of the entries, even if they are not identically distributed. Combining Theorem 2.2 of Latala (2005) with the concentration Theorem 2.3 by Meckes (2004) proves Theorem 2.4 that we need ${ }^{1}$.

Theorem 2.2 (Norm of Random Matrices Latala (2005)). For any finite $n \times m$ matrix $A$ of independent mean zero r.v.'s $a_{i, j}$ we have, for an absolute constant $C$,

$$
\begin{equation*}
\mathbf{E}\left\|\left(a_{i, j}\right)\right\| \leq C\left(\max _{i} \sqrt{\sum_{j} \mathbf{E} a_{i, j}^{2}}+\max _{j} \sqrt{\sum_{i} \mathbf{E} a_{i, j}^{2}}+\left(\sum_{i, j} \mathbf{E} a_{i, j}^{4}\right)^{\frac{1}{4}}\right) \tag{2.3}
\end{equation*}
$$

Theorem 2.3 (Concentration of Largest Singular Value: Bounded Range Meckes (2004)). For any finite $n \times m$, where $n \leq m$, matrix $A$, such that entries $a_{i, j}$ are independent r.v. supported in an interval of length at most $D$, then, for all $t$,

$$
\begin{equation*}
\operatorname{Pr}\left[\left|s_{1}(A)-\mathbb{M} s_{1}(A)\right| \geq t\right] \leq 4 e^{-t^{2} / 4 D^{2}} \tag{2.4}
\end{equation*}
$$

Theorem 2.4 (Largest Singular Value of a Mean-zero Random Matrix). For any finite $n \times K$, where $n \leq K$, matrix $A$, such that entries $a_{i, j}$ are independent mean zero r.v. supported in an interval of length at most $D$, with fourth moment upper bounded by $B$, then

$$
\begin{equation*}
\operatorname{Pr}\left[s_{1}(A) \geq C B^{1 / 4} \sqrt{K}+4 D \sqrt{\pi}+t\right] \leq 4 e^{-t^{2} / 4} \tag{2.5}
\end{equation*}
$$

for all $t$. Hence $\|A\| \leq C_{1} B^{1 / 4} \sqrt{K}$ for an absolute constant $C_{1}$.

[^1]
### 2.1. Generating a large gap in $s_{1}(\mathcal{X}), s_{2}(\mathcal{X})$

In order to apply Lemma 2.1 to the top two singular vectors of $X$ and $\mathcal{X}$ through

$$
\begin{align*}
\left\|u_{1}-\bar{u}_{1}\right\|_{2} & \leq \frac{4 s_{1}(X-\mathcal{X})}{\left|s_{1}(\mathcal{X})-s_{2}(\mathcal{X})\right|}  \tag{2.6}\\
\left\|u_{2}-\bar{u}_{2}\right\|_{2} & \leq \frac{4 s_{1}(X-\mathcal{X})}{\min \left(\left|s_{1}(\mathcal{X})-s_{2}(\mathcal{X})\right|,\left|s_{2}(\mathcal{X})\right|\right)} \tag{2.7}
\end{align*}
$$

we need to first bound $\left|s_{1}(\mathcal{X})-s_{2}(\mathcal{X})\right|$ away from zero, since otherwise, RHSs on both (2.6) and (2.7) become unbounded. We then analyze

$$
\operatorname{gap}(2, \mathcal{X})=\min \left(\left|s_{1}(\mathcal{X})-s_{2}(\mathcal{X})\right|,\left|s_{2}(\mathcal{X})\right|\right)
$$

Let us first define values $a, b, c$ that we use throughout the rest of the paper:

$$
\begin{equation*}
a=\sum_{k=1}^{K}\left(\mu_{1}^{k}\right)^{2}, \quad b=\sum_{k=1}^{K} \mu_{1}^{k} \mu_{2}^{k}, \quad c=\sum_{k=1}^{K}\left(\mu_{2}^{k}\right)^{2} . \tag{2.8}
\end{equation*}
$$

For the following analysis, we can assume that $a, b, c \in[K / 4, K]$, given that $X$ is normalized in Procedure Classify.

We first show that normalization of $X$ as described in Procedure Classify guarantees that not only $\left|s_{1}(\mathcal{X})-s_{2}(\mathcal{X})\right| \neq 0$, but there also exists a $\Theta(\sqrt{N K})$ amount of gap between $s_{1}(\mathcal{X})$ and $s_{2}(\mathcal{X})$ in Proposition 2.5:

$$
\begin{equation*}
\operatorname{gap}(\mathcal{X}):=\left|s_{1}(\mathcal{X})-s_{2}(\mathcal{X})\right|=\Theta(\sqrt{N K}) \tag{2.9}
\end{equation*}
$$

Proposition 2.5. For a normalized random matrix $X$, its expected value matrix $\mathcal{X}$ satisfies $\frac{4 c_{0} \sqrt{2 N K}}{5} \leq \operatorname{gap}(\mathcal{X}) \leq \sqrt{2 N K}$, where $c_{0}=\frac{|b| \sqrt{a c}}{K(a+c)}$ is a constant, given that $a, b, c \in[K / 4, K]$ as defined in (2.8). In addition,

$$
\begin{equation*}
\sqrt{\frac{K N}{4}} \leq s_{1}(\mathcal{X}) \leq \sqrt{2 N K}, \text { and } \sqrt{\frac{N K}{2}} \leq s_{1}(\mathcal{X})+s_{2}(\mathcal{X}) \leq \sqrt{2 N K} \tag{2.10}
\end{equation*}
$$

We next state a few important results that justify Procedure Classify. Note that the left singular vectors $\bar{u}_{i}, \forall i$ of $\mathcal{X}$ are of the form $\left[x_{i}, \ldots, x_{i}, y_{i}, \ldots, y_{i}\right]^{T}$ :

$$
\begin{equation*}
\bar{u}_{1}=\left[x_{1}, \ldots, x_{1}, y_{1}, \ldots, y_{1}\right]^{T}, \quad \text { and } \bar{u}_{2}=\left[x_{2}, \ldots, x_{2}, y_{2}, \ldots, y_{2}\right]^{T} \tag{2.11}
\end{equation*}
$$

where $x_{i}$ repeats $N_{1}$ times and $y_{i}$ repeats $N_{2}$ times. We first show Proposition 2.6 regarding signs of $x_{i}, y_{i}, i=1,2$, followed by a lemma bounding the separation of $x_{2}, y_{2}$. We then state the key Separation Lemma that allows us to conclude that least one of top two left singular vectors of $X$ can be used to classify data at each round. It can be extended to cases when $k>2$.

Proposition 2.6. Let $b$ as defined in (2.8): when $b>0$, entries $x_{1}$, $y_{1}$ in $\bar{u}_{1}$ have the same sign while $x_{2}, y_{2}$ in $\bar{u}_{2}$ have opposite signs.

Lemma 2.7. $\left|x_{2}-y_{2}\right|^{2} \leq \frac{C_{\max }}{2 N}$ where $C_{\max }=\left(\sqrt{\frac{1}{\omega_{1}}}+\sqrt{\frac{1}{\omega_{2}}}\right)^{2} \leq \frac{4}{\omega_{\min }} ;\left|x_{2}\right|^{2} \geq$ $\frac{C_{x \min }}{2 N}$ where $C_{x \text { min }}=\frac{\omega_{2}}{4 \omega_{1}^{2}+\omega_{1} \omega_{2}} ;\left|y_{2}\right|^{2} \geq \frac{C_{y \text { min }}}{2 N}$ where $C_{y \text { min }}=\frac{\omega_{1}}{4 \omega_{2}^{2}+\omega_{1} \omega_{2}}$.

Lemma 2.8 (Separation Lemma). $K \gamma=s_{1}(\mathcal{X})^{2}\left(x_{1}-y_{1}\right)^{2}+s_{2}(\mathcal{X})^{2}\left(x_{2}-y_{2}\right)^{2}$.
Proof. Let $\Delta:=P_{1}-P_{2}$ as in Theorem 1.2 , and $\vec{b}=[1,0, \ldots, 0,-1,0, \ldots, 0]^{T}$, where 1 appears in the first and -1 appears in the $N_{1}+1^{\text {st }}$ positions. Then $\Delta=X^{T} \vec{b}=\left[\mu_{1}^{1}-\mu_{2}^{1}, \mu_{1}^{2}-\mu_{2}^{2}, \ldots, \mu_{1}^{K}-\mu_{2}^{K}\right]$. Given $\mathcal{X}=s_{1}(\mathcal{X}) \bar{u}_{1} \bar{v}_{1}^{T}+s_{2}(\mathcal{X}) \bar{u}_{2} \bar{v}_{2}^{T}$, we thus rewrite $\Delta$ as: $\Delta=\mathcal{X}^{T} \vec{b}=s_{1}(\mathcal{X}) \bar{v}_{1} \bar{u}_{1}^{T} \vec{b}+s_{2}(\mathcal{X}) \bar{v}_{2} \bar{u}_{2}^{T} \vec{b}=s_{1}(\mathcal{X}) \bar{v}_{1}\left(x_{1}-\right.$ $\left.y_{1}\right)+s_{2}(\mathcal{X}) \bar{v}_{2}\left(x_{2}-y_{2}\right)$. The lemma follows from the fact that $\|\Delta\|_{2}=\sqrt{K \gamma}$ and $\bar{v}_{1}, \bar{v}_{2}$ are orthonormal.

Combining Proposition 2.6, Lemma 2.7, (2.10), and Lemma 2.8, we have
Corollary 2.9. $s_{2}(\mathcal{X}) \leq \frac{\sqrt{2 N K \gamma}}{\sqrt{c_{x} \text { min }}+\sqrt{c_{y \text { min }}}}$, and hence $\operatorname{gap}(2, \mathcal{X})=\min \left(s_{2}(\mathcal{X})\right.$, $\left.\left|s_{1}(\mathcal{X})-s_{2}(\mathcal{X})\right|\right)=s_{2}(\mathcal{X})$ for a sufficiently small $\gamma$.

In Section 2.2, we first prove a proposition regarding $a, b, c$ as defined in (2.8). We next provide the proof for Theorem 2.4 regarding the largest singular value of $(X-\mathcal{X})$. In Section 2.4, we show that the probability of error at each round for each individual is at most $f=1 / 10$, given the sample size $n$ as specified in Theorem 1.2. Hence by taking majority vote over the different runs for each sample, our algorithm will find the correct partition with probability $1-1 / n^{2}$, given that at each round we take a set of $K>n$ independent features.

### 2.2. Detailed analysis for the simple algorithm

Throughout the rest of the paper, we use $X, Y, H$ to represent random matrices, where $H=X X^{T}$ and $Y=\left[\begin{array}{cc}0 & X \\ X^{T} & 0\end{array}\right]$. We use $\mathcal{X}, \mathcal{Y}, \mathcal{H}$ to represent the corresponding static matrices. Let us substitute $a, b, c$ in $\mathcal{H}=\mathcal{X} \mathcal{X}^{T}$, where the blocks in $\mathcal{H}$ from top to bottom and from left to right are of size: $N_{1} \times N_{1}, N_{1} \times$ $N_{2}, N_{2} \times N_{1}$ and $N_{2} \times N_{2}$ respectively:

$$
\mathcal{H}=\mathcal{X X}^{T}=\left[\begin{array}{cccccc}
a & \ldots & a & b & \ldots & b  \tag{2.12}\\
\ldots & & & & & \\
a & \ldots & a & b & \ldots & b \\
b & \ldots & b & c & \ldots & c \\
\ldots & \ldots & b & c & \ldots & c
\end{array}\right]_{2 N \times 2 N}
$$

Proposition 2.10. For any choices of $\mu_{i}^{k}, a c \geq b^{2}$; By definition,

$$
\begin{equation*}
a+c-2 b=\sum_{i=1}^{K} \alpha_{k}^{2}, \text { where } \alpha_{k}=\left|\mu_{1}^{k}-\mu_{2}^{k}\right| . \tag{2.13}
\end{equation*}
$$

Proof. $a+c-2 b=\sum_{k} \alpha_{k}^{2}$ holds by definition.

$$
\begin{aligned}
a c-b^{2}= & \sum_{k=1}^{K}\left(\mu_{1}^{k}\right)^{2} \sum_{k=1}^{K}\left(\mu_{2}^{k}\right)^{2}-\left(\sum_{k=1}^{K}\left(\mu_{1}^{k} \mu_{2}^{k}\right)\right)^{2} \\
= & \sum_{k=1}^{K}\left(\mu_{1}^{k} \mu_{1}^{k}\right)^{2}+\sum_{j \neq k}\left(\left(\mu_{1}^{k} \mu_{2}^{j}\right)^{2}+\left(\mu_{1}^{j} \mu_{2}^{k}\right)^{2}\right) \\
& -\left(\sum_{k=1}^{K}\left(\mu_{1}^{k} \mu_{1}^{k}\right)^{2}+\sum_{j \neq k} 2 \mu_{1}^{k} \mu_{2}^{k} \mu_{1}^{j} \mu_{2}^{j}\right) \\
= & \sum_{j \neq k}\left(\mu_{1}^{k} \mu_{2}^{j}\right)^{2}+\left(\mu_{1}^{j} \mu_{2}^{k}\right)^{2}-2 \mu_{1}^{k} \mu_{2}^{k} \mu_{1}^{j} \mu_{2}^{j}=\sum_{j \neq k}\left(\mu_{1}^{k} \mu_{2}^{j}-\mu_{1}^{j} \mu_{2}^{k}\right)^{2} \geq 0
\end{aligned}
$$

Remark 2.11. Both matrices of $\mathcal{X}$ and $\mathcal{X} \mathcal{X}^{T}$ have rank at most two. When $a c=b^{2}, \mathcal{H}$ has rank 1 .

### 2.3. Proof of Theorem 2.4

By having an upper bound on both maximum variance and fourth moment of any entry, we have the following corollary of Theorem 2.2.

Corollary 2.12 (Largest Singular Value: Bounded Fourth Moment Latala (2005)). For any finite $n \times m$, where $n \leq m$, matrix of independent mean zero r.v.'s $a_{i, j}$, such that the maximum variances of any entry is at most $\sigma^{2}$, and each entry has a finite fourth moment $B$ we have

$$
\begin{equation*}
\mathbf{E}\left\|\left(a_{i, j}\right)\right\| \leq C\left(\sigma(\sqrt{m}+\sqrt{n})+(m n B)^{1 / 4}\right) \leq C B^{1 / 4} \sqrt{m} \tag{2.14}
\end{equation*}
$$

for an absolute constant $C$.
Remark 2.13. The requirement that $\sigma^{2}$ is upper bounded is not essential. The conclusion in Corollary 2.12 works so long as fourth moment is bounded by $B$.

Let $\mathbb{M} s_{1}(A)$ be the median of $s_{1}(A)$. Following a calculation from Meckes (2004), we have

$$
\begin{aligned}
\left|\mathbf{E}\left[s_{1}(A)\right]-\mathbb{M} s_{1}(A)\right| & \leq \mathbf{E}\left[\left|s_{1}(A)-\mathbb{M} s_{1}(A)\right|\right] \\
& =\int_{0}^{\infty} \operatorname{Pr}\left[\left|s_{1}(A)-\mathbb{M} s_{1}(A)\right| \geq t\right] d t \\
& \leq 4 \int_{0}^{\infty} e^{-t^{2} / 4 D^{2}} d t=4 D \sqrt{\pi}
\end{aligned}
$$

where $D \leq 1$ for Bernoulli random variables that we consider. This allows us to conclude Theorem 2.4.

### 2.4. Correctness of classification for the simple algorithm

We now prove correctness of our algorithm. We first show how to choose $T$ for Procedure Classify. Let $B$ denote the fourth moment bound for a single random variable in the mean zero random matrix $X-\mathcal{X}$; for the type of normalized Bernoulli r.v.s that we care about, $\sqrt{B}$ is in the order of $\sigma^{2}$, where $\sigma^{2}$ is defined in Theorem 1.2.

Let $N \gamma$ be a large enough constant. Let $s_{1}(X-\mathcal{X}) \leq C_{0} \sqrt{K}$, where $C_{0}=$ $C_{1} B^{1 / 4}$ as defined in Theorem 2.4 and let the threshold

$$
\begin{equation*}
T=\sqrt{C_{3} K N \gamma} \geq 15 C_{0} \sqrt{K} \tag{2.15}
\end{equation*}
$$

which requires that

$$
\begin{equation*}
C_{3} N \gamma \geq 225 C_{0}^{2}, \text { where } C_{3} \text { satisfies } \tag{2.16}
\end{equation*}
$$

Following Lemma A.4, (A.1), (A.3), and Proposition A.1, we have

$$
\begin{equation*}
\left|s_{2}(\mathcal{X})-s_{2}(X)\right| \leq s_{1}(X-\mathcal{X}) \leq C_{0} \sqrt{K} \tag{2.17}
\end{equation*}
$$

We have two cases,

1. When $s_{2}(X) \leq T$, by Lemma 2.7 and the fact that $s_{2}(\mathcal{X}) \leq s_{2}(X)+$ $s_{1}(X-\mathcal{X}) \leq T+C_{0} \sqrt{K} \leq \frac{16 T}{15}$, we have

$$
\begin{equation*}
s_{2}(\mathcal{X})^{2}\left|x_{2}-y_{2}\right|^{2} \leq \frac{256 T^{2}}{225} \frac{C_{\max }}{2 N} \leq \frac{128 C_{3} K \gamma C_{\max }}{225} \leq \frac{128 C_{3} K \gamma 4}{225 \omega_{\min }} \tag{2.18}
\end{equation*}
$$

for $C_{\max }$ as defined in Lemma 2.7. We want $s_{2}(\mathcal{X})^{2}\left|x_{2}-y_{2}\right|^{2} \leq \frac{3 K \gamma}{4}$. This holds so long as $\frac{128 C_{3} K \gamma C_{\max }}{225} \leq \frac{128 C_{3} K \gamma 4}{225 \omega_{\min }} \leq \frac{3 K \gamma}{4}$, which is true if

$$
\begin{equation*}
C_{3} \leq \frac{675 \omega_{\min }}{2048} ; \text { thus we take } C_{3}=\frac{675 \omega_{\min }}{2048} \text { from this point on. } \tag{2.19}
\end{equation*}
$$

It follows from Lemma 2.8 that $s_{1}(\mathcal{X})^{2}\left|x_{1}-y_{1}\right|^{2} \geq \frac{K \gamma}{4}$. Hence by (2.10)

$$
\begin{equation*}
\left|x_{1}-y_{1}\right| \geq \frac{\sqrt{K \gamma}}{2 s_{1}(\mathcal{X})} \geq \frac{\sqrt{K \gamma}}{2 \sqrt{2 N K}} \geq \frac{1}{2} \sqrt{\frac{\gamma}{2 N}} \tag{2.20}
\end{equation*}
$$

Thus the condition of Theorem 2.14 holds with $c_{2}=\frac{1}{2}$, so long as

$$
\begin{equation*}
N \gamma \geq \frac{2048 C_{1}^{2} \sqrt{B}}{3 \omega_{\min }} \tag{2.21}
\end{equation*}
$$

due to (2.16) and (2.19); This is a weaker condition than (2.32) for $f<\frac{1}{2}$.
2. When $s_{2}(X) \geq T$, we have $s_{2}(\mathcal{X}) \geq s_{2}(X)-s_{1}(X-\mathcal{X}) \geq T-C_{0} \sqrt{K} \geq$ $\frac{14 T}{15}$; This satisfies the condition of Theorem 2.18, with $c_{3}=\frac{14 \sqrt{C_{3}}}{15}=$ $\frac{7}{16} \sqrt{\frac{3 \omega_{\min }}{2}}$.

Let us first denote the first singular vector $u_{1}$ and its "noise" vector $\epsilon$ as follows:
$u_{1}^{T}=\left(x+\delta_{1}, \ldots, x+\delta_{N_{1}}, y+\tau_{1}, \ldots, y+\tau_{N_{2}}\right), \epsilon^{T}=\left(\delta_{1}, \ldots, \delta_{N_{1}}, \tau_{1}, \ldots, \tau_{N_{2}}\right)$.
It turns out that we only need to use the mixture mean

$$
\begin{equation*}
M=\frac{\sum_{i=1}^{N_{1}}\left(x+\delta_{i}\right)+\sum_{i=1}^{N_{2}}\left(y+\tau_{i}\right)}{2 N} \tag{2.22}
\end{equation*}
$$

to decide which side to put a node, i.e., to partition $j \in[2 N]$ according to $u_{1, j}<M$ or $u_{1, j} \geq M$, given that $N_{1} / N_{2}$ is a constant; Misclassifying any entry will contribute $\Omega\left(\frac{\gamma}{2 N}\right)$ amount to $\left\|\bar{u}_{1}-u_{1}\right\|_{2}^{2}$.
Theorem 2.14. Assume w.l.o.g. that $N_{1} \leq N_{2}$ and $2 N \leq K$. Let $\omega_{1}=N_{1} / 2 N$ and $\omega_{2}=N_{2} / 2 N$. Suppose $\left|x_{1}-y_{1}\right| \geq c_{2} \sqrt{\frac{\gamma}{2 N}}$ for some constant $c_{2}=\frac{1}{2}$. By requiring $N \geq \frac{2048 C_{1}^{2} \sqrt{B}}{3 \gamma \omega_{\min }}$ as in (2.21), and

$$
\begin{equation*}
N_{1} \geq \frac{2 c_{1}^{2} \sigma^{2}}{f c_{2}^{2} \gamma \omega_{1} \omega_{2}}, \text { or equivalently } 2 N \geq \frac{2 c_{1}^{2} \sigma^{2}}{f c_{2}^{2} \gamma \omega_{2} \omega_{1}^{2}}=\frac{25 C_{1}^{2} \sqrt{B}}{f c_{2}^{2} \gamma \omega_{2} \omega_{1}^{2}}, \tag{2.23}
\end{equation*}
$$

where $c_{1}=\frac{5 C_{1} B^{1 / 4}}{\sqrt{2} c_{0} \sigma}$ for $C_{1}$ specified in Theorem 2.4 and $c_{0}$ specified in Proposition 2.5, we can classify the two population using the mixture mean $M$ with the error factor at most $f$ for $N_{1}, N_{2}$ respectively whp.

By Lemma 2.1 and Theorem 2.4, we immediately have the following claim.
Claim 2.15. For $c_{1}$ chosen as in Theorem 2.14, $\|\epsilon\|_{2}^{2}=\sum_{i=1}^{N_{1}} \delta_{i}^{2}+\sum_{i=1}^{N_{2}} \tau_{i}^{2} \leq$ $\frac{c_{1}^{2} \sigma^{2}}{N}$.
Proof. Given that $c_{1}=\frac{5 C_{1} \frac{1}{\sqrt{B}}}{\sqrt{2} c_{0} \sigma}$ such that $C_{1}$ appears in Theorem 2.4 and $c_{0}$ appears in Proposition 2.5,

$$
\begin{aligned}
\sqrt{\sum_{i=1}^{N_{1}} \delta_{i}^{2}+\sum_{i=1}^{N_{2}} \tau_{i}^{2}} & =\left\|u_{1}-\bar{u}_{1}\right\|_{2} \sim 2 \theta_{1} \sim 2 \sin \left(\theta_{1}\right) \\
& \leq \frac{4 s_{1}(X-\mathcal{X})}{\operatorname{gap}(1, \mathcal{X})} \leq \frac{4 C_{1} \sqrt[1]{B} \sqrt{B} \sqrt{K}}{4 c_{0} \sqrt{2 N K} / 5}=\frac{c_{1} \sigma}{\sqrt{N}} .
\end{aligned}
$$

This allows us to conclude the claim.
We need the following lemma, proof of which appears in Appendix B.
Lemma 2.16. Assume that $2 N \leq K$ and Condition (2.23) in Theorem 2.14, we have

$$
\begin{equation*}
|M-x| \geq \frac{N_{2}(1-\sqrt{\gamma})|y-x|}{2 N},|y-M| \geq \frac{N_{1}(1-\sqrt{\gamma})|y-x|}{2 N} . \tag{2.24}
\end{equation*}
$$

Proof of Theorem 2.14. Recall that the largest $\bar{u}_{1}, \bar{u}_{2}$ have the form of $[x, \ldots, x$, $y, \ldots, y]$, where $x$ repeats $N_{1}$ times and $y$ repeats $N_{2}$ times; hence w.l.o.g., assume that $x<y$, we have

$$
\begin{align*}
& \forall i, \text { s.t. } x+\delta_{i}>M, \text { it contributes } \delta_{i}^{2}>|M-x|^{2} \geq \approx \frac{N_{2}^{2} c_{2}^{2} \gamma}{8 N^{3}} \text { to }\|\epsilon\|_{2}^{2},  \tag{2.25}\\
& \forall i, \text { s.t. } y-\tau_{i}<M, \text { it contributes } \delta_{i}^{2}>|M-y|^{2} \geq \approx \frac{N_{2}^{2} c_{2}^{2} \gamma}{8 N^{3}} \text { to }\|\epsilon\|_{2}^{2} . \tag{2.26}
\end{align*}
$$

Hence the total number of entries that goes above $M$ from $P_{1}$, and those goes below $M$ from $P_{2}$ can not be too many since their total contribution is upper bounded by $\|\epsilon\|_{2}^{2}=\left\|u_{1}-\bar{u}_{1}\right\|_{2}^{2}$. Let $\ell_{1}$ be the number of misclassified entries from $N_{1}$, i.e., those described in (2.25), by Lemma 2.1,

$$
\begin{equation*}
\ell_{1} \frac{N_{2}^{2} c_{2}^{2} \gamma}{8 N^{3}} \leq \ell_{1}|M-x|^{2} \leq\|\epsilon\|_{2}^{2} \leq \frac{c_{1}^{2} \sigma^{2}}{N} \tag{2.27}
\end{equation*}
$$

Thus given that $N_{1} \geq \frac{8 c_{1}^{2} \sigma^{2}}{f c_{2}^{2} \gamma} \geq \frac{2 c_{1}^{2} \sigma^{2}}{f \omega_{2}^{2} c_{2}^{2} \gamma}$; hence it suffices to guarantee that $\ell_{1} \leq$ $\frac{2 c_{1}^{2} \sigma^{2}}{\omega_{2}^{2} c_{2}^{2} \gamma} \leq f N_{1}$.

We next bound the number of entries from $P_{2}$ that goes below $M$, which can not be too many either; let $\ell_{2}$ be the number of misclassified entries from $P_{2}$,

$$
\begin{equation*}
\ell_{2} \frac{N_{1}^{2} c_{2}^{2} \gamma}{8 N^{3}} \leq \ell_{2}|M-y|^{2} \leq\|\epsilon\|_{2}^{2} \leq \frac{c_{1}^{2} \sigma^{2}}{N} \tag{2.28}
\end{equation*}
$$

hence by requiring

$$
\begin{equation*}
N_{2} \geq \frac{2 c_{1}^{2} \sigma^{2}}{f \omega_{1}^{2} c_{2}^{2} \gamma} \tag{2.29}
\end{equation*}
$$

it suffices to guarantee that $\ell_{2} \leq \frac{2 c_{1}^{2} \sigma^{2}}{\omega_{2}^{2} c_{2}^{2} \gamma} \leq f N_{2}$.
Condition (2.29) is equivalent to

$$
\begin{equation*}
N_{1}=\frac{N_{2} \omega_{1}}{\omega_{2}} \geq \frac{2 c_{1}^{2} \sigma^{2}}{f \omega_{1} \omega_{2} c_{2}^{2} \gamma} \tag{2.30}
\end{equation*}
$$

Thus by requiring

$$
\begin{equation*}
N_{1} \geq \frac{2 c_{1}^{2} \sigma^{2}}{f c_{2}^{2} \gamma \omega_{2} \omega_{1}} \tag{2.31}
\end{equation*}
$$

we have satisfied all requirements.
Combining Lemma 2.1 and Corollary 2.9, we have
Claim 2.17. Given that $s_{2}(\mathcal{X}) \geq c_{3} \sqrt{K N \gamma},\left\|u_{2}-\bar{u}_{2}\right\|_{2}^{2} \leq \frac{16 s_{1}(X-\mathcal{X})^{2}}{s_{2}(\mathcal{X})^{2}} \leq$ $\frac{16 C_{0}^{2} K}{c_{3}^{2} K N \gamma}$.

This allows us to prove the following theorem. Let the classification error factor be the number of misclassified individuals from one group over total amount of people in that group.

Theorem 2.18. Assume $N_{1} \leq N_{2}$ and $2 N \leq K$. Let $\omega_{1}=N_{1} / 2 N$ and $\omega_{2}=$ $N_{2} / 2 N$. Let $s_{2}(\mathcal{X}) \geq c_{3} \sqrt{K N \gamma}$, where $c_{3}=\frac{7}{16} \sqrt{\frac{3 \omega_{\min }}{2}}$ and $\omega_{\min }$ is the minimum possible weight allowed by the algorithm. By requiring

$$
\begin{equation*}
2 N \geq \frac{360 C_{0}^{2}}{\omega_{\min } f \gamma}\left(\frac{\omega_{2}}{\omega_{1}}+1\right)=\Theta\left(\frac{\sigma^{2}}{f \gamma \omega_{\min } \omega_{1}}\right) \tag{2.32}
\end{equation*}
$$

we can classify the two population using 0 to separate components in $u_{2}$, with error factor at most $f$ for both $P_{1}, P_{2} w h p$.
Proof. Let $\ell_{1}, \ell_{2}$ be the number of misclassified entries from $P_{1}$ and $P_{2}$ respectively; they each contribute at least $\frac{C_{x \min }}{2 N}$, and $\frac{C_{y \min }}{2 N}$ amount to $\left\|u_{2}-\bar{u}_{2}\right\|_{2}$, and hence by Claim 2.17,

$$
\begin{equation*}
\ell_{1} \frac{C_{x \min }}{2 N} \leq\left\|u_{2}-\bar{u}_{2}\right\|_{2}^{2} \leq \frac{16 C_{0}^{2} K}{c_{3}^{2} K N \gamma} \leq \frac{16 C_{0}^{2}}{c_{3}^{2} N \gamma} \tag{2.33}
\end{equation*}
$$

Hence $\ell_{1} \leq \frac{32 C_{0}^{2}}{c_{3}^{2} \gamma C_{x \text { min }}} \leq f N_{1}$ given that $N \geq \frac{16 C_{0}^{2}}{c_{3}^{2} f \gamma}\left(4 \frac{\omega_{1}}{\omega_{2}}+1\right)$.
Similarly, by Claim 2.17, we have $\ell_{2} \frac{C_{x \text { min }}}{2 N} \leq\left\|u_{2}-\bar{u}_{2}\right\|_{2}^{2}$ and thus $\ell_{2} \leq$ $\frac{32 C_{0}^{2}}{c_{3}^{2} \gamma C_{y \text { min }}} \leq \frac{N_{2}}{f}$ so long as $N \geq \frac{16 C_{0}^{2}}{c_{3}^{2} f \gamma}\left(4 \frac{\omega_{2}}{\omega_{1}}+1\right)$; the bound on $2 N$ follows by plugging in $c_{3}=\frac{7}{16} \sqrt{\frac{3 \omega_{\min }}{2}}$.

Finally,
Theorem 2.19. Given a set of $n \geq \Omega\left(\frac{\sigma^{2}}{\gamma f \omega \omega_{\min }}\right)$ individuals, by trying Procedure Classify for $\log n$ rounds, with probability of error at each round for each individual being $f=1 / 10$, where each round we take a set of $K>n$ independent features, and by taking majority vote over the different runs for each sample, our algorithm will find the correct partition with probability $1-1 / n^{2}$.

Proof. A sample is put in the wrong side with a probability $1 / 10$ at each round. Let $\mathcal{E}_{i}$ be the event that sample $i$ is misclassified for more than $\log n$ times, thus $\operatorname{Pr}\left[\mathcal{E}_{i}\right]=\left(\frac{1}{10}\right)^{\log n} \leq 1 / n^{3.32}$; hence by union bound, with probability $1-1 / n^{2}$, none of the $2 N$ individuals is misclassified.

## 3. The algorithm Partition

### 3.1. Preliminaries

Let $V=\{1, \ldots, n\}$ be the set of all $n$ individuals, and let $\psi: V \rightarrow\{1, \ldots, k\}$ be the map that assigns to each individual the population it belongs to. Set $V_{t}=\psi^{-1}(t)$ and $N_{t}=\left|V_{t}\right|$. Moreover, let $\mathbb{E}=\left(\mathbb{E}_{v i}\right)_{1 \leq v \leq n, 1 \leq i \leq K}$ be the $n \times K$
matrix with entries $\mathbb{E}_{v i}=p_{\psi(v)}^{i}$. For any $1 \leq t \leq k$ we let $\mathbb{E}^{V_{t}}=\left(p_{t}^{l}\right)_{l=1, \ldots, K}$ be the row of $\mathbb{E}$ corresponding to any $v \in V_{t}$. In addition, let $A=\left(a_{v i}\right)$ denote the empirical $n \times K$ input matrix. Thus, the entries of $\mathbb{E}$ equal the expectations of the entries of $A$.

As in Theorem 1.3, we let

$$
\gamma=K^{-1} \min _{1 \leq i<j \leq k}\left\|\mathbb{E}^{V_{i}}-\mathbb{E}^{V_{j}}\right\|^{2}, \quad \Gamma=K \gamma .
$$

Further, set $\lambda=\sqrt{K} \sigma$. Then the assumption from Theorem 1.3 can be rephrased as

$$
\begin{equation*}
n_{\min } \Gamma>C_{k} \lambda^{2} \text { and } \Gamma>K^{-1} \tag{3.1}
\end{equation*}
$$

where $C_{k}$ signifies a sufficiently large number that depends on $k$ only (the precise value of $C_{k}$ will be specified implicitly in the course of the analysis). As in the previous section, by repeating the partitioning process $\log n$ times, we may restrict our attention to the problem of classifying a constant fraction of the individuals correctly. That is, it is sufficient to establish the following claim.
Claim 3.1. There is a polynomial time algorithm Partition that satisfies the following. Suppose that (3.1) is true. Then whp Partition $(A, k)$ outputs a partition $\left(S_{1}, \ldots, S_{k}\right)$ of $V$ such that there exists a permutation $\sigma$ such that

$$
\sum_{i=1}^{k}\left|V_{i} \triangle S_{\sigma(i)}\right|<0.001 n_{\min } .
$$

Let $X=\left(x_{i j}\right)_{1 \leq i \leq n, 1 \leq j \leq K}$ be a $n \times K$ matrix. By $X_{i}$ we denote the $i^{\prime}$ th row $\left(X_{i 1}, \ldots, X_{i K}\right)$ of $X$. Moreover, we let

$$
\|X\|=\max _{\xi \in \mathbf{R}^{K}:\|\xi\|=1}\|X \xi\|
$$

signify the operator norm of $X$. A rank $k$ approximation of $X$ is a matrix $\widehat{X}$ of rank at most $k$ such that for any $n \times K$ matrix $Y$ of rank at most $k$ we have $\|X-\widehat{X}\| \leq\|X-Y\|$. Given $X$, a rank $k$ approximation $\widehat{X}$ can be computed as follows. Letting $\rho=\operatorname{rank}(X)$, we compute the singular value decomposition

$$
X=\sum_{i=1}^{\rho} \lambda_{i} \xi_{i} \eta_{i}^{T}
$$

here $\left(\xi_{i}\right)_{1 \leq i \leq \rho}$ is an orthonormal family in $\mathbf{R}^{n},\left(\eta_{i}\right)_{1 \leq i \leq \rho}$ is an orthonormal family in $\overline{\mathbf{R}}^{\bar{K}}$, and we assume that the singular values $\lambda_{i}$ are in decreasing order (i.e., $\lambda_{1} \geq \cdots \geq \lambda_{\rho}$ ). This can be accomplished in polynomial time within any numerical precision. Then $\widehat{X}=\sum_{i=1}^{\min \{k, \rho\}} \lambda_{i} \xi_{i} \eta_{i}^{T}$ is easily verified to be a rank $k$ approximation.

In addition to the operator norm, we are going to work with the Frobenius norm

$$
\|X\|_{F}=\sqrt{\sum_{i=1}^{n} \sum_{j=1}^{K} x_{i j}^{2}} .
$$

Although the following fact is well known, we provide its proof for completeness.
Lemma 3.2. If $X$ has rank $k$, then $\|X\|_{F}^{2} \leq k\|X\|^{2}$.
Proof. Let $X=\sum_{i=1}^{k} \lambda_{i} \xi_{i} \eta_{i}^{T}$ be a singular value decomposition as above. Then

$$
\|X\|_{F}^{2}=\sum_{i, j} x_{i j}^{2}=\sum_{i, j=1}^{k} \lambda_{i} \lambda_{j}\left\langle\xi_{i}, \xi_{j}\right\rangle\left\langle\eta_{i}, \eta_{j}\right\rangle .
$$

Since $\xi_{1}, \ldots, \xi_{k}$ and $\eta_{1}, \ldots, \eta_{k}$ are orthonormal families, we have $\left\langle\xi_{i}, \xi_{j}\right\rangle=$ $\left\langle\eta_{i}, \eta_{j}\right\rangle=1$ if $i=j$ and $\left\langle\xi_{i}, \xi_{j}\right\rangle=\left\langle\eta_{i}, \eta_{j}\right\rangle=0$ if $i \neq j$. Hence, $\|X\|_{F}^{2}=\sum_{i=1}^{k} \lambda_{i}^{2}$. This implies the assertion, because $\lambda_{i} \leq\|X\|$ for all $1 \leq i \leq k$.

### 3.2. Description of the algorithm

Algo 3.3. Partition $(A, k)$
Input: A $n \times K$ matrix $A$ and the parameter $k$. Output: A partition $S_{1}, \ldots, S_{k}$ of $V$.

1. Compute a rank $k$ approximation $\widehat{A}$ of $A$.

$$
\text { For } j=1, \ldots, 2 \log K \text { do }
$$

Let $\Gamma_{j}=K 2^{-j}$ and compute $Q^{(j)}(v)=\left\{w \in V:\left\|\widehat{A}_{w}-\widehat{A}_{v}\right\|^{2} \leq 0.01 \Gamma_{j}^{2}\right\}$ for all $v \in V$.
Then, determine sets $Q_{1}^{(j)}, \ldots, Q_{k}^{(j)}$ as follows: for $i=1, \ldots, k$ do
3. $\quad$ Pick $v \in V \backslash \bigcup_{l=1}^{i-1} Q_{l}^{(j)}$ such that $\left|Q^{(j)}(v) \backslash \bigcup_{l=1}^{i-1} Q_{l}^{(j)}\right|$ is maximum. Set $Q_{i}^{(j)}=Q^{(j)}(v) \backslash \bigcup_{l=1}^{i-1} Q_{l}^{(j)}$ and $\xi_{i}^{(j)}=\frac{1}{\left|Q_{i}^{(j)}\right|} \sum_{w \in Q_{i}^{(j)}} \widehat{A}_{w}$.
4. Partition the entire set $V$ as follows: first, let $S_{i}^{(j)}=Q_{i}^{(j)}$ for all $1 \leq i \leq k$.

Then, add each $v \in V \backslash \bigcup_{l=1}^{k} Q_{l}^{(j)}$ to a set $S_{i}^{(j)}$ such that $\left\|\widehat{A}_{v}-\xi_{i}^{(j)}\right\|$ is minimum.
Set $r_{j}=\sum_{i=1}^{k} \sum_{v \in S_{i}^{(j)}}\left\|\widehat{A}_{v}-\xi_{i}^{(j)}\right\|^{2}$.
5. Let $J$ be such that $r^{*}=r_{J}$ is minimum. Return $S_{1}^{(J)}, \ldots, S_{k}^{(J)}$.

The basic idea behind Partition is to classify each individual $v \in V$ according to its row vector $\widehat{A}_{v}$ in the rank $k$ approximation $\widehat{A}$. That is, two individuals $v, w$ are deemed to belong to the same population iff $\left\|\widehat{A}_{v}-\widehat{A}_{w}\right\|^{2} \leq 0.01 \Gamma^{2}$. Hence, Partition tries to determine sets $S_{1}, \ldots, S_{k}$ such that for any two $v, w$ in the same set $S_{j}$ the distance $\left\|\widehat{A}_{v}-\widehat{A}_{w}\right\|$ is small. To justify this approach, we show that $\widehat{A}$ is "close" to the expectation $\mathbb{E}$ of $A$ in the following sense.
Lemma 3.4. There is a constant $C>0$ such that $\sum_{v \in V}\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|^{2} \leq C k \lambda^{2}$ whp.
Proof. Since both $\widehat{A}$ and $\mathbb{E}$ have rank at most $k$, and as therefore $\widehat{A}-\mathbb{E}$ has rank at most $2 k$, Lemma 3.2 yields

$$
\sum_{v \in V}\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|^{2}=\|\widehat{A}-\mathbb{E}\|_{F}^{2} \leq 2 k\|\widehat{A}-\mathbb{E}\|^{2}
$$

Furthermore, $\|\widehat{A}-\mathbb{E}\| \leq\|\widehat{A}-A\|+\|\mathbb{E}-A\| \leq 2\|\mathbb{E}-A\|$, because $\widehat{A}$ is a rank $k$ approximation of $A$. As Theorem 2.4 implies that $\|A-\mathbb{E}\|^{2} \leq C \lambda^{2} / 8$ for a certain constant $C>0$, we thus obtain $\sum_{v \in V}\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|^{2} \leq 8 k\|A-\mathbb{E}\|^{2} \leq C k \lambda^{2}$.

Observe that Lemma 3.4 implies that for most $v$ we have $\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|^{2} \leq 10^{-6} \Gamma$, say. For letting $z=\left|\left\{v:\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|^{2}>10^{-6} \Gamma\right\}\right|$, we get

$$
10^{-6} \Gamma z \leq \sum_{v \in V}\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|^{2} \leq C k \lambda^{2}
$$

whence $z \ll n_{\text {min }}$ due to our assumption that $n_{\min } \Gamma \gg k \lambda^{2}$. Thus, most rows of $\widehat{A}$ are close to the corresponding rows of the expected matrix $\mathbb{E}$. Therefore, the separation assumption $n>\frac{C_{k} \sigma^{2}}{\gamma \omega}$ from Theorem 1.3 implies that for most pairs of elements in different classes $v \in V_{i}, w \in V_{j}$ the squared distance $\left\|\widehat{A}_{v}-\widehat{A}_{w}\right\|^{2}$ will be large (at least $0.99 \Gamma$, say). By contrast, for most pairs $u, v \in V_{i}$ of elements belonging to the same class $\left\|\widehat{A}_{v}-\widehat{A}_{w}\right\|^{2}$ will be small (at most $0.01 \Gamma$, say), because $\mathbb{E}_{v}=\mathbb{E}_{w}$.

As the above discussion indicates, if the algorithm were given $\Gamma$ as an input parameter, the procedure described in Steps $2-4$ (with $\Gamma_{j}$ replaced by $\Gamma$ ) would yield the desired partition of $V$. The procedure described in Steps $2-4$ is very similar to the spectral partitioning algorithm from McSherry (2001).

However, since $\Gamma$ is not given to the algorithm as an input parameter, PartiTION has to estimate $\Gamma$ on its own. (This is the new aspect here in comparison to McSherry (2001), and this fact necessitates a significantly more involved analysis.) To this end, the outer loop goes through $2 \log K$ "candidate values" $\Gamma_{j}$. These values are then used to obtain partitions $Q_{1}^{(j)}, \ldots, Q_{k}^{(j)}$ in Steps 2-4. More precisely, Step 2 uses $\Gamma_{j}$ to compute for each $v \in V$ the set $Q(v)$ of elements $w$ such that $\left\|\widehat{A}_{w}-\widehat{A}_{v}\right\| \leq 0.01 \Gamma_{j}^{2}$. Then, Step 3 tries to compute "big" disjoint $Q_{1}^{(j)}, \ldots, Q_{k}^{(j)}$, where each $Q_{i}^{(j)}$ results from some $Q\left(v_{i}\right)$. Further, Step 4 assigns all elements $v$ not covered by $Q_{1}^{(j)}, \ldots, Q_{k}^{(j)}$ to that $Q_{i}^{(j)}$ whose "center vector" $\xi_{i}^{(j)}$ is closest to $\widehat{A}_{v}$. In addition, Step 4 computes an "error parameter" $r_{j}$. Finally, Step 5 outputs the partition that minimizes the error parameter $r_{j}$.

Thus, we need to show that eventually picking the partition whose error term $r_{j}$ is minimum yields a good approximation to the ideal partition $V_{1}, \ldots, V_{k}$. The basic reason why this is true is that the "empirical" mean $\xi_{i}^{(j)}$ should approximate the expectation $\mathbb{E}^{V_{i}}$ for class $V_{i}$ well iff $Q_{i}^{(j)}$ is a good approximation of $V_{i}$. Hence, if $Q_{1}^{(j)}, \ldots, Q_{k}^{(j)}$ is "close" to $V_{1}, \ldots, V_{k}$, then

$$
r_{j}=\sum_{i=1}^{k} \sum_{v \in S_{i}^{(j)}}\left\|\widehat{A}_{v}-\xi_{i}^{(j)}\right\|^{2}
$$

will be about as small as $\|\widehat{A}-\mathbb{E}\|_{F}^{2}$ (cf. Lemma 3.4). In fact, the following lemma shows that if $\Gamma_{j}$ is "close" to the ideal $\Gamma$, then $r_{j}$ will be small.

Lemma 3.5. If $\frac{1}{2} \Gamma \leq \Gamma_{j} \leq \Gamma$, then $r_{j} \leq C_{0} k^{3} \lambda^{2}$ for a certain constant $C_{0}>0$.
We defer the proof of Lemma 3.5 to Section 3.3. Furthermore, the next lemma shows that any partition such that $r_{j}$ is small yields a good approximation to $V_{1}, \ldots, V_{k}$.

Lemma 3.6. Let $S_{1}, \ldots, S_{k}$ be a partition and $\xi_{1}, \ldots, \xi_{k}$ a sequence of vectors such that $\sum_{i=1}^{k} \sum_{v \in S_{i}}\left\|\xi_{i}-\widehat{A}_{v}\right\|^{2} \leq C_{0} k^{3} \lambda^{2}$. Then there is a bijection $\Xi:\{1, \ldots, k\} \rightarrow\{1, \ldots, k\}$ such that the following holds.

1. $\left\|\xi_{i}-\mathbb{E}^{V_{\Xi(i)}}\right\|^{2} \leq 0.001 \Gamma^{2}$ for all $i=1, \ldots, k$, and
2. $\sum_{i=1}^{k}\left|S_{i} \triangle V_{\Xi(i)}\right|<0.001 n_{\text {min }}$.

The proof of Lemma 3.6 can be found in Section 3.4.
Proof of Claim 3.1. Since the rank $k$ approximation $\widehat{A}$ can be computed in polynomial time (within any numerical precision), Partition is a polynomial time algorithm. Hence, we just need to show that $K^{-1} \leq \Gamma \leq K$; for then Partition will eventually try a $\Gamma_{j}$ such that $\frac{1}{2} \Gamma \leq \Gamma_{j} \leq \Gamma$, so that the claim follows from Lemmas 3.5 and 3.6. To see that $K^{-1} \leq \Gamma \leq K$, recall that we explicitly assume that $\Gamma>K^{-1}$. Furthermore, all entries of the vectors $\mathbb{E}^{V_{i}}$ lie between 0 and 1 , whence $\Gamma=\max _{i<j}\left\|\mathbb{E}^{V_{i}}-\mathbb{E}^{V_{j}}\right\|^{2} \leq K$.

### 3.3. Proof of Lemma 3.5

Suppose that $\frac{1}{2} \Gamma \leq \Gamma_{j} \leq \Gamma$. To ease up the notation, we omit the superscript $j$; thus, we let $S_{i}=S_{i}^{(j)}, Q_{i}=Q_{i}^{(j)}$ for $1 \leq i \leq k$, and $Q(v)=Q^{(j)}(v)$ for $v \in V$ (cf. Steps 2-4 of Partition). The following lemma shows that there is a permutation $\pi$ such that $\xi_{i}$ is "close" to $\mathbb{E}^{V_{\pi(i)}}$ for all $1 \leq i \leq k$, and that the sets $Q_{i}$ are "not too small".

Lemma 3.7. Suppose that $\frac{1}{2} \Gamma \leq \Gamma_{j} \leq \Gamma$. There is a bijection $\pi:\{1, \ldots, k\} \rightarrow$ $\{1, \ldots, k\}$ such that for each $1 \leq i \leq k$ we have $\left|Q_{i}\right| \geq \frac{1}{2}\left|V_{\pi(i)}\right|$ and $\| \xi_{i}-$ $\mathbb{E}^{V_{\pi(i)}} \|^{2} \leq 0.1 \Gamma$.
Proof. For $1 \leq i \leq k$ we choose $\pi(i)$ so that $\left|Q_{i} \cap V_{\pi(i)}\right|$ is maximum. We shall prove below that for all $1 \leq l \leq k$ we have

$$
\begin{align*}
\| \xi_{l}-\mathbb{E}^{V_{\pi(l)} \|^{2}} & \leq 0.1 \Gamma  \tag{3.2}\\
\left|Q_{l}\right| & \geq \max \left\{\left|V_{i}\right|: i \in\{1, \ldots, k\} \backslash \pi(\{1, \ldots, l-1\})\right\}-0.01 n_{\min },  \tag{3.3}\\
\left|Q_{l} \cap V_{\pi(l)}\right| & \geq\left|Q_{l}\right|-0.01 n_{\min } \tag{3.4}
\end{align*}
$$

These three inequalities imply the assertion. To see that $\pi$ is a bijection, let us assume that $\pi(l)=\pi\left(l^{\prime}\right)$ for two indices $1 \leq l<l^{\prime} \leq k$. Indeed, suppose that $l=\min \pi^{-1}(l)$. Then $\left|Q_{l}\right| \geq\left|V_{\pi(l)}\right|-0.01 n_{\min }$ by (3.3), and thus $\left|V_{\pi(l)} \backslash Q_{l}\right| \leq$ $0.1 n_{\text {min }}$ by (3.4). As $Q_{l} \cap Q_{l^{\prime}}=\emptyset$ by construction, we obtain the contradiction

$$
0.99 n_{\min } \stackrel{(3.3)}{\leq}\left|Q_{l^{\prime}}\right| \stackrel{(3.4)}{\leq} 1.1\left|Q_{l^{\prime}} \cap V_{\pi(l)}\right| \leq 1.1\left|V_{\pi(l)} \backslash Q_{l}\right| \leq 0.11 n_{\min }
$$

Finally, as $\pi$ is bijective, (3.3) entails that $\left|Q_{l}\right| \geq 0.9 V_{\pi(l)}$ for all $1 \leq l \leq k$. Hence, due to (3.4) we obtain $\left|Q_{l} \cap V_{l}\right| \geq 0.9\left|Q_{l}\right| \geq \frac{1}{2}\left|V_{\pi(l)}\right|$, as desired.

The remaining task is to establish (3.2)-(3.4). We proceed by induction on $l$. Thus, let us assume that (3.2)-(3.4) hold for all $l<L$; we are to show that then (3.2)-(3.4) are true for $l=L$ as well. As a first step, we establish (3.3). To this end, consider a class $V_{i}$ such that $i \notin \pi(\{1, \ldots, L-1\})$ and let $Z_{i}=\{v \in$ $\left.V_{i}:\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|^{2} \leq 0.001 \Gamma\right\}$. Then $0.001 \Gamma\left(\left|V_{i}\right|-\left|Z_{i}\right|\right) \leq \sum_{v \in V_{i} \backslash Z_{i}}\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|^{2} \leq$ $\|\widehat{A}-\mathbb{E}\|_{F}^{2} \leq C k \lambda^{2}$ (cf. Lemma 3.4) whence the assumption (3.1) on $\Gamma$ yields

$$
\begin{equation*}
\left|Z_{i}\right| \geq\left|V_{i}\right|-0.01 n_{\min } \tag{3.5}
\end{equation*}
$$

provided that $C_{k}$ is sufficiently large. Moreover, for all $v \in Z_{i}$ we have

$$
\begin{equation*}
Q(v)=\left\{w \in V:\left\|\widehat{A}_{v}-\widehat{A}_{w}\right\|^{2} \leq 0.01 \Gamma_{j}\right\} \supset Z_{i} \tag{3.6}
\end{equation*}
$$

because we are assuming that $\Gamma_{j} \geq \Gamma / 2$. In addition, let $w \in Q_{l}$ for some $l<L$; since our choice of $i$ ensures that $v \in V_{i} \neq V_{\pi(l)}$, we have

$$
\begin{equation*}
\sqrt{\Gamma} \leq\left\|\mathbb{E}^{V_{\pi(l)}}-\mathbb{E}_{v}\right\| \leq\left\|\mathbb{E}_{v}-\widehat{A}_{v}\right\|+\left\|\widehat{A}_{w}-\widehat{A}_{v}\right\|+\left\|\xi_{l}-\widehat{A}_{w}\right\|+\left\|\xi_{l}-\mathbb{E}^{V_{\pi(l)}}\right\| \tag{3.7}
\end{equation*}
$$

Now, the construction in Step 3 of Partition ensures that $\left\|\widehat{A}_{w}-\xi_{l}\right\| \leq 0.1 \sqrt{\Gamma}$. Furthermore, $\left\|\xi_{l}-\mathbb{E}^{V_{\pi(l)}}\right\| \leq \sqrt{\Gamma} / 3$ by induction (cf. (3.2)), and $\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\| \leq$ $0.1 \sqrt{\Gamma}$, because $v \in Z_{i}$. Hence, (3.7) entails that $\left\|\widehat{A}_{w}-\widehat{A}_{v}\right\|>0.1 \sqrt{\Gamma}$, so that $w \notin Q(v)$. Consequently, (3.6) yields

$$
\begin{equation*}
Z_{i} \cap Q_{l}=\emptyset \text { for all } l<L \tag{3.8}
\end{equation*}
$$

Finally, let $v_{L}$ signify the element chosen by Step 3 of Partition to construct $Q_{L}$. Then by construction $\left|Q_{L}\right|=\left|Q\left(v_{L}\right) \backslash \bigcup_{l=1}^{L-1} Q_{l}\right| \geq\left|Q(v) \backslash \bigcup_{l=1}^{L-1} Q_{l}\right|$. Therefore,

$$
\left|Q_{L}\right| \geq\left|Q(v) \backslash \bigcup_{l=1}^{L-1} Q_{l}\right| \stackrel{(3.6),(3.8)}{\geq}\left|Z_{i}\right| \stackrel{(3.5)}{\geq}\left|V_{i}\right|-0.01 n_{\mathrm{min}}
$$

As this estimate holds for all $i \notin \pi(\{1, \ldots, L-1\}),(3.3)$ follows.
Thus, we know that $Q_{L}$ is "big". As a next step, we prove (3.4), i.e., we show that $Q_{L}$ "mainly" consists of vertices in $V_{\pi(L)}$. To this end, let $1 \leq i \leq k$ be such that $\left\|\mathbb{E}^{V_{i}}-\widehat{A}_{v_{L}}\right\|$ is minimum. Let $Y=Q_{L} \backslash V_{i}$. Then for all $w \in Y$ we have $\left\|\mathbb{E}_{w}-\widehat{A}_{v_{L}}\right\| \geq\left\|\mathbb{E}^{V_{i}}-\widehat{A}_{v}\right\|$. Further, since $\sqrt{\Gamma} \leq\left\|\mathbb{E}_{w}-\mathbb{E}^{V_{i}}\right\| \leq \| \mathbb{E}_{w}-$ $\widehat{A}_{v_{L}}\|+\| \mathbb{E}^{V_{i}}-\widehat{A}_{v_{L}}\|\leq 2\| \mathbb{E}_{w}-\widehat{A}_{v_{L}} \|$, we conclude that $\left\|\mathbb{E}_{w}-\widehat{A}_{v_{L}}\right\|^{2} \geq \frac{1}{4} \Gamma$. On the other hand, as $w \in Q_{L}$, we have $\left\|\widehat{A}_{w}-\widehat{A}_{v_{L}}\right\|^{2} \leq 0.01 \Gamma$. Therefore, we obtain $\left\|\widehat{A}_{w}-\mathbb{E}_{w}\right\|^{2} \geq 0.1 \Gamma$ for all $w \in Y$, so that

$$
\begin{equation*}
0.1|Y| \Gamma \leq \sum_{w \in Y}\left\|\widehat{A}_{w}-\mathbb{E}_{w}\right\|^{2} \leq\|\widehat{A}-\mathbb{E}\|_{F}^{2} \stackrel{\text { Lemma }}{\leq}{ }^{3.4} c_{k} \lambda^{2} \tag{3.9}
\end{equation*}
$$

Hence, due to our assumption (3.1) on $\Gamma$, (3.9) yields that $|Y|<0.01 n_{\text {min }}$. Consequently, (3.3) entails that $\left|V_{i} \cap Q_{L}\right| \geq 0.99\left|Q_{L}\right|$, so that $i=\pi(L)$. Hence,
we obtain $\left|Q_{L} \cap V_{\pi(L)}\right|=\left|Q_{L} \cap V_{i}\right|=\left|Q_{L} \backslash Y\right| \geq\left|Q_{L}\right|-0.01 n_{\min }$, thereby establishing (3.4).

Finally, to show (3.2), we note that by construction $\left\|\xi_{L}-\widehat{A}_{v_{L}}\right\|^{2} \leq 0.01 \Gamma$ and $\left\|\widehat{A}_{w}-\widehat{A}_{v_{L}}\right\|^{2} \leq 0.01 \Gamma$ for all $w \in Q_{L} \cap V_{\pi(L)}$ (cf. Step 3 of Partition). Therefore,

$$
\begin{align*}
\left|Q_{L} \cap V_{\pi(L)}\right| \cdot & \left\|\mathbb{E}_{\pi(L)}-\xi_{L}\right\|^{2} \\
\quad \leq & 3 \sum_{w \in Q_{L} \cap V_{\pi(L)}}\left\|\xi_{L}-\widehat{A}_{v_{L}}\right\|^{2}+\left\|\widehat{A}_{w}-\widehat{A}_{v_{L}}\right\|^{2}+\left\|\widehat{A}_{w}-\widehat{\mathbb{E}}_{\pi(L)}\right\|^{2} \\
\quad \leq & 0.06 \Gamma\left|Q_{L} \cap V_{\pi(L)}\right|+3\|\widehat{A}-\mathbb{E}\|_{F}^{2} \\
\quad \leq & 0.06 \Gamma\left|Q_{L} \cap V_{\pi(L)}\right|+3 c_{k} \lambda^{2} .
\end{align*}
$$

Since $\left|Q_{L} \cap V_{\pi(L)}\right| \geq 0.9 n_{\min }$ due to (3.3) and (3.4), (3.10) entails that $\| \mathbb{E}_{\pi(L)}-$ $\xi_{L} \|^{2} \leq 0.07 \Gamma+\frac{6 c_{k} \lambda^{2}}{n_{\min }} \leq 0.1 \Gamma$. Thus, (3.2) follows.

In the sequel, we shall assume without loss of generality that the map $\pi$ from Lemma 3.7 is just the identity, i.e., $\pi(i)=i$ for all $i$. Bootstrapping on the estimate $\left\|\xi_{i}-\mathbb{E}^{V_{i}}\right\|^{2} \leq 0.1 \Gamma$ for $1 \leq i \leq k$ from Lemma 3.7, we derive the following stronger estimate.
Corollary 3.8. For all $1 \leq i \leq k$ we have $\left\|\xi_{i}-\mathbb{E}^{V_{i}}\right\|^{2} \leq 100\left|Q_{i}\right|^{-1} \sum_{v \in Q_{i}} \| \widehat{A}_{v}-$ $\mathbb{E}_{v} \|^{2}$.
Proof. By the Cauchy-Schwarz inequality,

$$
\begin{equation*}
\left\|\xi_{i}-\mathbb{E}^{V_{i}}\right\|=\left|Q_{i}\right|^{-1}\left\|\sum_{v \in Q_{i}} \widehat{A}_{v}-\mathbb{E}^{V_{i}}\right\| \leq\left|Q_{i}\right|^{-1 / 2}\left[\sum_{v \in Q_{i}}\left\|\widehat{A}_{v}-\mathbb{E}^{V_{i}}\right\|^{2}\right]^{1 / 2} . \tag{3.11}
\end{equation*}
$$

Furthermore, as $\left\|\xi_{i}-\mathbb{E}^{V_{i}}\right\|^{2} \leq 0.1 \Gamma$ by Lemma 3.7, for all $v \in Q_{i} \backslash V_{i}$ we have

$$
\begin{equation*}
\left\|\widehat{A}_{v}-\mathbb{E}^{V_{i}}\right\|^{2} \leq 2\left(\left\|\widehat{A}_{v}-\xi_{i}\right\|^{2}+\left\|\xi_{i}-\mathbb{E}^{V_{i}}\right\|^{2}\right) \leq \Gamma^{1 / 3}, \tag{3.12}
\end{equation*}
$$

because the construction of $Q_{i}$ in Step 3 of Partition ensures that $\left\|\widehat{A}_{v}-\xi_{i}\right\|^{2} \leq$ $0.01 \Gamma$. Hence, as $\left\|\mathbb{E}_{v}-\mathbb{E}^{V_{i}}\right\|^{2} \geq \Gamma$, (3.12) implies that $\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\| \geq 0.1\left\|\widehat{A}_{v}-\mathbb{E}^{V_{i}}\right\|$. Therefore, the assertion follows from (3.11).
Corollary 3.9. For all $v \in S_{i} \backslash V_{i}$ we have $\left\|\widehat{A}_{v}-\xi_{i}\right\| \leq 3\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|$.
Proof. Let $i \neq l$ and consider a $v \in S_{i} \cap V_{l}$. We shall establish below that

$$
\begin{equation*}
\left\|\widehat{A}_{v}-\xi_{i}\right\| \leq\left\|\widehat{A}_{v}-\xi_{l}\right\| . \tag{3.13}
\end{equation*}
$$

Then by Lemma 3.7 $\left\|\widehat{A}_{v}-\xi_{i}\right\| \leq\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|+\left\|\mathbb{E}_{v}-\xi_{l}\right\| \leq\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|+\sqrt{\Gamma} / 3$, and thus $\sqrt{\Gamma} \leq\left\|\mathbb{E}_{v}-\mathbb{E}^{V_{i}}\right\| \leq\left\|\widehat{A}_{v}-\xi_{i}\right\|+\left\|\xi_{i}-\mathbb{E}^{V_{i}}\right\|+\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\| \leq 2\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|+\frac{2}{3} \sqrt{\Gamma}$.

Consequently, we obtain $\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\| \geq \frac{1}{6} \sqrt{\Gamma}$, so that the assertion follows from the estimate

$$
\left\|\widehat{A}_{v}-\xi_{i}\right\| \stackrel{(3.13)}{\leq} \quad\left\|\widehat{A}_{v}-\xi_{l}\right\| \leq\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|+\left\|\mathbb{E}_{v}-\xi_{l}\right\| .
$$

Finally, we prove (3.13). If $v \in S_{i} \cap V_{l} \backslash Q_{i}$, then the construction of $S_{i}$ in Step 4 of Partition guarantees that $\left\|\widehat{A}_{v}-\xi_{i}\right\| \leq\left\|\widehat{A}_{v}-\xi_{l}\right\|$, as claimed. Thus, assume that $v \in Q_{i} \cap V_{l}$. Then

$$
\begin{aligned}
\left\|\widehat{A}_{v}-\xi_{i}\right\| \leq 0.15 \sqrt{\Gamma} & \text { [by the definition of } Q_{i} \text { in Step 3] } \\
\max \left\{\left\|\xi_{i}-\mathbb{E}^{V_{i}}\right\|,\left\|\xi_{l}-\mathbb{E}_{v}\right\|\right\} \leq \frac{1}{3} \sqrt{\Gamma} & \text { [by Lemma 3.7], } \\
\left\|\mathbb{E}^{V_{i}}-\mathbb{E}_{v}\right\| \geq \sqrt{\Gamma} &
\end{aligned}
$$

Therefore, if $\left\|\widehat{A}_{v}-\xi_{l}\right\|<\left\|\widehat{A}_{v}-\xi_{i}\right\|$, then we would arrive at the contradiction

$$
\begin{aligned}
\sqrt{\Gamma} & \leq\left\|\mathbb{E}^{V_{i}}-\mathbb{E}_{v}\right\| \leq\left\|\mathbb{E}^{V_{i}}-\xi_{i}\right\|+\left\|\mathbb{E}_{v}-\xi_{l}\right\|+\left\|\xi_{i}-\xi_{l}\right\| \\
& \leq \frac{2}{3} \sqrt{\Gamma}+\left\|\widehat{A}_{v}-\xi_{i}\right\|+\left\|\widehat{A}_{v}-\xi_{l}\right\|<\frac{2}{3} \sqrt{\Gamma}+2\left\|\widehat{A}_{v}-\xi_{i}\right\| \leq 0.99 \sqrt{\Gamma}
\end{aligned}
$$

Thus, we conclude that $\left\|\widehat{A}_{v}-\xi_{l}\right\| \geq\left\|\widehat{A}_{v}-\xi_{i}\right\|$, thereby completing the proof.
Proof of Lemma 3.5. Since $\left|Q_{i}\right| \geq \frac{1}{2}\left|V_{i}\right|$ by Lemma 3.7, we have the estimate

$$
\begin{align*}
\sum_{i=1}^{k} \sum_{w \in S_{i} \cap V_{i}}\left\|\widehat{A}_{w}-\xi_{i}\right\|^{2} & \leq 2 \sum_{i=1}^{k} \sum_{w \in S_{i} \cap V_{i}}\left[\left\|\widehat{A}_{w}-\mathbb{E}_{w}\right\|^{2}+\left\|\mathbb{E}_{w}-\xi_{i}\right\|^{2}\right] \\
& \stackrel{\text { Cor. } 3.8}{\leq} 2\|\widehat{A}-\mathbb{E}\|_{F}^{2}+200 \sum_{i=1}^{k} \frac{\left|S_{i} \cap V_{i}\right|}{\left|Q_{i}\right|} \sum_{v \in Q_{i}}\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|^{2} \\
& \leq 500\|\widehat{A}-\mathbb{E}\|_{F}^{2} \tag{3.14}
\end{align*}
$$

Furthermore, by Corollary 3.9

$$
\begin{equation*}
\sum_{i=1}^{k} \sum_{v \in S_{i} \backslash V_{i}}\left\|\widehat{A}_{v}-\xi_{i}\right\|^{2} \leq 9 \sum_{i=1}^{k} \sum_{v \in S_{i} \backslash V_{i}}\left\|\widehat{A}_{v}-\mathbb{E}_{v}\right\|^{2} \leq 9\|\widehat{A}-\mathbb{E}\|_{F}^{2} \tag{3.15}
\end{equation*}
$$

Since $\|\widehat{A}-\mathbb{E}\|_{F}^{2} \leq C k \lambda^{2}$ by Lemma 3.4, the bounds (3.14) and (3.15) imply the assertion.

### 3.4. Proof of Lemma 3.6

Set $S_{a b}=S_{a} \cap V_{b}$ for $1 \leq a, b \leq k$. Moreover, for each $1 \leq a \leq k$ let $1 \leq \pi(a) \leq k$ be such that $\left\|\mathbb{E}^{V_{\pi(a)}}-\xi_{a}\right\|$ is minimum. Then for all $b \neq \pi(a)$ we have

$$
\begin{equation*}
\sqrt{\Gamma} \leq\left\|\mathbb{E}^{V_{\pi(a)}}-\mathbb{E}^{V_{b}}\right\| \leq\left\|\mathbb{E}^{V_{\pi(a)}}-\xi_{a}\right\|+\left\|\mathbb{E}^{V_{b}}-\xi_{a}\right\| \leq 2\left\|\mathbb{E}^{V_{b}}-\xi_{a}\right\| \tag{3.16}
\end{equation*}
$$

so that $\left\|\mathbb{E}^{V_{b}}-\xi_{a}\right\| \geq \sqrt{\Gamma} / 2$. Therefore, by our assumption that $\sum_{i=1}^{k} \sum_{v \in S_{i}} \| \xi_{i}-$ $\widehat{A}_{v} \|^{2} \leq C_{0} k^{3} \lambda^{2}$, we have

$$
\begin{align*}
\frac{\Gamma}{4} \sum_{a=1}^{k} \sum_{1 \leq b \leq k: b \neq \pi(a)}\left|S_{a b}\right| & \leq \sum_{a, b=1}^{k}\left|S_{a b}\right| \cdot\left\|\mathbb{E}^{V_{b}}-\xi_{a}\right\|^{2} \\
& \leq 2 \sum_{a, b=1}^{k} \sum_{v \in S_{a b}}\left\|\mathbb{E}_{v}-\widehat{A}_{v}\right\|^{2}+\left\|\widehat{A}_{v}-\xi_{a}\right\|^{2} \\
& \leq \quad 2\|\widehat{A}-\mathbb{E}\|_{F}^{2}+2 \sum_{a, b=1}^{k} \sum_{v \in S_{a b}}\left\|\widehat{A}_{v}-\xi_{a}\right\|^{2} \\
& \begin{array}{c}
\text { Lemma } 3.4 \\
\leq
\end{array} \tag{3.17}
\end{align*} C_{0} k^{3} \lambda^{2}+2 C_{0} k^{3} \lambda^{2} \leq C_{0}^{2} k^{3} \lambda^{2} .
$$

Hence,

$$
\begin{equation*}
\sum_{a=1}^{k}\left|S_{a} \triangle V_{\pi(a)}\right|=\sum_{1 \leq a, b \leq k: b \neq \pi(a)} 2\left|S_{a b}\right| \leq \frac{8 c_{0}^{2} k^{3} \lambda^{2}}{\Gamma} \leq 0.001 n_{\min } \tag{3.18}
\end{equation*}
$$

provided that $C_{k}$ is sufficiently large (cf. (3.1)). Combining (3.17) and (3.18), we obtain $\frac{n_{\text {min }}}{2}\left\|\mathbb{E}^{V_{\pi(a)}}-\xi_{a}\right\|^{2} \leq\left|S_{a} \cap V_{\pi(a)}\right| \cdot\left\|\mathbb{E}_{\pi(a)}-\xi_{a}\right\|^{2} \leq c_{0}^{2} k^{3} \lambda^{2}$, whence

$$
\begin{equation*}
\left\|\mathbb{E}_{\pi(a)}-\xi_{a}\right\|^{2} \leq \frac{2 c_{0}^{2} k^{3} \lambda^{2}}{n_{\min }} \leq 0.001 \Gamma \quad \text { for all } 1 \leq a \leq k \tag{3.19}
\end{equation*}
$$

provided that $C_{k}$ is large enough. Thus, we have established the first two parts of the lemma. In addition, observe that (3.18) implies that $\pi$ is bijective (because the sets $S_{1}, \ldots, S_{k}$ are pairwise disjoint and $\left|V_{a}\right| \geq n_{\text {min }}$ for all $1 \leq a \leq k$ ). Finally, the third assertion follows from the estimate

$$
\begin{aligned}
\sum_{a, b=1}^{k}\left|S_{a b}\right| \cdot\left\|\mathbb{E}^{V_{\pi(a)}}-\mathbb{E}^{V_{\pi(b)}}\right\|^{2} & \leq 2 \sum_{a, b=1}^{k}\left|S_{a b}\right|\left(\left\|\mathbb{E}^{V_{\pi(a)}}-\xi_{a}\right\|^{2}+\left\|\mathbb{E}^{V_{\pi(b)}}-\xi_{a}\right\|^{2}\right) \\
& \stackrel{(3.16)}{\leq} 8 \sum_{a, b=1}^{k}\left|S_{a b}\right| \cdot\left\|\mathbb{E}^{V_{\pi(b)}}-\xi_{a}\right\|^{2} \\
& \stackrel{(3.17)}{\leq} 8 C_{0}^{2} k^{3} \lambda^{2} \leq 0.001 \Gamma n_{\min }
\end{aligned}
$$

where we assume once more that $C_{k}$ is sufficiently large.

## 4. Experiments

We illustrate the effectiveness of spectral techniques using simulations. In particular, we explore the case when we have a mixture of two populations; we show that when $N K>1 / \gamma^{2}$ and $K>1 / \gamma$, either the first or the second left


Fig 1. Plots show success rate as a function of $N$ for several values of $K$, when $\gamma=$ $(0.04)^{2}$. Each point is an average over 100 trials. Horizontal lines ("oracles") indicate the information-theoretically best possible success rate for that value of $K$ (how well one could do if one knew in advance which features satisfied $p_{1}^{i}>p_{2}^{i}$ and which satisfied $p_{1}^{i}<p_{2}^{i}$; they are not exactly horizontal because they are also an average over 100 runs). Vertical bars indicate the value of $N$ for which $N K=1 / \gamma^{2}$.
singular vector of $X$ shows an approximately correct partitioning, meaning that the success rate is well above $1 / 2$. The entry-wise expected value matrix $\mathcal{X}$ is: among $K / 2$ features, $p_{1}^{i}>p_{2}^{i}$ and for the other half, $p_{1}^{i}<p_{2}^{i}$ such that $\forall i$, $p_{1}^{i}, p_{2}^{i} \in\left\{\frac{1+\alpha}{2}+\frac{\epsilon}{2}, \frac{1-\alpha}{2}+\frac{\epsilon}{2}\right\}$, where $\epsilon=0.1 \alpha$. Hence $\gamma=\alpha^{2}$. We report results on balanced cases only, but we do observe that unbalanced cases show similar tradeoffs. For each population $P$, the success rate is defined as the number of individuals that are correctly classified, i.e., they belong to a group that $P$ is the majority of that group, versus the size of the population $|P|$.

Each point on the SVD curve corresponds to an average rate over 100 trials. Since we are interested in exploring the tradeoffs of $N, K$ in all ranges (e.g., when $N \ll K$ or $N \gg K$ ), rather than using the threshold $T$ in Procedure Classify that is chosen in case both $N, K>1 / \gamma$, to decide which singular vector to use, we try both $u_{1}$ and $u_{2}$ and use the more effective one to measure the success rate at each trial. For each data point, the distribution of $X$ is fixed across all trials and we generate an independent $X_{2 N \times K}$ for each trial to measure success rate based on the more effective classifier between $u_{1}$ and $u_{2}$.

One can see from the plot that when $K<1 / \gamma$, i.e., when $K=200$ and 400, no matter how much we increase $N$, the success rate is consistently low. Note that $50 / 100$ of success rate is equivalent to a total failure. In contrast, when $N$ is
smaller than $1 / \gamma$, as we increase $K$, we can always classify with a high success rate, where in general, $N K>1 / \gamma^{2}$ is indeed necessary to see a high success rate. In particular, the curves for $K=5000,2500,1250$ show the sharpness of the threshold behavior for increasing sample size $n$ from below $1 / K \gamma^{2}$ to above. For each curve, we also compute the best possible classification one could hope to make if one knew in advance which features satisfied $p_{1}^{i}>p_{2}^{i}$ and which satisfied $p_{1}^{i}<p_{2}^{i}$. These are the horizontal(ish) dotted lines above each curve. The fact that the solid curves are approaching these information-theoretic upper bounds shows that the spectral technique is correctly using the available information.

## Acknowledgements

We thank John Lafferty, Frank McSherry, Roman Vershynin and Larry Wasserman for many helpful discussions on this work.

## Appendix A: More Proofs for the simple algorithm classify

## A.1. Proof of Lemma 2.1

Let $u_{1}, \ldots, u_{n}, v_{1}, \ldots, v_{n}$ be the $n$ left and right singular vectors of $X$, corresponding to $s_{1}(X) \geq s_{2}(X) \geq \cdots \geq s_{n}(X)$, we have for $\forall i,\left\|u_{i}\right\|_{2}=1,\left\|v_{i}\right\|_{2}=1$ such that $X^{T} u_{i}=s_{i}(X) v_{i}$ and $X v_{i}=s_{i}(X) u_{i}$.

Before we prove Lemma 2.1, given an $n \times K$ matrix $X$, where $n<K$, let us first define $H=X X^{T}$ and a block matrix

$$
Y=\left[\begin{array}{cc}
0 & X  \tag{A.1}\\
X^{T} & 0
\end{array}\right]_{(2 N+K) \times(2 N+K)} .
$$

Recall that singular values of a real $n \times K$ matrix $X$ are exactly the nonnegative square roots of the $n$ largest eigenvalues of $H=X X^{T}$, i.e, $s_{i}(X)=$ $\sqrt{\lambda_{i}(H)}, \forall i=1, \ldots, n$, given that

$$
\begin{equation*}
H u_{i}=X X^{T} u_{i}=s_{i}(X) X v_{i}=s_{i}^{2}(X) u_{i} \tag{A.2}
\end{equation*}
$$

Hence the left singular vectors $u_{1}, \ldots, u_{n}$ of $X$ are eigenvectors of $H$ corresponding to $\lambda_{i}(H)=s_{i}^{2}(X)$.

We next show that the first $n$ eigenvalues of $Y$ and their corresponding eigenvectors:

$$
Y\left[\begin{array}{c}
u_{i}  \tag{A.3}\\
v_{i}
\end{array}\right]=\left[\begin{array}{cc}
0 & X \\
X^{T} & 0
\end{array}\right]\left[\begin{array}{c}
u_{i} \\
v_{i}
\end{array}\right]=\left[\begin{array}{c}
X v_{i} \\
X^{T} u_{i}
\end{array}\right]=\left[\begin{array}{c}
s_{i}(X) u_{i} \\
s_{i}(X) v_{i}
\end{array}\right]=s_{i}(X)\left[\begin{array}{c}
u_{i} \\
v_{i}
\end{array}\right]
$$

and hence
Proposition A.1. The largest $n$ eigenvalues of $Y$ are $s_{1}(X), \ldots, s_{n}(X)$ with corresponding eigenvectors $\left[u_{i}, v_{i}\right], \forall i=1, \ldots, n$, where $u_{i}, v_{i}, \forall i$, are left and right singular vectors of $X$ corresponding to $s_{i}(X)$.

In fact both $\pm s_{i}(X)$ are eigenvalues of $Y$, which is irrelevant.
Proof of Lemma 2.1. We first state a theorem, whose statement appears in a lecture note by Spielman (2002), with a slight modification (off by a factor on RHS). Our proof for this theorem is included here for completeness. It is known that for any real symmetric matrix, there exist a set of $n$ orthonormal eigenvectors.

Theorem A. 2 (Modified Version of Spielman (2002)). For A and M being two symmetric matrices and $E=M-A$. Let $\lambda_{1}(A) \geq \lambda_{2}(A) \geq \cdots \geq$ $\lambda_{n}(A)$ be eigenvalues of $A$, with orthonormal eigenvectors $v_{1}, v_{2}, \ldots, v_{n}$ and let $\lambda_{1}(M) \geq \lambda_{2}(M) \geq \cdots \geq \lambda_{n}(M)$ be eigenvalues of $M$ and $w_{1}, w_{2}, \ldots, w_{n}$ be the corresponding orthonormal eigenvectors of $M$, with $\theta_{i}=\angle\left(v_{i}, w_{i}\right)$. Then

$$
\begin{equation*}
\theta_{i} \sim \sin \left(\theta_{i}\right) \leq \frac{\left\|E-\Delta_{i} I\right\|_{2}}{\operatorname{gap}(i, A)} \leq \frac{\|E\|_{2}+\left|\Delta_{i}\right|}{\operatorname{gap}(i, A)} \leq \frac{2\|E\|_{2}}{\operatorname{gap}(i, A)} \tag{A.4}
\end{equation*}
$$

where $\operatorname{gap}(i, A)=\min _{j \neq i}\left|\lambda_{i}(A)-\lambda_{j}(A)\right|$ and $\Delta_{i}=\lambda_{i}(M)-\lambda_{i}(A)$.
Let us apply Theorem A. 2 to the symmetric matrix $Y$ in (A.1). In particular, we only compare the first $n$ eigenvectors of $Y$ of $\mathcal{Y}$. For the numerator of RHS of (A.4), we have $E=Y-\mathcal{Y}$, and $\|E\|_{2}=\|Y-\mathcal{Y}\|_{2}=s_{1}(Y-\mathcal{Y})$ by a derivation similar to (A.3), where eigenvectors of $E$ are concatenations of left and right singular vectors of $X-\mathcal{X}$; For the denominator, we have by Proposition A.1, $\operatorname{gap}(i, \mathcal{Y})=\min _{j \neq i}\left|\lambda_{i}(\mathcal{Y})-\lambda_{j}(\mathcal{Y})\right|=\min _{j \neq i}\left|s_{i}(\mathcal{X})-s_{j}(\mathcal{X})\right|$.

We first prove the following claim.
Claim A.3. For any symmetric $n \times n$ matrix $A$, let $\lambda_{i}, \forall i=1, \ldots, n$ be eigenvalues of $A$ with orthonormal eigenvectors $v_{1}, v_{2}, \ldots, v_{n}$, for all $y \perp v_{i}$,

$$
\left\|\left(A-\lambda_{i}\right) y\right\|_{2} \geq \min _{j \neq i}\left|\lambda_{i}-\lambda_{j}\right|\|y\|_{2} .
$$

Proof. Let us first assume $y \perp v_{i}$ and write $y=\sum_{j=1, j \neq i}^{n} c_{j} v_{j}$, thus we have $\|y\|_{2}=\sqrt{\sum_{j=1, j \neq i}^{n} c_{j}^{2}}$ and

$$
\begin{aligned}
\left\|\left(A-\lambda_{i}\right) y\right\|_{2} & =\left\|\sum_{j=1, j \neq i}^{n} c_{j}\left(A-\lambda_{i}\right) v_{j}\right\|_{2} \\
& =\left\|\sum_{j=1, j \neq i}^{n} c_{j}\left(\lambda_{j}-\lambda_{i}\right) v_{j}\right\|_{2}=\sqrt{\sum_{j=1, j \neq i}^{n} c_{j}^{2}\left|\lambda_{j}-\lambda_{i}\right|^{2}} \\
& \geq \min _{j \neq i}\left|\lambda_{i}-\lambda_{j}\right| \sqrt{\sum_{j=1, j \neq i}^{n} c_{j}^{2}}=\min _{j \neq i}\left|\lambda_{i}-\lambda_{j}\right|\|y\|_{2} .
\end{aligned}
$$

Proof of Theorem A.2. Let us construct a vector $y$ that is orthogonal to $v_{i}$ as follows:

$$
y=w_{i}-\left(v_{i}^{T} w_{i}\right) v_{i}
$$

By Claim A.3, we have

$$
\left\|\left(A-\lambda_{i}(A)\right) y\right\|_{2} \geq \min _{j \neq i}\left|\lambda_{i}(A)-\lambda_{j}(A)\right|\|y\|_{2}
$$

and hence

$$
\|y\|_{2} \leq \frac{\left\|\left(A-\lambda_{i}(A)\right) y\right\|_{2}}{\min _{j \neq i}\left|\lambda_{i}(A)-\lambda_{j}(A)\right|}
$$

On the other hand,

$$
\begin{aligned}
\left\|\left(A-\lambda_{i}(A)\right) y\right\|_{2} & =\left\|\left(A-\lambda_{i}(A)\right)\left(w_{i}-\left(v_{i}^{T} w_{i}\right) v_{i}\right)\right\|_{2} \\
& =\left\|\left(A-\lambda_{i}(A)\right) w_{i}\right\|_{2} \\
& =\left\|\left(M-E-\lambda_{i}(A)\right) w_{i}\right\|_{2} \\
& =\left\|\left(\lambda_{i}(M)-\lambda_{i}(A)\right) w_{i}-E w_{i}\right\|_{2} \\
& =\left\|\left(\Delta_{i} I-E\right) w_{i}\right\|_{2} \leq\left\|E-\Delta_{i} I\right\|_{2} \\
& \leq\|E\|_{2}+\left|\Delta_{i}\right|
\end{aligned}
$$

Finally, given that $\|w\|_{2}=1$,

$$
\begin{aligned}
\sin \left(\theta_{i}\right) & =\frac{\|y\|_{2}}{\|w\|_{2}} \leq \frac{\left\|\left(A-\lambda_{i}(A)\right) y\right\|_{2}}{\min _{j \neq i}\left|\lambda_{i}(A)-\lambda_{j}(A)\right|} \\
& \leq \frac{\|E\|_{2}+\left|\Delta_{i}\right|}{\operatorname{gap}(i, A)}
\end{aligned}
$$

Lemma A.4. $\forall i=1, \ldots, n,\left|\Delta_{i}\right| \leq\|E\|_{2}$.
Proof. Let $S_{j}$ be a subspace of dimension $j$. Recall the following definition of $\lambda_{i}$ for a matrix:

$$
\begin{equation*}
\lambda_{i}(M)=\inf _{S_{N-i+1}} \sup _{x \in S_{N-i+1},\|x\|_{2}=1} x^{T} M x \tag{A.5}
\end{equation*}
$$

In the following, let $S_{N-i+1}^{v}$ be the subspace that is orthogonal to the subset of orthonormal eigenvectors $v_{1}, \ldots, v_{i-1}$ of symmetric matrix $A$. Note that this is the $N-i+1$ dimensional subspace that achieves the minimum of the maximum of $v^{T} A v$ over all unit-length vectors $v$ in the particular subspace. We have

$$
\begin{aligned}
\lambda_{i}(M) & =\inf _{S_{N-i+1}} \sup _{x \in S_{N-i+1},\|x\|_{2}=1} x^{T} M x \leq \sup _{x \in S_{N-i+1}^{v},\|x\|_{2}=1} x^{T} M x \\
& \leq \sup _{x \in S_{N-i+1}^{v},\|x\|_{2}=1} x^{T}(A+E) x \\
& \leq \sup _{v \in S_{N-i+1}^{v},\|v\|_{2}=1} v^{T} A v+\sup _{x \in R^{n},\|x\|_{2}=1}\left|x^{T} E x\right| \\
& =\lambda_{i}(A)+\|E\|_{2}
\end{aligned}
$$

For the other direction, let $S_{N-i+1}^{w}$ be the subspace that is orthogonal to the subset of orthonormal eigenvectors $w_{1}, \ldots, w_{i-1}$ of symmetric matrix $M$. Note that this is the $N-i+1$ dimensional subspace that achieves the minimum of the maximum of $w^{T} M w$ over all unit-length vectors $w$ in the particular subspace. We have

$$
\begin{aligned}
\lambda_{i}(A) & =\inf _{S_{N-i+1}} \sup _{x \in S_{N-i+1},\|x\|_{2}=1} x^{T} A x \leq \sup _{x \in S_{N-i+1}^{w},\|x\|_{2}=1} x^{T} A x \\
& \leq \sup _{x \in S_{N-i+1}^{w},\|x\|_{2}=1} x^{T}(M+(-E)) x \\
& \leq \sup _{w \in S_{N-i+1}^{w},\|w\|_{2}=1} w^{T} M w+\sup _{x \in R^{n},\|x\|_{2}=1} x^{T}(-E) x \\
& \leq \sup _{w \in S_{N-i+1}^{w},\|w\|_{2}=1} w^{T} M w+\sup _{x \in R^{n},\|x\|_{2}=1}\left|x^{T}(-E) x\right| \\
& =\lambda_{i}(M)+\|E\|_{2},
\end{aligned}
$$

where $\|E\|_{2}=\|-E\|_{2}$. Thus $-\|E\|_{2} \leq \lambda_{i}(M)-\lambda_{i}(A) \leq\|E\|_{2}$, and $\left|\Delta_{i}\right| \leq$ $\|E\|_{2}$.

Therefore, $\sin \left(\theta_{i}\right) \leq \frac{\|E\|_{2}+\left|\Delta_{i}\right|}{\operatorname{gap}(i, A)} \leq \frac{2\|E\|_{2}}{\operatorname{gap}(i, A)}$.

## A.2. Some Propositions regarding the static matrices

For static matrix $\mathcal{H}=\mathcal{X} \mathcal{X}^{T}$ and $\mathcal{Y}=\left[\begin{array}{cc}0 & \mathcal{X} \\ \mathcal{X}^{T} & 0\end{array}\right]$, we define

$$
\begin{aligned}
\operatorname{gap}(\mathcal{H}) & =\left|\lambda_{1}(\mathcal{H})-\lambda_{2}(\mathcal{H})\right|, \\
\operatorname{gap}(\mathcal{Y}) & =\left|\lambda_{1}(\mathcal{Y})-\lambda_{2}(\mathcal{Y})\right|=\frac{\operatorname{gap}(\mathcal{H})}{\lambda_{1}(\mathcal{Y})+\lambda_{2}(\mathcal{Y})},
\end{aligned}
$$

Proposition A.5. For static matrix $\mathcal{Y}$, let $\operatorname{gap}(\mathcal{Y})=\left|\lambda_{1}(\mathcal{Y})-\lambda_{2}(\mathcal{Y})\right|=\frac{\operatorname{gap}(\mathcal{H})}{\lambda_{1}(\mathcal{Y})+\lambda_{2}(\mathcal{Y})}$, we have

$$
\begin{aligned}
\sqrt{\max \left\{N_{1} a, N_{2} c\right\}} & \leq \lambda_{1}(\mathcal{Y}) \leq \sqrt{N_{1} a+N_{2} c}, \\
\sqrt{N_{1} a+N_{2} c} & \leq \lambda_{1}(\mathcal{Y})+\lambda_{2}(\mathcal{Y}) \leq \sqrt{2\left(N_{1} a+N_{2} c\right)}, \\
\sqrt{\frac{N_{1} N_{2}\left(a c-b^{2}\right)}{N_{1} a+N_{2} c}} & \leq \lambda_{2}(\mathcal{Y}) \leq \sqrt{\frac{2 N_{1} N_{2}\left(a c-b^{2}\right)}{N_{1} a+N_{2} c}}, \\
\operatorname{gap}(\mathcal{Y})=\Theta\left(\frac{\operatorname{gap}(\mathcal{H})}{\sqrt{N_{1} a+N_{2} c}}\right) & =\Theta\left(\frac{\sqrt{\left(N_{1} a+N_{2} c\right)^{2}-4 N_{1} N_{2}\left(a c-b^{2}\right)}}{\sqrt{N_{1} a+N_{2} c}}\right) .
\end{aligned}
$$

Proof. We first show the following:

Proposition A.6. For static matrix $\mathcal{H}=\mathcal{X} \mathcal{X}^{T}$ as in (2.12), Let $\lambda_{1}(\mathcal{H}), \lambda_{2}(\mathcal{H})$ be the non-zero eigenvalues of $\mathcal{H}$, and denote $\operatorname{gap}(\mathcal{H})=\left|\lambda_{1}(\mathcal{H})-\lambda_{2}(\mathcal{H})\right|$.

$$
\begin{align*}
\lambda_{1}(\mathcal{H}) & =\frac{N_{1} a+N_{2} c+\sqrt{\left(N_{1} a-N_{2} c\right)^{2}+4 N_{1} N_{2} b^{2}}}{2}  \tag{A.6}\\
\lambda_{2}(\mathcal{H}) & =\frac{N_{1} a+N_{2} c-\sqrt{\left(N_{1} a-N_{2} c\right)^{2}+4 N_{1} N_{2} b^{2}}}{2}  \tag{A.7}\\
\left|N_{1} a-N_{2} c\right| & \leq \operatorname{gap}(\mathcal{H}) \leq N_{1} a+N_{2} c \tag{A.8}
\end{align*}
$$

where $\lambda_{2}(\mathcal{H})=0$, when $a c=b^{2}$ and $\operatorname{gap}(\mathcal{H})=N_{1} a+N_{2} c$.
Proof. Let $\mathcal{H}=\mathcal{X} \mathcal{X}^{T}$. The rank of $\mathcal{H}$ is at most 2. Therefore there exist at most two non-zero eigenvalues $\lambda_{1}, \lambda_{2}$ for $\mathcal{H}$, with corresponding nonzero eigenvectors $v_{1}, v_{2}$ being constant on each population. This is true because if we multiply $\mathcal{H}=\mathcal{X} \mathcal{X}^{T}$ by a permutation matrix $P$ to exchange two rows among the same population, we have $P \mathcal{H} v_{i}=\lambda_{i} P v_{i}, \forall i=1,2$; given that $P \mathcal{H} v_{i}=\mathcal{H} v_{i}$, we deduce that $P v_{i}=v_{i}$ for non-zero $\lambda_{i}$. Hence $v_{i}$ must be constant on each population.

Let the top two eigenvector $v_{1}, v_{2}$ be of form $[x, \ldots, x, y, \ldots, y]$, where $x$ repeats $N_{1}$ times and $y$ repeats $N_{2}$ times; Note that they corresponds to $\bar{u}_{1}$ and $\bar{u}_{2}$ of $\mathcal{X}$ following a derivation similar to (A.2).

We thus have the following equations:

$$
\begin{align*}
N_{1} a x+N_{2} b y & =\lambda x  \tag{A.9}\\
N_{1} b x+N_{2} c y & =\lambda y \tag{A.10}
\end{align*}
$$

which can be written in a matrix form:

$$
\left[\begin{array}{cc}
N_{1} a-\lambda & N_{2} b \\
N_{1} b & N_{2} c-\lambda
\end{array}\right]\left[\begin{array}{l}
x \\
y
\end{array}\right]=0
$$

Given that

$$
\left[\begin{array}{l}
x \\
y
\end{array}\right] \neq 0
$$

the matrix is not one-to-one and therefore

$$
D\left[\begin{array}{cc}
N_{1} a-\lambda & N_{2} b \\
N_{1} b & N_{2} c-\lambda
\end{array}\right]=0
$$

By solving $\left(N_{1} a-\lambda\right)\left(N_{2} c-\lambda\right)-N_{1} N_{2} b^{2}=0$, we get $\lambda_{1}(\mathcal{H}), \lambda_{2}(\mathcal{H})$ and $\operatorname{gap}(\mathcal{H})$. We next derive an upper bound on $\operatorname{gap}(\mathcal{H})$.

$$
\begin{align*}
\operatorname{gap}(\mathcal{H}) & =\sqrt{\left(N_{1} a-N_{2} c\right)^{2}+4 N_{1} N_{2} b^{2}}  \tag{A.11}\\
& =\sqrt{\left(N_{1} a+N_{2} c\right)^{2}-4 N_{1} N_{2} a c+4 N_{1} N_{2} b^{2}}  \tag{A.12}\\
& \leq \sqrt{\left(N_{1} a+N_{2} c\right)^{2}}  \tag{A.13}\\
& \leq N_{1} a+N_{2} c \tag{A.14}
\end{align*}
$$

where $a, c \geq 0$ and $a c \geq b^{2}$ as in Proposition 2.10. Hence $\operatorname{gap}(\mathcal{H}) \geq\left|N_{1} a-N_{2} c\right|$, given that $b^{2} \geq 0$.

Thus we have

$$
\begin{align*}
\max \left\{N_{1} a, N_{2} c\right\} & \leq \lambda_{1}(\mathcal{H}) \leq N_{1} a+N_{2} c  \tag{A.15}\\
0 & \leq \lambda_{2}(\mathcal{H}) \leq \min \left\{N_{1} a, N_{2} c\right\}  \tag{A.16}\\
\lambda_{1}(\mathcal{H})+\lambda_{2}(\mathcal{H}) & =N_{1} a+N_{2} c  \tag{A.17}\\
\lambda_{1}(\mathcal{H}) \lambda_{2}(\mathcal{H}) & =N_{1} N_{2}\left(a c-b^{2}\right) \tag{A.18}
\end{align*}
$$

Given two largest eigenvalues of $\mathcal{Y}, \lambda_{1}(\mathcal{Y})=\sqrt{\lambda_{1}(\mathcal{H})}$ and $\lambda_{2}(\mathcal{Y})=\sqrt{\lambda_{1}(\mathcal{H})}$ for $\mathcal{Y}=\left[\begin{array}{cc}0 & \mathcal{X} \\ \mathcal{X}^{T} & 0\end{array}\right]$, we get all inequalities, by Proposition A. 6 and the following: $\sqrt{\left(\lambda_{1}(\mathcal{Y})^{2}+\lambda_{2}(\mathcal{Y})^{2}\right)} \leq \lambda_{1}(\mathcal{Y})+\lambda_{2}(\mathcal{Y})=\sqrt{2\left(\lambda_{1}(\mathcal{Y})^{2}+\lambda_{2}(\mathcal{Y})^{2}\right)}$.

## A.3. Proofs of Proposition 2.5 and 2.6

Proof of Proposition 2.5. We rewrite Proposition A. 5 given that, for a normalized $\mathcal{X}, \operatorname{gap}(\mathcal{H}) \geq \frac{8 c_{0} N K}{5}$, as Proposition A. 8 and $\lambda_{j}(\mathcal{Y})=s_{j}(\mathcal{X})$. In particular,

$$
\begin{aligned}
\operatorname{gap}(\mathcal{X}) & =\operatorname{gap}(\mathcal{Y})=\frac{\operatorname{gap}(\mathcal{H})}{\lambda_{1}(\mathcal{Y})+\lambda_{2}(\mathcal{Y})} \\
& \geq \frac{\operatorname{gap}(\mathcal{H})}{\sqrt{N_{1} a+N_{2} c}} \geq \frac{8 c_{0} N K}{5 \sqrt{2 N K}} \\
& \geq \frac{4 \sqrt{2 N K}}{5}
\end{aligned}
$$

For the upper bound on $\operatorname{gap}(\mathcal{X})$, we have that

$$
\begin{aligned}
\operatorname{gap}(\mathcal{X}) & =\operatorname{gap}(\mathcal{Y})=\frac{\operatorname{gap}(\mathcal{H})}{\lambda_{1}(\mathcal{Y})+\lambda_{2}(\mathcal{Y})} \\
& \leq \frac{N_{1} a+N_{2} c}{\sqrt{N_{1} a+N_{2} c}} \leq \sqrt{2 N K}
\end{aligned}
$$

Definition A.7. For our application, we have $\forall k, 1 \geq p_{1}^{k}, p_{2}^{k} \geq 0$, and

$$
\mathcal{X}=\left[\begin{array}{cccc}
\frac{1+p_{1}^{1}}{2} & \frac{1+p_{1}^{2}}{2} & \ldots & \frac{1+p_{1}^{K}}{2} \\
\frac{1+p_{1}^{1}}{2} & \frac{1+p_{1}^{2}}{2} & \ldots & \frac{1+p_{1}^{K}}{2} \\
\frac{1+p_{2}^{1}}{2} & \frac{1+p_{2}^{2}}{2} & \ldots & \frac{1+p_{2}^{K}}{2} \\
\ldots & & & \\
\frac{1+p_{2}^{1}}{2} & \frac{1+p_{2}^{2}}{2} & \ldots & \frac{1+p_{2}^{K}}{2}
\end{array}\right]_{2 N \times K}
$$

It is easy to see that with this normalized random matrix, $\lambda_{1}(\mathcal{H})=\lambda_{2}(\mathcal{H})$ is not possible, given that $a, b, c \in[K / 4, K]$; furthermore, $\operatorname{gap}(\mathcal{H})=\Theta(N K)$ as in the Proposition A.8.

Proposition A.8. Given $\mathcal{H}=\mathcal{X} \mathcal{X}^{T}$ and $a, b, c$ as in (2.8) for any expected value mean matrix $\mathcal{X}$, which is not necessarily normalized,

$$
\operatorname{gap}(\mathcal{H})=\sqrt{\left(N_{1} a-N_{2} c\right)^{2}+4 N_{1} N_{2} b^{2}} \geq \frac{8 c_{0} N K}{5}
$$

where $c_{0}=\frac{|b| \sqrt{a c}}{K(a+c)}$.
Hence for a normalized $\mathcal{X}$, gap $(\mathcal{H})=\Theta(N K)$ given that $a, b, c \in[K / 4, K]$.
Proof. For a tighter lower bound of gap $(\mathcal{H})$ than the obvious $\left|N_{1} a-N_{2} c\right|$, let us assume w.l.o.g. that $N_{2} c \geq N_{1} a$. Thus we have

$$
\begin{equation*}
N_{2} \geq 2 N \frac{a}{a+c} \tag{A.19}
\end{equation*}
$$

We differentiate two cases:

- Balanced case: $N_{1} a \geq \frac{4}{25} N_{2} c$.
- Imbalanced case: $N_{1} a \leq \frac{4}{25} N_{2} c$.

For balanced case: we have $N_{1} \geq \frac{4}{25} \frac{N_{2} c}{a}$ and hence

$$
\begin{aligned}
\operatorname{gap}(\mathcal{H}) & \geq \sqrt{4 N_{1} N_{2} b^{2}} \geq \frac{4 N_{2}|b|}{5} \sqrt{\frac{c}{a}} \\
& \geq \frac{8 N|b|}{5} \frac{a}{a+c} \sqrt{\frac{c}{a}} \geq \frac{8 N|b|}{5} \frac{\sqrt{a c}}{a+c} \\
& =\frac{8 c_{0} N K}{5}
\end{aligned}
$$

where $N_{2} \geq 2 N \frac{a}{a+c}$ as in (A.19).
For the imbalanced case, given that $\sqrt{a c} \geq|b|$ by Proposition 2.10,

$$
\begin{aligned}
\operatorname{gap}(\mathcal{H}) & \geq \sqrt{\left(N_{1} a-N_{2} c\right)^{2}} \geq \frac{21}{25} N_{2} c \\
& \geq \frac{42}{25} \frac{N a c}{a+c} \geq \frac{8}{5} \frac{N|b| \sqrt{a c}}{a+c} \geq \frac{8 c_{0} N K}{5}
\end{aligned}
$$

Finally, for a normalized random matrix $X$ and its $\mathcal{X}$, we have $c_{0}$ being a constant and combing with the upper bound of $\operatorname{gap}(\mathcal{H}) \leq N_{1} a+N_{2} c \leq 2 N K$ concludes that $\operatorname{gap}(\mathcal{H})=\Theta(N K)$.
Proof of Proposition 2.6. By (A.2), $\bar{u}_{1}, \bar{u}_{2}$ are the first and second eigenvectors of $\mathcal{H}$ corresponding to $\lambda_{1}(\mathcal{H})$ and $\lambda_{2}(\mathcal{H})$. Let $x, y$ be entries that correspond to $P_{1}, P_{2}$ respectively in the first or second eigenvectors of $\mathcal{H}$. By (A.9) and (A.10), we have

$$
\frac{y}{x}=\frac{\lambda-N_{1} a}{N_{2} b}=\frac{N_{1} b}{\lambda-N_{2} c} .
$$

In addition, given any $b \neq 0$, we have $\operatorname{gap}(\mathcal{H})>\left|N_{1} a-N_{2} c\right|$ and hence $\lambda_{1}(\mathcal{H})>$ $\max \left\{N_{1} a, N_{2} c\right\}>\lambda_{2}(A)$. Therefore, for $b>0, \frac{y}{x} \geq 0$ for first eigenvector and $\leq 0$ for $v_{2}$. and for $b<0$, it is the opposite.

## A.4. Proof of Lemma 2.7

Proof of Lemma 2.7. We first show that $\left|x_{2}\right|,\left|y_{2}\right|$ are within a constant factor of each other, given that $\omega_{1} / \omega_{2}=\frac{N_{1}}{N_{2}}$ is a constant.
Proposition A.9. For a normalized $\mathcal{X}$, where $N_{1}, N_{2}, a, b \neq 0, x_{2}, y_{2}$ in the second top left singular vector $\bar{u}_{2}$ satisfy

$$
\frac{2 N_{2}}{N_{1}} \geq \frac{\left|x_{2}\right|}{\left|y_{2}\right|} \geq \frac{N_{2}}{2 N_{1}}
$$

Proof. By (A.9) and given the upper bound on $\operatorname{gap}(\mathcal{H})$ in (A.8),

$$
\begin{equation*}
\frac{\left|y_{2}\right|}{\left|x_{2}\right|}=\frac{N_{1} a-\lambda_{2}}{N_{2} b}=\frac{N_{1} a-N_{2} c+\operatorname{gap}(\mathcal{H})}{2 N_{2} b} \leq \frac{N_{1} a}{N_{2} b} \tag{A.20}
\end{equation*}
$$

and hence $\frac{\left|x_{2}\right|}{\left|y_{2}\right|} \geq \frac{N_{2} b}{N_{1} a}$. By (A.10) and (A.8), we have

$$
\begin{equation*}
\frac{\left|x_{2}\right|}{\left|y_{2}\right|}=\frac{N_{2} c-\lambda_{2}}{N_{1} b}=\frac{N_{2} c-N_{1} a+\operatorname{gap}(\mathcal{H})}{2 N_{1} b} \leq \frac{N_{2} c}{N_{1} b} \tag{A.21}
\end{equation*}
$$

We finish the proof by observing that $\frac{1}{2} \leq \frac{a}{b} \leq 2, \quad \frac{1}{2} \leq \frac{c}{b} \leq 2$ :, due to the fact that $\frac{1}{2} \leq \frac{\mu_{1}^{j}}{\mu_{2}^{j}} \leq 2, \forall j=1, \ldots, K$ for $\mu_{i}^{j} \in[1 / 2,1]$ in a normalized $\mathcal{X}$, and the following lemma:
Lemma A.10. If $0<c_{\min } \leq \frac{a_{i}}{b_{i}} \leq c_{\max }, \forall i=1, \ldots, n$, where $a_{i}, b_{i}>0$, then $c_{\text {min }} \leq \frac{\sum_{i=1}^{n} a_{i}}{\sum_{i=1}^{n} b_{i}} \leq c_{\text {max }}$.
Proof. $c_{\text {min }}=\frac{\sum_{i=1}^{n} c_{\min } b_{i}}{\sum_{i=1}^{n} b_{i}} \leq \frac{\sum_{i=1}^{n} a_{i}}{\sum_{i=1}^{n} b_{i}} \leq \frac{\sum_{i=1}^{n} c_{\max } b_{i}}{\sum_{i=1}^{n} b_{i}}=c_{\text {max }}$.

Let $x=x_{2}$ and $y=y_{2}$. By Proposition A.9, $|y| \leq \frac{2 N_{1}|x|}{N_{2}}$ and

$$
1=N_{1} x^{2}+N_{2} y^{2} \leq N_{1} x^{2}+N_{2}\left(\frac{2|x| N_{1}}{N_{2}}\right)^{2} \leq x^{2}\left(\frac{4 N_{1}^{2}+N_{1} N_{2}}{N_{2}}\right)
$$

hence for $C_{x \text { min }}=\frac{\omega_{2}}{4 \omega_{1}^{2}+\omega_{1} \omega_{2}}$,

$$
\begin{equation*}
x^{2} \geq \frac{\omega_{2}}{4 \omega_{1}^{2}+\omega_{1} \omega_{2}} \frac{1}{2 N} \tag{A.22}
\end{equation*}
$$

Looking in the other direction, by Proposition A.9, $|x| \leq \frac{2|y| N_{2}}{N_{1}}$,

$$
1=N_{1} x^{2}+N_{2} y^{2} \leq N_{2} y^{2}+N_{1}\left(\frac{2|y| N_{2}}{N_{1}}\right)^{2} \leq y^{2}\left(\frac{N_{2} N_{1}+4 N_{2}^{2}}{N_{1}}\right)
$$

and hence for a given $C_{y \text { min }}=\frac{\omega_{1}}{4 \omega_{2}^{2}+\omega_{1} \omega_{2}}$,

$$
|y|^{2} \geq \frac{\omega_{1}}{4 \omega_{2}^{2}+\omega_{1} \omega_{2}} \frac{1}{2 N} .
$$

On the other hand, by Proposition A.9, we have $|y| \geq \frac{N_{1}|x|}{2 N_{2}}$, we have

$$
1=N_{1} x^{2}+N_{2} y^{2} \geq N_{1} x^{2}+N_{2}\left(\frac{|x| N_{1}}{2 N_{2}}\right)^{2} \geq x^{2}\left(\frac{N_{1}^{2}+4 N_{1} N_{2}}{4 N_{2}}\right),
$$

and thus $x^{2} \leq \frac{4 \omega_{2}}{\omega_{1}^{2}+4 \omega_{1} \omega_{2}} \frac{1}{2 N}$. Looking in the other direction, by Proposition A.9, $|x| \geq \frac{|y| N_{2}}{2 N_{1}}$,

$$
1=N_{1} x^{2}+N_{2} y^{2} \geq N_{2} y^{2}+N_{1}\left(\frac{|y| N_{2}}{2 N_{1}}\right)^{2} \geq y^{2}\left(\frac{4 N_{2} N_{1}+N_{2}^{2}}{4 N_{1}}\right),
$$

and hence $|y|^{2} \leq \frac{4 \omega_{1}}{\omega_{2}^{2}+4 \omega_{1} \omega_{2}} \frac{1}{2 N}$. Hence we have that

$$
\begin{aligned}
|x-y|^{2}=(|x|+|y|)^{2} & \leq\left(\sqrt{\frac{4 \omega_{2}}{\omega_{1}^{2}+4 \omega_{1} \omega_{2}} \frac{1}{2 N}}+\sqrt{\frac{4 \omega_{1}}{\omega_{2}^{2}+4 \omega_{1} \omega_{2}} \frac{1}{2 N}}\right)^{2} \\
& \leq \frac{1}{2 N}\left(\sqrt{\frac{4 \omega_{2}}{\omega_{1}^{2}+4 \omega_{1} \omega_{2}}}+\sqrt{\frac{4 \omega_{1}}{\omega_{2}^{2}+4 \omega_{1} \omega_{2}}}\right)^{2}
\end{aligned}
$$

and $C_{\max }=\left(\sqrt{\frac{1}{\omega_{1}}}+\sqrt{\frac{1}{\omega_{2}}}\right)^{2}$. Hence

$$
C_{\max } \leq\left(\sqrt{\frac{4 \omega_{2}}{\omega_{1}^{2}+4 \omega_{1} \omega_{2}}}+\sqrt{\frac{4 \omega_{1}}{\omega_{2}^{2}+4 \omega_{1} \omega_{2}}}\right)^{2} \leq\left(\sqrt{\frac{1}{\omega_{1}}}+\sqrt{\frac{1}{\omega_{2}}}\right)^{2} .
$$

## Appendix B: Proof of Lemma 2.16

Recall that the largest left singular vectors $u_{1}, u_{2}$ has the form of $[x, \ldots, x, y, \ldots, y]$, where $x$ repeats $N_{1}$ times and $y$ repeats $N_{2}$ times. Proof of Lemma 2.16. Let us define the following random variables,

$$
\begin{equation*}
\delta=\frac{1}{N_{1}} \sum_{i=1}^{N_{1}} \delta_{i}, \quad \tau=\frac{1}{N_{2}} \sum_{i=1}^{N_{2}} \tau_{i} \tag{B.1}
\end{equation*}
$$

such that by Claim 2.15,

$$
\begin{aligned}
& |\delta|=\left|\frac{1}{N_{1}} \sum_{i=1}^{N_{1}} \delta_{i}\right| \leq \frac{1}{N_{1}} \sum_{i=1}^{N_{1}}\left|\delta_{i}\right| \leq \frac{\sqrt{N_{1} \sum_{i=1}^{N_{1}} \delta_{i}^{2}}}{N_{1}} \leq \frac{c_{1} \sigma}{\sqrt{N_{1} N}} \\
& |\tau|=\left|\frac{1}{N_{2}} \sum_{i=1}^{N_{2}} \tau_{i}\right| \leq \frac{1}{N_{2}} \sum_{i=1}^{N_{2}}\left|\tau_{i}\right| \leq \frac{\sqrt{N_{2} \sum_{i=1}^{N_{2}} \tau_{i}^{2}}}{N_{2}} \leq \frac{c_{1} \sigma}{\sqrt{N_{2} N}}
\end{aligned}
$$

and hence

$$
\begin{equation*}
\max \left(\left|N_{1} \delta\right|,\left|N_{2} \tau\right|\right) \leq \frac{c_{1} \sigma \sqrt{N_{2}}}{\sqrt{N}} \tag{B.2}
\end{equation*}
$$

given that we always assume that $N_{2}>N_{1}$. A natural classifier to separate individuals would be: $\frac{x+y}{2}$ when we use $u_{1}$; but we do not have access to $x$ and $y$. Recall that

$$
M=\frac{\sum_{i=1}^{N_{1}}\left(x+\delta_{i}\right)+\sum_{i=1}^{N_{2}}\left(y+\tau_{i}\right)}{2 N}=\frac{N_{1} x+N_{2} y}{2 N}+\frac{N_{1} \delta+N_{2} \tau}{2 N} .
$$

We are now ready to show that when $N_{1}, N_{2}$ are large enough, we see enough separation between the mixture sample mean and both $x$ and $y$. We first prove the following claims.
Claim B.1. $x N_{1} \delta+y N_{2} \tau=-\frac{\|\epsilon\|_{2}^{2}}{2}$.
Proof. This claim is obvious given that $\left\|u_{1}\right\|_{2}=\left\|\bar{u}_{1}\right\|_{2}=1$, and $\bar{u}_{1}, u_{1}, \epsilon$ all being real vectors, $\left\|u_{1}\right\|_{2}^{2}=\left\|\bar{u}_{1}\right\|_{2}^{2}+\|\epsilon\|_{2}^{2}+2\left\langle\bar{u}_{1}, \epsilon\right\rangle=\left\|\bar{u}_{1}\right\|_{2}^{2}+\|\epsilon\|_{2}^{2}+2 x N_{1} \delta+$ $y N_{2} \tau$.

We next use $\frac{1}{\sqrt{2 N}}\left|x N_{1} \delta+y N_{2} \tau\right|$ to obtain a bound on $\left|\frac{N_{1} \delta+N_{2} \tau}{2 N}\right|$, given that

$$
\begin{equation*}
\frac{1}{\sqrt{2 N}}\left|x N_{1} \delta+y N_{2} \tau\right| \leq \frac{\|\epsilon\|_{2}^{2}}{2 \sqrt{2 N}} \leq \frac{c_{1}^{2} \sigma^{2}}{2 N \sqrt{2 N}} \tag{B.3}
\end{equation*}
$$

Claim B.2. Let $N_{1} \leq N_{2}$, and $\omega_{1}=\frac{N_{1}}{2 N}$ and $\omega_{2}=\frac{N_{2}}{2 N}$, and given that $N_{1} \geq$ $\max \left(\frac{2 c_{1}^{2} \sigma^{2}}{c_{2} \gamma}, \frac{8 c_{1}^{2} \sigma^{2}}{\gamma} \frac{\omega_{2}}{\omega_{1}}\right)$, we have

$$
\begin{equation*}
\left|\frac{N_{1} \delta+N_{2} \tau}{2 N}\right| \leq \frac{N_{1}|y-x| \sqrt{\gamma}}{2 N} \tag{B.4}
\end{equation*}
$$

Proof. We next derive a bound on $\frac{N_{1} \delta+N_{2} \tau}{2 N}$. By Separation Lemma 2.8, we have $|x-y|=c_{2} \sqrt{\frac{\gamma}{2 N}}$ for a constant $c_{2}=1 / 2$, and thus we have $\frac{\max (x, y)}{\sqrt{2 N}}>\frac{1}{2 N}$. Therefore,

$$
\begin{aligned}
& \frac{\left|x N_{1} \delta+y N_{2} \tau\right|}{\sqrt{2 N}} \\
& =\frac{\left|\max (x, y)\left(N_{1} \delta+N_{2} \tau\right)+(x-\max (x, y)) N_{1} \delta+(y-\max (x, y)) N_{2} \tau\right|}{\sqrt{2 N}} \\
& \quad \geq \frac{|\max (x, y)|\left(N_{1} \delta+N_{2} \tau\right)}{\sqrt{2 N}}-\frac{|x-y| \max \left(\left|N_{1} \delta\right|,\left|N_{2} \tau\right|\right)}{\sqrt{2 N}}
\end{aligned}
$$

Thus we have, given (B.2), (B.3) and (B.5),

$$
\begin{aligned}
\left|\frac{N_{1} \delta+N_{2} \tau}{2 N}\right| & \leq \frac{|\max (x, y)|}{\sqrt{2 N}}\left|\left(N_{1} \delta+N_{2} \tau\right)\right| \\
& \leq \frac{\left|x N_{1} \delta+y N_{2} \tau\right|}{\sqrt{2 N}}+\frac{|x-y| \max \left(\left|N_{1} \delta\right|,\left|N_{2} \tau\right|\right)}{\sqrt{2 N}} \\
& \leq \frac{c_{1}^{2} \sigma^{2}}{2 N \sqrt{2 N}}+\frac{c_{2} \sqrt{\gamma}}{2 N} c_{1} \sigma \sqrt{\frac{N_{2}}{N}} \\
& \leq \frac{N_{1} c_{2} \gamma}{2 N \sqrt{2 N}} \leq \frac{N_{1}|y-x| \sqrt{\gamma}}{2 N}
\end{aligned}
$$

where

$$
\begin{array}{r}
\frac{c_{1}^{2} \sigma^{2}}{2 N \sqrt{2 N}}<\frac{N_{1} c_{2} \gamma}{4 N \sqrt{2 N}}, \text { holds so long as } N_{1} \geq \frac{2 c_{1}^{2} \sigma^{2}}{c_{2} \gamma}, \text { and } \\
\frac{c_{2} \sqrt{\gamma} c_{1} \sigma}{2 N} \sqrt{\frac{N_{2}}{N}}<\frac{N_{1} c_{2} \gamma}{4 N \sqrt{2 N}} . \text { holds so long as } N_{1} \geq \frac{8 c_{1}^{2} \sigma^{2}}{\gamma} \frac{\omega_{2}}{\omega_{1}}, \text { so that } \tag{B.6}
\end{array}
$$

$$
N_{1} \geq \frac{2 \sqrt{2} c_{1} \sigma \sqrt{N_{2}}}{\sqrt{\gamma}}
$$

Both conditions are guaranteed by (2.23) in Theorem 2.14.
This allows us to conclude that

$$
\begin{aligned}
\left|\frac{N_{1} x+N_{1} \delta+N_{2} y+N_{2} \tau}{2 N}-\frac{N_{1} x+N_{2} y}{2 N}\right| & =\left|\frac{N_{1} \delta+N_{2} \tau}{2 N}\right| \\
& \leq\left(\frac{\min \left\{N_{1}, N_{2}\right\} \sqrt{\gamma}}{2 N}\right)|x-y|
\end{aligned}
$$

Given that $|y-x|=c_{2} \sqrt{\gamma} / \sqrt{2 N}$ as shown in the Separation Lemma 2.8, we have

$$
\begin{aligned}
\left|\frac{N_{1} x+N_{1} \delta+N_{2} y+N_{2} \tau}{2 N}-x\right| & =\left|\frac{N_{2}(y-x)}{2 N}+\frac{N_{1} \delta+N_{2} \tau}{2 N}\right| \\
& \geq\left|\frac{N_{2}(y-x)}{2 N}\right|-\left|\frac{N_{1} \delta+N_{2} \tau}{2 N}\right| \\
& \geq \frac{N_{2}|y-x|}{2 N}-\frac{\min \left\{N_{1}, N_{2}\right\} \sqrt{\gamma}|y-x|}{2 N} \\
& \geq \frac{(1-\sqrt{\gamma}) N_{2}|y-x|}{2 N},
\end{aligned}
$$

and similarly,

$$
\begin{aligned}
\left|y-\frac{N_{1} x+N_{1} \delta+N_{2} y+N_{2} \tau}{2 N}\right| & =\left|\frac{N_{1}(y-x)}{2 N}-\frac{N_{1} \delta+N_{2} \tau}{2 N}\right| \\
& \geq\left|\frac{N_{1}(y-x)}{2 N}\right|-\left|\frac{N_{1} \delta+N_{2} \tau}{2 N}\right| \\
& \geq \frac{N_{1}|y-x|}{2 N}-\frac{\min \left\{N_{1}, N_{2}\right\} \sqrt{\gamma}|y-x|}{2 N} \\
& \geq \frac{(1-\sqrt{\gamma}) N_{1}|y-x|}{2 N} .
\end{aligned}
$$

## References

Achlioptas, D. and McSherry, F. (2005). On spectral learning of mixtures of distributions. In Proceedings of the 18th Annual COLT. (Version in http://www.cs.ucsc.edu/~optas/papers/). MR2203280
Arora, S. and Kannan, R. (2001). Learning mixtures of arbitrary gaussians. In Proceedings of 33rd ACM Symposium on Theory of Computing. MR2120323
Chaudhuri, K., Halperin, E., Rao, S. and Zhou, S. (2007). A rigorous analysis of population stratification with limited data. In Proceedings of the 18th ACM-SIAM SODA.
Coja-Oghlan, A. (2006). An adaptive spectral heuristic for partitioning random graphs. In Proceedings of the 33rd ICALP. MR2305568
Cryan, M. (1999). Learning and approximation Algorithms for Problems motivated by evolutionary trees. Ph.D. thesis, University of Warwick.
Cryan, M., Goldberg, L. and Goldberg, P. (2002). Evolutionary trees can be learned in polynomial time in the two state general markov model. SIAM J. of Computing 31 375-397. MR1861281

Dasgupta, A., Hopcroft, J., Kleinberg, J. and Sandler, M. (2005). On learning mixtures of heavy-tailed distributions. In Proceedings of the 46 th IEEE FOCS.
Dasqupta, S. (1999). Learning mixtures of gaussians. In Proceedings of the 40th IEEE Symposium on Foundations of Computer S cience. MR1917603
Dasgupta, S. and Schulman, L. J. (2000). A two-round variant of em for gaussian mixtures. In Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence (UAI).
Feldman, J., O’Donnell, R. and Servedio, R. (2005). Learning mixtures of product distributions over discrete domains. In Proceedings of the 46 th IEEE FOCS.
Feldman, J., O’Donnell, R. and Servedio, R. (2006). PAC learning mixtures of Gaussians with no separation assumption. In Proceedings of the 19th Annual COLT.
Fiedler, M. (1973). Algebraic connectivity of graphs. Czechoslovak Mathematical Journal 298-305. MR0318007

FJallstrom, P. (1998). Algorithms for graph partitioning: a survey. Tech. rep., Linkoping University Electroni Press.
Freund, Y. and Mansour, Y. (1999). Estimating a mixture of two product distributions. In Proceedings of the 12th Annual COLT.
Kannan, R., Salmasian, H. and Vempala, S. (2005). The spectral method for general mixture models. In Proc. of the 18th Annual COLT. MR2203279
Kearns, M., Mansour, Y., Ron, D., Rubinfeld, R., Schapir, R. and Sellie, L. (1994). On the learnability of discrete distributions. In Proceedings of the 26th ACM STOC.
Latala, R. (2005). Some estimates of norms of random matrices. In Proceedings of the American Mathematical Society, vol. 133. MR2111932
McSherry, F. (2001). Spectral partitioning of random graphs. In Proceedings of the $42 n d$ IEEE Symposium on Foundations of Computer Science. MR1948742
Meckes, M. (2004). Concentration of norms and eigenvalues of random matrices. J. Funct. Anal. 211 508-524. MR2057479
Mossel, E. and Roch, S. (2005). Learning nonsinglar phylogenies and hidden markov models. In Proceedings of the 37th ACM STOC. MR2181638
Patterson, N., Price, A. and Reich, D. (2006). Population structure and eigenanalysis. PLoS Genet 2. Doi:10.1371/journal.pgen. 0020190.
Price, A., Patterson, N., Plenge, R., Weinblatt, M., Shadick, N. and Reich, D. (2006). Principal components analysis corrects for stratification in genome-wide association studies. nature genetics 38 904-909.
Pritchard, J. K., Stephens, M. and Donnelly, P. (2000). Inference of population structure using multilocus genotype data. Genetics 155 954-959.
Spielman, D. (2002). The behavior of algorithms in practice. Lecture notes.
Vempala, V. and Wang, G. (2002). A spectral algorithm of learning mixtures of distributions. In Proceedings of the 43rd IEEE FOCS.
Vu, V. (2005). Spectral norm of random matrices. In Proceedings of 37 th ACM STOC. MR2181644
Zhou, S. (2006). Routing, Disjoint Paths, and Classification. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA. CMU Technical Report, CMU-PDL-06-109.


[^0]:    *Supported in part by the NSF under grant CCF-0514922
    $\dagger$ Supported by the German Research Foundation under grant CO 646
    ${ }^{\ddagger}$ Supported in part by the NSF under grant CCF-0502793
    §Supported in part by the NSF under grants CCF-0625879 and CNF-0435382

[^1]:    ${ }^{1}$ One can also obtain an upper bound of $O(\sqrt{n+K})$ on $s_{1}(A)$ using a theorem on by Vu (2005), through the construction a $(n+K) \times(n+K)$ square matrix out of $A$.

