APPROXIMATION ALGORITHMS FOR SPARSE PRINCIPAL COMPONENT ANALYSIS

Anonymous authors
Paper under double-blind review

ABSTRACT

Principal component analysis (PCA) is a widely used dimension reduction technique in machine learning and multivariate statistics. To improve the interpretability of PCA, various approaches to obtain sparse principal direction loadings have been proposed, which are termed Sparse Principal Component Analysis (SPCA). In this paper, we present three provably accurate, polynomial time, approximation algorithms for the SPCA problem, without imposing any restrictive assumptions on the input covariance matrix. The first algorithm is based on randomized matrix multiplication; the second algorithm is based on a novel deterministic thresholding scheme; and the third algorithm is based on a semidefinite programming relaxation of SPCA. All algorithms come with provable guarantees and run in low-degree polynomial time. Our empirical evaluations confirm our theoretical findings.

1 INTRODUCTION

Principal Component Analysis (PCA) and the related Singular Value Decomposition (SVD) are fundamental data analysis and dimension reduction tools in a wide range of areas including machine learning, multivariate statistics and many others. They return a set of orthogonal vectors of decreasing importance that are often interpreted as fundamental latent factors that underlie the observed data. Even though the vectors returned by PCA and SVD have strong optimality properties, they are notoriously difficult to interpret in terms of the underlying processes generating the data (Mahoney & Drineas, 2009), since they are linear combinations of *all* available data points or *all* available features. The concept of Sparse Principal Components Analysis (SPCA) was introduced in the seminal work of (d'Aspremont et al., 2007), where sparsity constraints were enforced on the singular vectors in order to improve interpretability. A prominent example where sparsity improves interpretability is document analysis, where sparse principal components can be mapped to specific topics by inspecting the (few) keywords in their support (d'Aspremont et al., 2007; Mahoney & Drineas, 2009; Papailiopoulos et al., 2013).

Formally, given a positive semidefinite (PSD) matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, SPCA can be defined as follows:[1]

$$\mathcal{Z}^* = \max_{\mathbf{x} \in \mathbb{R}^n,\ \|\mathbf{x}\|_2 \leq 1} \mathbf{x}^\top \mathbf{A} \mathbf{x}, \qquad \text{subject to } \|\mathbf{x}\|_0 \leq k. \tag{1}$$

In the above formulation, $\mathbf{A}$ is a covariance matrix representing, for example, all pairwise feature or object similarities for an underlying data matrix. Therefore, SPCA can be applied for either the object or feature space of the data matrix, while the parameter k controls the sparsity of the resulting vector and is part of the input. Let $\mathbf{x}^*$ denote a vector that achieves the optimal value $\mathcal{Z}^*$ in the above formulation. Then intuitively, the optimization problem of eqn. (1) seeks a *sparse*, unit norm vector $\mathbf{x}^*$ that maximizes the data variance.

It is well-known that solving the above optimization problem is NP-hard (Moghaddam et al., 2006a) and that its hardness is due to the sparsity constraint. Indeed, if the sparsity constraint was removed, then the resulting optimization problem can be easily solved by computing the top left or right singular vector of $\mathbf{A}$ and its maximal value $\mathcal{Z}^*$ is equal to the top singular value of $\mathbf{A}$.

Notation. We use bold letters to denote matrices and vectors. For a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, we denote its (i, j)-th entry by $A_{i,j}$; its i-th row by $\mathbf{A}_{i*}$ and its j-th column by $\mathbf{A}_{*j}$; its 2-norm by

[1] Recall that the p-th power of the ℓ_p norm of a vector $\mathbf{x} \in \mathbb{R}^n$ is defined as $\|\mathbf{x}\|_p^p = \sum_{i=1}^n |\mathbf{x}_i|^p$ for $0 < p < \infty$. For $p = 0$, $\|x\|_0$ is a semi-norm denoting the number of non-zero entries of x.

$\|\mathbf{A}\|_2 = \max_{\mathbf{x}\in\mathbb{R}^n,\ \|\mathbf{x}\|_2=1} \|\mathbf{A}\mathbf{x}\|_2$; and its (squared) Frobenius norm by $\|\mathbf{A}\|_F^2 = \sum_{i,j} A_{i,j}^2$. We use the notation $\mathbf{A} \succeq 0$ to denote that the matrix $\mathbf{A}$ is symmetric positive semidefinite (PSD) and $\mathrm{Tr}(\mathbf{A}) = \sum_i A_{i,i}$ to denote its trace, which is also equal to the sum of its singular values. Given a PSD matrix $\mathbf{A} \in \mathbb{R}^{n\times n}$, its Singular Value Decomposition is given by $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^T$, where $\mathbf{U}$ is the matrix of left/right singular vectors and $\mathbf{\Sigma}$ is the diagonal matrix of singular values.

1.1 Our Contributions

We present three algorithms for SPCA and associated quality-of-approximation results (Theorems 2.2, 3.1, and 4.1). All three algorithms are simple, intuitive, and run in $\mathcal{O}\left(n^{3.5}\right)$ or less time. They return a vector that is provably sparse and, when applied to the input covariance matrix $\mathbf{A}$, provably captures a fraction of the optimal solution $\mathcal{Z}^*$. We note that in all three algorithms, the output vector has a sparsity that depends on k (the target sparsity of the original SPCA problem of eqn. (1)) and ϵ (an accuracy parameter between zero and one).

The first algorithm is based on randomized, approximate matrix multiplication: it randomly (but non-uniformly) selects a subset of $\mathcal{O}\left(k/\epsilon^2\right)$ columns of $\mathbf{A}^{1/2}$ (the square root of the PSD matrix $\mathbf{A}$) and computes its top right singular vector. The output of this algorithm is precisely this singular vector, padded with zeros to become a vector in $\mathbb{R}^n$. It turns out that this simple algorithm, which, surprisingly has not been analyzed in prior work, returns an $\mathcal{O}\left(k/\epsilon^2\right)$ sparse vector $\mathbf{y} \in \mathbb{R}^n$ that satisfies (with constant probability that can be amplified as desired, see Section 2 for details):

$$\mathbf{y}^\top \mathbf{A}\mathbf{y} \geq \frac{1}{2}\mathcal{Z}^* - \epsilon\sqrt{\mathcal{Z}^*} \cdot \sqrt{\frac{\mathrm{Tr}(\mathbf{A})}{k}}.$$

Notice that the above bound depends on both $\mathcal{Z}^*$ and it square root and therefore is not a relative error bound. The second term scales as a function of the trace of $\mathbf{A}$ divided by k, which depends on the properties of the matrix $\mathbf{A}$ and the target sparsity.

The second algorithm is a deterministic thresholding scheme. It computes a small number of the top singular vectors of the matrix $\mathbf{A}$ and then applies a deterministic thresholding scheme on those singular vectors to (eventually) construct a sparse vector $\mathbf{z} \in \mathbb{R}^n$ that satisfies $\mathbf{z}^\top \mathbf{A}\mathbf{z} \geq (1/2)\mathcal{Z}^* - (3/2)\epsilon\,\mathrm{Tr}(\mathbf{A})$. Our analysis provides unconditional guarantees for the accuracy of the solution of this simple thresholding scheme. To the best of our knowledge, no such analyses have appeared in prior work (see Section 1.2 for details). The error bound of the second algorithm is weaker than the one provided by the first algorithm, but the second algorithm is deterministic and does not need to compute the square root (i.e., all singular vectors and singular values) of the matrix $\mathbf{A}$.

Our third algorithm provides novel bounds for the following standard convex relaxation of the problem of eqn. (1).

$$\max_{\mathbf{Z}\in\mathbb{R}^{n\times n},\ \mathbf{Z}\succeq 0} \mathrm{Tr}(\mathbf{A}\mathbf{Z}) \text{ s.t. } \mathrm{Tr}(\mathbf{Z}) \leq 1 \text{ and } \sum |Z_{i,j}| \leq k. \tag{2}$$

It is well-known that the optimal solution to eqn. (2) is at least the optimal solution to eqn. (1). We present a novel, two-step rounding scheme that converts the optimal solution matrix $\mathbf{Z} \in \mathbb{R}^{n\times n}$ to a vector $\mathbf{z} \in \mathbb{R}^n$ that has expected sparsity[2] $\tilde{\mathcal{O}}\left(k^2/\epsilon^2\right)$ and satisfies $\mathbf{z}^\top \mathbf{A}\mathbf{z} \geq \gamma_Z(1-\epsilon)\cdot\mathcal{Z}^* - \epsilon$. Here, γ_Z is a constant that precisely depends on the top singular value of $\mathbf{Z}$, the condition number of $\mathbf{Z}$, and the extent to which the SDP relaxation of eqn. (2) is able to capture the original problem (see Theorem 4.1 and the following discussion for details). To the best of our knowledge, this is the first analysis of a rounding scheme for the convex relaxation of eqn. (2) that does not assume a specific model for the covariance matrix $\mathbf{A}$.

Applications to Sparse Kernel PCA. Our algorithms have immediate applications to sparse kernel PCA (SKPCA), where the input matrix $\mathbf{A} \in \mathbb{R}^{n\times n}$ is instead implicitly given as a kernel matrix whose entry (i,j) is the value $k(i,j) := \langle \phi(\mathbf{X}_{i*}), \phi(\mathbf{X}_{j*})\rangle$ for some kernel function ϕ that implicitly maps an observation vector into some high-dimensional feature space. Although $\mathbf{A}$ is not explicit,

[2]For simplicity of presentation and following the lines of (Fountoulakis et al., 2017), we assume that the rows and columns of the matrix $\mathbf{A}$ have unit norm; this assumption was not necessary for the previous two algorithms and can be removed as in (Fountoulakis et al., 2017). We are also hiding a poly-logarithmic factor for simplicity, hence the $\tilde{O}(\cdot)$ notation. See Theorem 4.1 for a detailed statement.

we can query all $\mathcal{O}\left(n^2\right)$ entries of $\mathbf{A}$ using $\mathcal{O}\left(n^2\right)$ time assuming an oracle that computes the kernel function k. We can then subsequently apply our SPCA algorithms and achieve polynomial time runtime with the same approximation guarantees.

1.2 Prior work

SPCA was formally introduced by (d'Aspremont et al., 2007); however, previously studied PCA approaches based on rotating (Jolliffe, 1995) or thresholding (Cadima & Jolliffe, 1995) the top singular vector of the input matrix seemed to work well, at least in practice, given sparsity constraints. Following (d'Aspremont et al., 2007), there has been an abundance of interest in SPCA. (Jolliffe et al., 2003) considered LASSO (SCoTLASS) on an ℓ_1 relaxation of the problem, while (Zou & Hastie, 2005) considered a non-convex regression-type approximation, penalized similar to LASSO. Additional heuristics based on LASSO (Ando et al., 2009) and non-convex ℓ_1 regularizations (Zou & Hastie, 2005; Zou et al., 2006; Sriperumbudur et al., 2007; Shen & Huang, 2008) have also been explored. Random sampling approaches based on non-convex ℓ_1 relaxations (Fountoulakis et al., 2017) have also been studied; we highlight that unlike our approach, (Fountoulakis et al., 2017) solved a non-convex relaxation of the SPCA problem and thus perhaps relied on locally optimal solutions. Additionally, (Moghaddam et al., 2006b) considered a branch-and-bound heuristic motivated by greedy spectral ideas. (Journée et al., 2010; Papailiopoulos et al., 2013; Kuleshov, 2013; Yuan & Zhang, 2013) further explored other spectral approaches based on iterative methods similar to the power method. (Yuan & Zhang, 2013) specifically designed a sparse PCA algorithm with early stopping for the power method, based on the target sparsity.

Another line of work focused on using semidefinite programming (SDP) relaxations (d'Aspremont et al., 2007; d'Aspremont et al., 2008; Amini & Wainwright, 2009). Notably, (Amini & Wainwright, 2009) achieved provable theoretical guarantees regarding the SDP and thresholding approach of (d'Aspremont et al., 2007) in a *specific*, high-dimensional spiked covariance model, in which a base matrix is perturbed by adding a sparse maximal eigenvector. In other words, the input matrix is the identity matrix plus a "spike", i.e., a sparse rank-one matrix.

Despite the variety of heuristic-based sparse PCA approaches, very few theoretical guarantees have been provided for SPCA; this is partially explained by a line of hardness-of-approximation results. The sparse PCA problem is well-known to be NP-Hard (Moghaddam et al., 2006a). (Magdon-Ismail, 2017) shows that if the input matrix is not PSD, then even the *sign* of the optimal value cannot be determined in polynomial time unless $\mathsf{P} = \mathsf{NP}$, ruling out any multiplicative approximation algorithm. In the case where the input matrix is PSD, (Chan et al., 2016) shows that it is NP-hard to approximate the optimal value up to multiplicative $(1 + \epsilon)$ error, ruling out any polynomial-time approximation scheme (PTAS). Moreover, they show Small-Set Expansion hardness for any polynomial-time constant factor approximation algorithm and also that the standard SDP relaxation might have an exponential gap.

We conclude by summarizing prior work that offers provable guarantees (beyond the work of (Amini & Wainwright, 2009)), typically given *some assumptions about the input matrix*. (d'Aspremont et al., 2014) showed that the SDP relaxation can be used to find provable bounds when the covariance input matrix is formed by a number of data points sampled from Gaussian models with a single sparse singular vector. (Papailiopoulos et al., 2013) presented a combinatorial algorithm that analyzed a specific set of vectors in a low-dimensional eigenspace of the input matrix and presented relative error guarantees for the optimal objective, given the assumption that the input covariance matrix has a decaying spectrum. (Asteris et al., 2011) gave a polynomial-time algorithm that solves sparse PCA *exactly* for input matrices of constant rank. (Chan et al., 2016) showed that sparse PCA can be approximated in polynomial time within a factor of $n^{-1/3}$ and also highlighted an additive PTAS of (Asteris et al., 2015) based on the idea of finding multiple disjoint components and solving bipartite maximum weight matching problems. This PTAS needs time $n^{\mathrm{poly}(1/\epsilon)}$, whereas all of our algorithms have running times that are a low-degree polynomial in n.

2 Sparse PCA via Randomized Matrix Multiplication

Our first algorithm for SPCA leverages primitives and ideas from randomized matrix multiplication (Drineas & Kannan, 2001; Drineas et al., 2006; Drineas & Mahoney, 2016; 2018; Woodruff, 2014). Let $\mathbf{P} \in \mathbb{R}^{m\times n}$ and $\mathbf{Q} \in \mathbb{R}^{n\times p}$ and recall that their product $\mathbf{PQ}$ equals $\mathbf{PQ} = \sum_{i=1}^{n} \mathbf{P}_{*i}\mathbf{Q}_{i*}$.

Recall that $\mathbf{P}_{*i}$ denotes the i-th column of $\mathbf{P}$ and $\mathbf{Q}_{i*}$ the i-th row of $\mathbf{Q}$. A well-known approach to approximate the product $\mathbf{PQ}$ is to sample a subset of columns of $\mathbf{P}$ (we will do this without replacement) and the corresponding rows of $\mathbf{Q}$ (Drineas & Mahoney, 2018). Formally, let the random variables $Z_i \overset{iid}{\sim} \text{Bernoulli}(p_i)$, $i = 1 \ldots n$, denote whether the i-th column of $\mathbf{P}$ and the i-th row of $\mathbf{Q}$ are sampled. Define the diagonal sampling-and-rescaling matrix $\mathbf{S} \in \mathbb{R}^{n \times n}$ by $\mathbf{S} \triangleq \operatorname{diag}\left\{ {}^{Z_1}/_{\sqrt{p_1}}, \ldots, {}^{Z_n}/_{\sqrt{p_n}} \right\}$. The sampling probabilities $\{p_i\}_{i=1}^n$ *do not* have to sum up to one. The number of sampled column/row pairs, denoted by s, satisfies $\mathbb{E}[s] = \sum_{i=1}^n p_i$. (See Algorithm 1 for details.) The next lemma (see Appendix A.1 for its proof) presents accuracy bounds

Algorithm 1 Construct sampling-and-rescaling matrix $\mathbf{S}$

Input: Probabilities $\tilde{p}_i$, $i = 1 \ldots n$ and integer $s \ll n$.
Output: Diagonal sampling-and-rescaling matrix $\mathbf{S} \in \mathbb{R}^{n \times n}$.

1: **for** $i = 1$ **to** n **do**
2: $\quad \mathbf{S}_{ii} \leftarrow \begin{cases} {}^{1}/_{\sqrt{p_i}}, & \text{with probability } p_i = \min\{s\tilde{p}_i, 1\} \\ 0, & \text{otherwise} \end{cases}$

when Algorithm 1 is used to approximate matrix multiplication.

Lemma 2.1 *Given matrices $\mathbf{P} \in \mathbb{R}^{m \times n}$ and $\mathbf{Q} \in \mathbb{R}^{n \times p}$, let $\mathbf{S} \in \mathbb{R}^{n \times n}$ be constructed using Algorithm 1 with $\tilde{p}_i = {}^{\|\mathbf{P}_{*i}\|_2^2}/_{\|\mathbf{P}\|_F^2}$, for $i = 1 \ldots n$. Then,*

$$\mathbb{E}\left[\|\mathbf{PS}^2\mathbf{Q} - \mathbf{PQ}\|_F^2\right] \leq \frac{1}{s} \|\mathbf{P}\|_F^2 \|\mathbf{Q}\|_F^2 \,. \tag{3}$$

Our SPCA algorithm uses the above primitive to approximate the product of the (square root) of the input matrix $\mathbf{A}$ and its top right singular vector $\mathbf{v}$. Thus, the proposed SPCA algorithm *sparsifies* the top right singular vector of $\mathbf{v}$ of $\mathbf{A}$ without losing too much of the variance that is captured by $\mathbf{v}$. Interestingly, this conceptually simple algorithm has not been formally analyzed in prior work. Algorithm 2 details our approach.

Algorithm 2 SPCA via randomized matrix multiplication

Input: $\mathbf{A} \in \mathbb{R}^{n \times n}$, sparsity parameter k, accuracy parameter $\epsilon \in (0, 1)$.
Output: $\mathbf{y} \in \mathbb{R}^n$ satisfying $\mathbb{E}[\|\mathbf{y}\|_2] \leq 1$ and $\mathbb{E}[\|\mathbf{y}\|_0] \leq k/\epsilon^2$.

1. $\mathbf{X} \leftarrow \mathbf{A}^{1/2}$;
2. Use Algorithm 1 to construct $\mathbf{S} \in \mathbb{R}^{n \times n}$ with $\tilde{p}_i = {}^{\|\mathbf{X}_{*i}\|_2^2}/_{\|\mathbf{X}\|_F^2}$ and $s = {}^{4k}/_{\epsilon^2}$;
3. Let $\mathbf{v} \in \mathbb{R}^n$ be the top right singular vector of $\mathbf{XS}$;
4. $\mathbf{y} \leftarrow \mathbf{Sv}$;

Theorem 2.2 *Let k be the sparsity parameter and $\epsilon \in (0, 1]$ be the accuracy parameter. Let $\mathbf{S} \in \mathbb{R}^{n \times n}$ be the sampling matrix of Lemma 2.1 with $s = {}^{4k}/_{\epsilon^2}$. Then, Algorithm 2 returns a vector $\mathbf{y}$ with expected sparsity at most s (i.e., $\mathbb{E}[\|\mathbf{y}\|_0] \leq s$) and expected two norm at most one (i.e., $\mathbb{E}\left[\|\mathbf{y}\|_2^2\right] \leq 1$) such that with probability at least ${}^1/_4$, we have*

$$\mathbf{y}^\top \mathbf{A} \mathbf{y} \geq {}^1/_2 \mathcal{Z}^* - \epsilon \sqrt{\mathcal{Z}^*} \cdot \sqrt{{}^{\operatorname{Tr}(\mathbf{A})}/_k}. \tag{4}$$

See Appendix A.1 for a proof of the above theorem. We note that the success probability of Algorithm 2 can be trivially amplified by repeating the algorithm t times and keeping the vector $\mathbf{y}$ that maximizes $\mathbf{y}^\top \mathbf{A} \mathbf{y}$. Then, the failure probability of the overall approach diminishes exponentially fast as a function of t to at most $({}^3/_4)^t$. Finally, the running time of Algorithm 2 is dominated by the computation of a square root of the matrix $\mathbf{A}$ in the first step, which takes $\mathcal{O}\left(n^3\right)$ time via the computation of the SVD of $\mathbf{A}$.

3 SPCA VIA THRESHOLDING

Our second algorithm is based on a thresholding scheme using the top ℓ right singular vectors of the PSD matrix $\mathbf{A}$. Given $\mathbf{A}$ and an accuracy parameter ϵ, our approach first computes $\mathbf{\Sigma}_\ell \in \mathbb{R}^{\ell \times \ell}$

(the diagonal matrix of the top ℓ singular values of $\mathbf{A}$) and $\mathbf{U}_\ell \in \mathbb{R}^{n\times\ell}$ (the matrix of the top ℓ right singular vectors of $\mathbf{A}$), for $\ell = 1/\epsilon$. Then, it *deterministically* selects a subset of $\mathcal{O}\left(k/\epsilon^2\right)$ columns of $\mathbf{\Sigma}_\ell^{1/2}\mathbf{U}_\ell^\top$ using a simple thresholding scheme based on the norms of the columns of $\mathbf{\Sigma}_\ell^{1/2}\mathbf{U}_\ell^\top$. (Recall that k is the sparsity parameter of the SPCA problem.) In the last step, it returns the *top right singular vector* of the matrix consisting of the chosen columns of $\mathbf{\Sigma}_\ell^{1/2}\mathbf{U}_\ell^\top$. Notice that this right singular vector is an $\mathcal{O}\left(k/\epsilon^2\right)$-dimensional vector, which is finally expanded to a vector in $\mathbb{R}^n$ by appropriate padding with zeros. This sparse vector is our approximate solution to the SPCA problem of eqn. (1).

This simple algorithm is somewhat reminiscent of prior thresholding approaches for SPCA. However, to the best of our knowledge, no provable a priori bounds were known for such algorithms without strong assumptions on the input matrix. This might be due to the fact that prior approaches focused on thresholding only the top right singular vector of $\mathbf{A}$, whereas our approach thresholds the top $\ell = 1/\epsilon$ right singular vectors of $\mathbf{A}$. This slight relaxation allows us to present provable bounds for the proposed algorithm.

In more detail, let the SVD of $\mathbf{A}$ be $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^T$. Let $\mathbf{\Sigma}_\ell \in \mathbb{R}^{\ell\times\ell}$ be the diagonal matrix of the top ℓ singular values and let $\mathbf{U}_\ell \in \mathbb{R}^{n\times\ell}$ be the matrix of the top ℓ right (or left) singular vectors. Let $R = \{i_1, \ldots, i_{|R|}\}$ be the set of indices of rows of $\mathbf{U}_\ell$ that have squared norms at least ϵ/k and let $\bar{R}$ be its complement. Here $|R|$ denotes the cardinality of the set R and $R \cup \bar{R} = \{1, \ldots, n\}$. Let $\mathbf{R} \in \mathbb{R}^{n\times|R|}$ be a sampling matrix that selects[3] the columns of $\mathbf{U}_\ell$ whose indices are in the set R. Given this notation, we are now ready to state Algorithm 3.

Algorithm 3 SPCA via thresholding

Input: $\mathbf{A} \in \mathbb{R}^{n\times n}$, sparsity k, error parameter $\epsilon > 0$.
Output: $\mathbf{y} \in \mathbb{R}^n$ such that $\|y\|_2 = 1$ and $\|\mathbf{y}\|_0 = k/\epsilon^2$.

1: $\ell \leftarrow 1/\epsilon$;
2: Compute $\mathbf{U}_\ell \in \mathbb{R}^{n\times\ell}$ (top ℓ left singular vectors of $\mathbf{A}$) and $\mathbf{\Sigma}_\ell \in \mathbb{R}^{\ell\times\ell}$ (square roots of the top ℓ singular values of $\mathbf{A}$);
3: Let $R = \{i_1, \ldots, i_{|R|}\}$ be the set of rows of $\mathbf{U}_\ell$ with squared norms at least ϵ/k and let $\mathbf{R} \in \mathbb{R}^{n\times|R|}$ be the associated sampling matrix (see text for details);
4: $\mathbf{y} \in \mathbb{R}^{|R|} \leftarrow \operatorname{argmax}_{\|\mathbf{x}\|_2=1} \left\|\mathbf{\Sigma}_\ell\mathbf{U}_\ell^\top\mathbf{R}\mathbf{x}\right\|_2^2$;
5: **return** $\mathbf{z} = \mathbf{R}\mathbf{y} \in \mathbb{R}^n$;

Notice that $\mathbf{R}\mathbf{y}$ satisfies $\|\mathbf{R}\mathbf{y}\|_2 = \|\mathbf{y}\|_2 = 1$ (since $\mathbf{R}$ has orthogonal columns) and $\|\mathbf{R}\mathbf{y}\|_0 = |R|$. Since R is the set of rows of $\mathbf{U}_\ell$ with squared norms at least ϵ/k and $\|\mathbf{U}_\ell\|_F^2 = \ell = 1/\epsilon$, it follows that $|R| \leq k/\epsilon^2$. Thus, the vector returned by Algorithm 3 has k/ϵ^2 sparsity and unit norm.

Theorem 3.1 *Let k be the sparsity parameter and $\epsilon \in (0, 1]$ be the accuracy parameter. Then, the vector $\mathbf{z} \in \mathbb{R}^n$ (the output of Algorithm 3) has sparsity k/ϵ^2, unit norm, and satisfies $\mathbf{z}^\top\mathbf{A}\mathbf{z} \geq (1/2)\mathcal{Z}^* - (3/2)\epsilon\,\mathrm{Tr}(\mathbf{A})$.*

We defer the proof of Theorem 3.1 to Appendix A.2. The running time of Algorithm 3 is dominated by the computation of the top ℓ singular vectors and singular values of the matrix $\mathbf{A}$. In practice, any iterative method, such as subspace iteration using a random initial subspace or the Krylov subspace of the matrix, can be used towards this end. However, our current analysis does not account for the inevitable approximation error incurred by such methods, which run in $\mathcal{O}\left(\mathrm{nnz}(\mathbf{A})\ell\right)$ time. One could always use the SVD of the full matrix $\mathbf{A}$ ($\mathcal{O}\left(n^3\right)$ time) to compute the top ℓ singular vectors and singular values of $\mathbf{A}$. Finally, we highlight that, as an intermediate step in the proof of Theorem 3.1, we need to prove the following lemma (see Appendix A.2 for its proof):

Lemma 3.2 *Let $\mathbf{A} \in \mathbb{R}^{n\times n}$ be a PSD matrix and $\mathbf{\Sigma} \in \mathbb{R}^{n\times n}$ (respectively, $\mathbf{\Sigma}_\ell \in \mathbb{R}^{\ell\times\ell}$) be the diagonal matrix of all (respectively, top ℓ) singular values and let $\mathbf{U} \in \mathbb{R}^{n\times n}$ (respectively, $\mathbf{U}_\ell \in \mathbb{R}^{n\times\ell}$) be the matrix of all (respectively, top ℓ) singular vectors. Then, for all unit vectors $\mathbf{x} \in \mathbb{R}^n$,*

$$\left\|\mathbf{\Sigma}_\ell^{1/2}\mathbf{U}_\ell^\top\mathbf{x}\right\|_2^2 \geq \left\|\mathbf{\Sigma}^{1/2}\mathbf{U}^\top\mathbf{x}\right\|_2^2 - \epsilon\,\mathrm{Tr}(\mathbf{A}).$$

[3] Each column of $\mathbf{R}$ has a single non-zero entry (set to one), corresponding to one of the $|R|$ selected columns. Formally, $\mathbf{R}_{i_t,t} = 1$ for $t = 1, \ldots, |R|$; all other entries of $\mathbf{R}$ are set to zero.

The above simple lemma is very much at the heart of our proof of Theorem 3.1 and, unlike prior work, allows us to provide provably accurate bounds for the thresholding Algorithm 3.

4 SPCA VIA A SEMIDEFINITE PROGRAMMING RELAXATION

Our third algorithm is based on the SDP relaxation of eqn. (2). Recall that solving eqn. (2) returns a PSD matrix $\mathbf{Z}^* \in \mathbb{R}^{n \times n}$ that, by the definition of the semidefinite programming relaxation, satisfies $\mathrm{Tr}(\mathbf{A}\mathbf{Z}^*) \geq \mathcal{Z}^*$, where $\mathcal{Z}^*$ is the true optimal solution of SPCA in eqn. (2).

We now need to convert the matrix $\mathbf{Z}^* \in \mathbb{R}^{n \times n}$ into a sparse vector that will be the output of our approximate SPCA algorithm *and* will satisfy certain accuracy guarantees. Towards that end, we employ a novel two-step rounding procedure. First, a critical observation is that generating a random Gaussian vector $\mathbf{g} \in \mathbb{R}^n$ and computing the vector $\mathbf{Z}^*\mathbf{g} \in \mathbb{R}^n$ results in an unbiased estimator for the trace of $(\mathbf{Z}^*)^\top \mathbf{A}\mathbf{Z}^*$ in the following sense: $\mathbb{E}\left[\mathbf{g}^\top (\mathbf{Z}^*)^\top \mathbf{A}\mathbf{Z}^*\mathbf{g}\right] = \mathrm{Tr}((\mathbf{Z}^*)^\top \mathbf{A}\mathbf{Z}^*)$.

Using von Neumann's trace inequality, we can prove that

$$\mathbb{E}\left[\mathbf{g}^\top (\mathbf{Z}^*)^\top \mathbf{A}\mathbf{Z}^*\mathbf{g}\right] = \mathrm{Tr}((\mathbf{Z}^*)^\top \mathbf{A}\mathbf{Z}^*) \geq \gamma_{Z^*} \cdot \mathrm{Tr}(\mathbf{A}\mathbf{Z}^*) \geq \gamma_Z \cdot \mathcal{Z}^*. \tag{5}$$

Here γ_{Z^*} is a constant that precisely depends on the top singular value of $\mathbf{Z}^*$, the condition number of $\mathbf{Z}^*$, and the extent to which the SDP relaxation of eqn. (2) is able to capture the original problem (see Theorem 4.1 for the exact expression of γ_{Z^*}). The above inequality implies that, at least in expectation, we could use the vector $\mathbf{Z}^*\mathbf{g}$ as a "rounding" of the output of the semidefinite programming relaxation. However, there is absolutely no guarantee that the vector $\mathbf{Z}^*\mathbf{g}$ is sparse. Thus, in order to sparsify $\mathbf{Z}^*\mathbf{g}$, we employ a separate sparsification procedure, where each entry of $\mathbf{Z}^*\mathbf{g}$ is kept (and rescaled) with probability proportional to its magnitude. This procedure is similar to the one proposed in (Fountoulakis et al., 2017) and guarantees that larger entries of $\mathbf{Z}^*\mathbf{g}$ are more likely to be kept, while smaller entries of $\mathbf{Z}^*\mathbf{g}$ are more likely to be set to zero, without too much loss in accuracy. We also note that to ensure a sufficiently high probability of success for the overall approach, we generate multiple Gaussian vectors and keep the one that maximizes the quantity $\mathbf{g}^\top (\mathbf{Z}^*)^\top \mathbf{A}\mathbf{Z}^*\mathbf{g}$. See Algorithm 4 and Algorithm 5 for a detailed presentation of our approach.

Algorithm 4 SPARSIFY

Input: $\mathbf{y} \in \mathbb{R}^n$ and sparsity parameter s.
Output: $\mathbf{z} \in \mathbb{R}^n$ with $\mathbb{E}\left[\|\mathbf{z}\|_0\right] \leq s$.

1: **for** $i = 1 \ldots n$ **do**
2: $\quad \mathbf{z}_i := \begin{cases} \frac{1}{p_i}\mathbf{y}_i, & \text{with probability } p_i = \min\left\{1, \frac{s|y_i|}{\|\mathbf{y}\|_1}\right\}, \\ 0 & \text{otherwise.} \end{cases}$

Algorithm 5 Rounding-based SPCA

Input: PSD matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, error tolerance $\epsilon > 0$, and sparsity parameter k.
Output: $\mathbf{x} \in \mathbb{R}^n$ with $\mathbb{E}\left[\|\mathbf{x}\|_0\right] = s$.

1: Let $\mathbf{Z}^*$ be the optimal solution to the relaxed SPCA problem of eqn. (2);
2: $M \leftarrow {}^{80}/_{\epsilon^2}$ and $s \leftarrow O\left(\frac{k^2 \log^{5/2}(1/\epsilon)}{\epsilon^2}\right)$; ▷See Theorem A.13 for the exact value of s.
3: Generate M random Gaussian vectors $g_1, \ldots, g_M$ in $\mathbb{R}^n$;
4: $\mathbf{y} \leftarrow \mathbf{Z}^*\mathbf{g}_j$, where $j \leftarrow \mathrm{argmax}_{i=1\ldots M}\, \mathbf{g}_i^\top (\mathbf{Z}^*)^\top \mathbf{A}\mathbf{Z}^*\mathbf{g}_i$;
5: $\mathbf{z} \leftarrow$ SPARSIFY$(\mathbf{y}, s)$;

The running time of the algorithm is dominated by the time needed to solve the semidefinite programming relaxation of eqn. (2), which, in our setting, is $\mathcal{O}\left(n^{3.5}\right)$ (Alizadeh, 1995). We do note that SDP solvers such as the one in (Alizadeh, 1995) return an additive error approximation to the optimal solution. However, the running time dependence of SDP solvers on the additive error γ is logarithimic in $1/\gamma$ and thus highly accurate approximations can be derived without a significant increase in the number of iterations of the solver. Thus, for the sake of clarity, we omit this additive error from the rest of the analysis.

Our main quality-of-approximation result for Algorithm 5 is Theorem 4.1. For simplicity of presentation, and following the lines of (Fountoulakis et al., 2017), we assume that all rows and columns

of $\mathbf{A}$ have been normalized to have unit norms. This assumption can be relaxed as in (Fountoulakis et al., 2017). In the statement of the theorem, we will use the notation $\mathbf{Z}_1$ to denote the best rank-one approximation to the matrix $\mathbf{Z}$.

Theorem 4.1 *Given a PSD matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, a sparsity parameter k, and an error tolerance $\epsilon > 0$, let $\mathbf{Z}$ be an optimal solution to the relaxed SPCA problem of eqn. (2). Assume that*

$$\mathrm{Tr}(\mathbf{A}\mathbf{Z}) \leq \alpha \, \mathrm{Tr}(\mathbf{A}\mathbf{Z}_1) \tag{6}$$

for some constant $\alpha \geq 1$. Then, Algorithm 5 outputs a vector $\mathbf{z} \in \mathbb{R}^n$ that, with probability at least $5/8$, satisfies $\mathbb{E}\left[\|\mathbf{z}\|_0\right] = \mathcal{O}\left(k^2 \log^{5/2}(1/\epsilon)/\epsilon^2\right)$, $\|\mathbf{z}\|_2 = \mathcal{O}\left(\sqrt{\log 1/\epsilon}\right)$, and

$$\mathbf{z}^\top \mathbf{A} \mathbf{z} \geq \gamma_Z (1 - \epsilon) \cdot \mathcal{Z}^* - \epsilon.$$

Here $\gamma_Z = \left(1 - \left(1 - \frac{1}{\kappa(\mathbf{Z})}\right)\left(1 - \frac{1}{\alpha}\right)\right) \sigma_1(\mathbf{Z})$ with $\sigma_1(\mathbf{Z})$ and $\kappa(\mathbf{Z})$ being the top singular value and condition number of $\mathbf{Z}$ respectively.

We now discuss the condition of eqn. (6) and the constant γ_Z. Our assumption simply says that much of the trace of the matrix $\mathbf{A}\mathbf{Z}$ should be captured by the trace of $\mathbf{A}\mathbf{Z}_1$, as quantified by the constant α. For example, if $\mathbf{Z}$ were a rank-one matrix, then the assumption would hold with $\alpha = 1$. As the trace of $\mathbf{A}\mathbf{Z}_1$ fails to approximate the trace of $\mathbf{A}\mathbf{Z}$ (which intuitively implies that the SDP relaxation of eqn. (2) did not sufficiently capture the original problem) the constant α increases and the quality of the approximation decreases. More precisely, first notice that the constant γ_Z is upper bounded by one, because $\sigma_1(\mathbf{Z}) \leq 1$ by the SDP relaxation. Second, the quality of the approximation increases as γ_Z approaches one. This happens if either the condition number of $\mathbf{Z}$ is close to one *or* if the constant α is close to one; at the same time, $\sigma_1(\mathbf{Z})$ also needs to be close to one. Clearly, these conditions are satisfied if $\mathbf{Z}$ is well approximated by $\mathbf{Z}_1$. In our experiments, we indeed observed that α is close to one and that the top singular value of $\mathbf{Z}$ is close to one, which imply that γ_Z is also close to one (Appendix, Table 6). The proof of Theorem 4.1 is delegated to Appendix A.3 (as Theorem A.13), but we outline here a summary of statements that lead to the final bound.

Lemma 4.2 *Let $\mathbf{y}$ and $\mathbf{z}$ be defined as in Algorithm 5. If $\|\mathbf{y}\|_1 \leq \alpha$ and $\|\mathbf{y}\|_2 \leq \beta$, then*

$$|\mathbf{y}^\top \mathbf{A} \mathbf{y} - \mathbf{z}^\top \mathbf{A} \mathbf{z}| \leq 2|\mathbf{y}^\top \mathbf{A}(\mathbf{y} - \mathbf{z})| + |(\mathbf{y} - \mathbf{z})^\top \mathbf{A}(\mathbf{y} - \mathbf{z})|.$$

Moreover, with probability at least $7/8$, we have $|\mathbf{y}^\top \mathbf{A}(\mathbf{y} - \mathbf{z})| \leq 4\alpha\beta/\sqrt{s}$ and $|(\mathbf{y} - \mathbf{z})^\top \mathbf{A}(\mathbf{y} - \mathbf{z})| \leq \sqrt{\frac{64\alpha^4}{s^2} + \frac{96\alpha\beta^3}{s}}$.

Lemma 4.3 *Let $M = 80/\epsilon^2$, $\alpha = k(1 + 2\sqrt{\log M})$, and $\beta = 2\sqrt{\log M}$. If the sparsity parameter s is set to $s = 450\alpha^2\beta^3/\epsilon^2$, then with probability at least $3/4$, we have $\mathbf{y}^\top \mathbf{A} \mathbf{y} \leq \mathbf{z}^\top \mathbf{A} \mathbf{z} + \epsilon$.*

Letting $\mathbf{y} = \mathbf{Z}^* \mathbf{g}$, we now conclude the proof by combining eqn. (5) with the above lemma to bound (at least in expectation) the accuracy of Algorithm 5. To get a high probability bound, we leverage a result by (Avron & Toledo, 2011) on estimating the trace of PSD matrices. This approach allows us to properly analyze step 4 of Algorithm 5, which uses multiple random Gaussian vectors to achieve measure concentration (see Appendix A.3 for details).

Finally, we can bound the ℓ_2 norm of the vector $\mathbf{z}$ of Algorithm 5 by proving that, with probability at least $3/4$, $\|\mathbf{z}\|_2 = \mathcal{O}\left(\sqrt{\log 1/\epsilon}\right)$. Notice that this slightly relaxes the requirement that $\mathbf{z}$ has unit norm; however, even for accuracy ϵ close to machine precision, $\sqrt{\log 1/\epsilon}$ is a small constant.

5 Experiments

First, in order to explore the sparsity patterns of the outputs of our algorithms and how close they are as compared to standard methods, we first apply our methods on the *pit props* data which was introduced in (Jeffers, 1967) and is a benchmark example used to test sparse PCA. It is a 13×13 correlation matrix, originally calculated from 180 observations with 13 explanatory variables. While the existing algorithms aimed to extract top 6 principal components, we restrict only to the top principal component. In particular, we are interested to apply our algorithms on the *Pit Props* matrix

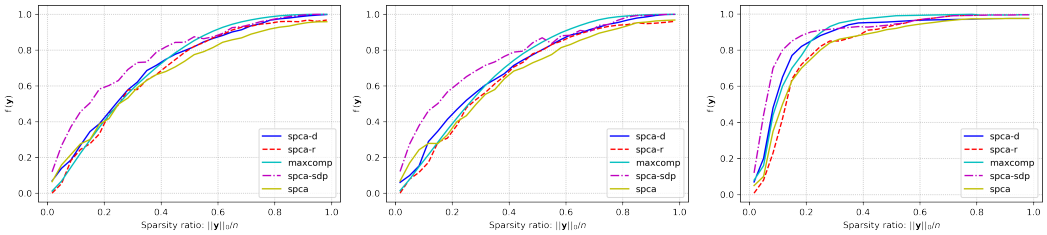

(a) CHR 1, $n = 37,493$ (b) CHR 2, $n = 40,844$ (c) Gene expression data

Fig. 1: Experimental results on real data: $f(\mathbf{y})$ vs. sparsity ratio.

	topdiam	length	moist	testsg	ovensg	ringtop	ringbut	bowmax	bowdist	whorls	clear	knots	diaknot	PVE	$\mathcal{Z}^*$
`spca-d` (Algorithm 3)	−0.420	−0.422	0	0	0	−0.296	−0.416	−0.305	−0.371	−0.394	0	0	0	30.71%	3.993
`spca-r` (Algorithm 2)	−0.448	−0.445	0	0	0	−0.379	−0.212	−0.179	−0.398	−0.471	0	0	0	28.65%	3.725
`spca-sdp` (Algorithm 5)	−0.424	−0.430	0	0	0	−0.268	−0.403	−0.313	−0.379	−0.399	0	0	0	30.74%	3.996
`spca` (Zou et al., 2006)	−0.477	−0.476	0	0	0.177	0	−0.250	−0.344	−0.416	−0.400	0	0	0	28.01%	3.641
`dspca` (d'Aspremont et al., 2007)	−0.491	−0.507	0	0	0	−0.067	−0.357	−0.234	−0.387	−0.409	0	0	0	29.39%	3.821

Table 1: Loadings, % of variace explained (PVE), and the objective function value for the first pc of the Pit Props data.

in a view to extract the top pc having a sparsity pattern similar to `spca` (Zou et al., 2006) or the `dspca` in (d'Aspremont et al., 2007), but explains more variability. Therefore, we take $k = 7$ and Table 1 shows that our methods perform very similar to `spca` or that of `dspca` with % of variace explained (PVE) uniformly better than `spca`. We also apply our algorithms on another benchmark artificial example from (Zou et al., 2006), please see Appendix B.3 for a detailed discussion.

Next, we further demonstrate the empirical performance of our algorithms on larger real-world datasets as well as on a synthetic dataset, similar to (Fountoulakis et al., 2017) (see Appendix B). We use human genetic data from the HGDP (Consortium, 2007) and HAPMAP (Li et al., 2008) (22 matrices, one for each chromosome). In addition, we also use a lung cancer gene expression dataset ($107 \times 22,215$) from (Landi et al., 2008) and a sparse document-term matrix ($2,858 \times 12,427$) created using the Text-to-Matrix Generator (TMG) (Zeimpekis & Gallopoulos, 2006) (see Appendix B).

Comparisons and metrics. We compare our Algorithm 2 (`spca-r`), Algorithm 3 (`spca-d`), and Algorithm 5 (`spca-sdp`) with the solutions returned by `spca` (Zou et al., 2006), as well as the simple `MaxComp` heuristic (Cadima & Jolliffe, 1995). We define the quantity $f(\mathbf{y}) = \mathbf{y}^\top \mathbf{A}\mathbf{y} / \|\mathbf{A}\|_2$ to measure the quality of an approximate solution $\mathbf{y} \in \mathbb{R}^n$ to the SPCA problem. Notice that $0 \leq f(\mathbf{y}) \leq 1$ for all $\mathbf{y}$ with $\|\mathbf{y}\|_2 \leq 1$. As $f(\mathbf{y})$ gets closer to one, the vector $\mathbf{y}$ captures more of the variance of the matrix $\mathbf{A}$ that corresponds to its top singular value and corresponding singular vector. Our goal is to identify a sparse vector $\mathbf{y}$ with $f(\mathbf{y}) \approx 1$. Since the outputs of our Algorithm 2 and Algorithm 5 may have norms that slightly exceed one and in order to have a fair comparison between different methods, we normalize our outputs in the same way as in (Fountoulakis et al., 2017) (Appendix B).

Results. In our experiments, for `spca-d`, `spca-r`, and `spca-sdp`, we fix the sparsity s to be equal to k, so that all algorithms return the same number of non-zero elements as `spca` and `maxcomp`. In Figures 1a-1c we evaluate the performance of the different SPCA algorithms by plotting $f(\mathbf{y})$ against $\|\mathbf{y}\|_0/n$, *i.e.,* the sparsity ratio of the output vector, on data from chromosome 1, chromosome 2, and the gene expression data. All our methods are *better than or at least comparable to* both `spca` and `maxcomp`. In particular, for the sparser output vectors, `spca-sdp` performs much better than both `spca` and `maxcomp`. However, in practice, the running time of the SDP relaxation is substantially higher than the other methods, which highlights the interesting trade-offs between the accuracy and computation discussed in (Amini & Wainwright, 2009). See Apeendix B for more experimental results.

6 Conclusion and Open Problems

We present three provably accurate, polynomial time, approximation algorithms for SPCA, without imposing restrictive assumptions on the input covariance matrix. Future directions include: *(i)* extend the proposed algorithms to handle more than one sparse singular vector by deflation or other strategies and *(ii)* explore matching lower bounds and/or improve the guarantees of Theorems 2.2, 3.1, and 4.1.

REFERENCES

Farid Alizadeh. Interior Point Methods in Semidefinite Programming with Applications to Combinatorial Optimization. *SIAM Journal on Optimization*, 5(1):13–51, 1995. 6

Arash A. Amini and Martin J. Wainwright. High-dimensional Analysis of Semidefinite Relaxations for Sparse Principal Components. *Annals of Statistics*, 37:2877–2921, 2009. 3, 8

Ei Ando, Toshio Nakata, and Masafumi Yamashita. Approximating the Longest Path Length of a Stochastic DAG by a Normal Distribution in Linear Time. *Journal of Discrete Algorithms*, 7(4): 420–438, 2009. 3

Megasthenis Asteris, Dimitris Papailiopoulos, and George N Karystinos. Sparse Principal Component of a Rank-deficient Matrix. In *2011 IEEE International Symposium on Information Theory Proceedings*, pp. 673–677, 2011. 3

Megasthenis Asteris, Dimitris Papailiopoulos, Anastasios Kyrillidis, and Alexandros G Dimakis. Sparse PCA via Bipartite Matchings. In *Advances in Neural Information Processing Systems*, pp. 766–774, 2015. 3

Haim Avron and Sivan Toledo. Randomized Algorithms for Estimating the Trace of an Implicit Symmetric Positive Semi-definite Matrix. *Journal of the ACM*, 58(2):8:1–8:34, 2011. 7, 19

Jorge Cadima and Ian T. Jolliffe. Loading and Correlations in the Interpretation of Principal Components. *Journal of Applied Statistics*, 22(2):203–214, 1995. 3, 8

Siu On Chan, Dimitris Papailliopoulos, and Aviad Rubinstein. On the Approximability of Sparse PCA. In *Proceedings of the 29th Conference on Learning Theory*, pp. 623–646, 2016. 3

International HapMap Consortium. A Second Generation Human Haplotype Map of over 3.1 Million SNPs. *Nature*, 449(7164):851, 2007. 8, 20

Alexandre d'Aspremont, Laurent El Ghaoui, Michael I. Jordan, and Gert R. G. Lanckriet. A Direct Formulation for Sparse PCA using Semidefinite Programming. *SIAM Review*, 49(3):434–448, 2007. 1, 3, 8, 21, 22

Petros Drineas and Ravi Kannan. Fast Monte-Carlo Algorithms for Approximate Matrix Multiplication. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pp. 452–461, 2001. 3

Petros Drineas and Michael W. Mahoney. RandNLA: Randomized Numerical Linear Algebra. *Communications of the ACM*, 59(6):80–90, 2016. 3

Petros Drineas and Michael W. Mahoney. *Lectures on Randomized Numerical Linear Algebra*, volume 25 of *The Mathematics of Data, IAS/Park City Mathematics Series*. American Mathematical Society, 2018. 3, 4

Petros Drineas, Ravi Kannan, and Michael W. Mahoney. Fast Monte Carlo Algorithms for Matrices II: Computing a Low-Rank Approximation to a Matrix. *SIAM Journal on Computing*, 36(1): 158–183, 2006. 3

Alexandre d'Aspremont, Francis Bach, and Laurent El Ghaoui. Optimal Solutions for Sparse Principal Component Analysis. *Journal of Machine Learning Research*, 9(Jul):1269–1294, 2008. 3

Alexandre d'Aspremont, Francis Bach, and Laurent El Ghaoui. Approximation Bounds for Sparse Principal Component Analysis. *Mathematical Programming*, 148(1-2):89–110, 2014. 3

Kimon Fountoulakis, Abhisek Kundu, Eugenia-Maria Kontopoulou, and Petros Drineas. A Randomized Rounding Algorithm for Sparse PCA. *ACM Transactions on Knowledge Discovery from Data*, 11(3):38:1–38:26, 2017. 2, 3, 6, 7, 8, 14, 15, 21

John NR Jeffers. Two case studies in the application of principal component analysis. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 16(3):225–236, 1967. 7

Ian T. Jolliffe. Rotation of principal components: Choice of Normalization Constraints. *Journal of Applied Statistics*, 22(1):29–35, 1995. 3

Ian T. Jolliffe, Nickolay T. Trendafilov, and Mudassir Uddin. A Modified Principal Component Technique Based on the LASSO. *Journal of Computational and Graphical Statistics*, 12(3): 531–547, 2003. 3

Michel Journée, Yurii Nesterov, Peter Richtárik, and Rodolphe Sepulchre. Generalized Power Method for Sparse Principal Component Analysis. *Journal of Machine Learning Research*, 11(2), 2010. 3

Volodymyr Kuleshov. Fast Algorithms for Sparse Principal Component Analysis Based on Rayleigh Quotient Iteration. In *Proceedings of the 30th International Conference on Machine Learning*, pp. 1418–1425, 2013. 3

Maria Teresa Landi, Tatiana Dracheva, Melissa Rotunno, Jonine D. Figueroa, Huaitian Liu, Abhijit Dasgupta, Felecia E. Mann, Junya Fukuoka, Megan Hames, Andrew W. Bergen, et al. Gene Expression Signature of Cigarette Smoking and Its Role in Lung Adenocarcinoma Development and Survival. *PloS one*, 3(2), 2008. 8, 20

Jun Z. Li, Devin M. Absher, Hua Tang, Audrey M. Southwick, Amanda M. Casto, Sohini Ramachandran, Howard M. Cann, Gregory S. Barsh, Marcus Feldman, Luigi L. Cavalli-Sforza, et al. Worldwide Human Relationships Inferred from Genome-Wide Patterns of Variation. *Science*, 319 (5866):1100–1104, 2008. 8, 20

Malik Magdon-Ismail. NP-Hardness and Inapproximability of Sparse PCA. *Information Processing Letters*, 126:35–38, 2017. 3

Michael W. Mahoney and P. Drineas. CUR Matrix Decompositions for Improved Data Analysis. In *Proceedings of the National Academy of Sciences*, pp. 697–702, 106 (3), 2009. 1

Baback Moghaddam, Yair Weiss, and Shai Avidan. Generalized Spectral Bounds for Sparse LDA. In *Proceedings of the 23rd International Conference on Machine learning*, pp. 641–648, 2006a. 1, 3

Baback Moghaddam, Yair Weiss, and Shai Avidan. Spectral Bounds for Sparse PCA: Exact and Greedy Algorithms. In *Advances in Neural Information Processing Systems*, pp. 915–922, 2006b. 3

Dimitris Papailiopoulos, Alexandros Dimakis, and Stavros Korokythakis. Sparse PCA through Low-rank Approximations. In *Proceedings of the 30th International Conference on Machine Learning*, pp. 747–755, 2013. 1, 3

Haipeng Shen and Jianhua Z. Huang. Sparse Principal Component Analysis via Regularized Low Rank Matrix Approximation. *Journal of Multivariate Analysis*, 99(6):1015–1034, 2008. 3

Bharath K. Sriperumbudur, David A. Torres, and Gert R.G. Lanckriet. Sparse Eigen Methods by D.C. Programming. In *Proceedings of the 24th International Conference on Machine Learning*, pp. 831–838, 2007. 3

Martin J. Wainwright. Lecture Notes, 2015. https://www.stat.berkeley.edu/~mjwain/stat210b/Chap2_TailBounds_Jan22_2015.pdf. 17

David P. Woodruff. Sketching as a Tool for Numerical Linear Algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1-2), 2014. 3

Xiao-Tong Yuan and Tong Zhang. Truncated Power Method for Sparse Eigenvalue Problems. *Journal of Machine Learning Research*, 14(Apr):899–925, 2013. 3

Dimitrios Zeimpekis and Efstratios Gallopoulos. TMG: A MATLAB Toolbox for Generating Term-Document Matrices from Text Collections. In *Grouping Multidimensional Data*, pp. 187–210. Springer, 2006. 8, 21

Hui Zou and Trevor Hastie. Regularization and Variable Selection via the Elastic Net. *Journal of the Royal Statistical Society: Series B*, 67(2):301–320, 2005. 3

Hui Zou, Trevor Hastie, and Robert Tibshirani. Sparse Principal Component Analysis. *Journal of Computational and Graphical Statistics*, 15(2):265–286, 2006. 3, 8, 21, 22, 23