# Large-Scale Sparse Principal Component Analysis with Application to Text Data


**Youwei Zhang**
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley
Berkeley, CA 94720
zyw@eecs.berkeley.edu

**Laurent El Ghaoui**
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley
Berkeley, CA 94720
elghaoui@eecs.berkeley.edu



## Abstract

Sparse PCA provides a linear combination of small number of features that maximizes variance across data. Although Sparse PCA has apparent advantages compared to PCA, such as better interpretability, it is generally thought to be computationally much more expensive. In this paper, we demonstrate the surprising fact that sparse PCA can be easier than PCA in practice, and that it can be reliably applied to very large data sets. This comes from a rigorous feature elimination pre-processing result, coupled with the favorable fact that features in real-life data typically have exponentially decreasing variances, which allows for many features to be eliminated. We introduce a fast block coordinate ascent algorithm with much better computational complexity than the existing first-order ones. We provide experimental results obtained on text corpora involving millions of documents and hundreds of thousands of features. These results illustrate how Sparse PCA can help organize a large corpus of text data in a user-interpretable way, providing an attractive alternative approach to topic models.


## 1   Introduction

The sparse Principal Component Analysis (Sparse PCA) problem is a variant of the classical PCA problem, which accomplishes a trade-off between the explained variance along a normalized vector, and the number of non-zero components of that vector.

Sparse PCA not only brings better interpretation [1], but also provides statistical regularization [2] when the number of samples is less than the number of features. Various researchers have proposed different formulations and algorithms for this problem, ranging from ad-hoc methods such as factor rotation techniques [3] and simple thresholding [4], to greedy algorithms [5, 6]. Other algorithms include SCoTLASS by [7], SPCA by [8], the regularized SVD method by [9] and the generalized power method by [10]. These algorithms are based on non-convex formulations, and may only converge to a local optimum. The $\ell_1$-norm based semidefinite relaxation DSPCA, as introduced in [1], does guarantee global convergence and as such, is an attractive alternative to local methods. In fact, it has been shown in [1, 2, 11] that simple ad-hoc methods, and the greedy, SCoTLASS and SPCA algorithms, often underperform DSPCA. However, the first-order algorithm for solving DSPCA, as developed in [1], has a computational complexity of $O(n^4\sqrt{\log n})$, with $n$ the number of


1


features, which is too high for many large-scale data sets. At first glance, this complexity estimate indicates that solving sparse PCA is much more expensive than PCA, since we can compute one principal component with a complexity of $O(n^2)$.

In this paper we show that solving DSPCA is in fact computationally easier than PCA, and hence can be applied to very large-scale data sets. To achieve that, we first view DSPCA as an approximation to a harder, cardinality-constrained optimization problem. Based on that formulation, we describe a safe feature elimination method for that problem, which leads to an often important reduction in problem size, prior to solving the problem. Then we develop a block coordinate ascent algorithm, with a computational complexity of $O(n^3)$ to solve DSPCA, which is much faster than the first-order algorithm proposed in [1]. Finally, we observe that real data sets typically allow for a dramatic reduction in problem size as afforded by our safe feature elimination result. Now the comparison between sparse PCA and PCA becomes $O(\hat{n}^3)$ v.s. $O(n^2)$ with $\hat{n} \ll n$, which can make sparse PCA surprisingly easier than PCA.

In Section 2, we review the $\ell_1$-norm based DSPCA formulation, and relate it to an approximation to the $\ell_0$-norm based formulation and highlight the safe feature elimination mechanism as a powerful pre-processing technique. We use Section 3 to present our fast block coordinate ascent algorithm. Finally, in Section 4, we demonstrate the efficiency of our approach on two large data sets, each one containing more than 100,000 features.

**Notation.**   $\mathcal{R}(Y)$ denotes the range of matrix $Y$, and $Y^\dagger$ its pseudo-inverse. The notation $\log$ refers to the extended-value function, with $\log x = -\infty$ if $x \leq 0$.

## 2   Safe Feature Elimination

**Primal problem.**   Given a $n \times n$ positive-semidefinite matrix $\Sigma$, the "sparse PCA" problem introduced in [1] is :

$$\phi = \max_Z \ \mathbf{Tr}\, \Sigma Z - \lambda \|Z\|_1 \ : \ Z \succeq 0, \ \ \mathbf{Tr}\, Z = 1 \tag{1}$$

where $\lambda \geq 0$ is a parameter encouraging sparsity. Without loss of generality we may assume that $\Sigma \succ 0$.

Problem (1) is in fact a relaxation to a PCA problem with a penalty on the cardinality of the variable:

$$\psi = \max_x \ x^T \Sigma x - \lambda \|x\|_0 \ : \ \|x\|_2 = 1 \tag{2}$$

Where $\|x\|_0$ denotes the cardinality (number of non-zero elemements) in $x$. This can be seen by first writing problem (2) as:

$$\max_Z \ \mathbf{Tr}\, \Sigma Z - \lambda \sqrt{\|Z\|_0} \ : \ Z \succeq 0, \ \ \mathbf{Tr}\, Z = 1, \mathbf{Rank}(Z) = 1$$

where $\|Z\|_0$ is the cardinality (number of non-zero elements) of $Z$. Since $\|Z\|_1 \leq \sqrt{\|Z\|_0} \|Z\|_F = \sqrt{\|Z\|_0}$, we obtain the relaxation

$$\max_Z \ \mathbf{Tr}\, \Sigma Z - \lambda \|Z\|_1 \ : \ Z \succeq 0, \mathbf{Tr}\, Z = 1, \mathbf{Rank}(Z) = 1$$

Further drop the rank constraint, leading to problem (1).

By viewing problem (1) as a convex approximation to the non-convex problem (2), we can leverage the safe feature elimination theorem first presented in [6, 12] for problem (2):

**Theorem 2.1** *Let* $\Sigma = A^T A$, *where* $A = (a_1, \ldots, a_n) \in \mathbf{R}^{m \times n}$. *We have*

$$\psi = \max_{\|\xi\|_2 = 1} \sum_{i=1}^n ((a_i^T \xi)^2 - \lambda)_+.$$

*An optimal non-zero pattern corresponds to indices* $i$ *with* $\lambda < (a_i^T \xi)^2$ *at optimum.*

We observe that the $i$-th feature is absent at optimum if $(a_i^T \xi)^2 \leq \lambda$ for every $\xi, \ \|\xi\|_2 = 1$. Hence, we can *safely* remove feature $i \in \{1, \ldots, n\}$ if

$$\Sigma_{ii} = a_i^T a_i < \lambda \tag{3}$$

2

A few remarks are in order. First, if we are interested in solving problem (1) as a relaxation to problem (2), we first calculate and rank all the feature variances, which takes $O(nm)$ and $O(n \log(n))$ respectively. Then we can safely eliminate any feature with variance less than $\lambda$. Second, the elimination criterion above is conservative. However, when looking for extremely sparse solutions, applying this safe feature elimination test with a large $\lambda$ can dramatically reduce problem size and lead to huge computational savings, as will be demonstrated empirically in Section 4. Third, in practice, when PCA is performed on large data sets, some similar variance-based criteria is routinely employed to bring problem sizes down to a manageable level. This purely heuristic practice has a rigorous interpretation in the context of *sparse* PCA, as the above theorem states explicitly the features that can be safely discarded.

## 3 Block Coordinate Ascent Algorithm

The first-order algorithm developed in [1] to solve problem (1) has a computational complexity of $O(n^4 \sqrt{\log n})$. With a theoretical convergence rate of $O(\frac{1}{\epsilon})$, the DSPCA algorithm does not converge fast in practice. In this section, we develop a block coordinate ascent algorithm with better dependence on problem size ($O(n^3)$), that in practice converges much faster.

**Failure of a direct method.** We seek to apply a "row-by-row" algorithm by which we update each row/column pair, one at a time. This algorithm appeared in the specific context of sparse covariance estimation in [13], and extended to a large class of SDPs in [14]. Precisely, it applies to problems of the form

$$\min_X \ f(X) - \beta \log \det X \ : \ L \le X \le U, \ \ X \succ 0, \tag{4}$$

where $X = X^T$ is a $n \times n$ matrix variable, $L, U$ impose component-wise bounds on $X$, $f$ is convex, and $\beta > 0$.

However, if we try to update the row/columns of $Z$ in problem (1), the trace constraint will imply that we never modify the diagonal elements of $Z$. Indeed at each step, we update only one diagonal element, and it is entirely fixed given all the other diagonal elements. The row-by-row algorithm does not directly work in that case, nor in general for SDPs with equality constraints. The authors in [14] propose an augmented Lagrangian method to deal with such constraints, with a complication due to the choice of appropriate penalty parameters. In our case, we can apply a technique resembling the augmented Lagrangian technique, without this added complication. This is due to the homogeneous nature of the objective function and of the conic constraint. Thanks to the feature elimination result (Thm. 2.1), we can always assume without loss of generality that $\lambda < \sigma_{\min}^2 := \min_{1 \le i \le n} \Sigma_{ii}$.

**Direct augmented Lagrangian technique.** We can express problem (1) as

$$\frac{1}{2}\phi^2 = \max_X \ \mathbf{Tr} \, \Sigma X - \lambda \|X\|_1 - \frac{1}{2}(\mathbf{Tr} \, X)^2 \ : \ X \succeq 0. \tag{5}$$

This expression results from the change of variable $X = \gamma Z$, with $\mathbf{Tr} \, Z = 1$, and $\gamma \ge 0$. Optimizing over $\gamma \ge 0$, and exploiting $\phi > 0$ (which comes from our assumption that $\lambda < \sigma_{\min}^2$), leads to the result, with the optimal scaling factor $\gamma$ equal to $\phi$. An optimal solution $Z^*$ to (1) can be obtained from an optimal solution $X^*$ to the above, via $Z^* = X^*/\phi$. (In fact, we have $Z^* = X^*/\mathbf{Tr}(X^*)$.)

To apply the row-by-row method to the above problem, we need to consider a variant of it, with a strictly convex objective. That is, we address the problem

$$\max_X \ \mathbf{Tr} \, \Sigma X - \lambda \|X\|_1 - \frac{1}{2}(\mathbf{Tr} \, X)^2 + \beta \log \det X, \ : \ X \succ 0, \tag{6}$$

where $\beta > 0$ is a penalty parameter. SDP theory ensures that if $\beta = \epsilon/n$, then a solution to the above problem is $\epsilon$-suboptimal for the original problem [15].

**Optimizing over one row/column.** Without loss of generality, we consider the problem of updating the last row/column of the matrix variable $X$. Partition the latter and the covariance matrix $S$ as

$$X = \left( \begin{array}{cc} Y & y \\ y^T & x \end{array} \right), \ \ \Sigma = \left( \begin{array}{cc} S & s \\ s^T & \sigma \end{array} \right),$$

3

where $Y, S \in \mathbf{R}^{(n-1)\times(n-1)}$, $y, s \in \mathbf{R}^{n-1}$, and $x, \sigma \in \mathbf{R}$. We are considering the problem above, where $Y$ is fixed, and $(y, x) \in \mathbf{R}^n$ is the variable. We use the notation $t := \mathbf{Tr}\, Y$.

The conic constraint $X \succ 0$ translates as $y^T Y^\dagger y \le x$, $y \in \mathcal{R}(Y)$, where $\mathcal{R}(Y)$ is the range of the matrix $Y$. We obtain the sub-problem

$$\psi := \max_{x,y} \left( \begin{array}{c} 2(y^T s - \lambda\|y\|_1) + (\sigma - \lambda)x - \frac{1}{2}(t+x)^2 \\ +\beta\log(x - y^T Y^\dagger y) \end{array} \right) \; : \; y \in \mathcal{R}(Y). \qquad (7)$$

**Simplifying the sub-problem.** We can simplify the above problem, in particular, avoid the step of forming the pseudo-inverse of $Y$, by taking the dual of problem (7).

Using the conjugate relation, valid for every $\eta > 0$:

$$\log \eta + 1 = \min_{z>0} \; z\eta - \log z,$$

and with $f(x) := (\sigma - \lambda)x - \frac{1}{2}(t+x)^2$, we obtain

$$\begin{aligned}
\psi + \beta &= \max_{y \in \mathcal{R}(Y)} 2(y^T s - \lambda\|y\|_1) + f(x) + \beta \min_{z>0} \left( z(x - y^T Y^\dagger y) - \log z \right) \\
&= \min_{z>0} \max_{y \in \mathcal{R}(Y)} 2(y^T s - \lambda\|y\|_1 - \beta z y^T Y^\dagger y) + \max_x (f(x) + \beta z x) - \beta \log z \\
&= \min_{z>0} h(z) + 2g(z)
\end{aligned}$$

where, for $z > 0$, we define

$$\begin{aligned}
h(z) &:= -\beta \log z + \max_x (f(x) + \beta z x) \\
&= -\beta \log z + \max_x \left( (\sigma - \lambda + \beta z)x - \frac{1}{2}(t+x)^2 \right) \\
&= -\frac{1}{2}t^2 - \beta \log z + \max_x \left( (\sigma - \lambda - t + \beta z)x - \frac{1}{2}x^2 \right) \\
&= -\frac{1}{2}t^2 - \beta \log z + \frac{1}{2}(\sigma - \lambda - t + \beta z)^2
\end{aligned}$$

with the following relationship at optimum:

$$x = \sigma - \lambda - t + \beta z. \qquad (8)$$

In addition,

$$\begin{aligned}
g(z) &:= \max_{y \in \mathcal{R}(Y)} y^T s - \lambda\|y\|_1 - \frac{\beta z}{2}(y^T Y^\dagger y) \\
&= \max_{y \in \mathcal{R}(Y)} y^T s + \min_{v\,:\,\|v\|_\infty \le \lambda} y^T v - \frac{\beta z}{2}(y^T Y^\dagger y) \\
&= \min_{v\,:\,\|v\|_\infty \le \lambda} \max_{y \in \mathcal{R}(Y)} \left( y^T(s+v) - \frac{\beta z}{2}(y^T Y^\dagger y) \right) \\
&= \min_{u\,:\,\|u-s\|_\infty \le \lambda} \max_{y \in \mathcal{R}(Y)} \left( y^T u - \frac{\beta z}{2}(y^T Y^\dagger y) \right) \\
&= \min_{u\,:\,\|u-s\|_\infty \le \lambda} \frac{1}{2\beta z} u^T Y u.
\end{aligned}$$

with the following relationship at optimum:

$$y = \frac{1}{\beta z} Y u. \qquad (9)$$

Putting all this together, we obtain the dual of problem (7): with $\psi' := \psi + \beta + \frac{1}{2}t^2$, and $c := \sigma - \lambda - t$, we have

$$\psi' = \min_{u,z} \frac{1}{\beta z} u^T Y u - \beta \log z + \frac{1}{2}(c + \beta z)^2 \; : \; z > 0, \; \|u - s\|_\infty \le \lambda.$$

Since $\beta$ is small, we can avoid large numbers in the above, with the change of variable $\tau = \beta z$:

$$\psi' - \beta \log \beta = \min_{u,\tau} \frac{1}{\tau} u^T Y u - \beta \log \tau + \frac{1}{2}(c + \tau)^2 \; : \; \tau > 0, \; \|u - s\|_\infty \le \lambda. \qquad (10)$$

4

**Solving the sub-problem.** Problem (10) can be further decomposed into two stages.

First, we solve the box-constrained QP

$$R^2 := \min_u \ u^T Y u \ : \ \|u - s\|_\infty \leq \lambda, \tag{11}$$

using a simple coordinate descent algorithm to exploit sparsity of $Y$. Without loss of generality, we consider the problem of updating the first coordinate of $u$. Partition $u$, $Y$ and $s$ as

$$u = \begin{pmatrix} \eta \\ \hat{u} \end{pmatrix}, \quad Y = \begin{pmatrix} y_1 & \hat{y}^T \\ \hat{y} & \hat{Y} \end{pmatrix}, \quad s = \begin{pmatrix} s_1 \\ \hat{s} \end{pmatrix},$$

Where, $\hat{Y} \in \mathbf{R}^{(n-2)\times(n-2)}$, $\hat{u}, \hat{y}, \hat{s} \in \mathbf{R}^{n-2}$, $y_1, s_1 \in \mathbf{R}$ are all fixed, while $\eta \in \mathbf{R}$ is the variable. We obtain the subproblem

$$\min_\eta \ y_1 \eta^2 + (2\hat{y}^T\hat{u})\eta \ : \ \|\eta - s_1\| \leq \lambda \tag{12}$$

for which we can solve for $\eta$ analytically using the formula given below.

$$\eta = \begin{cases} -\frac{\hat{y}^T\hat{u}}{y_1} & \text{if } \|s_1 + \frac{\hat{y}^T\hat{u}}{y_1}\| \leq \lambda, \ y_1 > 0, \\ s_1 - \lambda & \text{if } -\frac{\hat{y}^T\hat{u}}{y_1} < s_1 - \lambda, \ y_1 > 0 \text{ or if } \hat{y}^T\hat{u} > 0, \ y_1 = 0, \\ s_1 + \lambda & \text{if } -\frac{\hat{y}^T\hat{u}}{y_1} > s_1 + \lambda, \ y_1 > 0 \text{ or if } \hat{y}^T\hat{u} <= 0, \ y_1 = 0. \end{cases} \tag{13}$$

Next, we set $\tau$ by solving the one-dimensional problem:

$$\min_{\tau > 0} \ \frac{R^2}{\tau} - \beta \log \tau + \frac{1}{2}(c + \tau)^2.$$

The above can be reduced to a bisection problem over $\tau$, or by solving a polynomial equation of degree 3.

**Obtaining the primal variables.** Once the above problem is solved, we can obtain the primal variables $y, x$, as follows. Using formula (9), with $\beta z = \tau$, we set $y = \frac{1}{\tau} Y u$. For the diagonal element $x$, we use formula (8): $x = c + \tau = \sigma - \lambda - t + \tau$.

**Algorithm summary.** We summarize the above derivations in Algorithm 1. Notation: for any symmetric matrix $A \in \mathbf{R}^{n \times n}$, let $A_{\backslash i \backslash j}$ denote the matrix produced by removing row $i$ and column $j$. Let $A_j$ denote column $j$ (or row $j$) with the diagonal element $A_{jj}$ removed.

**Convergence and complexity.** Our algorithm solves DSPCA by first casting it to problem (6), which is in the general form (4). Therefore, the convergence result from [14] readily applies and hence every limit point that our block coordinate ascent algorithm converges to is the global optimizer. The simple coordinate descent algorithm solving problem (11) only involves a vector product and can take sparsity in $Y$ easily. To update each column/row takes $O(n^2)$ and there are $n$ such columns/rows in total. Therefore, our algorithm has a computational complexity of $O(Kn^3)$, where $K$ is the number of sweeps through columns. In practice, $K$ is fixed at a number independent of problem size (typically $K = 5$). Hence our algorithm has better dependence on the problem size compared to $O(n^4\sqrt{\log n})$ required of the first order algorithm developed in [1].

Fig 1 shows that our algorithm converges much faster than the first order algorithm. On the left, both algorithms are run on a covariance matrix $\Sigma = F^T F$ with $F$ Gaussian. On the right, the covariance matrix comes from a "spiked model" similar to that in [2], with $\Sigma = uu^T + VV^T/m$, where $u \in \mathbf{R}^n$ is the true sparse leading eigenvector, with $\mathbf{Card}(u) = 0.1n$, $V \in \mathbf{R}^{n \times m}$ is a noise matrix with $V_{ij} \sim \mathcal{N}(0, 1)$ and $m$ is the number of observations.

## 4 Numerical Examples

In this section, we analyze two publicly available large data sets, the NYTimes news articles data and the PubMed abstracts data, available from the UCI Machine Learning Repository [16]. Both

5

---
**Algorithm 1** Block Coordinate Ascent Algorithm
---
**Input:** The covariance matrix $\Sigma$, and a parameter $\rho > 0$.

1: Set $X^{(0)} = I$
2: **repeat**
3: &nbsp;&nbsp;&nbsp;&nbsp;**for** $j = 1$ to $n$ **do**
4: &nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;Let $X^{(j-1)}$ denote the current iterate. Solve the box-constrained quadratic program

$$R^2 := \min_u \; u^T X^{(j-1)}_{\backslash j \backslash j} u \; : \; \|u - \Sigma_j\|_\infty \leq \lambda$$

&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;using the coordinate descent algorithm
5: &nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;Solve the one-dimensional problem

$$\min_{\tau > 0} \; \frac{R^2}{\tau} - \beta \log \tau + \frac{1}{2}(\Sigma_{jj} - \lambda - \mathbf{Tr}\, X^{(j-1)}_{\backslash j \backslash j} + \tau)^2$$

&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;using a bisection method, or by solving a polynomial equation of degree 3.
6: &nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;&nbsp;First set $X^{(j)}_{\backslash j \backslash j} = X^{(j-1)}_{\backslash j \backslash j}$, and then set both $X^{(j)}$'s column $j$ and row $j$ using

$$X^{(j)}_j = \frac{1}{\tau} X^{(j-1)}_{\backslash j \backslash j} u$$

$$X^{(j)}_{jj} = \Sigma_{jj} - \lambda - \mathbf{Tr}\, X^{(j-1)}_{\backslash j \backslash j} + \tau$$

7: &nbsp;&nbsp;&nbsp;&nbsp;**end for**
8: &nbsp;&nbsp;&nbsp;&nbsp;Set $X^{(0)} = X^{(n)}$
9: **until** convergence
---
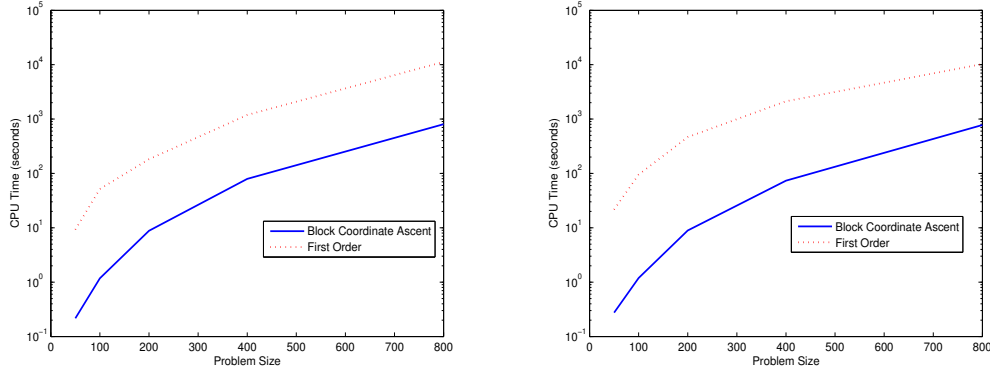
**Figure 1:** Speed comparisons between Block Coordinate Ascent and First-Order

text collections record word occurrences in the form of bag-of-words. The NYTtimes text collection contains $300,000$ articles and has a dictionary of $102,660$ unique words, resulting in a file of size 1 GB. The even larger PubMed data set has $8,200,000$ abstracts with $141,043$ unique words in them, giving a file of size 7.8 GB. These data matrices are so large that we cannot even load them into memory all at once, which makes even the use of classical PCA difficult. However with the pre-processing technique presented in Section 2 and the block coordinate ascent algorithm developed in Section 3, we are able to perform sparse PCA analysis of these data, also thanks to the fact that variances of words decrease drastically when we rank them as shown in Fig 2. Note that the feature elimination result only requires the computation of each feature's variance, and that this task is easy to parallelize.

By doing sparse PCA analysis of these text data, we hope to find interpretable principal components that can be used to summarize and explore the large corpora. Therefore, we set the target cardinality for each principal component to be 5. As we run our algorithm with a coarse range of $\lambda$ to search for

6
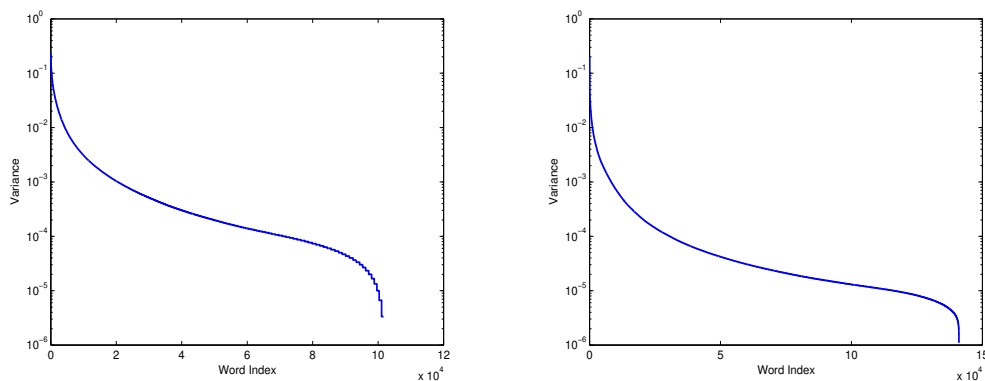
**Figure 2:** Sorted variances of 102,660 words in NYTimes (left) and 141,043 words in PubMed (right)

a solution with the given cardinality, we might end up accepting a solution with cardinality close, but not necessarily equal to, 5, and stop there to save computational time.

The top 5 sparse principal components are shown in Table 1 for NYTimes and in Table 2 for PubMed. Clearly the first principal component for NYTimes is about business, the second one about sports, the third about U.S., the fourth about politics and the fifth about education. Bear in mind that the NYTimes data from UCI Machine Learning Repository "have no class labels, and for copyright reasons no filenames or other document-level metadata" [16]. The sparse principal components still unambiguously identify and perfectly correspond to the topics used by *The New York Times* itself to classify articles on its own website.

**Table 1:** Words associated with the top 5 sparse principal components in NYTimes

| 1st PC (6 words) | 2nd PC (5 words) | 3rd PC (5 words) | 4th PC (4 words) | 5th PC (4 words) |
|---|---|---|---|---|
| million | point | official | president | school |
| percent | play | government | campaign | program |
| business | team | united_states | bush | children |
| company | season | u_s | administration | student |
| market | game | attack | | |
| companies | | | | |

After the pre-processing steps, it takes our algorithm around 20 seconds to search for a range of $\lambda$ and find one sparse principal component with the target cardinality (for the NYTimes data in our current implementation on a MacBook laptop with 2.4 GHz Intel Core 2 Duo processor and 2 GB memory).

**Table 2:** Words associated with the top 5 sparse principal components in PubMed

| 1st PC (5 words) | 2nd PC (5 words) | 3rd PC (5 words) | 4th PC (4 words) | 5th PC (4 words) |
|---|---|---|---|---|
| patient | effect | human | tumor | year |
| cell | level | expression | mice | infection |
| treatment | activity | receptor | cancer | age |
| protein | concentration | binding | maligant | children |
| disease | rat | | carcinoma | child |

A surprising finding is that the safe feature elimination test, combined with the fact that word variances decrease rapidly, enables our block coordinate ascent algorithm to work on covariance matrices of order at most $n = 500$, instead of the full order ($n = 102660$) covariance matrix for NYTimes, so as to find a solution with cardinality of around 5. In the case of PubMed, our algorithm only needs to work on covariance matrices of order at most $n = 1000$, instead of the full order ($n = 141,043$)

7

covariance matrix. Thus, at values of the penalty parameter $\lambda$ that target cardinality of 5 commands, we observe a dramatic reduction in problem sizes, about $150 \sim 200$ times smaller than the original sizes respectively. This motivates our conclusion that sparse PCA is in a sense, easier than PCA itself.

## 5   Conclusion

The safe feature elimination result, coupled with a fast block coordinate ascent algorithm, allows to solve sparse PCA problems for very large scale, real-life data sets. The overall method works especially well when the target cardinality of the result is small, which is often the case in applications where interpretability by a human is key. The algorithm we proposed has better computational complexity, and in practice converges much faster than, the first-order algorithm developed in [1]. Our experiments on text data also show that the sparse PCA can be a promising approach towards summarizing and organizing a large text corpus.

## References


[1] A. d'Aspremont, L. El Ghaoui, M. Jordan, and G. Lanckriet. A direct formulation of sparse PCA using semidefinite programming. *SIAM Review*, 49(3), 2007.

[2] A.A. Amini and M. Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. *The Annals of Statistics*, 37(5B):2877–2921, 2009.

[3] I. T. Jolliffe. Rotation of principal components: choice of normalization constraints. *Journal of Applied Statistics*, 22:29–35, 1995.

[4] J. Cadima and I. T. Jolliffe. Loadings and correlations in the interpretation of principal components. *Journal of Applied Statistics*, 22:203–214, 1995.

[5] B. Moghaddam, Y. Weiss, and S. Avidan. Spectral bounds for sparse PCA: exact and greedy algorithms. *Advances in Neural Information Processing Systems*, 18, 2006.

[6] A. d'Aspremont, F. Bach, and L. El Ghaoui. Optimal solutions for sparse principal component analysis. *Journal of Machine Learning Research*, 9:1269–1294, 2008.

[7] I. T. Jolliffe, N.T. Trendafilov, and M. Uddin. A modified principal component technique based on the LASSO. *Journal of Computational and Graphical Statistics*, 12:531–547, 2003.

[8] H. Zou, T. Hastie, and R. Tibshirani. Sparse Principal Component Analysis. *Journal of Computational & Graphical Statistics*, 15(2):265–286, 2006.

[9] Haipeng Shen and Jianhua Z. Huang. Sparse principal component analysis via regularized low rank matrix approximation. *J. Multivar. Anal.*, 99:1015–1034, July 2008.

[10] M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre. Generalized power method for sparse principal component analysis. *arXiv:0811.4724*, 2008.

[11] Y. Zhang, A. d'Aspremont, and L. El Ghaoui. Sparse PCA: Convex relaxations, algorithms and applications. In M. Anjos and J.B. Lasserre, editors, *Handbook on Semide nite, Cone and Polynomial Optimization: Theory, Algorithms, Software and Applications*. Springer, 2011. To appear.

[12] L. El Ghaoui. On the quality of a semidefinite programming bound for sparse principal component analysis. arXiv:math/060144, February 2006.

[13] O.Banerjee, L. El Ghaoui, and A. d'Aspremont. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *Journal of Machine Learning Research*, 9:485–516, March 2008.

[14] Zaiwen Wen, Donald Goldfarb, Shiqian Ma, and Katya Scheinberg. Row by row methods for semidefinite programming. Technical report, Dept of IEOR, Columbia University, 2009.

[15] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.

[16] A. Frank and A. Asuncion. UCI machine learning repository, 2010.



8