

2008/70



Generalized power method for sparse principal
component analysis

Michel Journée, Yurii Nesterov,
Peter Richtárik and Rodolphe Sepulchre

CORE

Voie du Roman Pays 34

B-1348 Louvain-la-Neuve, Belgium.

Tel (32 10) 47 43 04

Fax (32 10) 47 43 01

E-mail: corestat-library@uclouvain.be

<http://www.uclouvain.be/en-44508.html>

CORE DISCUSSION PAPER
2008/70

Generalized power method for sparse principal component analysis

Michel JOURNÉE¹, Yurii NESTEROV²,
Peter RICHTÁRIK³ and Rodolphe SEPULCHRE⁴

November 2008

Abstract

In this paper we develop a new approach to sparse principal component analysis (sparse PCA). We propose two single-unit and two block optimization formulations of the sparse PCA problem, aimed at extracting a single sparse dominant principal component of a data matrix, or more components at once, respectively. While the initial formulations involve nonconvex functions, and are therefore computationally intractable, we rewrite them into the form of an optimization program involving maximization of a convex function on a compact set. The dimension of the search space is decreased enormously if the data matrix has many more columns (variables) than rows. We then propose and analyze a simple gradient method suited for the task. It appears that our algorithm has best convergence properties in the case when either the objective function or the feasible set are strongly convex, which is the case with our single-unit formulations and can be enforced in the block case. Finally, we demonstrate numerically on a set of random and gene expression test problems that our approach outperforms existing algorithms both in quality of the obtained solution and in computational speed.

Keywords: sparse PCA, power method, gradient ascent, strongly convex sets, block algorithms.

¹ Department of Electrical Engineering and Computer Science, Université de Liège, B-4000 Liège, Belgium.
E-mail: m.journee@ulg.ac.be. M. Journée is a research fellow of the Belgian National Fund for Scientific Research (FNRS).

² Université catholique de Louvain, CORE and INMA, B-1348 Louvain-la-Neuve, Belgium.
E-mail: yurii.nesterov@uclouvain.be. This author is also member of ECORE, the newly created association between CORE and ECARES.

³ Université catholique de Louvain, CORE and INMA, B-1348 Louvain-la-Neuve, Belgium.
E-mail: peter.richtarik@uclouvain.be.

⁴ Department of Electrical Engineering and Computer Science, Université de Liège, B-4000 Liège, Belgium.
E-mail: r.sepulchre@ulg.ac.be.

The research results presented in this paper have been supported by a grant "Action de recherche concertée ARC 04/09-315" from the "Direction de la recherche scientifique – Communauté française de Belgique".

This paper presents research results of the Belgian Program on Interuniversity Poles of Attraction initiated by the Belgian State, Prime Minister's Office, Science Policy Programming. The scientific responsibility is assumed by the authors.

1 Introduction

Principal component analysis (PCA) is a well established tool for making sense of high dimensional data by reducing it to a smaller dimension. It has applications virtually in all areas of science—machine learning, image processing, engineering, genetics, neurocomputing, chemistry, meteorology, control theory, computer networks—to name just a few—where large data sets are encountered. It is important that having reduced dimension, the essential characteristics of the data are retained. If $A \in \mathbf{R}^{p \times n}$ is a matrix encoding p samples of n variables, with n being large, PCA aims at finding a few linear combinations of these variables, called *principal components*, which point in orthogonal directions explaining as much of the variance in the data as possible. If the variables contained in the columns of A are centered, then the classical PCA can be written in terms of the scaled *sample covariance matrix* $\Sigma = AA^T$ as follows:

$$\text{Find } z^* = \arg \max_{z^T z \leq 1} z^T \Sigma z. \quad (1.1)$$

Extracting one component amounts to computing the dominant eigenvector of Σ (or, equivalently, dominant right singular vector of A). Full PCA involves the computation of the singular value decomposition (SVD) of A . Principal components are, in general, combinations of all the input variables, i.e. the *loading vector* z^* is not expected to have many zero coefficients. In most applications, however, the original variables have concrete physical meaning and PCA then appears especially interpretable if the extracted components are composed only from a small number of the original variables. In the case of gene expression data, for instance, each variable represents the expression level of a particular gene. A good analysis tool for biological interpretation should be capable to highlight “simple” structures in the genome—structures expected to involve a few genes only—that explain a significant amount of the specific biological processes encoded in the data. Components that are linear combinations of a small number of variables are, quite naturally, usually easier to interpret. It is clear, however, that with this additional goal, some of the explained variance has to be sacrificed. The objective of *sparse principal component analysis* (sparse PCA) is to find a reasonable *trade-off* between these conflicting goals. One would like to explain *as much* variability in the data as possible, using components constructed from *as few* variables as possible. This is the classical trade-off between *statistical fidelity* and *interpretability*.

For about a decade, sparse PCA has been a topic of active research. Historically, the first suggested approaches were based on ad-hoc methods involving post-processing of the components obtained from classical PCA. For example, Jolliffe [11] considers using various rotation techniques to find sparse loading vectors in the subspace identified by PCA. Cadima et al. [5] propose to simply set to zero the PCA loadings which are in absolute value smaller than some threshold constant.

In recent years, more involved approaches have been put forward—approaches that consider the conflicting goals of explaining variability and achieving representation sparsity simultaneously. These methods usually cast the sparse PCA problem in the form of an optimization program, aiming at maximizing explained variance penalized for the number of non-zero loadings. For instance, the SCoTLASS algorithm proposed by Jolliffe et al. [12] aims at maximizing the Rayleigh quotient of the covariance matrix of the data under the ℓ_1 -norm based Lasso penalty [21]. Zou et al. [24] formulate sparse PCA as a regression-type optimization problem and impose

the Lasso penalty on the regression coefficients. d’Aspremont et al. [7] in their DSPCA algorithm exploit convex optimization tools to solve a convex relaxation of the sparse PCA problem. Shen and Huang [18] adapt the singular value decomposition (SVD) to compute low-rank matrix approximations of the data matrix under various sparsity-inducing penalties. Greedy methods, which are typical for combinatorial problems, have been investigated by Moghaddam et al. [14]. Finally, d’Aspremont et al. [6] propose a greedy heuristic accompanied with a certificate of optimality.

In many applications, several components need to be identified. The traditional approach consists of incorporating an existing single-unit algorithm in a deflation scheme, and computing the desired number of components sequentially (see, e.g., d’Aspremont et al. [7]). In the case of Rayleigh quotient maximization it is well-known that computing several components at once instead of computing them one-by-one by deflation with the classical power method might present better convergence whenever the largest eigenvalues of the underlying matrix are close to each other (see, e.g., Parlett [16]). Therefore, block approaches for sparse PCA are expected to be more efficient on ill-posed problems.

In this paper we consider two single-unit (Section 2.1 and 2.3) and two block formulations (Section 2.3 and 2.4) of sparse PCA, aimed at extracting m sparse principal components, with $m = 1$ in the former case and $p \geq m > 1$ in the latter. Each of these two groups comes in two variants, depending on the type of penalty we use to enforce sparsity—either ℓ_1 or ℓ_0 (cardinality).¹ While our basic formulations involve maximization of a *nonconvex* function on a space of dimension involving n , we construct *reformulations* that cast the problem into the form of maximization of a *convex* function on the unit Euclidean sphere in \mathbf{R}^p (in the $m = 1$ case) or the *Stiefel manifold*² in $\mathbf{R}^{p \times m}$ (in the $m > 1$ case). The advantage of the reformulation becomes apparent when trying to solve problems with many variables ($n \gg p$), since we manage to avoid searching a space of large dimension. At the same time, due to the convexity of the new cost function we are able to propose and *analyze* the iteration-complexity of a simple gradient-type scheme, which appears to be well suited for problems of this form. In particular, we study (Section 3) a first-order method for solving an optimization problem of the form

$$f^* = \max_{x \in \mathcal{Q}} f(x), \tag{P}$$

where \mathcal{Q} is a compact subset of a finite-dimensional vector space and f is convex. It appears that our method has best theoretical convergence properties when either f or \mathcal{Q} are strongly convex, which is the case in the single unit case (unit ball is strongly convex) and can be enforced in the block case by adding a strongly convex regularizing term to the objective function, constant on the feasible set. We do not, however, prove any results concerning the quality of the obtained solution. Even the goal of obtaining a local maximizer is in general unattainable, and we must be content with convergence to a stationary point.

In the particular case when \mathcal{Q} is the unit Euclidean ball in \mathbf{R}^p and $f(x) = x^T C x$ for some $p \times p$ symmetric positive definite matrix C , our gradient scheme specializes to the *power method*, which aims at maximizing the *Rayleigh quotient*

$$R(x) = \frac{x^T C x}{x^T x}$$

¹Our single-unit cardinality-penalized formulation is identical to that of d’Aspremont et al. [6].

²Stiefel manifold is the set of rectangular matrices with orthonormal columns.

and thus at computing the largest eigenvalue, and the corresponding eigenvector, of C .

By applying the proposed general gradient scheme to our sparse PCA reformulations of the form (P), we obtain algorithms (Section 4) with per-iteration computational cost $\mathcal{O}(npm)$.

We demonstrate on random Gaussian (Section 5.1) and gene expression data related to breast cancer (Section 5.2) that our methods are very efficient in practice. While achieving a balance between the explained variance and sparsity which is the same as or superior to the existing methods, they are faster, often converging before some of the other algorithms manage to initialize. Additionally, in the case of gene expression data our approach seems to extract components with strongest biological content.

Notation. For convenience of the reader, and at the expense of redundancy, some of the less standard notation below is also introduced at the appropriate place in the text where it is used. Parameters $m \leq p \leq n$ are actual values of dimensions of spaces used in the paper. In the definitions below, we use these actual values (i.e. n, p and m) if the corresponding object we define is used in the text exclusively with them; otherwise we make use of the dummy variables k (representing p or n in the text) and l (representing m, p or n in the text).

We will work with vectors and matrices of various sizes ($\mathbf{R}^k, \mathbf{R}^{k \times l}$). Given a vector $y \in \mathbf{R}^k$, its i^{th} coordinate is denoted by y_i . For a matrix $Y \in \mathbf{R}^{k \times l}$, y_i is the i^{th} column of Y and y_{ij} is the element of Y at position (i, j) .

By \mathbf{E} we refer to a finite-dimensional vector space; \mathbf{E}^* is its conjugate space, i.e. the space of all linear functionals on \mathbf{E} . By $\langle s, x \rangle$ we denote the action of $s \in \mathbf{E}^*$ on $x \in \mathbf{E}$. For a self-adjoint positive definite linear operator $G : \mathbf{E} \rightarrow \mathbf{E}^*$ we define a pair of norms on \mathbf{E} and \mathbf{E}^* as follows

$$\begin{aligned} \|x\| &\stackrel{\text{def}}{=} \langle Gx, x \rangle^{1/2}, \quad x \in \mathbf{E}, \\ \|s\|_* &\stackrel{\text{def}}{=} \langle s, G^{-1}s \rangle^{1/2}, \quad s \in \mathbf{E}^*. \end{aligned} \tag{1.2}$$

Although the theory in Section 3 is developed in this general setting, the sparse PCA applications considered in this paper require either the choice $\mathbf{E} = \mathbf{E}^* = \mathbf{R}^p$ (see Section 3.3 and problems (2.6) and (2.12) in Section 2) or $\mathbf{E} = \mathbf{E}^* = \mathbf{R}^{p \times m}$ (see Section 3.4 and problems (2.16) and (2.20) in Section 2). In both cases we will let G be the corresponding identity operator for which we obtain

$$\begin{aligned} \langle x, y \rangle &= \sum_i x_i y_i, \quad \|x\| = \langle x, x \rangle^{1/2} = \left(\sum_i x_i^2 \right)^{1/2} = \|x\|_2, \quad x, y \in \mathbf{R}^p, \text{ and} \\ \langle X, Y \rangle &= \text{Tr } X^T Y, \quad \|X\| = \langle X, X \rangle^{1/2} = \left(\sum_{ij} x_{ij}^2 \right)^{1/2} = \|X\|_F, \quad X, Y \in \mathbf{R}^{p \times m}. \end{aligned}$$

Thus in the vector setting we work with the *standard Euclidean norm* and in the matrix setting with the *Frobenius norm*. The symbol Tr denotes the trace of its argument.

Furthermore, for $z \in \mathbf{R}^n$ we write $\|z\|_1 = \sum_i |z_i|$ (ℓ_1 norm) and by $\|z\|_0$ (ℓ_0 “norm”) we refer to the number of nonzero coefficients, or *cardinality*, of z . By \mathbf{S}^p we refer to the space of all $p \times p$ symmetric matrices; \mathbf{S}_+^p (resp. \mathbf{S}_{++}^p) refers to the positive semidefinite (resp. definite)

cone. Eigenvalues of matrix Y are denoted by $\lambda_i(Y)$, largest eigenvalue by $\lambda_{\max}(Y)$. Analogous notation with the symbol σ refers to singular values.

By $\mathcal{B}^k = \{y \in \mathbf{R}^k \mid y^T y \leq 1\}$ (resp. $\mathcal{S}^k = \{y \in \mathbf{R}^k \mid y^T y = 1\}$) we refer to the unit Euclidean ball (resp. sphere) in \mathbf{R}^k . If we write \mathcal{B} and \mathcal{S} , then these are the corresponding objects in \mathbf{E} . The space of $n \times m$ matrices with unit-norm columns will be denoted by

$$[\mathcal{S}^n]^m = \{Y \in \mathbf{R}^{n \times m} \mid \text{Diag}(Y^T Y) = I_m\},$$

where $\text{Diag}(\cdot)$ represents the diagonal matrix obtained by extracting the diagonal of the argument. *Stiefel manifold* is the set of rectangular matrices of fixed size with orthonormal columns:

$$\mathcal{S}_m^p = \{Y \in \mathbf{R}^{p \times m} \mid Y^T Y = I_m\}.$$

For $t \in \mathbf{R}$ we will further write $\text{sign}(t)$ for the sign of the argument and $t_+ = \max\{0, t\}$.

2 Some formulations of the sparse PCA problem

In this section we propose four formulations of the sparse PCA problem, all in the form of the general optimization framework (P). The first two deal with the single-unit sparse PCA problem and the remaining two are their generalizations to the block case.

2.1 Single-unit sparse PCA via ℓ_1 -penalty

Let us consider the optimization problem

$$\phi_{\ell_1}(\gamma) \stackrel{\text{def}}{=} \max_{z \in \mathcal{B}^n} \sqrt{z^T \Sigma z} - \gamma \|z\|_1, \quad (2.1)$$

with sparsity-controlling parameter $\gamma \geq 0$ and sample covariance matrix $\Sigma = A^T A$.

The solution $z^*(\gamma)$ of (2.1) in the case $\gamma = 0$ is equal to the right singular vector corresponding to $\sigma_{\max}(A)$, the largest singular value of A . It is the first principal component of the data matrix A . The optimal value of the problem is thus equal to

$$\phi_{\ell_1}(0) = (\lambda_{\max}(A^T A))^{1/2} = \sigma_{\max}(A).$$

Note that there is no reason to expect this vector to be sparse. On the other hand, for large enough γ , we will necessarily have $z^*(\gamma) = 0$, obtaining maximal sparsity. Indeed, since

$$\max_{z \neq 0} \frac{\|Az\|_2}{\|z\|_1} = \max_{z \neq 0} \frac{\|\sum_i z_i a_i\|_2}{\|z\|_1} \leq \max_{z \neq 0} \frac{\sum_i |z_i| \|a_i\|_2}{\sum_i |z_i|} = \max_i \|a_i\|_2 = \|a_{i^*}\|_2,$$

we get $\|Az\|_2 - \gamma \|z\|_1 < 0$ for all nonzero vectors z whenever γ is chosen to be strictly bigger than $\|a_{i^*}\|_2$. From now on we will assume that

$$\gamma < \|a_{i^*}\|_2. \quad (2.2)$$

Note that there is a trade-off between the value $\|Az^*(\gamma)\|_2$ and the sparsity of the solution $z^*(\gamma)$. The penalty parameter γ is introduced to “continuously” interpolate between the two extreme cases described above, with values in the interval $[0, \|a_{i^*}\|_2)$. It depends on the particular

application whether sparsity is valued more than the explained variance, or vice versa, and to what extent. Due to these considerations, we will consider the solution of (2.1) to be a sparse principal component of A .

Reformulation. The reader will observe that the objective function in (2.1) is not convex, nor concave, and that the feasible set is of a high dimension if $p \ll n$. It turns out that these shortcomings are overcome by considering the following reformulation:

$$\begin{aligned} \phi_{\ell_1}(\gamma) &= \max_{z \in \mathcal{B}^n} \|Az\|_2 - \gamma \|z\|_1 \\ &= \max_{z \in \mathcal{B}^n} \max_{x \in \mathcal{B}^p} x^T Az - \gamma \|z\|_1 \end{aligned} \quad (2.3)$$

$$\begin{aligned} &= \max_{x \in \mathcal{B}^p} \max_{z \in \mathcal{B}^n} \sum_{i=1}^n z_i (a_i^T x) - \gamma \|z\|_1 \\ &= \max_{x \in \mathcal{B}^p} \max_{\bar{z} \in \mathcal{B}^n} \sum_{i=1}^n |\bar{z}_i| (|a_i^T x| - \gamma), \end{aligned} \quad (2.4)$$

where $z_i = \text{sign}(a_i^T x) \bar{z}_i$. In view of (2.2), there is some $x \in \mathcal{B}^p$ for which $a_i^T x > \gamma$. Fixing such x , solving the inner maximization problem for \bar{z} and then translating back to z , we obtain the closed-form solution

$$z_i^* = z_i^*(\gamma) = \frac{\text{sign}(a_i^T x) [|a_i^T x| - \gamma]_+}{\sqrt{\sum_{k=1}^n [|a_k^T x| - \gamma]_+^2}}, \quad i = 1, \dots, n. \quad (2.5)$$

Problem (2.4) can therefore be written in the form

$$\boxed{\phi_{\ell_1}^2(\gamma) = \max_{x \in \mathcal{S}^p} \sum_{i=1}^n [|a_i^T x| - \gamma]_+^2.} \quad (2.6)$$

Note that the objective function is differentiable and convex, and hence all local and global maxima must lie on the boundary, i.e., on the unit Euclidean sphere \mathcal{S}^p . Also, in the case when $p \ll n$, formulation (2.6) requires to search a space of a much lower dimension than the initial problem (2.1).

Sparsity. In view of (2.5), an optimal solution x^* of (2.6) defines a sparsity pattern of the vector z^* . In fact, the coefficients of z^* indexed by

$$\mathcal{I} = \{i \mid |a_i^T x^*| > \gamma\} \quad (2.7)$$

are active while all others must be zero. Geometrically, active indices correspond to the defining hyperplanes of the polytope

$$\mathcal{D} = \{x \in \mathbf{R}^p \mid |a_i^T x| \leq 1\}$$

that are (strictly) crossed by the line joining the origin and the point x^*/γ . Note that it is possible to say something about the sparsity of the solution even without the knowledge of x^* :

$$\gamma \geq \|a_i\|_2 \quad \Rightarrow \quad z_i^*(\gamma) = 0, \quad i = 1, \dots, n. \quad (2.8)$$

2.2 Single-unit sparse PCA via cardinality penalty

Instead of the ℓ_1 -penalization, the authors of [6] consider the formulation

$$\phi_{\ell_0}(\gamma) \stackrel{\text{def}}{=} \max_{z \in \mathcal{B}^n} z^T \Sigma z - \gamma \|z\|_0, \quad (2.9)$$

which directly penalizes the number of nonzero components (cardinality) of the vector z .

Reformulation. The reasoning of the previous section suggests the reformulation

$$\phi_{\ell_0}(\gamma) = \max_{x \in \mathcal{B}^p} \max_{z \in \mathcal{B}^n} (x^T A z)^2 - \gamma \|z\|_0, \quad (2.10)$$

where the maximization with respect to $z \in \mathcal{B}^n$ for a fixed $x \in \mathcal{B}^p$ has the closed form solution

$$z_i^* = z_i^*(\gamma) = \frac{[\text{sign}((a_i^T x)^2 - \gamma)]_+ a_i^T x}{\sqrt{\sum_{k=1}^n [\text{sign}((a_k^T x)^2 - \gamma)]_+ (a_k^T x)^2}}, \quad i = 1, \dots, n. \quad (2.11)$$

In analogy with the ℓ_1 case, this derivation assumes that

$$\gamma < \|a_{i^*}\|_2^2,$$

so that there is $x \in \mathcal{B}^n$ such that $(a_i^T x)^2 - \gamma > 0$. Otherwise $z^* = 0$ is optimal. Formula (2.11) is easily obtained by analyzing (2.10) separately for fixed cardinality values of z . Hence, problem (2.9) can be cast in the following form

$$\boxed{\phi_{\ell_0}(\gamma) = \max_{x \in \mathcal{S}^p} \sum_{i=1}^n [(a_i^T x)^2 - \gamma]_+}. \quad (2.12)$$

Again, the objective function is convex, albeit nonsmooth, and the new search space is of particular interest if $p \ll n$. A different derivation of (2.12) for the $n = p$ case can be found in [6].

Sparsity. Given a solution x^* of (2.12), the set of active indices of z^* is given by

$$\mathcal{I} = \{i \mid (a_i^T x^*)^2 > \gamma\}.$$

Geometrically, active indices correspond to the defining hyperplanes of the polytope

$$\mathcal{D} = \{x \in \mathbf{R}^p \mid |a_i^T x| \leq 1\}$$

that are (strictly) crossed by the line joining the origin and the point $x^*/\sqrt{\gamma}$. As in the ℓ_1 case, we have

$$\gamma \geq \|a_i\|_2^2 \quad \Rightarrow \quad z_i^*(\gamma) = 0, \quad i = 1, \dots, n. \quad (2.13)$$

2.3 Block sparse PCA via ℓ_1 -penalty

Consider the following block generalization of (2.3),

$$\phi_{\ell_1, m}(\gamma) \stackrel{\text{def}}{=} \max_{\substack{X \in \mathcal{S}_m^p \\ Z \in [\mathcal{S}^n]^m}} \text{Tr}(X^T A Z N) - \gamma \sum_{j=1}^m \sum_{i=1}^n |z_{ij}|, \quad (2.14)$$

where $\gamma \geq 0$ is a sparsity-controlling parameter and $N = \text{Diag}(\mu_1, \dots, \mu_m)$, with positive entries on the diagonal. The dimension m corresponds to the number of extracted components and is assumed to be smaller or equal to the rank of the data matrix, i.e., $m \leq \text{Rank}(A)$. It will be shown below that under some conditions on the parameters μ_i , the case $\gamma = 0$ recovers PCA. In that particular instance, any solution Z^* of (2.14) has orthonormal columns, although this is not explicitly enforced. For positive γ , the columns of Z^* are not expected to be orthogonal anymore. Most existing algorithms for computing several sparse principal components, for example those described by Zou et al. [24], d'Aspremont et al. [7] and Shen and Huang [18], also do not impose orthogonal loading directions. Simultaneously enforcing sparsity and orthogonality seems to be a hard (and perhaps questionable) task.

Reformulation. Since problem (2.14) is completely decoupled in the columns of Z , i.e.,

$$\phi_{\ell_1, m}(\gamma) = \max_{X \in \mathcal{S}_m^p} \sum_{j=1}^m \max_{z_j \in \mathcal{S}^n} \mu_j x_j^T A z_j - \gamma \|z_j\|_1,$$

the closed-form solution (2.5) of (2.3) is easily adapted to the block formulation (2.14):

$$z_{ij}^* = z_{ij}^*(\gamma) = \frac{\text{sign}(a_i^T x_j) [\mu_j |a_i^T x_j| - \gamma]_+}{\sqrt{\sum_{k=1}^n [\mu_j |a_k^T x_j| - \gamma]_+^2}}. \quad (2.15)$$

This leads to the reformulation

$$\boxed{\phi_{\ell_1, m}^2(\gamma) = \max_{X \in \mathcal{S}_m^p} \sum_{j=1}^m \sum_{i=1}^n [\mu_j |a_i^T x_j| - \gamma]_+^2}, \quad (2.16)$$

which maximizes a convex function $f : \mathbf{R}^{p \times m} \rightarrow \mathbf{R}$ on the Stiefel manifold \mathcal{S}_m^p .

Sparsity. A solution X^* of (2.16) again defines the sparsity pattern of the matrix Z^* : the entry z_{ij}^* is active if

$$\mu_j |a_i^T x_j^*| > \gamma,$$

and equal to zero otherwise. For $\gamma > \max_{i,j} \mu_j \|a_i\|_2$, the trivial solution $Z^* = 0$ is optimal.

Block PCA. For $\gamma = 0$, problem (2.16) can be equivalently written in the form

$$\phi_{\ell_1, m}^2(0) = \max_{X \in \mathcal{S}_m^p} \text{Tr}(X^T A A^T X N^2), \quad (2.17)$$

which has been well studied (see e.g., Brockett [3] and Absil et al. [1]). The solutions of (2.17) span the dominant m -dimensional invariant subspace of the matrix $A A^T$. Furthermore, if the

parameters μ_i are all distinct, the columns of X^* are the m dominant eigenvectors of AA^T , i.e., the m dominant left-eigenvectors of the data matrix A . The columns of the solution Z^* of (2.14) are thus the m dominant right singular vectors of A , i.e., the PCA loading vectors. Such a matrix N with distinct diagonal elements enforces the objective function in (2.17) to have isolated maximizers. In fact, if $N = I_m$, any point X^*U with X^* a solution of (2.17) and $U \in \mathcal{S}_m^m$ is also a solution of (2.17). In the case of sparse PCA, i.e., $\gamma > 0$, the penalty term enforces isolated maximizers. The technical parameter N will thus be set to the identity matrix in what follows.

2.4 Block sparse PCA via cardinality penalty

The single-unit cardinality-penalized case can also be naturally extended to the block case:

$$\phi_{\ell_0,m}(\gamma) \stackrel{\text{def}}{=} \max_{\substack{X \in \mathcal{S}_m^p \\ Z \in [\mathcal{S}^n]^m}} \text{Tr}(\text{Diag}(X^T A Z N)^2) - \gamma \|Z\|_0, \quad (2.18)$$

where $\gamma \geq 0$ is the sparsity inducing parameter and $N = \text{Diag}(\mu_1, \dots, \mu_m)$ with positive entries on the diagonal. In the case $\gamma = 0$, problem (2.20) is equivalent to (2.17) and therefore corresponds to PCA, provided that all μ_i are distinct.

Reformulation. Again, this block formulation is completely decoupled in the columns of Z ,

$$\phi_{\ell_0,m}(\gamma) = \max_{X \in \mathcal{S}_m^p} \sum_{j=1}^m \max_{z_j \in \mathcal{S}^n} (\mu_j x_j^T A z_j)^2 - \gamma \|z_j\|_0,$$

so that the solution (2.11) of the single unit case provides the optimal columns z_j :

$$z_{ij}^* = z_{ij}^*(\gamma) = \frac{[\text{sign}((\mu_j a_i^T x_j)^2 - \gamma)]_+ \mu_j a_i^T x_j}{\sqrt{\sum_{k=1}^n [\text{sign}((\mu_j a_k^T x_j)^2 - \gamma)]_+ \mu_j^2 (a_k^T x_j)^2}}. \quad (2.19)$$

The reformulation of problem (2.18) is thus

$$\boxed{\phi_{\ell_0,m}(\gamma) = \max_{X \in \mathcal{S}_m^p} \sum_{j=1}^m \sum_{i=1}^n [(\mu_j a_i^T x_j)^2 - \gamma]_+,} \quad (2.20)$$

which maximizes a convex function $f : \mathbf{R}^{p \times m} \rightarrow \mathbf{R}$ on the Stiefel manifold \mathcal{S}_m^p .

Sparsity. For a solution X^* of (2.20), the active entries z_{ij}^* of Z^* are given by the condition

$$(\mu_j a_i^T x_j^*)^2 > \gamma.$$

Hence for $\gamma > \max_{i,j} \mu_j \|a_i\|_2^2$, the optimal solution of (2.18) is $Z^* = 0$.

3 A gradient method for maximizing convex functions

By \mathbf{E} we denote an arbitrary finite-dimensional vector space; \mathbf{E}^* is its conjugate, i.e. the space of all linear functionals on \mathbf{E} . We equip these spaces with norms given by (1.2).

In this section we propose and analyze a simple gradient-type method for maximizing a convex function $f : \mathbf{E} \rightarrow \mathbf{R}$ on a compact set \mathcal{Q} :

$$\boxed{f^* = \max_{x \in \mathcal{Q}} f(x)}. \quad (\text{P})$$

Unless explicitly stated otherwise, we will *not* assume f to be differentiable. By $f'(x)$ we denote any subgradient of function f at x . By $\partial f(x)$ we denote its subdifferential.

At any point $x \in \mathcal{Q}$ we introduce some measure for the first-order optimality conditions:

$$\Delta(x) \stackrel{\text{def}}{=} \max_{y \in \mathcal{Q}} \langle f'(x), y - x \rangle.$$

Clearly, $\Delta(x) \geq 0$ and it vanishes only at the points where the gradient $f'(x)$ belongs to the normal cone to the set $\text{Conv}(\mathcal{Q})$ at x .³

We will use the following notation:

$$y(x) \stackrel{\text{def}}{\in} \text{Arg max}_{y \in \mathcal{Q}} \langle f'(x), y - x \rangle. \quad (3.1)$$

3.1 Algorithm

Consider the following simple algorithmic scheme.

Algorithm 1: Gradient scheme

input : Initial iterate $x_0 \in \mathbf{E}$.
output: x_k , approximate solution of (P)
begin
 $k \leftarrow 0$
 repeat
 $x_{k+1} \in \text{Arg max}\{f(x_k) + \langle f'(x_k), y - x_k \rangle \mid y \in \mathcal{Q}\}$
 $k \leftarrow k + 1$
 until a stopping criterion is satisfied
end

Note that for example in the special case $\mathcal{Q} = r \cdot \mathcal{S} \stackrel{\text{def}}{=} r \cdot \{x \in \mathbf{E} \mid \|x\| = r\}$ or $\mathcal{Q} = r \cdot \mathcal{B} \stackrel{\text{def}}{=} r \cdot \{x \in \mathbf{E} \mid \|x\| \leq r\}$, the main step of Algorithm 1 can be written in an explicit form:

$$y(x_k) = x_{k+1} = r \frac{G^{-1} f'(x_k)}{\|f'(x_k)\|_*}. \quad (3.2)$$

³The normal cone to the set $\text{Conv}(\mathcal{Q})$ at $x \in \mathcal{Q}$ is *smaller* than the normal cone to the set \mathcal{Q} . Therefore, the optimality condition $\Delta(x) = 0$ is *stronger* than the standard one.

3.2 Analysis

Our first convergence result is straightforward. Denote $\Delta_k \stackrel{\text{def}}{=} \min_{0 \leq i \leq k} \Delta(x_i)$.

Theorem 1 *Let sequence $\{x_k\}_{k=0}^{\infty}$ be generated by Algorithm 1 as applied to a convex function f . Then the sequence $\{f(x_k)\}_{k=0}^{\infty}$ is monotonically increasing and $\lim_{k \rightarrow \infty} \Delta(x_k) = 0$. Moreover,*

$$\Delta_k \leq \frac{f^* - f(x_0)}{k + 1}. \quad (3.3)$$

Proof:

From convexity of f we immediately get

$$f(x_{k+1}) \geq f(x_k) + \langle f'(x_k), x_{k+1} - x_k \rangle = f(x_k) + \Delta(x_k),$$

and therefore, $f(x_{k+1}) \geq f(x_k)$ for all k . By summing up these inequalities for $k = 0, 1, \dots, N-1$, we obtain

$$f^* - f(x_0) \geq f(x_k) - f(x_0) \geq \sum_{i=0}^k \Delta(x_i),$$

and the result follows. \square

For a sharper analysis, we need some technical assumptions on f and \mathcal{Q} .

Assumption 1 *The norms of the subgradients of f are bounded from below on \mathcal{Q} by a positive constant, i.e.*

$$\delta_f \stackrel{\text{def}}{=} \min_{\substack{x \in \mathcal{Q} \\ f'(x) \in \partial f(x)}} \|f'(x)\|_* > 0. \quad (3.4)$$

This assumption is not too binding because of the following result.

Proposition 2 *Assume that there exists a point $\bar{x} \notin \mathcal{Q}$ such that $f(\bar{x}) < f(x)$ for all $x \in \mathcal{Q}$. Then*

$$\delta_f \geq \left[\min_{x \in \mathcal{Q}} f(x) - f(\bar{x}) \right] / \left[\max_{x \in \mathcal{Q}} \|x - \bar{x}\| \right] > 0.$$

Proof:

Because f is convex, for any $x \in \mathcal{Q}$ we have

$$0 < f(x) - f(\bar{x}) \leq \langle f'(x), x - \bar{x} \rangle \leq \|f'(x)\|_* \cdot \|x - \bar{x}\|.$$

\square

For our next convergence result we need to assume either strong convexity of f or strong convexity of the set $\text{Conv}(\mathcal{Q})$.

Assumption 2 *Function f is strongly convex, i.e. there exists a constant $\sigma_f > 0$ such that for any $x, y \in \mathbf{E}$*

$$f(y) \geq f(x) + \langle f'(x), y - x \rangle + \frac{\sigma_f}{2} \|y - x\|^2. \quad (3.5)$$

Convex functions satisfy this inequality for *convexity parameter* $\sigma_f = 0$.

Assumption 3 *The set $\text{Conv}(\mathcal{Q})$ is strongly convex. This means that there exists a constant $\sigma_{\mathcal{Q}} > 0$ such that for any $x, y \in \text{Conv}(\mathcal{Q})$ and $\alpha \in [0, 1]$ the following inclusion holds:*

$$\alpha x + (1 - \alpha)y + \frac{\sigma_{\mathcal{Q}}}{2}\alpha(1 - \alpha)\|x - y\|^2 \cdot \mathcal{S} \subset \text{Conv}(\mathcal{Q}). \quad (3.6)$$

Convex sets satisfy this inclusion for *convexity parameter* $\sigma_{\mathcal{Q}} = 0$. It can be shown (see Appendix), that level sets of strongly convex functions with Lipschitz continuous gradient are again strongly convex. An example of such a function is the simple quadratic $x \mapsto \|x\|^2$. The level sets of this function correspond to Euclidean balls of varying sizes.

As we will see in Theorem 4, a better analysis of Algorithm 1 is possible if $\text{Conv}(\mathcal{Q})$, the convex hull of the feasible set of problem (P), is strongly convex. Note that in the case of the two formulations (2.6) and (2.12) of the sparse PCA problem, the feasible set \mathcal{Q} is the unit Euclidean sphere. Since the convex hull of the unit sphere is the unit ball, which is a strongly convex set, the feasible set of our sparse PCA formulations satisfies Assumption 3.

In the special case $\mathcal{Q} = r \cdot \mathcal{S}$ for some $r > 0$, there is a simple proof that Assumption 3 holds with $\sigma_{\mathcal{Q}} = \frac{1}{r}$. Indeed, for any $x, y \in \mathbf{E}$ and $\alpha \in [0, 1]$, we have

$$\begin{aligned} \|\alpha x + (1 - \alpha)y\|^2 &= \alpha^2\|x\|^2 + (1 - \alpha)^2\|y\|^2 + 2\alpha(1 - \alpha)\langle Gx, y \rangle \\ &= \alpha\|x\|^2 + (1 - \alpha)\|y\|^2 - \alpha(1 - \alpha)\|x - y\|^2. \end{aligned}$$

Thus, for $x, y \in r \cdot \mathcal{S}$ we obtain:

$$\|\alpha x + (1 - \alpha)y\| = [r^2 - \alpha(1 - \alpha)\|x - y\|^2]^{1/2} \leq r - \frac{1}{2r}\alpha(1 - \alpha)\|x - y\|^2.$$

Hence, we can take $\sigma_{\mathcal{Q}} = \frac{1}{r}$.

The relevance of Assumption 3 is justified by the following technical observation.

Proposition 3 *Let Assumption 3 be satisfied. Then for any $x \in \mathcal{Q}$ the following holds:*

$$\Delta(x) \geq \frac{\sigma_{\mathcal{Q}}}{2}\|f'(x)\|_* \cdot \|y(x) - x\|^2. \quad (3.7)$$

Proof:

Fix an arbitrary $x \in \mathcal{Q}$. Note that

$$\langle f'(x), y(x) - y \rangle \geq 0, \quad y \in \text{Conv}(\mathcal{Q}).$$

We will use this inequality for

$$y = y_{\alpha} \stackrel{\text{def}}{=} x + \alpha(y(x) - x) + \frac{\sigma_{\mathcal{Q}}}{2}\alpha(1 - \alpha)\|y(x) - x\|^2 \cdot \frac{G^{-1}f'(x)}{\|f'(x)\|_*}, \quad \alpha \in [0, 1].$$

In view of Assumption 3, $y_{\alpha} \in \text{Conv}(\mathcal{Q})$. Therefore,

$$0 \geq \langle f'(x), y_{\alpha} - y(x) \rangle = (1 - \alpha)\langle f'(x), x - y(x) \rangle + \frac{\sigma_{\mathcal{Q}}}{2}\alpha(1 - \alpha)\|y(x) - x\|^2 \cdot \|f'(x)\|_*.$$

Since α is an arbitrary value from $[0, 1]$, the result follows. \square

We are now ready to refine our analysis of Algorithm 1.

Theorem 4 (Convergence) *Let f be convex and let Assumption 1 and at least one of Assumptions 2 and 3 be satisfied. If $\{x_k\}$ is the sequence of points generated by Algorithm 1, then*

$$\sum_{k=0}^N \|x_{k+1} - x_k\|^2 \leq \frac{2(f^* - f(x_0))}{\sigma_{\mathcal{Q}}\delta_f + \sigma_f}. \quad (3.8)$$

Proof:

Indeed, in view of our assumptions and Proposition 3, we have

$$f(x_{k+1}) - f(x_k) \geq \Delta(x_k) + \frac{\sigma_f}{2} \|x_{k+1} - x_k\|^2 \geq \frac{1}{2}(\sigma_{\mathcal{Q}}\delta_f + \sigma_f) \|x_{k+1} - x_k\|^2.$$

□

We cannot in general guarantee that the algorithm will converge to a unique local maximizer. In particular, if started from a local minimizer, the method will not move away from this point. However, the above statement guarantees that the set of its limit points is connected and all of them satisfy the first-order optimality condition.

3.3 Maximization with spherical constraints

Consider $\mathbf{E} = \mathbf{E}^* = \mathbf{R}^p$ with $G = I_p$ and $\langle s, x \rangle = \sum_i s_i x_i$, and let

$$\mathcal{Q} = r \cdot \mathcal{S}^p = \{x \in \mathbf{R}^p \mid \|x\| = r\}.$$

Problem (P) takes on the form:

$$f^* = \max_{x \in r \cdot \mathcal{S}^p} f(x).$$

Since \mathcal{Q} is strongly convex ($\sigma_{\mathcal{Q}} = \frac{1}{r}$), Theorem 4 is meaningful for any convex function f ($\sigma_f \geq 0$). We have already noted (see (3.2)) that the main step of Algorithm 1 can be written down explicitly. Note that the single-unit sparse PCA formulations (2.6) and (2.12) conform to this setting. The following examples illustrate the connection to classical algorithms.

Example 5 (Power method) *In the special case of a quadratic objective function*

$$f(x) = \frac{1}{2}x^T Cx$$

for some $C \in \mathbf{S}_{++}^p$ on the unit sphere ($r = 1$), we have

$$f^* = \frac{1}{2}\lambda_{\max}(C),$$

and Algorithm 1 is equivalent to the power iteration method for computing the largest eigenvalue of C (Golub and Van Loan [9]). Hence for $\mathcal{Q} = \mathcal{S}^p$, we can think of our scheme as a generalization of the power method. Indeed, our algorithm performs the following iteration:

$$x_{k+1} = \frac{Cx_k}{\|Cx_k\|}, \quad k \geq 0.$$

Note that both δ_f and σ_f are equal to the smallest eigenvalue of C , and hence the right-hand side of (3.8) is equal to

$$\frac{\lambda_{\max}(C) - x_0^T Cx_0}{2\lambda_{\min}(C)}. \quad (3.9)$$

Example 6 (Shifted power method) *If C is not positive semidefinite in the previous example, the objective function is not convex and our results are not applicable. However, this complication can be circumvented by instead running the algorithm with the shifted quadratic function*

$$\hat{f}(x) = \frac{1}{2}x^T(C + \omega I_p)x,$$

where $\omega > 0$ satisfies $\hat{C} = \omega I_p + C \in \mathbf{S}_{++}^p$. On the feasible set, this change only adds a constant term to the objective function. The method, however, produces different sequence of iterates. Note that the constants δ_f and σ_f are also affected and, correspondingly, the estimate (3.9).

3.4 Maximization with orthonormality constraints

Consider $\mathbf{E} = \mathbf{E}^* = \mathbf{R}^{p \times m}$, the space of $p \times m$ real matrices, with $m \leq p$. Note that for $m = 1$ we recover the setting of the previous section. We assume this space is equipped with the trace inner product: $\langle X, Y \rangle = \text{Tr}(X^T Y)$. The induced norm, denoted by $\|X\|_F \stackrel{\text{def}}{=} \langle X, X \rangle^{1/2}$, is the Frobenius norm (we let G be the identity operator). We can now consider various feasible sets, the simplest being a ball or a sphere. Due to nature of applications in this paper, let us concentrate on the situation when \mathcal{Q} is a special subset of the sphere with radius $r = \sqrt{m}$, the Stiefel manifold \mathcal{S}_m^p :

$$\mathcal{Q} = \mathcal{S}_m^p = \{X \in \mathbf{R}^{p \times m} \mid X^T X = I_m\}.$$

Problem (P) then takes on the following form:

$$f^* = \max_{X \in \mathcal{S}_m^p} f(X).$$

Note that $\text{Conv}(\mathcal{Q})$ is not strongly convex ($\sigma_{\mathcal{Q}} = 0$), and hence Theorem 4 is meaningful only if f is strongly convex ($\sigma_f > 0$). At every iteration, the algorithm needs to maximize a linear function over the Stiefel manifold. The following standard result shows how this can be done.

Proposition 7 *Let $C \in \mathbf{R}^{p \times m}$, with $m \leq p$, and denote by $\sigma_i(C)$, $i = 1, \dots, m$, the singular values of C . Then*

$$\max_{X \in \mathcal{S}_m^p} \langle C, X \rangle = \text{Tr}[(C^T C)^{1/2}] = \sum_{i=1}^m \sigma_i(C), \quad (3.10)$$

and a maximizer X^* is given by the U factor in the polar decomposition of C :

$$C = UP, \quad U \in \mathcal{S}_m^p, \quad P \in \mathbf{S}_+^m.$$

If C is of full rank, then we can take $X^* = C(C^T C)^{-1/2}$.

Proof:

Existence of the polar factorization in the nonsquare case is covered by Theorem 7.3.2 in Horn and Johnson [10]. Let $C = V\Sigma W^T$ be the singular value decomposition of A ; that is, V is $p \times p$ orthonormal, W is $m \times m$ orthonormal, and Σ is $p \times m$ diagonal with values $\sigma_i(A)$ on the

diagonal. Then

$$\begin{aligned}
\max_{X \in \mathcal{S}_m^p} \langle C, X \rangle &= \max_{X \in \mathcal{S}_m^p} \langle V \Sigma W^T, X \rangle \\
&= \max_{X \in \mathcal{S}_m^p} \text{Tr} \Sigma (W^T X^T V) \\
&= \max_{Z \in \mathcal{S}_m^p} \text{Tr} \Sigma Z^T = \max_{Z \in \mathcal{S}_m^p} \sum_{i=1}^m \sigma_i(C) z_{ii} \leq \sum_i^m \sigma_i(C).
\end{aligned}$$

The third equality follows since the function $X \mapsto V^T X W$ maps \mathcal{S}_m^p onto itself. It remains to note that

$$\langle C, U \rangle = \text{Tr} P = \sum_i \lambda_i(P) = \sum_i \sigma_i(P) = \text{Tr}(P^T P)^{1/2} = \text{Tr}(C^T C)^{1/2} = \sum_i \sigma_i(C),$$

Finally, in the full rank case we have $\langle C, X^* \rangle = \text{Tr} C^T C (C^T C)^{-1/2} = \text{Tr}(C^T C)^{1/2}$. □

In the sequel, the symbol $\text{Uf}(C)$ will be used to denote the U factor of the polar decomposition of matrix $C \in \mathbf{R}^{p \times m}$, or equivalently, $\text{Uf}(C) = C(C^T C)^{-1/2}$ if C is of full rank. In view of the above result, the main step of Algorithm 1 can be written in the form

$$x_{k+1} = \text{Uf}(f'(x_k)). \quad (3.11)$$

Note that the block sparse PCA formulations (2.16) and (2.20) conform to this setting. Here is one more example:

Example 8 (Rectangular Procrustes Problem) *Let $C, X \in \mathbf{R}^{p \times m}$ and $D \in \mathbf{R}^{p \times p}$ and consider the following problem:*

$$\min\{\|C - DX\|_F^2 \mid X^T X = I_m\}. \quad (3.12)$$

Since $\|C - DX\|_F^2 = \|C\|_F^2 + \langle DX, DX \rangle - 2\langle CD, X \rangle$, by a similar shifting technique as in the previous example we can cast problem (3.12) in the following form

$$\max\{\omega \|X\|_F^2 - \langle DX, DX \rangle + 2\langle CD, X \rangle \mid X^T X = I_m\}.$$

For $\omega > 0$ large enough, the new objective function will be strongly convex. In this case our algorithm becomes similar to the gradient method proposed by Fraikin et al. [8].

The standard Procrustes problem in the literature is a special case of (3.12) with $p = m$.

4 Algorithms for sparse PCA

The application of our general method (Algorithm 1) to the four sparse PCA formulations of Section 2, i.e., (2.6), (2.12), (2.16) and (2.20), leads to Algorithms 2, 3, 4 and 5 below, that

provide a locally optimal pattern of sparsity for a matrix $Z \in [\mathcal{S}^n]^m$.⁴ This pattern is defined as a matrix $P \in \mathbf{R}^{n \times m}$ such that $p_{ij} = 0$ if the loading z_{ij} is active and $p_{ij} = 1$ otherwise. So P is an indicator of the coefficients of Z that are zeroed by our method. The computational complexity of the single-unit algorithms (Algorithms 2 and 3) is $\mathcal{O}(np)$ operations per iteration. The block algorithms (Algorithms 4 and 5) have complexity $\mathcal{O}(npm)$ per iteration.

4.1 Methods for pattern-finding

Algorithm 2: Single-unit sparse PCA method based on the ℓ_1 -penalty (2.6)

input : Data matrix $A \in \mathbf{R}^{p \times n}$
Sparsity-controlling parameter $\gamma \geq 0$
Initial iterate $x \in \mathcal{S}^p$

output: A locally optimal sparsity pattern P

begin

repeat

$x \leftarrow \sum_{i=1}^n [|a_i^T x| - \gamma]_+ \text{sign}(a_i^T x) a_i$

$x \leftarrow \frac{x}{\|x\|}$

until a stopping criterion is satisfied

Construct vector $P \in \mathbf{R}^n$ such that $\begin{cases} p_i = 0 & \text{if } |a_i^T x| > \gamma \\ p_i = 1 & \text{otherwise.} \end{cases}$

end

Algorithm 3: Single-unit sparse PCA algorithm based on the ℓ_0 -penalty (2.12)

input : Data matrix $A \in \mathbf{R}^{p \times n}$
Sparsity-controlling parameter $\gamma \geq 0$
Initial iterate $x \in \mathcal{S}^p$

output: A locally optimal sparsity pattern P

begin

repeat

$x \leftarrow \sum_{i=1}^n [\text{sign}((a_i^T x)^2 - \gamma)]_+ a_i^T x a_i$

$x \leftarrow \frac{x}{\|x\|}$

until a stopping criterion is satisfied

Construct vector $P \in \mathbf{R}^n$ such that $\begin{cases} p_i = 0 & \text{if } (a_i^T x)^2 > \gamma \\ p_i = 1 & \text{otherwise.} \end{cases}$

end

⁴This section discusses the general block sparse PCA problem. The single-unit case corresponds to the particular case $m = 1$.

Algorithm 4: Block Sparse PCA algorithm based on the ℓ_1 -penalty (2.16)

input : Data matrix $A \in \mathbf{R}^{p \times n}$
Sparsity-controlling parameter $\gamma \geq 0$
Initial iterate $X \in \mathcal{S}_m^p$
output: A locally optimal sparsity pattern P
begin
 repeat
 for $j = 1, \dots, m$ **do**
 $x_j \leftarrow \sum_{i=1}^n [|a_i^T x_j| - \gamma]_+ \text{sign}(a_i^T x) a_i$
 $X \leftarrow \text{Uf}(X)$
 until a stopping criterion is satisfied
 Construct matrix $P \in \mathbf{R}^{n \times m}$ such that $\begin{cases} p_{ij} = 0 & \text{if } |a_i^T x_j| > \gamma \\ p_{ij} = 1 & \text{otherwise.} \end{cases}$
end

Algorithm 5: Block Sparse PCA algorithm based on the ℓ_0 -penalty (2.20)

input : Data matrix $A \in \mathbf{R}^{p \times n}$
Sparsity-controlling parameter $\gamma \geq 0$
Initial iterate $X \in \mathcal{S}_m^p$
output: A locally optimal sparsity pattern P
begin
 repeat
 for $j = 1, \dots, m$ **do**
 $x_j \leftarrow \sum_{i=1}^n [\text{sign}((a_i^T x_j)^2 - \gamma)]_+ a_i^T x_j a_i$
 $X \leftarrow \text{Uf}(X)$
 until a stopping criterion is satisfied
 Construct matrix $P \in \mathbf{R}^{n \times m}$ such that $\begin{cases} p_{ij} = 0 & \text{if } (a_i^T x_j)^2 > \gamma \\ p_{ij} = 1 & \text{otherwise.} \end{cases}$
end

4.2 Post-processing

Once a “good” sparsity pattern P has been identified, the active entries of Z still have to be filled. To this end, we consider the optimization problem,

$$(X^*, Z^*) \stackrel{\text{def}}{=} \arg \max_{\substack{X \in \mathcal{S}_m^p \\ Z \in [\mathcal{S}^n]^m \\ Z_P = 0}} \text{Tr}(X^T A Z N), \quad (4.1)$$

where Z_P denotes the entries of Z that are constrained to zero and $N = \text{Diag}(\mu_1, \dots, \mu_m)$ with strictly positive μ_i . Problem (4.1) assigns the active part of the loading vectors Z to maximize the variance explained by the resulting components. By $Z_{\bar{P}}$, we refer to the complement of Z_P , i.e., to the active entries of Z . In the single-unit case $m = 1$, an explicit solution of (4.1) is

available,

$$\begin{aligned} X^* &= u, \\ Z_{\bar{P}}^* &= v \text{ and } Z_P^* = 0, \end{aligned} \quad (4.2)$$

where σuv^T with $\sigma > 0$, $u \in \mathcal{B}^p$ and $v \in \mathcal{B}^{\|\bar{P}\|_0}$ is a rank one singular value decomposition of the matrix $A_{\bar{P}}$, that corresponds to the submatrix of A containing the columns related to the active entries.

Although an exact solution of (4.1) is hard to compute in the block case $m > 1$, a local maximizer can be efficiently computed by optimizing alternatively with respect to one variable while keeping the other ones fixed. The following lemmas provide an explicit solution to each of these subproblems.

Lemma 9 For a fixed $Z \in [\mathcal{S}^n]^m$, a solution X^* of

$$\max_{X \in \mathcal{S}_m^p} \text{Tr}(X^T AZN)$$

is provided by the U factor of the polar decomposition of the product AZN .

Proof:

See Proposition 7. □

Lemma 10 The solution

$$Z^* \stackrel{\text{def}}{=} \arg \max_{\substack{Z \in [\mathcal{S}^n]^m \\ Z_P = 0}} \text{Tr}(X^T AZN), \quad (4.3)$$

is at any point $X \in \mathcal{S}_m^p$ defined by the two conditions $Z_{\bar{P}}^* = (A^T XND)_{\bar{P}}$ and $Z_P^* = 0$, where D is a positive diagonal matrix that normalizes each column of Z^* to unit norm, i.e.,

$$D = \text{Diag}(NX^T AA^T XN)^{-\frac{1}{2}}.$$

Proof:

The Lagrangian of the optimization problem (4.3) is

$$\mathcal{L}(Z, \Lambda_1, \Lambda_2) = \text{Tr}(X^T AZN) - \text{Tr}(\Lambda_1(Z^T Z - I_m)) - \text{Tr}(\Lambda_2^T Z),$$

where the Lagrangian multipliers $\Lambda_1 \in \mathbf{R}^{m \times m}$ and $\Lambda_2 \in \mathbf{R}^{n \times m}$ have the following properties: Λ_1 is an invertible diagonal matrix and $(\Lambda_2)_{\bar{P}} = 0$. The first order optimality conditions of (4.3) are thus

$$\begin{aligned} A^T XN - 2Z\Lambda_1 - \Lambda_2 &= 0 \\ \text{Diag}(Z^T Z) &= I_m \\ Z_P &= 0. \end{aligned}$$

Hence, any stationary point Z^* of (4.3) satisfies $Z_{\bar{P}}^* = (A^T XND)_{\bar{P}}$ and $Z_P^* = 0$, where D is a diagonal matrix that normalizes the columns of Z^* to unit norm. The second order optimality condition imposes the diagonal matrix D to be positive. Such a D is unique and given by $D = \text{Diag}(NX^T AA^T XN)^{-\frac{1}{2}}$. □

Algorithm 6: Alternating optimization scheme for solving (4.1)

input : Data matrix $A \in \mathbf{R}^{p \times n}$
Sparsity pattern $P \in \mathbf{R}^{n \times m}$
Matrix $N = \text{Diag}(\mu_1, \dots, \mu_m)$
Initial iterate $X \in \mathcal{S}_m^p$
output: A local minimizer (X, Z) of (4.1)
begin
 repeat
 $Z \leftarrow A^T X N$
 $Z \leftarrow Z \text{Diag}(Z^T Z)^{-\frac{1}{2}}$
 $Z_P \leftarrow 0$
 $X \leftarrow \text{Uf}(AZN)$
 until a stopping criterion is satisfied
end

The alternating optimization scheme is summarized in Algorithm 6, which computes a local solution of (4.1). It should be noted that Algorithm 6 is a postprocessing heuristic that, strictly speaking, is required only for the ℓ_1 block formulation (Algorithm 4). In fact, since the cardinality penalty only depends on the sparsity pattern P and not on the actual values assigned to $Z_{\bar{P}}$, a solution (X^*, Z^*) of Algorithms 3 or 5 is also a local maximizer of (4.1) for the resulting pattern P . This explicit solution provides a good alternative to Algorithm 6. In the single unit case with ℓ_1 penalty (Algorithm 2), the solution (4.2) is available.

4.3 Sparse PCA algorithms

To sum up, in this paper we propose four sparse PCA algorithms, each combining a method to identify a “good” sparsity pattern with a method to fill the active entries of the m loading vectors. They are summarized in Table 1.⁵

	Computation of P	Computation of $Z_{\bar{P}}$
GPower $_{\ell_1}$	Algorithm 2	Equation (4.2)
GPower $_{\ell_0}$	Algorithm 3	Equation (2.11)
GPower $_{\ell_1, m}$	Algorithm 4	Algorithm 6
GPower $_{\ell_0, m}$	Algorithm 5	Equation (2.19)

Table 1: New algorithms for sparse PCA.

4.4 Deflation scheme.

For the sake of completeness, we recall a classical deflation process for computing m sparse principal components with a single-unit algorithm (d’Aspremont et al. [7]). Let $z \in \mathbf{R}^n$ be

⁵Our algorithms are named GPower where the “G” stands for *generalized* or *gradient*.

a unit-norm sparse loading vector of the data A . Subsequent directions can be sequentially obtained by computing a dominant sparse component of the residual matrix $A - xz^T$, where $x = Az$ is the vector that solves

$$\min_{x \in \mathbf{R}^p} \|A - xz^T\|_F.$$

5 Numerical experiments

In this section, we evaluate the proposed power algorithms against existing sparse PCA methods. Three competing methods are considered in this study: a greedy scheme aimed at computing a local maximizer of (2.9) (d’Aspremont et al. [6]), the SPCA algorithm (Zou et al. [24]) and the sPCA-rSVD algorithm (Shen and Huang [18]). We do not include the DSPCA algorithm (d’Aspremont et al. [7]) in our numerical study. This method solves a convex relaxation of the sparse PCA problem and has a large computational complexity of $\mathcal{O}(n^3)$ compared to the other methods. Table 2 lists the considered algorithms.

GPower $_{\ell_1}$	Single-unit sparse PCA via ℓ_1 -penalty
GPower $_{\ell_0}$	Single-unit sparse PCA via ℓ_0 -penalty
GPower $_{\ell_1, m}$	Block sparse PCA via ℓ_1 -penalty
GPower $_{\ell_0, m}$	Block sparse PCA via ℓ_0 -penalty
Greedy	Greedy method
SPCA	SPCA algorithm
rSVD $_{\ell_1}$	sPCA-rSVD algorithm with an ℓ_1 -penalty (“soft thresholding”)
rSVD $_{\ell_0}$	sPCA-rSVD algorithm with an ℓ_0 -penalty (“hard thresholding”)

Table 2: Sparse PCA algorithms we compare in this section.

These algorithms are compared on random data (Section 5.1) as well as on real data (Section 5.2). All numerical experiments are performed in MATLAB. Our implementations of the GPower algorithms are initialized at a point for which the associated sparsity pattern has *at least one* active element. In case of the single-unit algorithms, such an initial iterate $x \in \mathcal{S}^p$ is chosen parallel to the column of A with the largest norm, i.e.,

$$x = \frac{a_{i^*}}{\|a_{i^*}\|_2}, \quad \text{where } i^* = \arg \max_i \|a_i\|_2. \quad (5.1)$$

For the block GPower algorithms, a suitable initial iterate $X \in \mathcal{S}_m^p$ is constructed in a block-wise manner as $X = [x|X_\perp]$, where x is the unit-norm vector (5.1) and $X_\perp \in \mathcal{S}_{m-1}^p$ is orthogonal to x , i.e., $x^T X_\perp = 0$. We stop the GPower algorithms once the relative change of the objective function is small:

$$\frac{f(x_{k+1}) - f(x_k)}{f(x_k)} \leq \epsilon = 10^{-4}.$$

MATLAB implementations of the SPCA algorithm and the greedy algorithm have been rendered available by Zou et al. [24] and d’Aspremont et al. [6]. We have, however, implemented the sPCA-rSVD algorithm on our own (Algorithm 1 in Shen and Huang [18]), and use it with the same stopping criterion as for the GPower algorithms. This algorithm initializes with the best rank-one approximation of the data matrix. This is done with the `svds` function in MATLAB.

Given a data matrix $A \in \mathbf{R}^{p \times n}$, the considered sparse PCA algorithms provide m unit-norm sparse loading vectors stored in the matrix $Z \in [\mathcal{S}^n]^m$. The samples of the associated components are provided by the m columns of the product AZ . The variance explained by these m components is an important comparison criterion of the algorithms. In the simple case $m = 1$, the variance explained by the component Az is

$$\text{Var}(z) = z^T A^T A z.$$

When z corresponds to the first principal loading vector, the variance is $\text{Var}(z) = \sigma_{\max}(A)^2$. In the case $m > 1$, the derived components are likely to be correlated. Hence, summing up the variance explained individually by each of the components overestimates the variance explained simultaneously by all the components. This motivates the notion of *adjusted variance* proposed by Zou et al. [24]. The adjusted variance of the m components $Y = AZ$ is defined as

$$\text{AdjVar } Z = \text{Tr } R^2,$$

where $Y = QR$ is the QR decomposition of the components sample matrix Y ($Q \in \mathcal{S}_m^p$ and R is an $m \times m$ upper triangular matrix).

5.1 Random test problems

All random data matrices $A \in \mathbf{R}^{p \times n}$ considered in this section are generated according to a Gaussian distribution, with zero mean and unit variance.

Trade-off curves. Let us first compare the single-unit algorithms, which provide a unit-norm sparse loading vector $z \in \mathbf{R}^n$. We first plot the variance explained by the extracted component against the cardinality of the resulting loading vector z . For each algorithm, the sparsity-inducing parameter is incrementally increased to obtain loading vectors z with a cardinality that decreases from n to 1. The results displayed in Figure 1 are averages of computations on 100 random matrices with dimensions $p = 100$ and $n = 300$. The considered sparse PCA methods aggregate in two groups: GPower_{ℓ_1} , GPower_{ℓ_0} , Greedy and rSVD_{ℓ_0} outperform the SPCA and the rSVD_{ℓ_1} approaches. It seems that these latter methods perform worse because of the ℓ_1 penalty term used in them. If one, however, post-processes the active part of z according to (4.2), as we do in GPower_{ℓ_1} , all sparse PCA methods reach the same performance.

Controlling sparsity with γ . Among the considered methods, the greedy approach is the only one to directly control the cardinality of the solution, i.e., the desired cardinality is an input of the algorithm. The other methods require a parameter controlling the trade-off between variance and cardinality. Increasing this parameter leads to solutions with smaller cardinality, but the resulting number of nonzero elements can not be precisely predicted. In Figure 2, we plot the average relationship between the parameter γ and the resulting cardinality of the loading vector z for the two algorithms GPower_{ℓ_1} and GPower_{ℓ_0} . In view of (2.8) (resp. (2.13)), the entries i of the loading vector z obtained by the GPower_{ℓ_1} algorithm (resp. the GPower_{ℓ_0} algorithm) satisfying

$$\|a_i\|_2 \leq \gamma \quad (\text{resp. } \|a_i\|_2^2 \leq \gamma) \quad (5.2)$$

have to be zero. Taking into account the distribution of the norms of the columns of A , this provides for every γ a theoretical upper bound on the expected cardinality of the resulting vector z .




Figure 1: **Trade-off curves** between explained variance and cardinality. The vertical axis is the ratio $\text{Var}(z_{\text{sPCA}})/\text{Var}(z_{\text{PCA}})$, where the loading vector z_{sPCA} is computed by sparse PCA and z_{PCA} is the first principal loading vector. The considered algorithms aggregate in two groups: GPower_{ℓ_1} , GPower_{ℓ_0} , Greedy and rSVD_{ℓ_0} (top curve), and SPCA and rSVD_{ℓ_1} (bottom curve). For a fixed cardinality value, the methods of the first group explain more variance. Postprocessing algorithms SPCA and rSVD_{ℓ_1} with equation (4.2), results, however, in the same performance as the other algorithms.

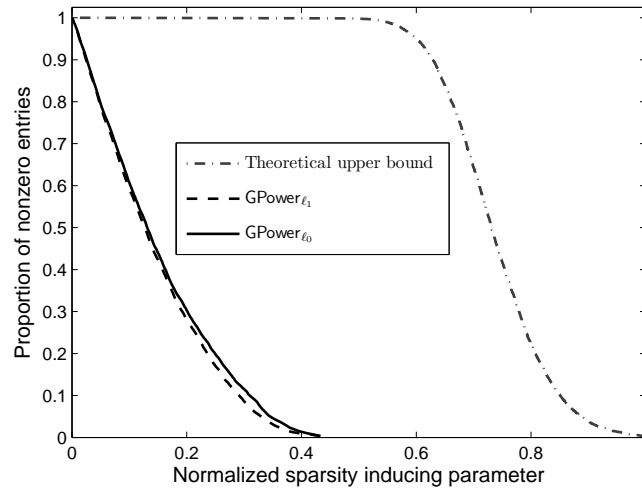


Figure 2: Dependence of cardinality on the value of the sparsity-inducing parameter γ . In case of the GPower_{ℓ_1} algorithm, the horizontal axis shows $\gamma/\|a_{i^*}\|_2$, whereas for the GPower_{ℓ_0} algorithm, we use $\sqrt{\gamma}/\|a_{i^*}\|_2$. The theoretical upper bound is therefor identical for both methods. The plots are averages based on 100 test problems of size $p = 100$ and $n = 300$.

Greedy versus the rest. The considered sparse PCA methods feature different empirical computational complexities. In Figure 3, we display the average time required by the sparse PCA

algorithms to extract one sparse component from Gaussian matrices of dimensions $p = 100$ and $n = 300$. One immediately notices that the greedy method slows down significantly as cardinality increases, whereas the speed of the other considered algorithms does not depend on cardinality. Since on average Greedy is much slower than the other methods, even for low cardinalities, we discard it from all following numerical experiments.




Figure 3: The computational complexity of Greedy grows significantly if it is set out to output a loading vector of increasing cardinality. The speed of the other methods is unaffected by the cardinality target.

Speed and scaling test. In Tables 3 and 4 we compare the speed of the remaining algorithms. Table 3 deals with problems with a fixed aspect ratio $n/p = 10$, whereas in Table 4, p is fixed at 500, and exponentially increasing values of n are considered. For the GPower_{ℓ_1} method, the sparsity inducing parameter γ was set to 10% of the upper bound $\gamma_{\max} = \|a_{i^*}\|_2$. For the GPower_{ℓ_0} method, γ was set to 1% of $\gamma_{\max} = \|a_{i^*}\|_2^2$ in order to aim for solutions of comparable cardinalities (see (5.2)). These two parameters have also been used for the rSVD_{ℓ_1} and the rSVD_{ℓ_0} methods, respectively. Concerning SPCA, the sparsity parameter has been chosen by trial and error to get, on average, solutions with similar cardinalities as obtained by the other methods. The values displayed in Tables 3 and 4 correspond to the average running times of the algorithms on 100 test instances for each problem size. In both tables, the new methods GPower_{ℓ_1} and GPower_{ℓ_0} are the fastest. The difference in speed between GPower_{ℓ_1} and GPower_{ℓ_0} results from different approaches to fill the active part of z : GPower_{ℓ_1} requires to compute a rank-one approximation of a submatrix of A (see Equation (4.2)), whereas the explicit solution (2.11) is available to GPower_{ℓ_0} . The linear complexity of the algorithms in the problem size n is clearly visible in Table 4.

Different convergence mechanisms. Figure 4 illustrates how the trade-off between explained variance and sparsity evolves in the time of computation for the two methods GPower_{ℓ_1} and rSVD_{ℓ_1} . In case of the GPower_{ℓ_1} algorithm, the initialization point (5.1) provides a good approximation of the final cardinality. This method then works on maximizing the variance while keeping the sparsity at a low level throughout. The rSVD_{ℓ_1} algorithm, in contrast, works

$p \times n$	100×1000	250×2500	500×5000	750×7500	1000×10000
GPower_{ℓ_1}	0.10	0.86	2.45	4.28	5.86
GPower_{ℓ_0}	0.03	0.42	1.21	2.07	2.85
SPCA	0.24	2.92	14.5	40.7	82.2
rSVD_{ℓ_1}	0.21	1.45	6.70	17.9	39.7
rSVD_{ℓ_0}	0.20	1.33	6.06	15.7	35.2

Table 3: Average computational time for the extraction of one component (in seconds).

$p \times n$	500×1000	500×2000	500×4000	500×8000	500×16000
GPower_{ℓ_1}	0.42	0.92	2.00	4.00	8.54
GPower_{ℓ_0}	0.18	0.42	0.96	2.14	4.55
SPCA	5.20	7.20	12.0	22.6	44.7
rSVD_{ℓ_1}	1.20	2.53	5.33	11.3	26.7
rSVD_{ℓ_0}	1.09	2.26	4.85	10.5	24.6

Table 4: Average computational time for the extraction of one component (in seconds).

in two steps. First, it maximizes the variance, without enforcing sparsity. This corresponds to computing the first principal component and requires thus a first run of the algorithm with random initialization and a sparsity inducing parameter set at zero. In the second run, this parameter is set to a positive value and the method works to rapidly decrease cardinality at the expense of only a modest decrease in explained variance. So, the new algorithm GPower_{ℓ_1} performs faster primarily because it combines the two phases into one, simultaneously optimizing the trade-off between variance and sparsity.




Figure 4: Evolution of the variance (solid lines and left axis) and cardinality (dashed lines and right axis) in time of computation for the methods GPower_{ℓ_1} and rSVD_{ℓ_1} on a test problem with $p = 250$ and $n = 2500$. The vertical axis is the ratio $\text{Var}(z_{\text{sPCA}})/\text{Var}(z_{\text{PCA}})$, where the loading vector z_{sPCA} is computed by sparse PCA and z_{PCA} is the first principal loading vector. The rSVD_{ℓ_1} algorithm first solves unconstrained PCA, whereas GPower_{ℓ_1} immediately optimizes the trade-off between variance and sparsity.

Extracting more components. Similar numerical experiments, which include the methods $\text{GPower}_{\ell_1, m}$ and $\text{GPower}_{\ell_0, m}$, have been conducted for the extraction of more than one component. A deflation scheme is used by the non-block methods to sequentially compute m components. These experiments lead to similar conclusions as in the single-unit case, i.e, the methods GPower_{ℓ_1} , GPower_{ℓ_0} , $\text{GPower}_{\ell_1, m}$, $\text{GPower}_{\ell_0, m}$ and rSVD_{ℓ_0} outperform the SPCA and rSVD_{ℓ_1} approaches in terms of variance explained at a fixed cardinality. Again, these last two methods can be improved by postprocessing the resulting loading vectors with Algorithm 6, as it is done for $\text{GPower}_{\ell_1, m}$. The average running times for problems of various sizes are listed in Table 5. The new power-like methods are significantly faster on all instances.

$p \times n$	50×500	100×1000	250×2500	500×5000	750×7500
GPower_{ℓ_1}	0.22	0.56	4.62	12.6	20.4
GPower_{ℓ_0}	0.06	0.17	2.15	6.16	10.3
$\text{GPower}_{\ell_1, m}$	0.09	0.28	3.50	12.4	23.0
$\text{GPower}_{\ell_0, m}$	0.05	0.14	2.39	7.7	12.4
SPCA	0.61	1.47	13.4	48.3	113.3
rSVD_{ℓ_1}	0.30	1.15	7.92	37.4	97.4
rSVD_{ℓ_0}	0.28	1.10	7.54	34.7	85.7

Table 5: Average computational time for the extraction of $m = 5$ components (in seconds).

5.2 Analysis of gene expression data

Gene expression data results from DNA microarrays and provide the expression level of thousands of genes across several hundreds of experiments. The interpretation of these huge databases remains a challenge. Of particular interest is the identification of genes that are systematically coexpressed under similar experimental conditions. We refer to Riva et al. [17] and references therein for more details on microarrays and gene expression data. PCA has been intensively applied in this context (e.g., Alter et al. [2]). Further methods for dimension reduction, such as independent component analysis (Liebermeister [13]) or nonnegative matrix factorization (Brunet et al. [4]), have also been used on gene expression data. Sparse PCA, which extracts components involving a few genes only, is expected to enhance interpretation.

Data sets. The results below focus on four major data sets related to breast cancer. They are briefly detailed in Table 6. Each sparse PCA algorithm computes ten components from these data sets.

Study	Samples (p)	Genes (n)	Reference
Vijver	295	13319	van de Vijver et al. [22]
Wang	285	14913	Wang et al. [23]
Naderi	135	8278	Naderi et al. [15]
JRH-2	101	14223	Sotiriou et al. [19]

Table 6: Breast cancer cohorts.

Speed. The average computational time required by the sparse PCA algorithms on each data set is displayed in Table 7. The indicated times are averages on all the computations performed to obtain cardinality ranging from n down to 1.

	Vijver	Wang	Naderi	JRH-2
GPower $_{\ell_1}$	7.72	6.96	2.15	2.69
GPower $_{\ell_0}$	3.80	4.07	1.33	1.73
GPower $_{\ell_1,m}$	5.40	4.37	1.77	1.14
GPower $_{\ell_0,m}$	5.61	7.21	2.25	1.47
SPCA	77.7	82.1	26.7	11.2
rSVD $_{\ell_1}$	46.4	49.3	13.8	15.7
rSVD $_{\ell_0}$	46.8	48.4	13.7	16.5

Table 7: Average computational times (in seconds).

Trade-off curves. Figure 5 plots the proportion of adjusted variance versus the cardinality for the “Vijver” data set. The other data sets have similar plots. As for the random test problems, this performance criterion does not discriminate among the different algorithms. All methods have in fact the same performance, provided that the SPCA and rSVD $_{\ell_1}$ approaches are used with postprocessing by Algorithm 6.




Figure 5: **Trade-off curves** between explained variance and cardinality (case of the “Vijver” data). The vertical axis is the ratio $\text{AdjVar}(Z_{\text{sPCA}})/\text{AdjVar}(Z_{\text{PCA}})$, where the loading vectors Z_{sPCA} are computed by sparse PCA and Z_{PCA} are the m first principal loading vectors.

Interpretability. A more interesting performance criterion is to estimate the biological interpretability of the extracted components. The *pathway enrichment index* (PEI) proposed by Teschendorff et al. [20] measures the statistical significance of the overlap between two kinds of gene sets. The first sets are inferred from the computed components by retaining the most expressed genes, whereas the second sets result from biological knowledge. For instance, metabolic pathways provide sets of genes known to participate together when a certain biological function is required. An alternative is given by the regulatory motifs: genes tagged with an identical motif are likely to be coexpressed. One expects sparse PCA methods to recover some of these biologically significant sets. Table 8 displays the PEI based on 536 metabolic pathways related

to cancer. The PEI is the fraction of these 536 sets presenting a statistically significant overlap with the genes inferred from the sparse principal components. The values in Table 8 correspond to the largest PEI obtained among all possible cardinalities. Similarly, Table 9 is based on 173 motifs. More details on the selected pathways and motifs can be found in Teschendorff et al. [20]. This analysis clearly indicates that the sparse PCA methods perform much better than PCA in this context. Furthermore, the new **GPower** algorithms, and especially the block formulations, provide largest PEI values for both types of biological information. In terms of biological interpretability, they systematically outperform previously published algorithms.

	Vijver	Wang	Naderi	JRH-2
PCA	0.0728	0.0466	0.0149	0.0690
GPower $_{\ell_1}$	0.1493	0.1026	0.0728	0.1250
GPower $_{\ell_1}$	0.1250	0.1250	0.0672	0.1026
GPower $_{\ell_1,m}$	0.1418	0.1250	0.1026	0.1381
GPower $_{\ell_0,m}$	0.1362	0.1287	0.1007	0.1250
SPCA	0.1362	0.1007	0.0840	0.1007
rSVD $_{\ell_1}$	0.1213	0.1175	0.0914	0.0914
rSVD $_{\ell_0}$	0.1175	0.0970	0.0634	0.1063

Table 8: PEI-values based on a set of 536 cancer-related pathways.

	Vijver	Wang	Naderi	JRH-2
PCA	0.0347	0	0.0289	0.0405
GPower $_{\ell_1}$	0.1850	0.0867	0.0983	0.1792
GPower $_{\ell_0}$	0.1676	0.0809	0.0925	0.1908
GPower $_{\ell_1,m}$	0.1908	0.1156	0.1329	0.1850
GPower $_{\ell_0,m}$	0.1850	0.1098	0.1329	0.1734
SPCA	0.1734	0.0925	0.0809	0.1214
rSVD $_{\ell_1}$	0.1387	0.0809	0.1214	0.1503
rSVD $_{\ell_0}$	0.1445	0.0867	0.0867	0.1850

Table 9: PEI-values based on a set of 173 motif-regulatory gene sets.

6 Conclusion

We have proposed two single-unit and two block formulations of the sparse PCA problem and constructed reformulations with several favorable properties. First, the reformulated problems are of the form of maximization of a convex function on a compact set, with the feasible set being either a unit Euclidean sphere or the Stiefel manifold. This structure allows for the design and iteration complexity analysis of a simple gradient scheme which applied to our sparse PCA setting results in four new algorithms for computing sparse principal components of a matrix $A \in \mathbf{R}^{p \times n}$. Second, our algorithms appear to be faster if either the objective function or the feasible set are strongly convex, which holds in the single-unit case and can be enforced in the block case. Third, the dimension of the feasible sets does not depend on n but on p and on the number m of components to be extracted. This is a highly desirable property if $p \ll n$. Last but not least, on random and real-life biological data, our methods systematically outperform the existing algorithms both in speed and trade-off performance. Finally, in the case of the

biological data, the components obtained by our block algorithms deliver the richest biological interpretation as compared to the components extracted by the other methods.

7 Appendix A

In this appendix we characterize a class of functions with strongly convex level sets. First we need to collect some basic preliminary facts. All the inequalities of Proposition 11 are well-known in the literature.

Proposition 11 (i) *If f is a strongly convex function with convexity parameter σ_f , then for all x, y and $0 \leq \alpha \leq 1$,*

$$f(\alpha x + (1 - \alpha)y) \leq \alpha f(x) + (1 - \alpha)f(y) - \frac{\sigma_f}{2}\alpha(1 - \alpha)\|x - y\|^2. \quad (7.1)$$

(ii) *If f is a convex differentiable function and its gradient is Lipschitz continuous with constant L_f , then for all x and h ,*

$$f(x + h) \leq f(x) + \langle f'(x), h \rangle + \frac{L_f}{2}\|h\|^2, \quad (7.2)$$

and

$$\|f'(x)\|_* \leq \sqrt{2L_f(f(x) - f_*)}, \quad (7.3)$$

where $f_* \stackrel{\text{def}}{=} \min_{x \in \mathbf{E}} f(x)$.

We are now ready for the main result of this section.

Theorem 12 (Strongly convex level sets) *Let $f : \mathbf{E} \rightarrow \mathbf{R}$ be a nonnegative strongly convex function with convexity parameter $\sigma_f > 0$. Also assume f has a Lipschitz continuous gradient with Lipschitz constant $L_f > 0$. Then for any $\omega > 0$, the set*

$$\mathcal{Q}_\omega \stackrel{\text{def}}{=} \{x \mid f(x) \leq \omega\}$$

is strongly convex with convexity parameter

$$\sigma_{\mathcal{Q}_\omega} = \frac{\sigma_f}{\sqrt{2\omega L_f}}.$$

Proof:

Consider any $x, y \in \mathcal{Q}_\omega$, scalar $0 \leq \alpha \leq 1$ and let $z_\alpha = \alpha x + (1 - \alpha)y$. Notice that by convexity, $f(z_\alpha) \leq \omega$. For any $u \in \mathbf{E}$,

$$\begin{aligned} f(z_\alpha + u) &\stackrel{(7.2)}{\leq} f(z_\alpha) + \langle f'(z_\alpha), u \rangle + \frac{L_f}{2}\|u\|^2 \\ &\leq f(z_\alpha) + \|f'(z_\alpha)\|\|u\| + \frac{L_f}{2}\|u\|^2 \\ &\stackrel{(7.3)}{\leq} f(z_\alpha) + \sqrt{2L_f(f(z_\alpha) - f_*)}\|u\| + \frac{L_f}{2}\|u\|^2 \\ &= \left(\sqrt{f(z_\alpha)} + \sqrt{\frac{L_f}{2}}\|u\| \right)^2 \\ &\stackrel{(7.1)}{\leq} \left(\sqrt{\omega - \beta} + \sqrt{\frac{L_f}{2}}\|u\| \right)^2, \end{aligned}$$

where

$$\beta = \frac{\sigma_f}{2}\alpha(1-\alpha)\|x-y\|^2. \quad (7.4)$$

In view of (3.6), it remains to show that the last displayed expression is bounded above by ω whenever u is of the form

$$u = \frac{\sigma_{\mathcal{Q}_\omega}}{2}\alpha(1-\alpha)\|x-y\|^2 s = \frac{\sigma_f}{2\sqrt{2\omega L_f}}\alpha(1-\alpha)\|x-y\|^2 s, \quad (7.5)$$

for some $s \in \mathcal{S}$. However, this follows directly from concavity of the scalar function $g(t) = \sqrt{t}$:

$$\begin{aligned} \sqrt{\omega - \beta} &= g(\omega - \beta) \leq g(\omega) - \langle g'(\omega), \beta \rangle \\ &= \sqrt{\omega} - \frac{\beta}{2\sqrt{\omega}} \\ &\stackrel{(7.4)}{\leq} \sqrt{\omega} - \frac{\sigma_f}{4\sqrt{\omega}}\alpha(1-\alpha)\|x-y\|^2 \\ &\stackrel{(7.5)}{\leq} \sqrt{\omega} - \sqrt{\frac{L_f}{2}}\|u\|. \end{aligned}$$

□

Example 13 Let $f(x) = \|x\|^2$. Note that $\sigma_f = L_f = 2$. If we let $\omega = r^2$, then

$$\mathcal{Q}_\omega = \{x \mid f(x) \leq \omega\} = \{x \mid \|x\| \leq r\} = r \cdot \mathcal{B}.$$

We have shown before (see the discussion immediately following Assumption 3), that the strong convexity parameter of this set is $\sigma_{\mathcal{Q}_\omega} = \frac{1}{r}$. Note that we recover this as a special case of Theorem 12:

$$\sigma_{\mathcal{Q}_\omega} = \frac{\sigma_f}{\sqrt{2\omega L_f}} = \frac{1}{r}.$$

References

- [1] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, January 2008.
- [2] O. Alter, P. O. Brown, and D. Botstein. Generalized singular value decomposition for comparative analysis of genome-scale expression data sets of two different organisms. *Proc Natl Acad Sci USA*, 100(6):3351–3356, 2003.
- [3] R. W. Brockett. Dynamical systems that sort lists, diagonalize matrices and solve linear programming problems. *Linear Algebra Appl.*, 146:79–91, 1991.
- [4] J. P. Brunet, P. Tamayo, T. R. Golub, and J. P. Mesirov. Metagenes and molecular pattern discovery using matrix factorization. *Proc Natl Acad Sci USA*, 101(12):4164–4169, 2004.
- [5] J. Cadima and I. T. Jolliffe. Loadings and correlations in the interpretation of principal components. *Journal of Applied Statistics*, 22:203–214, 1995.

- [6] A. d’Aspremont, F. R. Bach, and L. El Ghaoui. Optimal solutions for sparse principal component analysis. *Journal of Machine Learning Research*, 9:1269–1294, 2008.
- [7] A. d’Aspremont, L. El Ghaoui, M. I. Jordan, and G. R. G. Lanckriet. A direct formulation for sparse PCA using semidefinite programming. *Siam Review*, 49:434–448, 2007.
- [8] C. Fraikin, Yu. Nesterov, and P. Van Dooren. A gradient-type algorithm optimizing the coupling between matrices. *Linear Algebra and its Applications*, 429(5-6):1229–1242, 2008.
- [9] G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, 1996.
- [10] R. A. Horn and C. A. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, UK, 1985.
- [11] I. T. Jolliffe. Rotation of principal components: choice of normalization constraints. *Journal of Applied Statistics*, 22:29–35, 1995.
- [12] 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(3):531–547, 2003.
- [13] W. Liebermeister. Linear modes of gene expression determined by independent component analysis. *Bioinformatics*, 18(1):51–60, 2002.
- [14] B. Moghaddam, Y. Weiss, and S. Avidan. Spectral bounds for sparse PCA: Exact and greedy algorithms. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 915–922. MIT Press, Cambridge, MA, 2006.
- [15] A. Naderi, A. E. Teschendorff, N. L. Barbosa-Morais, S. E. Pinder, A. R. Green, D. G. Powe, J. F. R. Robertson, S. Aparicio, I. O. Ellis, J. D. Brenton, and C. Caldas. A gene expression signature to predict survival in breast cancer across independent data sets. *Oncogene*, 26:1507–1516, 2007.
- [16] Beresford N. Parlett. *The symmetric eigenvalue problem*. Prentice-Hall Inc., Englewood Cliffs, N.J., 1980. Prentice-Hall Series in Computational Mathematics.
- [17] A. Riva, A.-S. Carpentier, B. Torrèsani, and A. Hénaut. Comments on selected fundamental aspects of microarray analysis. *Computational Biology and Chemistry*, 29(5):319–336, 2005.
- [18] 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.
- [19] C. Sotiriou, P. Wirapati, S. Loi, A. Harris, S. Fox, J. Smeds, H. Nordgren, P. Farmer, V. Praz, B. Haibe-Kains, C. Desmedt, D. Larsimont, F. Cardoso, H. Peterse, D. Nuyten, M. Buyse, M. J. Van de Vijver, J. Bergh, M. Piccart, and M. Delorenzi. Gene expression profiling in breast cancer: understanding the molecular basis of histologic grade to improve prognosis. *J Natl Cancer Inst*, 98(4):262–272, 2006.
- [20] A. Teschendorff, M. Journée, P.-A. Absil, R. Sepulchre, and C. Caldas. Elucidating the altered transcriptional programs in breast cancer using independent component analysis. *PLoS Computational Biology*, 3(8):1539–1554, 2007.

- [21] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58(2):267–288, 1996.
- [22] M. J. van de Vijver, Y. D. He, L. J. van't Veer, H. Dai, A. A. Hart, D. W. Voskuil, G. J. Schreiber, J. L. Peterse, C. Roberts, M. J. Marton, M. Parrish, D. Atsma, A. Witteveen, A. Glas, L. Delahaye, T. van der Velde, H. Bartelink, S. Rodenhuis, E. T. Rutgers, S. H. Friend, and R. Bernards. A gene-expression signature as a predictor of survival in breast cancer. *N Engl J Med*, 347(25):1999–2009, 2002.
- [23] Y. Wang, J. G. Klijn, Y. Zhang, A. M. Sieuwerts, M. P. Look, F. Yang, D. Talantov, M. Timmermans, M. E. Meijer-van Gelder, J. Yu, T. Jatkoe, E. M. Berns, D. Atkins, and J. A. Foekens. Gene-expression profiles to predict distant metastasis of lymph-node-negative primary breast cancer. *Lancet*, 365(9460):671–679, 2005.
- [24] H. Zou, T. Hastie, and R. Tibshirani. Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, 15(2):265–286, 2006.

Recent titles

CORE Discussion Papers

- 2008/32. Jean J. GABSZEWICZ, Didier LAUSSEL and Nathalie SONNAC. The TV news scheduling game when the newscaster's face matters.
- 2008/33. Didier LAUSSEL and Joana RESENDE. Does the absence of competition *in* the market foster competition *for* the market? A dynamic approach to aftermarkets.
- 2008/34. Vincent D. BLONDEL and Yurii NESTEROV. Polynomial-time computation of the joint spectral radius for some sets of nonnegative matrices.
- 2008/35. David DE LA CROIX and Clara DELAVALLADE. Democracy, rule of law, corruption incentives and growth.
- 2008/36. Jean J. GABSZEWICZ and Joana RESENDE. Uncertain quality, product variety and price competition. 2008/37. Gregor ZOETTL. On investment decisions in liberalized electricity markets: the impact of price caps at the spot market.
- 2008/38. Helmuth CREMER, Philippe DE DONDER, Dario MALDONADO and Pierre PESTIEAU. Habit formation and labor supply.
- 2008/39. Marie-Louise LEROUX and Grégory PONTIERE. Optimal tax policy and expected longevity: a mean and variance approach.
- 2008/40. Kristian BEHRENS and Pierre M. PICARD. Transportation, freight rates, and economic geography.
- 2008/41. Gregor ZOETTL. Investment decisions in liberalized electricity markets: A framework of peak load pricing with strategic firms.
- 2008/42. Raouf BOUCEKKINE, Rodolphe DESBORDES and Hélène LATZER. How do epidemics induce behavioral changes?
- 2008/43. David DE LA CROIX and Marie VANDER DONCKT. Would empowering women initiate the demographic transition in least-developed countries?
- 2008/44. Geoffrey CARUSO, Dominique PEETERS, Jean CAVAILHES and Mark ROUNSEVELL. Space-time patterns of urban sprawl, a 1D cellular automata and microeconomic approach.
- 2008/45. Taoufik BOUEZMARNI, Jeroen V.K. ROMBOUTS and Abderrahim TAAMOUTI. Asymptotic properties of the Bernstein density copula for dependent data.
- 2008/46. Joe THARAKAN and Jean-Philippe TROPEANO. On the impact of labor market matching on regional disparities.
- 2008/47. Shin-Huei WANG and Cheng HSIAO. An easy test for two stationary long processes being uncorrelated via AR approximations.
- 2008/48. David DE LA CROIX. Adult longevity and economic take-off: from Malthus to Ben-Porath.
- 2008/49. David DE LA CROIX and Gregory PONTIERE. On the Golden Rule of capital accumulation under endogenous longevity.
- 2008/50. Jean J. GABSZEWICZ and Skerdilajda ZANAJ. Successive oligopolies and decreasing returns.
- 2008/51. Marie-Louise LEROUX, Pierre PESTIEAU and Grégory PONTIERE. Optimal linear taxation under endogenous longevity.
- 2008/52. Yuri YATSENKO, Raouf BOUCEKKINE and Natali HRITONENKO. Estimating the dynamics of R&D-based growth models.
- 2008/53. Roland Iwan LUTTENS and Marie-Anne VALFORT. Voting for redistribution under desert-sensitive altruism.
- 2008/54. Sergei PEKARSKI. Budget deficits and inflation feedback. 2008/55. Raouf BOUCEKKINE, Jacek B. KRAWCZYK and Thomas VALLEE. Towards an understanding of tradeoffs between regional wealth, tightness of a common environmental constraint and the sharing rules.
- 2008/56. Santanu S. DEY. A note on the split rank of intersection cuts.
- 2008/57. Yu. NESTEROV. Primal-dual interior-point methods with asymmetric barriers.
- 2008/58. Marie-Louise LEROUX, Pierre PESTIEAU and Gregory PONTIERE. Should we subsidize longevity?
- 2008/59. J. Roderick McCORIE. The role of Skorokhod space in the development of the econometric analysis of time series.

Recent titles

CORE Discussion Papers - continued

- 2008/60. Yu. NESTEROV. Barrier subgradient method.
- 2008/61. Thierry BRECHET, Johan EYCKMANS, François GERARD, Philippe MARBAIX, Henry TULKENS and Jean-Pascal VAN YPERSELE. The impact of the unilateral EU commitment on the stability of international climate agreements.
- 2008/62. Giorgia OGGIONI and Yves SMEERS. Average power contracts can mitigate carbon leakage.
- 2008/63. Jean-Sébastien TANCREZ, Philippe CHEVALIER and Pierre SEMAL. A tight bound on the throughput of queueing networks with blocking.
- 2008/64. Nicolas GILLIS and François GLINEUR. Nonnegative factorization and the maximum edge biclique problem.
- 2008/65. Geir B. ASHEIM, Claude D'ASPREMONT and Kuntal BANERJEE. Generalized time-invariant overtaking.
- 2008/66. Jean-François CAULIER, Ana MAULEON and Vincent VANNETELBOSCH. Contractually stable networks.
- 2008/67. Jean J. GABSZEWICZ, Filomena GARCIA, Joana PAIS and Joana RESENDE. On Gale and Shapley '*College admissions and stability of marriage*'.
- 2008/68. Axel GAUTIER and Anne YVRANDE-BILLON. Contract renewal as an incentive device. An application to the French urban public transport sector.
- 2008/69. Yuri YATSENKO and Natali HRITONENKO. Discrete-continuous analysis of optimal equipment replacement.
- 2008/70. Michel JOURNÉE, Yurii NESTEROV, Peter RICHTÁRIK and Rodolphe SEPULCHRE. Generalized power method for sparse principal component analysis.

Books

- Y. Pochet and L. Wolsey (eds.) (2006), *Production planning by mixed integer programming*. New York, Springer-Verlag.
- P. Pestieau (ed.) (2006), *The welfare state in the European Union: economic and social perspectives*. Oxford, Oxford University Press.
- H. Tulkens (ed.) (2006), *Public goods, environmental externalities and fiscal competition*. New York, Springer-Verlag.
- V. Ginsburgh and D. Throsby (eds.) (2006), *Handbook of the economics of art and culture*. Amsterdam, Elsevier.
- J. Gabszewicz (ed.) (2006), *La différenciation des produits*. Paris, La découverte.
- L. Bauwens, W. Pohlmeier and D. Veredas (eds.) (2008), *High frequency financial econometrics: recent developments*. Heidelberg, Physica-Verlag.
- P. Van Hentenrycke and L. Wolsey (eds.) (2007), *Integration of AI and OR techniques in constraint programming for combinatorial optimization problems*. Berlin, Springer.

CORE Lecture Series

- C. GOURIÉROUX and A. MONFORT (1995), *Simulation Based Econometric Methods*.
- A. RUBINSTEIN (1996), *Lectures on Modeling Bounded Rationality*.
- J. RENEGAR (1999), *A Mathematical View of Interior-Point Methods in Convex Optimization*.
- B.D. BERNHEIM and M.D. WHINSTON (1999), *Anticompetitive Exclusion and Foreclosure Through Vertical Agreements*.
- D. BIENSTOCK (2001), *Potential function methods for approximately solving linear programming problems: theory and practice*.
- R. AMIR (2002), *Supermodularity and complementarity in economics*.
- R. WEISMANTEL (2006), *Lectures on mixed nonlinear programming*.