

Sparse Generalized Eigenvalue Problem Via Smooth Optimization

Junxiao Song, Prabhu Babu, and Daniel P. Palomar, *Fellow, IEEE*

Abstract—In this paper, we consider an ℓ_0 -norm penalized formulation of the generalized eigenvalue problem (GEP), aimed at extracting the leading sparse generalized eigenvector of a matrix pair. The formulation involves maximization of a discontinuous nonconcave objective function over a nonconvex constraint set, and is therefore computationally intractable. To tackle the problem, we first approximate the ℓ_0 -norm by a continuous surrogate function. Then an algorithm is developed via iteratively majorizing the surrogate function by a quadratic separable function, which at each iteration reduces to a regular generalized eigenvalue problem. A preconditioned steepest ascent algorithm for finding the leading generalized eigenvector is provided. A systematic way based on smoothing is proposed to deal with the “singularity issue” that arises when a quadratic function is used to majorize the nondifferentiable surrogate function. For sparse GEPs with special structure, algorithms that admit a closed-form solution at every iteration are derived. Numerical experiments show that the proposed algorithms match or outperform existing algorithms in terms of computational complexity and support recovery.

Index Terms—Minorization-maximization, sparse generalized eigenvalue problem, sparse PCA, smooth optimization.

I. INTRODUCTION

THE generalized eigenvalue problem (GEP) for matrix pair (\mathbf{A}, \mathbf{B}) is the problem of finding a pair (λ, \mathbf{x}) such that

$$\mathbf{A}\mathbf{x} = \lambda\mathbf{B}\mathbf{x}, \quad (1)$$

where $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$, $\lambda \in \mathbb{R}$ is called the generalized eigenvalue and $\mathbf{x} \in \mathbb{R}^n, \mathbf{x} \neq \mathbf{0}$ is the corresponding generalized eigenvector. When \mathbf{B} is the identity matrix, the problem in (1) reduces to the simple eigenvalue problem.

GEP is extremely useful in numerous applications of high dimensional data analysis and machine learning. Many widely used data analysis tools, such as principle component analysis (PCA) and canonical correlation analysis (CCA), are special instances of the generalized eigenvalue problem [1], [2]. In these applications, usually $\mathbf{A} \in \mathbb{S}^n$, $\mathbf{B} \in \mathbb{S}_{++}^n$ (i.e., \mathbf{A} is symmetric

and \mathbf{B} is positive definite) and only a few of the largest generalized eigenvalues are of interest. In this case, all generalized eigenvalues λ and generalized eigenvectors \mathbf{x} are real and the largest generalized eigenvalue can be formulated as the following optimization problem

$$\lambda_{\max}(\mathbf{A}, \mathbf{B}) = \max_{\mathbf{x} \neq \mathbf{0}} \frac{\mathbf{x}^T \mathbf{A} \mathbf{x}}{\mathbf{x}^T \mathbf{B} \mathbf{x}}, \quad (2)$$

or equivalently

$$\lambda_{\max}(\mathbf{A}, \mathbf{B}) = \max_{\mathbf{x}} \{\mathbf{x}^T \mathbf{A} \mathbf{x} : \mathbf{x}^T \mathbf{B} \mathbf{x} = 1\}. \quad (3)$$

Despite the simplicity and popularity of the tools based on GEP, there is a potential problem: in general the eigenvector is not expected to have many zero entries, which makes the result difficult to interpret, especially when dealing with high dimensional data. An ad hoc approach to fix this problem is to set the entries with absolute values smaller than a threshold to zero. This thresholding approach is frequently used in practice, but it is found to be potentially misleading, since no care is taken on how well the artificially enforced sparsity fits the original data [3]. Obviously, approaches that can simultaneously produce accurate and sparse models are more desirable.

This has motivated active research in developing methods that enforce sparsity on eigenvectors, and many approaches have been proposed, especially for the simple sparse PCA case. For instance, Zou, Hastie, and Tibshirani [4] first recast the PCA problem as a ridge regression problem and then imposed ℓ_1 -norm penalty to encourage sparsity. In [5], d’Aspremont *et al.* proposed a convex relaxation for the sparse PCA problem based on semidefinite programming (SDP) and Nesterov’s smooth minimization technique was applied to solve the SDP. Shen and Huang [6] exploited the connection of PCA with singular value decomposition (SVD) of the data matrix and extracted the sparse principal components (PCs) through solving a regularized low rank matrix approximation problem. Journée *et al.* [7] rewrote the sparse PCA problem in the form of an optimization problem involving maximization of a convex function on a compact set and the simple gradient method was then applied. Although derived differently, the resulting algorithm GPower turns out to be identical to the rSVD algorithm in [6], except for the initialization and post-processing phases. Very recently, Luss and Teboulle [8] introduced an algorithm framework, called ConGradU, based on the well-known conditional gradient algorithm, that unifies a variety of seemingly different algorithms, including the GPower method and the rSVD method. Based on ConGradU, the authors also proposed a new algorithm for the ℓ_0 -constrained sparse PCA formulation.

Manuscript received June 13, 2014; revised November 13, 2014; accepted January 04, 2015. Date of publication January 21, 2015; date of current version February 23, 2015. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Rui Zhang. This work was supported by the Hong Kong RGC 617312 research grant.

The authors are with the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology (HKUST), Hong Kong (e-mail: jsong@ust.hk; eprabhubabu@ust.hk; palomar@ust.hk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2015.2394443

Among the aforementioned algorithms for sparse PCA, rSVD, GPower and ConGradU are very efficient and require only matrix vector multiplications at every iteration, thus can be applied to problems of extremely large size. But these algorithms are not well suited for the case where \mathbf{B} is not the identity matrix, for example, the sparse CCA problem, and direct application of these algorithms does not yield a simple closed-form solution at each iteration any more. To deal with this problem, [9] suggested that good results could still be obtained by substituting in the identity matrix for \mathbf{B} and, in [8], the authors proposed to substitute the matrix \mathbf{B} with its diagonal instead. In [10], [11], an algorithm was proposed to solve the problem with the general \mathbf{B} (to the best of our knowledge, this is the only one) based on D.C. (difference of convex functions) programming and minorization-maximization. The resulting algorithm requires computing a matrix pseudoinverse and solving a quadratic program (QP) at every iteration when \mathbf{A} is symmetric and positive semidefinite, and in the case where \mathbf{A} is just symmetric it needs to solve a quadratically constrained quadratic program (QCQP) at each iteration. It is computationally intensive and not amenable to problems of large size. The same algorithm can also be applied to the simple sparse PCA problem by simply restricting \mathbf{B} to be the identity matrix, and in this special case only one matrix vector multiplication is needed at every iteration and it is shown to be comparable to the GPower method regarding the computational complexity.

In this paper, we adopt the MM (majorization-minimization or minorization-maximization) approach to develop efficient algorithms for the sparse generalized eigenvalue problem. In fact, all the algorithms that can be unified by the ConGradU framework can be seen as special cases of the MM method. Since the ConGradU framework is based on maximizing a convex function over a compact set via linearizing the convex objective, and the linear function is just a special minorization function of the convex objective. Instead of only considering linear minorization function, in this paper we consider quadratic separable minorization that is related to the well known iteratively reweighted least squares (IRLS) algorithm [12]. By applying quadratic minorization functions, we turn the original sparse generalized eigenvalue problem into a sequence of regular generalized eigenvalue problems and an efficient preconditioned steepest ascent algorithm for finding the leading generalized eigenvector is provided. We call the resulting algorithm IRQM (iteratively reweighted quadratic minorization); it is in spirit similar to IRLS which solves the ℓ_1 -norm minimization problem by solving a sequence of least squares problems. Algorithms of the IRLS type often suffer from the infamous “singularity issue”, i.e., when using quadratic majorization functions for nondifferentiable functions, the variable may get stuck at a nondifferentiable point [13]. To deal with this “singularity issue”, we propose a systematic way via smoothing the nondifferentiable surrogate function, which is inspired by Nesterov’s smooth minimization technique for nonsmooth convex optimization [14], although in our case the surrogate function is nonconvex. The smoothed problem is shown to be equivalent to a problem that maximizes a convex objective over a convex constraint set and the convergence of the IRQM algorithm to a stationary point of the equivalent problem is

proved. For some sparse generalized eigenvalue problems with special structure, more efficient algorithms are also derived which admit a closed-form solution at every iteration.

The remaining sections of the paper are organized as follows. In Section II, the problem formulation of the sparse generalized eigenvalue problem is presented and the surrogate functions that will be used to approximate ℓ_0 -norm are discussed. In Section III, we first give a brief review of the MM framework and then algorithms based on the MM framework are derived for the sparse generalized eigenvalue problems in general and with special structure. A systematic way to deal with the “singularity issue” arising when using quadratic minorization functions is also proposed. In Section IV, the convergence of the proposed MM algorithms is analyzed. Section V presents numerical experiments and the conclusions are given in Section VI.

Notation: \mathbf{R} and \mathbf{C} denote the real field and the complex field, respectively. $\text{Re}(\cdot)$ and $\text{Im}(\cdot)$ denote the real and imaginary part, respectively. \mathbf{R}^n (\mathbf{R}_+^n , \mathbf{R}_{++}^n) denotes the set of (non-negative, strictly positive) real vectors of size n . \mathbf{S}^n (\mathbf{S}_+^n , \mathbf{S}_{++}^n) denotes the set of symmetric (positive semidefinite, positive definite) $n \times n$ matrices defined over \mathbf{R} . Boldface upper case letters denote matrices, boldface lower case letters denote column vectors, and italics denote scalars. The superscripts $(\cdot)^T$ and $(\cdot)^H$ denote transpose and conjugate transpose, respectively. $X_{i,j}$ denotes the $(i$ -th, j -th) element of matrix \mathbf{X} and x_i denotes the i -th element of vector \mathbf{x} . $\mathbf{X}_{i,:}$ denotes the i -th row of matrix \mathbf{X} , $\mathbf{X}_{:,j}$ denotes the j -th column of matrix \mathbf{X} . $\text{diag}(\mathbf{X})$ is a column vector consisting of all the diagonal elements of \mathbf{X} . $\text{Diag}(\mathbf{x})$ is a diagonal matrix formed with \mathbf{x} as its principal diagonal. Given a vector $\mathbf{x} \in \mathbf{R}^n$, $\|\mathbf{x}\|_0$ denotes the number of non-zero elements of \mathbf{x} , $\|\mathbf{x}\|_p := (\sum_{i=1}^n |x_i|^p)^{1/p}$, $0 < p < \infty$. \mathbf{I}_n denotes an $n \times n$ identity matrix. $\text{sgn}(x)$ denotes the sign function, which takes $-1, 0, 1$ if $x < 0, x = 0, x > 0$, respectively.

II. PROBLEM FORMULATION

Given a symmetric matrix $\mathbf{A} \in \mathbf{S}^n$ and a symmetric positive definite matrix $\mathbf{B} \in \mathbf{S}_{++}^n$, the main problem of interest is the following ℓ_0 -norm regularized generalized eigenvalue problem

$$\begin{aligned} & \underset{\mathbf{x}}{\text{maximize}} && \mathbf{x}^T \mathbf{A} \mathbf{x} - \rho \|\mathbf{x}\|_0 \\ & \text{subject to} && \mathbf{x}^T \mathbf{B} \mathbf{x} = 1, \end{aligned} \quad (4)$$

where $\rho > 0$ is the regularization parameter. This formulation is general enough and includes some sparse PCA and sparse CCA formulations in the literature as special cases.

The problem (4) involves the maximization of a non-concave discontinuous objective over a nonconvex set, thus really hard to deal with directly. The intractability of the problem is not only due to the nonconvexity, but also due to the discontinuity of the cardinality function in the objective. A natural approach to deal with the discontinuity of the ℓ_0 -norm is to approximate it by some continuous function. It is easy to see that the ℓ_0 -norm can be written as

$$\|\mathbf{x}\|_0 = \sum_{i=1}^n \text{sgn}(|x_i|).$$




Fig. 1. Three surrogate functions $g_p(x)$ that are used to approximate $\text{sgn}(|x|)$, $p = 0.2$.

Thus, to approximate $\|\mathbf{x}\|_0$, we may just replace the problematic $\text{sgn}(|x_i|)$ by some nicer surrogate function $g_p(x_i)$, where $p > 0$ is a parameter that controls the approximation. In this paper, we will consider the class of continuous even functions defined on \mathbf{R} , which are differentiable everywhere except at 0 and concave and monotone increasing on $[0, +\infty)$ and $g_p(0) = 0$. In particular, we will consider the following three surrogate functions:

- 1) $g_p(x) = |x|^p$, $0 < p \leq 1$
- 2) $g_p(x) = \log(1 + |x|/p)/\log(1 + 1/p)$, $p > 0$
- 3) $g_p(x) = 1 - e^{-|x|/p}$, $p > 0$.

The first is the p -norm-like measure (with $p \leq 1$) used in [15], [16], which is shown to perform well in promoting sparse solutions for compressed sensing problems. The second is the penalty function used in [11] for sparse generalized eigenvalue problem and when used to replace the ℓ_1 -norm in basis pursuit, it leads to the well known iteratively reweighted ℓ_1 -norm minimization algorithm [17]. The last surrogate function is used in [18] for feature selection problems, which is different from the first two surrogate functions in the sense that it has the additional property of being a lower bound of the function $\text{sgn}(|x|)$. To provide an intuitive idea about how these surrogate functions look like, they are plotted in Fig. 1 for fixed $p = 0.2$.

By approximating $\|\mathbf{x}\|_0$ with $\sum_{i=1}^n g_p(x_i)$, the original problem (4) is approximated by the following problem

$$\begin{aligned} & \underset{\mathbf{x}}{\text{maximize}} \quad \mathbf{x}^T \mathbf{A} \mathbf{x} - \rho \sum_{i=1}^n g_p(x_i) \\ & \text{subject to} \quad \mathbf{x}^T \mathbf{B} \mathbf{x} = 1. \end{aligned} \quad (5)$$

With the approximation, the problem (5) is still a nonconvex nondifferentiable optimization problem, but it is a continuous problem now in contrast to the original problem (4). In the following section, we will concentrate on the approximate problem (5) and develop fast algorithms to solve it based on the MM (majorization-minimization or minorization-maximization) scheme.

Note that for simplicity of exposition, we focus on real-valued matrices throughout the paper. However, the techniques developed in this paper can be adapted for complex-valued matrix pair (\mathbf{A}, \mathbf{B}) , with \mathbf{A} being an $n \times n$ Hermitian matrix and \mathbf{B}

being an $n \times n$ Hermitian positive definite matrix. One approach is to transform the problem to a real-valued one by defining

$$\begin{aligned} \tilde{\mathbf{x}} &= [\text{Re}(\mathbf{x})^T, \text{Im}(\mathbf{x})^T]^T \\ \tilde{\mathbf{A}} &= \begin{bmatrix} \text{Re}(\mathbf{A}) & -\text{Im}(\mathbf{A}) \\ \text{Im}(\mathbf{A}) & \text{Re}(\mathbf{A}) \end{bmatrix}, \tilde{\mathbf{B}} = \begin{bmatrix} \text{Re}(\mathbf{B}) & -\text{Im}(\mathbf{B}) \\ \text{Im}(\mathbf{B}) & \text{Re}(\mathbf{B}) \end{bmatrix}. \end{aligned}$$

In this approach, the cardinality is considered for the real and imaginary part of the vector \mathbf{x} separately. A more natural approach is to consider directly the complex-valued version of the ℓ_0 -norm regularized generalized eigenvalue problem

$$\begin{aligned} & \underset{\mathbf{x} \in \mathbf{C}^n}{\text{maximize}} \quad \mathbf{x}^H \mathbf{A} \mathbf{x} - \rho \|\mathbf{x}\|_0 \\ & \text{subject to} \quad \mathbf{x}^H \mathbf{B} \mathbf{x} = 1, \end{aligned}$$

where the ℓ_0 -norm can still be written as $\|\mathbf{x}\|_0 = \sum_{i=1}^n \text{sgn}(|x_i|)$, but with $|x_i|$ being the modulus of x_i now. Notice that the three surrogate functions $g_p(x)$ used to approximate $\text{sgn}(|x|)$ are all functions of $|x|$, by taking $|x|$ as the modulus of a complex number, the surrogate functions are directly applicable to the complex case. The quadratic minorization function that will be described in Section III can also be constructed similarly in the complex case and at each iteration of the resulting algorithm we still need to solve a regular generalized eigenvalue problem but with complex-valued matrices.

III. SPARSE GENERALIZED EIGENVALUE PROBLEM VIA MM SCHEME

A. The MM Method

The MM method refers to the majorization-minimization method or the minorization-maximization method, which is a generalization of the well known expectation maximization (EM) algorithm. It is an approach to solve optimization problems that are too difficult to solve directly. The principle behind the MM method is to transform a difficult problem into a series of simple problems. Interested readers may refer to [19] and references therein for more details (recent generalizations include [20], [21]).

Suppose we want to minimize $f(\mathbf{x})$ over $\mathcal{X} \in \mathbf{R}^n$. Instead of minimizing the cost function $f(\mathbf{x})$ directly, the MM approach optimizes a sequence of approximate objective functions that majorize $f(\mathbf{x})$. More specifically, starting from a feasible point $\mathbf{x}^{(0)}$, the algorithm produces a sequence $\{\mathbf{x}^{(k)}\}$ according to the following update rule

$$\mathbf{x}^{(k+1)} \in \arg \min_{\mathbf{x} \in \mathcal{X}} u(\mathbf{x}, \mathbf{x}^{(k)}), \quad (6)$$

where $\mathbf{x}^{(k)}$ is the point generated by the algorithm at iteration k , and $u(\mathbf{x}, \mathbf{x}^{(k)})$ is the majorization function of $f(\mathbf{x})$ at $\mathbf{x}^{(k)}$. Formally, the function $u(\mathbf{x}, \mathbf{x}^{(k)})$ is said to majorize the function $f(\mathbf{x})$ at the point $\mathbf{x}^{(k)}$ provided

$$u(\mathbf{x}, \mathbf{x}^{(k)}) \geq f(\mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{X}, \quad (7)$$

$$u(\mathbf{x}^{(k)}, \mathbf{x}^{(k)}) = f(\mathbf{x}^{(k)}). \quad (8)$$

In other words, function $u(\mathbf{x}, \mathbf{x}^{(k)})$ is a global upper bound for $f(\mathbf{x})$ and coincides with $f(\mathbf{x})$ at $\mathbf{x}^{(k)}$.

It is easy to show that with this scheme, the objective value is decreased monotonically at every iteration, i.e.,

$$f(\mathbf{x}^{(k+1)}) \leq u(\mathbf{x}^{(k+1)}, \mathbf{x}^{(k)}) \leq u(\mathbf{x}^{(k)}, \mathbf{x}^{(k)}) = f(\mathbf{x}^{(k)}). \quad (9)$$

The first inequality and the third equality follow from the properties of the majorization function, namely (7) and (8) respectively and the second inequality follows from (6).

Note that with straightforward changes, similar scheme can be applied to maximization. To maximize a function $f(\mathbf{x})$, we need to minorize it by a surrogate function $u(\mathbf{x}, \mathbf{x}^{(k)})$ and maximize $u(\mathbf{x}, \mathbf{x}^{(k)})$ to produce the next iterate $\mathbf{x}^{(k+1)}$. A function $u(\mathbf{x}, \mathbf{x}^{(k)})$ is said to minorize the function $f(\mathbf{x})$ at the point $\mathbf{x}^{(k)}$ if $-u(\mathbf{x}, \mathbf{x}^{(k)})$ majorizes $-f(\mathbf{x})$ at $\mathbf{x}^{(k)}$. This scheme refers to minorization-maximization and similarly it is easy to shown that with this scheme the objective value is increased at each iteration.

B. Quadratic Minorization Function

Having briefly introduced the general MM framework, let us return to the approximate sparse generalized eigenvalue problem (SGEP) in (5). To apply the MM scheme, the key step is to find an appropriate minorization function for the objective of (5) at each iteration such that the resulting problem is easy to solve. To construct such a minorization function, we keep the quadratic term $\mathbf{x}^T \mathbf{A} \mathbf{x}$ and only minorize the penalty term $-\rho \sum_{i=1}^n g_p(x_i)$ (i.e., majorize $\rho \sum_{i=1}^n g_p(x_i)$). More specifically, at iteration k we majorize each of the surrogate functions $g_p(x_i)$, $i = 1, \dots, n$ at $x_i^{(k)}$ with a quadratic function $w_i^{(k)} x_i^2 + c_i^{(k)}$, where the coefficients $w_i^{(k)}$ and $c_i^{(k)}$ are determined by the following two conditions (for $x_i^{(k)} \neq 0$):

$$g_p(x_i^{(k)}) = w_i^{(k)} (x_i^{(k)})^2 + c_i^{(k)}, \quad (10)$$

$$g'_p(x_i^{(k)}) = 2w_i^{(k)} x_i^{(k)}, \quad (11)$$

i.e., the quadratic function coincides with the surrogate function $g_p(x_i)$ at $x_i^{(k)}$ and is also tangent to $g_p(x_i)$ at $x_i^{(k)}$. Due to the fact that the surrogate functions of interest are differentiable and concave for $x_i > 0$ (also for $x_i < 0$), the second condition implies that the quadratic function is a global upper bound of the surrogate function $g_p(x_i)$. Then the objective of (5), i.e., $\mathbf{x}^T \mathbf{A} \mathbf{x} - \rho \sum_{i=1}^n g_p(x_i)$, is minorized by the quadratic function $\mathbf{x}^T \mathbf{A} \mathbf{x} - \rho \sum_{i=1}^n (w_i^{(k)} x_i^2 + c_i^{(k)})$, which can be written more compactly as

$$\mathbf{x}^T (\mathbf{A} - \rho \text{Diag}(\mathbf{w}^{(k)})) \mathbf{x} - \rho \sum_{i=1}^n c_i^{(k)}, \quad (12)$$

where $\mathbf{w}^{(k)} = [w_1^{(k)}, \dots, w_n^{(k)}]^T$.

Example 1: To compute the quadratic function $w_i^{(k)} x_i^2 + c_i^{(k)}$ that majorizes the surrogate function $g_p(x_i) = |x_i|^p$, $0 < p \leq 1$ at $x_i^{(k)} \neq 0$, we have the following two equations corresponding to (10) and (11) respectively:

$$|x_i^{(k)}|^p = w_i^{(k)} (x_i^{(k)})^2 + c_i^{(k)}, \quad (13)$$

$$\text{sgn}(x_i^{(k)}) p |x_i^{(k)}|^{p-1} = 2w_i^{(k)} x_i^{(k)}. \quad (14)$$




Fig. 2. The function $g_p(x_i) = |x_i|^p$ with $p = 0.5$ and its quadratic majorization function at $x_i^{(k)} = 2$.

By solving (13) and (14), we can get the quadratic majorization function

$$u(x_i, x_i^{(k)}) = \frac{p}{2} |x_i^{(k)}|^{p-2} x_i^2 + \left(1 - \frac{p}{2}\right) |x_i^{(k)}|^p, \quad (15)$$

which is illustrated in Fig. 2 with $p = 0.5$ and $x_i^{(k)} = 2$.

In fact, the idea of majorizing some penalty functions by quadratic separable functions is well known in robust regression [22]. It was first proposed to solve the absolute deviations curve fitting problem (i.e., regression with ℓ_1 -norm cost function) via iteratively solving a series of weighted least squares problems and the resulting algorithm is known as the iteratively reweighted least squares (IRLS) algorithm [12]. Later the idea was applied in signal processing for sparse signal reconstruction in [15], [23]. Recently, the IRLS approach has also been applied in Compressed Sensing [16]. Algorithms based on this idea often have the infamous singularity issue [13], that is, the quadratic majorization function $w_i^{(k)} x_i^2 + c_i^{(k)}$ is not defined at $x_i^{(k)} = 0$, due to the nondifferentiability of $g_p(x_i)$ at $x_i = 0$. For example, considering the quadratic majorization function of $g_p(x_i) = |x_i|^p$, $0 < p \leq 1$ in (15), it is easy to see that the coefficient $\frac{p}{2} |x_i^{(k)}|^{p-2}$ is not defined at $x_i^{(k)} = 0$. To tackle this problem, the authors of [13] proposed to define the majorization function at the particular point $x_i^{(k)} = 0$ as

$$u(x_i, 0) = \begin{cases} +\infty, & x_i \neq 0 \\ 0, & x_i = 0, \end{cases}$$

which implies that once $x_i^{(k)} = 0$, the next iteration will also be zero, i.e., $x_i^{(k+1)} = 0$. This may impact the convergence of the algorithm to a minimizer of the objective, if the corresponding element of the minimizer is in fact not zero. Another common approach for dealing with this singularity issue is to incorporate a small $\epsilon > 0$, for example, for the quadratic majorization

TABLE I
SMOOTH APPROXIMATION $g_p^\epsilon(x_i)$ OF THE SURROGATE FUNCTIONS $g_p(x_i)$ AND THE
QUADRATIC MAJORIZATION FUNCTIONS $u(x_i, x_i^{(k)}) = w_i^{(k)} x_i^2 + c_i^{(k)}$ AT $x_i^{(k)}$

Surrogate function $g_p(x_i)$	Smooth approximation $g_p^\epsilon(x_i)$	$w_i^{(k)}$
$ x_i ^p, 0 < p \leq 1$	$\begin{cases} \frac{p}{2} \epsilon^{p-2} x_i^2, & x_i \leq \epsilon \\ x_i ^p - (1 - \frac{p}{2}) \epsilon^p, & x_i > \epsilon \end{cases}$	$\begin{cases} \frac{p}{2} \epsilon^{p-2}, & x_i^{(k)} \leq \epsilon \\ \frac{p}{2} x_i^{(k)} ^{p-2}, & x_i^{(k)} > \epsilon \end{cases}$
$\log(1 + x_i /p) / \log(1 + 1/p), p > 0$	$\begin{cases} \frac{x_i^2}{2\epsilon(p+\epsilon) \log(1+1/p)}, & x_i \leq \epsilon \\ \frac{\log(1+ x_i /p) - \log(1+\epsilon/p) + \frac{\epsilon}{2(p+\epsilon)}}{\log(1+1/p)}, & x_i > \epsilon \end{cases}$	$\begin{cases} \frac{1}{2\epsilon(p+\epsilon) \log(1+1/p)}, & x_i^{(k)} \leq \epsilon \\ \frac{1}{2 \log(1+1/p) x_i^{(k)} (x_i^{(k)} + p)}, & x_i^{(k)} > \epsilon \end{cases}$
$1 - e^{- x_i /p}, p > 0$	$\begin{cases} \frac{e^{-\epsilon/p}}{2p\epsilon} x_i^2, & x_i \leq \epsilon \\ -e^{- x_i /p} + (1 + \frac{\epsilon}{2p}) e^{-\epsilon/p}, & x_i > \epsilon \end{cases}$	$\begin{cases} \frac{e^{-\epsilon/p}}{2p\epsilon}, & x_i^{(k)} \leq \epsilon \\ \frac{e^{- x_i^{(k)} /p}}{2p x_i^{(k)} }, & x_i^{(k)} > \epsilon \end{cases}$




Fig. 3. The function $g_p(x) = |x|^p$ with $p = 0.5$ and its smooth approximation $g_p^\epsilon(x)$ with $\epsilon = 0.05$.

function of $g_p(x_i) = |x_i|^p, 0 < p \leq 1$ in (15), replace the coefficient $\frac{p}{2} |x_i^{(k)}|^{p-2}$ by

$$w_i^{(k)} = \frac{p}{2} \left(\left(x_i^{(k)} \right)^2 + \epsilon \right)^{\frac{p-2}{2}},$$

which is the so called damping approach used in [16]. A potential problem of this approach is that although ϵ is small, we have no idea how it will affect the convergence of the algorithm, since the corresponding quadratic function is no longer a majorization function of the surrogate function.

C. Smooth Approximations of Non-Differentiable Surrogate Functions

To tackle the singularity issue that arose during the construction of the quadratic minorization function in (12), in this subsection we propose to incorporate a small $\epsilon > 0$ in a more systematic way.

The idea is to approximate the non-differentiable surrogate function $g_p(x)$ with a differentiable function of the following form

$$g_p^\epsilon(x) = \begin{cases} ax^2, & |x| \leq \epsilon \\ g_p(x) - b, & |x| > \epsilon, \end{cases}$$

which aims at smoothening the non-differentiable surrogate function around zero by a quadratic function. To make the function $g_p^\epsilon(x)$ continuous and differentiable at $x = \pm\epsilon$, the following two conditions are needed $a\epsilon^2 = g_p(\epsilon) - b$, $2a\epsilon = g_p'(\epsilon)$, which lead to $a = \frac{g_p'(\epsilon)}{2\epsilon}$, $b = g_p(\epsilon) - \frac{g_p'(\epsilon)}{2}\epsilon$. Thus, the smooth approximation of the surrogate function that will be employed is

$$g_p^\epsilon(x) = \begin{cases} \frac{g_p'(\epsilon)}{2\epsilon} x^2, & |x| \leq \epsilon \\ g_p(x) - g_p(\epsilon) + \frac{g_p'(\epsilon)}{2}\epsilon, & |x| > \epsilon, \end{cases} \quad (16)$$

where $\epsilon > 0$ is a constant parameter.

Example 2: The smooth approximation of the function $g_p(x) = |x|^p, 0 < p \leq 1$ is

$$g_p^\epsilon(x) = \begin{cases} \frac{p}{2} \epsilon^{p-2} x^2, & |x| \leq \epsilon \\ |x|^p - (1 - \frac{p}{2}) \epsilon^p, & |x| > \epsilon. \end{cases} \quad (17)$$

The case with $p = 0.5$ and $\epsilon = 0.05$ is illustrated in Fig. 3. When $p = 1$, the smooth approximation (17) is the well known Huber penalty function and the application of Huber penalty as smoothed absolute value has been used in [24] to derive fast algorithms for sparse recovery.

With this smooth approximation, the problem (5) becomes the following smoothed one:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{maximize}} \quad \mathbf{x}^T \mathbf{A} \mathbf{x} - \rho \sum_{i=1}^n g_p^\epsilon(x_i) \\ & \text{subject to} \quad \mathbf{x}^T \mathbf{B} \mathbf{x} = 1, \end{aligned} \quad (18)$$

where $g_p^\epsilon(\cdot)$ is the function given by (16). We now majorize the smoothed surrogate functions $g_p^\epsilon(x_i)$ with quadratic functions and the coefficients of the quadratic majorization functions are summarized in Table I (we have omitted the coefficient $c_i^{(k)}$ since it is irrelevant for the MM algorithm). Notice that the quadratic functions $w_i^{(k)} x_i^2 + c_i^{(k)}$ are now well defined. Thus, the smooth approximation we have constructed can be viewed as a systematic way to incorporate a small $\epsilon > 0$ to deal with the singularity issue of IRLS type algorithms.

Although there is no singularity issue when applying the quadratic minorization to the smoothed problem (18), a natural question is if we solve the smoothed problem, what can we say about the solution with regard to the original problem (5). In the following, we present some results which aim at answering the question.

Lemma 3: Let $g_p(\cdot)$ be a continuous even function defined on \mathbf{R} , differentiable everywhere except at zero, concave and monotone increasing on $[0, +\infty)$ with $g_p(0) = 0$. Then the smooth approximation $g_p^\epsilon(x)$ defined by (16) is a global lower bound of $g_p(x)$ and $g_p^\epsilon(x) + g_p(\epsilon) - \frac{g_p'(\epsilon)}{2}\epsilon$ is a global upper bound of $g_p(x)$.

Proof: The lemma is quite intuitive and the proof is omitted. ■

Proposition 4: Let $f(\mathbf{x})$ be the objective of the problem (5), with $g_p(\cdot)$ as in Lemma 3 and $f_\epsilon(\mathbf{x})$ be the objective of the smoothed problem (18) with $g_p^\epsilon(\cdot)$ as in (16). Let \mathbf{x}^* be the optimal solution of the problem (5) and \mathbf{x}_ϵ^* be the optimal solution of the smoothed problem (18). Then $0 \leq f(\mathbf{x}^*) - f(\mathbf{x}_\epsilon^*) \leq \rho n(g_p(\epsilon) - \frac{g_p'(\epsilon)}{2}\epsilon)$ and $\lim_{\epsilon \downarrow 0} \rho n(g_p(\epsilon) - \frac{g_p'(\epsilon)}{2}\epsilon) = 0$.

Proof: See Appendix A. ■

Proposition 4 gives a suboptimality bound on the solution of the smoothed problem (18) in the sense that we can solve the original problem (5) to a very high accuracy by solving the smoothed problem (18) with a small enough ϵ (say $\epsilon \leq 10^{-6}$, and $\epsilon = 10^{-8}$ is used in our simulations). Of course, in general it is hard to solve either problem (5) or (18) to global maximum or even local maximum, since both of them are nonconvex. But from this point, there may be advantages in solving the smoothed problem with the smoothing parameter ϵ decreasing gradually, since choosing a relatively large ϵ at the beginning can probably smoothen out some undesirable local maxima, so that the algorithm can escape from these local points. This idea has been used with some success in [16], [24]. A decreasing scheme of ϵ will be considered later in the numerical simulations.

D. Iteratively Reweighted Quadratic Minorization

With the quadratic minorization function constructed and the smoothing technique used to deal with the singularity issue, we are now ready to state the overall algorithm for the approximate SGEP in (5).

First, we approximate the non-differentiable surrogate functions $g_p(x_i)$ by smooth functions $g_p^\epsilon(x_i)$, which leads to the smoothed problem (18). Then at iteration k , we construct the quadratic minorization function $\mathbf{x}^T(\mathbf{A} - \rho \text{Diag}(\mathbf{w}^{(k)}))\mathbf{x} - \rho \sum_{i=1}^n c_i^{(k)}$ of the objective via majorizing each smoothed surrogate function $g_p^\epsilon(x_i)$ at $x_i^{(k)}$ by a quadratic function $w_i^{(k)}x_i^2 + c_i^{(k)}$ and solve the following minorized problem (with the constant term ignored)

$$\begin{aligned} & \underset{\mathbf{x}}{\text{maximize}} \quad \mathbf{x}^T(\mathbf{A} - \rho \text{Diag}(\mathbf{w}^{(k)}))\mathbf{x} \\ & \text{subject to} \quad \mathbf{x}^T \mathbf{B} \mathbf{x} = 1, \end{aligned} \quad (19)$$

which is to find the leading generalized eigenvector of the matrix pair $(\mathbf{A} - \rho \text{Diag}(\mathbf{w}^{(k)}), \mathbf{B})$, where $\mathbf{w}^{(k)} = [w_1^{(k)}, \dots, w_n^{(k)}]^T$ and $w_i^{(k)}, i = 1, \dots, n$ are given in Table I. The method is summarized in Algorithm 1 and we will refer to it as IRQM (Iterative Reweighted Quadratic Minorization), since it is based on iteratively minorizing the penalty function with reweighted quadratic function.

Algorithm 1: IRQM — Iteratively Reweighted Quadratic Minorization algorithm for the sparse generalized eigenvalue problem (5).

Require: $\mathbf{A} \in \mathbf{S}^n, \mathbf{B} \in \mathbf{S}_{++}^n, \rho > 0, \epsilon > 0$

- 1: Set $k = 0$, choose $\mathbf{x}^{(0)} \in \{\mathbf{x} : \mathbf{x}^T \mathbf{B} \mathbf{x} = 1\}$
 - 2: **repeat**
 - 3: Compute $\mathbf{w}^{(k)}$ according to Table I.
 - 4: $\mathbf{x}^{(k+1)} \leftarrow$ leading generalized eigenvector of the matrix pair $(\mathbf{A} - \rho \text{Diag}(\mathbf{w}^{(k)}), \mathbf{B})$
 - 5: $k \leftarrow k + 1$
 - 6: **until** convergence
 - 7: **return** $\mathbf{x}^{(k)}$
-

At every iteration of the proposed IRQM algorithm, we need to find the generalized eigenvector of the matrix pair $(\mathbf{A} - \rho \text{Diag}(\mathbf{w}^{(k)}), \mathbf{B})$ corresponding to the largest generalized eigenvalue. Since $\mathbf{B} \in \mathbf{S}_{++}^n$, a standard approach for this problem is to transform it to a standard eigenvalue problem via the Cholesky decomposition of \mathbf{B} . Then standard algorithms, such as power iterations (applied to a shifted matrix) and Lanczos method can be used. The drawback of this approach is that a matrix factorization is needed, making it less attractive when this factorization is expensive. Besides, as some $w_i^{(k)}$ become very large, the problem is highly ill-conditioned and standard iterative algorithms may suffer from extremely slow convergence.

To overcome these difficulties, we provide a preconditioned steepest ascent method, which is matrix factorization free and employs preconditioning to deal with the ill-conditioning problem. Let us derive the steepest ascent method without preconditioning first. The key step is to reformulate the leading generalized eigenvalue problem as maximizing the Rayleigh quotient

$$R(\mathbf{x}) = \frac{\mathbf{x}^T \tilde{\mathbf{A}} \mathbf{x}}{\mathbf{x}^T \mathbf{B} \mathbf{x}} \quad (20)$$

over the domain $\mathbf{x} \neq \mathbf{0}$, where $\tilde{\mathbf{A}} = \mathbf{A} - \rho \text{Diag}(\mathbf{w}^{(k)})$. Let $\mathbf{x}^{(l)}$ be the current iterate, the gradient of $R(\mathbf{x})$ at $\mathbf{x}^{(l)}$ is

$$\frac{2}{\mathbf{x}^{(l)T} \mathbf{B} \mathbf{x}^{(l)}} \left(\tilde{\mathbf{A}} \mathbf{x}^{(l)} - R(\mathbf{x}^{(l)}) \mathbf{B} \mathbf{x}^{(l)} \right),$$

which is an ascent direction of $R(\mathbf{x})$ at $\mathbf{x}^{(l)}$. Let $\mathbf{r}^{(l)} = \tilde{\mathbf{A}} \mathbf{x}^{(l)} - R(\mathbf{x}^{(l)}) \mathbf{B} \mathbf{x}^{(l)}$, the steepest ascent method searches along the line $\mathbf{x}^{(l)} + \tau \mathbf{r}^{(l)}$ for a τ that maximizes the Rayleigh quotient $R(\mathbf{x}^{(l)} + \tau \mathbf{r}^{(l)})$. Since the Rayleigh quotient $R(\mathbf{x}^{(l)} + \tau \mathbf{r}^{(l)})$ is a scalar function of τ , the maximum will be achieved either at points with zero derivative or as τ goes to infinity. Setting the derivative of $R(\mathbf{x}^{(l)} + \tau \mathbf{r}^{(l)})$ with respect to τ equal to 0, we can get the following quadratic equation

$$a\tau^2 + b\tau + c = 0, \quad (21)$$

where

$$\begin{aligned} a &= \mathbf{r}^{(l)T} \tilde{\mathbf{A}} \mathbf{r}^{(l)} \mathbf{x}^{(l)T} \mathbf{B} \mathbf{r}^{(l)} - \mathbf{r}^{(l)T} \mathbf{B} \mathbf{r}^{(l)} \mathbf{x}^{(l)T} \tilde{\mathbf{A}} \mathbf{r}^{(l)} \\ b &= \mathbf{r}^{(l)T} \tilde{\mathbf{A}} \mathbf{r}^{(l)} \mathbf{x}^{(l)T} \mathbf{B} \mathbf{x}^{(l)} - \mathbf{r}^{(l)T} \mathbf{B} \mathbf{r}^{(l)} \mathbf{x}^{(l)T} \tilde{\mathbf{A}} \mathbf{x}^{(l)} \\ c &= \mathbf{r}^{(l)T} \tilde{\mathbf{A}} \mathbf{x}^{(l)} \mathbf{x}^{(l)T} \mathbf{B} \mathbf{x}^{(l)} - \mathbf{r}^{(l)T} \mathbf{B} \mathbf{x}^{(l)} \mathbf{x}^{(l)T} \tilde{\mathbf{A}} \mathbf{x}^{(l)}. \end{aligned}$$

Algorithm 2: Preconditioned steepest ascent method for problem (19).

Require: $\mathbf{A} \in \mathbf{S}^n$, $\mathbf{B} \in \mathbf{S}_{++}^n$, $\mathbf{w}^{(k)}$, $\rho > 0$

- 1: Set $l = 0$, choose $\mathbf{x}^{(0)} \in \{\mathbf{x} : \mathbf{x}^T \mathbf{B} \mathbf{x} = 1\}$
- 2: Let $\tilde{\mathbf{A}} = \mathbf{A} - \rho \text{Diag}(\mathbf{w}^{(k)})$
- 3: **repeat**
- 4: $R(\mathbf{x}^{(l)}) = \mathbf{x}^{(l)T} \tilde{\mathbf{A}} \mathbf{x}^{(l)} / \mathbf{x}^{(l)T} \mathbf{B} \mathbf{x}^{(l)}$
- 5: $\mathbf{r}^{(l)} = \tilde{\mathbf{A}} \mathbf{x}^{(l)} - R(\mathbf{x}^{(l)}) \mathbf{B} \mathbf{x}^{(l)}$
- 6: $\mathbf{r}^{(l)} = \mathbf{P} \mathbf{r}^{(l)}$
- 7: $\mathbf{x} = \mathbf{x}^{(l)} + \tau \mathbf{r}^{(l)}$, with τ chosen to maximize $R(\mathbf{x}^{(l)} + \tau \mathbf{r}^{(l)})$
- 8: $\mathbf{x}^{(l+1)} = \mathbf{x} / \sqrt{\mathbf{x}^T \mathbf{B} \mathbf{x}}$
- 9: $l = l + 1$
- 10: **until** convergence
- 11: **return** $\mathbf{x}^{(l)}$

Let us denote $B_{xx} = \mathbf{x}^{(l)T} \mathbf{B} \mathbf{x}^{(l)}$, $B_{rr} = \mathbf{r}^{(l)T} \mathbf{B} \mathbf{r}^{(l)}$ and $B_{xr} = \mathbf{x}^{(l)T} \mathbf{B} \mathbf{r}^{(l)}$, by direct computation we have

$$b^2 - 4ac = \left(B_{xx} B_{rr} \left(R(\mathbf{r}^{(l)}) - R(\mathbf{x}^{(l)}) \right) - 2B_{xr} \left\| \mathbf{r}^{(l)} \right\|_2^2 \right)^2 + 4 \left\| \mathbf{r}^{(l)} \right\|_2^4 (B_{xx} B_{rr} - B_{xr}^2).$$

According to Cauchy-Schwartz inequality, it is easy to see that $B_{xx} B_{rr} - B_{xr}^2 \geq 0$, thus $b^2 - 4ac \geq 0$, which implies that the (21) has one or two real roots. By comparing the Rayleigh quotient $R(\mathbf{x}^{(l)} + \tau \mathbf{r}^{(l)})$ at the roots of (21) with $R(\mathbf{r}^{(l)})$ (the Rayleigh quotient corresponding to $\tau \rightarrow \infty$), we can determine the steepest ascent. It is worth noting that the coefficients of the (21) can be computed by matrix-vector multiplications and inner products only, thus very efficient.

Though the per-iteration computational complexity of this steepest ascent method is very low, it may converge very slow, especially when some $w_i^{(k)}$ become very large. To accelerate the convergence, we introduce a preconditioner here. Preconditioning is an important technique in iterative methods for solving large system of linear equations, for example the widely used preconditioned conjugate gradient method. It can also be used for eigenvalue problems. In the steepest ascent method, to introduce a positive definite preconditioner \mathbf{P} , we simply multiply the residual $\mathbf{r}^{(l)}$ by \mathbf{P} . The steepest ascent method with preconditioning for the leading generalized eigenvalue problem (19) is summarized in Algorithm 2. To use the algorithm in practice, the preconditioner \mathbf{P} remains to be chosen. For the particular problem of interest, we choose a diagonal \mathbf{P} as follows

$$\mathbf{P} = \begin{cases} \text{Diag}(\rho \mathbf{w}^{(k)} + |\text{diag}(\mathbf{A})|)^{-1}, & \frac{\rho \|\mathbf{w}^{(k)}\|_2}{\|\text{diag}(\mathbf{A})\|_2} > 10^2 \\ \mathbf{I}_n. & \text{otherwise.} \end{cases}$$

In other words, we apply a preconditioner only when some elements of $\mathbf{w}^{(k)}$ become relatively large. Since the preconditioner \mathbf{P} we choose here is positive definite, the direction $\mathbf{P} \mathbf{r}^{(l)}$ is still an ascent direction and the algorithm is still monotonically

increasing. For more details regarding preconditioned eigen-solvers, the readers can refer to the book [25].

In practice, the preconditioned steepest ascent method usually converges to the leading generalized eigenvector, but it is not guaranteed in principle, since the Rayleigh quotient is not concave. But note that the descent property (9) of the majorization-minimization scheme depends only on decreasing $u(\mathbf{x}, \mathbf{x}^{(k)})$ and not on minimizing it. Similarly, for the minorization-maximization scheme used by Algorithm 1, to preserve the ascent property, we only need to increase the objective of (19) at each iteration, rather than maximizing it. Since the steepest ascent method increases the objective at every iteration, thus when it is applied (initialized with the solution of previous iteration) to compute the leading generalized eigenvector at each iteration of Algorithm 1, the ascent property of Algorithm 1 can be guaranteed.

E. Sparse GEP With Special Structure

Until now we have considered the sparse GEP in the general case with $\mathbf{A} \in \mathbf{S}^n$, $\mathbf{B} \in \mathbf{S}_{++}^n$ and derived an iterative algorithm IRQM. If we assume more properties or some special structure for \mathbf{A} and \mathbf{B} , then we may derive simpler and more efficient algorithms. In the following, we will consider the case where $\mathbf{A} \in \mathbf{S}_+^n$ and $\mathbf{B} = \text{Diag}(\mathbf{b})$, $\mathbf{b} \in \mathbf{R}_{++}^n$. Notice that although this is a special case of the general sparse GEP, it still includes the sparse PCA problem as a special case where $\mathbf{B} = \mathbf{I}_n$.

We first present two results that will be used when deriving fast algorithms for this special case.

Proposition 5: Given $\mathbf{a} \in \mathbf{R}^n$, $\mathbf{w}, \mathbf{b} \in \mathbf{R}_{++}^n$, $\rho > 0$, let $\mathcal{I}_{\min} = \arg \min \{\rho w_i / b_i : i \in \{1, \dots, n\}\}$, $\mu_{\min} = -\min \{\rho w_i / b_i : i \in \{1, \dots, n\}\}$ and $s = \sum_{i \notin \mathcal{I}_{\min}} \frac{b_i a_i^2}{(\mu_{\min} b_i + \rho w_i)^2}$. Then the problem

$$\begin{aligned} & \underset{\mathbf{x}}{\text{maximize}} && 2\mathbf{a}^T \mathbf{x} - \rho \mathbf{x}^T \text{Diag}(\mathbf{w}) \mathbf{x} \\ & \text{subject to} && \mathbf{x}^T \text{Diag}(\mathbf{b}) \mathbf{x} = 1 \end{aligned} \quad (22)$$

admits the following solution:

- If $\exists i \in \mathcal{I}_{\min}$, such that $a_i^2 > 0$ or $s > 1$, then

$$x_i^* = \frac{a_i}{\mu b_i + \rho w_i}, \quad i = 1, \dots, n,$$

where $\mu > \mu_{\min}$ is given by the solution of the scalar equation

$$\sum_{i=1}^n \frac{b_i a_i^2}{(\mu b_i + \rho w_i)^2} = 1.$$

- Otherwise,

$$x_i^* = \begin{cases} a_i / (\mu_{\min} b_i + \rho w_i), & i \notin \mathcal{I}_{\min} \\ \sqrt{(1-s)/b_i}, & i = \max\{i : i \in \mathcal{I}_{\min}\} \\ 0, & \text{otherwise.} \end{cases}$$

Proof: See Appendix B. ■

Proposition 6: Given $\mathbf{a} \in \mathbf{R}^n$ with $|a_1| \geq \dots \geq |a_n|$ and $\rho > 0$, then the problem

$$\begin{aligned} & \underset{\mathbf{x}}{\text{maximize}} && \mathbf{a}^T \mathbf{x} - \rho \|\mathbf{x}\|_0 \\ & \text{subject to} && \|\mathbf{x}\|_2 = 1 \end{aligned} \quad (23)$$

admits the following solution:

- If $|a_1| \leq \rho$, then

$$x_i^* = \begin{cases} \text{sgn}(a_1), & i = 1 \\ 0, & \text{otherwise.} \end{cases} \quad (24)$$

- Otherwise,

$$x_i^* = \begin{cases} a_i / \sqrt{\sum_{j=1}^s a_j^2}, & i \leq s \\ 0, & \text{otherwise,} \end{cases}$$

where s is the largest integer p that satisfies the following inequality

$$\sqrt{\sum_{i=1}^p a_i^2} > \sqrt{\sum_{i=1}^{p-1} a_i^2} + \rho. \quad (25)$$

Proof: See Appendix C. ■

Let us return to the problem. In this special case, the smoothed problem (18) reduces to

$$\begin{aligned} & \underset{\mathbf{x}}{\text{maximize}} \quad \mathbf{x}^T \mathbf{A} \mathbf{x} - \rho \sum_{i=1}^n g_p^\epsilon(x_i) \\ & \text{subject to} \quad \mathbf{x}^T \text{Diag}(\mathbf{b}) \mathbf{x} = 1. \end{aligned} \quad (26)$$

The previously derived IRQM algorithm can be used here, but in that iterative algorithm, we need to find the leading generalized eigenvector at each iteration, for which another iterative algorithm is needed. By exploiting the special structure of this case, in the following we derive a simpler algorithm that at each iteration has a closed-form solution.

Notice that, in this case, $\mathbf{A} \in \mathbf{S}_+^n$, the first term $\mathbf{x}^T \mathbf{A} \mathbf{x}$ in the objective is convex and can be minorized by its tangent plane $2\mathbf{x}^T \mathbf{A} \mathbf{x}^{(k)}$ at $\mathbf{x}^{(k)}$. So instead of only minorizing the second term, we can minorize both terms. This suggests solving the following minorized problem at iteration k :

$$\begin{aligned} & \underset{\mathbf{x}}{\text{maximize}} \quad 2\mathbf{x}^T \mathbf{A} \mathbf{x}^{(k)} - \rho \mathbf{x}^T \text{Diag}(\mathbf{w}^{(k)}) \mathbf{x} \\ & \text{subject to} \quad \mathbf{x}^T \text{Diag}(\mathbf{b}) \mathbf{x} = 1, \end{aligned} \quad (27)$$

where $\mathbf{x}^{(k)}$ is the solution at iteration k and $\mathbf{w}^{(k)}$ is computed according to Table I. The problem is a nonconvex QCQP, but by letting $\mathbf{a} = \mathbf{A} \mathbf{x}^{(k)}$ and $\mathbf{w} = \mathbf{w}^{(k)}$, we know from Proposition 5 that it can be solved in closed-form. The iterative algorithm for solving problem (26) is summarized in Algorithm 3.

Algorithm 3: The MM algorithm for problem (26).

Require: $\mathbf{A} \in \mathbf{S}_+^n$, $\mathbf{b} \in \mathbf{R}_{++}^n$, $\rho > 0$, $\epsilon > 0$

- 1: Set $k = 0$, choose $\mathbf{x}^{(0)} \in \{\mathbf{x} : \mathbf{x}^T \text{Diag}(\mathbf{b}) \mathbf{x} = 1\}$
- 2: **repeat**
- 3: $\mathbf{a} = \mathbf{A} \mathbf{x}^{(k)}$
- 4: Compute $\mathbf{w}^{(k)}$ according to Table I.
- 5: Solve the following problem according to Proposition 5 and set the solution as $\mathbf{x}^{(k+1)}$:

$$\max_{\mathbf{x}} \{2\mathbf{a}^T \mathbf{x} - \rho \mathbf{x}^T \text{Diag}(\mathbf{w}^{(k)}) \mathbf{x} : \mathbf{x}^T \text{Diag}(\mathbf{b}) \mathbf{x} = 1\}$$

- 6: $k = k + 1$
 - 7: **until** convergence
 - 8: **return** $\mathbf{x}^{(k)}$
-

In fact, in this special case, we can apply the MM scheme to solve the original problem (4) directly, without approximating $\|\mathbf{x}\|_0$, i.e., solving

$$\begin{aligned} & \underset{\mathbf{x}}{\text{maximize}} \quad \mathbf{x}^T \mathbf{A} \mathbf{x} - \rho \|\mathbf{x}\|_0 \\ & \text{subject to} \quad \mathbf{x}^T \text{Diag}(\mathbf{b}) \mathbf{x} = 1. \end{aligned} \quad (28)$$

First, we define a new variable $\tilde{\mathbf{x}} = \text{Diag}(\mathbf{b})^{\frac{1}{2}} \mathbf{x}$ and using the fact $\|\text{Diag}(\mathbf{b})^{-\frac{1}{2}} \tilde{\mathbf{x}}\|_0 = \|\tilde{\mathbf{x}}\|_0$, the problem can be rewritten as

$$\begin{aligned} & \underset{\tilde{\mathbf{x}}}{\text{maximize}} \quad \tilde{\mathbf{x}}^T \tilde{\mathbf{A}} \tilde{\mathbf{x}} - \rho \|\tilde{\mathbf{x}}\|_0 \\ & \text{subject to} \quad \tilde{\mathbf{x}}^T \tilde{\mathbf{x}} = 1, \end{aligned} \quad (29)$$

where $\tilde{\mathbf{A}} = \text{Diag}(\mathbf{b})^{-\frac{1}{2}} \mathbf{A} \text{Diag}(\mathbf{b})^{-\frac{1}{2}}$.

Now the idea is to minorize only the quadratic term by its tangent plane, while keeping the ℓ_0 -norm. Given $\tilde{\mathbf{x}}^{(k)}$ at iteration k , linearizing the quadratic term yields

$$\begin{aligned} & \underset{\tilde{\mathbf{x}}}{\text{maximize}} \quad 2\tilde{\mathbf{x}}^T \tilde{\mathbf{A}} \tilde{\mathbf{x}}^{(k)} - \rho \|\tilde{\mathbf{x}}\|_0 \\ & \text{subject to} \quad \|\tilde{\mathbf{x}}\|_2 = 1, \end{aligned} \quad (30)$$

which has a closed-form solution. To see this, we first define $\mathbf{a} = 2\tilde{\mathbf{A}} \tilde{\mathbf{x}}^{(k)}$ and sort the entries of vector \mathbf{a} according to the absolute value (only needed for entries with $|a_i| > \rho$) in descending order, then Proposition 6 can be readily applied to obtain the solution. Finally we need to reorder the solution back to the original ordering. This algorithm for solving problem (28) is summarized in Algorithm 4.

It is worth noting that although the derivations of Algorithms 3 and 4 require \mathbf{A} to be symmetric positive semidefinite, the algorithms can also be used to deal with the more general case $\mathbf{A} \in \mathbf{S}^n$. When the matrix \mathbf{A} in problem (26) or (28) is not positive semidefinite, we can replace \mathbf{A} with $\mathbf{A}_\alpha = \mathbf{A} + \alpha \text{Diag}(\mathbf{b})$, with $\alpha \geq -\lambda_{\min}(\mathbf{A})/b_{\min}$ such that $\mathbf{A}_\alpha \in \mathbf{S}_+^n$, where $\lambda_{\min}(\mathbf{A})$ is the smallest eigenvalue of matrix \mathbf{A} and b_{\min} is the smallest entry of \mathbf{b} . Since the additional term $\alpha \mathbf{x}^T \text{Diag}(\mathbf{b}) \mathbf{x}$ in the objective is just a constant α over the constraint set, it is easy to see that after replacing \mathbf{A} with \mathbf{A}_α the resulting problem is equivalent to the original one. Then the Algorithm 3 or 4 can be readily applied.

Algorithm 4: The MM algorithm for problem (28).

Require: $\mathbf{A} \in \mathbf{S}_+^n$, $\mathbf{b} \in \mathbf{R}_{++}^n$, $\rho > 0$

- 1: Set $k = 0$, choose $\tilde{\mathbf{x}}^{(0)} \in \{\tilde{\mathbf{x}} : \tilde{\mathbf{x}}^T \tilde{\mathbf{x}} = 1\}$
- 2: Let $\tilde{\mathbf{A}} = \text{Diag}(\mathbf{b})^{-\frac{1}{2}} \mathbf{A} \text{Diag}(\mathbf{b})^{-\frac{1}{2}}$
- 3: **repeat**
- 4: $\mathbf{a} = 2\tilde{\mathbf{A}} \tilde{\mathbf{x}}^{(k)}$
- 5: Sort \mathbf{a} with the absolute value in descending order.
- 6: Compute $\tilde{\mathbf{x}}^{(k+1)}$ according to Proposition 6:

$$\tilde{\mathbf{x}}^{(k+1)} \in \arg \max_{\tilde{\mathbf{x}}} \{\mathbf{a}^T \tilde{\mathbf{x}} - \rho \|\tilde{\mathbf{x}}\|_0 : \|\tilde{\mathbf{x}}\|_2 = 1\}$$

- 7: Reorder $\tilde{\mathbf{x}}^{(k+1)}$
 - 8: $k = k + 1$
 - 9: **until** convergence
 - 10: **return** $\mathbf{x} = \text{Diag}(\mathbf{b})^{-\frac{1}{2}} \tilde{\mathbf{x}}^{(k)}$
-

IV. CONVERGENCE ANALYSIS

The algorithms proposed in this paper are all based on the minorization-maximization scheme, thus according to Subsection III.A, we know that the sequence of objective values evaluated at $\{\mathbf{x}^{(k)}\}$ generated by the algorithms is non-decreasing. Since the constraint sets in our problems are compact, the sequence of objective values is bounded. Thus, the sequence of objective values is guaranteed to converge to a finite value. The monotonicity makes MM algorithms very stable. In this section, we will analyze the convergence property of the sequence $\{\mathbf{x}^{(k)}\}$ generated by the algorithms.

Let us consider the IRQM algorithm in Algorithm 1, in which the minorization-maximization scheme is applied to the smoothed problem (18). In the problem, the objective is neither convex nor concave and the constraint set is also nonconvex. But as we shall see later, after introducing a technical assumption on the surrogate function $g_p(x)$, the problem is equivalent to a problem which maximizes a convex function over a convex set and we will prove that the sequence generated by the IRQM algorithm converges to the stationary point of the equivalent problem. The convergence of the Algorithm 3 can be proved similarly, since the minorization function applied can also be convexified. First, let us give the assumption and present some results that will be useful later.

Assumption 1: The surrogate function $g_p(x)$ is twice differentiable on $(0, +\infty)$ and its gradient $g'_p(x)$ is convex on $(0, +\infty)$.

It is easy to verify that the three surrogate functions listed in Table I all satisfy this assumption. With this assumption, the first result shows that the smooth approximation $g'_p(x)$ we have constructed is Lipschitz continuously differentiable.

Lemma 7: Let $g_p(\cdot)$ be a continuous even function defined on \mathbf{R} , differentiable everywhere except at zero, concave and monotone increasing on $[0, +\infty)$ with $g_p(0) = 0$. Let Assumption 1 be satisfied. Then the smooth approximation $g'_p(x)$ defined by (16) is Lipschitz continuously differentiable with Lipschitz constant $L = \max\{\frac{g'_p(\epsilon)}{\epsilon}, |g''_p(\epsilon)|\}$.

Proof: See Appendix D. ■

Next, we recall a useful property of Lipschitz continuously differentiable functions [26].

Proposition 8: If $f : \mathbf{R}^n \rightarrow \mathbf{R}$ is Lipschitz continuously differentiable on a convex set C with some Lipschitz constant L , then $\varphi(\mathbf{x}) = f(\mathbf{x}) + \frac{\alpha}{2}\mathbf{x}^T\mathbf{x}$ is a convex function on C for every $\alpha \geq L$.

The next result then follows, showing that the smoothed problem (18) is equivalent to a problem in the form of maximizing a convex function over a compact set.

Lemma 9: There exists $\alpha > 0$ such that $\mathbf{x}^T(\mathbf{A} + \alpha\mathbf{B})\mathbf{x} - \rho \sum_{i=1}^n g_p^\epsilon(x_i)$ is convex and the problem

$$\begin{aligned} & \underset{\mathbf{x}}{\text{maximize}} \quad \mathbf{x}^T(\mathbf{A} + \alpha\mathbf{B})\mathbf{x} - \rho \sum_{i=1}^n g_p^\epsilon(x_i) \\ & \text{subject to} \quad \mathbf{x}^T\mathbf{B}\mathbf{x} = 1 \end{aligned} \quad (31)$$

is equivalent to the problem (18) in the sense that they admit the same set of optimal solutions.

Proof: From Lemma 7, it is easy to see that $\mathbf{x}^T\mathbf{A}\mathbf{x} - \rho \sum_{i=1}^n g_p^\epsilon(x_i)$ is Lipschitz continuously differentiable. As-

sume the Lipschitz constant of its gradient is L , then according to Proposition 8, $\mathbf{x}^T(\mathbf{A} + \frac{L}{2}\mathbf{I})\mathbf{x} - \rho \sum_{i=1}^n g_p^\epsilon(x_i)$ is convex. Since \mathbf{B} is positive definite, $\lambda_{\min}(\mathbf{B}) > 0$. By choosing $\alpha \geq \frac{L}{2\lambda_{\min}(\mathbf{B})}$, we have that $\mathbf{x}^T(\alpha\mathbf{B} - \frac{L}{2}\mathbf{I})\mathbf{x}$ is convex. The sum of the two convex functions, i.e., $\mathbf{x}^T(\mathbf{A} + \alpha\mathbf{B})\mathbf{x} - \rho \sum_{i=1}^n g_p^\epsilon(x_i)$, is convex.

Since the additional term $\alpha\mathbf{x}^T\mathbf{B}\mathbf{x}$ is just a constant α over the constraint set $\mathbf{x}^T\mathbf{B}\mathbf{x} = 1$, it is obvious that any solution of problem (31) is also a solution of problem (18) and vice versa. ■

Generally speaking, maximizing a convex function over a compact set remains a hard nonconvex problem. There is some consolation, however, according to the following result in convex analysis [27].

Proposition 10: Let $f : \mathbf{R}^n \rightarrow \mathbf{R}$ be a convex function. Let $S \in \mathbf{R}^n$ be an arbitrary set and $\text{conv}(S)$ be its convex hull. Then

$$\sup\{f(\mathbf{x})|\mathbf{x} \in \text{conv}(S)\} = \sup\{f(\mathbf{x})|\mathbf{x} \in S\},$$

where the first supremum is attained only when the second (more restrictive) supremum is attained.

According to Proposition 10, we can further relax the constraint $\mathbf{x}^T\mathbf{B}\mathbf{x} = 1$ in problem (31) to $\mathbf{x}^T\mathbf{B}\mathbf{x} \leq 1$, namely, the problem

$$\begin{aligned} & \underset{\mathbf{x}}{\text{maximize}} \quad \mathbf{x}^T(\mathbf{A} + \alpha\mathbf{B})\mathbf{x} - \rho \sum_{i=1}^n g_p^\epsilon(x_i) \\ & \text{subject to} \quad \mathbf{x}^T\mathbf{B}\mathbf{x} \leq 1 \end{aligned} \quad (32)$$

is still equivalent to problem (18) in the sense that they admit the same set of optimal solutions.

Let us denote the objective function of problem (32) by $f_\epsilon^\alpha(\mathbf{x})$ and define $\mathcal{B} = \{\mathbf{x} \in \mathbf{R}^n | \mathbf{x}^T\mathbf{B}\mathbf{x} \leq 1\}$, then a point \mathbf{x}^* is referred to as a stationary point of problem (32) if

$$\nabla f_\epsilon^\alpha(\mathbf{x}^*)^T(\mathbf{x} - \mathbf{x}^*) \leq 0, \quad \forall \mathbf{x} \in \mathcal{B}. \quad (33)$$

Theorem 11: Let $\{\mathbf{x}^{(k)}\}$ be the sequence generated by the IRQM algorithm in Algorithm 1. Then every limit point of the sequence $\{\mathbf{x}^{(k)}\}$ is a stationary point of the problem (32), which is equivalent¹ to the problem (18).

Proof: Denote the objective function of the problem (18) by $f_\epsilon(\mathbf{x})$ and its quadratic minorization function at $\mathbf{x}^{(k)}$ by $q(\mathbf{x}|\mathbf{x}^{(k)})$, i.e., $q(\mathbf{x}|\mathbf{x}^{(k)}) = \mathbf{x}^T(\mathbf{A} - \rho\text{Diag}(\mathbf{w}^{(k)}))\mathbf{x}$. Denote $\mathcal{S} = \{\mathbf{x} \in \mathbf{R}^n | \mathbf{x}^T\mathbf{B}\mathbf{x} = 1\}$ and $\mathcal{B} = \{\mathbf{x} \in \mathbf{R}^n | \mathbf{x}^T\mathbf{B}\mathbf{x} \leq 1\}$. According to the general MM framework, we have

$$f_\epsilon(\mathbf{x}^{(k+1)}) \geq q(\mathbf{x}^{(k+1)}|\mathbf{x}^{(k)}) \geq q(\mathbf{x}^{(k)}|\mathbf{x}^{(k)}) = f_\epsilon(\mathbf{x}^{(k)}),$$

which means $\{f_\epsilon(\mathbf{x}^{(k)})\}$ is a non-decreasing sequence.

Assume that there exists a converging subsequence $\mathbf{x}^{(k_j)} \rightarrow \mathbf{x}^\infty$, then

$$\begin{aligned} q(\mathbf{x}^{(k_j+1)}|\mathbf{x}^{(k_j+1)}) &= f_\epsilon(\mathbf{x}^{(k_j+1)}) \geq f_\epsilon(\mathbf{x}^{(k_j+1)}) \\ &\geq q(\mathbf{x}^{(k_j+1)}|\mathbf{x}^{(k_j)}) \geq q(\mathbf{x}|\mathbf{x}^{(k_j)}), \quad \forall \mathbf{x} \in \mathcal{S}. \end{aligned}$$

Letting $j \rightarrow +\infty$, we obtain

$$q(\mathbf{x}^\infty|\mathbf{x}^\infty) \geq q(\mathbf{x}|\mathbf{x}^\infty), \quad \forall \mathbf{x} \in \mathcal{S}.$$

¹The equivalence of the problem (32) and (18) is in the sense that they have the same set of optimal solutions, but they may have different stationary points. The convergence to a stationary point of problem (32) does not imply the convergence to a stationary point of problem (18).

It is easy to see that we can always find $\alpha > 0$ such that $q_\alpha(\mathbf{x}|\mathbf{x}^\infty) = q(\mathbf{x}|\mathbf{x}^\infty) + \alpha \mathbf{x}^T \mathbf{B} \mathbf{x}$ is convex and \mathbf{x}^∞ is still a global maximizer of $q_\alpha(\mathbf{x}|\mathbf{x}^\infty)$ over \mathcal{S} . Due to Lemma 9, we can always choose α large enough such that $f_\epsilon^\alpha(\mathbf{x}) = f_\epsilon(\mathbf{x}) + \alpha \mathbf{x}^T \mathbf{B} \mathbf{x}$ is also convex. By Proposition 10, we have

$$q_\alpha(\mathbf{x}^\infty|\mathbf{x}^\infty) \geq q_\alpha(\mathbf{x}|\mathbf{x}^\infty), \forall \mathbf{x} \in \mathcal{B},$$

i.e., \mathbf{x}^∞ is a global maximizer of $q_\alpha(\mathbf{x}|\mathbf{x}^\infty)$ over the convex set \mathcal{B} . As a necessary condition, we get

$$\nabla q_\alpha(\mathbf{x}^\infty|\mathbf{x}^\infty)^T (\mathbf{x} - \mathbf{x}^\infty) \leq 0, \forall \mathbf{x} \in \mathcal{B}.$$

Since $\nabla f_\epsilon^\alpha(\mathbf{x}^\infty) = \nabla q_\alpha(\mathbf{x}^\infty|\mathbf{x}^\infty)$ by construction, we obtain

$$\nabla f_\epsilon^\alpha(\mathbf{x}^\infty)^T (\mathbf{x} - \mathbf{x}^\infty) \leq 0, \forall \mathbf{x} \in \mathcal{B},$$

implying that \mathbf{x}^∞ is a stationary point of the problem (32), which is equivalent to the problem (18) according to Lemma 9 and Proposition 10. ■

We note that in the above convergence analysis of Algorithm 1, the leading generalized eigenvector is assumed to be computed exactly at each iteration. Recall that the Algorithm 2 is not guaranteed to converge to the leading generalized eigenvector in principle, so if it is applied to compute the leading generalized eigenvector, the convergence of Algorithm 1 to a stationary point is no longer guaranteed.

V. NUMERICAL EXPERIMENTS

To compare the performance of the proposed algorithms with existing ones on the sparse generalized eigenvalue problem (SGEP) and some of its special cases, we present some experimental results in this section. All experiments were performed on a PC with a 3.20 GHz i5-3470 CPU and 8 GB RAM.

A. Sparse Generalized Eigenvalue Problem

In this subsection, we evaluate the proposed IRQM algorithm for the sparse generalized eigenvalue problem in terms of computational complexity and the ability to extract sparse generalized eigenvectors. The benchmark method considered here is the DC-SGEP algorithm proposed in [10], [11], which is based on D.C. (difference of convex functions) programming and minorization-maximization (to the best of our knowledge, this is the only algorithm proposed for this case). The problem that DC-SGEP solves is just (5) with the surrogate function $g_p(x) = \log(1 + |x|/p) / \log(1 + 1/p)$, but the equality constraint $\mathbf{x}^T \mathbf{B} \mathbf{x} = 1$ is relaxed to $\mathbf{x}^T \mathbf{B} \mathbf{x} \leq 1$. The DC-SGEP algorithm requires solving a convex quadratically constrained quadratic program (QCQP) at each iteration, which is solved by the solver Mosek² in our experiments. In the experiments, the stopping condition is $|f(\mathbf{x}^{(k+1)}) - f(\mathbf{x}^{(k)})| / \max(1, |f(\mathbf{x}^{(k)})|) \leq 10^{-5}$ for both algorithms. For the proposed IRQM algorithm, the smoothing parameter is set to be $\epsilon = 10^{-8}$.

1) *Computational Complexity*: In this subsection, we compare the computational complexity of the proposed IRQM Algorithm 1 with the DC-SGEP algorithm. The surrogate function $g_p(x) = \log(1 + |x|/p) / \log(1 + 1/p)$ is used for both algorithms






Fig. 4. Average running time versus problem size. Each curve is an average of 100 random trials.

in this experiment. The preconditioned steepest ascent method given in Algorithm 2 is applied to compute the leading generalized eigenvector at every iteration of the IRQM algorithm. To illustrate the effectiveness of the preconditioning scheme employed in Algorithm 2, we also consider computing the leading generalized eigenvector by invoking Algorithm 2 but without preconditioning, i.e., setting $\mathbf{P} = \mathbf{I}_n$. The data matrices $\mathbf{A} \in \mathbf{S}^n$ and $\mathbf{B} \in \mathbf{S}_{++}^n$ are generated as $\mathbf{A} = \mathbf{C} + \mathbf{C}^T$ and $\mathbf{B} = \mathbf{D}^T \mathbf{D}$, with $\mathbf{C} \in \mathbf{R}^{n \times n}$, $\mathbf{D} \in \mathbf{R}^{1.2n \times n}$ and the entries of both \mathbf{C} and \mathbf{D} independent, identically distributed and following $\mathcal{N}(0, 1)$. For both algorithms, the initial point $\mathbf{x}^{(0)}$ is chosen randomly with each entry following $\mathcal{N}(0, 1)$ and then normalized such that $(\mathbf{x}^{(0)})^T \mathbf{B} \mathbf{x}^{(0)} = 1$. The parameter p of the surrogate function is chosen to be 1 and the regularization parameter is $\rho = 0.1$.

The computational time for problems with different sizes are shown in Fig. 4. The results are averaged over 100 independent trials. From Fig. 4, we can see that the preconditioning scheme is indeed important for the efficiency of Algorithm 2 and the proposed IRQM algorithm is much faster than the DC-SGEP algorithm. It is worth noting that the solver Mosek which is used to solve the QCQPs for the DC-SGEP algorithm is well known for its efficiency, while the IRQM algorithm is entirely implemented in Matlab. The lower computational complexity of the IRQM algorithm, compared with the DC-SGEP algorithm, attributes to both the lower per iteration computational complexity and the faster convergence. To show this, the evolution of the objective function for one trial with $n = 100$ is plotted in Fig. 5 and we can see that the proposed IRQM algorithm takes much fewer iterations to converge. One may also notice that the two algorithms converge to the same objective value, but this does not hold in general since the problem is nonconvex.

2) *Random Data With Underlying Sparse Structure*: In this section, we generate random matrices $\mathbf{A} \in \mathbf{S}^n$ and $\mathbf{B} \in \mathbf{S}_{++}^n$ such that the matrix pair (\mathbf{A}, \mathbf{B}) has a few sparse generalized eigenvectors. To achieve this, we synthesize the data through the generalized eigenvalue decomposition $\mathbf{A} = \mathbf{V}^{-T} \text{Diag}(\mathbf{d}) \mathbf{V}^{-1}$ and $\mathbf{B} = \mathbf{V}^{-T} \mathbf{V}^{-1}$, where the first k columns of $\mathbf{V} \in \mathbf{R}^{n \times n}$ are pre-specified sparse vectors and the remaining columns are

²Mosek, available at <http://www.mosek.com/>

Fig. 5. Evolution of the objective function for one trial with $n = 100$.Fig. 6. Chance of exact recovery versus regularization parameter ρ . Parameter $p = 1$ is used for the surrogate functions.

generated randomly, \mathbf{d} is the vector of the generalized eigenvalues.



Here, we choose $n = 100$ and $k = 2$, where the two sparse generalized eigenvectors are specified as follows

$$\begin{cases} V_{i,1} = \frac{1}{\sqrt{5}} & \text{for } i = 1, \dots, 5, \\ V_{i,1} = 0 & \text{otherwise,} \end{cases}$$

$$\begin{cases} V_{i,2} = \frac{1}{\sqrt{5}} & \text{for } i = 6, \dots, 10, \\ V_{i,2} = 0 & \text{otherwise,} \end{cases}$$

and the generalized eigenvalues are chosen as

$$\begin{cases} d_1 = 10 \\ d_2 = 8 \\ d_i = 12, & \text{for } i = 3, 4, 5 \\ d_i \sim \mathcal{N}(0, 1), & \text{otherwise.} \end{cases}$$

Fig. 7. Chance of exact recovery versus regularization parameter ρ . Parameter $p = 0.3$ is used for the surrogate functions.Fig. 8. Chance of exact recovery versus regularization parameter ρ . Parameter $p = 0.1$ is used for the surrogate functions.

We generate 200 pairs of (\mathbf{A}, \mathbf{B}) as described above and employ the algorithms to compute the leading sparse generalized eigenvector $\mathbf{x}_1 \in \mathbf{R}^{100}$, which is hoped to be close to $\mathbf{V}_{:,1}$. The underlying sparse generalized eigenvector $\mathbf{V}_{:,1}$ is considered to be successfully recovered when $\min(\|\mathbf{x}_1 - \mathbf{V}_{:,1}\|_2, \|\mathbf{x}_1 + \mathbf{V}_{:,1}\|_2) \leq 0.01$. For the proposed IRQM algorithm, all the three surrogate functions listed in Table I are considered and we call the resulting algorithms “IRQM-log”, “IRQM-Lp” and “IRQM-exp”, respectively. For all the algorithms, the initial point $\mathbf{x}^{(0)}$ is chosen randomly. Regarding the parameter p of the surrogate function, three values, namely 1, 0.3 and 0.1, are compared. The corresponding performance along the whole path of the regularization parameter ρ is plotted in Figs. 6, 7 and 8, respectively.

From Fig. 6, we can see that for the case $p = 1$, the best chance of exact recovery achieved by the three IRQM algorithms are very close and all higher than that achieved by the DC-SGEP algorithm. From Figs. 7 and 8, we can see that as p becomes smaller, the best chance of exact recovery achieved by IRQM-exp, IRQM-log and DC-SGEP stay almost the same as in the case $p = 1$ (in fact decrease a little bit when $p = 0.1$), but the performance of IRQM-Lp degrades a lot. This may be explained by the fact that as p becomes smaller, the surrogate function $|x|^p$ tends to the function $\text{sgn}(|x|)$ much faster than the other two surrogate functions. So when $p = 0.1$ for example, it is much more pointed and makes the algorithm easily get stuck at some local point. In this sense, the log-based and exp-based surrogate functions seem to be better choices as they are not so sensitive to the choice of p .

3) *Decreasing Scheme of the Smoothing Parameter ϵ* : As have been discussed in the end of Section III.C, choosing a relatively large smoothing parameter ϵ at the beginning and decreasing it gradually may probably lead to better performance than the fixed ϵ scheme. In this section, we consider such a decreasing scheme and compare its performance with the fixed ϵ scheme in which the smoothing parameter is fixed to be $\epsilon = 10^{-8}$. The decreasing scheme that we will adopt is inspired by the continuation approach in [24]. The idea is to apply the IRQM algorithm to a succession of problems with decreasing smoothing parameters $\epsilon^{(0)} > \epsilon^{(1)} > \dots > \epsilon^{(T)}$ and solve the intermediate problems with less accuracy, where T is the number of decreasing steps. More specifically, at step $t = 0, \dots, T$, we apply the IRQM algorithm with smoothing parameter $\epsilon^{(t)}$ and stopping criterion $|f(\mathbf{x}^{(k+1)}) - f(\mathbf{x}^{(k)})| / \max(1, |f(\mathbf{x}^{(k)})|) \leq \sqrt{\epsilon^{(t)}}/10$ and then decrease the smoothing parameter for the next step by $\epsilon^{(t+1)} = \gamma \epsilon^{(t)}$ with $\gamma = (\epsilon^{(T)}/\epsilon^{(0)})^{1/T}$. At each step the IRQM algorithm is initialized with the solution of the previous step. The initial smoothing parameter is chosen as $\epsilon^{(0)} = \|\mathbf{x}^{(0)}\|_\infty/4$, where $\mathbf{x}^{(0)}$ is the random initial point and the minimum smoothing parameter is set as $\epsilon^{(T)} = 10^{-8}$, which is the parameter used in the fixed ϵ scheme. The number of decreasing steps is set to $T = 5$ in our experiment.

The remaining settings are the same as in the previous subsection and the log-based surrogate function with parameter $p = 0.3$ is used for the IRQM algorithm. The performance of the two schemes are shown in Fig. 9. From the figure, we can see that the decreasing scheme of the smoothing parameter achieves a higher chance of exact recovery.

B. Sparse Principal Component Analysis ($\mathbf{B} = \mathbf{I}_n$)

In this section, we consider the special case of the sparse generalized eigenvalue problem in which the matrix \mathbf{B} is the identity matrix, i.e., the sparse PCA problem, which has received most of the recent attention in the literature. In this case, the matrix \mathbf{A} is usually a (scaled) covariance matrix. Although there exists a vast literature on sparse PCA, most popular algorithms are essentially variations of the generalized power method (GPower) proposed in [7]. Thus we choose the GPower methods, namely GPower $_{\ell_1}$ and GPower $_{\ell_0}$, as the benchmarks in this section. The Matlab code of the GPower algorithms was downloaded from the authors' website. For the proposed




Fig. 9. Chance of exact recovery versus regularization parameter p . The log-based surrogate function with parameter $p = 0.3$ is used.

Algorithm 3, the surrogate function is chosen to be $g_p(x) = |x|$, such that the penalty function is just the ℓ_1 -norm, which is the same as in GPower $_{\ell_1}$. We call the resulting algorithm “MM $_{\ell_1}$ ” and the Algorithm 4 is referred to as “MM $_{\ell_0}$ ” in this section.

Note that for GPower methods, direct access to the original data matrix \mathbf{C} is required. When only the covariance matrix is available, a factorization of the form $\mathbf{A} = \mathbf{C}^T \mathbf{C}$ is needed (e.g., by eigenvalue decomposition or by Cholesky decomposition). If the data matrix \mathbf{C} is of size $m \times n$, then the per-iteration computational cost is $\mathcal{O}(mn)$ for all the four algorithms under consideration.

1) *Computational Complexity*: In this subsection, we compare the computational complexity of the four algorithms mentioned above, i.e., GPower $_{\ell_1}$, GPower $_{\ell_0}$, MM $_{\ell_1}$ and MM $_{\ell_0}$. The data matrix $\mathbf{C} \in \mathbf{R}^{n \times n}$ is generated randomly with the entries independent, identically distributed and following $\mathcal{N}(0, 1)$. The stopping condition is set to be $|f(\mathbf{x}^{(k+1)}) - f(\mathbf{x}^{(k)})| / \max(1, |f(\mathbf{x}^{(k)})|) \leq 10^{-5}$ for all the algorithms. The smoothing parameter for algorithm MM $_{\ell_1}$ is fixed to be $\epsilon = 10^{-8}$. The regularization parameter ρ is chosen such that the solutions of the four algorithms exhibit similar cardinalities (with about 5% nonzero entries).

The average running time over 100 independent trials for problems with different sizes are shown in Fig. 10. From the figure, we can see that the two ℓ_0 -norm penalized methods are faster than the two ℓ_1 -norm penalized methods and the proposed MM $_{\ell_0}$ is the fastest among the four algorithms, especially for problems of large size. For the two ℓ_1 -norm penalized methods, the proposed MM $_{\ell_1}$ is slower than GPower $_{\ell_1}$, which may result from the fact that MM $_{\ell_1}$ minorizes both the quadratic term and the ℓ_1 penalty term while GPower $_{\ell_1}$ keeps the ℓ_1 penalty term. It is worth noting that the GPower $_{\ell_1}$ is specialized for ℓ_1 penalty, while Algorithm 3 can also deal with various surrogate functions other than the ℓ_1 penalty.

2) *Random Data Drawn From a Sparse PCA Model*: In this subsection, we follow the procedure in [6] to generate random data with a covariance matrix having sparse eigenvectors. To




Fig. 10. Average running time versus problem size. Each curve is an average of 100 random trials.

achieve this, we first construct a covariance matrix through the eigenvalue decomposition $\mathbf{A} = \mathbf{V}\text{Diag}(\mathbf{d})\mathbf{V}^T$, where the first k columns of $\mathbf{V} \in \mathbf{R}^{n \times n}$ are pre-specified sparse orthonormal vectors. A data matrix $\mathbf{C} \in \mathbf{R}^{m \times n}$ is then generated by drawing m samples from a zero-mean normal distribution with covariance matrix \mathbf{A} , that is, $\mathbf{C} \sim \mathcal{N}(\mathbf{0}, \mathbf{A})$.

Following the settings in [7], we choose $n = 500$, $k = 2$, and $m = 50$, where the two orthonormal eigenvectors are specified as follows

$$\begin{cases} V_{i,1} = \frac{1}{\sqrt{10}} & \text{for } i = 1, \dots, 10, \\ V_{i,1} = 0 & \text{otherwise,} \\ V_{i,2} = \frac{1}{\sqrt{10}} & \text{for } i = 11, \dots, 20, \\ V_{i,2} = 0 & \text{otherwise.} \end{cases}$$

The eigenvalues are fixed at $d_1 = 400$, $d_2 = 300$ and $d_i = 1$, for $i = 3, \dots, 500$.

We randomly generate 500 data matrices $\mathbf{C} \in \mathbf{R}^{m \times n}$ and employ the four algorithms to compute the leading sparse eigenvector $\mathbf{x}_1 \in \mathbf{R}^{500}$, which is hoped to recover $\mathbf{V}_{:,1}$. We consider the underlying sparse eigenvector $\mathbf{V}_{:,1}$ is successfully recovered when $|\mathbf{x}_1^T \mathbf{V}_{:,1}| > 0.99$. The chance of successful recovery over a wide range of regularization parameter ρ is plotted in Fig. 11. The horizontal axis shows the normalized regularization parameter, that is $\rho / \max_i \|\mathbf{C}_{:,i}\|_2$ for GPower_{ℓ_1} and $\rho / \max_i \|\mathbf{C}_{:,i}\|_2^2$ for GPower_{ℓ_0} and MM_{ℓ_0} . For MM_{ℓ_1} algorithm, we use $\rho / (2\|\mathbf{C}^T \mathbf{C}\|_{\infty,2})$, where $\|\cdot\|_{\infty,2}$ is the operator norm induced by $\|\cdot\|_{\infty}$ and $\|\cdot\|_2$. From the figure, we can see that the highest chance of exact recovery achieved by the four algorithms is the same and for all algorithms it is achieved over a relatively wide range of ρ .

3) *Gene Expression Data*: DNA microarrays allow measuring the expression level of thousands of genes at the same time and this opens the possibility to answer some complex biological questions. But the amount of data created in an experiment is usually large and this makes the interpretation of these data challenging. PCA has been applied as a tool in the studies of gene expression data and their interpretation [28]. Naturally,




Fig. 11. Chance of exact recovery versus normalized regularization parameter.

sparse PCA, which extracts principal components with only a few nonzero elements can potentially enhance the interpretation.

In this subsection, we test the performance of the algorithms on gene expression data collected in the breast cancer study by Bild *et al.* [29]. The data set contains 158 samples over 12625 genes, resulting in a 158×12625 data matrix. Fig. 12 shows the explained variance versus cardinality for five algorithms, including the simple thresholding scheme. The proportion of explained variance is computed as $\mathbf{x}_{\text{SPCA}}^T (\mathbf{C}^T \mathbf{C}) \mathbf{x}_{\text{SPCA}} / \mathbf{x}_{\text{PCA}}^T (\mathbf{C}^T \mathbf{C}) \mathbf{x}_{\text{PCA}}$, where \mathbf{x}_{SPCA} is the sparse eigenvector extracted by sparse PCA algorithms, \mathbf{x}_{PCA} is the true leading eigenvector and \mathbf{C} is the data matrix. The simple thresholding scheme first computes the regular principal component \mathbf{x}_{PCA} and then keeps a required number of entries with largest absolute values. From the figure, we can see that the proportion of variance being explained increases as the cardinality increases as expected. For a fixed cardinality, the two GPower algorithms and the two proposed MM algorithms can explain almost the same amount of variance, all higher than the simple thresholding scheme, especially when the cardinality is small.

VI. CONCLUSION

We have developed an efficient algorithm IRQM that allows to obtain sparse generalized eigenvectors of a matrix pair. After approximating the ℓ_0 -norm penalty by some nonconvex surrogate functions, the minorization-maximization scheme is applied and the sparse generalized eigenvalue problem is turned into a sequence of regular generalized eigenvalue problems. The convergence to a stationary point is proved. Numerical experiments show that the proposed IRQM algorithm outperforms an existing algorithm based on D.C. programming in terms of both computational cost and support recovery. For sparse generalized eigenvalue problems with special structure (but still including sparse PCA as a special instance), two more efficient algorithms that have a closed-form solution at every iteration are derived again based on the minorization-maximization scheme. On both




Fig. 12. Trade-off curves between explained variance and cardinality.

synthetic random data and real-life gene expression data, the two algorithms are shown experimentally to have similar performance to the state-of-the-art.

APPENDIX A

PROOF OF PROPOSITION 4

Proof: From Lemma 3, it is easy to show that

$$f_\epsilon(\mathbf{x}) \geq f(\mathbf{x}) \geq f_\epsilon(\mathbf{x}) - \rho n \left(g_p(\epsilon) - \frac{g'_p(\epsilon)}{2} \epsilon \right), \quad \forall \mathbf{x} \in \mathbf{R}^n. \quad (34)$$

Since problems (5) and (18) have the same constraint set, we have

$$f_\epsilon(\mathbf{x}_\epsilon^*) \geq f(\mathbf{x}^*). \quad (35)$$

From the fact that \mathbf{x}^* is a global maximizer of problem (5), we know

$$f(\mathbf{x}^*) \geq f(\mathbf{x}_\epsilon^*). \quad (36)$$

Combining (34), (35) and (36), yields

$$f_\epsilon(\mathbf{x}_\epsilon^*) \geq f(\mathbf{x}^*) \geq f(\mathbf{x}_\epsilon^*) \geq f_\epsilon(\mathbf{x}_\epsilon^*) - \rho n \left(g_p(\epsilon) - \frac{g'_p(\epsilon)}{2} \epsilon \right).$$

Thus,

$$0 \leq f(\mathbf{x}^*) - f(\mathbf{x}_\epsilon^*) \leq \rho n \left(g_p(\epsilon) - \frac{g'_p(\epsilon)}{2} \epsilon \right).$$

Since $g_p(\cdot)$ is concave and monotone increasing on $[0, +\infty)$, it is easy to show that $g_p(\epsilon) \geq g'_p(\epsilon)\epsilon \geq 0$, for any $\epsilon > 0$. Hence

$$g_p(\epsilon) \geq g_p(\epsilon) - \frac{g'_p(\epsilon)}{2} \epsilon \geq 0. \quad (37)$$

Since $g_p(\cdot)$ is continuous and monotone increasing on $[0, +\infty)$ and $g_p(0) = 0$, we have $\lim_{\epsilon \downarrow 0} g_p(\epsilon) = 0$. Together with (37), we can conclude that

$$\lim_{\epsilon \downarrow 0} \left(g_p(\epsilon) - \frac{g'_p(\epsilon)}{2} \epsilon \right) = 0.$$

Since ρ and n are constants, the proof is complete. \blacksquare

APPENDIX B

PROOF OF PROPOSITION 5

Proof: First notice that the problem (22) is a nonconvex QCQP but with only one constraint, thus the strong duality holds [30], [31]. The optimality conditions for this problem are

$$(\mu \text{Diag}(\mathbf{b}) + \rho \text{Diag}(\mathbf{w}))\mathbf{x} = \mathbf{a} \quad (38)$$

$$\mathbf{x}^T \text{Diag}(\mathbf{b})\mathbf{x} = 1 \quad (39)$$

$$\mu \text{Diag}(\mathbf{b}) + \rho \text{Diag}(\mathbf{w}) \succeq \mathbf{0}. \quad (40)$$

Let us define

$$\mathcal{I}_{\min} = \arg \min \{ \rho w_i / b_i : i \in \{1, \dots, n\} \} \quad (41)$$

and

$$\mu_{\min} = - \min \{ \rho w_i / b_i : i \in \{1, \dots, n\} \}. \quad (42)$$

Then the third optimality condition (40) is just $\mu \geq \mu_{\min}$, since $w_i > 0$, $\rho > 0$, $b_i > 0$. Let us consider the optimality condition in two different cases:

- 1) $\mu > \mu_{\min}$. In this case, $\mu \text{Diag}(\mathbf{b}) + \rho \text{Diag}(\mathbf{w}) \succ \mathbf{0}$, from the first optimality condition (38) we get

$$\mathbf{x} = (\text{Diag}(\mu \mathbf{b} + \rho \mathbf{w}))^{-1} \mathbf{a}. \quad (43)$$

Substituting it into the second optimality condition (39), yields

$$\sum_{i \in \mathcal{I}_{\min}} \frac{b_i a_i^2}{(\mu b_i + \rho w_i)^2} + \sum_{i \notin \mathcal{I}_{\min}} \frac{b_i a_i^2}{(\mu b_i + \rho w_i)^2} = 1, \quad (44)$$

and it is easy to see that the left hand side is monotonically decreasing for $\mu \in (\mu_{\min}, +\infty)$. If $\exists i \in \mathcal{I}_{\min}$, such that $a_i^2 > 0$, then the left hand side of (44) tends to $+\infty$ as $\mu \rightarrow \mu_{\min}$. Notice that the left hand side goes to 0 as $\mu \rightarrow +\infty$, thus we are guaranteed to find a $\mu \in (\mu_{\min}, +\infty)$ satisfying (44). In practice, we may use bisection method to find the value of μ . If $a_i^2 = 0$, $\forall i \in \mathcal{I}_{\min}$, there still exists a $\mu > \mu_{\min}$ that satisfies (44) if and only if

$$\sum_{i \notin \mathcal{I}_{\min}} \frac{b_i a_i^2}{(\mu_{\min} b_i + \rho w_i)^2} > 1. \quad (45)$$

If (45) does not hold, it implies $\mu = \mu_{\min}$.

- 2) $\mu = \mu_{\min}$. In this case, we cannot compute \mathbf{x} via (43) anymore. Then to obtain \mathbf{x} , we first notice from (38) that

$$x_i = \frac{a_i}{\mu_{\min} b_i + \rho w_i}, \forall i \notin \mathcal{I}_{\min}. \quad (46)$$

Then, according to (39), for $x_i, i \in \mathcal{I}_{\min}$, they just need to satisfy the following equation

$$\sum_{i \in \mathcal{I}_{\min}} b_i x_i^2 = 1 - \sum_{i \notin \mathcal{I}_{\min}} \frac{b_i a_i^2}{(\mu_{\min} b_i + \rho w_i)^2}. \quad (47)$$

When $\text{card}(\mathcal{I}_{\min}) > 1$, (47) has infinite number of solutions and we may choose arbitrary one. ■

APPENDIX C

PROOF OF PROPOSITION 6

Proof: The problem (23) can be rewritten as

$$\text{maximize}_{s \in \{1, \dots, n\}} \left\{ -\rho s + \max_{\mathbf{x}} \{ \mathbf{a}^T \mathbf{x} : \|\mathbf{x}\|_2 = 1, \|\mathbf{x}\|_0 \leq s \} \right\}.$$

The inner maximization has a closed-form solution

$$x_i^* = \begin{cases} a_i / \sqrt{\sum_{j=1}^s a_j^2}, & i \leq s \\ 0 & \text{otherwise,} \end{cases}$$

then the problem becomes

$$\text{maximize}_{s \in \{1, \dots, n\}} \left\{ -\rho s + \sqrt{\sum_{i=1}^s a_i^2} \right\}.$$

It's easy to see that the optimal s is the largest integer p that satisfies the following inequality

$$\sqrt{\sum_{i=1}^p a_i^2} > \sqrt{\sum_{i=1}^{p-1} a_i^2} + \rho. \quad (48)$$

By squaring both sides of this inequality, we get

$$a_p^2 > \rho^2 + 2\rho \sqrt{\sum_{i=1}^{p-1} a_i^2},$$

which means $|a_p| > \rho$ is a necessary condition for (48) to be satisfied. Thus, in practice to find the largest integer p that satisfies (48) we only need to check for all p 's with $|a_p| > \rho$. If $0 < |a_1| \leq \rho$, it is easy to see that the solution of the problem (23) is given by (24). ■

APPENDIX D

PROOF OF LEMMA 7

Proof: From the way $g_p^\epsilon(x)$ is constructed, it is continuously differentiable. It remains to show that the gradient

$$(g_p^\epsilon)'(x) = \begin{cases} \frac{g_p'(\epsilon)}{\epsilon} x, & |x| \leq \epsilon \\ g_p'(x), & |x| > \epsilon \end{cases}$$

is Lipschitz continuous. From the fact that $g_p(x)$ is concave and monotone increasing on $(0, +\infty)$, we know that $g_p'(x)$ is non-increasing on $(0, +\infty)$ and $g_p'(x) > 0$. Since $g_p^\epsilon(x)$ is an even function, $(g_p^\epsilon)'(x)$ is odd. Thus, $(g_p^\epsilon)'(x)$ is non-increasing on

$(-\infty, -\epsilon)$, linearly increasing on $[-\epsilon, \epsilon]$ and non-increasing on $(\epsilon, +\infty)$. In addition, $(g_p^\epsilon)'(x) \leq 0$ when $x < 0$ and $(g_p^\epsilon)'(x) \geq 0$ when $x > 0$. With $(g_p^\epsilon)'(x)$ having these properties, to show the Lipschitz continuity of $(g_p^\epsilon)'(x)$, it is sufficient to show that $(g_p^\epsilon)'(x)$ is Lipschitz continuous on $[-\epsilon, \epsilon]$ and $(\epsilon, +\infty)$ respectively.

On $[-\epsilon, \epsilon]$, $(g_p^\epsilon)'(x) = \frac{g_p'(\epsilon)}{\epsilon} x$, which is Lipschitz continuous with Lipschitz constant $\frac{g_p'(\epsilon)}{\epsilon}$.

On $(\epsilon, +\infty)$, $(g_p^\epsilon)'(x) = g_p'(x)$, from Assumption 1 we know that $g_p'(x)$ is convex and differentiable on $(0, +\infty)$. Since $g_p'(x)$ is also non-increasing, we can conclude that $g_p''(x) \leq 0$ and is non-decreasing on $(0, +\infty)$. Thus, on $(\epsilon, +\infty)$, $|g_p''(x)|$ is bounded by $|g_p''(\epsilon)|$ and the Lipschitz continuity of $(g_p^\epsilon)'(x)$ on $(\epsilon, +\infty)$ follows. ■

REFERENCES

- [1] I. T. Jolliffe, *Principal Component Analysis*. New York, NY, USA: Springer-Verlag, 2002.
- [2] H. Hotelling, "Relations between two sets of variates," *Biometrika*, vol. 28, no. 3-4, pp. 321-377, Dec. 1936.
- [3] J. Cadima and I. T. Jolliffe, "Loading and correlations in the interpretation of principle components," *J. Appl. Statist.*, vol. 22, no. 2, pp. 203-214, 1995.
- [4] H. Zou, T. Hastie, and R. Tibshirani, "Sparse principal component analysis," *J. Computat. Graph. Statist.*, vol. 15, no. 2, pp. 265-286, 2006.
- [5] A. d'Aspremont, L. E. Ghaoui, M. I. Jordan, and G. R. Lanckriet, "A direct formulation for sparse PCA using semidefinite programming," *SIAM Rev.*, vol. 49, no. 3, pp. 434-448, 2007.
- [6] H. Shen and J. Z. Huang, "Sparse principal component analysis via regularized low rank matrix approximation," *J. Multivar. Anal.*, vol. 99, no. 6, pp. 1015-1034, Jul. 2008.
- [7] M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre, "Generalized power method for sparse principal component analysis," *J. Mach. Learn. Res.*, vol. 11, pp. 517-553, Mar. 2010.
- [8] R. Luss and M. Teboulle, "Conditional gradient algorithms for rank-one matrix approximations with a sparsity constraint," *SIAM Rev.*, vol. 55, no. 1, pp. 65-98, 2013.
- [9] D. M. Witten, R. Tibshirani, and T. Hastie, "A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis," *Biostatistics*, vol. 10, no. 3, pp. 515-534, Jul. 2009.
- [10] B. K. Sriperumbudur, D. A. Torres, and G. R. G. Lanckriet, "Sparse eigen methods by D.C. programming," in *Proc. 24th Int. Conf. Mach. Learn.*, New York, NY, USA, 2007, pp. 831-838.
- [11] B. K. Sriperumbudur, D. A. Torres, and G. R. G. Lanckriet, "A majorization-minimization approach to the sparse generalized eigenvalue problem," *Mach. Learn.*, vol. 85, no. 1-2, pp. 3-39, Oct. 2011.
- [12] E. Schlossmacher, "An iterative technique for absolute deviations curve fitting," *J. Amer. Statist. Assoc.*, vol. 68, no. 344, pp. 857-859, Dec. 1973.
- [13] M. A. Figueiredo, J. M. Bioucas-Dias, and R. D. Nowak, "Majorization-minimization algorithms for wavelet-based image restoration," *IEEE Trans. Image Process.*, vol. 16, no. 12, pp. 2980-2991, Dec. 2007.
- [14] Y. Nesterov, "Smooth minimization of non-smooth functions," *Math. Program.*, vol. 103, no. 1, pp. 127-152, May 2005.
- [15] I. F. Gorodnitsky and B. D. Rao, "Sparse signal reconstruction from limited data using FOCUSS: A re-weighted minimum norm algorithm," *IEEE Trans. Signal Process.*, vol. 45, no. 3, pp. 600-616, Mar. 1997.
- [16] R. Chartrand and W. Yin, "Iteratively reweighted algorithms for compressive sensing," in *Proc. 33rd Int. Conf. Acoust., Speech Signal Process.*, Las Vegas, NV, USA, Apr. 2008, pp. 3869-3872.
- [17] E. J. Candes, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *J. Fourier Anal. Appl.*, vol. 14, no. 5-6, pp. 877-905, Dec. 2008.
- [18] O. L. Mangasarian, "Machine learning via polyhedral concave minimization," in *Applied Mathematics and Parallel Computing*. New York, NY, USA: Springer-Verlag, 1996, pp. 175-188.

- [19] D. R. Hunter and K. Lange, "A tutorial on MM algorithms," *Amer. Statist.*, vol. 58, no. 1, pp. 30–37, Feb. 2004.
- [20] M. Razaviyayn, M. Hong, and Z.-Q. Luo, "A unified convergence analysis of block successive minimization methods for nonsmooth optimization," *SIAM J. Optimiz.*, vol. 23, no. 2, pp. 1126–1153, 2013.
- [21] G. Scutari, F. Facchinei, P. Song, D. P. Palomar, and J.-S. Pang, "Decomposition by partial linearization: Parallel optimization of multiagent systems," *IEEE Trans. Signal Process.*, vol. 62, no. 3, pp. 641–656, Feb. 2014.
- [22] P. W. Holland and R. E. Welsch, "Robust regression using iteratively reweighted least-squares," *Commun. Statist.-Theory Methods*, vol. 6, no. 9, pp. 813–827, Jan. 1977.
- [23] B. D. Rao and K. Kreutz-Delgado, "An affine scaling methodology for best basis selection," *IEEE Trans. Signal Process.*, vol. 47, no. 1, pp. 187–200, Jan. 1999.
- [24] S. Becker, J. Bobin, and E. J. Candès, "NESTA: A fast and accurate first order method for sparse recovery," *SIAM J. Imag. Sci.*, vol. 4, no. 1, pp. 1–39, 2011.
- [25] Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, "Templates for the solution of algebraic eigenvalue problems: A practical guide," in *Proc. SIAM*, 2000, vol. 11.
- [26] S. Zlobec, "On the Liu-Floudas convexification of smooth programs," *J. Global Optimiz.*, vol. 32, no. 3, pp. 401–407, Jul. 2005.
- [27] R. T. Rockafellar, *Convex Analysis*. Princeton, NJ, USA: Princeton Univ. Press, 1997, vol. 28.
- [28] A. Riva, A.-S. Carpentier, B. Torrèsani, and A. Hénaut, "Comments on selected fundamental aspects of microarray analysis," *Computat. Biol. Chem.*, vol. 29, no. 5, pp. 319–336, Oct. 2005.
- [29] A. H. Bild *et al.*, "Oncogenic pathway signatures in human cancers as a guide to targeted therapies," *Nature*, vol. 439, no. 7074, pp. 353–357, Jan. 2006.
- [30] Y. Huang and S. Zhang, "Complex matrix decomposition and quadratic programming," *Math. Operat. Res.*, vol. 32, no. 3, pp. 758–768, 2007.
- [31] Y. Huang and D. P. Palomar, "Rank-constrained separable semidefinite programming with applications to optimal beamforming," *IEEE Trans. Signal Process.*, vol. 58, no. 2, pp. 664–678, Feb. 2010.



Junxiao Song received the B.S. degree in control science and engineering from Zhejiang University (ZJU), Hangzhou, China, in 2011.

Currently, he is pursuing the Ph.D. degree with the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology (HKUST). His research interests are in convex optimization and efficient algorithms, with applications in signal processing and financial engineering.

Prabhu Babu received the Ph.D. degree in electrical engineering from the Uppsala University, Sweden, in 2012.

He is currently a Research Associate with the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology (HKUST), Hong Kong.



Daniel P. Palomar (S'99–M'03–SM'08–F'12) received the electrical engineering and Ph.D. degrees from the Technical University of Catalonia (UPC), Barcelona, Spain, in 1998 and 2003, respectively.

He is a Professor in the Department of Electronic and Computer Engineering, the Hong Kong University of Science and Technology (HKUST), Hong Kong, which he joined in 2006. Since 2013, he has been a Fellow of the Institute for Advance Study (IAS) at HKUST. He had previously held several research appointments, namely, at King's College London (KCL), London, UK; Technical University of Catalonia (UPC), Barcelona; Stanford University, Stanford, CA; Telecommunications Technological Center of Catalonia (CTTC), Barcelona; Royal Institute of Technology (KTH), Stockholm, Sweden; University of Rome "La Sapienza", Rome, Italy; and Princeton University, Princeton, NJ. His current research interests include applications of convex optimization theory, game theory, and variational inequality theory to financial systems and communication systems.

Dr. Palomar is a recipient of a 2004/2006 Fulbright Research Fellowship, the 2004 Young Author Best Paper Award by the IEEE Signal Processing Society, the 2002/2003 Best Ph.D. prize in information technologies and communications by the Technical University of Catalonia (UPC), the 2002/2003 Rosina Ribalta first prize for the Best Doctoral Thesis in information technologies and communications by the Epson Foundation, and the 2004 prize for the Best Doctoral Thesis in Advanced Mobile Communications by the Vodafone Foundation and COIT. He has been an Associate Editor of the IEEE TRANSACTIONS ON INFORMATION THEORY, an Associate Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING, a Guest Editor of the IEEE SIGNAL PROCESSING MAGAZINE 2010 Special Issue on "Convex Optimization for Signal Processing," the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS 2008 Special Issue on "Game Theory in Communication Systems," and the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS 2007 Special Issue on "Optimization of MIMO Transceivers for Realistic Communication Networks."