
An Inverse-free Truncated Rayleigh-Ritz Method for Sparse Generalized Eigenvalue Problem

Yunfeng Cai and Ping Li

Cognitive Computing Lab

Baidu Research

No. 10 Xibeiwang East Road, Beijing 100085, China

10900 NE 8th St. Bellevue, WA 98004, USA

{caiyunfeng, liping11}@baidu.com

Abstract

This paper considers the sparse generalized eigenvalue problem (SGEP), which aims to find the leading eigenvector with at most k nonzero entries. SGEP naturally arises in many applications in machine learning, statistics, and scientific computing, for example, the sparse principal component analysis (SPCA), the sparse discriminant analysis (SDA), and the sparse canonical correlation analysis (SCCA). In this paper, we focus on the development of a three-stage algorithm named *inverse-free truncated Rayleigh-Ritz method (IFTRR)* to efficiently solve SGEP. In each iteration of IFTRR, only a small number of matrix-vector products is required. This makes IFTRR well-suited for large scale problems. Particularly, a new truncation strategy is proposed, which is able to find the support set of the leading eigenvector effectively. Theoretical results are developed to explain why IFTRR works well. Numerical simulations demonstrate the merits of IFTRR.

1 Introduction

Given a matrix pair (\tilde{A}, \tilde{B}) , where \tilde{A}, \tilde{B} are both p -by- p symmetric matrices and \tilde{B} is (semi) positive definite, the sparse generalized eigenvalue problem (sparse GEP, or SGEP) aims to maximize the Rayleigh quotient $\frac{v^T \tilde{A} v}{v^T \tilde{B} v}$ with $v \in \mathbb{R}^p$ having no more than k nonzero entries, where $k \ll p$. Mathematically, SGEP can be formulated as the following optimization problem:

$$\max_{v \in \mathbb{R}^p} \frac{v^T \tilde{A} v}{v^T \tilde{B} v}, \quad \text{subject to } \|v\|_0 \leq k, \quad (1)$$

where $\|v\|_0$ denotes the ℓ_0 -norm of v , which is the number of nonzero entries of v . In many applications, such as sparse principle component analysis (SPCA) (Zou et al., 2006), sparse discriminant analysis (SDA) (Clemmensen et al., 2011), and sparse canonical correlation analysis (SCCA) (Witten et al., 2009) in high dimensional statistical analysis, the matrices \tilde{A} and \tilde{B} usually can be decomposed as follows:

$$\tilde{A} = A + E, \quad \tilde{B} = B + F, \quad (2)$$

where A, B are both symmetric, B is positive definite, E, F are symmetric perturbations due to finite sample estimation.

Next, we consider the following two concrete examples of SGEP, which arise from Sparse Fisher's discriminant analysis (SFDA) and SCCA, respectively.

Example 1 Given a data matrix $X \in \mathbb{R}^{n \times p}$ (i.e., n observations with p features), each row belongs to one of K classes. Denote $C_k \subset \{1, 2, \dots, n\}$ the indices of the observations in the k -th class, $n_k = |C_k|$, $\bar{x}_k = \sum_{i \in C_k} \frac{X_{(i,:)}}{n_k}$. Then we may use

$$\begin{aligned} \tilde{\Sigma}_b &= \sum_{k=1}^K \frac{n_k \bar{x}_k^T \bar{x}_k}{n}, \\ \tilde{\Sigma}_w &= \frac{1}{n} \sum_{k=1}^K \sum_{i \in C_k} (X_{(i,:)} - \bar{x}_k)^T (X_{(i,:)} - \bar{x}_k), \end{aligned}$$

as the estimators for the between class variance Σ_b and the within class variance Σ_w , respectively. SFDA intends to find a sparse leading discriminant vector v that maximizes $\frac{v^T \tilde{\Sigma}_b v}{v^T \tilde{\Sigma}_w v}$, which can be formulated as an SGEP with $\tilde{A} = \tilde{\Sigma}_b, \tilde{B} = \tilde{\Sigma}_w$.

Example 2 Let $X \in \mathbb{R}^{p/2}$ and $Y \in \mathbb{R}^{p/2}$ be two random variables, $\Sigma_{xx}, \Sigma_{yy}, \Sigma_{xy}$ be the covariance matrices for X, Y , the cross-covariance matrix between X and Y , respectively, $\hat{\Sigma}_{xx}, \hat{\Sigma}_{yy}, \hat{\Sigma}_{xy}$ be estimators

for Σ_{xx} , Σ_{yy} , Σ_{xy} , respectively. SCCA aims to maximize $v_x^T \hat{\Sigma}_{xy} v_y$, subject to $v_x^T \hat{\Sigma}_{xx} v_x = 1$, $v_y^T \hat{\Sigma}_{yy} v_y = 1$, $\|v_x\|_0 \leq s_x$, $\|v_y\|_0 \leq s_y$, where s_x and s_y are two small integers. Such a problem can be reformulated as an SGEP with $\tilde{A} = \begin{bmatrix} 0 & \hat{\Sigma}_{xy} \\ \hat{\Sigma}_{xy}^T & 0 \end{bmatrix}$, $\tilde{B} = \begin{bmatrix} \hat{\Sigma}_{xx} & 0 \\ 0 & \hat{\Sigma}_{yy} \end{bmatrix}$, $v = \begin{bmatrix} v_x \\ v_y \end{bmatrix}$.

Computational challenges. The SGEP can be computationally challenging. Recall that in our model we have $\tilde{A} = A + E$, $\tilde{B} = B + F$. Denote the eigenvalues and the corresponding eigenvectors of $Av = \lambda Bv$ by $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p$ and v_1, v_2, \dots, v_p , respectively. We call λ_1 the leading eigenvalue of (A, B) , v_1 the leading eigenvector of (A, B) , and (λ_1, v_1) the leading eigenpair of (A, B) . In this paper, the leading eigenvector v_1 is assumed to be sparse, i.e., $\|v_1\|_0 \ll p$. Overall, the task of SGEP is essentially to find an approximation of v_1 via (\tilde{A}, \tilde{B}) , without knowing E, F .

Due to finite number of samples, the perturbations E and F may be large, and consequently, the leading eigenvector of (\tilde{A}, \tilde{B}) may not be a good approximation of the true v_1 . Furthermore, in a high dimensional setting, \tilde{A} and \tilde{B} can be both ill-conditioned (or singular), meaning that infinity eigenvalue ∞ and indeterminate eigenvalue $0/0$ occur. In the presence of rounding-off error, numerical algorithms may fail to detect the singularity due to ill-conditioning. As a result, one may not be able to compute accurate or meaningful eigenvalues and eigenvectors. In fact, the optimization problem (1) is essentially a subset selection problem, which is known to be NP-hard (Moghaddam et al., 2005, 2006).

Here, it is worth mentioning that, when $B = \tilde{B} = I_p$, the problem (1) is reduced to the so-called *sparse eigenvalue problem* (SEP), also known as *sparse principal component analysis* (SPCA). Obviously SGEP can be substantially more challenging than SEP (SPCA).

Related work. A variety of numerical methods have been proposed for SEP. Existing algorithms of SEP are mostly optimization approaches, which are based on relaxation, or penalization, or both. The ℓ_1 -norm relaxation, inspired by LASSO, is first studied in (Jolliffe et al., 2003) and called SCoTLASS. In (Witten et al., 2009), a penalized matrix decomposition method is proposed for computing a low rank approximation of a matrix, where ℓ_1 -norm relaxation is used to encourage sparsity. In (d’Aspremont et al., 2007), a convex relaxation for the ℓ_1 constrained PCA is introduced and solved by semidefinite programming. In (Journée et al., 2010), a generalized power (GPower) method is proposed for SPCA, where ℓ_0/ℓ_1 penalization is used. In (Luss and Teboulle, 2013), based on the well-known conditional gradient algorithm, a framework called ConGradU is proposed, which unifies a variety of algorithms. In (Yuan and Zhang, 2013), a truncated power method (TPower) is proposed, which adopts the

power method to update the approximate eigenvector, followed by a truncation procedure that keeps a few largest magnitude entries of the approximate eigenvector and truncates the remaining entries to zero. Also, see (d’Aspremont et al., 2008; Moghaddam et al., 2006) for other greedy algorithms proposed for SPCA.

The SGEP, compared with SEP, is less investigated, especially for the large ill-conditioned problems. SEP solvers such as GPower, ConGradU, and TPower only require matrix-vector product (MVP) operations and hence are efficient for large problems. For SGEP, however, those methods are no longer directly applicable. In (Sriperumbudur et al., 2011), SGEP is framed as a difference of convex functions program and solved via a sequence of convex programs where the majorization-minimization method is used. In (Song et al., 2015), SGEP is transformed into a sequence of regular GEP via quadratic minorization functions, and the preconditioned steepest ascent method is used to find the leading eigenpair. In (Safou et al., 2018), a general framework called sparse estimation with linear programming is proposed, where the leading eigenpair of (\tilde{A}, \tilde{B}) is used to simplify the constraint. In (Tan et al., 2018), the truncated Rayleigh flow method (RIFLE) is proposed, where the approximate eigenvector is updated via fixed step size steepest ascent method, and followed by simple truncation.

Our proposal – IFTRR. We propose an inverse-free truncated Rayleigh-Ritz (IFTRR) method for SGEP. The classical Rayleigh-Ritz method is an approximate algorithm for computing eigenvalue equations (Demmel, 1997; Stewart, 2001). IFTRR has three major steps: first, the approximated eigenvector is updated via an inverse-free generalized eigensolver; second, with the help of the updated eigenvector, a truncation procedure is used to find the support set for the approximate eigenvector in the next iteration; third, a small GEP is solved, and the approximate eigenvector is updated.

In the implementation of IFTRR, only matrix-vector product (MVP) is required, and hence the method is inherently suited for large scale problems. Furthermore, IFTRR is applicable for ill-conditioned or singular \tilde{A}, \tilde{B} . Additionally, the proposed truncation procedure, which is based on “eigenvalue increment” (see Section 2), is able to find the support set of the leading eigenvector effectively. Our numerical experiments (Section 3) show that IFTRR usually converges in a few iterations.

Notation. The symbol \otimes denotes the Kronecker product. The calligraphic letter \mathcal{I} is usually used to denote an index set, $|\mathcal{I}|$ denotes the cardinality of \mathcal{I} , e.g., $\mathcal{I} = \{i_1, i_2, \dots, i_s\}$, where i_1, i_2, \dots, i_s are distinct integers, then $|\mathcal{I}| = s$. Let $a = [a_1, a_2, \dots, a_p]^T \in \mathbb{R}^p$, $A = [a_{jk}] \in \mathbb{R}^{p \times p}$, $a_{\mathcal{I}}, A_{\mathcal{I}}$

stand for $[\alpha_{i_1}, \alpha_{i_2}, \dots, \alpha_{i_s}] \in \mathbb{R}^s$ and $[a_{i_j i_k}] \in \mathbb{R}^{s \times s}$, respectively. For $w = [w_1, w_2, \dots, w_p]^T \in \mathbb{R}^p$, $\text{supp}(w)$ denotes the index set of all nonzero entries of w , $\text{supp}(w, k)$ denotes the index set of the k largest magnitude entries of w , i.e., $\text{supp}(w) \triangleq \{i \mid w_i \neq 0\}$, $\text{supp}(w, k) \triangleq \{i_1, \dots, i_k \mid |w_{i_1}| \geq |w_{i_2}| \geq \dots \geq |w_{i_p}|\}$. $A_{(j,:)}$ and $A_{(:,k)}$ denote the j -th row and k -th column of A , respectively. I_p is the $p \times p$ identity matrix, and e_j is its j -th column. For symmetric definite matrix pairs (A, B) and (\tilde{A}, \tilde{B}) , we denote $\rho(v) = \frac{v^T A v}{v^T B v}$, $\tilde{\rho}(v) = \frac{v^T \tilde{A} v}{v^T \tilde{B} v}$. The i th largest eigenvalue of (A, B) is denoted by $\lambda_i(A, B)$.

2 The Inverse Free Truncated Rayleigh-Ritz Method (IFTRR)

In this section, we present the inverse-free truncated Rayleigh-Ritz (IFTRR) method for solving the sparse generalized eigenvalue problem (SGEP). We first explain the intuition behind the development of the algorithm, before we present the details of the algorithm.

The mechanism of IFTRR is deceptively simple: given an approximate eigenvector, update the vector via certain eigensolvers, then sparsify the resulting vector via a truncation procedure.

2.1 Eigensolvers for Generalized Eigenvalue Problem (GEP)

For large scale generalized eigenvalue problem (GEP) $\tilde{A}v = \lambda \tilde{B}v$, where $\tilde{A}, \tilde{B} \in \mathbb{R}^{p \times p}$ are symmetric and \tilde{B} is (semi) positive definite, iterative methods are usually used to compute its a few largest (or smallest) eigenvalues and the corresponding eigenvectors. Simply speaking, these iterative methods consist of two major steps: first, determine a subspace of \mathbb{R}^p ; second, update the approximate eigenpairs via the Rayleigh-Ritz procedure. Some detailed discussions follow.

The subspace. The most popular subspace for the eigenvalue problem is the Krylov subspace. Given a matrix $T \in \mathbb{R}^{p \times p}$ and a nonzero vector $v \in \mathbb{R}^p$, the order- m Krylov subspace is defined as

$$\mathcal{K}_m(T, v) \triangleq \text{span}\{v, Tv, T^2v, \dots, T^{m-1}v\}. \quad (3)$$

An orthonormal basis of $\mathcal{K}_m(T, v)$ can be obtained via the Arnoldi iteration, e.g., (Stewart, 2001, Chapter 5), (Demmel, 1997, Chapter 6). Here we would like to emphasize that since the Arnoldi iteration only requires matrix vector product (MVP) Tv , it is unnecessary to formulate T explicitly and a subroutine that computes Tv is sufficient. For GEP, T is usually set as $T = \tilde{B}^{-1}\tilde{A}$. As \tilde{B}^{-1} is involved, the implementation of the MVP Tv requires the MVP $u = \tilde{A}v$ and also solving the linear system $\tilde{B}z = u$ for z . When the matrix size p is

large, it is expensive to solve the linear system $\tilde{B}z = u$. More importantly, when \tilde{B} is ill-conditioned (or even singular), solving $\tilde{B}z = u$ is not only expensive but also prone to large numerical errors. In (Golub and Ye, 2002), Golub and Ye propose to use $\mathcal{K}_m(\tilde{A} - \rho\tilde{B}, v)$ to solve GEP, where $\rho \in \mathbb{R}$ is a shift. There are also other subspaces that can be used to solve GEP, e.g., the Davidson method (Davidson, 1975) and the Jacobi-Davidson method (Sleijpen and Van der Vorst, 2000). In this paper, we use the Krylov subspace $\mathcal{K}_m(\tilde{A} - \rho\tilde{B}, v)$, mainly due to its simplicity, scalability, and most importantly, it is inverse-free, since only MVP $(\tilde{A} - \rho\tilde{B})v$ is required.

The Rayleigh-Ritz procedure. Given an m -dimensional subspace \mathcal{V}_m of \mathbb{R}^p ($m \ll p$), let V_m be an orthonormal basis of \mathcal{V}_m . The Rayleigh-Ritz procedure has three steps: First, project the GEP $\tilde{A}v = \lambda \tilde{B}v$ onto \mathcal{V}_m , which yields a small GEP $(V^T \tilde{A} V)y = \mu(V^T \tilde{B} V)y$; Second, solve the eigenpairs (μ_i, y_i) for $i = 1, \dots, m$ of the small GEP; Third, compute $(\tilde{\lambda}_i, \tilde{v}_i) = (\mu_i, V_m y_i)$ for $i = 1, \dots, m$. Then $(\tilde{\lambda}_i, \tilde{v}_i)$'s, often referred to as *Ritz pairs*, are used as approximate eigenpairs of the original GEP $\tilde{A}v = \lambda \tilde{B}v$.

2.2 The Truncation Procedure

Let (ρ, w) be a Ritz pair, which is used to approximate the leading eigenpair of (\tilde{A}, \tilde{B}) . In general, w is dense. So, it is natural to sparsify w since we are solving a sparse vector. A popular way to accomplish this task is the so-called *truncation*, where the entries of w are truncated to zeros except for the first k largest magnitude entries. However, this truncation procedure is found to be potentially misleading (Cadima and Jolliffe, 1995). In this paper, we propose to do the ‘‘truncation procedure’’ as follows:

1. Find a permutation $\{i_1, i_2, \dots, i_p\}$ of $\{1, 2, \dots, p\}$ such that $|w_{i_1}| \geq |w_{i_2}| \geq \dots \geq |w_{i_p}|$.
2. For $s = s_1, s_1 + 1, \dots, s_2$, set $\mathcal{J} = \{i_1, i_2, \dots, i_s\}$ and compute the leading eigenpair of the small GEP $(\tilde{A}_{\mathcal{J}}, \tilde{B}_{\mathcal{J}})$, denoted by (ρ_s, z_s) , $s_1 < s_2$ are user-prescribed integers.
3. Determine the smallest s such that $\rho_{s_2} - \rho_s \leq \text{tol}$, where $\text{tol} > 0$ is a small real number.
4. Set $\mathcal{I} = \{i_1, i_2, \dots, i_s\}$, $\hat{v}_{\mathcal{I}} = z_s$ and $\hat{v}_{\mathcal{I}^c} = 0$.

The above truncation procedure is based on this observation: for any $\mathcal{J} \supseteq \mathcal{I} = \text{supp}(v_1)$, we have

$$\lambda_1 = \max_{z \neq 0} \frac{z^T A_{\mathcal{I}} z}{z^T B_{\mathcal{I}} z} \leq \max_{z \neq 0} \frac{z^T A_{\mathcal{J}} z}{z^T B_{\mathcal{J}} z} \leq \max_{z \neq 0} \frac{z^T A z}{z^T B z} = \lambda_1.$$

Thus, $\max_{z \neq 0} \frac{z^T A_{\mathcal{J}} z}{z^T B_{\mathcal{J}} z} \equiv \lambda_1$ as long as $\mathcal{J} \supseteq \mathcal{I}$. In other words, when \mathcal{J} is a superset of \mathcal{I} , the leading eigenvalue

of $(A_{\mathcal{J}}, B_{\mathcal{J}})$ remains a constant. Therefore, for (\tilde{A}, \tilde{B}) , we also expect that when \mathcal{J} contains \mathcal{I} and $|\mathcal{J}| \ll p$, the leading eigenvalue of $(\tilde{A}_{\mathcal{J}}, \tilde{B}_{\mathcal{J}})$ slightly changes.

Let $\hat{w}_{\mathcal{I}} = w$, $\hat{w}_{\mathcal{I}^c} = 0$. We prefer \hat{v} rather than the simple truncated vector \hat{w} simply because

$$\tilde{\rho}(\hat{v}) = \max_{\text{supp}(v) \subset \mathcal{I}} \tilde{\rho}(v) \geq \tilde{\rho}(\hat{w}),$$

i.e., the target value of $\tilde{\rho}(v)$ at $v = \hat{v}$ is no less than that at $v = \hat{w}$.

In order to compare the simple truncation method with our ‘‘eigenvalue increment’’ method, we take an approximate leading eigenvector $w = [w_1, \dots, w_p]^T$ of (\tilde{A}, \tilde{B}) from section 3.1 ($p = 1000$, $n = 200$, $s = 6$). We sort the entries of w such that $|w_{i_1}| \geq \dots \geq |w_{i_p}|$, and compute $\rho_s = \lambda_1(\tilde{A}_{\mathcal{J}_s}, \tilde{B}_{\mathcal{J}_s})$, where $\mathcal{J}_s = \{i_1, \dots, i_s\}$.




Figure 1: Absolute entries vs. eigenvalues.

In Figure 1, we plot the top 25 entries $|w_{i_1}|, \dots, |w_{i_{25}}|$ and the 25 eigenvalues ρ_1, \dots, ρ_{25} . We can see that there is no obvious gap among $|w_5|, \dots, |w_{18}|$, therefore, it is difficult to determine a proper k to do the simple truncation. On the other hand, $\rho_{13}, \dots, \rho_{25}$ almost remain unchanged, then we may use $\mathcal{I} = \{i_1, \dots, i_{13}\}$ as the support set for a new approximate eigenvector. In other words, $w_{i_{14}}, \dots, w_{i_{1000}}$ are truncated to zeros. As a matter of fact, for this example, the true support set of v_1 is contained in $\{i_1, \dots, i_{13}\}$.

The above example (and many others) show that our truncation procedure is effective to find the true support set of v_1 . As a result, the proposed algorithm usually converges in a few steps.

2.3 Algorithm

We present the IFTRR method in Algorithm 1. Basically, IFTRR consists of three parts: The first part (lines 5 to 7) updates the leading eigenvector of (\tilde{A}, \tilde{B}) via an eigensolver. The second part determines an index set \mathcal{J} , which serves as the support of $v^{(t)}$ (lines 8 to 13), then compute $v^{(t)}$ (line 17). The third part (line 20) determines \mathcal{I} and compute the final solution \tilde{v} with $\text{supp}(\tilde{v}) = \mathcal{I}$. The first two parts generate a

sequence $\{(\rho^{(t)}, v^{(t)})\}_t$, and the last part is the final update for the solution. Some implementation details and discussions of the algorithm follow.

Algorithm 1 The proposed inverse-free truncated Rayleigh-Ritz (IFTRR) method.

- 1: **Input:** \tilde{A}, \tilde{B} , an integer k for sparsity level, an integer m for Krylov subspace dimension, and a randomly generated initial guess $v^{(0)} \in \mathbb{R}^p$.
- 2: **Output:** An approximate solution \tilde{v} to (1).
- 3: Set $t = 0$, $\rho^{(0)} = \tilde{\rho}(v^{(0)})$, $s_1 = k$ and $s_2 = k + \Delta k$; (e.g., $\Delta k = 20$ or 30 .)
- 4: **while** unconverged and $t < \text{itermax}$ **do**
- 5: Compute an orthonormal $Q \in \mathbb{R}^{p \times m}$ such that $\text{span}(Q) = \mathcal{K}_m(\tilde{A} - \rho^{(t)}\tilde{B}, v^{(t)})$;
- 6: Solve the leading eigenvector \tilde{y} of $(Q^T \tilde{A} Q, Q^T \tilde{B} Q)$;
- 7: Set $w = Q\tilde{y}$, $w = w/\|w\|_2$;
- 8: Find a permutation $\{i_1, i_2, \dots, i_p\}$ of $\{1, 2, \dots, p\}$ such that $|w_{i_1}| \geq |w_{i_2}| \geq \dots \geq |w_{i_p}|$;
- 9: **for** $s = s_1, s_1 + 1, \dots, s_2$ **do**
- 10: Set $\mathcal{J} = \{i_1, i_2, \dots, i_s\}$;
- 11: Solve the leading eigenpair of $(\tilde{A}_{\mathcal{J}}, \tilde{B}_{\mathcal{J}})$, denote it by (ρ_s, z_s) ;
- 12: **end for**
- 13: Find the smallest $s \in \{s_1, s_1 + 1, \dots, s_2\}$ such that $\rho_{s_2} - \rho_s \leq (s_2 - s) \times \text{tol}$;
- 14: Set $\rho^{(t+1)} = \rho_s$, $v_{\mathcal{J}}^{(t+1)} = z_s/\|z_s\|_2$, $v_{\mathcal{J}^c}^{(t+1)} = 0$, where $\mathcal{J} = \{i_1, i_2, \dots, i_s\}$;
- 15: Set $t = t + 1$;
- 16: **end while**
- 17: Let $\mathcal{I} = \text{supp}(v^{(t)}, k)$, solve the leading eigenvector z of $(\tilde{A}_{\mathcal{I}}, \tilde{B}_{\mathcal{I}})$, set $\tilde{v}_{\mathcal{I}} = z/\|z\|_2$, $\tilde{v}_{\mathcal{I}^c} = 0$.

Convergence test. The maximum number of iterations itermax is set as 100. The algorithm converges if one of the following conditions is satisfied:

- (i) $\frac{\|(\tilde{A} - \rho^{(t)}\tilde{B})v^{(t)}\|_2}{\|\tilde{A}\|_2 + |\rho^{(t)}|\|\tilde{B}\|_2} < \text{tol}_1$, where tol_1 is a tolerance and is set as 0.01 in our experiments. This condition indicates that $(\rho^{(t)}, v^{(t)})$ is a good approximation of an eigenpair of (\tilde{A}, \tilde{B}) .
- (ii) $|\rho^{(t)} - \rho^{(t-1)}| < \text{tol}_2$, where tol_2 is a prescribed tolerance, say 10^{-3} . This condition indicates that the value of $\rho^{(t)}$ stagnates, hence we may take $\{\rho^{(t)}\}$ as a converged sequence.

Solving a sequence of GEPs. At first glance, a sequence of generalized eigenvalue problems needs to be solved on lines 9 to 12. But notice that

$$\rho_{s+1} = \max_{z \neq 0} \frac{z^T \tilde{A}_{\mathcal{J}_2} z}{z^T \tilde{B}_{\mathcal{J}_2} z} \geq \frac{\begin{bmatrix} z_s \\ 0 \end{bmatrix}^T \tilde{A}_{\mathcal{J}_2} \begin{bmatrix} z_s \\ 0 \end{bmatrix}}{\begin{bmatrix} z_s \\ 0 \end{bmatrix}^T \tilde{B}_{\mathcal{J}_2} \begin{bmatrix} z_s \\ 0 \end{bmatrix}} = \rho_s,$$

where $\mathcal{J}_1 = \{i_1, i_2, \dots, i_s\}$, $\mathcal{J}_2 = \{i_1, i_2, \dots, i_s, i_{s+1}\}$.

Thus $\{\rho_s\}_{s=s_1}^{s_2}$ is a non-decreasing sequence and we use the idea of *bisection* to find the desired s on line 13:

```

Set  $a = s_1, b = s_2$ ;
Compute  $\rho_a, \rho_b$ ;
While  $b - a > 1$  do
    Set  $s = \frac{a+b + \text{mod}(a+b, 2)}{2}$ ;
    Compute  $\rho_s$ ;
    If  $\rho_{s_2} - \rho_s \leq \text{tol}$ , set  $b = s, \rho_b = \rho_s$ ;
    Otherwise, set  $a = s, \rho_a = \rho_s$ .
End while
    
```

As a result, there are approximately $\log_2(s_2 - s_1 + 1)$ small GEPs rather than $s_2 - s_1 + 1$.

Dealing with ill-conditioning. When \tilde{B} is singular, $\tilde{B}_{\mathcal{J}}$ can also be singular. As a result, infinity or indeterminate eigenvalues, which are sensitive to perturbations (Bai et al., 2000), occur, then it will be difficult to determine which eigenpair is the leading one. In our implementation, we use the following procedure as a cure for singular $\tilde{B}_{\mathcal{J}}$: First, compute the QR decomposition of $\tilde{B}_{\mathcal{J}}$ with column pivoting (Van Loan and Golub, 2012, Chapter 5.4.2) $\tilde{B}_{\mathcal{J}}\Pi = UR$, where Π is a permutation matrix, U is an orthogonal matrix and R is upper triangular with its diagonal entries non-negative and non-increasing; Then let $\text{tol}_3 > 0$ be a user given threshold, say $\text{tol}_3 = 10^{-9}$. Whenever $R_{(i,i)} < \text{tol}_3 \times R_{(1,1)}$, we remove the corresponding index in \mathcal{J} . For the resulting \mathcal{J} , $\tilde{B}_{\mathcal{J}}$ is good conditioned.

Table 1: Computational complexity of Algorithm 1. See the next paragraph “**Computational complexity.**” for the explanation of “.”.

LINE NO.	OPERATION	COMPLEXITY
5	MVP	-
5	ORTHOGONALIZATION	$\mathcal{O}(m^2p)$
6	EIGENVALUE PROB.	$\mathcal{O}(m^3)$
8	SORTING	$\mathcal{O}(p \log p)$
11, 17	EIGENVALUE PROB.	$\mathcal{O}(s^3)$

Computational complexity. In Table 1, we list the computational complexity of the major steps of the IFTRR method. The symbol “.” means the computational complexity is different case by case: (1) when \tilde{A} and \tilde{B} are available, the operation $(\tilde{A} - \rho^{(t)}\tilde{B})v$ requires $\mathcal{O}(mp^2)$ FLOPS if \tilde{A} and \tilde{B} are dense, $\mathcal{O}(m(\text{nnz}(\tilde{A}) + \text{nnz}(\tilde{B})))$ FLOPS if \tilde{A} and \tilde{B} are sparse; (2) when \tilde{A} and \tilde{B} are unavailable directly, the operations $\tilde{A}v$ and $\tilde{B}v$ are carried out via a data matrix $X \in \mathbb{R}^{n \times p}$, then the operation $(\tilde{A} - \rho^{(t)}\tilde{B})v$ in general requires $\mathcal{O}(mnp)$ FLOPS if \tilde{A} and \tilde{B} are dense, $\mathcal{O}(m \text{nnz}(X))$ FLOPS if X are sparse.¹ Since $s \ll p$ and $m \ll p$, the overall complexity of IFTRR is dominated by line 5.

¹For example, consider $X \in \mathbb{R}^{n \times p}$, where n is the number of observations and p is the number of features. The sample

2.4 Convergence

Before the study of the convergence, we give some definitions and preliminary lemmas.

The *angle* between two vectors $x, y \in \mathbb{R}^p$ is defined as $\theta(x, y) \triangleq \arccos \frac{|x^T y|}{\|x\| \|y\|}$. Define the *Crawford number* for a definite-symmetric matrix pair (A, B) as

$$c(A, B) \triangleq \min_{\|x\|_2=1} \sqrt{(x^T A x)^2 + (x^T B x)^2}.$$

The following lemma tells that when (A, B) is slightly perturbed, the changes of eigenvalues are small.

Lemma 1 (Van Loan and Golub, 2012, Theorem 8.7.3) *Suppose (A, B) is a symmetric-definite pair with eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p$, E and F are symmetric p -by- p matrices that satisfy*

$$\epsilon = \sqrt{\|E\|_2^2 + \|F\|_2^2} < c(A, B).$$

Then $(A + E, B + F)$ is also a symmetric-definite pair with eigenvalues $\tilde{\lambda}_1 \geq \tilde{\lambda}_2 \geq \dots \geq \tilde{\lambda}_p$ that satisfy

$$|\arctan(\lambda_i) - \arctan(\tilde{\lambda}_i)| \leq \arctan(\epsilon/c(A, B)),$$

for $i = 1, 2, \dots, p$.

Let $(\lambda^{(t)}, u^{(t)})$ be the current guess of the largest eigenpair of (A, B) and $Q_t \in \mathbb{R}^{p \times m}$ be an orthonormal basis for $\mathcal{K}_m(A - \lambda^{(t)}B, u^{(t)})$. Then $(\lambda^{(t+1)}, u^{(t+1)})$ can be obtained via the Rayleigh-Ritz procedure. Specifically, let (ω, y) be the largest eigenpair of $(Q_t^T A Q_t, Q_t^T B Q_t)$, then $\lambda^{(t+1)} = \omega$, $u^{t+1} = Q_t y$. By (Golub and Ye, 2002, Theorem 3.4), we have the following lemma.

Lemma 2 *Let the eigenvalues of $A - \lambda^{(t)}B$ be $\sigma_p \leq \dots \leq \sigma_2 < \sigma_1$. Assume $\lambda_2 < \lambda^{(t)} < \lambda_1$. Then*

$$\lambda_1 - \lambda^{(t+1)} \leq (\lambda_1 - \lambda^{(t)})\epsilon_m^2 + \mathcal{O}((\lambda_1 - \rho^{(t)})^{\frac{3}{2}}),$$

where $\epsilon_m = \min_{p \in \mathcal{P}_m, p(\sigma_1)=1} \max_{i \neq 1} |p(\sigma_i)| \leq 2 \left(\frac{1 - \sqrt{\psi}}{1 + \sqrt{\psi}} \right)^m$ with \mathcal{P}_m denoting the set of all polynomials of degree not greater than m , $\psi = \frac{\sigma_1 - \sigma_2}{\sigma_1 - \sigma_q}$.

Let $\rho^{(t)}, v^{(t)}$ be obtained via Algorithm 1. Denote $\mathcal{J}_t = \text{supp}(v^{(t)})$. Here, we consider an alternative way to update $\rho^{(t)}, v^{(t)}$:

(i) Compute an orthonormal basis for $\mathcal{K}_m(\tilde{A}_{\mathcal{J}_{t+1}} - \rho^{(t)}\tilde{B}_{\mathcal{J}_{t+1}}, v_{\mathcal{J}_{t+1}}^{(t)})$, denote it as $Q_{\mathcal{J}_{t+1}}$;

correlation matrix is $C = \frac{1}{n}(X - \mathbf{1}_n \bar{x}^T)^T (X - \mathbf{1}_n \bar{x}^T)$, where $\mathbf{1}_n = [1, \dots, 1]^T \in \mathbb{R}^n$, $\bar{x} = [\bar{x}_1, \dots, \bar{x}_p]^T$, \bar{x}_i is the mean of the i -th column of X . Then MVP $u = Cv$ is computed as $w = Xv - (\bar{x}^T v)\mathbf{1}_n$, $u = X^T w - (\mathbf{1}_n^T w)\bar{x}$, which requires $\mathcal{O}(np)$ if X is dense, and $\mathcal{O}(\text{nnz}(X))$ if X is sparse.

- (ii) Solve the largest eigenpair of $(Q_{\mathcal{J}_{t+1}}^T \tilde{A}_{\mathcal{J}_{t+1}} Q_{\mathcal{J}_{t+1}}, Q_{\mathcal{J}_{t+1}}^T \tilde{B}_{\mathcal{J}_{t+1}} Q_{\mathcal{J}_{t+1}})$, and denote it by (ω, y) ;
- (iii) Set $\hat{\rho}^{(t+1)} = \omega$, $\hat{v}_{\mathcal{J}_{t+1}}^{(t+1)} = Q_{\mathcal{J}_{t+1}} y$, $\hat{v}_{\mathcal{J}_{t+1}^c}^{(t+1)} = 0$.

Note that the above Rayleigh-Ritz procedure is only for the purpose of analyzing the convergence, is not applicable in practice since \mathcal{J}_{t+1} is unknown. Also note that the existence of $Q_{\mathcal{J}_{t+1}}$ implicitly requires that $m \leq |\mathcal{J}_{t+1}|$, meaning that the dimension of the Krylov subspace should not exceed sparsity level of the approximate eigenvector.

By Lemma 2, we have the following result.

Lemma 3 Let $\mathcal{J}_t = \text{supp}(v^{(t)})$, $\ell = |\mathcal{J}_{t+1}|$, the eigenvalues of $(\tilde{A}_{\mathcal{J}_{t+1}}, \tilde{B}_{\mathcal{J}_{t+1}})$ be $\lambda_{1,t+1} \geq \dots \geq \lambda_{\ell,t+1}$, the eigenvalues of $\tilde{A}_{\mathcal{J}_{t+1}} - \rho^{(t)} \tilde{B}_{\mathcal{J}_{t+1}}$ be $\sigma_\ell \leq \dots \leq \sigma_2 < \sigma_1$. Assume $\lambda_{2,t+1} < \rho^{(t)} < \lambda_{1,t+1}$. Then

$$\lambda_{1,t+1} - \hat{\rho}^{(t+1)} \leq (\lambda_{1,t+1} - \rho^{(t)}) \epsilon_m^2 + \mathcal{O}((\lambda_{1,t+1} - \rho^{(t)})^{\frac{3}{2}}),$$

where ϵ_m is the same as in Lemma 2.

Recall that $\rho^{(t+1)}$ is the largest eigenvalue of $(\tilde{A}_{\mathcal{J}}, \tilde{B}_{\mathcal{J}})$, $\hat{\rho}^{(t+1)}$ is the largest eigenvalue of $(Q_{\mathcal{J}_{t+1}}^T \tilde{A}_{\mathcal{J}_{t+1}} Q_{\mathcal{J}_{t+1}}, Q_{\mathcal{J}_{t+1}}^T \tilde{B}_{\mathcal{J}_{t+1}} Q_{\mathcal{J}_{t+1}})$, and $\mathcal{J}_{t+1} \subset \mathcal{J}$. Then we have

Lemma 4 It holds that $\rho^{(t+1)} \geq \hat{\rho}^{(t+1)}$.

Define

$$\eta_s^{(2)} \triangleq \max_{|\mathcal{J}| \leq s} \lambda_2(\tilde{A}_{\mathcal{J}}, \tilde{B}_{\mathcal{J}}), \quad (4a)$$

$$\eta_{s,\ell}^{(1)} \triangleq \max_{\substack{|\mathcal{J}| \leq s \\ |\mathcal{J} \cap \text{supp}(v_1)| \leq \ell}} \lambda_1(\tilde{A}_{\mathcal{J}}, \tilde{B}_{\mathcal{J}}). \quad (4b)$$

Combining Lemmas 3 and 4, we have

Theorem 1 Let $\mathcal{J}_t = \text{supp}(v^{(t)})$, $s = \sup_t |\mathcal{J}_t|$ and $k = |\text{supp}(v_1)|$. For any $\mathcal{J} \subset [p]$ with $|\mathcal{J}| = s$, denote the eigenvalues of $\tilde{A}_{\mathcal{J}} - \rho^{(t)} \tilde{B}_{\mathcal{J}}$ by $\sigma_{1,\mathcal{J}} > \sigma_{2,\mathcal{J}} \geq \dots \geq \lambda_{s,\mathcal{J}}$,

$$\psi_* = \min_{|\mathcal{J}|=s} \frac{\sigma_{1,\mathcal{J}} - \sigma_{2,\mathcal{J}}}{\sigma_{1,\mathcal{J}} - \sigma_{s,\mathcal{J}}}, \quad \epsilon_* = 2 \left(\frac{1 - \sqrt{\psi_*}}{1 + \sqrt{\psi_*}} \right)^m.$$

If $\eta_{s,k-1}^{(1)} \geq \rho^{(t)} > \eta_s^{(2)}$ and $|\mathcal{J}_t \cap \text{supp}(v_1)| < k$, then there exists a $\lambda_{1,t+1} \in (\rho^{(t)}, +\infty)$ such that

$$\lambda_{1,t+1} - \rho^{(t+1)} \leq (\lambda_{1,t+1} - \rho^{(t)}) \epsilon_*^2 + \mathcal{O}((\lambda_{1,t+1} - \rho^{(t)})^{\frac{3}{2}}).$$

Asymptotically,

$$\rho^{(t+1)} \gtrsim \rho^{(t)} + (\lambda_{1,t+1} - \rho^{(t)})(1 - \epsilon_*^2).$$

Remark 1 Let $s > k = |\text{supp}(v_1)|$. Assuming that for any $|\mathcal{J}| \leq s$, $\|\tilde{A}_{\mathcal{J}} - A_{\mathcal{J}}\|^2 + \|\tilde{B}_{\mathcal{J}} - B_{\mathcal{J}}\|^2$ is small, by Lemma 1, we know that $\lambda_i(\tilde{A}_{\mathcal{J}}, \tilde{B}_{\mathcal{J}}) \approx \lambda_i(A_{\mathcal{J}}, B_{\mathcal{J}})$, for all $i = 1, 2, \dots, s$. By interlacing property (e.g., (Van Loan and Golub, 2012, Theorem 8.1.7)), $\lambda_2(A_{\mathcal{J}}, B_{\mathcal{J}}) \leq \lambda_2$. Therefore, we have

$$\eta_s^{(2)} \lesssim \lambda_2 < \lambda_1 \approx \eta_{s,k}^{(1)},$$

i.e., the gap between $\eta_s^{(2)}$ and $\eta_{s,k}^{(1)}$ is larger than that between λ_2 and λ_1 . In fact, in practice, the former is much larger than the latter.

Remark 2 Let $k = |\text{supp}(v_1)|$. Intuitively, we also expect a gap between $\eta_{s,k-1}^{(1)}$ and $\eta_{s,k}^{(1)}$, the larger the gap is, the easier the problem is. Otherwise, when the gap is sufficiently small, the problem has two “solutions”: one is approximately v_1 , the other is v' such that $\tilde{\rho}(v') = \eta_{s,k-1}^{(1)}$; they are both sparse, and $\tilde{\rho}(v_1) \approx \tilde{\rho}(v')$. In addition, notice that $\|v'\|_0 = k-1 < k = \|v_1\|_0$, consequently, we can not expect to find good approximation of the true solution v_1 , since v' is a “better” solution, in the sense that it is sparser and $\tilde{\rho}(v_1) \approx \tilde{\rho}(v')$.

Theorem 1 tells that as long as $\text{supp}(v_1)$ is not a subset of \mathcal{J}_t and $\rho^{(t)} > \eta_s^{(2)}$, $\{\rho^{(t)}\}_t$ is asymptotically nondecreasing. Theorem 2 below tells that once $\rho^{(t)}$ is larger than $\eta_{s,k-1}^{(1)}$, $\text{supp}(v_1)$ is a subset of \mathcal{J}_t .

Theorem 2 If $\rho^{(t)} > \eta_{s,k-1}^{(1)}$ with $k = |\text{supp}(v_1)|$, then $\text{supp}(v_1) \subset \mathcal{J}_t$.

The following theorem tells that when $\text{supp}(v_1) \subset \text{supp}(v^{(t)})$, then $\rho^{(t)}$ is close to λ_1 and $\theta(v^{(t)}, v_1)$ is small, i.e., they are approximations of λ_1 and v_1 .

Theorem 3 Let $\mathcal{J}_t = \text{supp}(v^{(t)})$. Denote $c_{\mathcal{J}} = c(A_{\mathcal{J}_t}, B_{\mathcal{J}_t})$, $E_{\mathcal{J}} = \tilde{A}_{\mathcal{J}_t} - A_{\mathcal{J}_t}$, $F_{\mathcal{J}} = \tilde{B}_{\mathcal{J}_t} - B_{\mathcal{J}_t}$, and $\epsilon_{\mathcal{J}} = \sqrt{\|E_{\mathcal{J}}\|_2^2 + \|F_{\mathcal{J}}\|_2^2}$. Assume $\text{supp}(v_1) \subset \mathcal{J}_t$.

(a) If $\epsilon_{\mathcal{J}} < c_{\mathcal{J}}$, then

$$|\arctan(\rho^{(t)}) - \arctan(\lambda_1)| \leq \arctan(\epsilon_{\mathcal{J}}/c_{\mathcal{J}}).$$

(b) Furthermore, if $|\rho^{(t)}| \epsilon_{\mathcal{J}} < c_{\mathcal{J}}$ and $\rho^{(t)}$ is simple, then

$$\sin \theta(v^{(t)}, v_1) \leq \frac{\|\tilde{B}_{\mathcal{J}_t}\|_2 \delta + \sqrt{1 + \mu^2} \epsilon_{\mathcal{J}}}{g} = \mathcal{O}(\epsilon_{\mathcal{J}}),$$

where g is the smallest nonzero singular value of $\tilde{A}_{\mathcal{J}_t} - \rho^{(t)} \tilde{B}_{\mathcal{J}_t}$, $\delta = \frac{(1 + (\rho^{(t)})^2) \epsilon_{\mathcal{J}}}{c_{\mathcal{J}} - |\rho^{(t)}| \epsilon_{\mathcal{J}}}$.

3 Numerical Experiments

To illustrate the behavior of the IFTRR method and compare it with existing methods for the SGEP, this section presents some numerical experiments.

3.1 Sparse Canonical Correlation Analysis

Recall Example 2. In our simulations, we set $\Sigma_{xx} = \Sigma_{yy} = I_5 \otimes D$, $D = [d_{jl}] \in \mathbb{R}^{p/10 \times p/10}$ is a Toeplitz matrix with $d_{jl} = 0.8^{|j-l|}$. Let v_x^* be collinear with $\sum_{j=1}^{s/2} e_{5j-4}$ and $(v_x^*)^\top \Sigma_{xx} v_x^* = 1$, where s is a small even integer. Set $v_y^* = v_x^*$, $\Sigma_{xy} = 0.9 \times \Sigma_{xx} v_x^* (v_y^*)^\top \Sigma_{yy}$ (low rank case), or $\Sigma_{xy} = 0.9 \times \Sigma_{xx} v_x^* (v_y^*)^\top \Sigma_{yy} + 0.1 \times \Sigma_{xx} V_x^* (V_y^*)^\top \Sigma_{yy}$ (approximate low rank case), where V_x^*, V_y^* are random matrices such that $(V_x^*)^\top \Sigma_{xx} V_x^* = I_{p/2}$, $(V_y^*)^\top \Sigma_{yy} V_y^* = I_{p/2}$.

We perform the IFTRR method for 200 times under the following settings:

1. Low rank Σ_{xy} , $p = 1000$, $s = 6$, different numbers of samples $n = 100, 200, 300, 400$;
2. Low rank Σ_{xy} , $p = 1000$, $n = 400$, different numbers of sparsity levels $s = 6, 10, 14, 18$;
3. Same as setting 1 except that Σ_{xy} is approximate low rank;
4. Same as setting 2 except that Σ_{xy} is approximate low rank.
5. Low rank Σ_{xy} , $s = 6$, $p = 5000$, different numbers of samples $n = 2000, 4000, \dots, 10000$;
6. Low rank Σ_{xy} , $s = 6$, $n = 4000$, different numbers of features $p = 2000, 4000, \dots, 10000$.

The performance of the method is evaluated in terms of the angle between $v_1 = \begin{bmatrix} v_x^* \\ v_y^* \end{bmatrix}$ and the computed \tilde{v} , and also success rate – we say the returned \tilde{v} is a success if $\text{supp}(v_1) = \text{supp}(\tilde{v})$. The results are reported in Figures 2 and 3. We can see from these figures that (i) for reasonable large n , \tilde{v} (returned by the IFTRR method) is a good approximation of v_1 , the larger n is, the smaller the angle is, and the larger the success rate is; (ii) for different sparsity levels s , \tilde{v} is also a good approximation of v_1 , the smaller s is, the smaller the angle is, and the larger the success rate is; (iii) the results for the low rank case are better than that for the approximate low rank case. The above numerical results indicate that the IFTRR method gives a better result when the number of samples is sufficiently large and the leading eigenvector v_1 is sufficiently sparse.

In Figure 4, we give the boxplots of the CPU time for settings 5 and 6. We can see that with a fixed number of features, the CPU time increases almost linearly with respect to the number of samples; with a fixed number of samples, the CPU time increases almost linearly with respect to the number of features. This confirms that the computational cost of the IFTRR method is dominated by MVP, which is $\mathcal{O}(np)$.

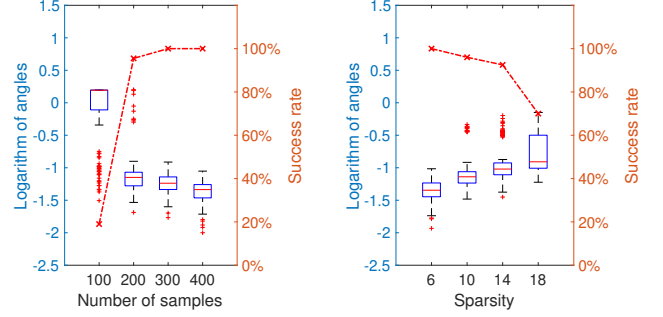


Figure 2: Accuracy and success rate, low rank case, from left to right, settings 1 to 2.

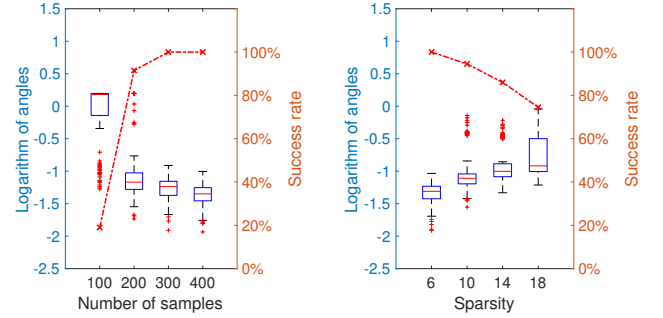


Figure 3: Accuracy and success rate, approximate low rank case, from left to right, settings 3 to 4.




Figure 4: CPU time, from left to right, settings 5 to 6.

3.2 Sparse Fisher's Discriminant Analysis

Recall Example 1. In our simulations, for $k = 1, 2, \dots, K$, we set $\bar{x}_{kj} = \frac{2k-2}{K+2}$ for $j = 2, 4, \dots, 40$, $\bar{x}_{kj} = 0$ otherwise. The data matrix X is generated via $X_{(:,i)} \sim N(\bar{x}_k, \Sigma)$ for $i \in \mathcal{C}_k$, where $\Sigma = I_5 \otimes D$, $D = [d_{jl}] \in \mathbb{R}^{p/5 \times p/5}$ is a Toeplitz matrix with $d_{jl} = 0.8^{|j-l|}$.

Table 2: Results of Misclassification error and number of selected features.

	K	GLMNET	D(M)SDA	RIFLE	IFTRR
ERR.	2	32	29	15	14(4)
FEAT.		88	105	42	42(1)
ERR.	4	495	247	192	103(11)
FEAT.		54	102	42	42(1)

Fix $p = 500$, for $K = 2, 4$, using 400 training samples, we perform GLMNET (Friedman et al., 2010),

DSDA/MSDA (Mai et al., 2015, 2012), RIFLE (Tan et al., 2018) and the IFTRR method. The results are then used to classify 1000 test samples, the misclassification error (denoted by ERR.) and number of selected features (denoted by FEAT.) are recorded. ERR. and FEAT. are averaged over 200 independently generated datasets and reported in Table 2. There the number in the brackets is the corresponding standard error, and all numbers are rounded to the nearest integers. Table 2 shows that for $K = 2$, the misclassification errors of RIFLE and the IFTRR method are comparable and lower than the other two methods; and for $K = 4$, the IFTRR method has the lowest misclassification error.

3.3 Sparse Sliced Inverse Regression

Consider the sparse sliced inverse regression for the model $Y = f(v_1^T X, \dots, v_k^T X, \epsilon)$, where Y is a univariate response, X is d -dimensional covariates, ϵ is the stochastic error independent of X , f is the link function, which is unknown. Under regularity conditions, the subspace spanned by v_1, \dots, v_k can be identified via solving an SGEP with $\tilde{A} = \hat{\Sigma}_{\mathbb{E}(X|Y)}$, $\tilde{B} = \hat{\Sigma}_x$, where $\hat{\Sigma}_x$ is the sample covariance matrix of X , $\hat{\Sigma}_{\mathbb{E}(X|Y)}$ is the sample covariance matrix of the conditional expectation $\mathbb{E}(X|Y)$, which is $\hat{\Sigma}_{\mathbb{E}(X|Y)} = \hat{\Sigma}_x - \frac{1}{n_1+n_2} \sum_{k=1}^2 n_k \hat{\Sigma}_{x,k}$, where n_k is the number of samples for class k , $\hat{\Sigma}_{x,k}$ is the sample covariance matrix for class k , for $k = 1, 2$. See (Chen et al., 2010; Li, 1991, 2007; Tan et al., 2018) and reference therein for more details.

Let $v^{(t)}$ be an approximate solution to SGEP. $Xv^{(t)}$ is usually used as a predictor. Here we may also use the indices for the nonzero entries of $v^{(t)}$ to select the features, and the features can be ranked by ordering the absolute values of the nonzero entries of $v^{(t)}$. Now we compare our method with feature selection methods – relief, mutinf, fsv and fisher, which are all available in the Feature Selection Library (Roffo, 2017; Roffo and Melzi, 2016; Roffo et al., 2017, 2015). The datasets are all downloaded from scikit-feature feature selection repository. Each dataset is randomly partitioned into a training set and a test set, and the test set size is approximately $0.2n$. A support vector machine (SVM) classifier is trained using the training set with only the selected features and then used to predict on the test set. The average accuracy (averaged over 10 independent runs) is plotted in Figure 5. We can see that our method is comparable with the other four methods for the first 5 datasets, outperforms the other four methods for the last dataset.

4 Conclusion

We have proposed the IFTRR method to solve the SGEP. The IFTRR method has the following advantages: Since only the MVP is required, the method is

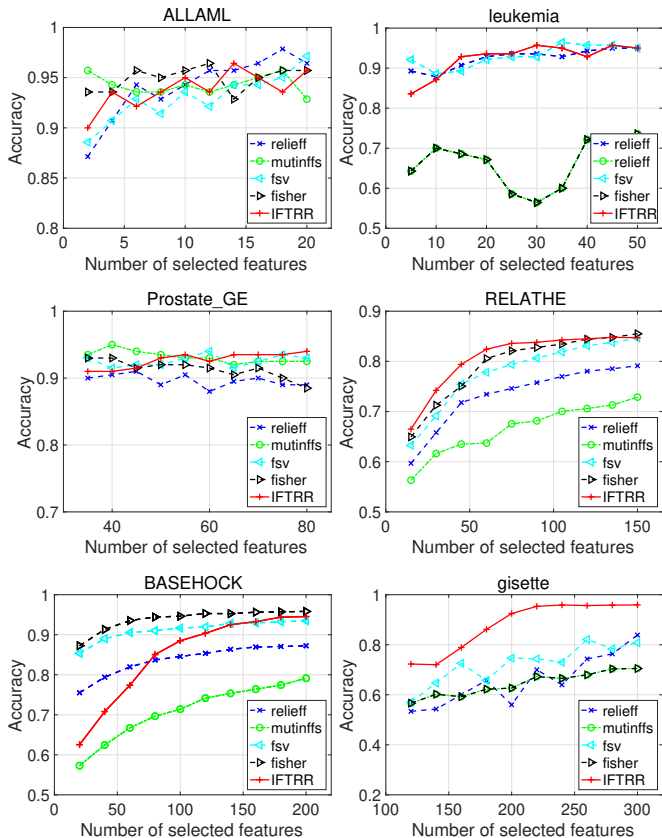


Figure 5: Accuracy vs. number of features.

suitable for large scale problem; A cure is incorporated into the IFTRR method, which makes it applicable for ill-conditioned or singular coefficient matrices \tilde{A} , \tilde{B} ; Based on “eigenvalue increment”, a new truncation procedure is proposed, which is able to find the support set of the leading eigenvector effectively, as a result, the IFTRR method usually converges in a few iterations. Numerical simulations show that the IFTRR method is effective and efficient, especially when the matrix size is large and the leading eigenvector is very sparse.

There are several future research topics for the IFTRR method. First, can we extend the IFTRR method to compute multiple leading sparse eigenvectors? Computing several leading eigenvectors one by one seems simple, how to compute them simultaneously is uneasy since the orthogonalization procedure, which is required for computing several eigenvectors simultaneously, usually destroys the sparsity. Second, how to truncate the eigenvector to ensure some structured sparsity, say group sparsity as in the group LASSO. Third, based on the IFTRR method, a general framework for solving SGEP can be obtained – `eigensolver + truncation`. Various eigensolvers together with certain truncation procedure can be tried to solve the SGEP. Which is the best choice, how is the convergence? More studies towards these directions are apparently required.

References

- Zhaojun Bai, James Demmel, Jack Dongarra, Axel Ruhe, and Henk van der Vorst. *Templates for the solution of algebraic eigenvalue problems: a practical guide*. SIAM, 2000.
- Jorge Cadima and Ian T Jolliffe. Loading and correlations in the interpretation of principle components. *J. Appl. Stat.*, 22(2):203–214, 1995.
- Xin Chen, Changliang Zou, R Dennis Cook, et al. Coordinate-independent sparse sufficient dimension reduction and variable selection. *The Annals of Statistics*, 38(6):3696–3723, 2010.
- Line Clemmensen, Trevor Hastie, Daniela Witten, and Bjarne Ersbøll. Sparse discriminant analysis. *Technometrics*, 53(4):406–413, 2011.
- A. d’Aspremont, L. El Ghaoui, M. Jordan, and G. Lanckriet. A direct formulation for sparse PCA using semidefinite programming. *SIAM Rev.*, 49(3):434–448, 2007.
- Ernest R. Davidson. The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices. *J. Comput. Phys.*, 17:87–94, 1975.
- James W Demmel. *Applied Numerical Linear Algebra*. SIAM, Philadelphia, PA, 1997.
- Alexandre d’Aspremont, Francis Bach, and Laurent El Ghaoui. Optimal solutions for sparse principal component analysis. *J. Machine Learning Res.*, 9(Jul):1269–1294, 2008.
- Jerome Friedman, Trevor Hastie, and Rob Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1):1, 2010.
- Gene H Golub and Qiang Ye. An inverse free preconditioned krylov subspace method for symmetric generalized eigenvalue problems. *SIAM J. Sci. Comput.*, 24(1):312–334, 2002.
- Ian T Jolliffe, Nickolay T Trendafilov, and Mudassir Uddin. A modified principal component technique based on the LASSO. *J. Comput. Graphical Statist.*, 12(3):531–547, 2003.
- Michel Journée, Yurii Nesterov, Peter Richtárik, and Rodolphe Sepulchre. Generalized power method for sparse principal component analysis. *J. Machine Learning Res.*, 11(Feb):517–553, 2010.
- Ker-Chau Li. Sliced inverse regression for dimension reduction. *J. Am. Stat. Assoc.*, 86(414):316–327, 1991.
- Lexin Li. Sparse sufficient dimension reduction. *Biometrika*, 94(3):603–613, 2007.
- Ronny Luss and Marc Teboulle. Conditional gradient algorithms for rank-one matrix approximations with a sparsity constraint. *SIAM Rev.*, 55(1):65–98, 2013.
- Qing Mai, Hui Zou, and Ming Yuan. A direct approach to sparse discriminant analysis in ultra-high dimensions. *Biometrika*, 99(1):29–42, 2012.
- Qing Mai, Yi Yang, and Hui Zou. Multiclass sparse discriminant analysis. *arXiv:1504.05845*, 2015.
- Baback Moghaddam, Yair Weiss, and Shai Avidan. Spectral bounds for sparse PCA: exact and greedy algorithms. In *Advances in Neural Information Processing Systems (NIPS)*, pages 915–922, Vancouver, Canada, 2005.
- Baback Moghaddam, Yair Weiss, and Shai Avidan. Generalized spectral bounds for sparse LDA. In *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML)*, pages 641–648, Pittsburgh, PA, 2006.
- Giorgio Roffo. Ranking to learn and learning to rank: On the role of ranking in pattern recognition applications. *arXiv:1706.05933*, 2017.
- Giorgio Roffo and Simone Melzi. Ranking to learn. In *International Workshop on New Frontiers in Mining Complex Patterns*, pages 19–35. Springer, 2016.
- Giorgio Roffo, Simone Melzi, and Marco Cristani. Infinite feature selection. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 4202–4210, Santiago, Chile, 2015.
- Giorgio Roffo, Simone Melzi, Umberto Castellani, and Alessandro Vinciarelli. Infinite latent feature selection: A probabilistic latent graph-based ranking approach. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1407–1415, Venice, Italy, 2017.
- Sandra E Safo, Jeongyoun Ahn, Yongho Jeon, and Sungkyu Jung. Sparse generalized eigenvalue problem with application to canonical correlation analysis for integrative analysis of methylation and gene expression data. *Biometrics*, 2018.
- Gerard LG Sleijpen and Henk A Van der Vorst. A Jacobi–Davidson iteration method for linear eigenvalue problems. *SIAM Rev.*, 42(2):267–293, 2000.
- Junxiao Song, Prabhu Babu, and Daniel P Palomar. Sparse generalized eigenvalue problem via smooth optimization. *IEEE Trans. Signal Process.*, 63(7):1627–1642, 2015.
- Bharath K Sriperumbudur, David A Torres, and Gert RG Lanckriet. A majorization-minimization approach to the sparse generalized eigenvalue problem. *Mach. Learn.*, 85(1-2):3–39, 2011.
- Gilbert W Stewart. *Matrix algorithms volume 2: eigen-systems*, volume 2. SIAM, 2001.

Kean Ming Tan, Zhaoran Wang, Han Liu, and Tong Zhang. Sparse generalized eigenvalue problem: Optimal statistical rates via truncated Rayleigh flow. *J. R. Statist. Soc. B*, 80(5):1057–1086, 2018.

Charles F Van Loan and Gene H Golub. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, 4th edition, 2012.

Daniela M Witten, Robert Tibshirani, and Trevor Hastie. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 10(3):515–534, 2009.

Xiao-Tong Yuan and Tong Zhang. Truncated power method for sparse eigenvalue problems. *J. Machine Learning Res.*, 14(Apr):899–925, 2013.

Hui Zou, Trevor Hastie, and Robert Tibshirani. Sparse principal component analysis. *J. Comput. Graphical Statist.*, 15(2):265–286, 2006.