

# Iteratively Reweighted Least Squares Algorithms for L1-Norm Principal Component Analysis

Young Woong Park  
Cox School of Business  
Southern Methodist University  
Dallas, Texas 75225  
Email: ywpark@smu.edu

Diego Klabjan  
Department of Industrial Engineering and  
and Management Sciences  
Northwestern University  
Evanston, Illinois 60208  
Email: d-klabjan@northwestern.edu

**Abstract**—Principal component analysis (PCA) is often used to reduce the dimension of data by selecting a few orthonormal vectors that explain most of the variance structure of the data. L1 PCA uses the L1 norm to measure error, whereas the conventional PCA uses the L2 norm. For the L1 PCA problem minimizing the fitting error of the reconstructed data, we propose an exact reweighted and an approximate algorithm based on iteratively reweighted least squares. We provide convergence analyses, and compare their performance against benchmark algorithms in the literature. The computational experiment shows that the proposed algorithms consistently perform best.

## I. INTRODUCTION

Principal component analysis (PCA) is a technique to find orthonormal vectors, which are a linear combination of the attributes of the data, that explain the variance structure of the data [12]. Since a few orthonormal vectors usually explain most of the variance, PCA is often used to reduce dimension of the data by keeping only a few of the orthonormal vectors. These orthonormal vectors are called *principal components* (PCs).

For dimensionality reduction, we are given target dimension  $p$ , the number of PCs. To measure accuracy, given  $p$  principal components, first, the original data is projected into the lower dimension using the PCs. Next, the projected data in the lower dimension is lifted to the original dimension using the PCs. Observe that this procedure causes loss of some information if  $p$  is smaller than the dimension of the original attribute space. The reconstruction error is defined by the difference between the projected-and-lifted data and the original data. To select the best  $p$  PCs, the following two objective functions are usually used:

- [P1] minimization of the reconstruction error,
- [P2] maximization of the variance of the projected data.

The traditional measure to capture the errors and variances in P1 and P2 is the  $L_2$  norm. For each observation, we have the vector of the reconstruction error and variance for P1 and P2, respectively. Then, the squared  $L_2$  norm of the vectors are added over all observations to define the total reconstruction error and variance for P1 and P2, respectively. In fact, in terms of a matrix norm, we optimize the Frobenius norm of the reconstruction error and projected data matrices for P1 and P2, respectively. With the  $L_2$  norm as the objective

function, P1 and P2 are actually equivalent. Further, P2 can be efficiently solved by singular value decomposition (SVD) of the data matrix or the eigenvalue decomposition (EVD) of the covariance matrix of the data. However, the  $L_2$  norm is usually sensitive to outliers. As an alternative, PCA based on  $L_1$  norm has been used to find more robust PCs.

For P1, instead of the  $L_2$  norm, we minimize the sum of the  $L_1$  norm of the reconstruction error vectors over all observations. A few heuristics have been proposed for this minimization problem. The heuristic proposed in [1] is based on a canonical correlation analysis. The iterative algorithm in [14] assumes that the projected-and-lifted data is a product of two matrices and is then iteratively optimizing by fixing one of the two matrices. The algorithm in [3] sequentially reduces the dimension by one. The algorithm is based on the observation that the projection from  $k$  to the best fit  $k - 1$  dimension can be found by solving several linear program (LP) problems for least absolute deviation regression. The algorithms in [14] and [3] actually try to find the best fitting subspace, where in the objective function the original data is approximated by the multiplication of two matrices, PC and score matrices. This approximation is not the same as the reconstructed matrix by PCs, while the ultimate goal is still minimizing the reconstruction errors.

The  $L_1$  norm for P2 has also been studied. This problem is often called the *projection pursuit  $L_1$ -PCA*. In this context, we maximize the sum of the  $L_1$  norm of the projected observation vectors over all observations. However, in contrast to the conventional  $L_2$  norm based PCA, the solutions of P1 and P2 with the  $L_1$  norm might not be same. The work in [9] studies  $L_1$ -norm based covariance matrix estimation, while the works in [6], [7], [15], [17] and [20] directly consider P2. The algorithm in [15] finds a local optimal solution by sequentially obtaining one PC that is orthogonal to the previously obtained PCs based on a greedy strategy. Recently, the work in [20] extended the algorithm in [15] using a non-greedy strategy. The works in [18] and [19] show that P2 with one PC is NP-hard when the number of observations and attributes are jointly arbitrarily large. The work in [18] provides a polynomial algorithm when the number of attributes is fixed.

As the objective functions are different, solving for P1 and P2 with different norms give different solutions in terms of the

signs and order of the PCs. In Figure 1, we present heat maps of PCs obtained by solving  $L_2$ -PCA, P1 with the  $L_1$  norm, and P2 with the  $L_1$  norm with  $p = 5$  for data set *cancer\_2* presented in Table I and used in the experiment in Section III-C. The three heat maps represent the matrices of PCs of  $L_2$ -PCA (left), P1 with the  $L_1$  norm (center), and P2 with the  $L_1$  norm (right). The rows and columns of each matrix represent original attributes and PCs, respectively. The blue, white, and red cells represent the intensity of positive, zero, and negative values. Note that both  $L_1$ -PCA variations give Attr 1 a large negative loading in either of the first or second PCs, whereas  $L_2$ -PCA gives Attr 1 a large positive loading in the third PC. Attr 9 has a large negative loading for the fifth PC of P1 with  $L_1$  norm, while  $L_2$ -PCA gives Attr 9 a large positive loading in the second PC.

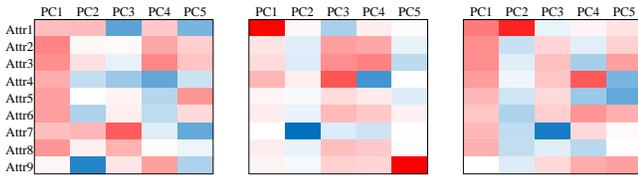


Fig. 1: Heat maps of 5 principal components by  $L_2$ -PCA (left), P1 with  $L_1$  norm (center), and P2 with  $L_1$  norm (right) for *cancer\_2* data set

In this paper, we propose two iterative algorithms for P1 with the  $L_1$  norm, and provide analytical convergence results. Although we propose iterative algorithms, our focus is the  $L_1$  objective function and thus our algorithms do not directly compare with iterative algorithms for the standard  $L_2$ -PCA such as algorithms based on EM [21] or NIPALS [10]. For P1 with the  $L_1$  norm, the first proposed algorithm, the exact reweighted version, is based on iteratively reweighted least squares (IRLS) that gives a weight to each observation. The weighted least square problem is then converted into the standard  $L_2$ -PCA problem with a weighted data matrix, and the algorithm iterates over different weights. We show that the algorithm gives convergent weights and the weighted data matrices have convergent eigenvalues. We also propose an approximate version in order to speed up the algorithm and show the convergent eigenvectors. The work in [8] provides an IRLS algorithm for sparse recovery. The reader is referred to [13] for a review of IRLS applied in different contexts. Recently, the work in [24] studied the minimization of a differentiable function of an orthogonal matrix. In the computational experiment, we compare our algorithms with benchmark algorithms in [3], [14], [15], [20]. The experiment shows that our algorithm for P1 regularly outperforms the benchmark algorithms with respect to the solution quality and its computational time is of the same order as the fastest among the other four. Even though  $L_1$ -PCA can be used for building robust PCs and is an alternative to other robust PCA approaches such as the work in [5], we limit the comparison for the  $L_1$ -PCA objective functions introduced in Section II,

as our goal is to directly optimize the  $L_1$  norms for P1.

Our contributions can be summarized as follows.

- 1) For P1 with the  $L_1$  norm, we propose the first IRLS algorithm in the literature, and show the convergence of the eigenvalues of the weighted matrices. The algorithm directly minimizes the reconstruction error, while the other benchmark algorithms primarily try to find the optimal subspace. An additional advantage of our algorithm is that it uses an  $L_2$ -PCA algorithm as a black box. Hence, by using a more scalable algorithm, the practical time complexity of our algorithm can further be reduced.
- 2) We propose an approximate version to speed up the algorithm and to guarantee convergent eigenvectors. The difference is that the approximate version uses a formula to update the eigenpairs when the changes in the weighted data matrices become relatively small.
- 3) The results of the computational experiment show that the proposed algorithms for P1 outperform the benchmark algorithms in most cases.

The rest of the paper is organized as follows. In Section II, we present the algorithms and show all analytic results. Then in Section III, the computational experiment is presented.

## II. ALGORITHMS FOR $L_1$ PCA

In this section, we present the algorithms for P1 and show the underlying analytic results. We use the following notation.

- $n$ : number of observations of the data
- $m$ : number of attributes of the data
- $p$ : number of principal components (target dimension)
- $I = \{1, \dots, n\}$ : index set of the observations
- $J = \{1, \dots, m\}$ : index set of the attributes
- $P = \{1, \dots, p\}$ : index set of the PCs
- $A \in \mathbb{R}^{n \times m}$ : data matrix with elements  $a_{ij}$  for  $i \in I, j \in J$
- $X \in \mathbb{R}^{m \times p}$ : principal components matrix with elements  $x_{jk}$  for  $j \in J, k \in P$
- $Y \in \mathbb{R}^{n \times p}$ : projected data matrix with elements  $y_{ik}$  for  $i \in I, k \in P$ , defined as  $Y = AX$
- $E \in \mathbb{R}^{n \times m}$ : reconstruction error matrix with elements  $e_{ij}$  for  $i \in I, j \in J$ , defined as  $E = A - YX^\top$
- $I_k \in \mathbb{R}^{k \times k}$ :  $k$  by  $k$  identity matrix

For a matrix  $R \in \mathbb{R}^{n_r \times m_r}$  with elements  $r_{ij}$ ,  $i = 1, \dots, n_r$ ,  $j = 1, \dots, m_r$ , we denote by

$$\|R\|_F = \sqrt{\sum_{i=1}^{n_r} \sum_{j=1}^{m_r} r_{ij}^2}$$

the Frobenius norm.

The conventional PCA problem, P1 with the  $L_2$  norm, can be written as

$$\min_{X \in \mathbb{R}^{m \times p}, X^\top X = I_p} \|A - AX X^\top\|_F^2. \quad (1)$$

Note that  $X^\top X = I_p$  is different from  $XX^\top \in \mathbb{R}^{m \times m}$  in the objective function.

We consider (1) with the  $L_1$  norm instead of the  $L_2$  norm in the objective function. The resulting P1 problem is written as

$$\min_{X \in \mathbb{R}^{m \times p}} \sum_{i \in I} \sum_{j \in J} |e_{ij}| \quad \text{s.t. } X^\top X = I_p, E = A - AXX^\top. \quad (2)$$

Next we present an iterative algorithm for (2) to minimize the reconstruction error. Instead of solving (2) directly, we iteratively solve a weighted version of (1) by giving a different weight for each observation.

We rewrite (2) in the following non-matrix form.

$$\min \sum_{i \in I} \sum_{j \in J} |e_{ij}| \quad (3a)$$

$$\text{s.t. } y_{ik} = \sum_{j \in J} a_{ij} x_{jk}, \quad i \in I, k \in P, \quad (3b)$$

$$e_{ij} = a_{ij} - \sum_{k \in P} y_{ik} x_{jk}, \quad i \in I, j \in J, \quad (3c)$$

$$\sum_{j \in J} x_{jk} x_{jq} = 0, \quad k \in P, q \in P, k < q, \quad (3d)$$

$$\sum_{j \in J} x_{jk} x_{jk} = 1, \quad k \in P, \quad (3e)$$

$$E, X, Y \text{ unconstrained.} \quad (3f)$$

Note that the corresponding  $L_2$ -PCA problem (1) can be also written as

$$\min_{E, X, Y} \sum_{i \in I} \sum_{j \in J} e_{ij}^2 \quad \text{s.t. } (3b) - (3f), \quad (4)$$

since the only difference between (1) and (2) is the objective function. However, there is no known algorithm that solves (3) optimally, whereas (4) can be solved by SVD or EVD. Hence, we want to take advantage of the fact that (4) can be optimally solved.

Let us consider a weighted version of (4):

$$g(w) = \min_{E, X, Y} \sum_{i \in I} \sum_{j \in J} w_i e_{ij}^2 \quad \text{s.t. } (3b) - (3f), \quad (5)$$

with  $w_i > 0$  for every  $i$  in  $I$ . Note that (4) and (5) are equivalent when  $w_i = 1$  for all  $i$  in  $I$ . However, solving (5) with non-constant  $w_i$ 's is not easy in its original form due to the orthogonality constraint. Instead, let us define weighted data matrix  $\bar{A} \in \mathbb{R}^{n \times m}$  with each element defined as

$$\bar{a}_{ij} = \sqrt{w_i} a_{ij}, \quad \text{for } i \in I, j \in J. \quad (6)$$

In the following proposition, we show that an optimal solution to (5) can be obtained by SVD of  $\bar{A}$ .

**Proposition 1.** Solving (5) with  $A$  is equivalent to solving (4) with  $\bar{A}$ .

*Proof.* Let  $(\bar{E}, \bar{X}, \bar{Y})$  be a solution to (4) with  $\bar{A}$ . We claim that, for any  $(\bar{E}, \bar{X}, \bar{Y})$ , there exists  $(E, X, Y)$  with the same objective function value for (5) with  $A$ , and vice versa. Let  $y_{ik} = \frac{1}{\sqrt{w_i}} \bar{y}_{ik}$ ,  $e_{ij} = \frac{1}{\sqrt{w_i}} \bar{e}_{ij}$ , and  $x_{jk} = \bar{x}_{jk}$ . We derive

$$\begin{aligned} y_{ik} &= \frac{\bar{y}_{ik}}{\sqrt{w_i}} = \frac{\sum_{j \in J} \bar{a}_{ij} \bar{x}_{jk}}{\sqrt{w_i}} = \frac{\sum_{j \in J} \sqrt{w_i} a_{ij} \bar{x}_{jk}}{\sqrt{w_i}}, \\ e_{ij} &= \frac{\bar{e}_{ij}}{\sqrt{w_i}} = \frac{\bar{a}_{ij} - \sum_{k \in P} \bar{y}_{ik} \bar{x}_{jk}}{\sqrt{w_i}} = a_{ij} - \sum_{k \in P} y_{ik} \bar{x}_{jk} \\ &= a_{ij} - \sum_{k \in P} y_{ik} x_{jk}, \\ \sum_{i=1}^n \sum_{j \in J} \bar{e}_{ij}^2 &= \sum_{i=1}^n \sum_{j \in J} w_i e_{ij}^2. \end{aligned}$$

Further, since  $\bar{x}_{jk} = x_{jk}$  for all  $j$  and  $k$ , orthonormality constraints (3d) and (3e) are automatically satisfied. Hence, solving (4) with  $\bar{A}$  is equivalent to solving (5) with  $A$ .  $\square$

Since now we know that (5) can be solved optimally, the remaining task is to define appropriate weights that give a good solution to (3). We first provide intuition behind our choices.

Let  $(E^*, X^*, Y^*)$  be an optimal solution to (3) and let

$$w_i^* = \begin{cases} \frac{\sum_{j \in J} |e_{ij}^*|}{\sum_{j \in J} (e_{ij}^*)^2}, & \text{if } \sum_{j \in J} (e_{ij}^*)^2 > 0, \\ M, & \text{if } \sum_{j \in J} (e_{ij}^*)^2 = 0, \end{cases} \quad (7)$$

be weights defined based on  $E^*$ , where  $M$  is a large number. Note that the value of  $M$ , for the case  $\sum_{j \in J} (e_{ij}^*)^2 = 0$ , does not affect the value of  $\sum_{i \in I} w_i^* \sum_{j \in J} (e_{ij}^*)^2$  because  $\sum_{j \in J} (e_{ij}^*)^2 = 0$  for the corresponding observation. However, considering the fact that we want to give less weight to the outliers in order to reduce their effect on the objective function, it is reasonable to assign a big number to the observations with zero error.

With  $w^*$  defined in (7), it is trivial to show

$$\sum_{i \in I} \sum_{j \in J} |e_{ij}^*| = \sum_{i \in I} w_i^* \sum_{j \in J} (e_{ij}^*)^2 \quad \text{and} \quad g(w^*) \leq \sum_{i \in I} \sum_{j \in J} |e_{ij}^*|. \quad (8)$$

The equality in (8) implies that, given  $w^*$  and  $E^*$ , the objective function value of (3) and (5) are equal. The inequality in (8) implies that, given  $w^*$ , the objective function value of (5) gives a lower bound on the optimal objective function value of (3). Hence, we aim to minimize the objective function of (3) by solving (5), hoping  $g(w^*)$  and  $\sum_{i \in I} \sum_{j \in J} |e_{ij}^*|$  are not far from each other.

The equality and inequality in (8) give motivation to use a weight formula similar to (7). Before presenting the weight update formula and the algorithm, let us define the following notation for the algorithm.

- $t$ : current iteration
- $w^t \in \mathbb{R}^{n \times 1}$ : weight vector used in iteration  $t$  with elements  $w_i^t$  for  $i \in I$
- $W_t \in \mathbb{R}^{n \times n}$ : diagonal matrix in iteration  $t$  with  $\sqrt{w_i^t}$ 's on the diagonal
- $A_t \in \mathbb{R}^{n \times m}$ : weighted data matrix defined in (6) with  $w^t$  in iteration  $t$ , defined as  $A_t = W_t A$
- $X_t \in \mathbb{R}^{m \times p}$ : the principal component matrix obtained by SVD of  $A_t$  in iteration  $t$
- $E_t \in \mathbb{R}^{n \times m}$ : reconstruction error matrix in iteration  $t$  with elements  $e_{ij}^t$  for  $i \in I, j \in J$ , defined as  $E_t = A - AX_t X_t^\top$
- $L2PCA(A_t, p)$ : subroutine that returns  $X_t \in \mathbb{R}^{m \times p}$  by solving (4) with  $A_t$
- $F(X_t)$ : objective function value of  $X_t$  for (3), defined as  $F(X_t) = \sum_{i \in I} \sum_{j \in J} |e_{ij}^t|$

- $F^{best}$ : current best objective function value
- $X^{best}$ : principal component matrix associated with  $F^{best}$

Note that  $X_t$  is obtained by SVD of  $A_t$ , but  $E_t = A - AX_tX_t^\top$  is based on  $A$  and is different from  $A_t - A_tX_tX_t^\top$ .

Motivated by (7), for iteration  $t + 1$ , we define

$$w_i^{t+1} = \begin{cases} \frac{\sum_{j \in J} |e_{ij}^t|}{\sum_{j \in J} (e_{ij}^t)^2}, & \text{if } \sum_{j \in J} (e_{ij}^t)^2 > 0, \\ M_t, & \text{if } \sum_{j \in J} (e_{ij}^t)^2 = 0, \end{cases} \quad (9)$$

where  $M_t = \max_{i \in I_t^+} \frac{\sum_{j \in J} |e_{ij}^t|}{\sum_{j \in J} (e_{ij}^t)^2}$  is the largest weight among the observations in  $I_t^+ = \{i \in I \mid \sum_{j \in J} (e_{ij}^t)^2 > 0\}$ . Using  $u^{t+1}$  in (9) for the weights is natural and we empirically observe that the algorithm is convergent. However, it is not trivial to show the convergence with  $u^{t+1}$  for (5). Hence, in order to show the convergence of the algorithm, we present a modified update formula based on  $u^{t+1}$ .

$$w_i^{t+1} = \begin{cases} w_i^t(1 - \beta^t), & \text{if } u_i^{t+1} < w_i^t(1 - \beta^t), \\ u_i^{t+1}, & \text{if } w_i^t(1 - \beta^t) \leq u_i^{t+1} \leq w_i^t(1 + \beta^t), \\ w_i^t(1 + \beta^t), & \text{if } u_i^{t+1} > w_i^t(1 + \beta^t), \end{cases} \quad (10)$$

where  $\beta \in (0, 1)$ . Note that  $\beta^t$  is the  $\beta$  to the power of  $t$  and is different from other superscript-containing notations such as  $w_i^t$  or  $u_i^t$ . The role of (10) is to enable bounds of the change for  $w^{t+1}$  from  $w^t$ . If  $u_i^{t+1}$  is too small compared to  $w_i^t$ , then  $w_i^{t+1}$  is assigned a value between  $u_i^{t+1}$  and  $w_i^t$ . If  $u_i^{t+1}$  is too large compare to  $w_i^t$ , then  $w_i^{t+1}$  obtains a value between  $u_i^{t+1}$  and  $w_i^t$ . Otherwise,  $w_i^{t+1}$  follows the weight formula in (9). Given  $\beta \in (0, 1)$ , we have  $\lim_{t \rightarrow \infty} \beta^t = 0$ , which implies  $\lim_{t \rightarrow \infty} w_i^t - w_i^{t+1} = 0$ . Further, since  $u$  and  $w$  are bounded above and below, we can show that  $w^t$  is convergent. By setting  $\beta$  close to 1, we would have  $w_i^{t+1} = u_i^{t+1}$  in most cases, as  $1 - \beta^t$  and  $1 + \beta^t$  are close to 0 and 2 for small values of  $t$ , i.e., early iterations. From all these facts and by using elementary mathematics, the following lemma follows.

**Lemma 1.** With  $w^t$  defined in (10),  $w^t$  and  $A_t$  are convergent in  $t$ .

We present the overall algorithmic framework in Algorithm 1. The algorithm requires data matrix  $A$ , target dimension  $p$ , and tolerance parameter  $\varepsilon$  as input. After the initialization of weights and the best objective function value  $F^{best}$  in Step 1, the while loop is executed until  $w^t$  and  $w^{t+1}$  are close enough. In each iteration of the while loop,  $A_t$  is constructed based on  $w_t$  and  $X_t$  is obtained by SVD of  $A_t$  (Steps 3 and 4). If  $X_t$  gives a lower objective function value than  $X^{best}$ , then  $X^{best}$  and  $F^{best}$  are updated. Recall that  $X_t$  is obtained by using  $A_t$ , but  $F(X_t)$  uses the original data matrix  $A$ . Each iteration ends after the update of weights in Step 6. Observe that the termination criteria in Step 2 solely depends on the convergence of  $w_t$ . Hence, the algorithm terminates in a finite number of iterations.

**Lemma 2.** Eigenvalues of  $A_t^\top A_t$  are convergent in  $t$ .

*Proof.* Recall that the weights  $w_t$  and weighted matrix  $A_t$  are convergent, which also implies  $A_t^\top A_t$  is also convergent.

---

### Algorithm 1 wPCA (Weight-based algorithm for $L_1$ -PCA)

---

**Require:**  $A$  (data),  $p$  (target dimension),  $\varepsilon$  (tolerance),  $\beta$   
**Ensure:** principal components  $X^{best} \in \mathbb{R}^{m \times p}$

- 1:  $t \leftarrow 1$ ,  $w_i^0 \leftarrow 2$ ,  $w_i^1 \leftarrow 1$ ,  $F^{best} \leftarrow \infty$
  - 2: **while**  $\|w^t - w^{t-1}\|_1 > \varepsilon$  **do**
  - 3:   set  $A_t$  based on  $w^t$
  - 4:    $X_t \leftarrow L2PCA(A_t, p)$
  - 5:   **if**  $F(X_t) < F^{best}$  **then**  $X^{best} \leftarrow X_t$ ,  $F^{best} \leftarrow F(X_t)$
  - 6:   update  $w_{t+1}$  by using (10) given  $\beta$
  - 7:    $t \leftarrow t + 1$
  - 8: **end while**
- 

Since the eigenvalues of symmetric matrices are point-wise convergent if the matrices are convergent [16], it is trivial to see that the eigenvalues of  $A_t^\top A_t$  are convergent.  $\square$

Hence, Algorithm 1 gives convergent eigenvalues. Although eigenvalues are convergent, it is not trivial to show the convergence of  $X_t$ . This is because even a slight change in a matrix can cause a change in an eigenvector, and eigenvalue-eigenvector pairs are not unique.

In order to provide convergent eigenpairs and accelerate the algorithm for large scale data, we use the first order eigenpair approximation formula from [23]. Let  $(X_{t-1}^i, \lambda_{t-1}^i)$  be an approximate eigenpair of  $A_{t-1}$  and  $A_t = A_{t-1} + \Delta_t$ . Then, the approximate eigenpairs of  $A_t$  can be obtained by

$$\lambda_t^i = \lambda_{t-1}^i + (X_{t-1}^i)^\top \Delta_t X_{t-1}^i, \quad (11)$$

$$X_t^i = X_{t-1}^i + \sum_{j \neq i} \left( \frac{(X_{t-1}^j)^\top \Delta_t X_{t-1}^i}{\lambda_{t-1}^i - \lambda_{t-1}^j} \right) X_{t-1}^j, \quad (12)$$

using the formula in [23]. The error is of the order of  $o(\|\Delta_t\|^2)$ . Let  $L2PCA\_Approx$  be a function that returns principal components by formula (11) - (12). The modified algorithm is presented in Algorithm 2.

---

### Algorithm 2 awPCA (Approximated weight-based algorithm for $L_1$ -PCA)

---

**Require:**  $A$  (data),  $p$  (target dimension),  $\varepsilon$  (tolerance),  $\beta$ ,  $\gamma$   
**Ensure:** principal components  $X^{best} \in \mathbb{R}^{m \times p}$

- 1:  $t \leftarrow 1$ ,  $w_i^0 \leftarrow 2$ ,  $w_i^1 \leftarrow 1$ ,  $F^{best} \leftarrow \infty$
  - 2: **while**  $\|w^t - w^{t-1}\|_1 > \varepsilon$  **do**
  - 3:   set  $A_t$  based on  $w^t$ ,  $\Delta_t \leftarrow A_t - A_{t-1}$
  - 4:   **if**  $\|\Delta_t\|^2 > \gamma \cdot \|A\|^2$  **then**  $X_t \leftarrow L2PCA(A_t, p)$
  - 5:   **Else**  $(X_t, \lambda_t) \leftarrow L2PCA\_approx(X_{t-1}, \lambda_{t-1}, p)$
  - 6:   **if**  $F(X_t) < F^{best}$  **then**  $X^{best} \leftarrow X_t$ ,  $F^{best} \leftarrow F(X_t)$
  - 7:   update  $w_{t+1}$  by using (10) given  $\beta$
  - 8:    $t \leftarrow t + 1$
  - 9: **end while**
- 

The difference is only in Lines 4 and 5. If the change in  $A_t$  is large (greater than  $\gamma \cdot \|A\|^2$ ), we use the original procedure  $L2PCA$ . If the change in  $A_t$  is small (less than or equal to  $\gamma \cdot \|A\|^2$ ), then we use the update formula (11) and (12). Algorithm 2 has the following convergence result.

**Proposition 2.** Eigenpairs of  $A_t^\top A_t$  in Algorithm 2 are convergent in  $t$ .

*Proof.* Note that the convergence of  $\Delta_t$  in Lemma 2 does not depend on how the eigenvectors are obtained and thus it holds whether we execute Lines 4 or 5 in each iteration. Hence, we have  $\lim_{t \rightarrow \infty} \Delta_t = 0$  in Algorithm 2.

Since  $\Delta_t$  converges to zero, after a certain number of iterations  $\bar{t}$ , we have  $\|\Delta_t\|^2 < \gamma \cdot \|A\|^2$  for all  $t > \bar{t}$ . From such large  $\bar{t}$ , the approximation rule applies. In [23], it is shown that

$$\begin{aligned} \bar{\lambda}_t^i &= \bar{\lambda}_{t-1}^i + (\bar{X}_{t-1}^i)^\top \Delta_t \bar{X}_{t-1}^i + o(\|\Delta_t\|^2), \\ \bar{X}_t^i &= \bar{X}_{t-1}^i + \sum_{j \neq i} \left( \frac{(\bar{X}_{t-1}^j)^\top \Delta_t \bar{X}_{t-1}^i}{\bar{\lambda}_{t-1}^i - \bar{\lambda}_{t-1}^j} \right) \bar{X}_{t-1}^j + o(\|\Delta_t\|^2), \end{aligned}$$

when  $(\bar{X}_{t-1}, \bar{\lambda}_{t-1})$  are the exact eigenpairs. Since all of the terms in the formula are bounded and  $\lim_{t \rightarrow \infty} \Delta_t = 0$ , we conclude that the eigenpairs of  $A_t^\top A_t$  are convergent.  $\square$

### III. COMPUTATIONAL EXPERIMENT

We compare the performance of the proposed and benchmark algorithms for varying instance sizes ( $m, n$ ) and number of PCs ( $p$ ). All experiments were performed on a personal computer with 8 GB RAM and Intel Core i7 (2.40GHz dual core). We implement Algorithms 1 and 2 in R [22], which we denote as **wPCA** and **awPCA**, respectively. The R script for **wPCA** and **awPCA** is available on a web site <sup>1</sup>. For the **awPCA** implementation, we use condition  $\|w^t - w^{t-1}\|_1 > \gamma \cdot \|w^t\|_1$  instead of  $\|\Delta_t\|^2 > \gamma \cdot \|A\|^2$ , to avoid unnecessary calculation of  $\|\Delta_t\|^2$  in Line 4 of Algorithm 2. The weight-based condition is similar to the original condition as the difference in the weight captures  $\Delta_t$ . For the experiment, we use parameters  $\varepsilon = 0.001$  and  $\beta = 0.99$  for **wPCA** and **awPCA** and  $\gamma = 0.1$  for **awPCA**, where the parameters are tuned based on pilot runs to balance the solution quality and execution time. We also set up maximum number of iterations to 200. We compare our algorithms with the algorithms in [3], [14], [15], [20]. The work in [4] provides R implementations of the algorithms in [3], [14], [15], which we denote as **Brooks**, **Ke**, **Kwak**, respectively. We implement the algorithm in [20] in R, which we denote as **Nie**.

Although algorithms **Ke** and **Brooks** are for (2) and **Kwak** and **Nie** are for the  $L_1$  norm version of **P2**, we evaluate the objective function value for all benchmark algorithms and compare them against our algorithms. Especially, **Kwak** and **Nie**, which solve different  $L_1$ -PCA problem, are included because

- we observed that **Kwak** and **Nie** are better than **Ke** and **Brooks** for (2) for some instances, and
- we found that **Kwak** and **Nie** are more scalable and solve larger data sets in a reasonable time in the experiment.

Therefore, we include **Kwak** and **Nie** for the comparison for solving **P1**.

It is worth to note that **Ke** and **Brooks** try to find the best fitting subspace, where definition of  $E$  in (2) is replaced by

$E = A - UX^\top$  and  $U \in R^{n \times p}$ . The optimal solutions of the two formulations may be different despite both minimizing the  $L_1$  distance from the original data.

Let  $F_{algo}$  represent the objective function value obtained by  $algo \in \{\text{Ke, Brooks, Kwak, Nie, wPCA, awPCA}\}$ , with respect to (2). For the comparison purposes for **awPCA**, we use the gap from the best objective function value defined as

$$\Delta_{algo} = \min \left\{ \frac{F_{algo}}{\min\{F_{\text{awPCA}}, F_{\text{Ke}}, F_{\text{Brooks}}, F_{\text{Kwak}}, F_{\text{Nie}}\}} - 1, 1 \right\},$$

for each  $algo \in \{\text{awPCA, Ke, Brooks, Kwak, Nie}\}$ . Similarly, for **wPCA**, we define

$$\Delta_{algo} = \min \left\{ \frac{F_{algo}}{\min\{F_{\text{wPCA}}, F_{\text{Ke}}, F_{\text{Brooks}}, F_{\text{Kwak}}, F_{\text{Nie}}\}} - 1, 1 \right\},$$

for each  $algo \in \{\text{wPCA, Ke, Brooks, Kwak, Nie}\}$ . Note that  $\Delta_{algo}$  represents the gap between  $algo$  and the best of all algorithms. Note also that we set up an upper bound of 1 for  $\Delta_{algo}$ . Hence, if the gap is larger than 1 (or 100%), then  $\Delta_{algo}$  is assigned value of 1 (or 100%).

For all of the instances used in the experiment, we first standardize each column and deal exclusively with the standardized data. Hence, the reconstruction errors are also calculated based on the standardized data.

In the computational experiment, we observed that  $\Delta_{\text{awPCA}}$  and  $\Delta_{\text{wPCA}}$  are very similar while the execution time of **awPCA** is much faster. Hence, in this section, we first focus on presenting the performance of **awPCA** against the benchmark algorithms and after on comparing the difference between **wPCA** and **awPCA**.

The rest of the section is organized as follows. In Section III-A, we present synthetic instance generation procedures and explain the instances from the UCI Machine Learning Repository [2]. In Sections III-B and III-C, we present the performance of **awPCA** for the synthetic and UCI instances, respectively. In Section III-D, we compare the performance of **wPCA** and **awPCA** for the UCI instances.

#### A. Instances

1) *Synthetic Instances:* In order to provide a systematic analysis, we generate synthetic instances with presence of outliers and various  $(m, n, r)$ , where  $m \in \{20, 50\}$ ,  $n \in \{100, 300\}$ , and  $r \in \{0, 0.1, 0.2, 0.3\}$ . For each  $(m, n, r)$ , we generate 5 distinguished instances. Hence, we have a total of 80 generated instances. The synthetic instances used in the experiment are available on a web site <sup>2</sup>. The detailed algorithm is presented at the end of this section. In the generation procedure,  $r\%$  of observations are generated to have a higher variance than the remaining normal observations. The instance generation algorithm needs additional parameter  $q$  (target rank), which we fix to 10 for the instances we generated in this experiment. Hence, the instances we use in the experiment are likely to have rank equal to 10. We consider different  $p$  values, where  $p \in \{8, 9, 10, 11, 12\}$ . Given that  $q = 10$ , we select  $p$  values around 10.

<sup>2</sup>[http://dynresmanagement.com/uploads/3/3/2/9/3329212/pca\\_instance\\_park\\_klabjan.zip](http://dynresmanagement.com/uploads/3/3/2/9/3329212/pca_instance_park_klabjan.zip)

<sup>1</sup><http://dynresmanagement.com/uploads/3/3/2/9/3329212/w1pca.zip>

The purpose of the instance generation algorithm is to generate instances with some of the observations as outliers, so that  $L_1$ -PCA solutions are more likely to be away from  $L_2$ -PCA solutions. In order to generate instances, we use the procedure described in [14] with a slight modification. The instances used in the experiment in [14] have fixed parameters and constant valued outliers to simulate data loss. To check the performance of the algorithms over various parameters and different (non-constant) patterns of outliers, we generate our own instances. In the generation procedure of [14], a matrix with a small fixed rank is generated and then extremely large constant values randomly replace the original data matrix. In their instances, outliers have the same value, which can be interpreted as data loss, but they do not consider outliers due to incorrect measurements or cases with only a few observations with outliers. Our procedure addresses all these cases.

We present the procedure in Algorithm 3.

---

**Algorithm 3** PCA instance generation

---

**Require:**  $m, n, q$  (target rank),  $r$  (% outliers)

**Ensure:**  $A \in \mathbb{R}^{n \times m}$

- 1: Generate random matrix  $P = [p_{ij}] \in \mathbb{R}^{n \times m}$  with  $p_{ij} \sim U(-100, 100)$
  - 2: **for** each row
  - 3:   Generate random number  $u_1 \sim U(0, 1)$
  - 4:   **if**  $u_1 < r$
  - 5:     **for** each column  $j \leq q$
  - 6:       Generate  $u_2 \sim U(0, 1)$
  - 7:       **if**  $u_2 < 0.1$ , **then**  $h_{ij} \sim N(0, 30)$
  - 8:       **else**  $h_{ij} \sim N(0, 1)$
  - 9:     **end for**
  - 10:   **else**  $h_{ij} \sim N(0, 1)$
  - 11: **end for**
  - 12: Obtain  $P = U\Sigma V^\top$ , the SVD of  $P$
  - 13: Construct  $A = (U[, 1 : q] + H)\Sigma[1 : q, 1 : q]V^\top[1 : q]$  with  $H = [h_{ij}] \in \mathbb{R}^{n \times q}$  generated in Steps 2 - 10
  - 14: Adjust  $A$  to have 0 mean for each column
- 

In Step 1, we first generate random matrix  $P$ , where each  $p_{ij}$  is from the uniform distribution between -100 and 100. Next in Steps 2 - 10, we generate random perturbation matrix  $H$  with approximately  $r$  percent of rows having extremely large perturbations, where each row has approximately 10% extreme value entries. After SVD of  $PU\Sigma V^\top$  in Step 12, data matrix  $A$  is generated, where  $U[, 1 : q]$  is the submatrix of  $U$  with the first  $q$  columns,  $\Sigma[1 : q, 1 : q]$  is the submatrix of  $\Sigma$  with the first  $q$  columns and  $q$  rows, and  $V^\top[1 : q]$  is the submatrix of  $V^\top$  with the first  $q$  columns. The final data matrix  $A$  is generated in Step 14 after adjusting it to have 0 column means.

2) *UCI instances:* We also consider classification datasets from the UCI Machine Learning Repository [2] and adjust them to create PCA instances. Based on the assumption that observations in the same class of a classification data set have similar attribute distributions, we consider each class

of the classification datasets. For each dataset, we partition the observations based on labels (classes). When there exist many labels, we select the top two labels with respect to the number of observations among all labels. For each partitioned data, labels and attributes with zero standard deviation (hence, meaningless) are removed and the matrix is standardized to have zero mean and unit standard deviation for each attribute.

In Table I, we list the PCA instances we used and the corresponding original dataset from [2]. In the first column, abbreviate names of the original data sets are presented. The full names of the data sets are Breast Cancer Wisconsin, Indian Liver Patient Dataset, Cardiotocography, Ionosphere, Connectionist Bench (Sonar), Landsat Satellite, Spambase, Magic Gamma Telescope, Page Blocks Classification, and Pen-Based Recognition of Handwritten Digits. Each PCA instance is classified as small or large based on  $m$  and  $n$ . If  $mn \leq 15,000$ , the instance is classified as small, otherwise, the instance is classified as large. In the last column in Table I, the small and large instances are indicated by  $S$  and  $L$ , respectively. For the large instances, only Kwak and Nie are compared with the proposed algorithms, due to scalability issues of the other benchmark algorithms.

TABLE I: PCA instances created based on the datasets from the UCI Machine Learning Repository [2]

Original dataset from UCI			PCA instance		
Name	$(m, n)$	Num labels	Name	$(m, n)$	size
cancer	(9,699)	2	cancer_2	(9,444)	S
			cancer_4	(9,239)	S
ilpd	(10,583)	2	ilpd_1	(10,416)	S
			ilpd_2	(10,167)	S
cardio	(21,2126)	10	cardio_1	(19,384)	S
			cardio_2	(19,579)	S
iono	(34,351)	2	iono_b	(33,126)	S
			iono_g	(32,225)	S
sonar	(60,208)	1	sonar_g	(60,111)	S
			sonar_r	(60,97)	S
landsat	(36,4435)	7	landsat_1	(36,1072)	L
			landsat_3	(36,961)	L
spam	(57,4601)	2	spam_0	(57,2788)	L
			spam_1	(57,1813)	L
magic	(10,19020)	2	magic_g	(10,12332)	L
			magic_h	(10,6688)	L
blocks	(10,5473)	5	blocks_1	(10,4913)	L
hand	(16,10992)	10	hand_0	(16,1142)	L
			hand_1	(16,1143)	L

*B. Performance of awPCA for Synthetic Instances*

In Table II, we present the result for awPCA for the synthetic instances. Although we created synthetic instances with varying  $r$  (% of outliers) values, we observed that the performances of the algorithms are very similar over different  $r$  values for each  $(m, n, p)$  triplet. Hence, in Table II, we present the average value over all  $r$ . That is, each row of the table is the average of 20 instances for the corresponding  $(m, n)$  pair given  $p$ . The first two columns are the instance size and number of PCs, the next five columns are  $\Delta_{algo}$  for all algorithms, and the last five columns are the execution times in seconds. For each row, the lowest  $\Delta_{algo}$  value among the five algorithms is boldfaced.

TABLE II: Performance of awPCA for synthetic instances

Instance ( $m, n$ )	$p$	Gap from the best ( $\Delta_{algo}$ )					Time (seconds)				
		awPCA	Ke	Brooks	Kwak	Nie	awPCA	Ke	Brooks	Kwak	Nie
(20, 100)	8	<b>1%</b>	6%	2%	12%	18%	0.0	0.7	1.3	0.0	0.0
	9	4%	22%	<b>3%</b>	16%	26%	0.0	0.5	1.3	0.0	0.0
	10	<b>0%</b>	<b>0%</b>	1%	7%	7%	0.0	0.4	1.3	0.0	0.0
	11	<b>0%</b>	69%	2%	7%	12%	0.0	0.4	1.3	0.0	0.0
	12	<b>0%</b>	70%	2%	7%	16%	0.0	0.4	1.2	0.0	0.0
(20, 300)	8	<b>2%</b>	8%	3%	10%	12%	0.0	4.9	9.5	0.0	0.1
	9	<b>3%</b>	10%	<b>3%</b>	13%	19%	0.0	2.8	9.6	0.0	0.1
	10	<b>0%</b>	<b>0%</b>	1%	3%	3%	0.0	1.2	9.6	0.0	0.1
	11	<b>0%</b>	9%	1%	3%	6%	0.0	1.2	9.6	0.0	0.1
	12	<b>0%</b>	7%	1%	3%	8%	0.0	1.2	9.6	0.0	0.1
(50, 100)	8	<b>1%</b>	3%	2%	13%	16%	0.0	3.8	19.2	0.0	0.0
	9	<b>1%</b>	7%	3%	15%	21%	0.0	2.9	19.1	0.0	0.0
	10	<b>0%</b>	<b>0%</b>	3%	7%	7%	0.0	2.9	19.3	0.0	0.0
	11	<b>0%</b>	100%	2%	7%	8%	0.0	4.3	19.2	0.0	0.0
	12	<b>0%</b>	100%	3%	7%	10%	0.0	4.8	19.2	0.0	0.0
(50, 300)	8	<b>1%</b>	3%	2%	9%	13%	0.1	27.0	227.4	0.0	0.1
	9	<b>2%</b>	5%	3%	10%	16%	0.1	22.0	226.5	0.0	0.2
	10	<b>0%</b>	<b>0%</b>	1%	3%	3%	0.0	18.7	226.7	0.0	0.2
	11	<b>0%</b>	86%	1%	3%	4%	0.0	23.1	228.1	0.0	0.2
	12	<b>0%</b>	98%	1%	3%	5%	0.0	27.2	226.0	0.0	0.2

Note that  $\Delta_{awPCA}$  values are near zero for all instances. Further,  $\Delta_{awPCA}$  has the lowest gaps (boldfaced numbers) for all instances among all algorithms except for one instance class. Brooks constantly gives the second best gaps while Ke gives 0% gaps for  $p = 10$ , third best gaps for  $p < 10$  and worst gaps for  $p > 10$ . Nie and Kwak generally give the similar result as they are designed to solve the same problem.

The execution times of the algorithms can be grouped into two groups: awPCA, Kwak, and Nie are in the faster group and Ke and Brooks are in the slower group. Ke and Brooks spend much more time on larger instances compared to the other three algorithms. Although it is not easy to compare, Kwak is the fastest among all algorithms, yet  $\Delta_{Kwak}$  is not as low as  $\Delta_{awPCA}$ . It is important to note that the difference in the execution time between Kwak and awPCA is negligible and awPCA is fastest among the algorithms designed to solve P1 with the  $L_1$  norm (i.e., Ke, Brooks, and awPCA).

### C. Performance of awPCA for UCI Instances

For each PCA instance in Table I, we execute the algorithms with various  $p$  values. The number of PCs  $p$  covers the entire spectrum  $0, \dots, m$  in increments of 2, 3, 5, or 10 depending on  $m$ : *cancer* and *ilpd* with  $p \in \{2, 4, 6, 8\}$ , *cardio* with  $p \in \{3, 6, 9, 12, 15\}$ , *iono* with  $p \in \{5, 10, 15, 20, 25, 30\}$ , *sonar* and *spam* with  $p \in \{10, 20, 30, 40, 50\}$ , *landsat* with  $p \in \{5, 10, \dots, 30, 35\}$ , *magic* and *block* with  $p \in \{1, 3, 5, 7, 9\}$ , and *hand* with  $p \in \{2, 4, \dots, 12, 14\}$ .

For the small UCI instances, we present heat maps of  $\Delta_{algo}$  and the execution times of all algorithms in Figures 2 and 3. In both figures, the numbers are  $\Delta_{algo}$  in percentage or execution time in seconds, a white cell implies near-zero  $\Delta_{algo}$  or near-zero execution time, and a dark gray cell implies the opposite. In Figure 2, awPCA is consistently best with Brooks usually being the second best algorithm. The value of  $\Delta_{awPCA}$  is zero except for a few cases. For such cases with  $\Delta_{awPCA} > 0$ , Brooks performs the best. The values of  $\Delta_{Ke}$ ,  $\Delta_{Kwak}$ , and  $\Delta_{Nie}$  tend to increase in  $p$ . In Figure 3, we observe the same trend from Section III-B: awPCA, Kwak, and Nie are in the

faster group and Ke and Brooks become slower as instance size increases.

For the large UCI instances, we only compare awPCA against Kwak and Nie, due to scalability issues of Ke and Brooks. Hence,  $\Delta_{algo}$  here is the gap from the best of awPCA, Kwak and Nie. In Figures 4 and 5, we present heat maps of  $\Delta_{algo}$  and the execution times of the three algorithms. Similar to Figures 2 and 3, a white cell implies a low value. In Figure 4, awPCA is consistently best except for four cases and even for the four cases  $\Delta_{awPCA}$  are very small. We observe that  $\Delta_{Kwak}$  and  $\Delta_{Nie}$  tend to increase in  $p$ , where  $\Delta_{Kwak}$  is slightly smaller than  $\Delta_{Nie}$  in general. In Figure 5, the execution time of Kwak is the fastest, and awPCA and Nie are of the same magnitude, although awPCA is slightly faster than Nie.

Based on the results for the UCI instances, we conclude that awPCA performs the best while the execution time of awPCA is of the same order or lower than the remaining algorithm.

### D. Comparison of wPCA and awPCA for UCI Instances

In this section, we compare wPCA and awPCA for the UCI instances in terms of solution quality ( $\Delta_{wPCA}$  and  $\Delta_{awPCA}$ ) and execution time. In Table III, we present the average performance of wPCA and awPCA for all  $p$  values considered in Section III-C. The fourth column is defined as  $diff = \Delta_{wPCA} - \Delta_{awPCA}$ , where a negative  $diff$  value implies that wPCA gives a better solution and near-zero  $diff$  value implies that  $\Delta_{wPCA}$  and  $\Delta_{awPCA}$  are similar. The seventh column is defined as  $ratio = \text{execution time of awPCA} / \text{execution time of wPCA}$ , where a less-than-one  $ratio$  value implies that awPCA is faster than wPCA. In Table III, we observe that  $\Delta_{wPCA}$  and  $\Delta_{awPCA}$  are very similar except for two instances (boldfaced values), while awPCA spends only 20% of the time of wPCA on average. Note also that awPCA is not always inferior to wPCA. Although it is rare, awPCA gives a better solution than wPCA for instances *magic\_h* and *cardio\_1*. In general, we found that  $\Delta_{awPCA}$  is very similar or slightly larger than  $\Delta_{wPCA}$ , while awPCA is much faster. On the other hand, we can also ignore the time difference if execution times are within a few seconds. The UCI instances *spam\_0*, *spam\_1*, *magic\_g*, and *magic\_h* with clear time difference between wPCA and awPCA have  $mn > 50000$ . Therefore, we recommend to use wPCA when the data size is small and awPCA when the data size is very large.

## IV. CONCLUSIONS

In this paper, we consider the  $L_1$ -PCA problem minimizing the  $L_1$  reconstruction errors and present iterative algorithms, wPCA and awPCA, where awPCA is an approximation version of wPCA developed to avoid computationally expensive operations of SVD. The core of the algorithms relies on an iteratively reweighted least squares scheme and the expressions in (8). Although the optimality of  $L_1$ -PCA was not able to be shown and remains unknown, we show that the eigenvalues of wPCA and awPCA converge and that eigenvectors of awPCA converge. In the computational experiment, we observe that awPCA outperforms all of the benchmark algorithms while

Instance	cancer_2	cancer_4	ilpd_1	ilpd_2	cardio_1	cardio_2	iono_r	iono_g	sonar_g	sonar_r
p	2 4 6 8	2 4 6 8	2 4 6 8	2 4 6 8	3 6 9 12 15 18	3 6 9 12 15 18	5 10 15 20 25 30	5 10 15 20 25 30	10 20 30 40 50	10 20 30 40 50
awPCA	0 0 15 61	0 0 3 0	3 0 35 1	1 0 0 4	0 0 0 0 3 0	0 0 0 0 7 0	1 0 0 0 0 2	0 0 0 0 3 14	0 0 0 0 0	0 0 0 0 0
Ke	15 67 100 100	0 0 5 41	0 8 52 100	0 12 26 90	3 10 20 22 14 100	0 7 19 43 39 100	0 0 2 2 4 47	0 26 32 51 96 100	0 3 4 19 46	0 1 6 35 67
Brooks	2 22 0 0	0 1 4 4	1 3 0 0	1 6 9 0	4 8 10 6 0 1	1 7 2 9 0 20	2 3 4 5 7 0	0 1 1 1 0 0	4 6 9 10 10	4 6 9 14 15
Kwak	1 41 38 100	0 1 5 16	3 9 49 45	2 4 12 26	4 7 14 15 13 0	3 8 8 15 9 1	2 3 4 4 5 8	5 5 5 5 4 15	3 5 5 9 11	3 5 6 10 11
Nie	0 39 69 100	1 1 0 60	12 12 46 56	6 1 28 13	6 10 15 12 19 13	1 11 13 19 37 6	3 4 5 11 10 13	30 15 13 19 36 98	4 14 17 35 57	7 10 16 33 57

Fig. 2: Heat map of the gap (%) from the best for small UCI instances

Instance	cancer_2	cancer_4	ilpd_1	ilpd_2	cardio_1	cardio_2	iono_r	iono_g	sonar_g	sonar_r
p	2 4 6 8	2 4 6 8	2 4 6 8	2 4 6 8	3 6 9 12 15 18	3 6 9 12 15 18	5 10 15 20 25 30	5 10 15 20 25 30	10 20 30 40 50	10 20 30 40 50
awPCA	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0 0	0 0 0 0 0 0	0 0 0 0 0 0	0 0 0 0 0 0	0 0 0 0 0	0 0 0 0 0
Ke	1 2 3 1	1 1 0 0	3 3 1 1	0 0 1 0	7 7 12 11 3 1	17 18 11 12 5 1	3 4 4 5 4 3	6 9 12 12 8 2	9 19 27 30 28	8 15 23 27 24
Brooks	2 2 2 2	1 1 1 1	3 3 3 3	1 1 1 1	23 21 19 23 23 21	39 40 41 37 38 37	10 10 11 10 11 12	33 29 26 29 26 32	45 47 53 51 54	36 35 36 36 36
Kwak	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0 0	0 0 0 0 0 0	0 0 0 0 0 0	0 0 0 0 0 0	0 0 0 0 0	0 0 0 0 0
Nie	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0 0 0	0 0 0 0 0 0	0 0 0 0 0 0	0 0 0 0 0 0	0 0 0 0 0	0 0 0 0 0

Fig. 3: Heat map of the execution times (seconds) for small UCI instances

Instance	landsat_1	landsat_3	spam_0	spam_1	magic_g	magic_h	blocks_1	hand_0	hand_1
p	5 10 15 20 25 30 35	5 10 15 20 25 30 35	10 20 30 40 50	10 20 30 40 50	3 5 7 9	3 5 7 9	3 5 7 9	2 4 6 8 10 12 14	2 4 6 8 10 12 14
awPCA	0 0 0 0 0 0	0 0 0 0 0 0	0 0 0 0 0 0	0 0 0 0 0 0	0 0 0 1	1 0 1 0	0 0 0 0	0 0 0 0 0 0	0 0 0 0 0 0
Kwak	1 1 0 1 1 2 0	1 1 1 1 2 5 0	34 69 100 69 100	24 71 100 100 100	0 1 1 4	0 4 0 10	5 0 5 17	1 1 1 7 3 1 1	2 2 3 4 3 0 1
Nie	4 3 3 5 7 13 38	2 2 3 6 6 8 30	38 76 100 83 100	24 73 100 100 100	3 3 11 0	3 18 4 3	6 0 7 5	2 5 4 11 13 4 28	3 2 7 9 10 20 14

Fig. 4: Heat map of the gap (%) from the best for large UCI instances

Instance	landsat_1	landsat_3	spam_0	spam_1	magic_g	magic_h	blocks_1	hand_0	hand_1
p	5 10 15 20 25 30 35	5 10 15 20 25 30 35	10 20 30 40 50	10 20 30 40 50	3 5 7 9	3 5 7 9	3 5 7 9	2 4 6 8 10 12 14	2 4 6 8 10 12 14
awPCA	2 1 1 1 1 0 0	0 0 1 0 1 1 0	5 6 6 5 5	2 3 3 4 5	14 17 18 14	5 15 5 4	3 3 3 2	0 0 1 0 0 0 0	0 0 0 0 0 0 0
Kwak	0 0 0 0 0 0 0	0 0 0 0 0 0 0	0 1 1 1 2	0 0 1 1 1	0 0 0 0	0 0 0 0	0 0 0 0 0	0 0 0 0 0 0 0	0 0 0 0 0 0 0
Nie	0 1 1 1 1 1 1	0 1 1 1 1 2 2	4 6 7 8 8	1 1 2 5 4	6 19 11 16	5 10 8 12	1 4 2 3	0 0 0 1 1 1 0	0 0 0 1 1 1 1

Fig. 5: Heat map of the execution times (seconds) for large UCI instances

TABLE III: Comparison of wPCA and awPCA for UCI instances

Instance	Gap from the best ( $\Delta_{algo}$ )			Time (seconds)		
	wPCA	awPCA	diff	wPCA	awPCA	ratio
cancer_2	15.7%	19.0%	<b>-3.3%</b>	1.7	0.1	0.1
cancer_4	0.8%	0.8%	0.0%	0.4	0.0	0.1
ilpd_1	1.0%	9.8%	<b>-8.7%</b>	0.3	0.1	0.3
ilpd_2	0.9%	1.2%	-0.3%	0.5	0.0	0.0
cardio_1	0.7%	0.6%	0.1%	0.3	0.1	0.3
cardio_2	1.0%	1.2%	-0.2%	1.1	0.1	0.1
iono_r	0.4%	0.4%	0.0%	0.4	0.0	0.1
iono_g	2.8%	2.8%	0.0%	0.2	0.0	0.2
sonar_g	0.1%	0.1%	0.0%	1.0	0.1	0.1
sonar_r	0.0%	0.0%	0.0%	0.9	0.1	0.1
landsat_1	0.0%	0.0%	0.0%	3.2	0.6	0.2
landsat_3	0.6%	0.6%	0.0%	5.7	0.4	0.1
spam_0	0.0%	0.0%	0.0%	16.7	5.2	0.3
spam_1	0.0%	0.0%	0.0%	18.9	3.2	0.2
magic_g	0.1%	0.2%	-0.1%	47.7	15.6	0.3
magic_h	0.6%	0.5%	0.1%	13.8	7.1	0.5
blocks_1	0.0%	0.0%	0.0%	1.8	0.3	0.2
hand_0	0.0%	0.0%	0.0%	1.8	0.3	0.2
hand_1	0.0%	0.0%	0.0%	2.2	0.3	0.1
average			-0.7%			0.2

the execution times are competitive. Out of the four algorithms designed to minimize the L1 reconstruction errors (Ke, Brooks, wPCA, awPCA), we observe that awPCA is the fastest algorithm with near-best solution qualities.

## REFERENCES

- [1] A. Baccini, P. Besse, and A. de Fuguerolles. A L1-norm PCA and a heuristic approach. In *Proceedings of the International Conference on Ordinal and Symbolic Data Analysis*, pages 359–368, March 1996.
- [2] K. Bache and M. Lichman. UCI machine learning repository, 2013.
- [3] J. Brooks, J. Dula, and E. Boone. A pure L1-norm principal component analysis. *Computational Statistics and Data Analysis*, 61:83–98, May 2013.
- [4] J. Brooks and S. Jot. pcaL1: An implementation in R of three methods for L1-norm principal component analysis. *Optimization Online*, 2012.
- [5] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust Principal Component Analysis? *Journal of the ACM*, 58(3), 2011.
- [6] V. Choulakian. L1-norm projection pursuit principal component analysis. *Computational Statistics and Data Analysis*, 50(6):1441–1451, March 2006.
- [7] C. Croux and A. Ruiz-Gazen. High breakdown estimators for principal components: the projection-pursuit approach revisited. *Journal of Multivariate Analysis*, 95(1):206 – 226, 2005.
- [8] I. Daubechies, R. DeVore, M. Fornasier, and C. S. Gntkr. Iteratively reweighted least squares minimization for sparse recovery. *Communications on Pure and Applied Mathematics*, 63(1):1–38, 2010.
- [9] J. Galpin and D. Hawkins. Methods of L1 estimation of a covariance matrix. *Computational Statistics and Data Analysis*, 5(4):305–319, September 1987.

- [10] P. Geladi and B. R. Kowalski. Partial least-squares regression: a tutorial. *Analytica Chimica Acta*, 185:1–17, 1986.
- [11] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, 2013.
- [12] I. Jolliffe. *Principal Component Analysis*. Springer, 2002.
- [13] M. Jorgensen. *Iteratively Reweighted Least Squares*. John Wiley & Sons, Ltd, 2006.
- [14] Q. Ke and T. Kanade. Robust L1 norm factorization in the presence of outliers and missing data by alternative convex programming. In *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 739–746, June 2005.
- [15] N. Kwak. Principal component analysis based on L1-norm maximization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(9):1672–1680, September 2008.
- [16] P. D. Lax. *Linear algebra and its applications*. John Wiley & Sons, 2007.
- [17] G. Li and Z. Chen. Projection-pursuit approach to robust dispersion matrices and principal components: primary theory and monte carlo. *Journal of the American Statistical Association*, 80(391):759–766, 1985.
- [18] P. P. Markopoulos, G. N. Karystinos, and D. A. Pados. Optimal algorithms for L1-subspace signal processing. *IEEE Transactions on Signal Processing*, 62(19):5046–5058, October 2014.
- [19] M. McCoy and J. A. Tropp. Two proposals for robust PCA using semidefinite programming. *Electronic Journal of Statistics*, 5:1123–1160, 2011.
- [20] F. Nie, H. Huang, C. Ding, D. Luo, and H. Wang. Robust principal component analysis with non-greedy L1-norm maximization. In *Proceeding of 22nd International Conference on Artificial Intelligence*, pages 1433–1438, June 2011.
- [21] S. Roweis. EM Algorithms for PCA and SPCA. In *Advances in Neural Information Processing Systems*, pages 626–632, 1998.
- [22] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014.
- [23] Y. Shmueli, G. Wolf, and A. Averbuch. Updating kernel methods in spectral decomposition by affinity perturbations. *Linear Algebra and its Applications*, 437:1356–1365, 2012.
- [24] Z. Wen and W. Yin. A feasible method for optimization with orthogonality constraints. *Mathematical Programming*, 142(1-2):397–434, 2013.