



Fast and Accurate Proper Orthogonal Decomposition using Efficient Sampling and Iterative Techniques for Singular Value Decomposition

V. CHARUMATHI, M. RAMAKRISHNA, and VINITA VASUDEVAN,

Indian Institute of Technology

In this article, we propose a computationally efficient iterative algorithm for proper orthogonal decomposition (POD) using random sampling based techniques. In this algorithm, additional rows and columns are sampled and a merging technique is used to update the dominant POD modes in each iteration. We derive bounds for the spectral norm of the error introduced by a series of merging operations. We use an existing theorem to get an approximate measure of the quality of subspaces obtained on convergence of the iteration. Results on various datasets indicate that the POD modes and/or the subspaces are approximated with excellent accuracy with a significant runtime improvement over computing the truncated SVD. We also propose a method to compute the POD modes of large matrices that do not fit in the RAM using this iterative sampling and merging algorithms.

CCS Concepts: • **Mathematics of computing** → **Computations on matrices**; • **Computing methodologies** → **Principal component analysis**;

Additional Key Words and Phrases: Iterative and incremental sampling, fast and accurate POD

ACM Reference format:

V. Charumathi, M. Ramakrishna, and Vinita Vasudevan. 2022. Fast and Accurate Proper Orthogonal Decomposition using Efficient Sampling and Iterative Techniques for Singular Value Decomposition. *ACM Trans. Math. Softw.* 48, 2, Article 16 (May 2022), 24 pages.
<https://doi.org/10.1145/3506691>

1 INTRODUCTION

Proper orthogonal decomposition (POD), also known as **Principal component analysis (PCA)** or Karhunen–Loeve transform is a dimensionality reduction technique used in a variety of applications including information compression, pattern mining, clustering and classification, facial recognition, missing data estimation, data-driven modeling, and Galerkin projections. PCA aims at finding an optimal set of basis vectors that captures energetically dominant features of the dataset. These basis vectors are also referred to as eigenfeatures or POD modes or principal directions/components (PCs). The assumption is that a few of these modes are sufficient to capture all significant features of the dataset i.e., the data is highly correlated and inherently low rank. In

Authors' Address: V. Charumathi and V. Vasudevan, Department of Aerospace Engineering Indian Institute of Technology Madras Chennai, Tamilnadu, India 600 036; emails: charumathivasudevan89@gmail.com, vinita@ee.iitm.ac.in; M. Ramakrishna, Department of Aerospace Engineering Indian Institute of Technology Madras Chennai, Tamilnadu, India 600 036; email: krishna@ae.iitm.ac.in.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

0098-3500/2022/05-ART16 \$15.00

<https://doi.org/10.1145/3506691>

many of these cases, the dataset is also large, often distributed over several computers and dense. The two methods commonly used to compute POD modes are via (a) eigendecomposition of $A^T A$ and (b) SVD of A . In both cases, a full decomposition is expensive and approximate methods are used to obtain the dominant modes. In this article, we explore the use of random sampling based techniques to obtain the POD modes accurately and in a computationally efficient manner. The way our data is arranged, these modes are the left singular vectors of the data matrix.

We use the following notation throughout the article, unless otherwise specified. Lower case letters are used to denote scalars. Lower case bold-face letters are used to denote a vector. Upper case letters are used to denote a matrix. The i th element of vector \mathbf{a} (bold font) is denoted as a_i (non-bold font). The element in the i th row and j th column of the matrix A is denoted by a_{ij} . \mathbf{a}^i denotes the i th row vector and \mathbf{a}_j denotes the j th column vector of matrix A . A_k denotes the best rank- k approximation of a matrix A and the approximation to A_k is denoted using tilde, as \tilde{A}_k . Let $A_k = U_k \Sigma_k V_k^T$ and $\tilde{A}_k = \tilde{U}_k \tilde{\Sigma}_k \tilde{V}_k^T$ be the SVDs of A_k and \tilde{A}_k . If A is obtained by mean centering rows of the data matrix, then U_k and \tilde{U}_k are the accurate and the approximate principal components (PCs or POD modes), respectively. $\|A\|_p$ denotes the p norm of matrix A . A is assumed to be an $m \times n$ matrix.

Algorithms that approximate A_k can be classified as random projection, random sampling based algorithms or a combination of both. Random projection based algorithms have a multiplicative error bound [Halko et al. 2011a; Sarlos 2006]. They have been used for PCA in Erichson et al. [2017], Halko et al. [2011b], Nguyen et al. [2009], and Rokhlin et al. [2010] with algorithmic optimizations to improve accuracy and numerical stability [Erichson et al. 2019; Gu 2015; Li et al. 2017; Rokhlin et al. 2010] and with out-of-core modifications in Bose et al. [2019] and Halko et al. [2011b] for large matrices. Yamazaki et al. [2017] use random projections to compute the SVD and random sampling (of rows) of additional data to incrementally update the SVD. Drineas and Mahoney [2016] also use a combination of projection and sampling techniques, where projection is used to precondition the matrix so that subsequent uniform sampling works well.

Random sampling methods are of two types: those that sample elements of the matrix and those that sample columns/rows of the matrix. The motivation behind sampling elements [Achlioptas and McSherry 2007] from the matrix is to make the matrix sparser, thereby reducing the time taken to compute SVD of the matrix. Sampling columns/rows is usually done using (a) uniform probability [Drineas et al. 2003; Yamazaki et al. 2017] (b) leverage scores [Cohen et al. 2015; Drineas et al. 2008; Li et al. 2013; Yamazaki et al. 2017] (c) column and/or row norms [Drineas et al. 2003, 2004, 2006; Frieze et al. 2004]. In all three cases, error bounds are derived assuming sampling is done with replacement. Uniform and row/column norm based sampling have additive error bounds, while leverage score based sampling has a multiplicative error bound. However, leverage score based sampling is computationally more involved since it requires an estimate of the singular vectors. While row/column norm based sampling require a maximum of two passes over the data to get the sampling probability distribution, uniform sampling requires none. But it can be inefficient as it is equally likely to sample unimportant columns or rows, as for example, zero columns/rows.

The results of the experimental evaluation done by Menon and Elkan [2011] indicates the column norm sampling method has low run times, with accuracies well below the error bounds for standard datasets. Additionally, random sampling of columns/rows has several advantages—it retains the sparsity of the matrix and is easily amenable to incremental improvement. Multiplicative error bounds using row/column norm based sampling can be obtained using adaptive sampling, which has several rounds of sampling based on row/column norms [Deshpande et al. 2006; Deshpande and Vempala 2006].

In this article, our focus is to obtain POD modes of large datasets accurately and in a computationally efficient manner using random sampling based techniques. We evaluated the

performance of sampling algorithms proposed in Drineas et al. [2006]. Even with tight error parameters, the top k POD modes computed using these sampling algorithms as well as the subspace spanned by these vectors have significant errors. To improve accuracy, we propose an iterative technique in which additional rows and columns are sampled and the dominant POD modes are updated using a **merge-and truncate (MAT)** operation [Vasudevan and Ramakrishna 2017] in each iteration. The iteration converges when the angle between the current and updated modes is less than a threshold. We also derive bounds for the spectral norm of the error introduced by a series of MAT operations. We use an existing theorem to get an approximate measure of the quality of subspaces obtained after convergence. Results on various datasets indicate that the POD modes and/or the subspaces are approximated with excellent accuracy with a significant runtime improvement over computing the truncated SVD. Using this iterative algorithm, we also demonstrate a method to compute POD modes of large matrices that do not fit in the RAM of the system.

The article is organized as follows. In Section 2, we discuss a few relevant sampling algorithms found in literature and their drawbacks. This motivates the need for an algorithm with more efficiency and better accuracy. Section 3 contains the proposed iterative algorithm, the error bound for MAT operations and the quality measure. In Section 4, we discuss the results of these algorithms as applied to different datasets. Section 5 concludes the article. Appendix A contains the experimental results for the datasets obtained using row/column norm based sampling algorithms (LTSVD, CTSVD) proposed in Drineas et al. [2006].

The following datasets were used for testing all the algorithms.

- (1) Faces data from Cambridge [2002], a collection of 400 images from 40 subjects, each of size 112×92 giving 400 images of dimension 10,304 each. The matrix is of size $10,304 \times 400$. This dataset will be referred to as Faces dataset in the article.
- (2) Velocity vector data of a flow generated using a CFD simulation on a grid of size 257×257 . The snapshot vectors are of dimension 132,098 and there are 1,024 of them giving a matrix of size $132,098 \times 1,024$. It will be referred as V_{2D} .
- (3) 2414 images of 38 subjects cropped to show only the face of the subject in different illumination conditions from Lee et al. [2005]. The images are of size 192×168 giving 2414 images of dimension 32,256 each. Therefore, the matrix size is $32,256 \times 2,414$. This dataset is referred to as **cropped Yale faces (CYF)**.
- (4) Faces data from Georgiades et al. [2001], a collection of 16,128 images from 28 subjects in different poses and illumination conditions. 18 of the 16,128 images in the database were discarded since they were corrupted. Each image is of size 480×640 . We have 16,110 images of dimension 307,200 each organized as a $307,200 \times 16,110$ matrix. This dataset is referred to as **Yale faces (YF)**.

All the datasets are mean-centered row-wise. The principal components are the left singular vectors. It is seen that all the datasets have more rows than columns ("tall and thin"). The first three datasets are small and used for validation of the algorithms. The last dataset is large and the mean-centered data requires 40GB of RAM.

2 EVALUATION OF EXISTING SAMPLING ALGORITHMS AND MOTIVATION

Algorithms to compute low rank approximations of a matrix using sampled columns (LTSVD)/rows and columns (CTSVD) have been proposed in Drineas et al. [2006]. The sampling is done with replacement and the sampled rows and columns are scaled to preserve the Frobenius norm of the matrix. For both algorithms, the authors show that bound for the error in the rank- k approximation

$$\|A - \tilde{U}_k \tilde{U}_k^T A\|_F^2 \leq \|A - A_k\|_F^2 + \epsilon \|A\|_F^2, \quad (1)$$

holds with a probability of at least $(1 - \delta)$. The sampling probability of the i th column, p_i and the number of columns sampled, c , in LTSVD is

$$p_i = \frac{|\mathbf{a}_i|^2}{\|\mathbf{A}\|_F^2}, \quad (2)$$

$$c = 4k(1 + \sqrt{8 \log(1/\delta)})^2 / \epsilon^2. \quad (3)$$

In the case of CTSVD, the sampling probabilities for columns and rows (p_i and q_i) are

$$p_i = \frac{|\mathbf{a}_i|^2}{\|\mathbf{A}\|_F^2}, \quad q_j = \frac{|\mathbf{d}^j|^2}{\|\mathbf{D}\|_F^2}, \quad (4)$$

Here, \mathbf{D} is the matrix containing the sampled and scaled columns. The number of rows and columns sampled (c and w) that need to be sampled are

$$c = w = k^2(1 + \sqrt{\log(2/\delta)})^2 / \epsilon^4, \quad (5)$$

ϵ and δ are the error parameters and the failure probability that need to be set. An experimental evaluation of the algorithm by Menon and Elkan [2011] concluded that, in practice, the error is much lower than the bound given in Equation (1).

In this article, our focus is accurate computation of the principal components, which are the left singular vectors. Therefore, we implemented the two algorithms in order to evaluate its performance with respect to the error in the singular vectors. In order to evaluate the accuracy of the modes, we used the following metrics:

- (1) The angle between the accurate and approximate left singular vectors ($\theta_i = \cos^{-1}(\mathbf{u}_i^T \tilde{\mathbf{u}}_i)$), denoted by mode angle.
- (2) The principal angles between the rank- k subspaces (ϕ_i). This is computed as the inverse cosine of the singular values of the matrix $\mathbf{U}_k^T \tilde{\mathbf{U}}_k$ [Björck and Golub 1973].

The pseudo-codes for our implementation and the detailed results are included in Appendix A. A summary of the results is as follows.

- (1) When k (and hence number of columns sampled) is small, the error in the singular values can be as high as 20% for some datasets. For larger values of k , this error is of the order of 5% or less. (see Figure 8 in Appendix A).
- (2) There is a large error in the computed POD modes. As seen in Figure 9, if the mode angle is greater than about 10° , the POD modes have erroneous features. The mode angles obtained using these algorithms are as large as 80° in some cases (see Figure 10). As expected, the accuracy of the dominant subspaces is better, with principal angles between subspaces less than 40° in most cases (see Figure 11). This large error occurs because each column is scaled by a factor that depends on the column norm. While this preserves the Frobenius norm, it tends to distort the subspace spanned by the columns. There is no significant improvement in accuracy, even when the error parameters are tightened.

One major problem with the algorithms is that, it is hard to a priori fix the values of the error parameters ϵ and δ (see Equations (3) and (5)). Some of these issues are resolved by the adaptive sampling algorithm proposed in Deshpande and Vempala [2006]. This algorithm also has a tighter multiplicative error bound. It is based on approximate volume sampling, which essentially involves multiple rounds of row sampling. In each round, the probability of picking a row is proportional to the squared distance from the span of previously sampled rows.

The steps involved in the algorithm are as follows:

- (1) Sample k rows from A based on the approximate volume sampling (see [Deshpande and Vempala 2006] for details) and compute the basis, Q , spanning these rows.
- (2) Find the space orthogonal to the space spanned by these basis vectors, $E = A - AQQ^T$.
- (3) Sample s rows from A based on row norms of E and compute Q , the basis spanning all sampled rows.
- (4) Repeat from (2) for $(k + 1) \log_2(k + 1)$ times.

In each round $2k$ rows are sampled, except for the last round where $16k/\epsilon$ rows are sampled. For $k = 10$ and $\epsilon = 0.5$, we would need 38 rounds of sampling and 1060 rows will be sampled to find the basis. Once the iteration is over, Q is an approximation to the basis for the row space. To obtain the approximate left-singular vectors, we need to additionally find the SVD of AQ . For the V_{2D} dataset, we see that AQ is approximately the same size as A and it is larger for the Faces dataset (for which you anyway cannot get more than 400 basis vectors for the column space).

The advantage of this method is that (a) the error bound is multiplicative and (b) it directly approximates the dominant left and right subspaces. The sampled rows are not scaled, thus giving a better approximation to the dominant right singular vectors.

The algorithm can be made more efficient if (a) there is some pruning of basis vectors in Q in each iteration, depending on whether it belongs to the dominant subspace or not and (b) computation of $E = A - AQQ^T$ is avoided. Since the memory required by AQQ^T is the same as A , the step becomes inefficient when the matrix sizes are large.

3 PROPOSED ITERATIVE SAMPLING ALGORITHM (ISMA)

In our algorithm, for improved computational efficiency and accuracy, we would like to have several rounds of column or column and row sampling rather than just row sampling. More importantly, in each round, we would also like to prune the basis vectors that are not in the dominant subspace. We would like to stop the iteration when additional sampling does not improve the quality of either the modes or the subspace spanned by the modes. All sampling is done with replacement and duplicates are removed.

We use the following steps in our iteration.

- (1) Sample a set of columns from A based on column norms. Optionally, also sample rows. We use Equations (3) and (5) to estimate the number of columns and rows to be sampled, respectively. Compute SVD of the sampled matrix. Truncate to get the dominant POD modes.
- (2) From the remaining columns, sample an additional set of columns or columns and rows. For this, we explore various sampling strategies. These are discussed in Section 3.1. Find distinct columns (optionally, if rows are sampled, remove duplicates and scale remaining rows), compute POD modes of the newly sampled matrix and use it to update the dominant modes.
- (3) Find the cosines between the previously computed and newly updated dominant modes or the principal cosines between the subspaces. If any one of the cosines is less than the desired value, τ , continue the iteration.

Algorithm 1 details the steps involved in the iteration. The main iteration is the procedure **Iterative Sampling and Merging Algorithm (ISMA)**. It uses three other procedures: GET UPDATE, SAMPLE AND SCALE UNIQUE ROWS and BLOCK MERGE. The procedure GET UPDATE returns the POD modes of the newly sampled columns. In case rows are also sampled (the parameter *rows* = 1), GET UPDATE in turn calls the procedure SAMPLE AND SCALE UNIQUE ROWS that performs a row sampling, removes duplicates and scales the remaining rows. BLOCK MERGE is used to update the dominant subspace. It is described further in Section 3.2.

ALGORITHM 1: ISMA; Inputs: matrix A , sampling probability \mathbf{p} , rank of approximation k , number of columns to be sampled c , rank of matrices after BLOCK MERGE r , rows (=1 if rows are also to be sampled) and tolerance τ .

```

1: procedure ISMA( $A, \mathbf{p}, k, c, w, r, \tau, rows$ )
2:    $S \leftarrow \{1, 2, \dots, n\}$ 
3:    $\tilde{U}, \tilde{\Sigma}, S \leftarrow \text{GET\_UPDATE}(A, S, 0, \mathbf{p}, c, w, rows)$ ;
4:   if  $rows == 1$  then ▷ Orthonormalise  $U$  when there is row sampling
5:      $Q, R \leftarrow \text{QR}(\tilde{U}); U, \Sigma, V \leftarrow \text{SVD}(R); \tilde{U} = QU$ 
6:   end if
7:    $\xi_i \leftarrow 0$  where  $i = 1, 2, \dots, k$ 
8:   while  $\exists \xi_i : \xi_i < \tau$  and  $S$  not empty do
9:      $p_i \leftarrow 1/s$ ; or  $p_i \leftarrow \|a_i - \tilde{U}\tilde{U}^T a_i\|_2^2 / \sum \|a_i - \tilde{U}\tilde{U}^T a_i\|_F^2$ , where  $i \in S$ 
10:     $\tilde{U}, \tilde{\Sigma}, S \leftarrow \text{GET\_UPDATE}(A, S, \tilde{U}_k, \mathbf{p}, c, w, rows)$ ;
11:     $\tilde{U}, \tilde{\Sigma} \leftarrow \text{BLOCK\_MERGE}(\tilde{U}, \tilde{\Sigma}, \tilde{U}, \tilde{\Sigma}, r)$ 
12:     $\xi_i \leftarrow \tilde{u}_i^T \tilde{u}_i$  or  $X, \Xi, Y \leftarrow \text{SVD}(\tilde{U}_k^T \tilde{U}_k)$ ;  $\xi_i \leftarrow \Xi_{ii}$  where  $i = 1, 2, \dots, k$ 
13:     $\tilde{U} \leftarrow \tilde{U}; \tilde{\Sigma} \leftarrow \tilde{\Sigma}$ 
14:  end while
15:   $Q, R \leftarrow \text{QR}(A^T \tilde{U})$  ▷ for computation of error measure
16:   $U_R \Sigma_R V_R^T \leftarrow \text{SVD}(R)$ 
17:   $\tilde{U}, \tilde{\Sigma} \leftarrow \tilde{U} V_R, \Sigma_R$ 
18:   $\tilde{V} \leftarrow Q U_R$  ▷ only computed when error measure is needed
19:  return  $\tilde{u}_i, \tilde{\sigma}_i, \tilde{v}_i$  where  $i = 1, 2, \dots, k$ 
20: end procedure
21: procedure GET_UPDATE( $A, S, U, \mathbf{p}, c, w, rows$ )
22:    $I_s \leftarrow s_i : \Pr[s_i = \alpha] = p_\alpha \in \mathbf{p}, s_i \in S$ 
23:    $I_c \leftarrow s_i : \forall \text{ distinct } s_i \in I_s; g \leftarrow \text{card}(I_c); S \leftarrow S \setminus I_c$ 
24:    $d_i \leftarrow a_{\eta_i}$ , where  $\eta_i \in I_c$  and  $i = 1, 2, \dots, g$ 
25:   if  $rows \neq 1$  then
26:      $\tilde{V} \tilde{\Sigma}^2 \tilde{V}^T \leftarrow \text{SVD}(D^T D)$ ;  $\tilde{u}_i \leftarrow D \tilde{v}_i / \tilde{\sigma}_i, i \in 1, 2, \dots, r$ 
27:     return  $\tilde{U}, \tilde{\Sigma}, S$ 
28:   else
29:      $q_i \leftarrow \|d^i\|_2^2; q_i \leftarrow q_i / \sum_{i=1}^m q_i$  or if  $U \neq 0$   $q_i \leftarrow |u^i|_2^2 / k$ 
30:      $W \leftarrow \text{SAMPLE AND SCALE UNIQUE ROWS}(D, \mathbf{q}, w)$ 
31:      $\tilde{V} \tilde{\Sigma}^2 \tilde{V}^T \leftarrow \text{SVD}(W^T W)$ ;  $\tilde{u}_i \leftarrow D \tilde{v}_i / \tilde{\sigma}_i, i \in 1, 2, \dots, r$ 
32:     return  $\tilde{U}, \tilde{\Sigma}, S$ 
33:   end if
34: end procedure
35: procedure SAMPLE AND SCALE UNIQUE ROWS( $C, \mathbf{q}, w$ )
36:    $\hat{I}_s \leftarrow \hat{s}_j$  containing  $w$  row indices sampled with probability  $\mathbf{q}$ 
37:    $I_r \leftarrow \hat{s}_j; \hat{t} \leftarrow \{\hat{t}_j, \text{ number of occurrences of } \hat{s}_j\} : \forall \text{ distinct } \hat{s}_j \in \hat{I}_s$ 
38:    $h \leftarrow \text{card}(I_r)$ 
39:    $w^i \leftarrow c^{\xi_i} \sqrt{\hat{t}_i / (w q_{\xi_i})}$ , where  $\xi_i \in I_r$  and  $i = 1, 2, \dots, h$ 
40:   return  $W$ 
41: end procedure

```

The matrix D in GET UPDATE contains the (distinct) sampled columns and the matrix W in SAMPLE AND SCALE UNIQUE ROWS is obtained after additionally sampling rows. If there is no row sampling, GET UPDATE returns the r dominant POD modes, \tilde{U}_r , and singular values which are computed using a SVD of $D^T D$. Otherwise, the approximate left singular vectors are computed using the right singular vectors of W (line 31).

When using right singular vectors of W as an approximation to those of D , the modes obtained may not be orthonormal. Therefore, for the first approximation of the modes, we orthonormalize using QR and SVD of a smaller matrix (lines 4–6 of Algorithm 1). In the subsequent iterations, orthonormalization is implicit in the BLOCK MERGE algorithm.

The iteration stops when the POD modes computed in two successive iterations are similar. Similarity is measured using cosines between the modes or principal cosines between subspaces spanned by the modes. After the iterations converge, the quality of the approximated subspace can be measured using Wedin's theorem. It requires estimation of the singular values and right singular vectors (lines 15–18 in Algorithm 1). We describe this measure in Section 3.3.

3.1 Sampling Strategies

3.1.1 Column Sampling. In the first round, we do a column norm based sampling, with probabilities as given in Equation (2), to get an approximation of the dominant modes. Additional columns are sampled in subsequent iterations in an attempt to correct the modes. The sampling strategy for the additional columns is a compromise between the number of iterations and the computational effort required in each iteration. We explored three sampling strategies for adding more columns.

- (a) Uniform sampling (UNF) is easy to implement as no extra computation for probabilities is necessary. Larger number of iterations are needed as there is no way to avoid sampling of unimportant or highly correlated columns.
- (b) L_2 norm sampling (L2N) has only simple arithmetic manipulation of probabilities from the second iteration but may repeatedly sample columns that are similar or highly correlated.
- (c) ORT requires us to project the matrix onto the current estimate of the dominant subspace (\tilde{U}) and sample based on norm of the component orthogonal to this subspace. Probability of sampling columns is computed as,

$$p_i = |a_i - \tilde{U}\tilde{U}^T a_i|_2^2 / \sum |a_i - \tilde{U}\tilde{U}^T a_i|_2^2, \quad (6)$$

Note that, it avoids sampling of highly correlated components and can lead to fewer iterations. However, the computational and memory requirements in each iteration is large, making each iteration slow and sometimes impractical for large datasets.

3.1.2 Row Sampling. We have tried two strategies for row sampling:

- (a) L_2 norm of rows and
- (b) **leverage scores (LS).**

Sampling based on L_2 norm of the rows has been used in Drineas et al. [2006] as well as in the adaptive sampling algorithm. Both LS based sampling and uniform sampling have been used by Yamazaki et al. [2017], who propose an incremental algorithm for SVD. Their algorithm reduces the size of the matrix by projecting columns and sampling rows. Computing leverage scores requires an estimate of the left singular vectors. Therefore, it is not done in the first round of sampling.

When rows are also sampled, we use the right singular vectors and the singular values of W as an approximation to those of D . These are used to compute the POD modes (see line 31 of Algorithm (1)). For this approximation to be accurate, the primary requirement is that the singular

values computed using W be accurate. In order to maintain the accuracy of the singular values, we scale the sampled rows as specified in the procedure **SAMPLE AND SCALE UNIQUE ROWS**. We remove duplicate rows that were sampled and scale the remaining rows with $\sqrt{\hat{t}_i/wq_i}$ where w is the number of rows sampled, \hat{t}_i is the number of occurrences of each sampled row, and q_i denotes sampling probability. This scale factor makes the computed SVD equal to the SVD obtained without removing duplicates. This is similar to what is done in Sun et al. [2007] for sampling columns.

3.2 Merge and Truncate (MAT) Operation

Our algorithm requires an efficient technique for estimation of POD vectors in each iteration. Algorithms to estimate these vectors have been proposed in (Bai et al. [2005], Liang et al. [2014], and Qu et al. [2002]). Also, several techniques have been proposed for incremental SVD [Baker et al. 2012; Brand 2002, 2006; Gu and Eisenstat 1993], which can be used to get the principal components. We decided to use the algorithm in Vasudevan and Ramakrishna [2017] (which is a generalization of the algorithm in Iwen and Ong [2016]) for the following reasons.

- (a) We directly obtain the left singular vectors as the output of the operation.
- (b) These left singular vectors are obtained using SVD and QR of smaller matrices. Therefore they are orthogonal by construction.

Let A be a matrix of size $m \times n$ that consists of sub-matrices X and Y containing n_1 and n_2 columns, respectively, where $n = n_1 + n_2$. Assume that the rank- r approximations of X and Y are $U_{1r}\Sigma_{1r}V_{1r}^T$ and $U_{2r}\Sigma_{2r}V_{2r}^T$. The rank- r approximation of A , \tilde{A}_r , can be computed from the individual SVDs by first merging the two SVDs and then truncating it to a rank- r approximation as follows.

The component of U_{2r} orthogonal to U_{1r} is $U_t = U_{2r} - U_{1r}(U_{1r}^T U_{2r})$. If $U_t = U_o R$ is the corresponding QR decomposition, we have

$$\begin{bmatrix} X & Y \end{bmatrix} \approx \begin{bmatrix} U_{1r} & U_o \end{bmatrix} \begin{bmatrix} \Sigma_{1r} & (U_{1r}^T U_{2r})\Sigma_{2r} \\ \mathbf{0} & R\Sigma_{2r} \end{bmatrix} \begin{bmatrix} V_{1r}^T & \mathbf{0} \\ \mathbf{0} & V_{2r}^T \end{bmatrix} = \begin{bmatrix} U_{1r} & U_o \end{bmatrix} E \begin{bmatrix} V_{1r}^T & \mathbf{0} \\ \mathbf{0} & V_{2r}^T \end{bmatrix} \quad (7)$$

Now E is a much smaller $(2r) \times (2r)$ matrix. If $E = U_E \Sigma_E V_E^T$, we get

$$\begin{bmatrix} X & Y \end{bmatrix} \approx \begin{bmatrix} U_{1r} & U_o \end{bmatrix} U_E \Sigma_E V_E^T \begin{bmatrix} V_{1r}^T & \mathbf{0} \\ \mathbf{0} & V_{2r}^T \end{bmatrix} = U \Sigma V^T, \quad (8)$$

where $U = [U_{1r} \ U_o]U_E$, $\Sigma = \Sigma_E$ and $V = \begin{bmatrix} V_{1r} & \mathbf{0} \\ \mathbf{0} & V_{2r} \end{bmatrix}V_E$. These singular values and vectors are once again truncated to get \tilde{U}_r and $\tilde{\Sigma}_r$. Algorithm (2) details the steps involved. In general, if there are P partitions, a series of MAT operations can be used to obtain the r left singular vectors. Since we are only interested in the left singular vectors, we do not compute V .

The penalty we pay for getting a more accurate estimate of the dominant subspace is finding the SVD of the newly sampled columns to do a MAT operation in each iteration. However, the number of additional sampled columns in each iteration is much lower than the total number of columns in the matrix. Therefore, a speedup is obtained since the complexity of SVD computation depends quadratically on the number of columns for “tall and thin” matrices.¹ Assuming there are P partitions and all n columns are sampled with same number of columns sampled in each iteration, the total number of flops required for computing a rank- r approximation of A using a series of MAT operations is $\approx \frac{14mn^2}{P} + \frac{192n^3}{P^2}$ where $r \ll n$ [Vasudevan and Ramakrishna 2017].

¹For “short and fat” matrices, we need to iteratively sample rows to get a runtime improvement when run on a single machine.

ALGORITHM 2: Pseudocode for block merge algorithm: merges left/right singular vectors of two adjacent blocks. Inputs: $U_1 \in \mathbb{R}^{m_1 \times n_1}$, $U_2 \in \mathbb{R}^{m_1 \times n_2}$, and required rank r .

```

1: procedure BLOCK MERGE( $U_1, \Sigma_1, U_2, \Sigma_2, r$ )
2:    $U_{1r}, \Sigma_{1r}, U_{2r}, \Sigma_{2r} \leftarrow \text{Truncate}(U_1, \Sigma_1, U_2, \Sigma_2)$  retaining only  $r$  values
3:    $U_t \leftarrow U_{2r} - U_{1r}(U_{1r}^T U_{2r}); U_0, R \leftarrow \text{QR}(U_t)$ 
4:    $E \leftarrow \begin{bmatrix} \Sigma_{1r} & (U_{1r}^T U_{2r}) \Sigma_{2r} \\ \mathbf{0} & R \Sigma_{2r} \end{bmatrix}; U_E \Sigma_E V_E^T \leftarrow \text{SVD}(E)$ 
5:    $U \leftarrow \begin{bmatrix} U_{1r} & U_0 \end{bmatrix} U_E; \Sigma \leftarrow \Sigma_E; r \leftarrow \min(r, \max(i : \sigma_i \neq 0))$ 
6:   Truncate  $U$  and  $\Sigma$  retaining only  $r$  values
7:   return  $U, \Sigma$ 
8: end procedure

```

3.2.1 Error Analysis for Series of MAT Operations. We derive an error bound in terms of the spectral norm of the error in the rank- r matrix obtained using a series of MAT operations. For this analysis, we use the following well known results from literature [Li 2014; Mathias 2014]. Let X be an $m \times n$ matrix and $q = \min(m, n)$.

- (1) If $X = U\Sigma V^T$, and $X_r = U_r \Sigma_r V_r^T$ is a rank- r approximation to X obtained by retaining the top r singular values, then, $\|X - X_r\|_2 = \sigma_{r+1}(X)$ and $\|X - X_r\|_2 \leq \|X - B_r\|_2$, where B_r is any rank- r matrix.
- (2) If Y is an $m \times t$ sub-matrix of X , then, $\sigma_{i+n-t}(X) \leq \sigma_i(Y) \leq \sigma_i(X)$
- (3) If $\hat{X} = X + \Delta X$ is a perturbation of X , then $|\sigma_i - \hat{\sigma}_i| \leq \|\Delta X\|_2$, $i = 1, 2, \dots, q$.

We use the following notation. Let the matrix $X = [X_1 \ X_2 \ \dots \ X_P]$ be divided into P partitions. Define $Y_0 = [X_1]$, $Y_1 = [Y_0 \ X_2]$ and so on. In general $Y_j = [Y_{j-1} \ X_{j+1}]$, so that $Y_{P-1} = X$. Let X_r denote the optimal rank- r approximation of X , and X_{ir} that of X_i . Define $Z_1 = [X_{1r} \ X_{2r}]$ and $Z_j = [Z_{j-1r} \ X_{j+1r}]$. Let $Z = Z_{P-1}$. Therefore the error $E = \|X - Z_r\|_2$, where Z_r are the optimal rank- r approximations of Z .

The bound for the error can be derived as follows.

$$\begin{aligned}
\|X - Z_r\|_2 &\leq \|X - Z\|_2 + \|Z - Z_r\|_2 \\
&\leq \|X - Z\|_2 + \|Z - X_r\|_2 \\
&\leq \|X - Z\|_2 + \|Z - X\|_2 + \|X_r - X\|_2 \\
&= \|X_r - X\|_2 + 2\|X - Z\|_2 \\
&= \sigma_{r+1}(X) + 2\|X - Z\|_2.
\end{aligned} \tag{9}$$

We find $\|X - Z\|_2$ iteratively. For the MAT operation on the first two partitions, we have

$$\begin{aligned}
\|Z_1 - Y_1\|_2 &\leq \|X_1 - X_{1r}\|_2 + \|X_2 - X_{2r}\|_2 \\
&= \sigma_{r+1}(X_1) + \sigma_{r+1}(X_2) \leq 2\sigma_{r+1}(X).
\end{aligned} \tag{10}$$

For subsequent partitions,

$$\begin{aligned}
\|Z_j - Y_j\|_2 &\leq \|Z_{j-1r} - Y_{j-1}\|_2 + \|X_{j+1} - X_{j+1r}\|_2 \\
&\leq \|Z_{j-1r} - Z_{j-1}\|_2 + \|Z_{j-1} - Y_{j-1}\|_2 + \|X_{j+1} - X_{j+1r}\|_2.
\end{aligned} \tag{11}$$

If $Z_i = Y_i + \Delta X$, then $\sigma_{r+1}(Z_i) \leq \sigma_{r+1}(Y_i) + \|Z_i - Y_i\|_2$. Therefore,

$$\begin{aligned}
\|Z_j - Y_j\|_2 &\leq \sigma_{r+1}(Y_{j-1}) + 2\|Z_{j-1} - Y_{j-1}\|_2 + \sigma_{r+1}(X_{j+1}) \\
&\leq \sigma_{r+1}(X) + 2\|Z_{j-1} - Y_{j-1}\|_2 + \sigma_{r+1}(X).
\end{aligned} \tag{12}$$

The last inequality follows since Y_{j-1} and X_{j+1} are submatrices of X . Since $\|X - Z\|_2 = \|Y_{P-1} - Z_{P-1}\|_2$, it can be computed iteratively using Equation (12). Therefore, the bound for error, $\|X - Z_r\|_2$, is $(2^{P+1} - 3)\sigma_{r+1}(X)$.² We note the following:

- (1) In the worst case, $\sigma_{r+1}(X_i)$ is close to $\sigma_{r+1}(X)$ and the error is close to $(2^{P+1} - 3)\sigma_{r+1}(X)$. But this typically happens when the size of the submatrix is large, which means P is small.
- (2) If there are many partitions, the number of MAT operations are large, but each $\sigma_{r+1}(X_i)$ is likely to be small. If Y is an $m \times t$ submatrix of X , then $\sigma_{r+1+n-t}(X) \leq \sigma_{r+1}(Y) \leq \sigma_{r+1}(X)$. Since each submatrix is thin and t is small, the lower bound is quite small. Hence the error in the average case is likely to be much smaller than in the worst case. In practice, we have found this is true.

Overall, if we need k POD vectors to be approximated well, then either the number of partitions should be small or we can set r to be larger than k . We set r to $3k$ to get good accuracies.

3.3 Quality of Converged Solution

In order to estimate the accuracy of the computed POD modes, we need a measure to quantify the error in the approximation. After convergence, we can get a bound on the error in the subspaces using Wedin's theorem [Wedin 1972]. Assume that at the end of the iteration, we get $\tilde{A}_r = \tilde{U}_r \tilde{\Sigma}_r \tilde{V}_r^T$, out of which we take the top k singular values and vectors. Define $R = A\tilde{V}_k - \tilde{U}_k \tilde{\Sigma}_k$ and $S = A^T \tilde{U}_k - \tilde{V}_k \tilde{\Sigma}_k$. Assume $\sin(\Theta)$ and $\sin(\Phi)$ are diagonal matrices containing the sine of principal angles between subspaces spanned by U_k and \tilde{U}_k and V_k and \tilde{V}_k respectively. The sines of the principal angles are a measure of the distance between the accurate and approximate dominant left and right subspaces.

Wedin's theorem [Li 1996; Stewart and Guang Sun 1990; Wedin 1972] gives a bound on the norms of the sine of principal angles in terms of the norms of R and S . The statement of the theorem is as follows:

$$\sqrt{\|\sin(\Theta)\|_F^2 + \|\sin(\Phi)\|_F^2} \leq \frac{\sqrt{\|R\|_F^2 + \|S\|_F^2}}{\omega}, \quad (13)$$

where $\omega \triangleq \min\{\min_{j=1, \dots, n-k} |\tilde{\sigma}_k - \sigma_{k+j}|, \tilde{\sigma}_k\} > 0$ and σ_i are the singular values of A . In order to compute ω , we need all the singular values of A , which are not available. However, at the end of the iteration, we have an approximation of $r (= 3k)$ singular values. Assuming that the k th and $(k+1)$ th singular values are approximated well (which was the case in all the datasets we have looked at), we estimate ω as $\min\{|\tilde{\sigma}_k - \tilde{\sigma}_{k+1}|, \tilde{\sigma}_k\}$. If our iteration converges to the correct subspaces, the right hand side should be close to zero. If it approaches $\sqrt{2k}$, either (a) additional columns need to be sampled or (b) the k th and $(k+1)$ th singular values are very close to each other and small errors in the singular values result in large values for the bound. In the limit when they are identical, any vector in the subspace is a POD mode and the bound can become arbitrary. In this case, a better measure of the quality of the solutions can be obtained after increasing k . Usually increasing k by one or two is sufficient.

Our main focus is the POD modes. However, in order to compute the bound, the singular values and right singular vectors are also required. Since we do not scale columns and use some subset of the columns for computation of the modes, the error in the singular values (not the vectors)

²The trend for the error is similar in the expression derived in Iwen and Ong [2016] (see Theorem 3 in the reference). They find a bound on the Frobenius norm of the error which grows as $(1 + \sqrt{2})^P$, which is slightly worse than ours. We used the spectral norm as it is more appropriate for noisy data as argued in Li et al. [2017]. Our derivation is also simpler and more specific to the case we are looking at, namely, an iterative algorithm which has a series of MAT operations.

Table 1. Average Time Taken (in Seconds) by SVD(A) and SVD- $A^T A$ for Different Datasets

Dataset	SVD(A)		SVD- $A^T A$	
	single thread	4 threads	single thread	4 threads
V_{2D}	25.28	13.94	8.39	2.75
Faces	0.35	0.17	0.13	0.069
Cropped Yale faces (CYF)	34.04	16.43	15.23	6.86
Yale faces (YF)	-	-	-	-

Note that SVD of YF could not be computed in our test system (32 GB) since it does not fit in the system's memory.

could be significant. This is especially true if the total number of columns sampled is much less than the number of columns in the matrix. The left singular vectors will be quite accurate if the columns corresponding to the dominant subspace are captured, but the singular values will not be accurate. If this is the case, we use the technique followed in the adaptive sampling method i.e., after the modes have converged, the singular values and the right singular vectors can be obtained by computing the SVD of $\tilde{U}_k^T A$ as detailed in lines 15–18 in Algorithm (1).

4 RESULTS

All algorithms are run in a 4 core Intel® Core™ i7-6700K processor that runs at a maximum clock frequency of 4GHz with hyper-threading on. The system has a 32 GB RAM. For the YF dataset, the mean centered data required 40 GB memory. Therefore for this dataset, the runs for the iterative algorithm were done in a different system, a 8 core Intel® Xeon® CPU E5-2650V2 processor that runs at a maximum clock frequency of 2.6GHz with hyper-threading on, with 64GB of RAM.

The algorithms are written in Python and run in Python version 3.5.3. It uses *numpy* version 1.12.1-3, *scipy* version 0.18.1-2 and *openBLAS* version 0.2.19-3, which is multi-threaded. Double precision was used for all computations. We used the linear algebra routines in Scipy, which are based on LAPACK for SVD computation. The random number generator in *numpy* is used. Runtime and speedup values are an average of five runs. We report speed-up obtained when *openBLAS* is run using a single thread, to give a measure of the operation count in each algorithm and four threads (since our system is a four core system), which gives an indication of how the runtime scales with number of threads. We have not explicitly parallelized our algorithm and any improvement with multiple threads is entirely due to *openBLAS*.

The accurate POD modes of A can be computed using SVD($A^T A$) as follows, (a) find right singular vectors and singular values of $A^T A$ using either SVD or eigendecomposition, (b) compute the k left singular values as $\mathbf{u}_i = A\mathbf{v}_i/\sigma_i$. We denote the algorithm as SVD- $A^T A$.

If A is a tall and skinny matrix, computing SVD(A) using SVD- $A^T A$ is faster. This is because SVD(A) takes approximately βmn^2 floating point operations (FLOPs). SVD- $A^T A$ requires mn^2 FLOPs to compute $A^T A$, $(\beta + 16)n^3$ to compute SVD($A^T A$) and another mn^2 operations to compute \mathbf{u}_i from $A\mathbf{v}_i/\sigma_i$. Assuming $\beta = 6$ [Björck 2015; Golub and Van Loan 2007], this gives $m/(m/3 + (1 + 8/3)n)$ speedup. The results in Table 1 are close to what is expected. Since SVD- $A^T A$ is faster, we benchmark the speedup of the algorithms with respect to time taken for computing k POD modes using SVD- $A^T A$.

4.1 Accuracy and Runtime of the Iterative Algorithms

We evaluated the two versions of ISMA algorithm (Algorithm (1))—an **iterative column sampling (ICS)** algorithm and an **iterative column and row sampling (ICRS)** algorithm. We tested the performance of the iterative algorithms with two different convergence criteria: cosines

Table 2. Sampling Strategies used for Sampling Rows and Columns by Iterative Sampling Algorithms

strategy	column sampling	row sampling (when applicable)
L2N	L_2 norm sampling of columns in all iterations	–
UNF	uniform sampling from second iteration	L_2 norm sampling
ORT	from second iteration, sampling based on orthogonal complement of A projected onto current approximation of modes	L_2 norm sampling
LS	uniform sampling from second iteration	leverage score sampling using current approximation of modes from second iteration

For all sampling strategies, in the first iteration, columns (and rows in the case of column and row sampling) are sampled based on their respective L_2 norms.

Table 3. Parameter Values (k, ϵ, δ) used in the Iterative Algorithms for Various Datasets

Case	V_{2D} , CYF k, ϵ, δ	Faces k, ϵ, δ	YF k, ϵ, δ
I1	10, 0.7, 0.6	10, 1, 0.75	20, 0.6, 0.6
I2	2, 0.7, 0.6	2, 0.7, 0.6	5, 0.6, 0.6

between the modes obtained in successive iterations and the principal cosines between the subspaces computed in successive iterations. As indicated previously, we tried different sampling strategies from the second iteration. Table 2 lists the different sampling strategies we use in this article.

Using ϵ and δ around 0.6 – 0.7 resulted in a reasonable tradeoff between the number of columns (and rows) sampled in each iteration and the number of iterations. For the Faces dataset, we chose a higher ϵ, δ since the number of columns in the matrix is quite small. ORT was not used in the iterative algorithms for YF, since the overhead for computing the orthogonal complement in each iteration proved to be large. Table 3 lists the parameter values used for the iterative algorithms. In all cases, the tolerance parameter for convergence τ , is set to 0.99.

Figure 1 shows the first k mode angles computed by the iterative algorithms using L2N, UNF, ORT, and LS. This is run with convergence of individual modes. As expected, the accuracies are significantly better than for a single round of sampling. We also ran the algorithm with convergence of the subspace spanned by modes as the criterion with the same value for τ . Figure 2 shows the principal angles between the subspaces spanned by the approximate modes and the modes obtained using the truncated SVD. The algorithms give similar results as in Figure 1 for all datasets. Moreover, it is clear that the accuracies are similar for all the sampling strategies, indicating it is sufficient to use uniform distribution after the first round of sampling. This is because UNF is the least complex and computationally more efficient than the other column sampling strategies.

To evaluate the quality of subspaces obtained, we computed the measure for error in the subspaces using Equation (13). Table 4(a) shows the error measure for various cases. In general, the error measure is small indicating good accuracies in the modes. In a few cases (V_{2D} ($k = 2$) and YF ($k = 20$)), the measure is large. As discussed, this could happen if σ_k and σ_{k+1} are clustered




Fig. 1. First k mode angles, θ_i , of different datasets computed by iterative algorithms based on convergence of individual modes. ICS: Column sampling, ICRS: Column and row sampling.




Fig. 2. First k principal angles, ϕ_i between the accurate subspace and the subspace obtained using the iterative algorithm. ICS: Column sampling, ICRS: Column and row sampling.

together. Small changes in the singular values can result in large variations in the measure (as seen with L2N and UNF). Figure 7 in Appendix A shows the dominant singular values of all our test datasets. It can be seen that σ_2 and σ_3 of V_{2D} are very close in value, whereas σ_4 is well separated. Therefore when k is increased to 3, the error measure reduces significantly. This can be seen in Table 4(b).

The other case for which the measure is large is for YF (UNF). In this case also, several singular values are clustered together as seen in Figure 7. The error drops when k is increased as seen in Table 4(b). If L2N is used for sampling, the measure is small. It turns out that all the columns were sampled in this case, leading to a low error. In general, it can be seen from Figures 1 and 2 that the accuracy of the modes are much better than indicated by the measure. Also note that, in most cases, our approximation of ω worked well. The average error in ω was found to be around 6%.

Figures 3 and 4 contain the speedup obtained when the iterative algorithms are run with convergence of individual modes and subspaces formed by the modes as the convergence criterion, respectively. Only the smaller datasets are included. YF is discussed later. Speedup is evaluated with respect to time taken for computing k POD modes using $SVD-A^T A$. We note that (a) very good speedup is obtained for CYF for both single and multi-threaded execution. The speedup is the least when ORT is used as the sampling strategy, which is as expected. Since the accuracies

Table 4. Error measure computed using POD vectors obtained from Iterative algorithms with convergence of individual modes.

(a) $\sqrt{\|R\|_F^2 + \|S\|_F^2} / \hat{\omega}$ where $\hat{\omega} \triangleq \min \{|\tilde{\sigma}_k - \tilde{\sigma}_{k+1}|, \tilde{\sigma}_k\}$.

k	Dataset	ICS				ICRS		$\sqrt{2k}$	$\frac{\sigma_k - \sigma_{k+1}}{\sigma_k}$
		L2N	UNF	ORT	UNF	ORT	LS		
10	faces	0.275	0.252	0.282	0.657	0.677	0.62	4.47	0.107
10	CYF	0.731	0.567	0.665	0.868	0.943	0.748	4.47	0.024
10	V_{2D}	0.119	0.123	0.096	0.158	0.132	0.154	4.47	0.11
20	YF	0.428	5.053	-	6.48	-	0.347	5.48	0.026
2	faces	0.191	0.23	0.22	0.081	0.25	0.268	2	0.272
2	CYF	0.07	0.055	0.083	0.055	0.097	0.051	2	0.624
2	V_{2D}	0.762	1.95	2.88	1.278	5.34	1.115	2	0.022
5	YF	2.08	1.71	-	1.394	-	1.43	3.16	0.073

(b) Error measure with value k increased by one for V_{2D} and YF.

k	Dataset	ICS			ICRS		
		L2N	UNF	ORT	UNF	ORT	LS
3	V_{2D}	0.237	0.32	0.42	0.356	0.523	0.453
21	YF	0.207	0.88	-	0.88	-	1.46

Values in bold are either greater than or close to the upper limit of $\sqrt{2k}$.

of all schemes are similar, it makes sense to use the UNF strategy as it gives maximum speedup. (b) For small datasets such as the FACES dataset, it is better to use SVD- $A^T A$. (c) V_{2D} has about 1,000 columns. The results show that it may be worthwhile to use the iterative algorithm if k is small.

We found that for $k = 10$ (I1), nearly all columns were sampled for the Faces, V_{2D} and CYF datasets when convergence was achieved, both in terms of mode and principal angles. For I2 ($k = 2$), the number of columns sampled was substantially lower than n in all cases and the speedup is much larger. For the V_{2D} and Faces datasets, the speedup is almost the same for both convergence criteria. However, for CYF, speed up obtained is substantially larger with convergence of subspace for $k = 2$. This is because the first two singular values are almost identical. Therefore, more columns are sampled to capture individual modes accurately than for subspace spanned by the modes.

As expected, in comparison to other sampling strategies, ORT is slower. UNF, LS, and L2N have very similar speed up for most cases. This reinforces our conclusion that sampling with uniform sampling from the second iteration is as good as other sampling strategies we tried.

As indicated earlier, mean centered dataset for YF occupies 40GB memory. We implemented Algorithm 1 in a system with 64 GB RAM and compared the results obtained using UNF with (a) SVD($A^T A$) and (b) projection algorithms proposed in Erichson et al. [2017], Halko et al. [2011a], and Rokhlin et al. [2010]. We implemented two versions of projection algorithms: one in which the rows are projected onto the random matrix (RSVD) [Halko et al. 2011a] and the other in which the columns are projected (COLRSVD) [Erichson et al. 2017; Rokhlin et al. 2010]. We report results for RSVD and COLRSVD for the YF dataset with three subspace iterations and $s = k + 5$ as suggested in Halko et al. [2011a]. The wall clock and computational times taken for YF dataset using these

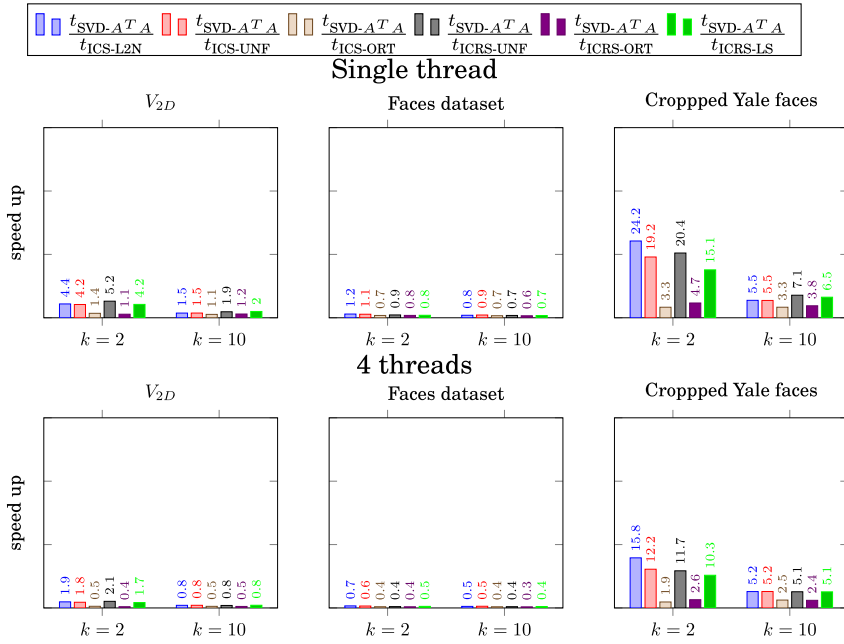


Fig. 3. Speed up for iterative algorithms based on convergence of individual modes. ICS: Column sampling, ICRS: Column and row sampling.

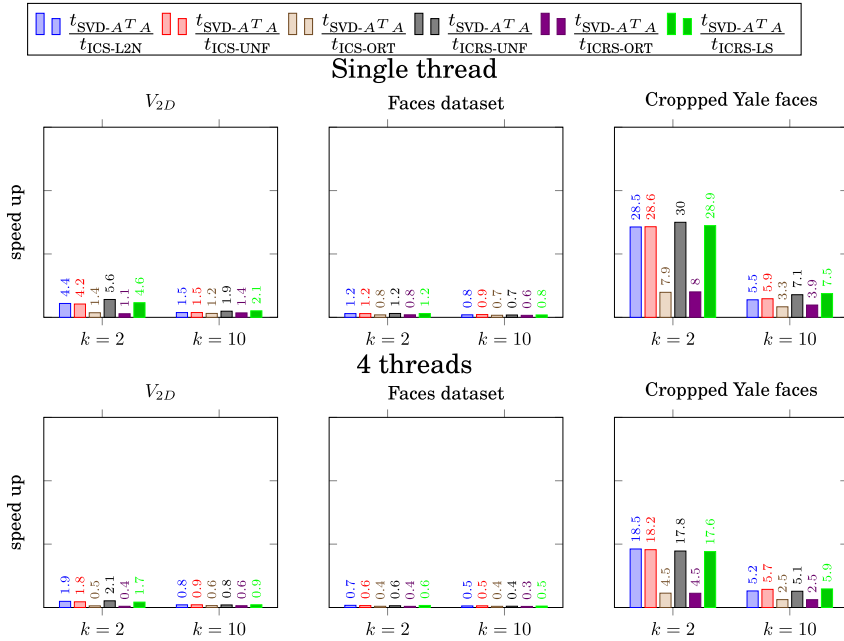


Fig. 4. Speed up for iterative algorithms based on convergence of subspaces formed by modes. ICS: Column sampling, ICRS: Column and row sampling.

Table 5. Wall Clock (WT) and Computational Times (CT) in Seconds (s) for Projection and Iterative Sampling Algorithms for YF Dataset

Case	Algorithm	single thread		8 threads	
		WT(s)	CT(s)	WT(s)	CT(s)
$k = 5$ (I2)	ICS-UNF	217.79	50.78	224.062	38.57
	ICRS-UNF	205.71	42.11	222.226	41.73
	RSVD	298.75	108.43	283.5	108.8
	COLRSVD	286.44	109.83	281.78	108.97
	SVD- $A^T A$	7151.07	6964.7	2149.77	1962.77
$k = 20$ (I1)	ICS-UNF	626.18	445.34	299.05	130.59
	ICRS-UNF	379.16	198.19	278.74	111
	RSVD	335.12	160.77	290.7	159.32
	COLRSVD	346.89	159.98	290.18	159.48
	SVD- $A^T A$	7403.33	7216.96	2171.17	1984.17

Algorithms were run in a 64 GB system since YF does not fit in our test system.

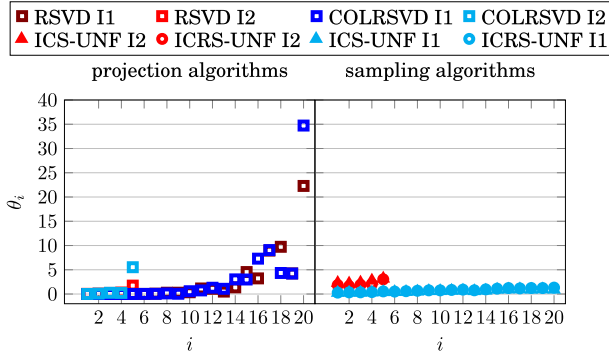


Fig. 5. First k mode angles, θ_i , computed by projection algorithms with three subspace iterations and iterative sampling algorithms using sampling strategy UNF.

algorithms is shown in Table 5. The wall clock time includes the time required to read the data from the disk. Figure 5 shows the mode angles obtained. We can see that for Yale faces dataset with $k = 5$, the projection algorithms and sampling algorithms have similar accuracy. Sampling algorithms are $1.4\times$ faster with respect to wall clock times and $2\times$ faster with respect to computational times. For $k = 20$, the wall clock times of projection and sampling algorithms are similar (except for ICS-UNF) but the projection algorithms have a large error in comparison to the iterative sampling algorithms. Also note that, the iterative sampling algorithms are much faster (7 – 30 \times) than truncated SVD.

4.2 Incremental-Iterative Algorithm for Matrices that do Not Fit in the RAM

As we already specified before, the mean-centered data matrix for YF dataset requires 40 GB memory and does not fit in a system that has 32 GB RAM. There are two possible ways to compute the POD modes of this dataset in this system:

- (1) Run the iterative algorithm (Algorithm (1)). In the first pass, data is read in chunks that fit in the memory to compute the sampling probabilities. In the next pass, the data is read again to form the sampled matrix. Assuming we have enough memory for the sampled matrix




Fig. 6. Left: mode angles of YF computed using convergence of individual modes. Right: principal angles computed using convergence of subspace formed by modes. Results are of the incremental sampling algorithms applied on YF dataset partitioned into four. Parameters are set as given in Table 3.

and its left singular vectors in the RAM, \tilde{U} and $\tilde{\Sigma}$ of the matrix are computed. Subsequent iterations may require one or both passes on data depending on the sampling strategies used. The iteration is continued until Algorithm (1) converges.

- (2) Do an incremental computation by running the iterative algorithm on each block of the partitioned data as follows. Load the block into memory and use ICS/ICRS to compute \tilde{U} , $\tilde{\Sigma}$ of the block. Use the MAT algorithm to get an updated \tilde{U} , $\tilde{\Sigma}$ of the dataset. Delete the partition and the sampled columns and continue the process until all blocks of data have been processed. Algorithm (3) details the steps involved.

The first alternative requires $i + 1$ passes of the matrix for ICS where i is the number of iterations for convergence. For ICRS, it takes a maximum of $3i + 1$ passes for i iterations. By passes, we mean accessing the entire matrix, though it may be done in blocks. The second alternative requires only one pass over the entire matrix irrespective of whether ICS or ICRS is used to find \tilde{U} , $\tilde{\Sigma}$ of the partition. Clearly, the first alternative is severely bottle-necked by disk access times. For YF, we used the second option and used an incremental computation after partitioning the matrix into four blocks. We ran the iterative algorithms on each block using the same parameters as given for YF in Table 3.

Figure 6 shows that both mode angles and the principal angles are less than 5° . Table 6 shows the wall clock time and computational time for ICS-UNF, ICRS-UNF with convergence of (a) modes and (b) subspace spanned by modes. It can be seen that for both convergence criteria, time taken is very similar. As expected, ICRS is faster than ICS. Therefore, we favor ICRS over ICS.

5 CONCLUSION

In this article, we proposed an iterative algorithm to improve the modes/subspaces spanned by the modes using a previously proposed MAT algorithm. Unlike the earlier methods used for multiple rounds of sampling, we estimate the POD modes in each iteration and stop sampling when no further improvement is obtained in either the modes or subspaces spanned by the modes. The algorithm resulted in good accuracies with significant improvement in runtime even if all columns are sampled when convergence is achieved. We found that using column-norm sampling in the first round and uniform sampling in subsequent rounds resulted in good speedups, with accuracy comparable to using norm of the orthogonal component of the sampled columns.

We proposed a measure using Wedin's theorem to quantify the accuracy in the computed subspaces. In most cases, the measure was much less than its upper limit, corroborating the

ALGORITHM 3: Incremental POD computed using the iterative sampling algorithms for large matrices that do not fit in system's RAM. Input: parameters necessary for iterative algorithms and number of blocks, t .

```

1: procedure INCREMENTAL POD( $A, t, k, c, w, r, \tau, rows$ )
2:    $B \leftarrow \{a_i : i \in 1, 2, \dots, \lfloor n/t \rfloor\}$ 
3:    $\tilde{U}, \tilde{\Sigma} \leftarrow \text{ISMA}(B, k, c, w, r, \tau, rows)$ 
4:   for  $i \in 1, \dots, (t-1)$  do
5:      $B \leftarrow \{a_i : i \in \lceil in/t \rceil, \dots, \lfloor (i+1)n/t \rfloor\}$ 
6:      $\hat{U}, \hat{\Sigma} \leftarrow \text{ISMA}(B, k, c, w, r, \tau, rows)$ 
7:      $\tilde{U}, \tilde{\Sigma} \leftarrow \text{BLOCK MERGE}(\tilde{U}, \tilde{\Sigma}, \hat{U}, \hat{\Sigma}, r)$ 
8:      $\tilde{U} \leftarrow \hat{U}; \tilde{\Sigma} \leftarrow \hat{\Sigma}$ 
9:   end for
10:  return  $\tilde{u}_i, \tilde{\sigma}_i$  where  $i = 1, 2, \dots, k$ 
11: end procedure

```

Table 6. Wall clock (WT) and computational times (CT) in seconds (s) for incremental computation of POD modes for Yale faces dataset run with a single thread in a 32GB system.

		Convergence of individual modes			
Case	sampling strategy	ICS		ICRS	
		WT(s)	CT(s)	WT(s)	CT(s)
$k = 5$ (I2)	UNF	146.73	51.14	134.09	35.92
$k = 20$ (I1)	UNF	302.06	205.27	217.22	121.41
		Convergence of subspace spanned by modes			
Case	sampling strategy	ICS		ICRS	
		WT(s)	CT(s)	WT(s)	CT(s)
$k = 5$ (I2)	UNF	146.59	50.87	125.76	30.69
$k = 20$ (I1)	UNF	299.73	204.11	211.72	116.45

fact that the approximated modes were very accurate. In few cases, the measure was large due to clustering of singular values. When the rank of approximation, k , was increased so that this clustering is avoided, the error measure reduced to a large extent. In general, the accuracy of the modes are much better than indicated by the measure.

For large matrices that do not fit in the RAM, we used the iterative algorithm on each partition of the data (that fit in the RAM) to approximate the dominant left singular vectors and singular values of the partition and then used MAT operation to get the modes of the entire data. As mentioned, computing the modes using incremental sampling and MAT is advantageous as it requires only one pass over the data. We obtained very good accuracy using this method when applied on our large dataset.

APPENDIX

A EXPERIMENTAL EVALUATION OF LTSVD AND CTSVD

We evaluated the column and column and row norm based sampling algorithms (referred to as LTSVD and CTSVD). Details are contained in Algorithms (4) and (5). The parameters to be set in the

algorithms are (a) ϵ, δ : error and failure probability parameters and (b) k : desired number of POD modes. Based on these parameters, number of columns/rows to be sampled, c/w are set according to Equations (3) and (5). We evaluated the algorithms based on (a) runtime and (b) accuracy for various sets of parameters.

ALGORITHM 4: LTSVD algorithm. Inputs: matrix $A \in \mathbb{R}^{m \times n}$, vector containing sampling probabilities \mathbf{p} , number of columns to be sampled c , and rank of approximation, k .

```

1: procedure LTSVD( $A, \mathbf{p}, k, c$ )
2:   for  $i \leftarrow 1, 2, \dots, c$  do
3:     Pick  $s_i \in \{1, 2, \dots, n\}$  with probability  $\Pr[s_i = \alpha] = p_\alpha, \alpha = 1, 2, \dots, n$ 
4:      $\mathbf{c}_i \leftarrow \mathbf{a}_{s_i} / \sqrt{c p_{s_i}}$ 
5:   end for
6:    $\tilde{\mathbf{V}} \tilde{\Sigma}^2 \tilde{\mathbf{V}}^T \leftarrow \text{SVD}(C^T C)$   $\triangleright \tilde{\mathbf{V}}, \tilde{\Sigma}$  are the right singular vectors and values of  $C$ 
7:    $\tilde{\mathbf{u}}_i \leftarrow C \tilde{\mathbf{v}}_i / \tilde{\sigma}_i$  where  $i = 1, 2, \dots, l$  and  $l = \min(k, r)$ ;  $r$ : rank of  $C$ 
8:   return  $\tilde{\mathbf{u}}_i, \tilde{\sigma}_i$  where  $i = 1, 2, \dots, l$ 
9: end procedure

```

A.1 Results

We would like the error parameters, ϵ and δ , to be small (< 1). However, since the number of columns sampled depends inversely on powers of ϵ , it increases rapidly and often exceeds the number of columns in the matrix. This limits the parameter values that can be used. To get some representative results, we looked at results for three cases namely, C1: $c \approx n$ except for the YF dataset, where it is lower, C2: $c \approx n/2$ (lower for CYF), and C3: $c \ll n$ (higher for YF). For V_{2D} , Faces and CYF, we used $k = 10, 5, 2$ for the three cases C1, C2, and C3, respectively. Since YF is a larger dataset, with a more gradual decay of singular values, we used $k = 20, 10, 5$ for the three cases. For these cases, $k, \epsilon, \delta, c, w$ are shown in Table 7.

Figure 7 shows the first 20 singular values of the datasets. V_{2D} has a dominant singular value, followed by a sharp decay. The most gradual decay of singular values is for the YF dataset, where the first two modes capture only 30% of the energy. For CYF the first two singular values are very close. For all datasets, but especially for CYF, there is a clustering of the singular values beyond the first 10 values. Note that the singular values of V_{2D} is plotted on a log scale, while others are on a linear scale.

We found that the error in approximating A_k is well below the theoretical error bound for the respective algorithms in all cases. Figure 8 shows the percentage error in the singular values.³ It is seen that the error is less than 5% for CYF and YF. The two dominant singular values are captured accurately in CYF even when the error parameter values are close to one. The error is slightly larger for V_{2D} and Faces dataset, but is still less than 20%, except for CTSVD (C3) for the V_{2D} dataset.

For many applications in fluid flow and in gene expression modeling, it is the POD modes rather than the singular values that are important. Figure 9 contains a few modes from the V_{2D} and FACES datasets. It is seen that there are visually apparent distortions in the approximated modes, even though the corresponding singular values match well. A measure of the accuracy of the POD modes is the cosine similarity between the approximate mode ($\tilde{\mathbf{u}}_i$) and the one obtained using a truncated SVD of the entire dataset (\mathbf{u}_i). Figure 10 shows the angle between the two. We refer to these angles as “mode angles”. It can be seen from Figure 10 that the modes are approximated very poorly. For

³In the case of CTSVD C1, we could not obtain 10 values with the filter (see line 13 of Algorithm (4)) and the values plotted are obtained without the filter.

Table 7. Parameter Values (k, ϵ, δ) used in LTSVD and CTSVD Algorithms for Various Datasets

Case	V_{2D} k, ϵ, δ	Faces k, ϵ, δ	CYF k, ϵ, δ	YF k, ϵ, δ
LTSVD				
C1	10, 0.7, 0.45	10, 0.75, 0.8	10, 0.46, 0.44	20, 0.35, 0.35
C2	5, 1, 0.1	5, 0.75, 0.75	5, 1, 0.1	10, 0.3, 0.25
C3	2, 1, 0.8	2, 1, 0.8	2, 1, 0.8	5, 0.3, 0.25
CTSVD				
C1	10, 1.04, 0.94	10, 1.3, 1	10, 0.9, 0.7	20, 0.83, 0.9
C2	5, 1, 0.1	5, 1, 0.1	5, 1, 1	10, 0.62, 0.9
C3	2, 1, 0.8	2, 1, 0.8	2, 1, 0.8	5, 0.5, 0.9




Fig. 7. First 20 singular values of datasets considered.

ALGORITHM 5: CTSVD Algorithm. Inputs: matrix $A \in \mathbb{R}^{m \times n}$, vector containing sampling probabilities for column, \mathbf{p} , number of columns c and rows w to be sampled, and rank of approximation, k .

```

1: procedure CTSVD( $A, \mathbf{p}, k, c, w$ )
2:   for  $i \leftarrow 1, 2, \dots, c$  do
3:     Pick  $s_i \in \{1, 2, \dots, n\}$  with probability  $\Pr[s_i = \alpha] = p_\alpha, \alpha = 1, 2, \dots, n$ 
4:      $\mathbf{c}_i \leftarrow \mathbf{a}_{s_i} / \sqrt{c p_{s_i}}$ 
5:   end for
6:    $q_i \leftarrow \|\mathbf{c}_i\|_2^2; q_i \leftarrow q_i / \sum_{i=1}^m q_i$ 
7:   for  $j \leftarrow 1, 2, \dots, w$  do
8:     Pick  $\hat{s}_j \in \{1, 2, \dots, m\}$  with probability  $\Pr[\hat{s}_j = \alpha] = q_\alpha, \alpha = 1, 2, \dots, m$ 
9:      $\mathbf{w}^j \leftarrow \mathbf{c}^{\hat{s}_j} / \sqrt{w q_{\hat{s}_j}}$ 
10:  end for
11:   $\tilde{\mathbf{V}} \tilde{\Sigma}^2 \tilde{\mathbf{V}}^T \leftarrow \text{SVD}(\mathbf{W}^T \mathbf{W})$ 
12:   $\gamma \leftarrow \epsilon / (100k)$ 
13:   $l \leftarrow \min(k, \max(r : \tilde{\sigma}_r^2 \geq \gamma \|\mathbf{W}\|_F^2))$ 
14:   $\tilde{\mathbf{u}}_i \leftarrow C \tilde{\mathbf{v}}_i / \tilde{\sigma}_i$  where  $i = 1, 2, \dots, l$ 
15:  return  $\tilde{\mathbf{u}}_i, \tilde{\sigma}_i$  where  $i = 1, 2, \dots, l$ 
16: end procedure

```

reference, the corresponding mode angle is indicated in Figure 9, clearly indicating that a large mode angle implies potentially larger distortions.

The mode angles are known to be sensitive to the clustering of singular values. This is the reason why the first two mode angles are large for CYF ($\sigma_1 \sim \sigma_2$). In the limiting case when




Fig. 8. Percentage error in the singular values of different datasets for different parameter values, as given in Table 7.

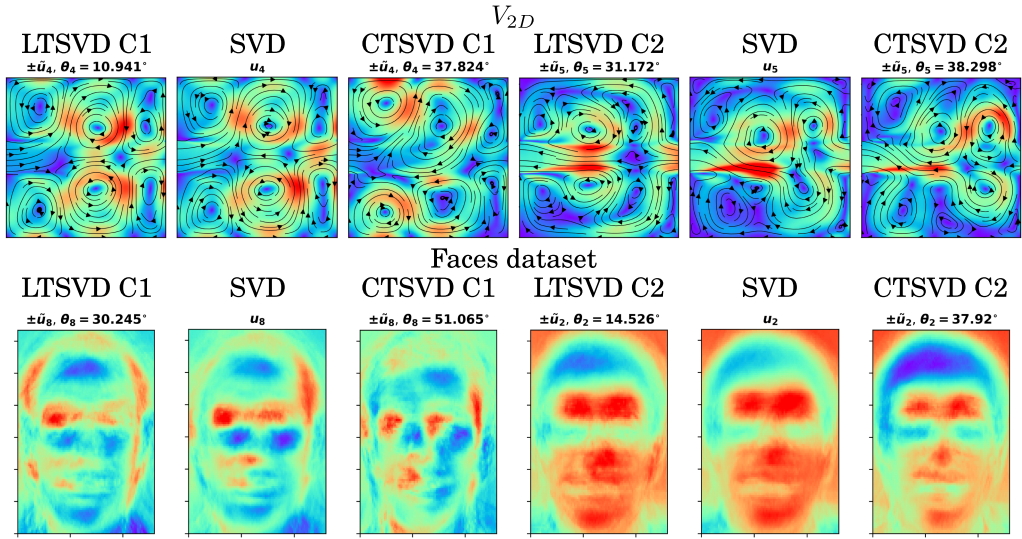


Fig. 9. POD modes of various datasets computed by SVD, LTSVD (Algorithm (4)) and CTSVD (Algorithm (5)). POD mode of V_{2D} consists of both the x and y components of velocity vector at each point in space. The color at each point denotes the magnitude of this vector at that point. Streamlines shown, denote the direction of the flow captured by the POD mode.

the singular values are identical, it is only the eigenspace that matters and not the eigenvectors. Also, for applications such as facial recognition, it is the space spanned by the modes rather than the modes themselves that are of interest. For this reason, we also measure the accuracy of the subspaces spanned by k singular vectors rather than the accuracy of each mode. A measure of this accuracy is the principal or canonical angles between the subspaces, whose cosines are the singular values of $\tilde{U}_k^T U_k$ [Björck and Golub 1973]. Figure 11 shows the k principal angles, ϕ_i , between the two k -dimensional subspaces. Although the accuracy of the subspace is better than the modes themselves, in some cases, there is significant error even for $c \approx n$. The first two principal angles for CYF match well, indicating that the subspace is captured more accurately than the modes themselves (due to almost identical singular values).




Fig. 10. First k mode angles, θ_i for different parameter values, as given in Table 7.




Fig. 11. First k principal angles, ϕ_i , for different datasets for different parameter values, as given in Table 7.

REFERENCES

- Dimitris Achlioptas and Frank McSherry. 2007. Fast computation of low-rank matrix approximations. *Journal of the ACM* 54, 2 (April 2007), 9–es. DOI:<https://doi.org/10.1145/1219092.1219097>
- Zheng-Jian Bai, Raymond H. Chan, and Franklin T. Luk. 2005. Principal component analysis for distributed data sets with updating. In *Advanced Parallel Processing Technologies*. Jiannong Cao, Wolfgang Nejdl, and Ming Xu (Eds.), Springer, Berlin, 471–483.
- C. G. Baker, K. A. Gallivan, and P. Van Dooren. 2012. Low-rank incremental methods for computing dominant singular subspaces. *Linear Algebra and its Applications* 436, 8 (2012), 2866–2888. DOI:<https://doi.org/10.1016/j.laa.2011.07.018>
- Åke Björck. 2015. *Numerical Methods in Matrix Computations*, Vol. 59. Springer.
- Ake Björck and Gene H. Golub. 1973. Numerical methods for computing angles between linear subspaces. *Mathematics of Computation* 27, 123 (1973), 579–594.
- Aritra Bose, Vassilis Kalantzis, Eugenia-Maria Kontopoulou, Mai Elkady, Peristera Paschou, and Petros Drineas. 2019. TeraPCA: A fast and scalable software package to study genetic variation in tera-scale genotypes. *Bioinformatics* 35, 19 (04 2019), 3679–3683. DOI:<https://doi.org/10.1093/bioinformatics/btz157>
- Matthew Brand. 2002. Incremental singular value decomposition of uncertain data with missing values. In *Proceedings of the European Conference on Computer Vision*. Springer, 707–720.
- Matthew Brand. 2006. Fast low-rank modifications of the thin singular value decomposition. *Linear Algebra and Its Applications* 415, 1 (2006), 20–30.
- AT&T Laboratories Cambridge. 2002. The Database of Faces. (2002). Retrieved 05 Jan., 2018 from <http://www.cl.cam.ac.uk/research/dtg/attarchive/facedatabase.html>.
- Michael B. Cohen, Yin Tat Lee, Cameron Musco, Christopher Musco, Richard Peng, and Aaron Sidford. 2015. Uniform sampling for matrix approximation. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science (ITCS'15)*. ACM, New York, NY, 181–190. DOI:<https://doi.org/10.1145/2688073.2688113>
- Amit Deshpande, Luis Rademacher, Santosh Vempala, and Grant Wang. 2006. Matrix approximation and projective clustering via volume sampling. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithm*.

- A. Deshpande and S. Vempala. 2006. Adaptive sampling and fast low-rank matrix approximation. In *Proceedings of the 9th International Conference on Approximation Algorithms for Combinatorial Optimization Problems, and 10th International Conference on Randomization and Computation (APPROX'06/RANDOM'06)*. Springer-Verlag, Berlin, 292–303. DOI:https://doi.org/10.1007/11830924_28
- Petros Drineas, Eleni Drinea, and Patrick S. Huggins. 2003. An experimental evaluation of a monte-carlo algorithm for singular value decomposition. In *Advances in Informatics*, Yannis Manolopoulos, Skevos Evripidou, and Antonis C. Kakas (Eds.). Springer, Berlin, 279–296.
- P. Drineas, A. Frieze, R. Kannan, S. Vempala, and V. Vinay. 2004. Clustering large graphs via the singular value decomposition. *Machine Learning* 56, 1–3 (June 2004), 9–33. DOI:<https://doi.org/10.1023/B:MACH.0000033113.59016.96>
- Petros Drineas, Ravi Kannan, and Michael W. Mahoney. 2006. Fast monte carlo algorithms for matrices II: Computing a low-rank approximation to a matrix. *SIAM Journal on Computing* 36, 1 (2006), 158–183.
- Petros Drineas and Michael W. Mahoney. 2016. RandNLA: Randomized numerical linear algebra. *Communications of the ACM* 59, 6 (May 2016), 80–90. DOI:<https://doi.org/10.1145/2842602>
- Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. 2008. Relative-error CUR matrix decompositions. *SIAM Journal on Matrix Analysis and Applications* 30, 2 (Sept. 2008), 844–881. DOI:<https://doi.org/10.1137/07070471X>
- N. B. Erichson, S. L. Brunton, and J. N. Kutz. 2017. Compressed singular value decomposition for image and video processing. In *Proceedings of the 2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*. 1880–1888.
- N. Benjamin Erichson, Sergey Voronin, Steven L. Brunton, and J. Nathan Kutz. 2019. Randomized Matrix Decompositions Using R. *Journal of Statistical Software* 89, 11 (2019). DOI: [10.18637/jss.v089.i](https://doi.org/10.18637/jss.v089.i)
- Alan Frieze, Ravi Kannan, and Santosh Vempala. 2004. Fast monte-carlo algorithms for finding low-rank approximations. *Journal of the ACM* 51, 6 (Nov. 2004), 1025–1041. DOI:<https://doi.org/10.1145/1039488.1039494>
- A. S. Georgiades, P. N. Belhumeur, and D. J. Kriegman. 2001. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 6 (2001), 643–660.
- Gene H. Golub and Charles F. Van Loan. 2007. *Matrix Computations* (3 ed.). Hindustan Book agency.
- Ming Gu. 2015. Subspace iteration randomization and singular value problems. *SIAM Journal on Scientific Computing* 37, 3 (2015), A1139–A1173.
- Ming Gu and Stanley C. Eisenstat. 1993. A stable and fast algorithm for updating the singular value decomposition. Technical Report. YALEU/DCS/RR-966. Dept. of Computer Science, Yale University.
- Nathan Halko, Per-Gunnar Martinsson, Yoel Shkolnisky, and Mark Tygert. 2011b. An algorithm for the principal component analysis of large data sets. *SIAM Journal on Scientific Computing* 33, 5 (Oct. 2011), 2580–2594. DOI:<https://doi.org/10.1137/100804139>
- N. Halko, P. G. Martinsson, and J. A. Tropp. 2011a. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review* 53, 2 (2011), 217–288.
- M. A. Iwen and B. W. Ong. 2016. A distributed and incremental SVD algorithm for agglomerative data analysis on large networks. *SIAM Journal on Matrix Analysis and Applications* 37, 4 (2016), 1699–1718.
- K. C. Lee, J. Ho, and D. Kriegman. 2005. Acquiring linear subspaces for face recognition under variable lighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27, 5 (2005), 684–698.
- Huamin Li, George C. Linderman, Arthur Szlam, Kelly P. Stanton, Yuval Kluger, and Mark Tygert. 2017. Algorithm 971: An implementation of a randomized algorithm for principal component analysis. *ACM Transactions on Mathematical Software* 43, 3 (Jan. 2017), 14 pages. DOI:<https://doi.org/10.1145/3004053>
- Mu Li, Gary L. Miller, and Richard Peng. 2013. Iterative row sampling. In *Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS'13)*. IEEE Computer Society, 127–136. DOI:<https://doi.org/10.1109/FOCS.2013.22>
- Ren-Cang Li. 1996. *Relative Perturbation Theory: (II) Eigenspace and Singular Subspace Variations*. Technical Report Lapack working note #85. Oak Ridge National Laboratories.
- Ren-Cang Li. 2014. Matrix perturbation theory. In *Handbook of Linear Algebra*. L. Hogben (Ed.), Chapman and Hall/CRC.
- Yingyu Liang, Maria-Florina F. Balcan, Vandana Kanchanapally, and David Woodruff. 2014. Improved distributed principal component analysis. In *Advances in Neural Information Processing Systems 27*. Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Eds.), Curran Associates, Inc., 3113–3121. Retrieved from <http://papers.nips.cc/paper/5619-improved-distributed-principal-component-analysis.pdf>.
- R. Mathias. 2014. Singular values and singular value inequalities. In *Handbook of Linear Algebra*. L. Hogben (Ed.), Chapman and Hall/CRC.
- Aditya Krishna Menon and Charles Elkan. 2011. Fast algorithms for approximating the singular value decomposition. *ACM Transactions on Knowledge Discovery from Data* 5, 2 (Feb. 2011), 36 pages. DOI:<https://doi.org/10.1145/1921632.1921639>
- Nam H. Nguyen, Thong T. Do, and Trac D. Tran. 2009. A fast and efficient algorithm for low-rank approximation of a matrix. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC'09)*. ACM, New York, NY, 215–224. DOI:<https://doi.org/10.1145/1536414.1536446>

- Yongming Qu, George Ostrouchov, Nagiza Samatova, and Al Geist. 2002. Principal component analysis for dimension reduction in massive distributed data sets. In *Proceedings of IEEE International Conference on Data Mining (ICDM)*, Vol. 1318. 1788.
- V. Rokhlin, A. Szlam, and M. Tygert. 2010. A randomized algorithm for principal component analysis. *SIAM Journal on Matrix Analysis and Applications* 31, 3 (2010), 1100–1124. DOI:<https://doi.org/10.1137/080736417>
- T. Sarlos. 2006. Improved approximation algorithms for large matrices via random projections. In *Proceedings of the 2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. 143–152. DOI:<https://doi.org/10.1109/FOCS.2006.37>
- G. W. Stewart and Ji guang Sun. 1990. *Matrix Perturbation Theory*. Academic Press.
- J. Sun, Yinglian Xie, Hui Zhang, and Christos Faloutsos. 2007. Less is more: Compact matrix decomposition for large sparse graphs. In *Proceedings of the 2007 SIAM International Conference on Data Mining*. DOI:<https://doi.org/10.1137/1.9781611972771.33>
- Vinita Vasudevan and M. Ramakrishna. 2017. A hierarchical singular value decomposition algorithm for low rank matrices. (October 2017). arXiv:1710.02812. Retrieved from <https://arxiv.org/abs/1710.02812>.
- Per-Ake Wedin. 1972. Perturbation bounds in connection with singular value decomposition. *BIT:Numerical Mathematics* 12, 1 (1972), 99–111.
- I. Yamazaki, S. Tomov, and J. Dongarra. 2017. Sampling algorithms to update truncated SVD. In *Proceedings of the 2017 IEEE International Conference on Big Data (Big Data)*. 817–826. DOI:<https://doi.org/10.1109/BigData.2017.8257997>

Received November 2020; revised July 2021; accepted December 2021