

# Randomized algorithms for distributed computation of principal component analysis and singular value decomposition

Huamin Li · Yuval Kluger · Mark Tygert

May 31, 2017

**Abstract** As illustrated via numerical experiments with an implementation in Spark (the popular platform for distributed computation), randomized algorithms provide solutions to two ubiquitous problems: (1) the distributed calculation of a full principal component analysis or singular value decomposition of a highly rectangular matrix, and (2) the distributed calculation of a low-rank approximation (in the form of a singular value decomposition) to an arbitrary matrix. Carefully honed algorithms yield results that are uniformly superior to those of the stock, deterministic implementations in Spark; for instance, whereas the stock software will without warning return left singular vectors that are far from numerically orthonormal, a significantly burnished randomized implementation generates left singular vectors that are numerically orthonormal to nearly the machine precision.

**Keywords** Spark · parallel · cluster · orthogonalization

**Mathematics Subject Classification (2010)** 65F15 · 68W20 · 65F25 · 65F30

---

Y. Kluger and H. Li were supported in part by United States National Institutes of Health grant 1R01HG008383-01A1. Y. Kluger is with the Program in Applied Mathematics, the Program in Biological and Biomedical Sciences, the Cancer Center, the Center for Medical Informatics, and the Department of Pathology in the School of Medicine at Yale University.

---

H. Li

Yale University, Program in Applied Mathematics, 51 Prospect St., New Haven, CT 06510  
E-mail: huamin.li@yale.edu

Y. Kluger

Yale University, School of Medicine, Department of Pathology, Suite 505L, 300 George St., New Haven, CT 06520  
E-mail: yuval.kluger@yale.edu

M. Tygert

Facebook Artificial Intelligence Research, 1 Facebook Way, Menlo Park, CA 94025  
E-mail: tygert@fb.com

## 1 Introduction

Singular value decomposition is the “Swiss Army Knife” and “Rolls Royce” of matrix decompositions due to its ubiquitous, unsurpassed utility, as observed by [10], quoting Dianne O’Leary. Indeed, singular value decomposition provides highly accurate solutions to most problems in numerical linear algebra, and is especially important for low-rank approximation and principal component analysis, which is among the most important methods in statistics, data analysis and analytics, and machine learning, as discussed by [12] and many others. The singular value decomposition of a matrix  $A$  consists of matrices  $U$ ,  $\Sigma$ , and  $V$  such that

$$A = U\Sigma V^*, \quad (1)$$

where the columns of  $U$  are orthonormal, as are the columns of  $V$ , the adjoint (conjugate transpose) of  $V$  is  $V^*$ , and  $\Sigma$  is diagonal and its entries are nonnegative. Accurate, efficient numerical calculation of the singular value decomposition is essential for many applications, including in the now common setting that the data being processed is distributed across multiple computers.

In fact, distributed calculations on clusters of computers are spawning whole industries devoted to software platforms such as Hadoop and Spark, as discussed by [15], [9], and many others; Spark is rapidly becoming a dominant software platform, enabling distributed machine learning via its MLlib library. The current version of Spark’s MLlib includes rudimentary routines for calculating the singular value decomposition; we aim to make improvements based on combining and honing both well-known and lesser-known numerical methods and tricks.

Below, we consider two problems that are of particular interest for distributed computation due to their computational tractability: {1} calculating the full singular value decomposition of a tall and skinny matrix (a matrix for which a full row can fit on a single machine), and {2} calculating a low-rank approximation to an arbitrary matrix such that the spectral norm of the difference between the approximation and the matrix being approximated is nearly as small as possible. These problems are relatively tractable since every matrix in the resulting decompositions has at least one dimension that is small enough to ensure that a full row can fit in memory on a single machine. The solutions below to the second problem {2} leverage the solutions to the first problem {1}, returning all results in the form of a singular value decomposition  $U\Sigma V^*$ , where the columns of  $U$  are orthonormal, as are the columns of  $V$ , and  $\Sigma$  is diagonal and its entries are nonnegative.

The rest of the present paper has the following structure: Section 2 addresses problem {1}, introducing Algorithms 1–4. Section 3 addresses problem {2}, introducing Algorithms 5 and 6, which leverage the Algorithms 1–4 discussed in Section 2. Section 4 summarizes our outlook. Appendix A corroborates the results presented earlier in the paper, regardless of the number of executors in our cluster of computers. Appendix B displays the times required to synthesize the matrices in our tests (just for reference, for comparative purposes). Neither appendix is integral to the main points of the present paper, and may be omitted.

*Remark 1* Throughout the present paper, the “working precision” refers to the machine precision adjusted to account for roundoff error. For instance, the working precision could be  $10^{-11}$  for double-precision floating-point arithmetic with

matrices of the sizes considered below (whereas the machine precision would be  $2.2 \times 10^{-16}$ ). When interpreting the tables, please note the italicized text in Table 1 for the heading  $\|A - U\Sigma V^*\|_2$ .

*Remark 2* The main purpose of our implementation is to add efficient principal component analysis and singular value decomposition to Spark, not to compute them as efficiently as could be possible (in fact, bypassing Spark may enhance efficiency). On the whole, the aspects of Spark unrelated to sophisticated mathematical algorithms tend to be more important than the parts dependent on such algorithms, even though the algorithmic aspects are the subject of the present paper. As much as we would like to think that our own contributions are the most important, we do realize that database management, stream processing, fault tolerance and recovery, ease of deployment and administration, coupling with other systems, etc. are typically far more important. Spark is becoming a dominant platform for machine learning at scale, so the purpose of our implementation is to enable big data analytics for Spark, whether or not Spark is the ideal platform for principal component analysis or singular value decomposition. In particular, Spark is likely to distribute data over clusters as appropriate for tasks other than principal component analysis and singular value decomposition, and our implementation must deal with the data as distributed however Spark sees fit. This remark is especially important in light of the findings of [9].

*Remark 3* The Spark implementation is available at <http://github.com/hl475/svd> (and we hope that the main branch of Spark will pull in these changes soon). For expository purposes, we also provide a serial implementation in Python 3, at <http://tygert.com/valid.tar.gz> (the algorithms of the present paper are meant for parallel computation, but we opt to provide the Python 3 codes in addition to the implementation for Spark, as the Python is far easier to read and run).

*Remark 4* Many of the algorithms discussed below are randomized. Exhaustive prior work, including that of [2], [4], [6], [11], [14], [18], and [19], demonstrates that the randomized methods are at least as accurate, reliable, and efficient as the more classical deterministic algorithms. The probability of obtaining results departing significantly from those observed is well known to be negligible, both empirically and theoretically (for rigorous proofs, see the work mentioned in the previous sentence).

## 2 Singular value decomposition of tall and skinny matrices

There are many ways to calculate the singular value decompositions of tall and skinny matrices (matrices for which a full row can fit on a single machine). This section compares different methods.

One method is perfectly numerically stable — producing results accurate to nearly the machine precision — but requires merging intermediate results through multiple levels of a dependency tree (such a merge is known as a “reduction”); this is the randomized method of [2] and Section 5 of [4] (which draws on the techniques of [6] and others). The randomization there eliminates the need for pivoting. The pseudocode in Algorithms 1 and 2 summarizes this randomized method, accounting for several considerations discussed shortly. Another method, based on

computing the Gram matrix  $A^*A$  of the matrix  $A$  being processed, loses half the digits in some cases, but can leverage extremely efficient accumulation/aggregation strategies with minimal blocking dependencies and synchronization requirements; this is the method of [17], [8], [18], and [19]. The pseudocode in Algorithms 3 and 4 summarizes this method, accounting for the following considerations. Both methods produce highest accuracy when running their orthonormalization of singular vectors twice in succession, though running twice is superfluous during all but the last of the subspace iterations of the randomized algorithms for low-rank approximation discussed in the following section. Running the orthonormalization twice ensures that the resulting singular vectors are numerically orthonormal, though running twice has little effect on the accuracy of their linear span (and little effect on the spectral norm of the difference between the calculated decomposition and the matrix  $A$  being processed).

A third method, based on “tournament pivoting,” is similar to the first method discussed above, being reasonably numerically stable while requiring the merging of intermediate results through multiple levels of a dependency/reduction tree; this is the method of [5] and others. This third method is deterministic (unlike the first method), but is otherwise more complicated, less efficient, less accurate, and weaker theoretically with regard to revealing the ranks of the rank-deficient matrices commonly encountered when using the singular value decomposition for dimension reduction; similar remarks apply to a fourth method, that of [3]. In the sequel, we consider only the first two methods discussed above. These two methods also happen to be the easiest to implement (which can be critical for future deployment and maintenance).

*Remark 5* For convenience, in our implementation of the first method mentioned above (the randomized method) we replaced the usual random Gaussian matrix with a product  $DFS\tilde{D}F\tilde{S}$ , where  $D$  and  $\tilde{D}$  are diagonal matrices whose diagonal entries are independent and identically distributed random numbers drawn uniformly from the unit circle in the complex plane,  $F$  is the discrete Fourier transform, and  $S$  and  $\tilde{S}$  are independent uniformly random permutations (calculated via the Fisher-Yates-Durstenfeld-Knuth shuffle of [7]). To process vectors of real numbers (rather than complex numbers), we partitioned the vectors into pairs of real numbers and viewed each pair as consisting of the real and imaginary parts of a complex number. We found empirically that chaining two products  $DFS$  into  $DFS\tilde{D}F\tilde{S}$  was sufficient; chaining a few (specifically, logarithmic in the number of columns of the matrix whose singular value decomposition is being computed) is rigorously known to be sufficient, as proven by [1]. Chaining several is affordable computationally but seems like overkill.

*Remark 6* When testing the second method mentioned above (the method based on the Gram matrix) we found that explicitly normalizing the left singular vectors improved accuracy significantly. Explicitly normalizing does require computing the Euclidean norms of the columns of the matrix of left singular vectors, but this costs substantially less than computing the Gram matrix in the first place.

The remainder of the present section gives empirical results on the first two methods discussed above and on the existing singular value decomposition of tall and skinny matrices in Spark’s MLlib, which is based on Gram matrices, similar to the second method mentioned above (though the existing Spark routine

lacks the refinement of Remark 6). We demonstrate the performance of the first two methods both when running their orthonormalization twice in succession and when running only once (with the resulting loss of accuracy and gain in speed). Our primary implementations of the algorithms in the present section process an IndexedRowMatrix from Spark’s MLlib.

Our software includes examples of matrices with many different distributions of singular values and singular vectors. For clarity of the presentation, the results we present below pertain to the following class of matrices — matrices for which the various algorithms produced accuracies near the worst that we encountered in our experiments:

$$A = U\Sigma V^*, \quad (2)$$

where  $U$  and  $V$  are  $m \times m$  and  $n \times n$  discrete cosine transforms, respectively, and  $\Sigma$  is the  $m \times n$  matrix whose entries are all zeros aside from the diagonal entries

$$\Sigma_{j,j} = \exp\left(\frac{j-1}{n-1} \cdot \ln\left(10^{-20}\right)\right) \quad (3)$$

for  $j = 1, 2, \dots, n$ . Note that, like  $A$  from (2) and (3), matrices arising from real data are often numerically rank-deficient; indeed, real data sets are often messy, with duplicate or nearly duplicate columns and rows, symmetries or near symmetries that limit the numerical rank, etc. Singular value decomposition and principal component analysis are very helpful for untangling the mess in real data, and certainly need to function reliably in such circumstances, circumstances such that the matrix being processed may be highly ill-conditioned.

The headings of the tables have the meanings detailed in Table 1. Our Spark environment is detailed in Table 2. We used many — 20 — iterations of the power method in order to ascertain the spectral-norm errors reported in the tables. The timings in the tables do not include the time spent checking the accuracy (we used so many power iterations just to be extra careful in providing highly accurate error estimates, in order to facilitate fully trustworthy comparisons of the different algorithms).

Tables 3, 4, and 5 report timings and errors for several experiments. The reconstruction errors  $\|A - U\Sigma V^*\|_2$  for Algorithms 1 and 2 (which are similar) are clearly superior to all those for Algorithms 3 and 4, which makes sense since the latter algorithms use the Gram matrix and can therefore lose half their digits. For the left singular vectors, the errors  $\text{MaxEntry}(|U^*U - I|)$  for Algorithms 2 and 4 (which are similarly good) are clearly superior to all those for Algorithms 1 and 3, which makes sense since the latter algorithms orthonormalize the singular vectors only once. For the right singular vectors, the error  $\text{MaxEntry}(|V^*V - I|)$  is near the machine precision ( $2.2 \times 10^{-16}$ ) for all algorithms. All together, then, Algorithm 2 is the most accurate of all, with all its errors approaching the machine precision adjusted for roundoff. However, on our cluster with our version of Spark, Algorithm 4 is somewhat faster than Algorithm 2; the reconstruction error  $\|A - U\Sigma V^*\|_2$  is somewhat worse for Algorithm 4 than for Algorithm 2, but may be acceptable in many circumstances.

As expected, the timings in Tables 3, 4, and 5 are roughly proportional to the numbers of rows in the matrices (the number of columns is fixed throughout these tables), and the errors adhere to the working precision mentioned in Remark 1 and in the pseudocodes for the algorithms.

### 3 Low-rank approximation of arbitrary matrices

As discussed by [11], randomized algorithms permit the efficient calculation of nearly optimal rank- $k$  approximations to a given  $m \times n$  matrix  $A$ , that is, of matrices  $U$ ,  $\Sigma$ , and  $V$  such that  $U$  is  $m \times k$ ,  $\Sigma$  is  $k \times k$ ,  $V$  is  $n \times k$ , the columns of  $U$  are orthonormal, as are the columns of  $V$ ,  $\Sigma$  is diagonal and its entries are nonnegative, and

$$\|A - U\Sigma V^*\|_2 \approx \sigma_{k+1}(A), \quad (4)$$

where  $\|A - U\Sigma V^*\|_2$  denotes the spectral norm of  $A - U\Sigma V^*$ , and  $\sigma_{k+1}(A)$  is the spectral-norm accuracy of the best approximation to  $A$  of rank at most  $k$  (which is also equal to the  $(k+1)$ st greatest singular value of  $A$ ).

Our codes implement Algorithms 4.4 and 5.1 of [11], duplicated here as Algorithms 5 and 6, respectively. The output of Algorithm 5 feeds into Algorithm 6. Algorithm 5 is based on tall-skinny matrix factorizations of the form  $Q \cdot R$ , where the columns of  $Q$  are orthonormal and  $R$  is square ( $R$  need not be triangular, however). Given a matrix factorization of the form  $U \cdot \Sigma \cdot V^*$ , where the columns of  $U$  are orthonormal, as are the columns of  $V$ , and where both  $\Sigma$  and  $V$  are square, we use  $Q = U$  and  $R = \Sigma V^*$  to obtain a factorization of the form  $Q \cdot R$ . In our implementations, we obtain matrix factorizations of the form  $U \cdot \Sigma \cdot V^*$  via the methods evaluated in Section 2 above. Below, we compare the results of using the two different methods evaluated in Section 2 above for the tall-skinny matrix factorizations required in Algorithm 5, always running Algorithm 5 and feeding its output into Algorithm 6.

We run the tall-skinny factorization twice in succession only for the very last step in Algorithm 5; the purpose of the earlier steps in Algorithm 5 is to track a subspace, and so long as the column spaces of the resulting matrices are accurate, then whether the columns are numerically orthonormal matters little (in fact, replacing  $Q$  with the lower triangular/trapezoidal factor  $L$  in an  $LU$  factorization is sufficient, as shown by [16] and [14] — the column space of  $L$  is the same as the column space of  $Q$ ).

The remainder of the present section gives empirical results on Algorithm 5 feeding into Algorithm 6, when using in Algorithm 5 the two different methods evaluated in Section 2. This section also presents empirical results on the existing implementation of low-rank approximation in Spark’s MLlib, which is based on the implicitly restarted Arnoldi method in ARPACK of [13]. Our implementations of the algorithms in the present section process a BlockMatrix from Spark’s MLlib.

Our software includes examples of matrices with many different distributions of singular values and singular vectors. For clarity of the presentation, the results we present below pertain to the following class of matrices — matrices for which the various algorithms produced accuracies near the worst that we encountered in our experiments:

$$A = U\Sigma V^*, \quad (5)$$

where  $U$  and  $V$  are  $m \times m$  and  $n \times n$  discrete cosine transforms, respectively, and  $\Sigma$  is the  $m \times n$  matrix whose entries are all zeros aside from the diagonal entries

$$\Sigma_{j,j} = \exp\left(\frac{j-1}{l-1} \cdot \ln(10^{-20})\right) \quad (6)$$

for  $j = 1, 2, \dots, l$ . Notice that (6) is the same as (3) when replacing  $l$  in (6) with  $n$ .

The headings of the tables have the meanings detailed in Table 1 (which defines  $l$ ). Our Spark environment is detailed in Table 2. We used many — 20 — iterations of the power method in order to ascertain the spectral-norm errors reported in the tables. The timings in the tables do not include the time spent checking the accuracy (we used so many power iterations just to be extra careful in providing highly accurate error estimates, in order to facilitate fully trustworthy comparisons of the different algorithms).

Tables 9 and 10 consider sizes of matrices that are too large for computing all possible singular values and singular vectors (rather than just a low-rank approximation) on our cluster with Spark. Table 9 indicates that, on our cluster with our version of Spark, the timings for Algorithm 6 using Algorithms 1 and 2 are similar to the timings for Algorithm 6 using Algorithms 3 and 4. At the same time, Table 10 indicates that the reconstruction error  $\|A - U\Sigma V^*\|_2$  for Algorithm 6 using Algorithms 1 and 2 is superior to the error for Algorithm 6 using Algorithms 3 and 4, while the other notions of accuracy are comparable for both. Thus, on our cluster with our version of Spark, Algorithm 6 using Algorithms 1 and 2 makes more sense than Algorithm 6 using Algorithms 3 and 4.

Tables 6, 7, and 8 correspond to Tables 3, 4, and 5. In accord with Tables 9 and 10, on our cluster with our version of Spark, the timings for Algorithm 6 using Algorithms 1 and 2 are similar to the timings for Algorithm 6 using Algorithms 3 and 4, while the reconstruction error  $\|A - U\Sigma V^*\|_2$  for Algorithm 6 using Algorithms 1 and 2 is superior to the error for Algorithm 6 using Algorithms 3 and 4, and the other notions of accuracy are comparable for both. Tables 6, 7, and 8 thus also show that, on our cluster with our version of Spark, Algorithm 6 using Algorithms 1 and 2 makes more sense than Algorithm 6 using Algorithms 3 and 4.

The timings in Table 9 and the errors in Table 10 are as expected, as are all the results in Tables 6, 7, and 8, with the timings roughly proportional to  $l$  times the numbers of entries in the matrices, and with the errors adhering to the working precision mentioned in Remark 1 and in the pseudocodes for the algorithms.

## 4 Conclusion

The numerical experiments reported above illustrate that the algorithms detailed in this paper outperform (or at least match) the stock implementations for Spark’s MLlib with respect to both accuracy and efficiency. As Spark’s library for machine learning migrates to the upcoming DataFrame format, it could incorporate these algorithms, as could other platforms for distributed computation. The key is the combination of various careful choices and tricks — convenient randomization, explicit normalization of singular vectors, single vs. double orthonormalization, highly accurate reconstruction vs. subspace tracking, etc.

---

**Algorithm 1:** Randomized singular value decomposition of tall and skinny matrices (from [2])

---

**Input:** A tall and skinny real matrix  $A$

**Output:** Real matrices  $U$ ,  $\Sigma$ , and  $V$  such that  $A = U\Sigma V^*$ , the columns of  $U$  are orthonormal, as are the columns of  $V$ , and  $\Sigma$  is diagonal and its entries are nonnegative

- 1 Apply an appropriately random orthogonal matrix  $\Omega$  (see Remark 5 regarding “appropriately random”) to every column of  $A^*$ , obtaining  $B = \Omega A^*$ .
  - 2 Using the TSQR method of [6], compute a factorization  $B^* = QR$ , where the columns of  $Q$  are orthonormal, and  $R$  is upper triangular.
  - 3 Discard the rows of  $R$  corresponding to diagonal entries which are zero, and discard the corresponding columns of  $Q$ , too (if working in finite-precision arithmetic, view any diagonal entry of  $R$  as numerically zero that is less than the first diagonal entry of  $R$  times the working precision).
  - 4 Calculate the singular value decomposition  $R = \tilde{U}\Sigma\tilde{V}^*$ , where the columns of  $\tilde{U}$  are orthonormal, as are the columns of  $\tilde{V}$ , and  $\Sigma$  is diagonal and its entries are nonnegative.
  - 5 Form  $U = Q\tilde{U}$ .
  - 6 Apply the inverse of the random orthogonal matrix  $\Omega$  from Step 1 to every column of  $\tilde{V}$ , obtaining  $V = \Omega^{-1}\tilde{V}$  (as  $\Omega$  is orthogonal,  $\Omega^{-1} = \Omega^*$ ).
- 

**Algorithm 2:** Randomized singular value decomposition of tall and skinny matrices (from [2]), with double orthonormalization

---

**Input:** A tall and skinny real matrix  $A$

**Output:** Real matrices  $U$ ,  $\Sigma$ , and  $V$  such that  $A = U\Sigma V^*$ , the columns of  $U$  are orthonormal, as are the columns of  $V$ , and  $\Sigma$  is diagonal and its entries are nonnegative

- 1 Apply an appropriately random orthogonal matrix  $\Omega$  (see Remark 5 regarding “appropriately random”) to every column of  $A^*$ , obtaining  $B = \Omega A^*$ .
  - 2 Using the TSQR method of [6], compute a factorization  $B^* = \tilde{Q}\tilde{R}$ , where the columns of  $\tilde{Q}$  are orthonormal, and  $\tilde{R}$  is upper triangular.
  - 3 Discard the rows of  $\tilde{R}$  corresponding to diagonal entries which are zero, and discard the corresponding columns of  $\tilde{Q}$ , too (if working in finite-precision arithmetic, view any diagonal entry of  $\tilde{R}$  as numerically zero that is less than the first diagonal entry of  $\tilde{R}$  times the working precision).
  - 4 Using the TSQR method of [6], compute a factorization  $\tilde{Q} = QR$ , where the columns of  $Q$  are orthonormal, and  $R$  is upper triangular.
  - 5 Discard the rows of  $R$  corresponding to diagonal entries which are zero, and discard the corresponding columns of  $Q$ , too (if working in finite-precision arithmetic, view any diagonal entry of  $R$  as numerically zero that is less than the first diagonal entry of  $R$  times the working precision).
  - 6 Form  $T = R\tilde{R}$ .
  - 7 Calculate the singular value decomposition  $T = \tilde{U}\Sigma\tilde{V}^*$ , where the columns of  $\tilde{U}$  are orthonormal, as are the columns of  $\tilde{V}$ , and  $\Sigma$  is diagonal and its entries are nonnegative.
  - 8 Form  $U = Q\tilde{U}$ .
  - 9 Apply the inverse of the random orthogonal matrix  $\Omega$  from Step 1 to every column of  $\tilde{V}$ , obtaining  $V = \Omega^{-1}\tilde{V}$  (as  $\Omega$  is orthogonal,  $\Omega^{-1} = \Omega^*$ ).
-



---

**Algorithm 3:** Gram-based singular value decomposition of tall and skinny matrices (from [19])

---

**Input:** A tall and skinny real matrix  $A$

**Output:** Real matrices  $U$ ,  $\Sigma$ , and  $V$  such that  $A = U\Sigma V^*$ , the columns of  $U$  are orthonormal, as are the columns of  $V$ , and  $\Sigma$  is diagonal and its entries are nonnegative

- 1 Form the Gram matrix  $B = A^*A$ .
  - 2 Calculate the eigendecomposition  $B = VD V^*$ , where the columns of  $V$  are orthonormal, and  $D$  is diagonal and its entries are nonnegative.
  - 3 Form  $\tilde{U} = AV$ .
  - 4 Set  $\tilde{\Sigma}$  to be the diagonal matrix whose diagonal entries are the Euclidean norms of the columns of  $\tilde{U}$ , in accord with Remark 6.
  - 5 Discard the columns and rows of  $\tilde{\Sigma}$  corresponding to diagonal entries which are zero, and discard the corresponding columns of  $\tilde{U}$  and  $V$ , too (if working in finite-precision arithmetic, view any entry of  $\tilde{\Sigma}$  as numerically zero that is less than the greatest entry of  $\tilde{\Sigma}$  times the square root of the working precision).
  - 6 Form  $U = \tilde{U}\tilde{\Sigma}^{-1}$ , in accord with Remark 6.
- 

**Algorithm 4:** Gram-based singular value decomposition of tall and skinny matrices (from [19]), with double orthonormalization

---

**Input:** A tall and skinny real matrix  $A$

**Output:** Real matrices  $U$ ,  $\Sigma$ , and  $V$  such that  $A = U\Sigma V^*$ , the columns of  $U$  are orthonormal, as are the columns of  $V$ , and  $\Sigma$  is diagonal and its entries are nonnegative

- 1 Form the Gram matrix  $B = A^*A$ .
  - 2 Calculate the eigendecomposition  $B = \tilde{V}\tilde{D}\tilde{V}^*$ , where the columns of  $\tilde{V}$  are orthonormal, and  $\tilde{D}$  is diagonal and its entries are nonnegative.
  - 3 Form  $\tilde{Y} = A\tilde{V}$ .
  - 4 Set  $\tilde{\Sigma}$  to be the diagonal matrix whose diagonal entries are the Euclidean norms of the columns of  $\tilde{Y}$ , in accord with Remark 6.
  - 5 Discard the columns and rows of  $\tilde{\Sigma}$  corresponding to diagonal entries which are zero, and discard the corresponding columns of  $\tilde{Y}$  and  $\tilde{V}$ , too (if working in finite-precision arithmetic, view any entry of  $\tilde{\Sigma}$  as numerically zero that is less than the greatest entry of  $\tilde{\Sigma}$  times the square root of the working precision).
  - 6 Form  $Y = \tilde{Y}\tilde{\Sigma}^{-1}$ , in accord with Remark 6.
  - 7 Form the Gram matrix  $Z = Y^*Y$ .
  - 8 Calculate the eigendecomposition  $Z = WDW^*$ , where the columns of  $W$  are orthonormal, and  $D$  is diagonal and its entries are nonnegative.
  - 9 Form  $\tilde{Q} = YW$ .
  - 10 Set  $T$  to be the diagonal matrix whose diagonal entries are the Euclidean norms of the columns of  $\tilde{Q}$ , in accord with Remark 6.
  - 11 Discard the columns and rows of  $T$  corresponding to diagonal entries which are zero, and discard the corresponding columns of  $\tilde{Q}$  and  $W$ , too (if working in finite-precision arithmetic, view any entry of  $T$  as numerically zero that is less than the greatest entry of  $T$  times the square root of the working precision).
  - 12 Form  $Q = \tilde{Q}T^{-1}$ , in accord with Remark 6.
  - 13 Form  $R = TW^*\tilde{\Sigma}\tilde{V}^*$ .
  - 14 Calculate the singular value decomposition  $R = P\Sigma V^*$ , where the columns of  $P$  are orthonormal, as are the columns of  $V$ , and  $\Sigma$  is diagonal and its entries are nonnegative.
  - 15 Form  $U = QP$ .
-

---

**Algorithm 5:** Randomized subspace iteration (Algorithm 4.4 of [11])
 

---

**Input:** A real  $m \times n$  matrix  $A$  and integers  $l$  and  $i$  such that  $0 < l < \min(m, n)$  and  $i \geq 0$ ; the number of iterations is  $i$

**Output:** A real  $m \times l$  matrix  $Q$  whose columns are orthonormal and whose range approximates the range of  $A$ , in the sense that the spectral norm  $\|A - QQ^*A\|_2$  is small

- 1 Form an  $n \times l$  matrix  $\tilde{Q}_0$  whose entries are independent and identically distributed centered Gaussian random variables.
  - 2 **for**  $j = 1$  **to**  $i$  **do**
  - 3     Form  $Y_j = A\tilde{Q}_{j-1}$ .
  - 4     Compute a factorization  $Y_j = Q_j R_j$ , where the columns of  $Q_j$  are orthonormal and  $R_j$  is square, using Algorithm 1 or Algorithm 3 (as described at the beginning of Section 3).
  - 5     Form  $\tilde{Y}_j = A^* Q_j$ .
  - 6     Compute a factorization  $\tilde{Y}_j = \tilde{Q}_j \tilde{R}_j$ , where the columns of  $\tilde{Q}_j$  are orthonormal and  $\tilde{R}_j$  is square, using Algorithm 1 or Algorithm 3 (as described at the beginning of Section 3).
  - 7 **end**
  - 8 Form  $Y = A\tilde{Q}_i$ .
  - 9 Compute a factorization  $Y = QR$ , where the columns of  $Q$  are orthonormal and  $R$  is square, using in this last step the double orthonormalization of Algorithms 2 and 4, not the single orthonormalization of Algorithms 1 and 3 (and, again, see the beginning of Section 3).
- 

---

**Algorithm 6:** Straightforward singular value decomposition (Algorithm 5.1 of [11])
 

---

**Input:** Matrices  $A$  and  $Q$  such that the spectral norm  $\|A - QQ^*A\|_2$  is small and the columns of  $Q$  are orthonormal (the matrix  $Q$  output from Algorithm 5 is suitable for the input here)

**Output:** Matrices  $U$ ,  $\Sigma$ , and  $V$  such that the spectral norm  $\|A - U\Sigma V^*\|_2$  is small, the columns of  $U$  are orthonormal, as are the columns of  $V$ , and  $\Sigma$  is diagonal and its entries are nonnegative

- 1 Form the matrix  $B = Q^*A$ .
  - 2 Compute a singular value decomposition,  $B = \tilde{U}\Sigma V^*$ , where the columns of  $\tilde{U}$  are orthonormal, as are the columns of  $V$ , and  $\Sigma$  is diagonal and its entries are nonnegative.
  - 3 Form the matrix  $U = Q\tilde{U}$ .
-

**Table 1** Meanings of the headings in the other tables

Heading	Meaning
$m$	number of rows in the matrix being decomposed or approximated
$n$	number of columns in the matrix being decomposed or approximated
$l$	rank of the approximation being constructed in Algorithms 5 and 6 (for the tables using Algorithms 5 and 6)
$i$	number of iterations used in Algorithm 5 (for the tables using Algorithms 5 and 6)
Algorithm	specifies the number of the algorithm used (or “pre-existing” for the original implementation in Spark)
Aux. Algs.	specifies the numbers of the auxiliary algorithms used in Algorithm 5 for the tables using Algorithms 5 and 6 (or “pre-existing” when using the original implementation in Spark)
CPU Time	sum over all CPU cores in all executors of the time in seconds spent actually processing
Wall-Clock	sum over all executors of the time in seconds that they were reserved
$\ A - U\Sigma V^*\ _2$	spectral norm of the discrepancy between the computed approximation $U\Sigma V^*$ and the matrix $A$ being decomposed or approximated; <i>please note that our setting for the working precision largely determines this error — see Remark 1 and the steps in the algorithms, “Discard. . . .”</i>
MaxEntry( $ U^*U - I $ )	maximal absolute value of the entries in the difference between $U^*U$ and the identity matrix $I$ , where $U\Sigma V^*$ is the computed approximation
MaxEntry( $ V^*V - I $ )	maximal absolute value of the entries in the difference between $V^*V$ and the identity matrix $I$ , where $U\Sigma V^*$ is the computed approximation

**Table 2** Settings for Spark

Parameter	Setting
spark.dynamicAllocation.maxExecutors	180
spark.executor.cores	30
spark.executor.memory	60g
rowsPerPart (in a BlockMatrix) <sup>†</sup>	1024
colsPerPart (in a BlockMatrix)	1024
Spark version	2.0.1
total machines available	200
BLAS-LAPACK library	Intel MKL

<sup>†</sup>This is also the number of rows in a block of the resilient distributed dataset that underlies an IndexedRowMatrix. Our software converts the matrix in formulae (2) and (5) from a BlockMatrix to an IndexedRowMatrix whenever necessary, which preserves the number of rows per block.

**Table 3**  $m = 1,000,000$ ;  $n = 2,000$ 

Algorithm	CPU Time	Wall-Clock	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1	1.48E+04	1.48E+04	9.76E-12	6.84E-06	3.51E-15
2	6.84E+04	9.01E+04	9.76E-12	6.44E-13	4.68E-15
3	1.33E+04	1.67E+04	9.92E-08	6.20E-04	1.73E-14
4	1.36E+04	2.52E+04	9.64E-07	1.10E-14	2.90E-15
pre-existing	1.12E+04	1.28E+04	1.83E-09	2.34E-00	3.12E-15

**Table 4**  $m = 100,000$ ;  $n = 2,000$ 

Algorithm	CPU Time	Wall-Clock	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1	1.59E+03	1.02E+03	9.76E-12	5.47E-06	3.22E-15
2	6.85E+03	3.39E+03	9.76E-12	6.85E-13	4.06E-15
3	1.32E+03	9.19E+02	9.92E-08	3.11E-04	1.22E-14
4	1.58E+03	1.30E+03	9.64E-07	6.66E-15	2.69E-15
pre-existing	1.27E+03	9.68E+02	2.75E-15	9.91E-01	2.50E-15

**Table 5**  $m = 10,000$ ;  $n = 2,000$ 

Algorithm	CPU Time	Wall-Clock	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1	3.86E+02	8.40E+01	9.76E-12	4.35E-06	3.55E-15
2	9.26E+02	1.42E+02	9.76E-12	7.67E-12	3.19E-15
3	2.52E+02	5.60E+01	9.92E-08	2.15E-04	1.82E-14
4	3.16E+02	8.40E+01	9.64E-07	6.66E-15	3.33E-15
pre-existing	2.15E+02	7.30E+01	1.89E-15	9.97E-01	2.57E-15

**Table 6** Algorithm 6;  $m = 1,000,000$ ;  $n = 2,000$ ;  $l = 20$ ;  $i = 2$ 

Aux. Algs.	CPU Time	Wall-Clock	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1 then 2	3.06E+03	8.80E+03	2.64E-12	4.44E-15	8.88E-16
3 then 4	2.80E+03	9.94E+03	4.83E-07	3.77E-15	5.55E-16
pre-existing	6.06E+03	1.16E+04	3.36E-10	1.00E-00	6.66E-16

**Table 7** Algorithm 6;  $m = 100,000$ ;  $n = 2,000$ ;  $l = 20$ ;  $i = 2$ 

Aux. Algs.	CPU Time	Wall-Clock	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1 then 2	3.28E+02	4.78E+02	2.64E-12	3.11E-15	1.44E-15
3 then 4	4.33E+02	4.71E+02	4.83E-07	1.55E-15	8.36E-16
pre-existing	6.17E+02	4.92E+02	3.36E-10	1.00E-00	4.44E-16

**Table 8** Algorithm 6;  $m = 10,000$ ;  $n = 2,000$ ;  $l = 20$ ;  $i = 2$ 

Aux. Algs.	CPU Time	Wall-Clock	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1 then 2	7.20E+01	7.50E+01	2.64E-12	2.22E-15	1.89E-15
3 then 4	8.00E+01	9.30E+01	4.83E-07	6.66E-16	6.66E-16
pre-existing	1.18E+02	9.40E+01	3.36E-10	1.00E-00	6.66E-16

**Table 9** Timings for Algorithm 6;  $l = 10$ ;  $i = 2$ 

Aux. Algs.	$m$	$n$	CPU Time	Wall-Clock
1 then 2	100,000	100,000	1.04E+04	4.88E+03
3 then 4	100,000	100,000	9.52E+03	7.41E+03
1 then 2	1,000,000	10,000	9.11E+03	1.05E+04
3 then 4	1,000,000	10,000	9.56E+03	1.01E+04
1 then 2	100,000	10,000	1.10E+03	5.40E+02
3 then 4	100,000	10,000	1.02E+03	4.93E+02

**Table 10** Errors for Algorithm 6;  $l = 10$ ;  $i = 2$ 

Aux. Algs.	$m$	$n$	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1 then 2	100,000	100,000	7.74E-12	6.66E-16	1.78E-15
3 then 4	100,000	100,000	2.15E-07	7.77E-16	1.33E-15
1 then 2	1,000,000	10,000	7.74E-12	3.00E-15	7.77E-16
3 then 4	1,000,000	10,000	2.15E-07	2.89E-15	4.44E-16
1 then 2	100,000	10,000	7.74E-12	1.22E-15	9.99E-16
3 then 4	100,000	10,000	2.15E-07	2.86E-16	4.44E-16

**Table 11**  $m = 1,000,000$ ;  $n = 2,000$ ; restricted to ten times fewer executors

Algorithm	CPU Time	Wall-Clock	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1	9.23E+03	4.72E+03	9.76E-12	6.21E-06	3.00E-15
2	5.91E+04	5.44E+04	9.76E-12	6.75E-13	3.06E-15
3	7.36E+03	4.14E+03	9.92E-08	6.13E-04	1.38E-14
4	1.00E+04	7.72E+03	9.64E-07	1.02E-14	2.69E-15
pre-existing	6.54E+03	3.56E+03	1.79E-09	3.17E-00	3.96E-15

**Table 12**  $m = 100,000$ ;  $n = 2,000$ ; restricted to ten times fewer executors

Algorithm	CPU Time	Wall-Clock	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1	1.74E+03	8.76E+02	9.76E-12	5.30E-06	3.33E-15
2	7.08E+03	3.74E+03	9.76E-12	4.93E-13	3.89E-15
3	1.26E+03	7.36E+02	9.92E-08	2.33E-04	1.87E-14
4	1.62E+03	1.01E+03	9.64E-07	5.33E-15	3.33E-15
pre-existing	1.27E+03	8.13E+02	2.15E-15	9.92E-01	2.32E-15

**Table 13**  $m = 10,000$ ;  $n = 2,000$ ; restricted to ten times fewer executors

Algorithm	CPU Time	Wall-Clock	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1	4.02E+02	9.80E+01	9.76E-12	5.80E-06	3.67E-15
2	8.69E+02	1.70E+02	9.76E-12	2.65E-11	3.88E-15
3	2.04E+02	6.70E+01	9.92E-08	2.55E-04	1.73E-14
4	2.26E+02	9.00E+01	9.64E-07	5.33E-15	2.89E-15
pre-existing	1.86E+02	9.50E+01	2.45E-15	9.96E-01	2.36E-15

### A Restricting to ten times fewer executors

Tables 11–13, Tables 14–16, and Tables 17 and 18 display results analogous to those in Tables 3–5, Tables 6–8, and Tables 9 and 10, but with the number of executors, `spark.dynamicAllocation.maxExecutors`, set to 18 (rather than 180). The results are broadly comparable to those presented earlier. This indicates how the timings scale with the number of machines. Of course, other processing in Spark (not necessarily related to principal component analysis or singular value decomposition) can benefit from having the data stored over more executors, and moving data around the cluster can dominate the overall timings in real-world usage (see also Remark 2 in the introduction of the present paper).

**Table 14** Alg. 6;  $m = 1,000,000$ ;  $n = 2,000$ ;  $l = 20$ ;  $i = 2$ ; restricted to  $10\times$  fewer executors

Aux. Algs.	CPU Time	Wall-Clock	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1 then 2	2.48E+03	4.44E+03	2.64E-12	4.88E-15	1.22E-15
3 then 4	2.33E+03	4.47E+03	4.83E-07	3.33E-15	6.66E-16
pre-existing	5.56E+03	6.84E+03	3.36E-10	1.00E-00	6.66E-16

**Table 15** Alg. 6;  $m = 100,000$ ;  $n = 2,000$ ;  $l = 20$ ;  $i = 2$ ; restricted to  $10\times$  fewer executors

Aux. Algs.	CPU Time	Wall-Clock	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1 then 2	3.99E+02	4.10E+02	2.64E-12	2.89E-15	1.55E-15
3 then 4	3.28E+02	4.05E+02	4.83E-07	2.44E-15	8.88E-16
pre-existing	6.31E+02	5.17E+02	3.36E-10	1.00E-00	8.88E-16

**Table 16** Alg. 6;  $m = 10,000$ ;  $n = 2,000$ ;  $l = 20$ ;  $i = 2$ ; restricted to  $10\times$  fewer executors

Aux. Algs.	CPU Time	Wall-Clock	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1 then 2	7.60E+01	9.80E+01	2.64E-12	2.66E-15	1.55E-15
3 then 4	6.30E+01	7.40E+01	4.83E-07	2.22E-15	1.55E-15
pre-existing	1.21E+02	9.80E+01	3.36E-10	1.00E-00	6.66E-16

**Table 17** Timings for Alg. 6;  $l = 10$ ;  $i = 2$ ; restricted to  $10\times$  fewer executors

Aux. Algs.	$m$	$n$	CPU Time	Wall-Clock
1 then 2	100,000	100,000	1.04E+04	6.07E+03
3 then 4	100,000	100,000	1.02E+04	6.28E+03
1 then 2	1,000,000	10,000	9.36E+03	5.93E+03
3 then 4	1,000,000	10,000	9.38E+03	6.77E+03
1 then 2	100,000	10,000	1.01E+03	5.19E+02
3 then 4	100,000	10,000	1.01E+03	5.04E+02

**Table 18** Errors for Alg. 6;  $l = 10$ ;  $i = 2$ ; restricted to  $10\times$  fewer executors

Aux. Algs.	$m$	$n$	$\ A - U\Sigma V^*\ _2$	MaxEntry( $ U^*U - I $ )	MaxEntry( $ V^*V - I $ )
1 then 2	100,000	100,000	7.74E-12	1.55E-15	1.78E-15
3 then 4	100,000	100,000	2.15E-07	8.88E-16	1.78E-15
1 then 2	1,000,000	10,000	7.74E-12	1.55E-15	6.66E-16
3 then 4	1,000,000	10,000	2.15E-07	1.11E-15	1.11E-16
1 then 2	100,000	10,000	7.74E-12	2.00E-15	8.88E-16
3 then 4	100,000	10,000	2.15E-07	8.88E-16	7.77E-16

**Table 19** Timings for generating (2)–(3)

$m$	$n$	CPU Time	Wall-Time
1,000,000	2,000	4.76E+03	3.91E+03
100,000	2,000	4.50E+02	2.48E+02
10,000	2,000	5.00E+01	2.60E+01

**Table 20** Timings for generating (5)–(6) with  $l = 20$ 

$m$	$n$	CPU Time	Wall-Time
1,000,000	2,000	5.61E+02	1.37E+03
100,000	2,000	6.30E+01	7.80E+01
10,000	2,000	8.00E+00	1.70E+01

**Table 21** Timings for generating (5)–(6) with  $l = 10$ 

$m$	$n$	CPU Time	Wall-Time
100,000	100,000	7.30E+01	7.60E+01
1,000,000	10,000	4.93E+02	1.79E+03
100,000	10,000	4.20E+01	5.20E+01

## B Timings for generating the test matrices

For comparative purposes, Tables 19–21 list the times required to generate (2)–(3) and (5)–(6) using the settings in Table 2.



## References

1. Ailon, N., Rauhut, H.: Fast and RIP-optimal transforms. *Discrete Comput. Geom.* **52**(4), 780–798 (2014)
2. Ballard, G., Demmel, J., Dumitriu, I.: Minimizing communication for eigenproblems and the singular value decomposition. Tech. Rep. UCB/EECS-2011-14, Dept. EECS, UC Berkeley (2011)
3. Benson, A.R., Gleich, D.F., Demmel, J.: Direct QR factorizations for tall-and-skinny matrices in MapReduce architectures. In: Proc. IEEE Internat. Conf. Big Data, pp. 264–272. IEEE (2013)
4. Demmel, J., Dumitriu, I., Holtz, O.: Fast linear algebra is stable. *Numerische Mathematik* **108**(1), 59–91 (2007)
5. Demmel, J., Grigori, L., Gu, M., Xiang, H.: Communication-avoiding rank-revealing QR factorization with column pivoting. *SIAM J. Matrix. Anal. Appl.* **36**(1), 55–89 (2015)
6. Demmel, J., Grigori, L., Hoemmen, M., Langou, J.: Communication-optimal parallel and sequential QR and LU factorizations. *SIAM J. Sci. Comput.* **34**(1), 206–239 (2012)
7. Durstenfeld, R.: Algorithm 235: random permutation. *Communications of the ACM* **7**(7), 420 (1964)
8. Fukaya, T., Nakatsukasa, Y., Yanagisawa, Y., Yamamoto, Y.: CholeskyQR2: a simple and communication-avoiding algorithm for computing a tall-skinny QR factorization on a large-scale parallel system. In: Proc. 5th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems, pp. 31–38. IEEE (2014)
9. Gittens, A., Devarakonda, A., Racah, E., Ringenburt, M., Gerhardt, L., Kottalam, J., Liu, J., Maschhoff, K., Canon, S., Chhugani, J., Sharma, P., Yang, J., Demmel, J., Harrell, J., Krishnamurthy, V., Mahoney, M.W., Prabhat: Matrix factorization at scale: a comparison of scientific data analytics in Spark and C+MPI using three case studies. In: Proc. 2016 IEEE International Conference on Big Data, pp. 204–213. IEEE (2016)
10. Golub, G.H., Mahoney, M.W., Drineas, P., Lim, L.: Bridging the gap between numerical linear algebra, theoretical computer science, and data applications. *SIAM News* **39**(8), 1–3 (2006)
11. Halko, N., Martinsson, P.G., Tropp, J.: Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review* **53**(2), 217–288 (2011)
12. Jolliffe, I.T.: *Principal Component Analysis*, 2nd edn. Springer Series in Statistics. Springer-Verlag, New York, NY (2002)
13. Lehoucq, R., Sorensen, D., Yang, C.: *ARPACK User’s Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*. SIAM, Philadelphia, PA (1998)
14. Li, H., Linderman, G., Szlam, A., Stanton, K., Kluger, Y., Tygert, M.: Algorithm 971: an implementation of a randomized algorithm for principal component analysis. *ACM Trans. Math. Software* **43**(3), 28:1–28:14 (2016)
15. Linden, A., Krensky, P., Hare, J., Idoine, C.J., Sicular, S., Vashisth, S.: Magic quadrant for data science platforms. Tech. Rep. G00301536, Gartner (2017)
16. Shabat, G., Shmueli, Y., Aizenbud, Y., Averbuch, A.: Randomized LU decomposition. *Appl. Comput. Harmon. Anal.* (2016). DOI 10.1016/j.acha.2016.04.006. To appear
17. Stathopoulos, A., Wu, K.: A block orthogonalization procedure with constant synchronization requirements. *SIAM J. Sci. Comput.* **23**(6), 2165–2182 (2002)
18. Yamazaki, I., Tomov, S., Dongarra, J.: Mixed-precision Cholesky QR factorization and its case studies on multicore CPU with multiple GPUs. *SIAM J. Sci. Comput.* **37**(3), C307–C330 (2015)
19. Yamazaki, I., Tomov, S., Dongarra, J.: Stability and performance of various singular value QR implementations on multicore CPU with a GPU. *ACM Trans. Math. Software* **43**(2), 10:1–10:18 (2016)