

ACCURATE PRINCIPAL COMPONENT ANALYSIS VIA A FEW ITERATIONS OF ALTERNATING LEAST SQUARES

ARTHUR SZLAM, ANDREW TULLOCH, AND MARK TYGERT

Abstract. A few iterations of alternating least squares with a random starting point provably suffice to produce nearly optimal spectral- and Frobenius-norm accuracies of low-rank approximations to a matrix; iterating to convergence is unnecessary. Thus, software implementing alternating least squares can be retrofitted via appropriate setting of parameters to calculate nearly optimally accurate low-rank approximations highly efficiently, with no need for convergence.

Key words. low-rank approximation, principal component analysis, alternating least squares, alternating minimization, randomized algorithm

AMS subject classifications. 65F30, 65F15, 65D99

1. Introduction. Low-rank approximations are popular throughout the sciences and engineering, often in the form of principal component analysis, and converting any low-rank approximation to a singular value decomposition or principal component analysis is trivial and efficient, as detailed, for example, by [2], [1], or [3]. To calculate an accurate approximation to a matrix A , we consider the low-rank approximations

$$(1) \quad A \approx S_0 T_0, \quad A \approx S_1 T_1, \quad A \approx S_2 T_2, \quad \dots,$$

with S_i being a tall and skinny matrix and T_i being a short and fat matrix, produced via iterations starting from S_0 — iterations called “alternating least squares” by [4] (among others): for each $i = 0, 1, 2, \dots$, having S_i already, we obtain T_i minimizing the norm

$$(2) \quad \|S_i T_i - A\|,$$

then, having T_i already, we obtain S_{i+1} minimizing the norm

$$(3) \quad \|S_{i+1} T_i - A\|,$$

where these norms denote the spectral or Frobenius norms (the Frobenius norm of a matrix is the square root of the sum of the squares of the absolute values of the entries of the matrix); specifically, we use the minimizers

$$(4) \quad T_i = S_i^{(-1)} A$$

and

$$(5) \quad S_{i+1} = A T_i^{(-1)}$$

for $i = 0, 1, 2, \dots$, where

$$(6) \quad S_i^{(-1)} = (S_i^* S_i)^{-1} S_i^*$$

and

$$(7) \quad T_i^{(-1)} = T_i^* (T_i T_i^*)^{-1}$$

(see Section 2 below for precise definitions, particularly for the inverse and pseudo-inverse). The appendix reviews the well-known fact that these minimize both the spectral and Frobenius norms.

Following [2], we demonstrate that the approximations attain high accuracy after just a few of these iterations. Specifically, the remainder of the present paper has the following structure: Section 2 sets notational conventions used throughout the paper. Via mathematical analysis, Section 3 proves the high accuracy. Section 4 illustrates the high accuracy via numerical examples with a Matlab prototype available at <http://tygert.com/software.html>

The accompanying prototype comes complete with a comprehensive collection of tests, but is all in Matlab. The present paper provides a fully rigorous basis for more general software packages implementing alternating least squares to be retrofitted via appropriate setting of parameters to calculate nearly optimally accurate low-rank approximations, with no need to wait for convergence.

2. Notation. This section sets our notational conventions. For any full-rank square matrix A , we use A^{-1} to denote the inverse of A . For any rank-deficient square matrix A , we use A^{-1} to denote the pseudoinverse of A ; the pseudoinverse of A is the matrix representing the inverse of A with its domain restricted to the row space of A (plus the identically zero map restricted to the null space of A). Needless to say, for a full-rank square matrix, the pseudoinverse is the same as the inverse.

For any matrix A , we denote by A^* the adjoint (that is, the conjugate transpose) of A , so that the spectral norm of A is given by the action of A on vectors via

$$(8) \quad \|A\|_2 = \sqrt{\max_{v: v^*v=1} v^*A^*Av},$$

and the Frobenius norm $\|A\|_F$ of A is the square root of the sum of the squares of the absolute values of the entries of A . The spectral and Frobenius norms of a vector viewed as a matrix with a single row or column are the same, and are also known as the Euclidean norm of the vector. A definition equivalent to (8) is

$$(9) \quad \|A\|_2 = \max_{v: \|v\|_2=1} \|Av\|_2,$$

where the norms of the vectors are the Euclidean norms.

3. Analysis of accuracy. This section demonstrates that the procedure — alternating least squares — described in the introduction produces a highly accurate approximation $S_i T_i$ to the given matrix A even for a small number i of iterations, provided that S_0 is one of the random matrices used by [2] (for example, the entries of S_0 can be independent and identically distributed standard normal variates). The demonstration is simply a reduction to the proof of accuracy for similar algorithms by [2]; we leave the brunt of the proof (together with a discussion of the intuitions behind the proof) to [2]. We begin by proving several lemmas.

The following lemma provides an explicit expression for S_i from the iterations in (4) and (5).

LEMMA 1. *Given matrices A and S_0 for the iterations in (4) and (5), the matrix S_i coming from those iterations can be expressed as*

$$(10) \quad S_i = (AA^*)^i S_0 B_0 B_1 B_2 \cdots B_{i-1}$$

for $i = 1, 2, 3, \dots$, where we define

$$(11) \quad B_i = (S_i^* AA^* S_i)^{-1} S_i^* S_i.$$

Proof. Combining (4)–(7) and (11) yields

$$(12) \quad S_{i+1} = AA^*S_i(S_i^*AA^*S_i)^{-1}S_i^*S_i = AA^*S_iB_i$$

for $i = 0, 1, 2, \dots$. Iterating the recurrence in (12) yields (10). \square

The following lemma follows straightforwardly from using singular value decompositions.

LEMMA 2. *Suppose that A , S , and T are matrices such that*

$$(13) \quad T = S^{(-1)}A,$$

where

$$(14) \quad S^{(-1)} = (S^*S)^{-1}S^*.$$

*Then, the ranks of S^*A , T , and AT^* are all equal.*

Proof. Combining (13) and (14) yields that

$$(15) \quad T = (S^*S)^{-1}S^*A$$

and

$$(16) \quad TA^* = (S^*S)^{-1}S^*AA^*.$$

We form the full singular value decompositions

$$(17) \quad A = U_A\Sigma_A V_A^*$$

and

$$(18) \quad S = U_S\Sigma_S V_S^*,$$

where U_A , V_A , U_S , and V_S are unitary, and all entries of Σ_A and Σ_S are nonnegative and are zero off the main diagonals. Combining (15)–(18) yields

$$(19) \quad S^*A = V_S\Sigma_S^*U_S^*U_A\Sigma_A V_A^*,$$

$$(20) \quad T = V_S\Sigma_S^{(-1)}U_S^*U_A\Sigma_A V_A^*,$$

and

$$(21) \quad TA^* = V_S\Sigma_S^{(-1)}U_S^*U_A\Sigma_A^{(2)}U_A^*,$$

where $\Sigma_S^{(-1)}$ is the same as Σ_S^* , but replacing its nonzero diagonal entries with their reciprocals, and where $\Sigma_A^{(2)} = \Sigma_A\Sigma_A^*$ is the square diagonal matrix with the squares of the diagonal entries of Σ_A on its diagonal.

Combining (19)–(21) and the fact that U_A , V_A , U_S , and V_S are unitary yields

$$(22) \quad \text{rank}(S^*A) = \text{rank}(\Sigma_S^*W\Sigma_A),$$

$$(23) \quad \text{rank}(T) = \text{rank}(\Sigma_S^{(-1)}W\Sigma_A),$$

and

$$(24) \quad \text{rank}(TA^*) = \text{rank}(\Sigma_S^{(-1)}W\Sigma_A^{(2)}),$$

where W is the unitary matrix

$$(25) \quad W = U_S^*U_A.$$

The claim stated in the lemma (that the ranks of S^*A , T , and AT^* are all equal) then follows from the combination of (22)–(24) and the facts that $\Sigma_S^{(-1)}W\Sigma_A$ is the same as $\Sigma_S^*W\Sigma_A$ with its rows rescaled by nonzero multiples (so that they have the same row space), and that (assuming A is square) $\Sigma_S^{(-1)}W\Sigma_A^{(2)}$ is the same as $\Sigma_S^{(-1)}W\Sigma_A$ with its columns rescaled by nonzero multiples (so that they have the same column space); of course, the rank of a matrix is equal to the dimension of its row space (which is the same as the dimension of its column space). If A is not square, then either $\Sigma_S^{(-1)}W\Sigma_A^{(2)}$ is the same as $\Sigma_S^{(-1)}W\Sigma_A$ augmented by columns of zeros and with its columns rescaled by nonzero multiples or $\Sigma_S^{(-1)}W\Sigma_A$ is the same as $\Sigma_S^{(-1)}W\Sigma_A^{(2)}$ augmented by columns of zeros and with its columns rescaled by nonzero multiples (so that again they have the same column space, which is the same as the column space of the original, unaugmented $\Sigma_S^{(-1)}W\Sigma_A$ or $\Sigma_S^{(-1)}W\Sigma_A^{(2)}$). \square

The adjoint of the preceding lemma is the following.

COROLLARY 3. *Suppose that A , S , and T are matrices such that*

$$(26) \quad S = AT^{(-1)},$$

where

$$(27) \quad T^{(-1)} = T^*(TT^*)^{-1}.$$

Then, the ranks of AT^ , S , and S^*A are all equal.*

The following lemma follows from using both Lemma 2 and Corollary 3 repeatedly.

LEMMA 4. *Given the matrices A , T_i , and S_{i+1} from (4) and (5), the ranks of S_0^*A , T_0 , AT_0^* , S_1 , S_1^*A , T_1 , AT_1^* , S_2 , S_2^*A , T_2 , AT_2^* , S_3 , S_3^*A , T_3 , AT_3^* , \dots are all equal.*

Proof. This lemma follows from induction on $i = 0, 1, 2, \dots$, using Lemma 2 and Corollary 3. \square

The following theorem follows from (10) and Lemma 4.

THEOREM 5. *Given matrices A and S_0 for the iterations in (4) and (5), the column space of the matrix S_i coming from those iterations is the same as the column space of $(AA^*)^iS_0$, for $i = 1, 2, 3, \dots$.*

Proof. As seen from (10), the column space of S_i is a subspace of the column space of $(AA^*)^iS_0$. Moreover, the row space of $(AA^*)^iS_0$ is a subspace of the row space of A^*S_0 , so

$$(28) \quad \text{rank}((AA^*)^iS_0) \leq \text{rank}(A^*S_0).$$

As already mentioned, the column space of S_i is a subspace of the column space of $(AA^*)^iS_0$; if the subspace were not the whole space, then

$$(29) \quad \text{rank}(S_i) < \text{rank}((AA^*)^iS_0),$$

and then combining (28) and (29) would yield

$$(30) \quad \text{rank}(S_i) < \text{rank}(A^*S_0) = \text{rank}(S_0^*A),$$

contradicting Lemma 4 . . . thus, the claim stated in the theorem must be true. \square

Finally, calculating T_i minimizing (2) constructs the best approximation S_iT_i to A such that the column space of the approximation lies in the column space of S_i — which is the same as the column space of $(AA^*)^iS_0$, as Theorem 5 proves — where “best” means minimizing the discrepancy in the spectral norm, which is the same as minimizing the discrepancy in the Frobenius norm, as reviewed in the appendix. This produces a highly accurate approximation S_iT_i to A even for a small number i of iterations, as proven by [2], provided that S_0 is one of the random matrices used by [2] (for example, the entries of S_0 can be independent and identically distributed standard normal variates) — iterating until convergence is unnecessary.

4. Numerical examples. This section presents several numerical experiments on an implementation in Matlab of the algorithm (alternating least squares) discussed in the introduction. Although the numerical experiments discussed here are somewhat limited in order to keep the presentation succinct, the codes together with software extensively testing them are available at <http://tygert.com/software.html>

We consider various values for positive integers m and n , as specified in the captions for Tables 1–3, and calculate rank- k approximations to the $m \times n$ matrix

$$(31) \quad A = F\Sigma G,$$

where F and G are $m \times m$ and $n \times n$ unitary discrete Fourier transforms, respectively, and Σ is an $m \times n$ matrix whose entries are all zeros except for the diagonal entries

$$(32) \quad \Sigma_{i,i} = \delta^{\lfloor i/2 \rfloor / (k/2)}$$

for $i = 1, 2, \dots, k$, and

$$(33) \quad \Sigma_{i,i} = \delta \cdot \frac{\min\{m, n\} - i}{\min\{m, n\} - k - 1}$$

for $i = k + 1, k + 2, \dots, \min\{m, n\}$ ($\lfloor i/2 \rfloor$ is the greatest integer less than or equal to $i/2$); the tables below specify various values for k and δ . Thus, the spectral norm of A is 1:

$$(34) \quad \|A\|_2 = 1.$$

The headings of Tables 1–3 have the following meanings:

- j is the number of iterations conducted.
- k is the rank of the approximation constructed — the number of columns in S_i from (2), which is also the number of rows in T_i from (2).
- δ is the spectral-norm accuracy of the best possible rank- k approximation.
- ϵ is the spectral-norm accuracy of the calculated rank- k approximation, with the spectral-norm accuracy computed via 100 iterations of the power method.
- t is the time in seconds required to compute the approximation (without using any fast Fourier transforms to leverage the special structure of the matrix A).

The tables illustrate the importance of using at least one (preferably two or more) iterations, as then the accuracy (ϵ) of the computed approximation is nearly the best possible (δ). The accuracies are indeed excellent, even with just a couple iterations. The timings scale as expected, roughly in proportion to the number of entries in the matrices; we used Matlab version R2015B on an Apple MacBook Pro with a 2.6 GHz Intel Core i7 processor.

TABLE 1
 $m = 2048, n = 4096$

j	k	δ	ϵ	t
0	2	1e-03	1.4e-02	5.8e-01
1	2	1e-03	1.0e-03	8.1e-01
2	2	1e-03	1.0e-03	1.2e+00
10	2	1e-03	1.0e-03	4.5e+00
0	10	1e-03	1.8e-02	7.5e-01
1	10	1e-03	1.2e-03	1.4e+00
2	10	1e-03	1.0e-03	2.1e+00
10	10	1e-03	1.0e-03	8.0e+00
0	2	1e-11	1.3e-10	4.1e-01
1	2	1e-11	1.0e-11	7.8e-01
2	2	1e-11	1.0e-11	1.2e+00
10	2	1e-11	1.0e-11	4.3e+00
0	10	1e-11	2.4e-10	7.5e-01
1	10	1e-11	1.0e-11	1.4e+00
2	10	1e-11	1.0e-11	2.1e+00
10	10	1e-11	1.0e-11	8.0e+00

TABLE 2
 $m = 4096, n = 4096$

j	k	δ	ϵ	t
0	2	1e-03	1.5e-02	9.6e-01
1	2	1e-03	1.0e-03	1.9e+00
2	2	1e-03	1.0e-03	2.8e+00
10	2	1e-03	1.0e-03	1.0e+01
0	10	1e-03	2.2e-02	1.8e+00
1	10	1e-03	1.3e-03	3.6e+00
2	10	1e-03	1.0e-03	5.4e+00
10	10	1e-03	1.0e-03	2.1e+01
0	2	1e-11	2.6e-10	9.6e-01
1	2	1e-11	1.0e-11	1.9e+00
2	2	1e-11	1.0e-11	2.9e+00
10	2	1e-11	1.0e-11	1.0e+01
0	10	1e-11	4.2e-10	1.8e+00
1	10	1e-11	1.0e-11	3.6e+00
2	10	1e-11	1.0e-11	5.5e+00
10	10	1e-11	1.0e-11	2.0e+01

TABLE 3
 $m = 4096, n = 8192$

j	k	δ	ϵ	t
0	2	1e-03	1.2e-02	2.0e+00
1	2	1e-03	1.0e-03	3.7e+00
2	2	1e-03	1.0e-03	5.7e+00
10	2	1e-03	1.0e-03	2.0e+01
0	10	1e-03	2.1e-02	3.8e+00
1	10	1e-03	1.4e-03	7.5e+00
2	10	1e-03	1.0e-03	1.1e+01
10	10	1e-03	1.0e-03	4.1e+01
0	2	1e-11	1.5e-10	1.8e+00
1	2	1e-11	1.0e-11	3.8e+00
2	2	1e-11	1.0e-11	5.6e+00
10	2	1e-11	1.0e-11	2.1e+01
0	10	1e-11	5.3e-10	3.6e+00
1	10	1e-11	1.0e-11	7.1e+00
2	10	1e-11	1.0e-11	1.1e+01
10	10	1e-11	1.0e-11	4.0e+01

Appendix A. Common minimizers for the spectral & Frobenius norms.

This appendix reviews the fact that, given matrices A and S , one matrix T minimizing the norm

$$(35) \quad \|ST - A\|,$$

with the norm being the spectral norm or the Frobenius norm, is

$$(36) \quad T = S^{(-1)}A,$$

where the so-called “pseudoinverse” of S is

$$(37) \quad S^{(-1)} = (S^*S)^{-1}S^*,$$

with the inverse and pseudoinverse defined in Section 2.

Indeed, for the spectral norm, for any T — not just that in (36),

$$(38) \quad \begin{aligned} \|ST - A\|_2^2 &= \|(ST - A)^*(ST - A)\|_2 \\ &= \left\| \left(\left(\frac{SS^{(-1)}}{I - SS^{(-1)}} \right) (ST - A) \right)^* \left(\frac{SS^{(-1)}}{I - SS^{(-1)}} \right) (ST - A) \right\|_2 \\ &= \left\| \left(\frac{SS^{(-1)}}{I - SS^{(-1)}} \right) (ST - A) \right\|_2^2 = \left\| \left(\frac{ST - SS^{(-1)}A}{SS^{(-1)}A - A} \right) \right\|_2^2. \end{aligned}$$

The definition of the spectral norm in (8) yields

$$(39) \quad \|SS^{(-1)}A - A\|_2^2 \leq \left\| \left(\frac{ST - SS^{(-1)}A}{SS^{(-1)}A - A} \right) \right\|_2^2 \leq \|ST - SS^{(-1)}A\|_2^2 + \|SS^{(-1)}A - A\|_2^2.$$

Combining (38) and (39) yields

$$(40) \quad \|SS^{(-1)}A - A\|_2 \leq \|ST - A\|_2 \leq \sqrt{\|ST - SS^{(-1)}A\|_2^2 + \|SS^{(-1)}A - A\|_2^2},$$

so that (35) is minimal for the spectral norm when

$$(41) \quad ST = SS^{(-1)}A.$$

For the Frobenius norm, for any T ,

$$(42) \quad \|ST - A\|_F^2 = \sum_{k=1}^n \|St^{[k]} - a^{[k]}\|_2^2,$$

where $t^{[1]}, t^{[2]}, \dots, t^{[n]}$ are the columns of T , and $a^{[1]}, a^{[2]}, \dots, a^{[n]}$ are the columns of A . Using the above result for the spectral norm with $t^{[k]}$ replacing T and with $a^{[k]}$ replacing A , the right-hand side of (42) is minimal when

$$(43) \quad St^{[k]} = SS^{(-1)}a^{[k]}$$

for $k = 1, 2, \dots, n$, which happens to be equivalent to (41) for the full A for all its columns simultaneously.

Thus, for both the spectral and Frobenius norms, (35) is minimal when (41) holds, and (41) clearly holds for T defined in (36).

[A similar argument uses the identity

$$\begin{aligned}
(44) \quad & (ST - A)^*(ST - A) - \left((I - SS^{(-1)})A \right)^* (I - SS^{(-1)})A \\
& = (ST - A)^*(ST - A) - A^*(I - SS^{(-1)})A = (ST - A)^*SS^{(-1)}(ST - A), \\
& = (S^*(ST - A))^* (S^*S)^{-1} (S^*(ST - A)),
\end{aligned}$$

the fact that the right-hand side of (44) is nonnegative definite, and the relations

$$(45) \quad \|ST - A\|_2^2 = \max_{v: v^*v=1} v^*(ST - A)^*(ST - A)v,$$

$$(46) \quad \|(I - SS^{(-1)})A\|_2^2 = \max_{v: v^*v=1} v^* \left((I - SS^{(-1)})A \right)^* (I - SS^{(-1)})Av,$$

$$(47) \quad \|ST - A\|_F^2 = \sum_{k=1}^n (e^{[k]})^*(ST - A)^*(ST - A)e^{[k]},$$

and

$$(48) \quad \|(I - SS^{(-1)})A\|_F^2 = \sum_{k=1}^n (e^{[k]})^* \left((I - SS^{(-1)})A \right)^* (I - SS^{(-1)})Ae^{[k]},$$

where $e^{[1]}, e^{[2]}, \dots, e^{[n]}$ are the unit basis vectors, with $e^{[k]}$ being the column vector of all zeros, except for its k th entry, which is 1.]

REFERENCES

- [1] M. GU, *Subspace iteration randomization and singular value problems*, SIAM J. Sci. Comput., 37 (2015), pp. A1139–A1173.
- [2] N. HALKO, P.-G. MARTINSSON, AND J. TROPP, *Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Review, 53 (2011), pp. 217–288.
- [3] D. WOODRUFF, *Sketching as a Tool for Numerical Linear Algebra*, vol. 10 of Foundations and Trends in Theoretical Computer Science, Now publishers, 2014.
- [4] F. W. YOUNG, Y. TAKANE, AND J. DE LEEUW, *The principal components of mixed measurement level multivariate data: an alternating least squares method with optimal scaling features*, Psychometrika, 43 (1978), pp. 279–281.