

Algorithm 805: Computation and Uses of the Semidiscrete Matrix Decomposition

TAMARA G. KOLDA

Sandia National Laboratories

and

DIANNE P. O'LEARY

University of Maryland

We present algorithms for computing a semidiscrete approximation to a matrix in a weighted norm, with the Frobenius norm as a special case. The approximation is formed as a weighted sum of outer products of vectors whose elements are ± 1 or 0, so the storage required by the approximation is quite small. We also present a related algorithm for approximation of a tensor. Applications of the algorithms are presented to data compression, filtering, and information retrieval; software is provided in C and in Matlab.

Categories and Subject Descriptors: D.3.2 [**Programming Languages**]: Language Classifications—*C*; *Matlab*; G.1.2 [**Numerical Analysis**]: Approximation; H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval

General Terms: Algorithms

Additional Key Words and Phrases: Singular value decomposition, semidiscrete decomposition, latent semantic indexing, compression, matrix decomposition

Kolda's work was performed while she was a postdoctoral fellow at Oak Ridge National Laboratory and was supported through the Applied Mathematical Sciences Research Program, Office of Energy Research, U.S. Department of Energy, under contract DE-AC05-96OR22464 with Lockheed Martin Energy Research Corporation. O'Leary's work was supported by the National Science Foundation under Grant CCR-97-32022 and by the Department Informatik, ETH Zürich, Switzerland. An earlier version of this paper appeared as University of Maryland Technical Report CS-TR-4012/UMIACS-TR-99-22 (April 1999) and Oak Ridge National Laboratory Technical Memorandum ORNL-13766 (April, 1999).

Authors' addresses: T. G. Kolda, Computational Science and Mathematics Research Department, Sandia National Laboratories, Livermore, CA 94551-9217; email: tgkolda@sandia.gov; D. P. O'Leary, Computer Science Department and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742; email: oleary@cs.umd.edu.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2001 ACM 0098-3500/00/0900-0415 \$5.00

1. INTRODUCTION

A semidiscrete decomposition (SDD) expresses a matrix as a weighted sum of outer products formed by vectors with entries constrained to be in the set $\mathcal{S} = \{-1, 0, 1\}$. O’Leary and Peleg [1983] introduced the SDD approximation in the context of image compression, and Kolda and O’Leary [1998; 1999b] used the SDD approximation for latent semantic indexing (LSI) in information retrieval; these applications are discussed in Section 4.

The primary advantage of SDD approximations over other types of matrix approximations such as the truncated singular value decomposition (SVD) is that, as we will demonstrate with numerical examples in Section 6, it typically provides a more accurate approximation for far less storage.

The SDD approximation of an $m \times n$ matrix A is a decomposition of the form

$$A_k = \underbrace{[x_1 x_2 \cdots x_k]}_{X_k} \underbrace{\begin{bmatrix} d_1 & 0 & \cdots & 0 \\ 0 & d_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & d_k \end{bmatrix}}_{D_k} \underbrace{\begin{bmatrix} y_1^T \\ y_2^T \\ \vdots \\ y_k^T \end{bmatrix}}_{Y_k^T} = \sum_{i=1}^k d_i x_i y_i^T.$$

Here each x_i is an m -vector with entries from the set $\mathcal{S} = \{-1, 0, 1\}$, each y_i is an n -vector with entries from the set \mathcal{S} , and each d_i is a positive scalar. We call this a k -term SDD approximation.

Although every matrix can be expressed as an mn -term SDD approximation

$$A = \sum_{i=1}^m \sum_{j=1}^n a_{ij} e_i e_j^T,$$

where e_k is the k th unit vector, the usefulness of an SDD approximation is in developing approximations that have far fewer terms, and we focus on algorithms for computing such approximating SDDs.

Since the storage requirement for a k -term SDD approximation is k floating-point numbers plus $k(m + n)$ entries from \mathcal{S} , it is inexpensive to store quite a large number of terms. For example, for a dense, single precision matrix of size $10,000 \times 10,000$, almost 80,000 SDD terms can be stored in the space of the original data, and almost 160,000 terms can be stored for a double precision matrix of the same size.

In Section 2, we discuss an algorithm for computing a *weighted* SDD approximation to a matrix. Section 3 focuses on the tensor SDD to form an approximation to an $m_1 \times m_2 \times \cdots \times m_n$ tensor. Applications of these algorithms are given in Section 4. A storage-efficient implementation for the SDD approximation is presented in Section 5. Numerical results with our software are presented in Section 6.

2. THE WEIGHTED SDD APPROXIMATION

Let $A \in \Re^{m \times n}$ be a given matrix, and let $W \in \Re^{m \times n}$ be a given matrix of nonnegative weights. The weighted approximation problem is to find a matrix $B \in \Re^{m \times n}$ that solves

$$\min \|A - B\|_W^2,$$

subject to some constraints on B . Here the *weighted norm* $\|\cdot\|_W$ is defined as

$$\|A\|_W^2 = \sum_{i=1}^m \sum_{j=1}^n a_{ij}^2 w_{ij}.$$

If all the weights are one, then the norm is the Frobenius norm.

2.1 Computing the Weighted SDD Approximation

The case where B is a low-rank matrix has been considered by Gabriel and Zamir [1979] and others, and they obtain a solution with some similarities to the truncated singular value decomposition, although computation is much more expensive. We show how to iteratively generate a weighted semidiscrete approximation of the form $dx y^T$. Let A_k denote the k -term approximation ($A_0 \equiv 0$). Let R_k be the *residual* at the k th step, i.e., $R_k = A - A_{k-1}$. Then the optimal choice of the next triplet (d_k, x_k, y_k) is the solution to the subproblem

$$\min F_k(d, x, y) \equiv \|R_k - dx y^T\|_W^2 \text{ subject to } x \in S^m, y \in S^n, d > 0. \quad (1)$$

Here $R_k \equiv A - \sum_{i=1}^{k-1} d_i x_i y_i^T$ is the residual matrix. As with the regular SDD, this is a mixed integer programming problem that can be rewritten as an integer program. First, recall the definition of the *Hadamard* or *elementwise product* of matrices, i.e., $(A \circ B)_{ij} = a_{ij} b_{ij}$.

THEOREM 1. *Solving the mixed integer program (1) is equivalent to solving the integer program*

$$\max \tilde{F}_k(x, y) \equiv \frac{[x^T(R_k \circ W)y]^2}{(x \circ x)^T W (y \circ y)} \text{ subject to } x \in S^m, y \in S^n. \quad (2)$$

PROOF. We can eliminate d in (1) as follows. First rewrite $F_k(d, x, y)$ as

$$F_k(d, x, y) = \|R_k\|_W^2 - 2dx^T(R_k \circ W)y + d^2(x \circ x)^T W (y \circ y). \quad (3)$$

At the optimal solution, $\partial F_k / \partial d = 0$, so the optimal value of d is given by

$$d^* = \frac{x^T(R_k \circ W)y}{(x \circ x)^T W (y \circ y)}.$$

Substituting d^* in (3) yields

$$F_k(d^*, x, y) = \|R_k\|_W^2 - \frac{(x^T(R_k \circ W)y)^2}{(x \circ x)^T W(y \circ y)}. \quad (4)$$

Thus solving (4) is equivalent to solving (2). \square

The integer program (2) has $3^{(m+n)}$ feasible points, so the cost of an exhaustive search for the optimal solution grows exponentially with m and n . Rather than doing this, we use an *alternating algorithm* to generate an *approximate* solution to the subproblem. Assuming that y is fixed, \tilde{F}_k can be written as

$$\tilde{F}_k(x, y) = \frac{(x^T s)^2}{(x \circ x)^T v}, \quad (5)$$

where $s \equiv (R_k \circ W)y$ and $v \equiv W(y \circ y)$. To determine the maximum, $2^m - 1$ possibilities must be checked. This can be reduced to just checking m possibilities.

THEOREM 2. *For the integer program (5), if it is known that x has exactly J nonzeros, then the solution is given by*

$$x_{ij} = \begin{cases} \text{sign}(s_{ij}) & \text{if } 1 \leq j \leq J \\ 0 & \text{if } J + 1 \leq j \leq m, \end{cases}$$

where the pairs of (s_i, v_i) have been sorted so that

$$\frac{|s_{i_1}|}{v_{i_1}} \geq \frac{|s_{i_2}|}{v_{i_2}} \geq \dots \geq \frac{|s_{i_m}|}{v_{i_m}}.$$

PROOF. First note that if s_i is zero, then a nonzero value of x_i cannot affect the numerator of \tilde{F} , and $x_i = 0$ minimizes the denominator, so $x_i = 0$ is optimal. If $v_i = 0$, then $s_i = 0$, so choose $x_i = 0$. Therefore, we need only consider indices for which s_i and v_i are nonzero, and without loss of generality, we will assume that the s_i are all positive and ordered so that $i_j = j, j = 1, \dots, m$.

We complete the proof by showing that if the optimal solution has nonzeros with indices in some set I , and if $q \in I$ and $p < q$, then $p \in I$.

Assume to the contrary, and partition I into $I_1 \cup I_2$, where indices in I_1 are less than p and those in I_2 are greater than p . The case $p = 1$ is left to the reader; here we assume $p > 1$.

For ease of notation, let

$$S_1 = \sum_{i \in I_1} s_i, \quad V_1 = \sum_{i \in I_1} v_i,$$

and define S_2 and V_2 analogously.

By the ordering of the ratios s/v , we know that $s_i v_p < s_p v_i$ for all $i \in I_2$; therefore,

$$S_2 v_p < s_p V_2. \quad (6)$$

Since I is optimal, we know that

$$\frac{S_1^2}{V_1} \leq \frac{(S_1 + S_2)^2}{V_1 + V_2},$$

therefore, by cross-multiplying and canceling terms, we obtain

$$S_1^2 V_2 \leq S_2^2 V_1 + 2S_1 S_2 V_1. \quad (7)$$

Similarly,

$$\frac{(S_1 + S_2 + s_p)^2}{V_1 + V_2 + v_p} \leq \frac{(S_1 + S_2)^2}{V_1 + V_2},$$

so

$$\begin{aligned} & s_p^2(V_1 + V_2) + 2S_1 s_p(V_1 + V_2) + 2S_2 s_p(V_1 + V_2) \\ & \leq S_1^2 v_p + S_2^2 v_p + 2S_1 S_2 v_p \\ & \leq S_1^2 v_p + S_2 s_p V_2 + 2S_1 s_p V_2 \quad \text{by (6)} \\ & \leq (S_2^2 V_1 + 2S_1 S_2 V_1) \frac{v_p}{V_2} + S_2 s_p V_2 + 2S_1 s_p V_2 \quad \text{by (7)} \\ & \leq (S_2 s_p V_1 + 2S_1 s_p V_1) + S_2 s_p V_2 + 2S_1 s_p V_2 \quad \text{by (6)}. \end{aligned}$$

Canceling terms in this inequality we obtain

$$s_p^2(V_1 + V_2) + S_2 s_p(V_1 + V_2) \leq 0,$$

a contradiction. \square

The algorithm for the weighted SDD approximation, shown in Figure 1, is the same as the algorithm for the regular SDD approximation [O'Leary and Peleg 1983] if W is the all ones matrix. The method will generate a weighted approximation A_k for which $k = k_{\max}$ or $\|A - A_k\| < \rho_{\min}$. The work of each inner iteration is controlled by the parameters l_{\max} , the maximum number of allowed inner iterations, and α_{\min} , the relative improvement threshold. The weighted SDD algorithm can be used for either dense or sparse matrices. In the case that A or W is sparse, the approximation

- (1) Let R_k denote the residual, and initialize $R_1 \leftarrow A$.
 Let $\rho_k = \|R_k\|_W^2$ be the norm of the residual, and initialize $\rho_1 \leftarrow \|R_1\|_W^2$.
 Let A_k denote the k -term approximation, and initialize $A_0 \leftarrow 0$.
 Choose k_{\max} , the maximum number of terms in the approximation.
 Choose ρ_{\min} , the desired accuracy of the approximation.
 Choose l_{\max} , the maximum allowable inner iterations.
 Choose α_{\min} , the minimum relative improvement, and set $\alpha > \alpha_{\min}$.
- (2) For $k = 1, 2, \dots, k_{\max}$, while $\rho_k > \rho_{\min}$, do
- (a) Choose y so that $(R_k \circ W)y \neq 0$.
 - (b) For $l = 1, 2, \dots, l_{\max}$, while $\alpha > \alpha_{\min}$, do
 - i. Set $s \leftarrow (R_k \circ W)y$, $v \leftarrow W(y \circ y)$.
 Solve $\max \frac{(x^T s)^2}{(x \circ x)^T v}$ s.t. $x \in \mathcal{S}^m$.
 - ii. Set $s \leftarrow (R_k \circ W)^T x$, $v \leftarrow W^T(x \circ x)$.
 Solve $\max \frac{(y^T s)^2}{(y \circ y)^T v}$ s.t. $y \in \mathcal{S}^n$.
 - iii. $\beta \leftarrow \frac{[x^T(R_k \circ W)y]^2}{(x \circ x)^T W(y \circ y)}$.
 - iv. If $l > 1$: $\alpha \leftarrow \frac{\beta - \bar{\beta}}{\bar{\beta}}$.
 - v. $\bar{\beta} \leftarrow \beta$.
 End l -loop.
 - (c) $x_k \leftarrow x$, $y_k \leftarrow y$.
 - (d) $d_k \leftarrow \frac{x_k^T (R_k \circ W) y_k}{(x \circ x)^T W(y \circ y)}$.
 - (e) $A_k \leftarrow A_{k-1} + d_k x_k y_k^T$.
 - (f) $R_{k+1} \leftarrow R_k - d_k x_k y_k^T$.
 - (g) $\rho_{k+1} \leftarrow \rho_k - \beta$.
- End
- k
- loop.

Fig. 1. Computing a weighted SDD approximation.

A_k in Step (2e) is usually not stored explicitly; rather, the individual elements (d_k, x_k, y_k) are stored. Similarly, $R_{k+1} \circ W$ in Step (2f) can be applied in Steps (2a), (2b(i)), (2b(ii)), (2b(iii)), and (2d) without explicitly forming it.

2.2 Convergence of the Weighted SDD Algorithm

We show that the weighted norm of the residual generated by the weighted SDD algorithm is strictly decreasing and, furthermore, the weighted SDD algorithm converges linearly to the original matrix.

LEMMA 1. *The residual matrices generated by the weighted SDD algorithm satisfy*

$$\|R_{k+1}\|_W < \|R_k\|_W \text{ for all } k \text{ such that } R_k \neq 0.$$

PROOF. At the end of the inner iterations, we are guaranteed to have found x_k and y_k such that $x_k^T(R_k \circ W)y_k > 0$. The result follows from (4). \square

As with the SDD algorithm, several different strategies can be used to initialize y in Step (2a) in the weighted SDD algorithm (Figure 1). Here, we only present these schemes briefly. The same convergence results hold, and the proofs are similar to those given for the SDD algorithm in Kolda [1997], as listed below.

- (1) MAX: Choose e_j such that j is the index of the column containing the largest magnitude entry in $R_k \circ R_k \circ W$.
- (2) CYC: Choose e_i where $i = (k \bmod n) + 1$.
- (3) THR: Accept a given unit vector only if it satisfies $\|R_k e_j\|_W^2 \geq \|R_k\|_W^2/n$.

THEOREM 3 [KOLDA 1997]. *The sequence $\{A_k\}$ generated by the SDD algorithm with MAX initialization converges to A in the Frobenius norm. Furthermore, the rate of convergence is at least linear.*

THEOREM 4 [KOLDA 1997]. *The sequence $\{A_k\}$ generated by the SDD algorithm with CYC initialization converges to A in the Frobenius norm. Furthermore, the rate of convergence is at least n -step linear.*

THEOREM 5 [KOLDA 1997]. *The sequence $\{A_k\}$ generated by the SDD algorithm with THR initialization converges to A in the Frobenius norm. Furthermore, the rate of convergence is at least linear.*

Note that although the SDD algorithm can be initialized using a discrete approximation to the SVD vector, this cannot be done in the weighted-SDD case, since there is no simple analog to the SVD.

3. THE TENSOR SDD APPROXIMATION

Let A be an $m_1 \times m_2 \times \dots \times m_n$ tensor over \mathfrak{R} . The order of A is n . The dimension of A is $m \equiv \prod_{j=1}^n m_j$, and m_j is the j th subdimension. An element of A is specified as

$$A_{i_1 i_2 \dots i_n}$$

where $i_j \in \{1, 2, \dots, m_j\}$ for $j = 1, \dots, n$. A matrix is a tensor of order two.

As with matrices, we may be interested in a storage-efficient approximation of a given tensor. We extend the notion of an SDD approximation to a tensor SDD approximation. First we define some notation for tensors.

3.1 Notation

If A and B are two tensors of the same size (i.e., the order n and all subdimensions m_j are equal), then the inner product of A and B is defined as

$$A \cdot B \equiv \sum_{i_1=1}^{m_1} \sum_{i_2=1}^{m_2} \dots \sum_{i_n=1}^{m_n} A_{i_1 i_2 \dots i_n} B_{i_1 i_2 \dots i_n}$$

We define the *norm* of A , $\|A\|$, to be

$$\|A\|^2 \equiv A \cdot A = \sum_{i_1=1}^{m_1} \sum_{i_2=1}^{m_2} \cdots \sum_{i_n=1}^{m_n} A_{i_1 i_2 \cdots i_n}^2.$$

Suppose B is an $m_1 \times \cdots \times m_{j-1} \times m_{j+1} \times \cdots \times m_n$ tensor of order $n - 1$. Then the i_j th ($1 \leq i_j \leq m_j$) element of the *contracted product* of A and B is defined as

$$(A \cdot B)_{i_j} \equiv \sum_{i_1=1}^{m_1} \cdots \sum_{i_{j-1}=1}^{m_{j-1}} \sum_{i_{j+1}=1}^{m_{j+1}} \cdots \sum_{i_n=1}^{m_n} A_{i_1 \cdots i_{j-1} i_{j+1} \cdots i_n} B_{i_1 \cdots i_{j-1} i_{j+1} \cdots i_n}.$$

A *decomposed* tensor is a tensor that can be written as

$$X = x^{(1)} \otimes x^{(2)} \otimes \cdots \otimes x^{(n)},$$

where $x^{(j)} \in \mathfrak{R}^{m_j}$ for $j = 1, \dots, n$. The vectors $x^{(j)}$ are called the *components* of X . In this case,

$$X_{i_1 i_2 \cdots i_n} = x_{i_1}^{(1)} x_{i_2}^{(2)} \cdots x_{i_n}^{(n)}.$$

LEMMA [KOLDA 2000]. *Let A be a tensor of order n and X a decomposed tensor of order n . Then*

$$A \cdot X = (A \cdot X^{(-j)}) \cdot x^{(j)},$$

where the notation $X^{(-j)}$ indicates X with the j th component removed, i.e.,

$$X^{(-j)} \equiv x^{(1)} \otimes \cdots \otimes x^{(j-1)} \otimes x^{(j+1)} \otimes \cdots \otimes x^{(n)}.$$

The notion of rank for tensors of order greater than two is a nontrivial matter (e.g., see Kolda [2000]), but a single decomposed tensor is always a tensor of rank one.

3.2 Definition of the Tensor SDD Approximation

Suppose we wish to approximate an n -dimensional tensor A as follows,

$$A \approx A_k \equiv \sum_{i=1}^k d_i X_i,$$

where $d_i > 0$ and X_i is a decomposed tensor whose components are restricted to $x_i^{(j)} \in \mathcal{S}^{m_j}$, with $\mathcal{S} = \{-1, 0, 1\}$. This is called a k -term *tensor SDD*.

The SDD representation is efficient in terms of storage. If the tensor A is dense, the total storage required for A is

$$\gamma \prod_{j=1}^n m_j,$$

where γ is the amount of storage required for each element of A . For example, if the elements of A are integer values between 0 and 255, then γ is one byte (8 bits). The storage required for the approximation A_k is

$$k \left(\alpha + \beta \sum_{j=1}^n m_j \right),$$

where α is the storage required for each d_k and is usually chosen to be equal to γ and β is the amount of storage required to store each element of \mathcal{S} , i.e., $\log_2 3$ bits. Since $k \ll \prod_{j=1}^n m_j$, the approximation generally requires significantly less storage than the original tensor.

3.3 Computing a Tensor SDD Approximation

As with the regular and weighted SDDs, a tensor SDD approximation can be constructed via a greedy algorithm. Each iteration, a new d and X are computed that are the solution to the following subproblem:

$$\min F_k(d, X) \equiv \|R_k - dX\|^2 \text{ subject to } d > 0, x^{(j)} \in \mathcal{S}^{m_j} \text{ for } j = 1, \dots, n, \quad (8)$$

where $R_k \equiv A - \sum_{i=1}^{k-1} d_i X_i$ denotes the k th residual matrix. This is a mixed integer programming problem, but it can be simplified to an integer program as demonstrated by the following theorem, a generalization of Theorem 1 (in the case where W is the ones matrix).

THEOREM 6. *Solving the mixed integer program (8) is equivalent to solving the integer program*

$$\max \tilde{F}(X) = \frac{(R \cdot X)^2}{\|X\|^2} \text{ subject to } x^{(j)} \in \mathcal{S}^{m_j} \text{ for } j = 1, \dots, n. \quad (9)$$

PROOF. See Kolda and O’Leary [1999a]. □

The integer programming problem (9) has $3^{m_1+m_2+\dots+m_n}$ possible solutions. To solve this problem approximately, an alternating algorithm will be used. The idea is the same as for the regular and weighed SDDs. Fix all the components of X except one, say $x^{(j)}$, and find the optimal $x^{(j)}$ under those conditions. Repeat this process for another value of j , continuing until improvement in the value of $\tilde{F}(X)$ stagnates.

Assume that all components of X are fixed except $x^{(j)}$. Then (9) reduces to

$$\max \frac{(s \cdot x^{(j)})^2}{\|x^{(j)}\|_2^2} \text{ subject to } x^{(j)} \in \mathcal{S}^{m_j},$$

where $s \equiv (R_k \cdot X^{(-j)}) / \|X^{(-j)}\|^2$. This is the same as the problem for the regular SDD, so we know how to solve it.

- (1) Let R_k denote the residual, and initialize $R_1 \leftarrow A$.
 Let $\rho_k = \|R_k\|^2$ be the norm of the residual, and initialize $\rho_1 \leftarrow \|R_1\|^2$.
 Let A_k denote the k -term approximation, and initialize $A_0 \leftarrow 0$.
 Choose k_{\max} , the maximum number of terms in the approximation.
 Choose ρ_{\min} , the desired accuracy of the approximation.
 Choose l_{\max} , the maximum allowable inner iterations.
 Choose α_{\min} , the minimum relative improvement, and set $\alpha > 2\alpha_{\min}$.
- (2) For $k = 1, 2, \dots, k_{\max}$, while $\rho_k > \rho_{\min}$, do
 - (a) Initialize $X = x^{(1)} \otimes x^{(2)} \otimes \dots \otimes x^{(n)}$ so that $R_k \cdot X \neq 0$.
 - (b) For $l = 1, 2, \dots, l_{\max}$, while $\alpha > \alpha_{\min}$, do
 - i. For $j = 1, 2, \dots, n$ do
 Set $s \leftarrow R_k \cdot X^{(-j)}$.
 Solve $\max \frac{(s^T x^{(j)})^2}{\|x^{(j)}\|_2^2}$ s.t. $x^{(j)} \in \mathcal{S}^{m_j}$.
 End j -loop.
 - ii. $\beta \leftarrow \frac{(R_k \cdot X)^2}{\|X\|^2}$.
 - iii. If $l > 1$: $\alpha \leftarrow \frac{\beta - \bar{\beta}}{\bar{\beta}}$.
 - iv. $\bar{\beta} \leftarrow \beta$.
 End l -loop.
 - (c) $X_k \leftarrow X$.
 - (d) $d_k \leftarrow \frac{R_k \cdot X_k}{\|X_k\|^2}$.
 - (e) $A_k \leftarrow A_{k-1} + d_k X_k$.
 - (f) $R_{k+1} \leftarrow R_k - d_k X_k$.
 - (g) $\rho_{k+1} \leftarrow \rho_k - \beta$.
 End k -loop.

Fig. 2. Computing a tensor SDD approximation.

The tensor SDD approximation algorithm is given in Figure 2. In Step (2a), X should be chosen so that $R_k \cdot X \neq 0$. Unless R_k is zero itself (in which case $A_{k-1} = A$), it is always possible to pick such an X . The for-loop in Step (2b(i)) does not need to go through the components of X in order. That loop could be replaced by “For $j = \pi(1), \pi(2), \dots, \pi(n)$ do,” where π is an n -permutation. Note that in each step of (2b(i)), the value of X may change and that the objective function is guaranteed to be at least as good as it was with the previous X .

3.4 Convergence of the Tensor SDD Algorithm

Like the SDD, the tensor SDD algorithm has the property that the norm of the residual decreases each outer iteration. Furthermore, we can prove convergence results similar to those for the SDD (proofs are omitted but are similar to those for the SDD) using each of the following starting strategies in Step (2a) of the tensor SDD algorithm:

- (1) MAX: Initialize $X = e_{j_1}^{(1)} \otimes e_{j_2}^{(2)} \otimes \cdots \otimes e_{j_n}^{(n)}$, where $R_{j_1 j_2, \dots, j_n}$ is the largest magnitude element of R .
- (2) CYC: Same idea as for the SDD, but now the cycle is $\prod_{j=2}^n m_j$ long.
- (3) THR: Choose $X^{(-1)} = e_{j_2}^{(2)} \otimes \cdots \otimes e_{j_n}^{(n)}$ (i.e., x with the first component removed) such that

$$\|(R \cdot X^{(-1)})\|_2^2 \geq \|R\|^2 \left/ \prod_{j=2}^n m_j \right.$$

Although an appropriate choice of $e_{j_2}^{(2)} \otimes \cdots \otimes e_{j_n}^{(n)}$ is guaranteed to exist, it may be difficult to find because of the large space of elements to search through.

4. APPLICATIONS

SDD approximations are useful in applications involving storage compression, data filtering, and feature extraction. As examples, we discuss in this section the use of an SDD approximation in image compression, chromosome classification, and latent semantic indexing of documents.

4.1 Data Compression via SDD and Weighted SDD Approximations

If a matrix consumes too much storage space, then an SDD approximation is one way to reduce the storage burden. For example, an SDD approximation can be used for image compression. The SDD was originally developed by O'Leary and Peleg [1983] for this application. If each pixel value (e.g., gray level) is stored as a matrix entry, then a k -term SDD of the resulting matrix can be stored as an approximation to the original image.

Other matrix approximation techniques have been used for image compression. The SVD [Golub and Van Loan 1989] provides a set of basis vectors that gives the optimal low-rank approximation in the sense of minimizing the sum squared errors (Frobenius norm). But these vectors are expensive to generate and take quite a bit of storage space ($n + m + 1$ floating-point elements per term, although it is possible to use lower precision). At the other extreme, predetermined basis vectors can be used (e.g., Haar basis or other wavelet bases). In this case, the basis vectors do not need to be explicitly stored, but the number of terms is generally much larger than for the SVD. Although the SDD algorithm chooses the basis vectors to fit the particular problem (like the SVD), it chooses them with restricted entries (like the wavelet bases), making the storage per term only $\log_2 3(n + m)$ bits plus one floating-point number.

Experiments using SDD approximations for images achieved 10 to 1 compression (using the SDD with run-length encoding) without visual degradation of the image [O'Leary and Peleg 1983].

If some portions of the image are more important than others (a face, rather than clothing, for example), then large weights (i.e., entries in W)

could be used to produce greater fidelity of the approximation in these important regions.

4.2 Data Filtering via SDD and Weighted SDD Approximations

The k -term approximations produced by the SDD algorithm can be thought of as filtered approximations, finding relations between the columns (or rows) of the matrix that are hidden by local variations. Thus, if we have many observations of the same vector-valued phenomenon, then an SDD approximation of the data can reveal the essential unchanging characteristics.

This fact has application in chromosome classification. Given a “training set” consisting of many observations of a given type of chromosome (e.g., a human X chromosome), an SDD approximation of this data extracts common characteristics, similar to a principal component analysis, but typically requiring less storage space. Then the idealized representation of this chromosome can be used to identify other chromosomes of the same type (*chromosome karyotyping*). Weights can be used to rate the importance of the different observations, or their degree of certainty.

For more information on this technique, see Conroy et al. [2000].

4.3 Feature Extraction via SDD and Weighted SDD Approximations

The low-rank approximations produced by the SDD algorithm extract features that are common among the columns (or rows) of the matrix. This task is addressed by latent semantic indexing (LSI) of documents. A database of documents can be represented by a term-document matrix, in which each matrix entry represents the importance of some term in a particular document. Documents can be clustered for retrieval based on common features. Standard algorithms use a low-rank SVD to extract these feature vectors, but the storage involved is often greater than that for the original matrix. In contrast, SDD approximations have been used by Kolda and O’Leary [1998; 1999b] to achieve similar retrieval performance at a much lower storage cost, and weights could be added to ensure that critical terms are weighted more heavily in the retrieval.

4.4 Image Compression with a Tensor SDD

A $n \times m$ pixel color image is usually broken into its composite colors for storage, i.e., red, green, and blue for an RGB image. This results in a $n \times m \times p$ image where p is the number of colors. The tensor SDD algorithm can be used to compress such an image and should be more efficient than individually compressing the matrix associated with each color.

Further, collections of gray or color images can be compressed using the tensor SDD algorithm since they form an $n \times m \times p \times q$ tensor where q is the number of images. If the images are related, then this should again be more efficient than compressing the images separately.

Table I. Bit Representation of \mathcal{S} -Values

\mathcal{S} -Value	Value Bit	Sign Bit
0	0	undef.
1	1	0
-1	1	1

5. IMPLEMENTATION DETAILS

We provide Matlab software for computing SDD and weighted SDD approximations to matrices and for computing SDD approximations to tensors.

We also provide C software for computing SDD approximations which can be extended to weighted or tensor SDDs. We do not provide the extension here since it will vary by application. The software we provide assumes that one has a sparse matrix stored in compressed sparse column (CSC) format (e.g., see Barrett et al. [1994, p. 65]) and is applicable to problems arising, for example, in latent semantic indexing in information retrieval.

In our discussion of the implementation details, we focus on the C data structures and implementation of the regular SDD algorithm that take advantage of the compressed storage.

5.1 Data Structures

An entry from the discrete set \mathcal{S} , referred to as an \mathcal{S} -value, can be stored using only $\log_2 3$ bits. We actually use two bits of storage per \mathcal{S} -value because it is advantageous in computations involving the \mathcal{S} -values (see Section 5.2) and requires only 26% more memory. The first bit is the *value* bit and is on if the \mathcal{S} -value is nonzero and off otherwise; the second bit is the *sign* bit and is on for an \mathcal{S} -value of -1 , off for 1 , and undefined for 0 (Table I). The undefined bits would not be stored if we were storing using only $\log_2 3$ bits per \mathcal{S} -value.

Each iteration, a new (d, x, y) triplet is computed. The x and y vectors of length m and n , respectively, are referred to as \mathcal{S} -vectors. In SDDPACK, we store each \mathcal{S} -vector's value and sign arrays packed into unsigned long integer arrays.

Suppose that we are working on a p -bit architecture (i.e., the length of a single word of memory is p bits). Then the memory allocated to the value array to hold m bits is $\lceil m/p \rceil$ words. Storage for the sign array is the same. An example of an \mathcal{S} -vector and its representation on an 8-bit architecture is given in Figure 3. Notice that extra bits in the last word of the array and sign bits associated with zero-values are undefined. Extra bits are ignored (i.e., masked to an appropriate value) in any calculations. We used an 8-bit example for simplicity; current architectures are generally 32- or 64-bit (Table II).

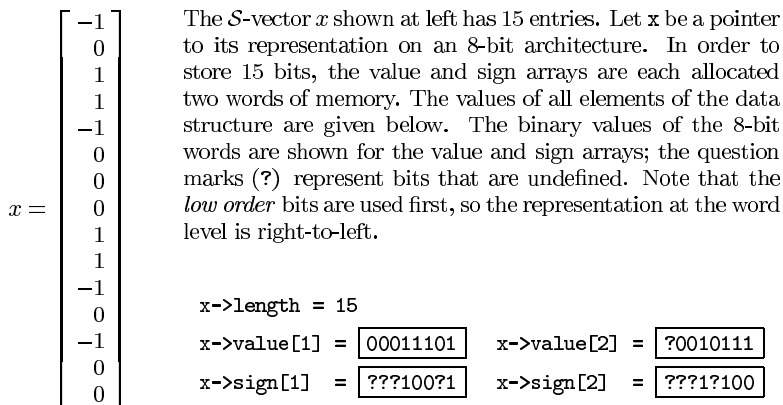


Fig. 3. Illustration of *svector* data structure.

Table II. Current Architectures

32-Bit	64-Bit
Sun Sparc Intel Pentium IBM RS6000	SGI Octane Dec Alpha

```

i = 10
index = i >> 3
mask = 1 << (i AND 7)
x->value[index] AND mask = 00000010
x->sign[index]   AND mask = 00000000
                    
```

Fig. 4. Looking up a value in a packed array.

5.2 Computations with Objects Using Packed Storage

Given an \mathcal{S} -vector in packed storage, we can look up the value in a particular entry as follows. If i is the desired entry, then the index into the packed array is $i \text{ div } p$, and the bit we want inside that word is $i \text{ mod } p$, and the desired bit can be masked off. We first do a mask on the appropriate word in the value array. If the result is zero, then entry i is zero, and we need do no further work. Otherwise, the entry is either +1 or -1, and we need to determine the sign. We mask off the appropriate word in the sign array. If that is zero, the entry is +1; otherwise, it is -1.

For example, Figure 4 shows how to look up entry 10 in the example in Figure 3. Here i is the desired entry. To compute `index`, the index into the packed array, divide by p , the number of bits per word. Since p is always a power of two, this can be accomplished by a right shift. In this example, we right shift 3 since $\log_2 8 = 3$. Given the correct index into the packed arrays, the correct bit inside the word is determined by a mod by p . Again, since p is always a power of two, we can use a shortcut by doing a logical AND with $p - 1$, in this example, 7. Then mask the appropriate word in

```

common = a->value AND b->value
oppsign = (a-> sign XOR b->sign) AND common
ip = popcount(common) - 2 popcount(oppsign)

```

Fig. 5. Inner product of two \mathcal{S} -vectors.

the value array. In this example, it is nonzero, so the entry is either $+1$ or -1 . Then mask the appropriate word in the sign array and determine that the entry is $+1$.

Note that the alignment of the value and sign arrays makes it easy to do individual lookups of values. If we did not store the “filler” bits in the sign array for the zero entries, the sign array would be much shorter, but we would have a difficult time knowing where in the sign array to look for the appropriate bit.

In the previous example, we saw how to look up a random entry in a packed array. Often we walk through an \mathcal{S} -vector in sequence. In that case, computations can be performed even more quickly by copying the current value and sign words into the register to be used p times and quickly updating the mask with just a single left shift. Every p entries, we reset the mask to one and swap the next value and sign words into the register.

The inner product between two \mathcal{S} -vectors, something that we require, can be computed as follows. The result is the number of nonzeros in common minus twice the number of common nonzeros with opposite signs. Pseudocode is given in Figure 5 for the inner product of two \mathcal{S} -vectors a and b . In practice, the logical ANDs and ORs are done on a word-by-word basis and the popcount (sum) is determined using a lookup table on a byte-by-byte basis. So, for computing the inner product of two m -long \mathcal{S} -vectors, the work required is $3\lceil m/p \rceil + 4\lceil m/8 \rceil$ and requires no multiplication.

Each iteration of the SDD calculation, the most expensive operations are the computations of $R_k y$ or $R_k^T x$ (Steps (2b(i)) and (2b(ii)) of the SDD Algorithm of Figure 1). We focus on the computation of $R_k y$ and explain the differences for the transpose at the conclusion. The residual breaks into two parts: the original matrix, A , and a $(k - 1)$ -term SDD approximation that we denote by XDY^T .

The computation $v = Ay$ is a sparse matrix times an \mathcal{S} -vector. The sparse matrix is stored in CSC format. We loop through the matrix columnwise, which means that we walk through the y -vector in sequence. If y_j is zero, then nothing is done with column j . Otherwise, we either add ($y_j = 1$) or subtract ($y_j = -1$) the entries in column j from the appropriate entries in the solution vector v .

The computation of $w = XDY^T y$ breaks down into three parts: $Y^T y$, $D(Y^T y)$, and $X(DY^T y)$. The first part is an \mathcal{S} -matrix times an \mathcal{S} -vector, which reduces to an inner product between two \mathcal{S} -vectors for each entry in the solution. The result of $Y^T y$ is an integer vector. The $D(Y^T y)$ is just a simple scaling operation, and the result is a real vector. The final product is $X(DY^T y)$, and in this case we walk through each bit in the X matrix

Table III. Test Matrices

Matrix	Rows	Cols	NNZ	Rank	Density(%)
bfw62a	62	62	450	62	11.7
impcol_c	137	137	411	137	2.2
west0132	132	132	414	132	2.4
watson2	66	67	409	66	9.2

column by column and take appropriate action. Again, only additions and subtractions are required, no multiplications.

In the case of the transpose computation, the main difference is in the computation of $A^T x$. Here, we are forced to use random access into x since A is stored in CSC format. The method for computing $(XDY^T)^T x$ is nearly identical to that described previously for $XDY^T y$, except that the roles of X and Y are swapped.

So, the only multiplications required in our computations are the diagonal scalings; everything else is additions and subtractions. Further, the pieces of an SDD approximation are small and fit well into cache.

6. NUMERICAL RESULTS

The computational experiments presented here are done in Matlab, with the Matlab code and examples provided in SDDPACK. In general, the C SDDPACK code should be used when speed and storage efficiency are concerns. No results are presented here for the weighted or tensor SDDs algorithms although MATLAB code for these decompositions are included.

We discuss previous research and present new results on the SDD algorithm and starting criteria as well as comparisons between SDD and SVD approximations.

Kolda [1997] presents comparisons of the various starting criteria on small, dense matrices. To summarize, the MAX, CYC, and THR techniques are nearly identical in performance. The SVD initialization typically results in fewer inner iterations per outer iteration, but the gain is offset by the expense of computing the starting vector.

Kolda and O’Leary [1998] compared SDD and SVD approximations for latent semantic indexing for information retrieval. At equal levels of retrieval performance, the SDD model required approximately 20 times less storage and performed queries about twice as fast. On the negative side, the SVD approximation can be computed about four times faster than the SDD approximation for equal performance levels. The SDD computations used option PER, as described subsequently—we may be able to improve the speed and performance by using option THR instead.

We compare various initialization strategies for the SDD algorithm on several sparse matrices from MatrixMarket; the test set is described in Table III. We test four different initialization strategies as listed below.

THR: Cycle through the unit vectors (starting where it left off at the previous iteration) until $\|R_k e_j\|_2^2 \geq \|R_k\|_F^2/n$, and set $y = e_j$. (Threshold)

Table IV. Comparison of Initialization Techniques for `bfw62a` and `impcol_c`

Init.	bfw62a			impcol_c		
	% Resid.	In. Its.	% Density	% Resid.	In. Its.	% Density
THR	28.19	3.69	9.33	3.53	2.58	1.79
CYC	25.54	3.73	9.55	7.86	3.47	6.47
ONE	22.86	6.81	41.13	36.93	5.95	24.32
PER	25.48	6.79	21.48	31.09	6.39	21.24

Table V. Comparison of Initialization Techniques for `west0132` and `watson2`

Init.	west0132			watson2		
	% Resid.	In. Its.	% Density	% Resid.	In. Its.	% Density
THR	0.00	5.62	1.95	16.99	3.02	3.87
CYC	0.01	3.25	1.68	20.51	2.76	4.17
ONE	0.01	5.64	11.97	78.74	5.42	18.94
PER	0.30	8.46	3.54	75.99	4.82	10.69

CYC: Initialize $y = e_i$, where $i = ((k - 1) \bmod n) + 1$. (Cycling)

ONE: Initialize y to the all ones vector. (Ones)

PER: Initialize y to a vector such that elements 1,101,201, . . . are one and the remaining elements are zero. (Periodic ones)

We do not test the MAX strategy because these matrices are sparse, so the residual is stored implicitly. The “ONE” and “PER” strategies are tested to compare to past initialization strategies used by O’Leary and Peleg [1983] and Kolda and O’Leary [1998]. The other parameters of the SDD algorithm are set as follows. We set k_{\max} to the rank of the matrix for comparison to the SVD which will have a zero residual at the point. We set $\rho_{\min} = 0$ which forces the SDD algorithm to continue until it has a zero residual or, the more likely scenario, $k = k_{\max}$. We set $\alpha_{\min} = 0.01$ which was chosen empirically to balance work and approximation quality. We set $l_{\max} = 100$ which is a large enough bound that it is never attained.

The performance of these four strategies on our four test matrices is shown in Tables IV and V. The tables compare the relative reduction in the residual (as a percentage), the average number of inner iterations (which includes the extra work for initialization in THR), and the density of the final factors (as a percentage). The initialization can have a dramatic affect on the residual after k terms. In the `impcol_c` and `watson2` matrices, THR and CYC are drastically better than ONE and PER. The number of inner iterations is lowest overall for CYC, with THR being a close second. In terms of density, THR and CYC are drastically better in every case, perhaps because the initial vector is sparse. It seems that the density of the factors may be somewhat related to the density of the original matrix. Overall, THR is best, with CYC a close second.

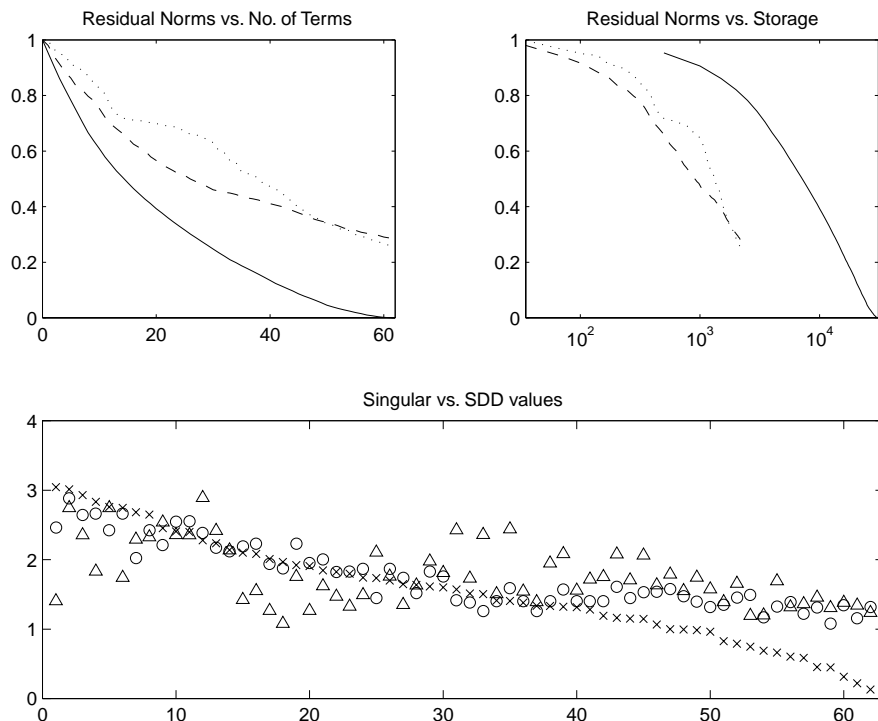


Fig. 6. Comparison of the SVD (solid line, x marks), SDD-THR (dashed line, o marks), and SDD-CYC (dotted line, triangle marks) on *bfw62a*.

In Figures 6–9, the SVD, SDD-THR, and SDD-CYC are compared. The results on *bfw62a* are given in Figure 6. The upper left plot shows a comparison of the relative residual ($\|R_k\|/\|R_0\|$) versus the number of terms. The SVD is the optimal decomposition for a fixed number of terms, so the SDD curves will lie above it. However, the SDD algorithm still gives good reduction in the residual, requiring only about twice as many terms as the SVD for the same level of reduction. SDD-THR gives a better residual than SDD-CYC until the last few terms, where SDD-CYC “catches up.” In the upper right, a plot of the residual versus the storage is shown; for the same level of reduction in the residual, the storage requirement for the SDD approximation is about one to two orders of magnitude less than for the SVD. In the bottom plot, the singular values and SDD values are shown, where the i th SDD value is defined as $\hat{d}_i = d_i\|x_i\|_2\|y_i\|_2$. Initially, the SDD values are smaller than the singular values because they cannot capture as much information; later, they are larger because they are capturing the information missed initially.

The `impcol_c` matrix has an interesting singular value pattern (see Figure 7): there is one isolated singular value at 11, a cluster of singular values at 3, and another cluster at 2. SDD-THR mimics the SVD closely because SDD-THR also finds one isolated singular SDD value, as many SDD values at 3, and almost as many SDD values at 2. SDD-CYC, on the

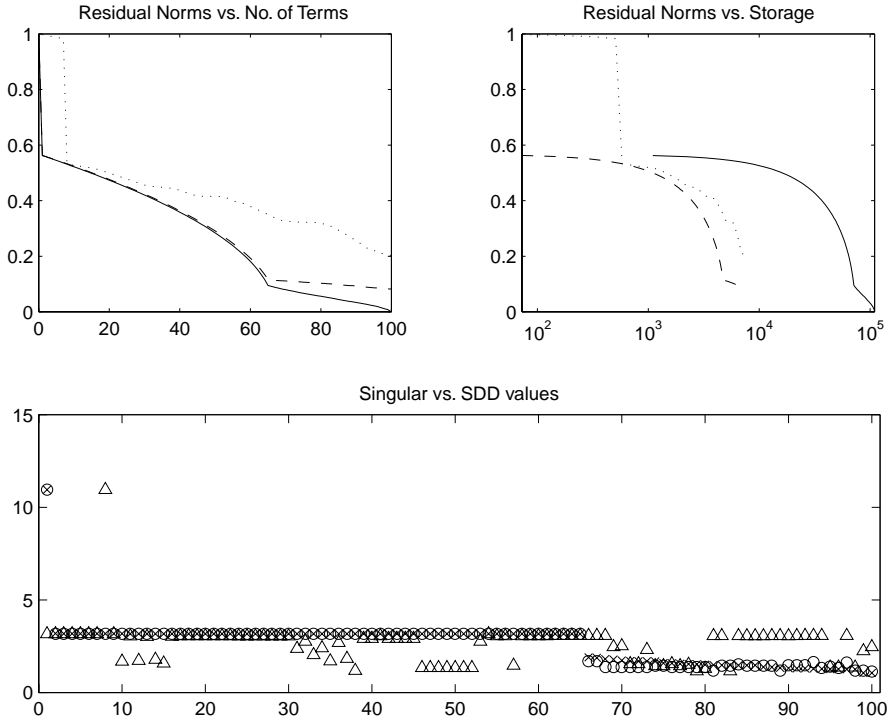


Fig. 7. Comparison of the SVD (solid line, x marks), SDD-THR (dashed line, o marks), and SDD-CYC (dotted line, triangle marks) on `impcol_c`.

other hand, has trouble mimicking singular values because it does not pick out the isolated value at first. Still, both SDD variants are superior to the SVD in terms of storage versus residual norm.

On `west0132` (see Figure 8), we see phenomena similar to that for `impcol_c`. SDD-THR finds isolated SDD values and quickly reduces the residual—almost as quickly as the SVD itself in terms of number of terms. SDD-CYC has more trouble isolating SDD values but eventually gets them as well. Here, SDD-THR is superior to the SVD in terms of storage, but SDD-CYC is not.

The last matrix, `watson2` (see Figure 9), most closely resembles `bfw62a` in the structure of its singular values, although `watson2` has three eigenvalues that are slightly isolated, and we can see that both SDD methods eventually pick out such values which results in the steeper drops in the residual curves. Again, SDD-THR does better than the SDD-CYC in all respects. SDD-THR requires about twice as many terms to get the same reduction in storage as the SVD, while using an order of magnitude less storage.

7. CONCLUSIONS

By presenting the code for computing SDD approximations, we hope to stimulate more uses of this storage-efficient matrix approximation method.

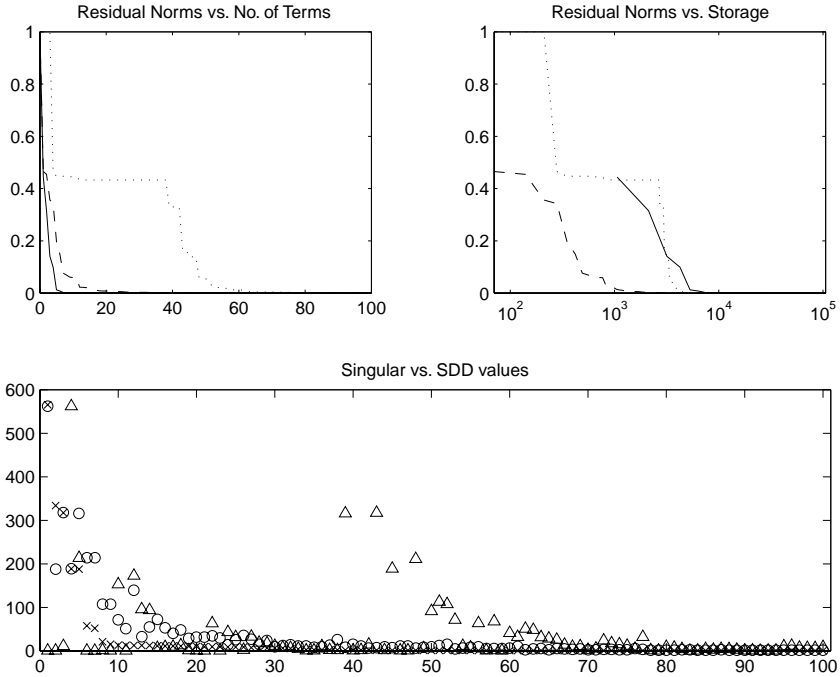


Fig. 8. Comparison of the SVD (solid line, x marks), SDD-THR (dashed line, o marks), and SDD-CYC (dotted line, triangle marks) on west0132.

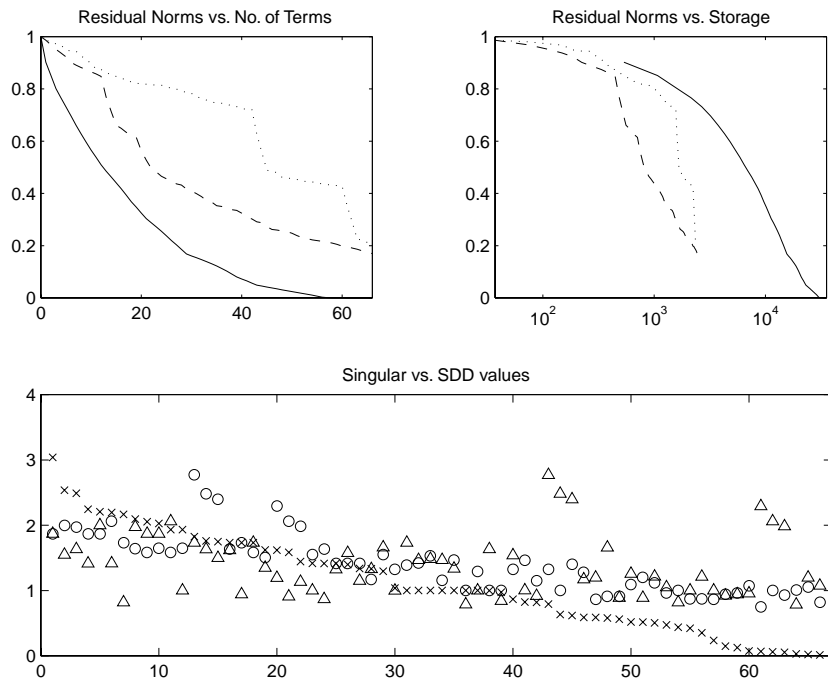


Fig. 9. Comparison of the SVD (solid line, x marks), SDD-THR (dashed line, o marks), and SDD-CYC (dotted line, triangle marks) on watson2.

ACKNOWLEDGMENTS

We are grateful to Professor Walter Gander and Professor Martin Gutnecht for their hospitality at ETH.

REFERENCES

- BARRETT, R., BERRY, M., CHAN, T., DEMMEL, J., DONATO, J., DONGARRA, J., ELJKHOUT, V., POZO, R., ROMINE, C., AND VAN DER VORST, H. 1994. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*. SIAM, Philadelphia, PA.
- CONROY, J., KOLDA, T. G., O'LEARY, D. P., AND O'LEARY, T. 2000. Chromosome identification. *Lab. Invest.*. To be published.
- GABRIEL, K. R. AND ZAMIR, S. 1979. Lower rank approximation of matrices by least squares with any choice of weights. *Technometrics* 21, 489–498.
- GOLUB, G. AND VAN LOAN, C. F. 1989. *Matrix Computations*. 2nd ed. Johns Hopkins University Press, Baltimore, MD.
- KOLDA, T. G. 1997. Limited-memory matrix methods with applications. Ph.D. Dissertation. University of Maryland at College Park, College Park, MD.
- KOLDA, T. G. 2000. Orthogonal rank decompositions for tensors. Tech. Rep. SAND2000-8566. Sandia National Laboratories, Livermore, CA.
- KOLDA, T. G. AND O'LEARY, D. P. 1998. A semidiscrete matrix decomposition for latent semantic indexing in information retrieval. *ACM Trans. Inf. Syst.* 16, 4, 322–346.
- KOLDA, T. G. AND O'LEARY, D. P. 1999a. Computation and uses of the semidiscrete matrix decomposition. Tech. Rep. CS-TR-4012 and UMIACS-TR-99-22. Department of Computer Science, University of Maryland, College Park, MD.
- KOLDA, T. G. AND O'LEARY, D. P. 1999b. Latent semantic indexing via a semi-discrete matrix decomposition. In *The Mathematics of Information Coding, Extraction and Distribution*, G. Cybenko, D. P. O'Leary, and J. Rissanen, Eds. IMA Volumes in Mathematics and Its Applications, vol. 107. Springer-Verlag, Vienna, Austria, 73–80.
- O'LEARY, D. P. AND PELEG, S. 1983. Digital image compression by outer product expansion. *IEEE Trans. Commun.* 31, 441–444.

Received: April 1999; revised: December 1999; accepted: May 2000