Available online at www.sciencedirect.com
SCIENCE DIRECT®

LINEAR ALGEBRA AND ITS APPLICATIONS

Linear Algebra and its Applications 415 (2006) 20–30

www.elsevier.com/locate/laa

Fast low-rank modifications of the thin singular value decomposition

Matthew Brand

MERL, 201 Broadway, Cambridge, MA 02139, USA

Received 28 May 2003; accepted 27 July 2005
Available online 27 September 2005
Submitted by Kugazov

Abstract

This paper develops an identity for additive modifications of a singular value decomposition (SVD) to reflect updates, downdates, shifts, and edits of the data matrix. This sets the stage for fast and memory-efficient sequential algorithms for tracking singular values and subspaces. In conjunction with a fast solution for the pseudo-inverse of a submatrix of an orthogonal matrix, we develop a scheme for computing a thin SVD of streaming data in a single pass with linear time complexity: A rank-r thin SVD of a $p \times q$ matrix can be computed in $O(pqr)$ time for $r \leqslant \sqrt{\min(p, q)}$.

AMS classification: 49M27; 15A18; 15A23; 65F20

Keywords: Singular value decomposition; Sequential updating; Subspace tracking

1. The singular value decomposition

The singular value decomposition (SVD) diagonalizes a real matrix $\mathbf{X} \in \mathbb{R}^{p\times q}$ via left and right rotations by orthonormal matrices $\mathbf{U} \in \mathbb{R}^{p\times p}$ and $\mathbf{V} \in \mathbb{R}^{q\times q}$, e.g.,

E-mail address: brand@merl.com
URL: http://www.merl.com/people/brand/

0024-3795/$ - see front matter © 2005 Elsevier Inc. All rights reserved.
doi:10.1016/j.laa.2005.07.021

$\mathbf{U}^\top\mathbf{X}\mathbf{V} = \mathbf{S}$ is diagonal and nonnegative. Equivalently, it decomposes $\mathbf{X}$ into a sum of rank-1 matrices generated by singular value *triplets*: $\mathbf{X} = \mathbf{U}\,\mathrm{diag}(\mathbf{s})\mathbf{V}^\top = \sum_i \mathbf{u}_i s_i \mathbf{v}_i^\top$ for singular values s_i on the diagonal of $\mathbf{S}$ and singular vectors $\mathbf{u}_i$ and $\mathbf{v}_i$ drawn from the columns of $\mathbf{U}$ and $\mathbf{V}$.

The rank-r *thin* SVD restricts this sum to the r triplets having the largest-magnitude singular values. We will write this $\mathbf{U}\,\mathrm{diag}(\mathbf{s})\mathbf{V}^\top \xleftarrow{r} \mathbf{X}$ with orthonormal subspace matrices $\mathbf{U} \in \mathbb{R}^{p\times r}$, $\mathbf{V} \in \mathbb{R}^{q\times r}$ and singular value vector $\mathbf{s} \in \mathbb{R}^r \geqslant 0$. In signal processing, this reduction of $\mathbf{X}$ to a product of thin matrices is interpreted as a form of lossy compression, with the subspace matrices acting as encoding and decoding operators. By the Schmidt (later Eckart–Young–Mirsky) theorem, the thin SVD is the optimal rank-r approximation of $\mathbf{X}$ under any unitarily invariant norm, including the Frobenius norm [1]. This licenses the additional interpretation of the thin SVD as a form of noise suppression, where $\mathbf{X}$ is presumed to be a low-rank data matrix containing measurements contaminated with additive Gaussian noise.

Computing a full SVD is fundamentally an $\mathrm{O}(pq \cdot \min(p, q))$-time problem, making decompositions of extremely large matrices infeasible. Shortly after the introduction of a practical algorithm for computing the SVD on digital computers in the 1960s [2], research turned to problems of faster methods for computing approximations such as the thin SVD, as well as updating an SVD to incorporate new data (e.g., [3,4]). In recent years the practical need for such methods has become acute and the literature has grown accordingly. Section 5 reviews the recent literature in light of the results presented below:

(1) A general identity for additive modifications of an SVD (Section 2).
(2) Specializations of this identity to give SVD updates, downdates, and rank-1 modifications with reduced computational complexity (Section 3).
(3) An expanded thin SVD and sequential updating scheme that offers a strictly linear-time thin SVD in a single pass through a data matrix (Section 4).

The last result has practical value in online settings where data must be incorporated into the SVD as it arrives, typically because the data is too large to be stored or even buffered. For example, many computer vision algorithms call for a "running" thin SVD of a video stream—effectively a data matrix with $\approx 10^5$ rows and an inexhaustible supply of columns. Financial transaction streams and network activity streams are even more demanding.

2. Additive modifications

Let real matrix $\mathbf{X} \in \mathbb{R}^{p\times q}$ have rank r and economy SVD $\mathbf{U}\mathbf{S}\mathbf{V}^\top = \mathbf{X}$ with $\mathbf{S} \in \mathbb{R}^{r\times r}$. Let $\mathbf{A} \in \mathbb{R}^{p\times c}$, $\mathbf{B} \in \mathbb{R}^{q\times c}$ be arbitrary matrices of rank c. We are interested in the SVD of the sum