Signal Processing 89 (2009) 857–875

Contents lists available at ScienceDirect

Signal Processing

journal homepage: www.elsevier.com/locate/sigpro

The Householder Compressor Theorem and its application in subspace tracking

Peter Strobach

AST-Consulting Inc., Bahnsteig 6, 94133 Röhrnbach, Germany

ARTICLE INFO

Article history:
Received 28 April 2008
Received in revised form
28 October 2008
Accepted 2 November 2008
Available online 7 November 2008

Keywords:
Subspace tracking
Eigenvectors
Cholesky
Householder

ABSTRACT

The energy in a symmetric $(r+1)\times(r+1)$ matrix can be transformed perfectly into a smaller $r\times r$ submatrix by means of a two-sided Householder transformation. The parameters of this Householder transformation are uniquely determined by the minor eigenvector of the larger matrix. This compressor is the key to a new type of square-root Householder subspace tracker which is optimal in terms of both complexity and performance. Computer experiments validate the theoretical findings.

1. Introduction

Updating algorithms for subspace and eigensystem tracking have been a subject of constant interest for decades. A review of early methods can be found in [1]. Most of these algorithms were obtained on the basis of sometimes heuristical considerations using concepts known from batch eigendecompositions or gradient/descent methods. This unsatisfactory situation also holds for many of the more recently presented subspace trackers (a list of references is intentionally omitted).

In this paper, we start up from scratch with a very basic consideration of the theoretical requirements for subspace tracking. It turns out that we need essentially a lossless compressor with a capability of mapping the energy of a square symmetric $(r+1)\times(r+1)$ matrix into a smaller submatrix of dimension $r\times r$. We show that the really perfect tool to achieve this is a properly designed two-sided Householder transformation. The parameters of the Householder matrices used in this transformation are exactly determined by the minor eigenvector of the given larger $(r+1)\times(r+1)$ matrix. The remaining "residual" in this compression is just the minor (smallest) eigenvalue of the larger matrix.

This result is manifested in the Householder Compressor Theorem, a formidable tool for developing high-performance subspace trackers. The theorem essentially says that an $N\times r$, $N>r$ subspace tracking problem can be reduced entirely to the "problem" of tracking the minor eigenvector of an $(r+1)\times(r+1)$ power matrix. This is easily and almost perfectly accomplished using a single inverse iteration per time step, because the magnitude ratio between the second-smallest and the smallest eigenvalue of the larger matrix will be very large in a problem of this kind. The rest of the tracker comprises just a set of suitable transformations.

This paper is organized as follows: in Section 2, we introduce the Householder Compressor Theorem. In Section 3, a square-root subspace tracker on this basis is developed. A complete quasicode table is provided. This tracker reaches the $3Nr$ lower bound in dominant complexity for an algorithm of this kind, and produces highest quality results. This is demonstrated with some computer experiments shown in Section 4. In Section 4,

E-mail address: peter_strobach@gmx.de
URL: http://www.ast-consulting.net

0165-1684/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.sigpro.2008.11.001

we also discuss some pressing problems of exception processing, such as idle mode and rank drop, which inevitably occur in the practical application of every tracker. In our new tracker, we can perfectly handle such exceptions, because we are working on triangular power matrix square-roots and employ no inverse matrix updating. Section 5 concludes this paper.

2. The Householder Compressor Theorem

In the course of this paper, we shall see that the central problem in subspace tracking is a partial compression problem of the following kind: Given a symmetric augmented $(r+1)\times(r+1)$ matrix $\bar{\mathbf{S}}$ with

$$\bar{\mathbf{S}} = \begin{bmatrix} \mathbf{S} & \mathbf{h} \\ \mathbf{h}^{\mathrm{T}} & X \end{bmatrix}, \tag{1}$$

we shall find an orthonormal compressor $\mathbf{C}$ with $\mathbf{C}\mathbf{C}^{\mathrm{T}} = \mathbf{C}^{\mathrm{T}}\mathbf{C} = \mathbf{I}$ so that the following partial compression is achieved:

$$\mathbf{C}^{\mathrm{T}}\bar{\mathbf{S}}\mathbf{C} = \begin{bmatrix} & & & 0 \\ & \mathbf{S}_c & & \vdots \\ & & & 0 \\ 0 & \cdots & 0 & \lambda_{min} \end{bmatrix}, \tag{2}$$

where λ_{min} is the minor eigenvalue of $\bar{\mathbf{S}}$. It is not difficult to see that this constitutes the maximum attainable partial compression, because the energy in the appended h-vectors is perfectly mapped into the $r\times r$ upper-left submatrix $\mathbf{S}_c$. The inevitably remaining residual is only the minor eigenvalue in the lower-right corner of the larger matrix after compression.

Suppose we have given the eigenvalue decomposition (EVD) of $\bar{\mathbf{S}}$ in the following form:

$$\bar{\mathbf{S}} = [\mathbf{U}\ \ \mathbf{u}] \begin{bmatrix} & & & 0 \\ & \mathbf{\Lambda} & & \vdots \\ & & & 0 \\ 0 & \cdots & 0 & \lambda_{min} \end{bmatrix} \begin{bmatrix} \mathbf{U}^{\mathrm{T}} \\ \mathbf{u}^{\mathrm{T}} \end{bmatrix}, \tag{3}$$

where $\mathbf{\Lambda}$ is the diagonal matrix of ordered dominant eigenvalues, $\mathbf{U}$ is the corresponding matrix of dominant eigenvectors, and $\mathbf{u}$ is the minor eigenvector of $\bar{\mathbf{S}}$.

It is not difficult to demonstrate that based on this information, the optimal partial compressor is given as follows:

$$\mathbf{C} = [\mathbf{U}\mathbf{\Theta}\ \ \mathbf{u}], \tag{4}$$

where $\mathbf{\Theta}$ denotes any arbitrary orthonormal rotor of dimension $r\times r$. For this $\mathbf{C}$, we simply obtain the result

$$\mathbf{C}^{\mathrm{T}}\bar{\mathbf{S}}\mathbf{C} = \begin{bmatrix} & & & 0 \\ & \mathbf{\Theta}^{\mathrm{T}}\mathbf{\Lambda}\mathbf{\Theta} & & \vdots \\ & & & 0 \\ 0 & \cdots & 0 & \lambda_{min} \end{bmatrix}. \tag{5}$$

Interestingly, $\mathbf{S}_c = \mathbf{\Theta}^{\mathrm{T}}\mathbf{\Lambda}\mathbf{\Theta}$ is in general a square symmetric matrix that turns into the diagonal matrix of dominant eigenvalues only in the special case of $\mathbf{\Theta} = \mathbf{I}$. At this point, it is important to realize that in subspace tracking, there will be no need for a diagonal $\mathbf{S}_c$. A square $\mathbf{S}_c$ will do the job as long as subspace tracking and not directly eigenvector tracking is the issue.

Hence in our case of subspace tracking, we will have the interesting freedom of choosing $\mathbf{\Theta}$ completely arbitrarily. Only the minor eigenvector $\mathbf{u}$ must be known. This gives rise to the following *Householder Compressor Theorem*:

Given a Householder matrix

$$\mathbf{H} = \mathbf{I} - 2\bar{\mathbf{v}}\bar{\mathbf{v}}^{\mathrm{T}}, \tag{6}$$

with

$$\bar{\mathbf{v}} = \begin{bmatrix} \mathbf{v} \\ \varphi \end{bmatrix}, \quad \mathbf{v}^{\mathrm{T}}\mathbf{v} + \varphi^2 = 1, \tag{7}$$

and the minor eigenvector $\mathbf{u}$ of $\bar{\mathbf{S}}$ partitioned as

$$\mathbf{u} = \begin{bmatrix} \mathbf{w} \\ \rho \end{bmatrix}, \tag{8}$$

we choose

$$\varphi = \sqrt{\frac{1-\rho}{2}}, \tag{9a}$$

$$\mathbf{v} = -\frac{1}{2\varphi}\mathbf{w}. \tag{9b}$$

With this special choice, the Householder reflector $\mathbf{H}$ will act as an optimal partial compressor on $\bar{\mathbf{S}}$ with the result:

$$\mathbf{H}\bar{\mathbf{S}}\mathbf{H} = \begin{bmatrix} & & & 0 \\ & \mathbf{S}_c & & \vdots \\ & & & 0 \\ 0 & \cdots & 0 & \lambda_{min} \end{bmatrix}, \tag{10}$$

where

$$\mathbf{S}_c = (\mathbf{I} - 2\mathbf{v}\mathbf{v}^{\mathrm{T}})\mathbf{S}(\mathbf{I} - 2\mathbf{v}\mathbf{v}^{\mathrm{T}}) - 2\varphi(\mathbf{I} - 2\mathbf{v}\mathbf{v}^{\mathrm{T}})\mathbf{h}\mathbf{v}^{\mathrm{T}} - 2\varphi\mathbf{v}\mathbf{h}^{\mathrm{T}}(\mathbf{I} - 2\mathbf{v}\mathbf{v}^{\mathrm{T}}) + 4\varphi^2 X\mathbf{v}\mathbf{v}^{\mathrm{T}}. \tag{11}$$

Proof. Observe that from (6) and (7), we can write

$$\mathbf{H} = \begin{bmatrix} \mathbf{I} - 2\mathbf{v}\mathbf{v}^{\mathrm{T}} & -2\varphi\mathbf{v} \\ -2\varphi\mathbf{v}^{\mathrm{T}} & 1 - 2\varphi^2 \end{bmatrix}. \tag{12}$$

If we choose the rightmost column vector of this Householder matrix equal to the minor eigenvector of $\bar{\mathbf{S}}$, namely

$$\begin{bmatrix} -2\varphi\mathbf{v} \\ 1 - 2\varphi^2 \end{bmatrix} = \mathbf{u} = \begin{bmatrix} \mathbf{w} \\ \rho \end{bmatrix}, \tag{13}$$

we can easily verify that this equation holds for the particular choice of $\mathbf{v}$ and φ as given in (9a) and (9b).

Consequently, we also find that

$$\begin{bmatrix} \mathbf{I} - 2\mathbf{v}\mathbf{v}^{\mathrm{T}} \\ -2\varphi\mathbf{v}^{\mathrm{T}} \end{bmatrix} = [\mathbf{U}\mathbf{\Theta}]. \tag{14}$$

The leading r column (or row) vectors of $\mathbf{H}$ can be posed as a rotated version of the leading eigenvectors of $\bar{\mathbf{S}}$ because $\mathbf{H}^{\mathrm{T}}\mathbf{H} = \mathbf{H}\mathbf{H} = \mathbf{I}$. Consequently, the compression in (10) will be the same as the compression in (5). □