



UNIVERSITY
OF WOLLONGONG
AUSTRALIA

University of Wollongong
Research Online

Faculty of Informatics - Papers (Archive)

Faculty of Engineering and Information Sciences

2001

Signal analysis using a multiresolution form of the singular value decomposition

Ramakrishna Kakarala

Motorola Australian Research Centre

Philip Ogunbona

University of Wollongong, philipo@uow.edu.au

Publication Details

Kakarala, R. & Ogunbona, P. (2001). Signal analysis using a multiresolution form of the singular value decomposition. *IEEE Transactions on Image Processing*, 10 (5), 724-735.

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library:
research-pubs@uow.edu.au

Signal analysis using a multiresolution form of the singular value decomposition

Abstract

This paper proposes a multiresolution form of the singular value decomposition (SVD) and shows how it may be used for signal analysis and approximation. It is well-known that the SVD has optimal decorrelation and subrank approximation properties. The multiresolution form of SVD proposed here retains those properties, and moreover, has linear computational complexity. By using the multiresolution SVD, the following important characteristics of a signal may be measured, at each of several levels of resolution: isotropy, sphericity of principal components, self-similarity under scaling, and resolution of mean-squared error into meaningful components. Theoretical calculations are provided for simple statistical models to show what might be expected. Results are provided with real images to show the usefulness of the SVD decomposition.

Keywords

value, form, signal, singular, decomposition, analysis, multiresolution

Disciplines

Physical Sciences and Mathematics

Publication Details

Kakarala, R. & Ogunbona, P. (2001). Signal analysis using a multiresolution form of the singular value decomposition. *IEEE Transactions on Image Processing*, 10 (5), 724-735.

Signal Analysis Using a Multiresolution Form of the Singular Value Decomposition

Ramakrishna Kakarala and Philip O. Ogunbona, *Senior Member, IEEE*

Abstract—This paper proposes a multiresolution form of the singular value decomposition (SVD) and shows how it may be used for signal analysis and approximation. It is well-known that the SVD has optimal decorrelation and subrank approximation properties. The multiresolution form of SVD proposed here retains those properties, and moreover, has linear computational complexity. By using the multiresolution SVD, the following important characteristics of a signal may be measured, at each of several levels of resolution: isotropy, sphericity of principal components, self-similarity under scaling, and resolution of mean-squared error into meaningful components. Theoretical calculations are provided for simple statistical models to show what might be expected. Results are provided with real images to show the usefulness of the SVD decomposition.

Index Terms—Karhunen–Loève transform, multivariate statistics, principal components analysis, singular value decomposition.

I. INTRODUCTION

THE singular value decomposition (SVD) of a real-valued $M \times N$ matrix X , with $M \leq N$, may be written

$$X = USV^t. \quad (1)$$

Here, U is an orthogonal $M \times M$ matrix whose columns (called the “left singular vectors”) are the eigenvectors of XX^t , V is a $N \times N$ matrix whose columns (the “right singular vectors”) are eigenvectors of X^tX , and S is the $M \times M$ diagonal matrix whose entries (the “singular values”) are the square roots of the corresponding eigenvalues of XX^t . Letting $\hat{X} = U^tX = SV^t$, the SVD may also be written $X = U\hat{X}$. As discussed in more detail in the following, this second form reveals a useful connection with recent research in signal-adapted filterbanks: U is essentially the decorrelating matrix obtained from an input signal’s measured second-order statistics, and \hat{X} contains the subband decomposition of the signal. This paper shows how, by recursively resampling and decomposing the largest rank-one matrix, the SVD may be developed into a multiresolution signal decomposition. Furthermore, it describes how the multiresolution SVD thereby obtained provides useful information for analyzing, and comparing, signals.

Manuscript received June 8, 2000; revised February 1, 2001. R. Kakarala was supported by Motorola. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Eric L. Miller.

R. Kakarala was with Motorola Australian Research Centre, Botany, NSW 2019, Australia, on leave from the University of Auckland, Auckland, Australia. He is now with Agilent Laboratories, Palo Alto, CA 94304 USA (e-mail: kakarala@labs.agilent.com; www.labs.agilent.com/personal/Ram_Kakarala).

P. O. Ogunbona is with the Motorola Australian Research Centre, Botany, NSW 2019, Australia (e-mail: pogunbon@arc.corp.mot.com).

Publisher Item Identifier S 1057-7149(01)03272-9.

The SVD is closely linked with the concepts of principal component analysis (PCA) and Karhunen–Loève transform (KLT). The relationships among SVD, PCA, and KLT are discussed in detail elsewhere [1], [2, Ch. 8], [3], but it is useful to review the concepts briefly here in order to establish relationships with recent research. Let the $M \times N$ matrix X be interpreted as a data matrix containing, for example, M measurements on each of N individuals. The data matrix is centered by removing the mean in each row from elements of that row; in matrix terms this is $\bar{X} = X(I_N - (1/N)e_N e_N^t)$, where I_N is the $N \times N$ identity, and e_N is the $N \times 1$ vector containing all ones. Then $\bar{X}\bar{X}^t$ is the scatter matrix of the data, and following an eigen-decomposition, it may be written $\bar{X}\bar{X}^t = U\Lambda U^t$, where Λ is the diagonal matrix of eigenvalues. The *principal components* are obtained by the transformation $Y = U^t\bar{X}$, showing that Y is the same as \hat{X} above, aside from mean removal. Since the scatter matrix is a scalar multiple of the covariance matrix, we may also interpret U as the KLT for the *sample* in X . Obviously, U is not the KLT for the population, if X is sampled from a random vector.

The links between the multiresolution SVD that we introduce below and other, well-known, multiresolution decompositions, including wavelets and filterbanks, may be examined by employing the system approach provided by Unser [4, p. 47]. Fig. 1(a) shows how a sequence of p -dimensional vectors, denoted $X(k)$, is generated from the input sequence $x(k)$. Specifically

$$X(k) = [x(kp) \quad x(kp-1) \quad \cdots \quad x(kp-p+1)]^t. \quad (2)$$

Suppose that a filtered vector output signal $\hat{X}(k)$ is constructed by the formula

$$\hat{X}(k) = \sum_{n=-\infty}^{\infty} H_n X(k-n). \quad (3)$$

Here, H_n is a sequence of $p \times p$ matrices. With the $p \times p$ transfer function matrix being denoted $H(z) = \sum_n H_n z^{-n}$, we obtain the z -transform representation

$$\hat{X}(z) = H(z)X(z) \quad (4)$$

where $\hat{X}(z)$ and $X(z)$ are, respectively, the z -transforms of the vector sequences $\hat{X}(k)$ and $X(k)$. Fig. 1(b) provides the system block diagram, which is simply a polyphase analysis filterbank [15]. A multiresolution decomposition is obtained from this system by recursively applying the analysis filterbank to one or more of the components of the output vector $\hat{X}(k)$. For example, a wavelet multiresolution decomposition is obtained by setting $p = 2$, and setting $H(z) = \text{diag}[h_{1p}(z), h_{1p}(z)]$,

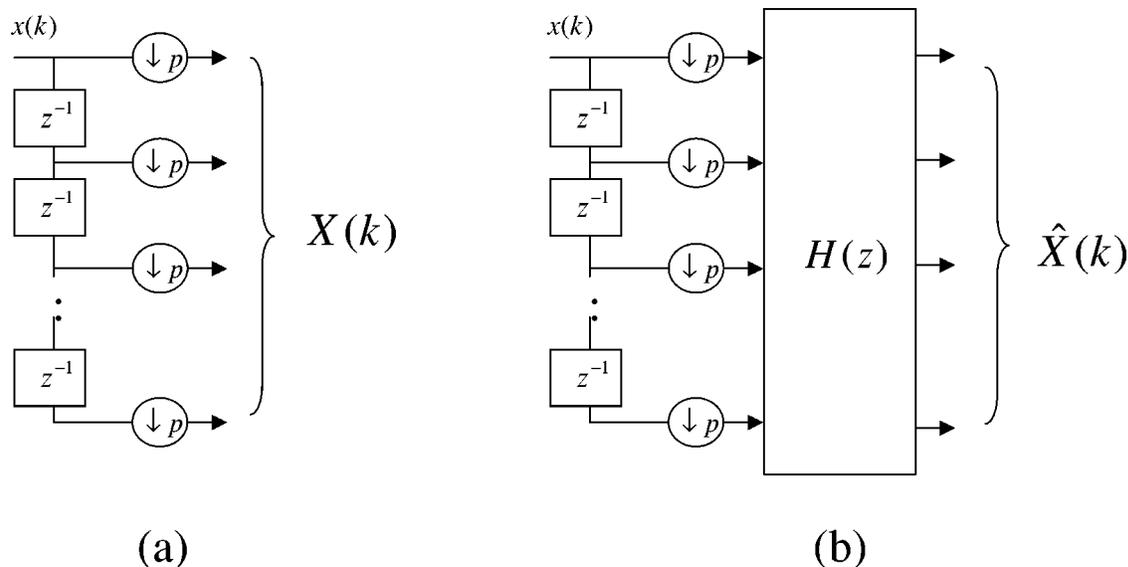


Fig. 1. (a) Signal $x(k)$ divided into nonoverlapping blocks of length p , each block contained in vector $X(k)$. (b) System diagram for a filter $H(z)$, operating on $X(k)$ and producing output $\hat{X}(k)$.

where $h_{lp}(z)$ and $h_{hp}(z)$ are, respectively, the transfer functions of suitably constructed lowpass and highpass filters. In the wavelet case, further levels of decomposition are obtained by recursively applying this filterbank to the output of the lowpass filter $h_{lp}(z)$.

A SVD may also be cast in this framework. Suppose we group M successive column vectors $X(k)$ into a $p \times M$ matrix, denoted X . Then, using (3) with $H_n = U^t$ for $n = 0$, and $H_n = 0$ otherwise, we obtain that $\hat{X} = U^t X$, where \hat{X} is the $p \times M$ matrix whose columns are M successive vectors $\hat{X}(k)$. If U is the matrix of left singular vectors of X , then $X = U\hat{X}$ is simply a form of SVD, as mentioned above. Therefore, the SVD may be viewed as a filterbank whose transfer function is a constant matrix, i.e., $H(z) = U^t$, where the matrix contains the signal's left singular vectors. The multiresolution SVD described below essentially repeats this construction to obtain successive levels of decomposition.

The idea of using the left singular vectors to filter the input signal has been proposed in other papers on filterbanks [4]–[7]. Those papers focus on how “optimal” filters may be designed for energy compaction, where optimality is defined as an ensemble-averaged measure. In this paper, we focus not on filter design, but rather on how a multiresolution SVD may be constructed and, more importantly, how it may be interpreted. Our approach focuses on deterministic signals, although the likelihood statistics proposed below implicitly assume normal distributions.

The organization of this paper is as follows. In Section III, we show how the SVD may be developed into a multiresolution analysis. The computational complexity of this decomposition is analyzed in Section IV. In Section V, the main results of this paper are presented:

- 1) concept of coding gain is interpreted in terms of the standard likelihood ratio for sphericity of principal components;

- 2) self-similarity of a signal at different scales is determined by checking for repeating eigen-decompositions;
- 3) mean squared error (MSE) between two signals is decomposed into a sum of MSEs over the principal components.

In Sections VI and VII, examples are given of signals and images analyzed using the proposed techniques.

II. NOTATION AND CONVENTIONS

Throughout this paper, the following notation and conventions are used. A N element row vector is indexed $[x(1), \dots, x(N)]$, and similarly a $M \times N$ matrix X is indexed

$$X = \begin{bmatrix} x(1, 1) & x(1, 2) & \cdots & x(1, N) \\ x(2, 1) & \vdots & \ddots & \vdots \\ x(M, 1) & \cdots & \cdots & x(M, N) \end{bmatrix}.$$

The k th row of matrix X is denoted $X(k, \cdot)$, and the k th column is denoted $X(\cdot, k)$. The methods described in this paper are for real-valued signals, but they extend to complex-valued signals by replacing every instance of transpose with conjugate-transpose.

Every positive semidefinite matrix W has an eigen-decomposition $W = U\Lambda U^t$, where U is an orthogonal matrix of eigenvectors, and Λ is the diagonal matrix of eigenvalues arranged in decreasing order. Noting that if u is an eigenvector of W with eigenvalue λ , then so is $-u$, we henceforth assume without loss of generality that the first nonzero element of each eigenvector is positive.

Henceforth, a matrix denoted S (possibly with subscripts) represents a diagonal matrix of singular values. The singular values are written $s(k)$ for $1 \leq k \leq M$, or sometimes $s(X; k)$ when the matrix needs to be identified. *Singular values are always assumed to be in arranged in decreasing order so that*

$s(1) \geq s(2) \geq \dots \geq s(M)$. Note that the SVD $X = USV^t$ may be written as a sum of outer products

$$X = \sum_{k=1}^M s(k)U(\cdot, k)V(\cdot, k)^t. \quad (5)$$

Each outer product $U(\cdot, k)V(\cdot, k)^t$ is a rank one matrix, and the partial sum

$$X^{(q)} = \sum_{k=1}^q s(k)U(\cdot, k)V(\cdot, k)^t \quad (6)$$

has rank q for $q \leq \text{rank}(X)$. This partial sum has an important approximation property [9], which may be described as follows. The Frobenius norm $\|\cdot\|_F$ of X is defined as

$$\|X\|_F = \sqrt{\sum_{k=1}^M \sum_{\ell=1}^N |x(k, \ell)|^2}.$$

Then, for any matrix Y with $\text{rank}(Y) = q \leq \text{rank}(X)$, we have that

$$\|X - Y\|_F \geq \|X - X^{(q)}\|_F = \sqrt{s(X; q+1)^2 + \dots + s(X; M)^2}. \quad (7)$$

In this sense, $X^{(q)}$ provides the best rank q approximation to X . Note that (7) implies

$$\|X\|_F = \sqrt{s(X; 1)^2 + \dots + s(X; M)^2}. \quad (8)$$

The $N \times N$ row centering matrix $H_N = I_N - (1/N)e_N e_N^t$, which appeared in Section I, is symmetric and idempotent, since $H_N H_N = H_N$. It is easily shown that the singular values of H_N are $s(1) = \dots = s(N-1) = 1$, and $s(N) = 0$.

Suppose now that $\bar{X} = XH_N$ is the mean corrected matrix for some $M \times N$ matrix X , with $M \leq N$. Let the SVD of \bar{X} be denoted $\bar{X} = AS(\bar{X})B^t$, with A the eigenvector matrix for $\bar{X}\bar{X}^t$; $S(\bar{X})$ the singular value matrix; and B the $N \times M$ matrix of eigenvectors of $\bar{X}^t\bar{X}$. As in (6), let

$$\bar{X}^{(q)} = \sum_{k=1}^q s(\bar{X}; k)A(\cdot, k)B(\cdot, k)^t. \quad (9)$$

From (7) it follows that $\bar{X}^{(q)}$ is the optimum rank q approximation to \bar{X} . Suppose now that the SVD of X is $X = USV^t$ and that $X^{(q)}$ is obtained as in (6). Using the results in [8, Ch. 3], we may establish some important facts about the respective SVDs of \bar{X} and X .

- 1) For $k = 1, \dots, M$, we have

$$s(\bar{X}; k) \leq s(X; k). \quad (10)$$

A proof is provided in the Appendix.

- 2) From (10), it follows that mean correction generally reduces the error in rank q approximations

$$\|\bar{X} - \bar{X}^{(q)}\|_F \leq \|X - X^{(q)}\|_F. \quad (11)$$

- 3) Let $\Upsilon = (1/N)X e_N e_N^t$ be the mean matrix (constant along rows). Then from (11), it follows that adding the

mean to the rank q approximation to \bar{X} , i.e., forming $\bar{X}^{(q)} + \Upsilon$, gives a better approximation to the original matrix X than simply using $X^{(q)}$ since

$$\|X - (\bar{X}^{(q)} + \Upsilon)\|_F = \|\bar{X} - \bar{X}^{(q)}\|_F \leq \|X - X^{(q)}\|_F. \quad (12)$$

Hence, from the viewpoint of approximation, it is better to remove the mean, form the rank q approximation to the corrected matrix, and add the mean back at the end. In [10, p. 593], it is shown that the columns of $\bar{X}^{(q)} + \Upsilon$ are the optimum q -dimensional *subspace* approximation to the corresponding columns of X .

III. MULTIREOLUTION FORM OF THE SVD

This section describes how the multiresolution SVD may be constructed. Recall that in the dyadic wavelet transform, the signal is filtered separately by low and high pass filters, and the output of each is decimated by a factor of two. This procedure is recursively repeated on the decimated *lowpass* output, until the desired level of decomposition is achieved.

The basic idea behind the multiresolution SVD is to replace filtering with SVD at each level of approximation. The idea is now described, initially for one-dimensional (1-D) signals, and for the dyadic case. Extensions to higher dimensions, and to p -adic decompositions, are described subsequently.

A. One-Dimensional Case

Let $X = [x(1), \dots, x(N)]$ represent a finite-extent, 1-D signal. Assume that N is divisible by 2^L for some $L \geq 1$. Let the data matrix at the first level, denoted X_1 , be constructed so that its top row contains the odd-numbered samples, and the bottom row contains the even-numbered samples

$$X_1 = \begin{bmatrix} x(1) & x(3) & \dots & x(N-1) \\ x(2) & x(4) & \dots & x(N) \end{bmatrix}. \quad (13)$$

The corresponding centered matrix is $\bar{X}_1 = X_1 H_{N/2}$. Let U_1 be the eigenvector matrix bringing the scatter matrix $T_1 = \bar{X}_1 \bar{X}_1^t$ into diagonal form: $U_1^t T_1 U_1 = S_1^2$, where $S_1^2 = \text{diag}\{s_1(1)^2, s_1(2)^2\}$ contains the squares of the two singular values, with $s_1(1) \geq s_1(2)$. Now let $\hat{X}_1 = U_1^t \bar{X}_1$, so that $\bar{X}_1 = U_1 \hat{X}_1$.

The top row of \hat{X}_1 , namely $\hat{X}_1(1, \cdot)$, contains the principal component that corresponds to the largest eigenvalue, and may be considered the “smooth” or “approximation” component when pairs of elements in X are analyzed. The bottom row contains the “detail” component that corresponds to the smallest eigenvalue. Let $\Phi_1 = \hat{X}_1(1, \cdot)$, and $\Psi_1 = \hat{X}_1(2, \cdot)$ represent the smooth and detail components, respectively. Note that Φ_1 and Ψ_1 are uncorrelated since the rows of \hat{X}_1 have zero mean, and moreover

$$\hat{X}_1 \hat{X}_1^t = S_1^2. \quad (14)$$

Hence the signal has been decomposed into uncorrelated smooth and detail components.

The next level of the multiresolution SVD repeats the procedure described above, but now using the smooth component Φ_1 in place of X . This procedure is repeated recursively L times.

The procedure may now be described formally. Let $\Phi_0(1, \cdot) = X$, so that the initial “smooth” component is the original signal. For each level ℓ , the vector Φ_ℓ has $N_\ell = N/2^\ell$ elements, which are denoted as follows:

$$\Phi_\ell = [\phi_\ell(1), \dots, \phi_\ell(N_\ell)].$$

The L -level dyadic SVD is defined by the following equations. For $\ell = 1, \dots, L-1$, set

$$X_\ell = \begin{bmatrix} \phi_{\ell-1}(1) & \phi_{\ell-1}(3) & \cdots & \phi_{\ell-1}(2N_\ell - 1) \\ \phi_{\ell-1}(2) & \phi_{\ell-1}(4) & \cdots & \phi_{\ell-1}(2N_\ell) \end{bmatrix} \quad (15)$$

$$\bar{X}_\ell = X_\ell H_{N_\ell} \quad (16)$$

$$T_\ell = \bar{X}_\ell \bar{X}_\ell^t = U_\ell S_\ell^2 U_\ell^t \quad (17)$$

$$\hat{X}_\ell = U_\ell^t \bar{X}_\ell \quad (18)$$

$$\Phi_\ell = \hat{X}_\ell(1, \cdot); \quad \Psi_\ell = \hat{X}_\ell(2, \cdot). \quad (19)$$

Note that in (17), we require the singular values to be arranged so that $s_\ell(1) \geq s_\ell(2)$.

A complete signal representation may be obtained from the dyadic multiresolution SVD as follows. To fully specify X , it is sufficient to store the lowest resolution smooth component Φ_L , and the detail components Ψ_ℓ , for $\ell = 1, \dots, L$. In addition, the mean vectors $\mu_\ell = (1/N)_\ell X_\ell e_{N_\ell}$, and the eigenvector matrices U_ℓ must also be stored. Hence, one may write the dyadic multiresolution SVD as the following transformation:

$$X \rightarrow \{\Phi_L, \{\Psi_\ell\}_{\ell=1}^L, \{U_\ell\}_{\ell=1}^L, \{\mu_\ell\}_{\ell=1}^L\}. \quad (20)$$

It is easy to see how X may be reconstructed from the right hand side, since each of the steps in (15)–(18) is reversible.

The same procedure may be applied without mean removal. Although mean removal improves approximation (see Section II), reasons to *not* remove the mean include: to reduce computation, and to resolve the mean-squared error between two signals into meaningful components (see Section V-D). Note that without mean removal, the components of \hat{X}_ℓ are not necessarily uncorrelated, only orthogonal.

A useful inequality for singular values at different resolution levels is

$$s_{\ell+1}(1)^2 + s_{\ell+1}(2)^2 \leq s_\ell(1)^2 \quad (21)$$

if mean correction is used. If the mean is not removed, then equality is obtained in (21). This inequality follows from (8), (10), and (14).

To illustrate the points above, we provide the following example.

Example 1: Let X be the eight-element vector $[1, 2, 3, 4, 5, 6, 7, 8]$. We compute its two-level dyadic SVD as follows. We have that

$$X_1 = \begin{bmatrix} 1 & 3 & 5 & 7 \\ 2 & 4 & 6 & 8 \end{bmatrix}.$$

The scatter matrix of the columns is

$$T_1 = \bar{X}_1 \bar{X}_1^t = \begin{bmatrix} 20 & 20 \\ 20 & 20 \end{bmatrix}.$$

The singular values of this matrix are $s_1(1) = \sqrt{40}$, $s_1(2) = 0$. The corresponding eigenvectors are $U_1(\cdot, 1) = (1/\sqrt{2})[1, 1]^t$, and $U_1(\cdot, 2) = (1/\sqrt{2})[1, -1]^t$. The transformation (18) produces

$$\hat{X}_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} -6 & -2 & 2 & 6 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

The next level operates on $\Phi_1 = \hat{X}_1(1, \cdot) = \frac{1}{\sqrt{2}}[-6, -2, 2, 6]$, and produces (with $\bar{X}_2 \bar{X}_2^t$ found to be a constant matrix with all entries equal to 16, and U_2 found to be the same as U_1)

$$\hat{X}_2 = \begin{bmatrix} -4 & 4 \\ 0 & 0 \end{bmatrix}.$$

Hence, the complete two-level ($L = 2$) SVD is

$$X \rightarrow \left\{ \begin{array}{l} [-4, 4], \{[0, 0], [0, 0, 0, 0]\}, \\ \left\{ \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \right\}, \\ \left\{ \frac{1}{\sqrt{2}} \begin{bmatrix} -4 \\ 0 \end{bmatrix}, \begin{bmatrix} 4 \\ 5 \end{bmatrix} \right\} \end{array} \right\}.$$

The procedure described may be generalized in several ways. First, it is possible to decompose the detail component Ψ_ℓ at each level ℓ into two further components, using the same method as for the smooth component. In this manner a general transformation (similar in spirit to the wavelet packet approach [11]) is obtained. Second, one may divide the original signal X into nonoverlapping blocks of length $p > 2$ (assuming that N is divisible by p^L). At each level, a $p \times p$ scatter matrix is formed, and its eigenvector matrix U is employed to decorrelate the centered block matrix using (18). The p principal components obtained may be ordered by the eigenvalue; the component corresponding to the maximum eigenvalue is essentially the “smooth” component, and this may be further decomposed to form the next level. Combinations of these generalizations are also possible.

B. Two and Higher Dimensions

The multiresolution SVD may be formulated in two and higher dimensions. We show here in detail how the formulation works in two dimensions, and indicate briefly the extension to higher dimensions.

Suppose that X is a $M \times N$ matrix. As in the 1-D case, we decompose x into $p \times q$ blocks, and analyze the covariance of the blocks. While p, q may be chosen arbitrarily, it is convenient to work with $p = q = 2$, which is perhaps the smallest truly two-dimensional (2-D) block format.

The first level of decomposition proceeds as follows. Divide X into nonoverlapping 2×2 blocks, and arrange each block into a 4×1 vector by stacking columns to form the data matrix X_1 . The blocks may be taken in transpose raster-scan manner, that is, proceeding downward first, and then to the right. Specifically, in terms of the elements of X , the first data matrix is shown in the equation at the bottom of the next page. The eigen-decomposition of the 4×4 scatter matrix is $T_1 = \bar{X}_1 \bar{X}_1^t = U_1 S_1^2 U_1^t$, in which we choose to have the singular values arranged in decreasing order: $s_1(1) \geq s_1(2) \geq s_1(3) \geq s_1(4)$.

As in the 1-D case, let $\hat{X}_1 = U_1^t \bar{X}_1$. Note that the top row of \hat{X}_1 corresponds to the largest eigenvalue, and therefore may be considered the “smooth” component inherent in the 2×2 blocks. The remaining rows contain the “detail” components, which may correspond to edges, or texture, in an image for example. The rows are orthogonal (and hence uncorrelated because they have zero mean), since $\hat{X}_1 \hat{X}_1^t = S_1^2$.

Note that the elements in each row may be rearranged to form a $M/2 \times N/2$ matrix. To proceed to the next stage of the decomposition, let Φ_1 denote the $M/2 \times N/2$ matrix formed by rearranging the row $\hat{X}_1(1, \cdot)$ into a matrix, by first filling in the columns, and next the rows, i.e., as shown in (22) at the bottom of the page. Essentially, Φ_1 forms the “smooth” image component inherent in the collection of 2×2 blocks. In a similar way, each of the three remaining rows $\hat{X}_1(2, \cdot)$, $\hat{X}_1(3, \cdot)$, $\hat{X}_1(4, \cdot)$ may be rearranged into $M/2 \times N/2$ matrices, which we denote respectively by $\Psi_1^{(1)}$, $\Psi_1^{(2)}$, and $\Psi_1^{(3)}$.

The next level of the transform now proceeds as above, but with X replaced by Φ_1 . The process is repeated recursively for L levels, using equations analogous to (15)–(18). The complete decomposition may be stated as follows:

$$X \rightarrow \left\{ \Phi_L, \{ \Psi_\ell^{(1)}, \Psi_\ell^{(2)}, \Psi_\ell^{(3)} \}_{\ell=1}^L, \{ U_\ell \}_{\ell=1}^L, \{ \mu_\ell \}_{\ell=1}^L \right\}. \quad (23)$$

As in the 1-D case, the decomposition may be performed without mean removal, with the same benefits and drawbacks. Furthermore, the inequality (21) extends to two dimensions, with

$$s_{\ell+1}(1)^2 + \dots + s_{\ell+1}(4)^2 \leq s_\ell(1)^2. \quad (24)$$

Again, equality is obtained if mean correction is not employed.

The following example illustrates the concepts previously discussed.

Example 2: The image “boats” in Fig. 2 is used to compute the two-level SVD that is shown in Fig. 3. The transformation matrices are as follows:

$$U_1 = \begin{bmatrix} 0.49 & 0.52 & 0.48 & 0.51 \\ 0.50 & 0.48 & -0.52 & -0.49 \\ 0.49 & -0.48 & 0.51 & -0.50 \\ 0.50 & -0.52 & -0.49 & 0.49 \end{bmatrix}$$

$$U_2 = \begin{bmatrix} 0.49 & 0.56 & 0.44 & 0.49 \\ 0.50 & -0.43 & 0.56 & -0.49 \\ 0.49 & 0.44 & -0.56 & -0.49 \\ 0.50 & -0.54 & -0.43 & 0.50 \end{bmatrix}.$$

Each of the columns of the above matrices may be rearranged into a 2×2 spatial filter, according to the pixel elements that they operate on. For example, the columns of U_1 may be rearranged into the 2×2 filters shown below [from left to right, rearrangements of $U_1(\cdot, 1)$, ..., $U_1(\cdot, 4)$]

$$\begin{bmatrix} 0.49 & 0.49 \\ 0.50 & 0.50 \end{bmatrix}, \quad \begin{bmatrix} 0.52 & -0.48 \\ 0.48 & -0.52 \end{bmatrix}$$

$$\begin{bmatrix} 0.48 & 0.51 \\ -0.52 & -0.49 \end{bmatrix}, \quad \begin{bmatrix} 0.51 & -0.50 \\ -0.49 & 0.49 \end{bmatrix}.$$

Note that $U_1(\cdot, 1)$ is a local average (smoothing), $U_1(\cdot, 2)$ is a vertical edge filter, $U_1(\cdot, 3)$ is a horizontal edge filter, and $U_1(\cdot, 4)$ is a diagonal edge filter. A similar interpretation may be given to the columns of U_2 . Note that for each ℓ , the leftmost eigenvector $U_\ell(\cdot, 1)$, which corresponds to the largest eigenvalue, is simply a local average: *hence, for this image, the “smooth” components Φ_1 and Φ_2 are obtained as a local average of adjacent pixels.* Furthermore, the detail components $\Psi_\ell^{(k)}$ for $k = 1, 2, 3$ essentially contain one of the following: horizontal edges, vertical edges, or diagonal edges. The pattern is visually apparent in Fig. 3.

$$X_1 = \begin{bmatrix} x(1, 1) & x(3, 1) & \dots & x(M-1, 1) & x(1, 3) & \dots & x(M-1, N-1) \\ x(2, 1) & x(4, 1) & \dots & x(M, 1) & x(2, 3) & \dots & x(M, N-1) \\ x(1, 2) & x(3, 2) & \dots & x(M-1, 2) & x(1, 4) & \dots & x(M-1, N) \\ x(2, 2) & x(4, 2) & \dots & x(M, 2) & x(2, 4) & \dots & x(M, N) \end{bmatrix}$$

$$\Phi_1 = \begin{bmatrix} \hat{x}_1(1, 1) & \hat{x}_1\left(1, \frac{M}{2} + 1\right) & \dots & \hat{x}_1\left(1, \frac{N}{2} \left(\frac{M}{2} - 1\right)\right) \\ \hat{x}_1(1, 2) & \vdots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ \hat{x}_1\left(1, \frac{M}{2}\right) & \dots & \dots & \hat{x}_1\left(1, \frac{N}{2} \frac{M}{2}\right) \end{bmatrix} \quad (22)$$

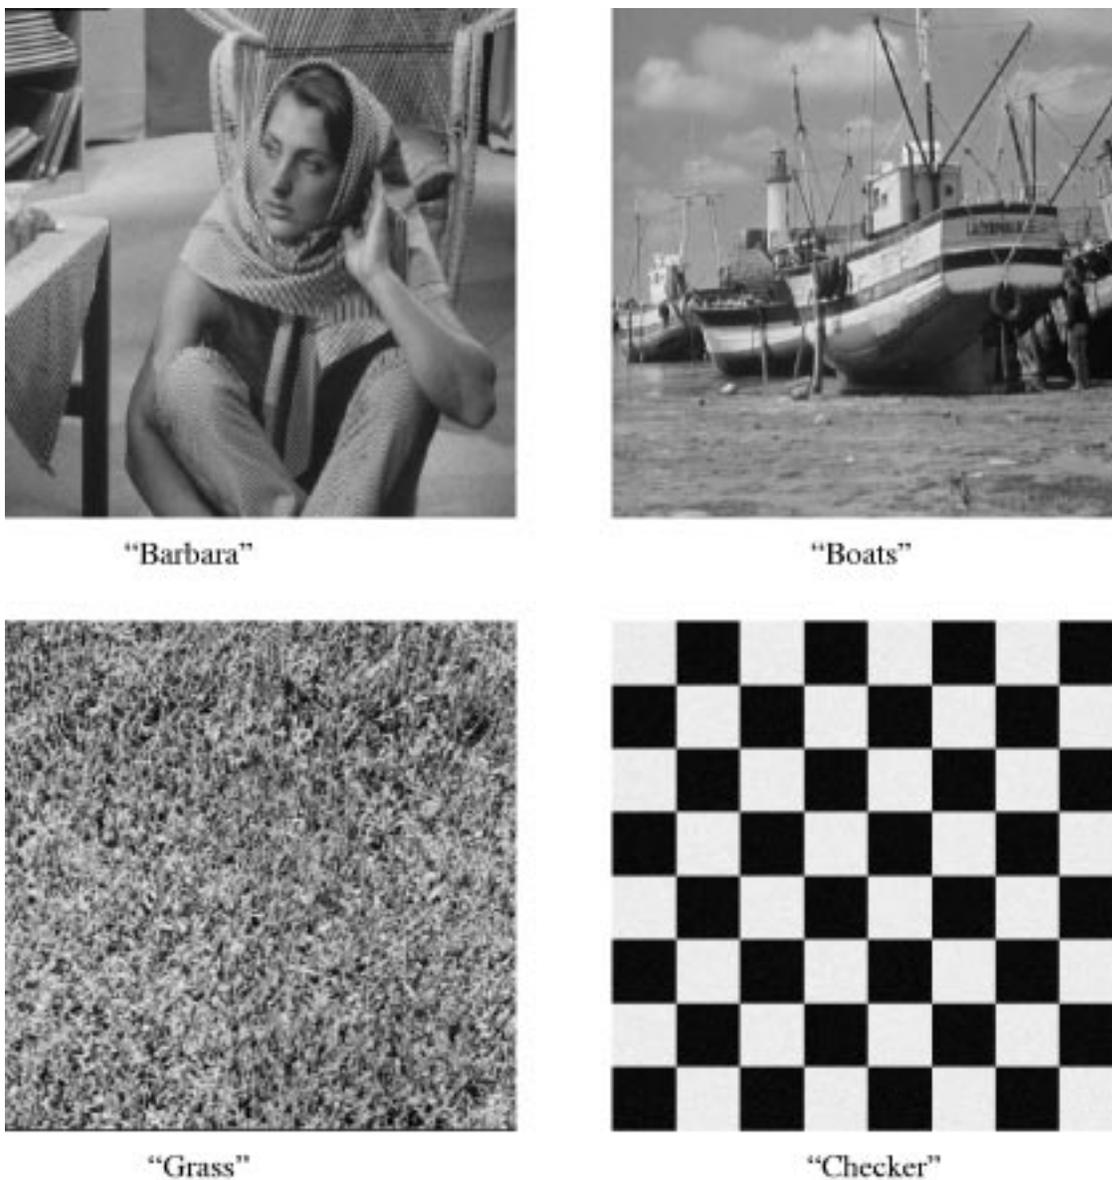


Fig. 2. Four test images used in the paper.

The emergence of edge filters in the eigenvectors in this example is an interesting phenomenon. It is essentially caused by two factors: 1) the existence of positive correlation between neighboring pixels and 2) the use of a 2×2 block size. The first factor, positive correlation, causes the largest eigenvalue's eigenvector to have all positive entries (the reason why is discussed in detail in Section V). Given that, the remaining eigenvectors have to have at least one negative entry in order to be orthogonal to the largest eigenvector. Now, the second factor, a 2×2 block size, means that the remaining eigenvectors are either edge or corner filters, depending on whether they have an even or odd number of negative entries. The emergence of edge rather than corner filters for this particular image indicates that the strength of edges exceeds that of corners in this case.

The same remarks concerning generalizations of the multiresolution SVD made in the 1-D case are applicable in two dimensions. Clearly, we may use any size block, and moreover, we may choose to decompose additional components in addition to (or instead of) the one having largest eigenvalue. In $n > 2$

dimensions, we may use blocks having n elements in each dimension, e.g., $3 \times 3 \times 3$ in three dimensions; of course, this may be generalized to any size blocks.

IV. COMPLEXITY ANALYSIS

Here, we show that a p -adic multiresolution SVD has $O(N)$ complexity for 1-D signals $X = [x(1), \dots, x(N)]$. A similar approach to the one presented below shows that in two-dimensions, the complexity is also linear in the number of samples (pixels). A count of the number of arithmetic operations is given to provide a more precise analysis of the computational burden.

In the p -adic decomposition, each level uses only $1/p$ of the number of samples as the next lower level. The maximum number of useful levels of decomposition with a length N signal is $L = \log_p(N) - 1$. Note that at level $\log_p(N)$, the matrix $\bar{X}_{\log_p(N)}$ contains only zeros after mean removal, so this level is not actually necessary.

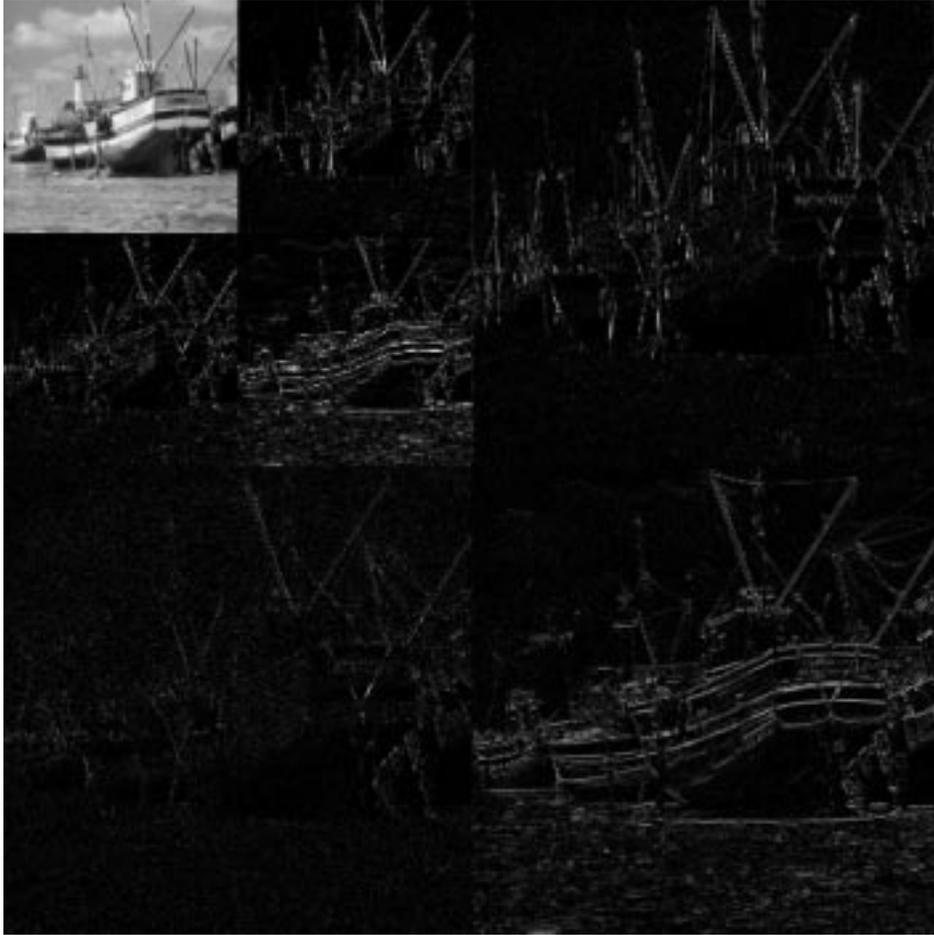


Fig. 3. Two-level SVD of the image “boats” is shown. The subimages are arranged in clockwise order of decreasing eigenvalue. To aid visibility, the subimages have been contrast-stretched so that the pixel values fit into the range $0, \dots, 255$. Moreover, the absolute value of the edge images is shown since the negative values are as significant as the positive ones.

Suppose that x has length $N = p^{L+1}$. We count the number of operations to compute \hat{X}_1 . To form the centered block matrix \bar{X}_1 , we note that each row has N/p elements. Removing the mean in each row requires $(N/p - 1)$ additions, one division, and N/p subtractions, for a total of $2N/p$ operations. Hence, the total cost for all p rows is $2N$ operations.

Next, to compute each element of the $p \times p$ scatter matrix T_1 , we need a total of N/p multiplications and $(N/p - 1)$ additions. Since the scatter matrix is symmetric, we need compute only $p(p + 1)/2$ entries. Hence, the total cost of computing the scatter matrix from \bar{X}_1 is $(2N - p)(p + 1/2)$ operations. Diagonalizing the scatter matrix may be done in a finite number of operations, which we denote Q_p . A rough estimate is $Q_p = 10p^3$ for $p > 2$ (estimated using the “flops” counter in MATLAB), which is $O(p^3)$; see also [3, p. 165]. Computing \hat{X}_1 requires p multiplications and $p - 1$ additions for each of the N elements, giving a total of $N(2p - 1)$ operations for the entire matrix. Adding all of these together, we find that the first level requires

$$\begin{aligned} & 2N + (2N - p)\frac{p+1}{2} + N(2p - 1) + Q_p \\ & = N(3p + 2) - \frac{p(p+1)}{2} + Q_p \end{aligned}$$

operations.

Let $C_p = Q_p - p(p + 1)/2$ be the cost independent of N . It is easy to see that computing the next level of decomposition requires exactly $(3p + 2)N/p + C_p$ operations, and in general, the ℓ th level requires $(3p + 2)(N/p^{\ell-1}) + C_p$ operations. Hence the total cost for L levels is

$$\begin{aligned} \sum_{\ell=1}^{L-1} (3p + 2)\frac{N}{p^{\ell-1}} + C_p & \leq N\frac{3p+2}{1-\frac{1}{p}} + LC_p \\ & = N\frac{3p^2 + 2p}{p-1} + (\log_p(N) - 1)C_p. \end{aligned} \quad (25)$$

For large N , the right hand side is bounded above by $5p^2N$, which is a linear function of N . Hence, the overall complexity is $O(N)$.

For comparison, an N -point FFT requires $2N \log_2(N)$ operations if N is a power of two. For $N = 256$, the FFT’s cost slightly exceeds the cost of computing the dyadic SVD up to seven levels.

V. INTERPRETATION OF DECOMPOSITION

The p -adic multiresolution SVD of a signal provides three matrices at each level: the singular value matrix S_ℓ , the decor-

relating filter matrix U_ℓ , and the orthogonal subband matrix \hat{X}_ℓ . Each of these sets of matrices provides useful information about the signal being analyzed. The singular values, for example, are proportional to the subband variances: the rows of \hat{X}_ℓ have zero mean, and furthermore their covariance matrix is $(1/N)\hat{X}_\ell\hat{X}_\ell^t = (1/N)S_\ell^2$. The objective of this section is to describe how the singular values, eigenvectors, and subbands may be interpreted and utilized.

The value of p discussed in this section refers to the number of elements in each block, regardless of dimension. For example, if the SVD is computed using 2×2 blocks on an image, then $p = 4$.

This section makes use of the log likelihood ratio for measuring the goodness of fit of data to a model. Although the likelihood ratio may, in principle, be formulated for any distribution, much simplification results if the data are normally distributed. Clearly, there are many instances where data are not normally distributed, and the application of the specific likelihood ratio derived under the normal assumption is questionable. However, it is still worthwhile to determine the form of the likelihood ratio for the normal case, because the form itself gives insight into the characteristics of the data that are worth measuring, such as coding gain (see Section V-B).

A. Eigenvectors as Decorrelating Filters

It is reasonable to wonder whether the principal component with the largest singular value is the “smooth” component, as it is called in Section III. Indeed, this terminology is only appropriate if the largest component eigenvector, which is the first column $U_\ell(1, \cdot)$ of the decorrelating matrix, is a lowpass filter. This need not be the case in general. However, it does occur when all of the entries of this eigenvector have the same sign; using the sign convention in Section II, we may suppose all the entries to be nonnegative, and the corresponding filter is therefore lowpass.

An important case when $U_\ell(\cdot, 1)$ is positive is when all of the entries in T_ℓ are positive. Then, the Perron–Frobenius theorem [12, p. 542] states that the largest eigenvalue has algebraic multiplicity one, exceeds in modulus all the other eigenvalues, and moreover has an associated eigenvector whose entries are all positive. Hence, for positively correlated data, we are guaranteed that the largest principal component is unique, and is extracted by a positively weighted moving average.

Example 2 shows that the leading eigenvector is nearly $(1/2)e_4$ for both U_1 and U_2 . In general, a measure of the data’s isotropy is the extent which the constant vector $(1/\sqrt{p})e_p$ fits as an eigenvector, corresponding to the largest eigenvalue, of the sample covariance matrix [13, p. 34]. Assuming normally distributed data, the goodness of fit may be measured from the likelihood ratio

$$\Gamma_{\text{ISO}} = N_\ell \left(\frac{s_\ell(1)^2}{p} e_p^t T_\ell^{-1} e_p + \frac{1}{s_\ell(1)^2 p} e_p^t T_\ell e_p - 2 \right). \quad (26)$$

Here, $N_\ell = N/p^\ell$ for length N data, and T_ℓ , $s_\ell(1)$ are, respectively, the scatter matrix and the largest singular value at the ℓ th level. It is known that, if the isotropy hypothesis is valid, then Γ_{ISO} is asymptotically distributed as χ^2 with $p - 1$ degrees of freedom (see [13, p. 63]). If we were attempting a formal hy-

pothesis test, then we would reject the isotropy model at the α level of significance if Γ_{ISO} exceeds the $(1 - \alpha) \times 100$ th percentile of the corresponding χ^2 distribution.

B. Coding Gain as Test of Sphericity

Suppose that a M -channel filterbank is employed to decompose a wide-sense stationary (WSS) input process. Let σ_k^2 denote the (ensemble) variance of the output of the k -th channel, for $k = 1, \dots, M$. The coding gain of a filterbank [6, Eq. (3)] is defined as the ratio of two errors: the numerator is the mean-squared reconstruction error when using direct quantization (PCM); and the denominator is the corresponding error when a filterbank’s output is quantized to the same bit rate. Modeling quantization as an additive WSS process, and assuming optimal bit allocation, it can be shown that coding gain is the ratio of the arithmetic mean and geometric means

$$G_{\text{SBC}} = \frac{\frac{1}{M} \sum_{k=1}^M \sigma_k^2}{\left(\prod_{k=1}^M \sigma_k^2 \right)^{1/M}}. \quad (27)$$

The well-known inequality between the arithmetic mean and the geometric mean shows that $G_{\text{SBC}} \geq 1$, with equality if and only if the variances are all equal. Filterbank design generally seeks to maximize coding gain.

The relationship between coding gain and the multiresolution SVD is now established. Suppose that at some level ℓ of a p -adic SVD decomposition, $p - k$ of the measured singular values in the matrix S_ℓ^2 are nearly equal. Under that supposition, if $k = 0$ then all of the singular values are nearly equal, and therefore “smooth” and “detail” components are not clearly identifiable in the signal. In such cases, one has to decide either to further decompose *all* of the components or none at all. This is particularly relevant in applications where the number of levels of decomposition are not known beforehand. If $k > 1$, then the last $p - k$ detail components, say, are equally significant, and should be treated in the same manner in further processing. The standard multivariate test for equally significant principal components (assuming normally distributed data) is now described [2, pp. 235–236]. Data with equally significant principal components are often described as *spherical*, and hence the corresponding test is known as a test of sphericity. For a data matrix having N_ℓ columns, the log likelihood ratio is

$$\Gamma_{\text{SP}} = \left(N_\ell - \frac{2p + 11}{6} \right) (p - k) \ln \frac{\frac{1}{p - k} \sum_{q=k+1}^p s_\ell(q)^2}{\left(\prod_{q=k+1}^p s_\ell(q)^2 \right)^{1/p - k}}. \quad (28)$$

It can be shown that Γ_{SP} is asymptotically distributed as a χ^2 distribution with $1/2(p - k + 2)(p - k - 1)$ degrees of freedom.

The likelihood ratio is based on the ratio of the arithmetic and geometric means of the measured squares of singular values. As mentioned above, these are proportional to the measured variances of the subbands. Hence, for normally distributed data,

the optimal statistic for determining whether smooth and detail components can be usefully separated is the measured coding gain.

C. Self-Similarity

The standard wavelet transform requires application of the same filterbank at each level. If the data are self-similar at different scales, then it is indeed appropriate to use the same decorrelating filterbank. However, it is an important question of how to test for self-similarity. One approach is to formulate the problem as one of common principal components (CPC) analysis [13]. In CPC, tests are constructed for the hypothesis that k sample covariance matrices have the same basis of eigenvectors. In our terminology, this is equivalent to the hypothesis that there exists a single orthogonal matrix W which can diagonalize each of the scatter matrices up to level k

$$W^t T_\ell W = S_\ell^2, \quad \ell = 1, \dots, k. \quad (29)$$

If this hypothesis holds, then the data may be said to have similar covariances at different levels of scale, which is an indication of "self-similarity."

One difficulty with this hypothesis is that the scatter matrix at level ℓ , namely T_ℓ , depends on the choice of $U_{\ell-1}$ to diagonalize the scatter matrix at the next lower level, $T_{\ell-1}$. There is no easy way to formulate this relationship. An alternative, and simpler, approach is to ask whether $U_{\ell-1}$ also diagonalizes T_ℓ at a given level ℓ . If this model, which we call the model of repeating principal components (RPC), is correct, then we have

$$U_{\ell-1}^t T_\ell U_{\ell-1} = S_\ell^2. \quad (30)$$

Using standard multivariate techniques [13, pp. 67–68], the log likelihood ratio statistic for testing this model is seen to be

$$\Gamma_{\text{RPC}} = N_\ell \log \frac{\det(\text{diag}(U_{\ell-1}^t T_\ell U_{\ell-1}))}{\det T_\ell} \quad (31)$$

where N_ℓ is the number of columns in \hat{X}_ℓ . The general theory of likelihood ratios [2, p. 124] shows that Γ_{RPC} is asymptotically distributed as χ^2 with $p(p-1)/2$ degrees of freedom.

D. Error Analysis

Both peak signal-to-noise ratio (PSNR) and mean-squared error (MSE) are widely used in image compression research as simple global measures of fidelity. Obviously, these global measures of error do not capture either the true nature or the visibility of the errors. Of great interest is how the error is distributed among the smooth, edge and texture regions of the image. For this, we may employ the multiresolution SVD, without mean removal.

Let Y be the result of lossy compression applied to a $M \times N$ image X . The MSE between X and Y is simply $(MN)^{-1} \|X - Y\|_F^2$. Suppose that the multiresolution SVD of X is computed, without mean removal. Let X_1 be the first data matrix (with dimensions $pq \times MN/pq$ for a $p \times q$ block decomposition). Now, realizing that the smooth and detail components Φ_1 , and $\Psi_1^{(k)}$, for $k = 1, \dots, pq-1$, are orthogonal, we may break the MSE into additive constituents along each of the components.

To do this, let U_1 be the eigenvector matrix diagonalizing $X_1 X_1^t$. Defining Y_1 to be the data matrix formed from Y (in

the same way that X_1 is formed from X), we construct a new matrix, denoted \tilde{Y}_1 , by the rule

$$\tilde{Y}_1 = U_1^t Y_1.$$

Note that the rows of \tilde{Y}_1 are not necessarily orthogonal, since U_1 is derived from the data in X . However, we have transformed Y_1 by employing the same spatial filters as used to obtain \hat{X}_1 , and therefore, we may compare corresponding components. This allows us to break up the MSE into the MSE within the span of \hat{X}_1 , and the MSE in the orthogonal complement of the span. The MSE within the span is simply the closest approximation to Y_1 from linear combinations of the p^2 orthogonal components (subbands) in \hat{X}_1 , and the MSE in the orthogonal complement is simply the residual.

The normalized orthogonal components lie in the columns of the V_1 matrix appearing in the SVD $X_1 = U_1 S_1 V_1^t$. Note that $V_1^t = S_1^{-1} \hat{X}_1$. Now, the projection of the error $\hat{X}_1 - \tilde{Y}_1$ onto the span of \hat{X}_1 is obtained from the $pq \times pq$ matrix P_E , which is defined as

$$P_E = (\hat{X}_1 - \tilde{Y}_1)(V_1)^t. \quad (32)$$

The projection onto the orthogonal complement of the column space of V_1 is determined by the matrix $I_{N_\ell} - V_1 V_1^t$. Hence, the total error may be written as the sum of components

$$\|X - Y\|_F^2 = \|P_E\|_F^2 + \|(\hat{X}_1 - \tilde{Y}_1)(I_{N_\ell} - V_1 V_1^t)\|_F^2. \quad (33)$$

The residual MSE is in the rightmost term. Noting that $\|P_E\|_F^2 = \text{trace}\{P_E P_E^t\}$, we may examine the magnitude of errors in the symmetric matrix $E = P_E P_E^t$. The diagonal entries of E , denoted $e(1, 1), \dots, e(pq, pq)$, reveal the magnitudes of error between Y and the original image X , divided among the components of X . For example $e(1, 1)$ measures the part of the MSE which lies in the smooth component Φ_1 , and $e(2, 2)$ is the error in Ψ_1 , the most significant detail component, and similarly for the remaining elements.

VI. CALCULATIONS FOR THE MARKOV-1 MODEL

To provide insight into the examples that follow in Section VII, we examine how the multiresolution SVD performs for a simple, but widely used, probabilistic image model. Suppose that we apply the decomposition to an image X whose pixel values are sampled from a zero-mean, wide-sense stationary random field, having the covariance

$$r(m, n) = c\rho^{|m|+|n|} \quad (34)$$

where c is an arbitrary positive constant, and $0 < \rho < 1$. This is the Markov-1 separable covariance model, often used in image processing for the analysis of transform coding [3, p. 508].

Assume that we perform a L -level multiresolution SVD of such an image, using 2×2 blocks as in Section III-B. Then the theoretical covariance matrix at level $\ell = 1$ is the Toeplitz matrix whose first row is $c[1, \rho, \rho, \rho^2]$. It is easy to see that the largest eigenvalue of this matrix is $(1 + \rho)^2$, with corresponding eigenvector $(1/2)e_4$. Hence, the smooth component Φ_1 is extracted by applying a 2×2 equally weighted local average to X , followed by decimation in both dimensions. Hence, the pixel

values in Φ_1 also have zero mean and form a wide-sense stationary random process. With some straightforward algebra, it may be shown that the covariance function of Φ_1 is

$$r_1(m, n) = c(1 + \rho)^2 \left(\frac{\rho(1 + \rho)}{2} \right)^{|m|+|n|}. \quad (35)$$

We see that Φ_1 also has a separable Markov-1 covariance. Therefore, it follows that Φ_ℓ , for $\ell > 1$, are Markov-1 covariance random fields. The Markov-1 structure is self-replicating under multiresolution SVD.

This random field has perfect isotropy in the sense defined above, since the largest eigenvector is $(1/2)c_4$ at every level ℓ . Moreover, the eigenvector matrix U_ℓ at every level ℓ is $U_H \otimes U_H$, where U_H is the Haar matrix

$$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$$

and \otimes denotes Kronecker product. Hence the principal components repeat exactly. However, it may be shown that the data become more spherical as ℓ increases. To see this, note that the sphericity is determined by

$$\frac{\frac{1}{4} \sum_{k=1}^4 s_\ell(k)^2}{\left(\prod_{k=1}^4 s_\ell(k) \right)^{1/4}} = \frac{\frac{1}{4} \text{trace}(C_\ell)}{\det(C_\ell)^{1/4}}$$

where C_ℓ is the theoretical covariance matrix at level ℓ . Because the Markov-1 structure is self-replicating, C_ℓ must be a Toeplitz matrix with first row $c_\ell[1, \rho_\ell, \rho_\ell, \rho_\ell^2]$. Note that $\det(C_\ell) = c_\ell^4(1 - \rho_\ell^2)^2$, which is a monotonically decreasing function of ρ_ℓ . From (35) we see that $\rho_{\ell+1} = \rho_\ell(\rho_\ell + 1)/2$, from which it follows that $\rho_{\ell+1} < \rho_\ell$. Hence the ratio of arithmetic to geometric means, which measures the departure from sphericity, decreases as the level increases.

VII. EXAMPLES

A four-level decomposition was performed on each of the four 512×512 images shown in Fig. 2. The images “boats” and “Barbara” are standard compression test images, and “grass” is one of the Brodatz textures [14]. The image “checker” is a synthetic image with uniformly distributed noise superimposed. The noise amplitude is 10% of the maximum grey level. In each case, the block size was 2×2 , and mean correction was used. The number of blocks at each level were $N_4 = 1024$, $N_3 = 4096$, $N_2 = 16384$, and $N_1 = 65536$.

The statistics Γ_{ISO} , Γ_{SP} , and Γ_{RPC} were computed to measure goodness of fit of, respectively, isotropy, sphericity, and common principal components.¹ As discussed in Section V, the aim is not to perform formal hypothesis testing, but rather to carry out exploratory data analysis.

A. Isotropy, Sphericity, and RPC

Table I shows the values of the isotropy statistic Γ_{ISO} . For comparison, the χ^2 distribution with $p - 1 = 3$ degrees of

¹Matlab code and test images are available at http://www.labs.agilent.com/personal/Ram_Kakarala/msvd.html.

TABLE I
VALUES OF ISOTROPY STATISTIC Γ_{ISO}

level\image	boat	barbara	grass	checker
1	6.3	74.4	209.8	2.74
2	7.2	5.7	76.6	2.26
3	13.7	1.7	17.5	2.13
4	5.2	6.4	2.5	4.18

TABLE II
VALUES OF SPHERICITY STATISTIC Γ_{SP} . MULTIPLY BY 10^3

level\image	boat	barbara	grass	checker
1	541.4	289.0	220.2	990.0
2	111.9	93.5	24.0	315.27
3	23.0	25.6	2.8	96.1
4	4.8	4.9	0.5	28.1

TABLE III
VALUES OF RPC STATISTIC Γ_{RPC}

level\image	boat	barbara	grass	checker
2	28.5	21.6	60.9	2.54
3	41.1	69.7	3.3	5.78
4	53.9	21.6	3.9	4.78

freedom has its 95th percentile at 7.81. It can be seen that the “boats” image fits the isotropy model well at levels 1, 2, and 4, and is at approximately the 99th percentile at level 3. The “Barbara” image has a grainy appearance due to having originally been scanned, and hence is significantly nonisotropic at level 1. However, at higher levels, the graininess does not play a role, and hence the image data fits the isotropy hypothesis. The rough texture in “grass” is far from isotropic at all levels, except for the fourth. At this level, a 2×2 block represents a 16×16 block in the original image, which contains the texturing element. The “checker” image is isotropic at all levels due to its simple square-wave structure.

Table II shows the values of the sphericity statistic Γ_{SP} with $k = 0$. For comparison, the χ^2 statistic with $(p+2)(p-1)/2 = 9$ degrees of freedom has its 95th percentile at 16.9. It can be seen that the data are highly nonspherical, as all values exceed 500; this implies a clear separation between smooth and detail components. Similar tests for $k > 0$ showed clear separations between the detail components, with the exception of the “checker” image. In that image, the first and second detail components $\Psi_\ell^{(1)}$ and $\Psi_\ell^{(2)}$, which correspond to the vertical and horizontal edge structures, are nearly equal in magnitude. This is reasonable given that the image is symmetric about a diagonal axis.

Table III shows the values of the RPC statistic Γ_{RPC} . As noted earlier, this statistic measures the extent to which the



Fig. 4. Image “boat” has been compressed to 0.5 bits/pixel. (a) Image “D97” uses the Daubechies (9, 7) wavelet and (b) image “HAAR” uses the HAAR wavelet.

scatter matrix $T_{\ell+1}$ is diagonalized by U_{ℓ} , the eigenbasis at the next lower level. For comparison, the χ^2 statistic with $p(p-1)/2 = 6$ degrees of freedom has its 95th percentile at 12.6. It can be seen that for the “boats” or “Barbara” images the repeating principal components model does not fit the data. For these images, it is truly necessary to change the decorrelating filters at each level of decomposition, up to level 4. However, for the synthetic “checker” image, and the “grass” texture image, the fit is very good, indicating self-similarity between levels 2 and 3 and levels 3 and 4.

B. Decomposition of MSE

The two images shown in Fig. 4 were obtained by compressing the “boats” image using two different wavelets. The image “D97” was compressed using the Daubechies (9, 7) wavelet and the image “HAAR” using the Haar wavelet. Both images were compressed to 0.5 bits per pixel, using the quantization scheme in the UICODER software package (see [15]). The PSNR of image “D97” is 32.4 dB, and the PSNR of image “HAAR” is 30.1 dB. On closer inspection, the blocking artifact in “HAAR” is visible.

Table IV shows values obtained from the decomposition of MSE as described in Section V-D. It can be seen that 98% of the total MSE for “D97” lies in the residual, and while 94% is the corresponding figure for “HAAR.” The blocking artifact that is visible in “HAAR” shows up as the slightly higher error in each of the edge components (vertical, horizontal, and diagonal), as well as in the substantially higher residual.

VIII. CONCLUSIONS

This paper has proposed a multiresolution form of the singular value decomposition for purposes of signal analysis and comparison. Methods for interpreting and using the singular values, eigenvectors, and subbands are discussed. In particular, the concept of coding gain was interpreted as a statistic for spherical principal components; the degree of self-similarity at

TABLE IV
DECOMPOSITION OF MSE INTO COMPONENTS

component\image	D97	HAAR
Φ_1 (smooth)	0.01	0.99
Ψ_1 (vertical edge)	0.11	0.73
Ψ_2 (horizontal edge)	0.06	0.81
Ψ_3 (diagonal edge)	0.22	1.27
Residual	37.1	59.4
Total MSE	37.5	63.2

adjacent levels was measured by the fit of repeating principal components; the mean squared error between two images was shown to be decomposable along the subbands.

It is worth considering how to compare the multiresolution SVD to existing multiresolution decompositions, e.g., wavelets. Comparison needs a basis. The multiresolution SVD provides an analysis tool to inquire into the properties (isotropy, sphericity, self-similarity) of signals. The basis for comparison is the extent that other transforms provide the same analysis tool. With the multiresolution SVD, we are able to talk in terms of signal properties because we are using a signal-adapted transform. Insofar as a nonsignal-adapted transform (such as wavelets) is used, there is no comparison—perhaps a contrast.

However, a valuable aspect of nonadapted transforms is that they provide a neutral basis for comparing different signals. To the extent that the basis is “natural,” signals may be meaningfully compared. For example, the Fourier basis is natural in the sense that it is motivated by physical considerations. There is no similarly motivated basis for multiresolution decompositions, and hence there will always be a multitude of approaches here.

We have not yet discussed the aspect of compression. Clearly, the SVD is attractive as a compression tool, since it yields op-

timal subrank approximation. However, a potential disadvantage of the SVD is that the decorrelating filters do not necessarily have linear phase. It would be interesting to investigate a “biorthogonal” form of the SVD, if one exists. Moreover, the effect of quantization on subrank approximation needs further study.

Finally, the multiresolution SVD may be viewed as a type of fast, approximate SVD. The relationships to other fast but approximate SVD algorithms, such as obtained by a Monte-Carlo approach [16], and the extent to which the actual SVD may be approximated by a multiresolution SVD, are topics deserving further study.

APPENDIX

Equation (10) may be shown as follows. Let \tilde{X} be the $N \times N$ (square) matrix obtained by placing X on top of $N - M$ rows of zeros. Then the *nonzero* singular values of \tilde{X} are easily seen to be exactly those of X . We may now use the result [8, p. 178] that for every pair of matrices A, B having the same dimensions

$$s(AB^t; k) \leq s(A; k)s(B; 1).$$

Now, letting $A = X$ and $B = H$, and using the fact that largest singular value of H is 1, we obtain (10).

ACKNOWLEDGMENT

The authors are grateful for the helpful suggestions of X. Zhang, W. Q. Li, S. Boland, M. M. Thomson, and F. Gu.

REFERENCES

- [1] J. J. Gerbrands, “On the relationships between SVD, KLT, and PCA,” *Pattern Recognit.*, vol. 14, pp. 375–381, 1981.
- [2] K. V. Mardia, J. T. Kent, and J. M. Bibby, *Multivariate Analysis*. London, U.K.: Academic, 1979.
- [3] A. K. Jain, *Digital Image Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1989.
- [4] M. Unser, “An extension of the Karhunen–Loève transform for wavelets and perfect reconstruction filterbank,” in *Proc. SPIE*, vol. 2034, A. F. Laine, Ed., 1993, pp. 45–56.
- [5] P. Moulin and M. K. Mihçak, “Theory and design of signal-adapted fir paraunitary filter banks,” *IEEE Trans. Signal Processing*, vol. 46, no. 4, pp. 920–929, 1998.
- [6] P. P. Vaidyanathan, “Theory of optimal orthonormal subband coders,” *IEEE Trans. Signal Processing*, vol. 46, pp. 1528–1543, June 1998.
- [7] A. Kiraç and P. P. Vaidyanathan, “Theory and design of optimum fir compaction filters,” *IEEE Trans. Signal Processing*, vol. 46, pp. 903–919, Apr. 1998.

- [8] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*. Cambridge, UK: Cambridge Univ. Press, 1991.
- [9] C. Eckart and G. Young, “The approximation of one matrix by another of lower rank,” *Psychometrika*, vol. 1, pp. 211–218, 1936.
- [10] C. R. Rao, *Linear Statistical Inference and its Applications*, 2nd ed. New York: Wiley, 1973.
- [11] R. R. Coifman and M. V. Wickerhauser, “Entropy based algorithms for best basis selection,” *IEEE Trans. Inform. Theory*, vol. 32, pp. 712–718, Mar. 1992.
- [12] P. Lancaster and M. Tismenetsky, *The Theory of Matrices*, 2nd ed. New York: Academic, 1985.
- [13] B. Flury, *Common Principal Components and Related Multivariate Models*. New York: Wiley, 1988.
- [14] [Online]. Available: <http://sipi.usc.edu/services/database/database.html>
- [15] G. Strang and T. Nguyen, *Wavelets and Filter Banks*. Cambridge, MA: Wellesley-Cambridge, 1996.
- [16] A. Frieze, R. Kannan, and S. Vempala, “Fast Monte-Carlo algorithms for finding low-rank approximations,” in *Proc. 1998 IEEE Symp. Foundations Computer Science*, pp. 370–378.



Ramakrishna Kakarala received the Ph.D. degree in mathematics from the University of California, Irvine, in 1992.

He is a Research Engineer with Agilent Laboratories, Palo Alto, CA. Prior to joining Agilent, he was a Lecturer in electrical and electronic engineering with the University of Auckland, Auckland, New Zealand.



Philip O. Ogunbona (SM'97) received the B.Sc. degree (with first class honors) in electronics and electrical engineering in 1981 from the University of Ife, Nigeria, and the Ph.D. degree in electrical engineering from Imperial College of Science, Technology and Medicine, London, U.K., in 1987.

After post-doctoral work in the Department of Computing, Imperial College of Science, Technology and Medicine, he joined STC Research Labs, Harlow, U.K., in 1989. He joined the Department of Electrical and Computer Engineering, University of

Wollongong, Australia, as a Lecturer in 1990. In 1998, he joined the Visual Information Processing Lab, Motorola Labs, Australian Research Centre, Botany, NSW, where he is now a Principal Research Engineer and leads the Digital Imaging and Signal Processing team. His research interests include image segmentation and analysis, multiresolution techniques, and image and video compression.

Dr. Ogunbona is an active member of the Australian Pattern Recognition Society and is the current Chairman of the New South Wales joint Chapter of the IEEE Communications and Signal Processing Societies.