

A Comparison of Eigendecomposition for Sets of Correlated Images at Different Resolutions

Kishor Saitwal and Anthony A. Maciejewski
Dept. of Electrical and Computer Eng.
Colorado State University
Fort Collins, CO 80523-1373, USA
Email: {Kishor.Saitwal, aam}@colostate.edu

Rodney G. Roberts
Dept. of Electrical and Computer Eng.
Florida A & M - Florida State University
Tallahassee, FL 32310-6046, USA
Email: rroberts@eng.fsu.edu

Abstract—Eigendecomposition is a common technique that is performed on sets of correlated images in a number of computer vision and robotics applications. Unfortunately, the computation of an eigendecomposition can become prohibitively expensive when dealing with very high resolution images. While reducing the resolution of the images will reduce the computational expense, it is not known how this will affect the quality of the resulting eigendecomposition. The work presented here proposes a framework for quantifying the effects of varying the resolution of images on the eigendecomposition that is computed from those images. Preliminary results show that an eigendecomposition from low-resolution images may be nearly as effective in some applications as those from high-resolution images.¹

I. INTRODUCTION

Eigendecomposition-based techniques play an important role in several computer vision applications, e.g., pattern recognition, image compression, image approximation, and object/pose detection. The advantage of these techniques, also referred to as subspace methods, are that they are purely appearance based and that they require few online computations. Various referred to as eigenspace methods, singular value decomposition (SVD) methods, principal component analysis methods, and Karhunen-Loeve transformation methods [1], they have been used extensively in a variety of applications such as face characterization [2] and recognition [3], lip-reading [4], [5], object recognition, pose detection, visual tracking, and inspection [6]-[9]. All of these applications are based on taking advantage of the fact that a set of highly correlated images can be approximately represented by a small set of eigenimages [10]. Once the set of principal eigenimages is determined, online computation using these eigenimages can be performed very efficiently. However, the offline calculation required to determine both the appropriate number of eigenimages as well as the eigenimages themselves can be prohibitively expensive.

¹This work was supported by the National Imagery and Mapping Agency under contract no. NMA201-00-1-1003 and through collaborative participation in the Robotics Consortium sponsored by the U. S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0012. The U. S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

The resolution of the given correlated images, in terms of the number of pixels, is one of the factors that greatly affects the amount of offline calculation required to compute an eigendecomposition. In particular, many common algorithms that compute the complete SVD of a general matrix require on the order of m^2n^2 flops, where m is the total number of pixels in a single image and n is the number of images. Most users of eigendecomposition techniques would like to use as large a resolution as is available for the original images in order to maintain as much information as possible; however, this frequently results in an impractical computational burden. Thus users are typically forced to downsample their images to a lower resolution using a “rule of thumb” or some *ad hoc* criterion to obtain a manageable level of computation. The purpose of the work described here is to quantify the tradeoff between the resolution of correlated images and the “quality” of their resulting eigendecomposition, in terms of measures that are relevant to the user’s motivation for performing an eigendecomposition.

The paper is organized as follows. In Section II, the fundamentals of applying eigendecomposition to related images are reviewed. Section III describes a method for preprocessing the singular values and singular vectors calculated at different resolutions so that their meaningful comparison can be performed. In Section IV, different error measures to compare the SVD at different resolutions are defined and the possibility of using a low-resolution SVD to solve problems associated with high-resolution images is explored. Section V explains the experimental results on a variety of image data sets with concluding remarks given in Section VI.

II. PRELIMINARIES

In this work, a grey-scale image is described by an $h \times h$ square array of pixels with intensity values normalized between 0 and 1. Thus, an image will be represented by a matrix $\mathcal{X} \in [0, 1]^{h \times h}$. Because sets of related images are considered in this paper, it is convenient to represent an image equivalently as a column vector, obtained simply by “row-scanning”, i.e., concatenating the rows to obtain

the *image vector* \mathbf{x} of length $m = h^2$ given by

$$\mathbf{x} = \text{vec}(\mathcal{X}^T).$$

The *image data matrix* of a set of images $\mathcal{X}_1, \dots, \mathcal{X}_n$ is an $m \times n$ matrix, denoted X , and defined as

$$X = [\mathbf{x}_1 \cdots \mathbf{x}_n],$$

where typically $m \gg n$. Only the case where n is fixed is considered in this paper, as opposed to cases where X is constantly updated with new images.

The SVD of X is given by

$$X = U\Sigma V^T, \quad (1)$$

where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ are orthogonal, and $\Sigma = [\Sigma_d \mathbf{0}]^T \in \mathbb{R}^{m \times n}$ where $\Sigma_d = \text{diag}(\sigma_1, \dots, \sigma_n)$ with $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$ and $\mathbf{0}$ is an n by $m - n$ zero matrix. The SVD of X plays a central role in several important imaging applications such as image compression and pattern recognition. The columns of U , denoted $\hat{\mathbf{u}}_i$, $i = 1, \dots, m$, are referred to as the eigenimages of X ; these are the eigenvectors of the correlation matrix of the image vector. The eigenimages provide an orthonormal basis for the columns of X , ordered in terms of importance; the corresponding singular values measure how “aligned” the columns of X are with the associated eigenimage. The components of the i^{th} column of V measure how much each individual image contributes to the i^{th} eigenimage.

III. COMPARING THE SVD AT DIFFERENT RESOLUTIONS

The singular value decomposition of correlated images at different resolutions give different U , Σ and V matrices. These different sets of matrices can be compared against each other to study the effect of different resolutions on the SVD. To perform a meaningful comparison, the singular values and singular vectors of the low-resolution image data matrix must be modified. To distinguish both the resolution and the size of a singular vector, we will use the notation ${}^{(p)}\mathbf{u}_{i(q)}$, where the preceding superscript p denotes the fact that the vector is associated with images of resolution $p \times p$ and the subscript q denotes the actual dimension of the vector \mathbf{u}_i .

A. Singular vectors

Recall that each eigenimage $\hat{\mathbf{u}}_i$ is an m -vector, where m is the number of pixels in an image. As the resolution of the images is varied, the dimension of each eigenimage will also change accordingly. Hence to compare the SVD at different resolutions, the eigenimages at lower resolutions must be enlarged to match in size with those at higher resolution. For our study, this is performed by using bicubic interpolation. Due to the enlargement and interpolation, the resulting eigenimages no longer have

unit norm and they must be renormalized. This is done using

$${}^{(p)}\hat{\mathbf{u}}_{i(q)} = \frac{{}^{(p)}\mathbf{u}_{i(q)}}{\|{}^{(p)}\mathbf{u}_{i(q)}\|}, \quad (2)$$

where ${}^{(p)}\mathbf{u}_{i(q)}$ represents the i^{th} low-resolution, i.e., $p \times p$, eigenimage enlarged to the size of the i^{th} high-resolution, i.e., $q \times q$, eigenimage. (The bicubic interpolation also affects the orthogonality of these singular vectors; however, this effect is small. For the examples used in this study, $\|I - U^T U\|$ is on the order of 10^{-3} for $q = 128$.)

Recall that each right singular vector $\hat{\mathbf{v}}_i$ is an n -vector, where n is the number of images. Thus the change in the resolution of the images does not affect the size of these vectors and the respective right singular vectors can be directly compared with each other, as long as the number of images per object remains fixed for all resolutions.

B. Singular values

The matrix Σ in (1) containing the singular values of X need not be resized before the comparison of singular values at different resolutions. However, due to the lower dimension of the low-resolution images, these values are scaled-down versions of those for the high-resolution images. Hence the singular values at low resolution must be scaled up properly before they can be compared with those at high resolution.

The singular values of the correlated image data set determine the scaling of the associated eigenimages. As indicated earlier in this section, the low-resolution eigenimages are enlarged to the size of high-resolution eigenimages. Hence each low-resolution singular value should be scaled using

$${}^{(p)}\tilde{\sigma}_i = {}^{(p)}\sigma_i \|{}^{(p)}\mathbf{u}_{i(q)}\|, \quad (3)$$

where ${}^{(p)}\sigma_i$ represents the i^{th} $p \times p$ resolution singular value that is to be compared with a higher resolution $q \times q$ singular value.

IV. DIFFERENCE MEASURES FOR SVD'S

Different measures are considered in this section to compare the SVD of the images at different resolutions. The error measures that directly compare the SVD at different resolutions are defined and the effect of these SVD's on common applications, like reconstruction and pose detection of the high-resolution images using their preprocessed SVD's at low resolution, is discussed.

A. Definition of error measures

- **Difference between singular values:** The simplest error measure considered is the difference between the corresponding singular values calculated for a set of correlated images at different resolutions. As described in Section III, the singular values at low

resolution are scaled before calculating this measure, which is given by

$$\Delta \sigma_i = {}^{(q)}\sigma_i - {}^{(p)}\tilde{\sigma}_i, \quad (4)$$

where ${}^{(q)}\sigma_i$ is the i^{th} singular value of the image data set at the higher resolution and ${}^{(p)}\tilde{\sigma}_i$ is the corresponding (scaled) singular value of the image data set at the lower resolution.

- **Angles between singular vectors:** Left and right singular vectors calculated for a set of correlated images at different resolutions can be compared by calculating the angle between the corresponding vectors. Right singular vectors at low resolution do not need any preprocessing, but the left singular vectors at low resolution need to be enlarged and renormalized before their comparison with their high-resolution counterparts. The angle between the two unit vectors can be calculated by performing their dot product. Thus the cosines of the angles between the singular vectors are given by

$$\Delta u_i = {}^{(q)}\hat{\mathbf{u}}_{i(q)}^T {}^{(p)}\hat{\mathbf{u}}_{i(q)} \quad (5)$$

and

$$\Delta v_i = {}^{(q)}\hat{\mathbf{v}}_i^T {}^{(p)}\hat{\mathbf{v}}_i, \quad (6)$$

where ${}^{(q)}\hat{\mathbf{u}}_{i(q)}$ and ${}^{(q)}\hat{\mathbf{v}}_i$ are the i^{th} left and right singular vectors, respectively, of the image data set at the higher resolution, and ${}^{(p)}\hat{\mathbf{v}}_i$ is the i^{th} right singular vector of the image data set at the lower resolution.

- **Rotation of subspaces:** The previous error measure calculates the angles between the individual singular vectors. However the i^{th} singular vector at low resolution may not be aligned with the i^{th} singular vector at high resolution. On the other hand, the subspaces containing the first k vectors may span the same vector space. Hence another error measure is used to calculate the rotation between these subspaces. The possibility that the data matrix $B \in \mathbb{R}^{m \times k}$ can be rotated into the data matrix $A \in \mathbb{R}^{m \times k}$ is explored [11] by solving the problem

$$r = \min_Q \|A - BQ\|_F, \quad (7)$$

where $\|\cdot\|_F$ represents the Frobenius norm and $Q \in \mathbb{R}^{k \times k}$ is an orthogonal matrix. The Q that minimizes $\|A - BQ\|_F$ can be calculated as follows:

- Form the matrix $C = B^T A$,
- Compute the SVD of C , i.e., $C = U_c \Sigma_c V_c^T$,
- Find the orthogonal matrix $Q = U_c V_c^T$,

where U_c and V_c are the matrices containing the left and right singular vectors of C , respectively, while

Σ_c is a diagonal matrix containing the singular values of C in descending order. The smaller the norm r , the closer A and B are to representing the same subspace. In our experiments, the subspaces A and B are given by

$$\begin{aligned} A &= {}^{(q)}U_{k(q)}, \\ B &= {}^{(p)}U_{k(q)}, \end{aligned}$$

where U_k is the matrix containing the first k eigenimages as its columns.

- **Angles between subspaces:** Two sets of eigenimages that represent two subspaces, both having dimension k , share k different *principal angles* [11] with the k^{th} *principal angle* giving the largest angle between the subspaces. Because the matrices representing the subspaces in our experiments are nearly orthogonal, this measure is closely related to the previous measure concerning the rotation of subspaces. In particular, the Σ_c matrix containing the singular values of C gives the *principal angles* between the subspaces, i.e.,

$$\text{diag}(\cos(\theta_1), \dots, \cos(\theta_k)) = \Sigma_c, \quad (8)$$

where θ_i is the i^{th} *principal angle*. If the maximum angle θ_k is close to 0, then the two subspaces are considered to be closely aligned with each other.

B. Effect on common applications

The error measures explained earlier in this section give a direct comparison between the SVD at different resolutions. Another way of comparing them is to apply them to real problems. Two such common applications, namely reconstruction of images and pose detection, are considered in this subsection.

- **Reconstruction of images:** The eigenimages calculated using the SVD of correlated images are sometimes used to form an approximation $\tilde{\mathbf{x}}^{(k)}$ of the original image vector \mathbf{x} that is given by

$$\tilde{\mathbf{x}}^{(k)} = \sum_{i=1}^k (\hat{\mathbf{u}}_i^T \mathbf{x}) \hat{\mathbf{u}}_i, \quad (9)$$

where k is the dimension of the eigenspace. The accuracy of this approximation can be evaluated using

$$\rho_k = \frac{\sqrt{\sum_{i=1}^k \|\hat{\mathbf{u}}_i^T \mathbf{x}\|^2}}{\|\mathbf{x}\|_F}, \quad (10)$$

where ρ_k is referred to as the energy recovery ratio [12]. The eigenimages at all resolutions (with proper preprocessing) are used to approximate the high-resolution images, with the results of using low-resolution eigenimages being compared to those when using the high-resolution eigenimages.

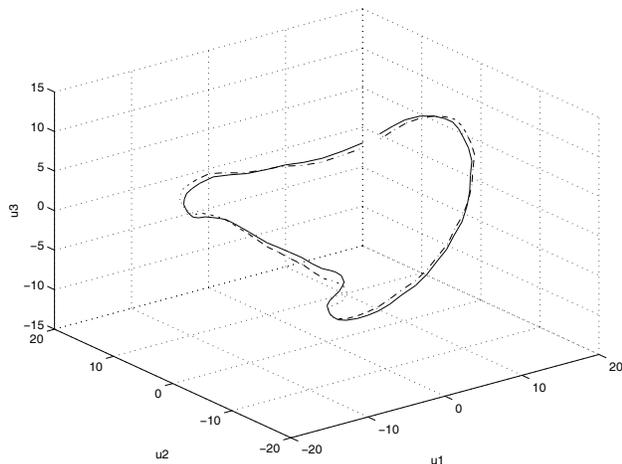


Fig. 1. This figure shows two manifolds consisting of the normalized training images (after subtracting the normalized average image from each of them) projected onto the subspace spanned by the first three eigenimages. The solid manifold represents the projection of the high-resolution images onto their eigenspace, while the dash-dotted manifold represents the projection of the same high-resolution images onto the enlarged and renormalized eigenimages calculated at low resolution.

- **Pose detection:** The standard application of eigen-decomposition to solve the pose-detection problem, when the training and test images are considered to have the same resolution, is explained in [12]. This procedure is modified in this paper so that the pose of the object in a high-resolution test image can be detected with the eigenspace representing all its possible orientations in low-resolution training images.

In any pose-detection problem using the SVD, a reduced-order representation of the object's orientation change is obtained by projecting the training images onto the space spanned by the dominant eigenimages, interpolating to obtain a manifold. (Fig. 1 shows two such manifolds, where the high-resolution images are projected onto two different subspaces spanning the first three eigenimages.)

The estimated pose of the object, $\hat{\phi}$, in the test image is given by

$$\hat{\phi} = \arg \min_{\phi} \frac{\|\mathbf{t}_{\phi} - \mathbf{p}\|}{\|\mathbf{t}_{\phi}\|}, \quad (11)$$

where \mathbf{t}_{ϕ} represents a point on the manifold of training images at orientation ϕ , \mathbf{p} represents the eigenspace projection of the test image and $\frac{\|\mathbf{t}_{\phi} - \mathbf{p}\|}{\|\mathbf{t}_{\phi}\|}$ is the normalized distance from \mathbf{p} to \mathbf{t}_{ϕ} .

A measure of the quality of this eigenspace for solving the pose-detection problem is given by $\sum_{i=1}^{n_t} |\phi_i - \hat{\phi}_i|$, where ϕ_i represents the correct pose of the i^{th} orientation of an object and n_t denotes the total number of test images.



Fig. 2. The 20 objects used for the experiments in this section.

The results analyzing the above-mentioned error measures, reconstruction procedures and pose detection procedures are presented in the next section.

V. EXPERIMENTAL RESULTS

Several experiments were performed and the different error measures (refer to Section IV) were used to compare the SVD of the correlated images at different resolutions. The reconstruction and pose detection procedures were also carried out for the high-resolution images using their preprocessed SVD at low resolution. A variety of objects (see Fig. 2) were used to perform these experiments. For each object, 360 training images were obtained by rotating the object by 1 degree between successive images. All training images were of size 128×128 pixels with 8 bits used to represent intensity. Two new image data sets were generated by reducing the original image resolution to 64×64 and 32×32 by using bicubic interpolation. Three different sets of SVD's were then calculated for these image data sets. The resultant singular values and singular vectors for the low-resolution images were preprocessed as described in Section III before they were compared with those for the high-resolution images.

The first set of tests was designed to evaluate the different error measures that compare the SVD at different resolutions. Fig. 3 shows different plots for these error measures calculated for object 2 (second object in the first row of objects in Fig. 2). These plots show the general behavior for most of the objects when comparing their SVD's. The plots of the difference between the singular values in the first row show a very slow and steady increase in the difference between the singular values at 128×128 and those at 64×64 . Given that the first few singular values are so large, the relative error is almost

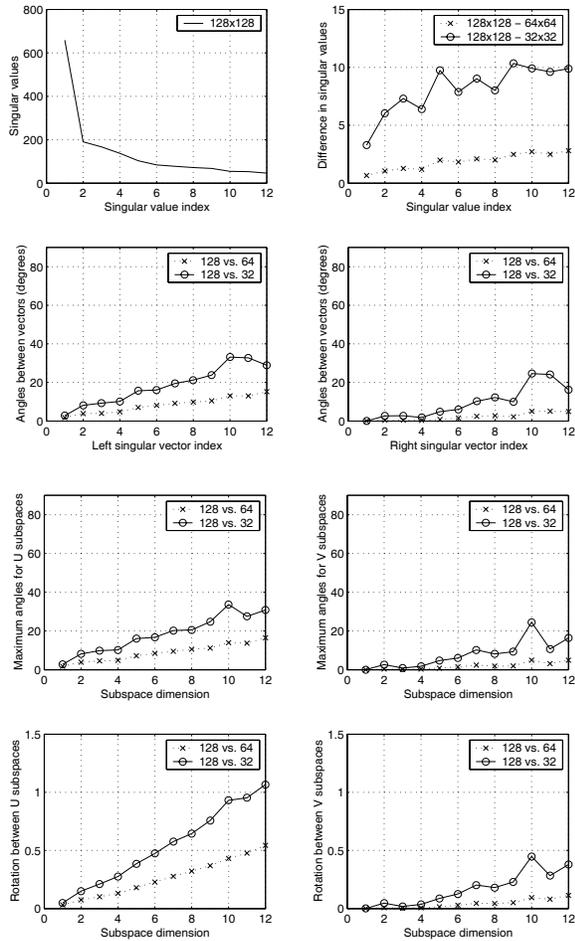


Fig. 3. This figure shows different plots of the error measures calculated for object 2 that represents the general behavior for most of the objects when the SVD at different resolutions were compared. The first row shows the plot of the singular values for 128×128 resolution images and its difference with the scaled singular values for 64×64 and 32×32 resolution images. The second row shows the plots for the angles between the respective left and right singular vectors (after preprocessing the left singular vectors at low resolution). The third row shows the maximum principal angles between the respective singular vector subspaces, when the subspace dimension was varied from 1 to 12, while the last row shows the minimum norm of the rotation required to rotate a low-resolution eigenspace into its corresponding high-resolution eigenspace, where the subspace dimension was varied from 1 to 12. All angles are plotted in degrees.

negligible. In particular, the relative error between these values varies from 0.1% ($k = 1$) to 6.06% ($k = 12$) indicating a good approximation of 128×128 singular values using the scaled 64×64 singular values. For 32×32 , the relative error varies from 0.5% ($k = 1$) to 21.37% ($k = 12$); however, the relative error in the first few singular values is again small enough for the low-resolution eigenspace to be comparable to the high-resolution eigenspace for the first few eigenimages.

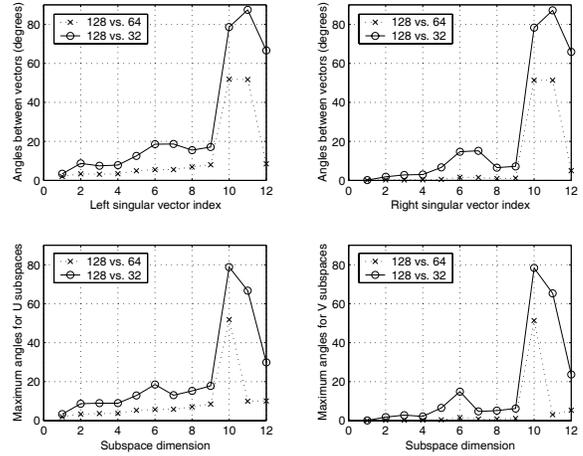


Fig. 4. This figure shows plots for the angles between singular vectors and for the maximum principal angles between the respective singular vector subspaces calculated for object 6. This example shows a case where the plots were different than the general trend of the plots for the objects as shown for object 2 shown in Fig. 3.

The angles between the respective singular vectors (second row in Fig. 3) for 64×64 are small for both the left and right singular vectors, while these angles for 32×32 are almost twice as big as those for 64×64 . There is a steady increase in the angles when the singular vector index increases. However, the large increase in the angle is not necessarily meaningful because the index is ill-conditioned when the singular values are nearly equal.

The maximum principal angles between the subspaces containing the singular vectors (third row in Fig. 3) show that the maximum principal angle, when the subspace dimension is twelve, is very close to the angle between the 12^{th} singular vectors. This indicates that these singular vectors span the same vector space when the appropriate dimension is used. The rotation indices between these subspaces (fourth row in Fig. 3) exhibit similar behavior to that of the maximum principal angles and show that the rotation of the 64×64 vector subspace into the 128×128 vector subspace is much closer than the rotation of the subspace associated with the 32×32 resolution images.

Fig. 4 gives an example of a set of subspace difference plots that varies from the general trend of such plots. These plots are for object 6 (first object in the second row of Fig. 2). The first row of plots in Fig. 4 shows drastic variations in the angles between the singular vectors after the 9^{th} vector. The angles for the 10^{th} and 11^{th} vectors for 64×64 jump to high values before coming down for the 12^{th} vectors. This happens because the 10^{th} and 11^{th} vectors for the 64×64 case and the 128×128 case are swapped. This can be explained by looking at the maximum principal angles plots in the second row of Fig. 4. When the subspace dimension is ten, the maximum principal angle is very large. However, after including the

TABLE I
RELATIVE RECONSTRUCTION AND POSE DETECTION RATIOS

Object	Reconstruction ratio		Pose detection ratio	
	64×64	32×32	64×64	32×32
1	0.9990	0.9967	0.9964	1.0000
2	0.9982	0.9938	1.0000	1.0000
3	0.9979	0.9921	1.0000	1.0035
4	0.9989	0.9955	1.0070	1.0000
5	0.9971	0.9916	1.0115	0.9885
6	0.9987	0.9950	1.0000	1.0000
7	0.9930	0.9861	1.1235	1.1070
8	0.9967	0.9915	0.9925	0.9887
9	0.9959	0.9891	0.9965	0.9965
10	0.9991	0.9973	1.0190	1.0190
11	0.9991	0.9966	1.0035	1.0035
12	0.9995	0.9973	0.9965	1.0000
13	0.9992	0.9962	1.0000	1.0000
14	0.9985	0.9941	1.0000	0.9965
15	0.9988	0.9958	1.0035	1.0035
16	0.9992	0.9956	1.0000	1.0000
17	0.9990	0.9958	1.0000	1.0000
18	0.9991	0.9953	1.0000	1.0000
19	0.9992	0.9970	1.0000	1.0000
20	0.9981	0.9925	1.0000	1.0000
Mean	0.9982	0.9942	1.0075	1.0053
Maximum	0.9995	0.9973	1.1235	1.1070
Minimum	0.9930	0.9861	0.9925	0.9885

11^{th} vector, it drops down indicating that the first eleven singular vectors for 64×64 essentially span the same vector space as those for 128×128 . On the other hand, the 10^{th} , 11^{th} , and 12^{th} vectors for the 32×32 resolution images and those for the 128×128 resolution images are swapped. However, once again the maximum principal angle decreases when the subspaces are compared with twelve singular vectors. Thus, the singular vectors at different resolutions span essentially the same vector space as long as the dimension of the subspace is selected appropriately.

This swapping of the low-resolution singular vectors is not a significant issue when these vectors are used for image reconstruction or pose detection. This is due to the fact that the associated singular values are nearly equal.

Reconstruction of the 128×128 resolution images was carried out using eigenimages at all resolutions. The average energy recovery ratio using an eigenspace dimension of twelve, i.e., ρ_{12} , was calculated for the image data matrices of all 20 objects. The first half of Table I gives the ratio of ρ_{12} using preprocessed low-resolution eigenimages to ρ_{12} using 128×128 resolution eigenimages. These ratios show that the reconstructions using low-resolution eigenimages are almost as good as those using high-resolution eigenimages. Fig. 5 provides a visual example of the relative quality of these different resolution reconstructions for object 2, which was selected because it has the reconstruction ratio that is closest to the mean value. The comparison plot in Fig. 5(a) shows

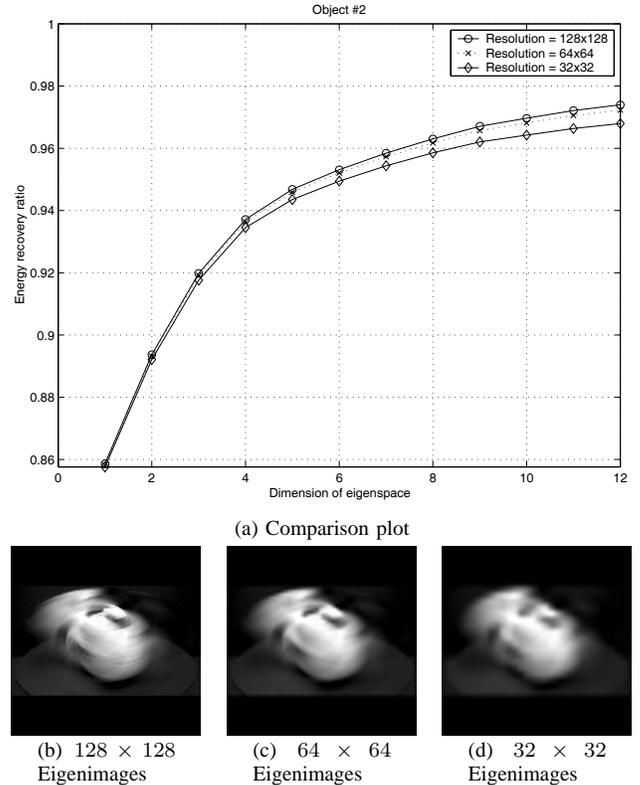


Fig. 5. Part (a) of this figure shows the energy recovery ratio plot for object 2. Parts (b), (c), and (d) show the reconstruction of a 128×128 image for object 2, image 1, using the first 12 eigenimages at three different resolution, namely 128×128 , 64×64 , and 32×32 . The eigenimages at low resolutions are enlarged and renormalized before using them for the reconstruction. This reconstruction represents the average case from Table I.

that the plot for the energy recovery ratio using all three resolution eigenspaces are very close together and the actual reconstruction of a single image (Fig. 5(b), (c), and (d)) indicates that all three reconstructions are visually very close to the original image of this object shown in Fig. 2.

The pose detection application was also evaluated for all 360 orientations of all 20 objects. The preprocessed 64×64 and 32×32 resolution eigenimages were used for pose detection of the 128×128 resolution images. A set of 72 images (with successive orientations 5 degrees apart) was used as training images for generating the manifold shown in Fig. 1 with the remaining images used as test images. Successful pose detection was defined as identifying the training image whose pose was closest to the pose of the test image. The total number of successful pose detections over all test images and all objects is thus a measure of how effective a given image resolution is for solving the pose-detection problem. The ratio of the number of successful pose detections at the lower resolutions to that of the original 128×128 resolution is

given in the second half of Table I. This table shows that eigenimages at 64×64 and at 32×32 (with mean ratios of 1.0075 and 1.0053, respectively) actually perform better than those at 128×128 . This indicates that there is little value in performing the more computationally expensive eigendecomposition at the higher resolution in this case.

VI. CONCLUSION

This paper has presented a framework for quantifying the tradeoff associated with performing eigendecomposition on correlated images at lower resolutions in order to mediate the high computational expense of performing these calculations at high resolutions. Although our initial results are preliminary, it appears that the quantitative difference between the SVD of image data matrices of quite different resolutions is surprisingly small. It is particularly encouraging that the resulting eigenspaces are very similar, because this is the main factor that determines the accuracy of such common applications as image reconstruction and object/pose detection.²

VII. REFERENCES

- [1] K. Fukunaga, *Introduction to Statistical Pattern Recognition*. London: Academic Press, 1990.
- [2] L. Sirovich and M. Kirby, "Low-dimensional procedure for the characterization of human faces," *Journal of Optical Society of America*, vol. 4, no. 3, pp. 519–524, March 1987.
- [3] M. Turk and A. Pentland, "Eigenfaces for recognition," *Journal of Cognitive Neuroscience*, vol. 3, no. 1, pp. 71–86, March 1991.
- [4] H. Murase and R. Sakai, "Moving object recognition in eigenspace representation: Gait analysis and lip reading," *Pattern Recognition Letters*, vol. 17, no. 2, pp. 155–162, Feb. 1996.
- [5] G. Chiou and J. N. Hwang, "Lipreading from color video," *IEEE Transactions on Image Processing*, vol. 6, no. 8, pp. 1192–1195, Aug. 1997.
- [6] H. Murase and S. K. Nayar, "Illumination planning for object recognition using parametric eigenspaces," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, no. 12, pp. 1219–1227, Dec. 1994.
- [7] H. Murase and S. K. Nayar, "Visual learning and recognition of 3-D objects from appearance," *International Journal on Computer Vision*, vol. 14, no. 1, pp. 5–24, Jan. 1995.
- [8] H. Murase and S. K. Nayar, "Detection of 3D objects in cluttered scenes using hierarchical eigenspace," *Pattern Recognition Letters*, vol. 18, no. 4, pp. 375–384, April 1997.
- [9] S. K. Nayar, S. A. Nene, and H. Murase, "Subspace method for robot vision," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 5, pp. 750–758, Oct. 1996.
- [10] C. Y. Chang, A. A. Maciejewski, and V. Balakrishnan, "Fast eigenspace decomposition of correlated images," *IEEE Transactions on Image Processing*, vol. 9, no. 11, pp. 1937–1949, Nov. 2000.
- [11] G. H. Golub and C. F. V. Loan, *Matrix Computations*. Baltimore, Maryland: Johns Hopkins, 1996.
- [12] C. Y. Chang, A. A. Maciejewski, V. Balakrishnan, and R. G. Roberts, "Eigendecomposition-based pose detection in the presence of occlusion," in *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 29 - Nov. 03 2001, pp. 569–576.

²The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U. S. Government.