

An Illustration of Eigenspace Decomposition for Illumination Invariant Pose Estimation

Randy C. Hoover*, Anthony A. Maciejewski*, Rodney G. Roberts⁺, and Ryan P. Hoppal*

*Dept. of Electrical and Computer Eng.
Colorado State University

Fort Collins, CO 80523-1373, USA

Email: {hoover, aam, rphoppal}@colostate.edu

⁺Dept. of Electrical and Computer Eng.

Florida A & M - Florida State University

Tallahassee, FL 32310-6046, USA

Email: rroberts@eng.fsu.edu

Abstract—Determining the pose of a three-dimensional object under unknown lighting conditions is a challenging problem. Eigenspace methods represent one computationally efficient method for doing illumination invariant pose estimation, and have been applied in a variety of application domains. Unfortunately, determining the appropriate eigenspace dimension, as well as the eigenspace itself, is computationally prohibitive for real-world applications. This paper presents a method to reduce this expense by using results from spectral theory. In particular, this paper shows that a set of images of an object under a wide range of illumination conditions and a fixed pose can be significantly reduced by projecting this data on to a few low-frequency spherical harmonics, producing a set of “harmonic images”. It is then shown that the dimensionality of the set of harmonic images can be further reduced by utilizing the Fast Fourier Transform. An eigendecomposition is then applied in the spectral domain thus relieving the computational burden. Experimental results are presented to compare the proposed algorithm to the true eigendecomposition, as well as assess the computational savings.

Index Terms—Eigenspace Decomposition, pose estimation, illumination variation, spherical harmonics, Fourier transforms.

I. INTRODUCTION

Over the last several decades, classification of three-dimensional (3-D) objects from two-dimensional (2-D) images under a wide range of illumination conditions has become an important issue in the computer vision community. Specific examples include face recognition, target tracking, automated assembly and inspection, robot localization, and human robot interaction. Subspace methods represent one computationally efficient approach for dealing with this class of problems. Subspace methods, also referred to as eigenspace methods, principal component analysis, or the Karhunen-Loeve transformation [1], [2], have been applied in a variety of application domains. All of these applications are based on the fact that a set of highly correlated images can be approximately represented by a small set of eigenimages [3]–[5]. Once the principal eigenimages of an image data set have been determined, using these eigenimages is computationally efficient for the on-line classification of 3-D objects.

Unfortunately, the off-line calculation for determining the appropriate subspace dimension, as well as the principal eigenimages themselves is computationally expensive. This drawback has been addressed using several different approaches based on either iterative power methods, conjugate gradient

algorithms, or eigenspace updating [6]–[8]. A fundamentally different approach was proposed by Chang *et al.* [9] where the authors show that the Fourier transform can be used to approximate the desired subspace dimension, as well as the principal eigenimages if the image data set is correlated in one-dimension. This result has recently been extended to correlation in two and three-dimensional orientations under ambient lighting conditions by utilizing spherical harmonics and Wigner- D functions in place of the Fourier transform [10]–[12].

It has been shown empirically that the set of images of a convex largely diffuse object under a wide range of illumination conditions and a fixed pose approximately lies within a nine-dimensional (9-D) linear subspace [13], [14]. This result was analytically verified in [15], [16] by expanding the Lambertian kernel in a series of spherical harmonics. The principle eigenmodes of the set of images can then be calculated by evaluating a set of spherical harmonics at the surface normals of the object and applying standard ray-tracing techniques. In the current work, we present an efficient method to compute the principal eigenimages when the set of images contains variation in both illumination and pose of the object. We treat each image of the object under a different illumination condition and fixed pose as a sample on the surface of the 2-sphere (S^2) and utilize spherical harmonics to reduce the dimensionality. Rather than attempting to evaluate the spherical harmonics at the surface normals however, we expand the image data itself in a series of spherical harmonics, resulting in a set of “harmonic images”. We show that for most objects, a significant amount of the energy in a set of images of an object under a wide range of illumination conditions and a fixed pose is captured by the first few low-frequency spherical harmonics. Furthermore, we place no restriction on the convexity or reflectance model of the object. We then show that the eigendecomposition of the entire data set (variation in illumination and pose) can be efficiently computed by applying Chang’s eigendecomposition algorithm to the resulting set of harmonic images.

The remainder of this paper is organized as follows. In Section II, the fundamentals needed to apply an eigendecomposition to a related image data set is explained, much of which is discussed in [9]. Section II also gives a brief

overview of related previous work. In Section III, we give a brief introduction to spherical harmonics and discuss the problem of reducing the dimensionality of the image data in the illumination dimension. In Section IV a brief overview of Chang’s eigendecomposition algorithm is provided. We then use this to develop a computationally efficient algorithm for estimating the principle eigenimages of a data set due to variation in illumination and pose in Section V. Section VI provides some experimental results with conclusions and future work discussed in Section VII.

II. PRELIMINARIES

A. Mathematical Description

In this work, a gray-scale image is described by an $h \times v$ array of square pixels with intensity values normalized between 0 and 1. Thus an image is represented by a matrix $\mathcal{X} \in [0, 1]^{h \times v}$. The image \mathcal{X} is then “row-scanned” to obtain the *image vector* \mathbf{f} of length $m = hv$, i.e., $\mathbf{f} = \text{vec}(\mathcal{X}^T) \in \mathbb{R}^{hv \times 1}$. 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 $X = [\mathbf{f}_1, \dots, \mathbf{f}_n]$, where typically $m > n$ with fixed n [9].

In this paper, we consider sets of images of a known (rigid) object under different pose and illumination conditions. The objects are placed at the center of an *illumination sphere*, while the camera is moved to discrete locations on a line of constant co-latitude. At each of the discrete locations, images of the object are captured under a dense but finite set of illumination conditions generated from a single point light source at infinity arriving from a distinct illumination direction. Under these assumptions, the image vector can be parameterized by $\mathbf{f} = \mathbf{f}(\boldsymbol{\xi}_i, r)$ where $r \in \{0, \dots, a-1\}$ represents the r^{th} pose of the object and $\boldsymbol{\xi}_i, i \in \{0, \dots, b-1\}$, is the unit vector pointing at the angle of co-latitude $\beta_i \in (0, \pi)$ measured down from the upper pole, and the angle of longitude $\alpha_i \in [0, 2\pi)$, which is the parameterization of the i^{th} direction of the point light source at each pose. An example of this procedure is shown in Fig. 1 where the illumination directions are determined using the Hierarchical Equal Area isoLatitude Pixelization (HEALPix) sampling pattern [17]. Using this notation, the image data matrix is constructed as

$$X = [\mathbf{f}(\boldsymbol{\xi}_0, 0), \mathbf{f}(\boldsymbol{\xi}_1, 0), \dots, \mathbf{f}(\boldsymbol{\xi}_{b-1}, 0), \\ \mathbf{f}(\boldsymbol{\xi}_0, 1), \mathbf{f}(\boldsymbol{\xi}_1, 1), \dots, \mathbf{f}(\boldsymbol{\xi}_{b-1}, 1), \dots, \\ \mathbf{f}(\boldsymbol{\xi}_0, a-1), \mathbf{f}(\boldsymbol{\xi}_1, a-1), \dots, \mathbf{f}(\boldsymbol{\xi}_{b-1}, a-1)], \quad (1)$$

where the first b columns of X correspond to a single pose of the object under b different illumination conditions. The average image vector is then subtracted from the image data matrix X to generate the zero mean image data matrix \hat{X} , which has the interpretation of an “unbiased” image data matrix.

One approach to computing the eigenimages of \hat{X} is to utilize the singular value decomposition (SVD). The *thin* SVD of \hat{X} is given by $\hat{X} = \hat{U}\hat{\Sigma}\hat{V}^T$ where $\hat{U} \in \mathbb{R}^{m \times n}$ is right orthogonal, i.e., $\hat{U}^T\hat{U} = I$, $\hat{V} \in \mathbb{R}^{n \times n}$ is orthogonal, and $\hat{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_n)$ with $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$. The columns of \hat{U} , denoted $\hat{\mathbf{u}}_i, i = 1, \dots, n$, are the left singular vectors or

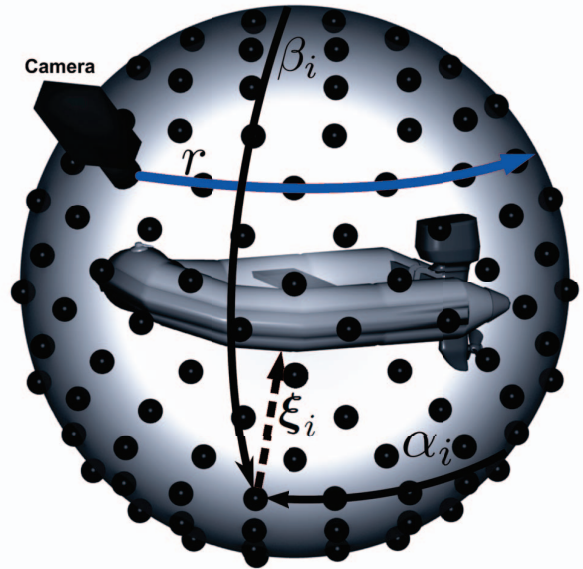


Fig. 1. A graphical depiction of the proposed method of acquiring images from a dense set of illumination conditions at each pose. The object is placed at the center of the illumination sphere with the camera moving along a line of constant co-latitude. The black dots on the sphere represent different illumination conditions. As the camera moves along the line of co-latitude, an image of the object is captured under each of the distinct illumination conditions.

eigenimages of \hat{X} , while the columns of \hat{V} , denoted $\hat{\mathbf{v}}_i, i = 1, \dots, n$ are referred to as the right singular vectors of \hat{X} . The left singular vectors of \hat{X} are computed as the eigenvectors of the covariance matrix $\hat{X}\hat{X}^T$. In practice, the left singular vectors $\hat{\mathbf{u}}_i$ are not known or computed exactly, and instead estimates $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k$, denoted \tilde{U}_k that form a k -dimensional basis, are used. The accuracy of a practical implementation of subspace methods then depends on three factors: the properties of \hat{X} , the dimension k , and the quality of the estimates \tilde{U}_k . The measure we will use for quantifying the quality of these estimates is referred to as the energy recovery ratio and defined as [9], [18]

$$\rho(\hat{X}, \tilde{U}_k) = \frac{\sum_{i=1}^k \|\tilde{\mathbf{u}}_i^T \hat{X}\|^2}{\|\hat{X}\|_F^2}. \quad (2)$$

Note that if the $\tilde{\mathbf{u}}_i$ are orthonormal, $\rho \leq 1$.

B. Related Work

The principal calculation required with subspace methods is the precomputation of estimates of the left singular vectors \tilde{U}_k of the $m \times n$ matrix \hat{X} . This is a very computationally expensive operation when m and n are large. This drawback has been addressed using several different approaches based on either iterative power methods, conjugate gradient algorithms, or eigenspace updating [6]–[8]. A fundamentally different

approach was proposed by Chang *et al.* [9] where the authors showed that if the image data matrix was correlated in one-dimension, then the right singular vectors are approximately spanned by a few low frequency Fourier harmonics. As a result, the Fourier transform of the image data matrix along the temporal dimension may be used to estimate the desired subspace dimension k , as well as the principal eigenimages \hat{U}_k (refer to Section IV). When variation in pose under ambient lighting conditions is considered, the results of [9] can be extended to correlation in higher dimensions by replacing Fourier harmonics with spherical harmonics and Wigner- D functions [10]–[12]. In the current work, we are interested in computing the eigenimages of a set of images containing variation on both pose as well as illumination conditions. To this end, we first show that the set of images of an object under different illumination conditions and a fixed pose can be represented by a small set of harmonic images using the spherical harmonic transform (SHT). An introduction to spherical harmonics and the SHT along with dimensionality reduction in the illumination dimension is presented in the next section.

III. SPHERICAL HARMONICS

A. Introduction

Spherical harmonics, and the SHT, have been applied to a variety of application domains over the last several decades. Specific examples include solving PDE's for weather and climate models [19], geophysics [20], [21], quantum mechanics [22], 3-D model retrieval [23], as well several applications in computer vision [10]–[12], [15], [16], [24]–[27].

B. Discrete Spherical Harmonic Transform

The development of a fast discrete SHT has been an active research area over the last decade [17], [19], [28], [29]. Analogous to the Fourier basis for functions defined on the line or circle, under proper normalization, spherical harmonics satisfy

$$\int_0^\pi \int_0^{2\pi} Y_{p_1, q_1}(Y_{p_2, q_2})^* d\alpha \sin(\beta) d\beta = \delta_{p_1 p_2} \delta_{q_1 q_2}, \quad (3)$$

where the superscript $*$ is the complex conjugate, i.e., they provide an orthonormal basis for functions defined on S^2 . Therefore, a real valued band-limited function $f(\xi_i, r)$ whose domain is $L^2(S^2)$ may be expanded in a series of spherical harmonics as

$$f(\xi_i, r) = \sum_{p=0}^{p_{\max}} \sum_{|q| \leq p} f_{p,q}^r Y_{p,q}(\xi_i), \quad (4)$$

where $L^2(\cdot)$ is the Hilbert space of square integrable functions, $f(\xi_i) \in [0, 1]$ is a single pixel of the image data vector $\mathbf{f}(\xi_i, r)$, $Y_{p,q}(\xi_i)$ is the spherical harmonic of degree p and order q , $f_{p,q}^r$ is the corresponding harmonic coefficient at pose r , and l_{\max} is related to the bandwidth of the function. Recall that $\mathbf{f}(\xi_i, r)$ is an image of the object under a illumination

condition ξ_i and pose r . The harmonic coefficients $f_{p,q}^r$ are calculated using

$$f_{p,q}^r = \frac{4\pi}{b} \sum_{i=0}^{b-1} w_i f(\xi_i, r) Y_{p,q}(\xi_i), \quad (5)$$

where w_i is a quadrature weight and $Y_{p,q}(\xi_i)$ is the real-valued spherical harmonic defined by

$$Y_{p,q}(\xi_i) = \begin{cases} \sqrt{2} \kappa_{p,q} \cos(q\alpha_i) P_{p,q}(x) & \text{if } q > 0 \\ \sqrt{2} \kappa_{p,q} \sin(|q|\alpha_i) P_{p,|q|}(x) & \text{if } q < 0 \\ \kappa_{p,0} P_{p,0}(x) & \text{if } q = 0 \end{cases} \quad (6)$$

In (6), $P_{p,q}(x)$ is the associated Legendre polynomial of degree p and order q , $x = \cos(\beta_i)$, and $\kappa_{p,q}$ is a normalization factor.

C. Dimensionality Reduction in the Illumination Dimension

As mentioned in Section I, we first reduce the dimensionality of the image data due to variation in illumination conditions under a fixed pose. For this development, notice that $f_{p,q}^r$ is the harmonic coefficient for a single pixel in the set of images due to a change in illumination conditions at the r^{th} pose. If all m pixels of the set of images due to a change in illumination conditions at the r^{th} pose are expanded using (5), then $\mathbf{f}_{p,q}^r \in \mathbb{R}^{m \times 1}$ represents a ‘‘harmonic image’’ of degree p and order q at pose r .

As discussed in [13]–[16], the set of images of a (largely diffuse) convex object under variations in illumination conditions and fixed pose approximately lie within a 9-D linear subspace. While we place no restrictions on the convexity or reflectance model of the objects in this work, we show that most of the energy of a set of images of an arbitrary object under variations in illumination conditions and fixed pose is captured by a small set of spherical harmonics. Therefore, by truncating the harmonic transform in (5) to $p = 2$, we obtain a set of nine harmonic images. These nine harmonic images are capable of recovering a significant amount of the total energy in the set of images when only variations in the illumination conditions are considered.

To illustrate this, CAD generated ray-traced images of 20 different objects were captured from 90 different poses with 48 different light source locations at each pose. An example image of each object is shown in Fig. 2 (the CAD models were provided by [30]). For each of the 90 different poses, the harmonic transform in (5) was used to reduce the dimensionality of the data from 48 images to 9, 16, 25, and 36 harmonic images, i.e., $p = 2, 3, 4,$ and 5 respectively. The harmonic images were then orthonormalized and the energy recovery ratio defined in (2) was used to compute how much energy each of the four subspaces are capable of recovering at each pose. The minimum amount of energy recovered across all 90 poses for all 20 objects is depicted in Fig. 3. Notice that with the exception of objects 17, 18, and 20, over 95% of the energy is recovered by the 9-dimensional subspace for all 90 poses. Furthermore, adding additional harmonic images does not significantly increase the amount of energy recovered.



Fig. 2. Ray-traced CAD models courtesy of Kator Legaz [30]. Each object is sampled as discussed in Section II at a resolution of 128×128 . The objects are ordered from left to right, then top to bottom.

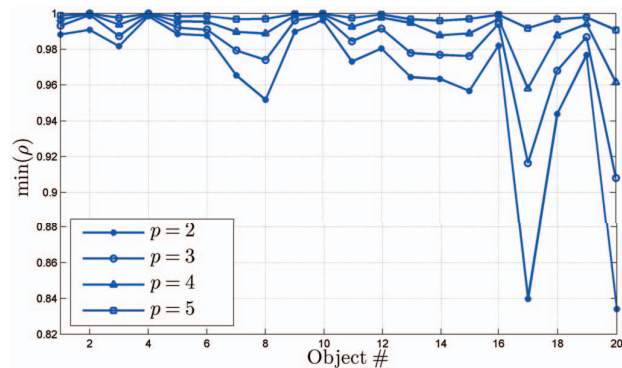


Fig. 3. The minimum amount of energy recovered by 9, 16, 25, and 36 orthonormalized harmonic images for each object under all 90 test poses. With the exception of objects 17, 18, and 20, over 95% of the energy is recovered by the 9-D linear subspace for all 90 poses.

IV. AN OVERVIEW OF CHANG’S ALGORITHM

A. Introduction

This section gives an overview of one of the fastest known algorithms for estimating the first k eigenimages of an image data set correlated in one-dimension, the details of which can be found in [9]. Consider capturing images of the object as discussed in Section II, under a single illumination condition at each pose. Capturing images in this manner results in a one-dimensionally correlated image data matrix X . Consider the special case where X is constructed such that the image \mathbf{f}_{i+1} is obtained from \mathbf{f}_i by a planar rotation of $2\pi/n$, then the correlation matrix $X^T X$ is a *circulant* matrix with circularly symmetric rows. For this special case, the eigendecomposition of $X^T X$ is given by the Discrete Fourier Transform (DFT), and an unordered SVD of X can be obtained for this special case by letting $V = H$, where H is the “real” Fourier matrix.

In other words, the right singular vectors of X in this case are given by pure sinusoids of frequencies that are multiples of $2\pi/n$ radians. The left singular vectors (eigenimages) can then be obtained by computing $XH = U\Sigma$ which can be computed efficiently using Fast Fourier Transform (FFT) techniques [9].

B. Chang’s Eigendecomposition Algorithm

While the above analysis does not hold for arbitrary image data sets, it has been shown in [9] that the analytical expressions for planar rotations serve as a good approximation for the eigendecomposition of image data sets correlated in one-dimension. In general, for image sequences correlated in one-dimension, the following two properties can be observed [9], [18].

- 1) The right singular vectors of X are well-approximated by sinusoids of frequencies that are multiples of $2\pi/n$ radians, and the power spectra of the right singular vectors consist of a narrow band around the corresponding dominant harmonics.
- 2) The dominant frequencies of the power spectra of the (ordered) singular vectors increase approximately linearly with their index.

These two properties indicate that the right singular vectors of an image data set correlated in one-dimension are approximately spanned by the first few low frequency harmonics. Therefore, by projecting the image data set X onto these first few low frequency harmonics and computing the eigendecomposition in the spectral domain, the computational expense associated with computing the SVD can be significantly reduced.

Chang’s algorithm makes use of the above two properties to estimate the subspace dimension k as well as the principal eigenimages \tilde{U}_k of the image data matrix X . It was shown in [9] that if the power spectra of the first j right singular vectors of X are restricted to the band $[0, 2\pi j/n]$, then for $\rho(X^T, H_j) \geq \mu$, the quantity $\rho(X, \tilde{U}_k)$ will exceed μ for some $k \leq j$, where H_j is the matrix containing the first j columns of H , and μ is a user specified value. This inequality shows that the energy recovery ratio as computed using the first few low frequency harmonics of H provides a lower bound on the energy recovery ratio as computed using the estimated eigenimages. Furthermore, this bound is shown to be extremely tight in most cases [9], with a tight upper bound given by the energy recovery ratio as computed by the “true” eigenimages. In other words, the first k estimated eigenimages \tilde{U}_k of the matrix product XH_j are shown to be very good estimates of U_k .

V. FAST EIGENDECOMPOSITION ALGORITHM

Our objective is to estimate the first k principal eigenimages \tilde{U}_k of \hat{X} such that $\rho(\hat{X}, \tilde{U}_k) \geq \mu$, where μ is the user specified energy recovery ratio. To this end, we make use of two observations, the first is that reducing the dimensionality of the data in the illumination dimension can be efficiently done using the analysis provided in Section III resulting in a set of harmonic images $\mathbf{f}_{p,q}^r$ at each pose. Note that each harmonic image corresponds to a spherical harmonic of degree

p and order q at each of the a poses. Therefore, each set of harmonic images corresponding to a given value of p and q across all r can be concatenated to form the matrix

$$\hat{X}_{p,q} = [\mathbf{f}_{p,q}^0, \mathbf{f}_{p,q}^1, \dots, \mathbf{f}_{p,q}^{a-1}]. \quad (7)$$

Furthermore, because the harmonic expansion is truncated, there will be nine such matrices in total, each of size $\mathbb{R}^{m \times r}$.

Each of the nine matrices in (7) now only contain variations due to a change in pose for a given spherical harmonic coefficient, and thus are correlated in a single dimension. Therefore, the second observation that can be made is that the dimensionality of the data in the temporal dimension can be reduced by applying the results observed by Chang *et al.* to each of the nine matrices $\hat{X}_{p,q}$. In other words, we can assume that the right singular vectors of $\hat{X}_{p,q}$ are well-approximated by a few low-frequency Fourier harmonics, and the FFT can be used to determine $\rho(\hat{X}_{p,q}^T H_{j_i}) \geq \mu_t$ for each of the nine p, q combinations, where μ_t is a user specified value for the energy recovery along the temporal dimension for each of the nine harmonics. Notice that $j_i, i = 1, 2, \dots, 9$, corresponds to the number of Fourier harmonics required for the i^{th} (p, q) combination to achieve the user specified energy recovery ratio μ_t . Let $Z_{j_i}^{p,q}$ denote the matrix $\hat{X}_{p,q} H_{j_i} \in \mathbb{R}^{m \times j_i}$ for each (p, q) combination, and construct the reduced order matrix

$$\bar{X} = [Z_{j_1}^{0,0}, Z_{j_2}^{1,-1}, Z_{j_3}^{1,0}, Z_{j_4}^{1,1}, Z_{j_5}^{2,-2}, \dots, Z_{j_9}^{2,2}], \quad (8)$$

that effectively recombines the image data due to variation in both illumination and pose into a single matrix. Note that the matrix \bar{X} has considerably fewer columns than that of \hat{X} . Furthermore, because most SVD algorithms require $\mathcal{O}(mn^2)$ flops, computing the dominant left singular vectors \tilde{U}_k of \bar{X} by means of the SVD results in excellent estimates of \hat{U}_k at a significant computational savings. The entire algorithm is summarized as follows:

EIGENDECOMPOSITION ALGORITHM SUMMARY

- 1) Use the SHT to compute the matrices $P_r = [\mathbf{f}_{0,0}^r, \mathbf{f}_{1,-1}^r, \mathbf{f}_{1,0}^r, \mathbf{f}_{1,1}^r, \mathbf{f}_{2,-2}^r, \dots, \mathbf{f}_{2,2}^r]$ for all r .
- 2) Construct the matrices $\hat{X}_{p,q}$ by concatenating each of the harmonic images $\mathbf{f}_{p,q}^r$ in P_r for each r as shown in (7).
- 3) For each of the nine matrices $\hat{X}_{p,q}$, determine the smallest number j_i such that $\rho(\hat{X}_{p,q}^T H_{j_i}) \geq \mu_t$, where μ_t is the user specified energy recovery ratio in the temporal dimension, and $i = 1, 2, \dots, 9$ corresponds to the i^{th} matrix $\hat{X}_{p,q}$.
- 4) Let $Z_{j_i}^{p,q}$ denote the matrix $\hat{X}_{p,q} H_{j_i}$ and construct the matrix \bar{X} defined in (8). Note that the matrices $Z_{j_i}^{p,q}$ can be efficiently computed using the FFT.
- 5) Compute the SVD of $\bar{X} = \tilde{U} \tilde{\Sigma} \tilde{V}$.
- 6) Return $\rho(\hat{X}, \tilde{U}_k) \geq \mu$. Where μ is the user specified energy recovery ratio.

VI. EXPERIMENTAL RESULTS

A. Test Data

The proposed algorithm detailed in Section V was tested on each of the objects in Fig. 2. Recall that each of the objects was sampled at a resolution of 128×128 from 90 different poses under 48 different light source locations at each pose. To accurately represent real objects using CAD models, the reflectance model used accounts for material properties such as surface roughness and surface hardness, and incorporates a mix of diffuse and specular reflection using the Cook-Torrance reflectance model [31]. The mean image was then subtracted to construct the image data matrix \hat{X} . The parameters used in the algorithm were $\mu_t = 0.95$ for the temporal reduction and $\mu = 0.8$ for the total energy recovered. The true SVD of the image data matrix \hat{X} was also computed using MATLAB for a ground truth comparison. The quality measures outlined in Section II were used to evaluate the accuracy of the estimated subspace.

B. Performance and Computational Savings

Table I shows the required subspace dimension k and the time required to estimate the first k left singular vectors \tilde{U}_k for each object in Fig. 2 to meet the user specified energy recovery ratio $\mu = 0.8$. This result is compared to the true SVD as computed by MATLAB. Table I also shows the column dimension of \bar{X} in step 5 of the proposed algorithm. Note that for the current test data, the number of columns in \bar{X} is 4320, whereas for all 20 objects in Fig. 2, the number of columns in \bar{X} never exceeds 576, thus resulting in significant computational savings.

TABLE I
THE REQUIRED SUBSPACE DIMENSION k AND THE TIME REQUIRED TO ESTIMATE THE FIRST k LEFT SINGULAR VECTORS FOR EACH OBJECT TO MEET THE USER SPECIFIED ENERGY RECOVERY RATIO $\mu = 0.8$. THE RESULTS ARE COMPARED AGAINST THE TRUE SVD USING MATLAB. THE TABLE ALSO SHOWS THE COLUMN DIMENSION OF \bar{X} .

Object	Dim. k		Time [min.]		Col. Dim. of \bar{X}
	True	Proposed	True	Proposed	
1	17	17	31.274	0.111	378
2	9	9	25.528	0.070	162
3	13	13	32.342	0.116	379
4	15	15	29.955	0.137	474
5	10	10	31.564	0.076	229
6	14	15	30.954	0.181	576
7	16	17	27.874	0.099	239
8	31	31	31.551	0.152	446
9	19	19	30.842	0.162	502
10	14	14	31.597	0.154	448
11	22	22	31.736	0.117	356
12	20	20	31.825	0.188	561
13	8	8	21.117	0.114	254
14	12	12	30.830	0.107	270
15	23	23	30.776	0.153	472
16	27	27	21.272	0.109	249
17	196	217	22.857	0.183	552
18	20	20	15.501	0.093	173
19	25	25	21.489	0.152	439
20	33	46	21.433	0.099	209
Mean			27.616	0.129	368.400
Min.			15.501	0.070	162
Max.			32.342	0.188	576

Fig. 4 shows the difference in energy recovered by the true SVD and the proposed algorithm for all objects in Fig. 2. As can be seen from the figure, with the exception of object 20, there is less than a 1% difference in how much energy the proposed algorithm is capable of recovering compared to the true SVD. The subspace dimension used for the energy calculation is outlined in column 2 of Table I. As is apparent from Table I and Fig. 4, the estimates \tilde{U}_k using the proposed algorithm, are very good approximations to the left singular vectors \hat{U}_k at a significant computational savings.




Fig. 4. The difference in energy recovered by the true SVD and the proposed algorithm for all objects in Fig. 2. The subspace dimension used for the calculation is listed in column 2 of Table I

VII. SUMMARY

This paper has presented a new algorithm to efficiently estimate the eigendecomposition of an image data set generated due to a change in both illumination and pose. The algorithm is based on using the spherical harmonic transform to reduce the dimensionality of the data due to a change in illumination conditions, generating a set of harmonic images. The harmonic images are then projected onto a few low-frequency Fourier harmonics for a reduction in data due to a change in pose. In addition to significant computational savings as compared to the direct SVD approach, it has been shown that the estimated eigenimages are very close to the true eigenimages as computed by the direct SVD.

REFERENCES

- [1] K. Fukunaga, *Introduction to Statistical Pattern Recognition*. London, U.K.: Academic, 1990.
- [2] J. J. Gerbrands, "On the relationships between SVD, KLT and PCA," *Pattern Recognition*, vol. 14, no. 1-6, pp. 375-381, Dec. 1981.
- [3] H. Murakami and V. Kumar, "Efficient calculation of primary images from a set of images," *IEEE Trans. PAMI*, vol. 4, no. 5, pp. 511-515, Sept. 1982.
- [4] M. Turk and A. Pentland, "Eigenfaces for recognition," *J. Cogn. Neurosci.*, vol. 3, no. 1, pp. 71-86, Mar. 1991.
- [5] H. Murase and S. K. Nayar, "Visual learning and recognition of 3-D objects from appearance," *Int. J. Comp. Vis.*, vol. 14, no. 1, pp. 5-24, Jan. 1995.
- [6] X. Yang, T. K. Sarkar, and E. Arvas, "A survey of conjugate gradient algorithms for solution of extreme eigen-problems for a symmetric matrix," *IEEE Trans. ASSP*, vol. 37, no. 10, pp. 1550-1556, Oct. 1989.
- [7] C. R. Vogel and J. G. Wade, "Iterative SVD-based methods for ill-posed problems," *SIAM J. Sci. Comput.*, vol. 15, no. 3, pp. 736-754, May 1994.
- [8] S. Chandrasekaran, B. Manjunath, Y. Wang, J. Winkler, and H. Zhang, "An eigenspace update algorithm for image analysis," *CVGIP: Graphic Models and Image Proc.*, vol. 59, no. 5, pp. 321-332, Sept. 1997.
- [9] C. Y. Chang, A. A. Maciejewski, and V. Balakrishnan, "Fast eigenspace decomposition of correlated images," *IEEE Trans. Image Proc.*, vol. 9, no. 11, pp. 1937-1949, Nov. 2000.
- [10] R. C. Hoover, A. A. Maciejewski, and R. G. Roberts, "Aerial pose detection of 3-D objects using hemispherical harmonics," in *IEEE SSIAI*, Santa Fe, NM, Mar. 2008, pp. 157-160.
- [11] R. C. Hoover, A. A. Maciejewski, and R. G. Roberts, "Pose detection of 3-D objects using S^2 -correlated images and discrete spherical harmonic transforms," in *IEEE Int. Conf. Robot. Automat.*, Pasadena, CA, May. 2008, pp. 993-998.
- [12] R. C. Hoover, A. A. Maciejewski, and R. G. Roberts, "Pose detection of 3-D objects using images sampled on $SO(3)$, spherical harmonics, and Wigner-D matrices," in *IEEE Conf. on Automat. Sci. and Engr.*, Washington DC, Aug 2008, pp. 47-52.
- [13] R. Epstein, P. Hallinan, and A. Yuille, "5 ± 2 eigenimages suffice: An empirical investigation of low-dimensional lighting models," in *IEEE Workshop on Physics-Based Vision*, Cambridge, MA, June 1995, pp. 108-116.
- [14] S. K. Nayar and H. Murase, "Dimensionality of illumination in appearance matching," in *IEEE Int. Conf. Robot. Automat.*, Minneapolis, MN, Apr. 1996, pp. 1326-1332.
- [15] R. Ramamoorthi, "Analytic PCA construction for theoretical analysis of lighting variability in images of a Lambertian object," *IEEE Trans. PAMI*, vol. 24, no. 10, pp. 1322-1333, Oct. 2002.
- [16] R. Basri and D. Jacobs, "Lambertian reflectance and linear subspaces," *IEEE Trans. PAMI*, vol. 25, pp. 383-390, Feb. 2003.
- [17] K. M. Górski, E. Hivon, A. L. Bandy, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, "HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere," *The Astrophysical Journal*, vol. 622, pp. 759-771, Apr. 2005.
- [18] K. Saitwal, A. A. Maciejewski, R. G. Roberts, and B. Draper, "Using the low-resolution properties of correlated images to improve the computational efficiency of eigenspace decomposition," *IEEE Trans. Image Proc.*, vol. 15, no. 8, pp. 2376-2387, Aug. 2006.
- [19] P. N. Swartztrauber and W. F. Spatz, "Generalized discrete spherical harmonic transforms," *J. of Comp. Phys.*, vol. 159, no. 2, pp. 213-230, Apr. 2000.
- [20] M. A. Wiecezorek and F. J. Simons, "Localized spectral analysis on the sphere," *Geoph. J. Int.*, vol. 162, no. 3, pp. 655-675, Sept. 2005.
- [21] F. J. Simons, F. A. Dahlen, and M. A. Wiecezorek, "Spatiospectral localization on a sphere," *SIAM Review*, vol. 48, no. 3, pp. 504-536, 2006.
- [22] D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonskii, *Quantum Theory of Angular Momentum*. Hackensack, NJ: World Scientific, 1988.
- [23] B. Bustos, D. A. Keim, D. Saupe, T. Schreck, and D. V. Vranic, "Feature-based similarity search in 3-D object databases," *ACM Comp. Surveys*, vol. 37, no. 4, pp. 345-387, Dec. 2005.
- [24] L. Qing, S. Shan, and W. Gao, "Eigen-harmonic faces: Face recognition under generic lighting," in *IEEE Int. Conf. Auto. Face Gesture Rec.*, Seoul, Korea, May 2004, pp. 296-301.
- [25] S. Romdhani, J. Ho, T. Vetter, and D. J. Kriegman, "Face recognition using 3-D models: Pose and illumination," *Proc. of the IEEE*, vol. 294, no. 11, pp. 1977-1999, Nov. 2006.
- [26] L. Zhang and D. Samaras, "Face recognition from a single training image under arbitrary unknown lighting using spherical harmonics," *IEEE Trans. PAMI*, vol. 28, no. 3, pp. 351-363, Mar. 2006.
- [27] A. Makadia and K. Daniilidis, "Rotation recovery from spherical images without correspondences," *IEEE Trans. PAMI*, vol. 28, no. 7, pp. 1170-1175, July 2006.
- [28] D. M. Healy Jr., D. Rockmore, P. Kostelec, and S. Moore, "FFTs for the 2-sphere-improvements and variations," *J. of Fourier Anal. and App.*, vol. 9, no. 4, pp. 341-385, July 2003.
- [29] P. J. Kostelec and D. N. Rockmore, "FFTs on the Rotation Group," Santa Fe Institute Working Papers Series Paper #03-11-060, 2003.
- [30] K. Legaz. (2007) Kator Legaz: 3-D model database for Blender. [Online]. Available: http://www.katorlegaz.com/3d_models/index.php
- [31] (2009) Blender: Open source 3-D modeling application. [Online]. Available: <http://www.blender.org/>