

Pose Detection of 3-D Objects Using S^2 -Correlated Images and Discrete Spherical Harmonic Transforms

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

*Dept. of Electrical and Computer Eng.
Colorado State University
Fort Collins, CO 80523-1373, USA
Email: {hoover, aam}@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—The pose detection of three-dimensional (3-D) objects from two-dimensional (2-D) images is an important issue in computer vision and robotics applications. Specific examples include automated assembly, automated part inspection, robotic welding, and human robot interaction, as well as a host of others. Eigendecomposition is a common technique for dealing with this issue and has been applied to sets of correlated images for this purpose. Unfortunately, for the pose detection of 3-D objects, a very large number of correlated images must be captured from many different orientations. As a result, the eigendecomposition of this large set of images is very computationally expensive. In this work, we present a method for capturing images of objects from many locations by sampling S^2 appropriately. Using this spherical sampling pattern, the computational burden of computing the eigendecomposition can be reduced by using the *Spherical Harmonic Transform* to “condense” information due to the correlation in S^2 . We propose a computationally efficient algorithm for approximating the eigendecomposition based on the spherical harmonic transform analysis. Experimental results are presented to compare and contrast the algorithm against the true eigendecomposition, as well as quantify the computational savings.

I. INTRODUCTION

Over the last several decades, pose detection of three-dimensional (3-D) objects from two-dimensional (2-D) images has become an important issue in computer vision and robotics applications. Specific examples include automated assembly, automated part inspection, robotic welding, and human robot interaction, as well as a host of others. 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], [4]. Once the principal eigenimages of an image data set have been determined, using these eigenimages is very 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 [5]–[7]. A fundamentally different approach was proposed by Chang *et al.* [8] where the authors show that the FFT may be used to approximate the desired subspace dimension, as well as the principal eigenimages if the image data set is correlated in one dimension. In [9] the computational efficiency of Chang’s algorithm is increased further by using the low resolution properties of the image data set.

In this paper, we exploit the work of [8] and extend this to full 3-D object classification using theory from spectral analysis. In general, if recognition and pose detection of 3-D objects is desired, the training image data set needs to contain views of the object from many different orientations. The approach taken here to capture images of the object from a large number of orientations is to consider an imaginary unit sphere with the object placed at the sphere’s center. If a specified number of samples are defined on the sphere, the object’s image data set may be generated by capturing images of the object at each of the specified sample points. Once the image data set is constructed in this manner, we take advantage of the spherical sampling pattern by using the *Spherical Harmonic Transform* (SHT) in place of the FFT to capture the frequency information of this data set. We then propose a computationally efficient algorithm for approximating the subspace dimension, as well as the principal eigenimages using the resulting SHT analysis.

The rest of this paper is organized as follows. In Section II, we explain the fundamentals needed to apply an eigendecomposition to a related image data set, much of which is discussed in [8]. In Section III, we give a brief introduction to spherical harmonics and the SHT. In Section IV we propose a computationally efficient algorithm for approximating the subspace dimension, as well as the principal eigenimages. Experimental results are given in Section V, with conclusions and future work outlined in Section VI.

This work was supported in part by the Missile Defense Agency under contract no. HQ0006-05-C-0035. Approved for Public Release 07-MDA-2981 (12 DEC 07).

II. PRELIMINARIES

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 will be represented by a matrix $\mathcal{X} \in [0, 1]^{h \times v}$. Because sets of related images are considered in this paper, the *image vector* \mathbf{f} of length $m = hv$ is obtained by “row-scanning” an image into a column vector, i.e., $\mathbf{f} = \text{vec}(\mathcal{X}^T)$. 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 $X = [\mathbf{f}_1, \dots, \mathbf{f}_n]$, where typically $m > n$ with fixed n . Because we will be sampling images on the unit sphere, it should be noted that $\mathbf{f} = \mathbf{f}(\boldsymbol{\xi}_p)$ where $\boldsymbol{\xi}_p$, $p \in \{0, \dots, n-1\}$ is the unit vector pointing at the angle of co-latitude $\theta_p \in [0, \pi]$ measured down from the upper pole, and the angle of longitude $\phi_p \in [0, 2\pi)$, which is the parameterization of the sphere in spherical coordinates.

The singular value decomposition (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 columns of U , denoted \mathbf{u}_i , $i = 1, \dots, m$, are referred to as the left singular vectors or eigenimages of X , while the columns of V , denoted \mathbf{v}_i , $i = 1, \dots, n$ are referred to as the right singular vectors of X . The left singular vectors (or eigenimages) of X can be interpreted as the eigenvectors of the covariance matrix of the image vector. The eigenimages provide an orthonormal basis for the column space 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 in X contributes to the i^{th} eigenimage.

In practice, the left singular vectors \mathbf{u}_i are not known or computed exactly, and instead estimates $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k$ which form a k -dimensional basis are used. The accuracy of a practical implementation of subspace methods then depends on three factors: the properties of X , the dimension k , and the quality of the estimates $\tilde{\mathbf{u}}_i$. The measure we will use for quantifying this accuracy is the “energy recovery ratio” denoted ρ , and defined as

$$\rho(X, \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k) = \frac{\sum_{i=1}^k \|\tilde{\mathbf{u}}_i^T X\|^2}{\|X\|_F^2} \quad (2)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. Note that if the $\tilde{\mathbf{u}}_i$ are orthonormal, $\rho \leq 1$, and for any given k achieves a maximum value of $(\sum_{i=1}^k \sigma_i^2) / (\sum_{i=1}^n \sigma_i^2)$ when $\text{span}(\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k) = \text{span}(\mathbf{u}_1, \dots, \mathbf{u}_k)$. It should also be noted that if the columns of X are intensity normalized (which is typical in subspace methods, see for instance [3]), then $\|X\|_F^2$ is simply equal to the number of images in X .

The principal calculation required with subspace methods is the precomputation of estimates of the left singular vectors

$\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k$ of the $m \times n$ matrix X . This is a very computationally expensive operation when m and n are large. Reducing this computational expense by exploiting correlation between images has been the topic of much previous research [5]–[9]. In [8], Chang *et al.* showed that if the image data matrix was correlated in one dimension, then the FFT may be used to approximate the desired subspace dimension k , as well as the principal eigenimages $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k$. This result is based on the fact that for a large class of objects, most of the energy in the image data matrix X is contained in the low frequency harmonics of $\text{FFT}(X)$. Because our objects are no longer correlated in one dimension however, the direct FFT approach is no longer effective. Furthermore, because we are sampling on the sphere, the FFT is multivalued at the poles and thus not well posed in the latitude direction. Our solution to this problem follows the work of Chang, however we use the SHT in place of the FFT to calculate the frequency information of this data set. Using the resulting SHT analysis, it is possible to approximate the desired subspace dimension k , as well as the principal eigenimages of the image data matrix X . A brief introduction to spherical harmonics and the SHT is presented in the next section.

III. SPHERICAL HARMONICS

A. Introduction

Spherical harmonics have been applied to a variety of problems that arise on the surface of the unit sphere (denoted as the 2-sphere or S^2). They have been used for solving PDE’s in spherical geometry for weather and climate models [10], geophysics [11], [12], quantum mechanics [13], [14], as well as a host of other related applications. Over the last decade, spherical harmonics have been gaining popularity in the computer vision and computer graphics arena. Spherical harmonics have been applied to several computer vision applications with unknown lighting [15]–[18], as well as 3-D model retrieval [19], [20], and 3-D shape descriptors [21]. Spherical harmonics have also been applied to rotation estimation and convolution of spherical images [22].

In the context of modern computer vision applications, the information received is in digital format, thus the development of a discrete SHT is needed to process the data. This topic has been addressed in various ways dating back to the 1800’s [23]. Swarztrauber *et al.* proposed a method for computing the discrete SHT using the Gauss-Legendre quadrature [24], [25], as well as an efficient method for computing the quadrature weights and points [26]. Alternatively, Driscoll and Healy proposed a method for computing the discrete SHT using an equi-angular grid of Chebyshev nodes [27], [28]. In [29], Górski *et al.* propose the Hierarchical Equal Area isoLatitude Pixelization (HEALPix) which has the advantage that the samples have equal area weighting over S^2 , as a result, they don’t oversample the polar regions. In [30], the authors compare and contrast the three above mentioned tessellations and it was determined that for this particular application, the HEALPix tessellation performed the best in terms of better angular resolution in sampling, and

better estimation of the eigenspace. For the aforementioned reasons, in this paper, we use the HEALPix tessellation to define the sampling pattern over S^2 used for the construction of a discrete SHT (see Fig. 1 for an example of the HEALPix sampling pattern).

Spherical harmonics, typically denoted Y_l^m , are the angular solutions to Laplace's equation in spherical coordinates, and have the factorization

$$Y_l^m(\theta, \phi) = \kappa_l^m P_l^m(\cos(\theta)) e^{jm\phi}, \quad (3)$$

where $P_l^m(\cos(\theta))$ is the associated Legendre polynomial of degree l and order m , and $\kappa_{l,m}$ is a normalization constant, equal to

$$\kappa_l^m = \sqrt{\left(\frac{2l+1}{4\pi}\right) \frac{(l-|m|)!}{(l+|m|)!}}. \quad (4)$$

Using the above normalization, the spherical harmonics satisfy the condition

$$\int_{\theta=0}^{\pi} \int_{\phi=0}^{2\pi} Y_{l_1}^{m_1} (Y_{l_2}^{m_2})^* \sin(\theta) d\theta d\phi = \delta_{l_1 l_2} \delta_{m_1 m_2}, \quad (5)$$

where the superscript $*$ is the complex conjugate, i.e., they form an orthonormal basis for S^2 . As a result, any square integrable function $f(\theta, \phi) \in L^2(S^2)$ may be projected onto this basis, where L^2 is the Hilbert space of square integrable functions.

B. Discrete Spherical Harmonic Transform

A real valued band-limited function $f(\xi_p)$ whose domain is $L^2(S^2)$ may be represented by its spherical harmonic expansion as

$$f(\xi_p) = \sum_{l=0}^{l_{\max}} \sum_{|m| \leq l} f_l^m Y_l^m(\xi_p) \quad (6)$$

where $f(\xi_p) \in [0, 1]$ is a single pixel of the image data vector $\mathbf{f}(\xi_p)$. Once again we remind the reader that ξ_p , $p \in \{0, \dots, n-1\}$ is the unit vector pointing at the angle of colatitude $\theta_p \in [0, \pi]$ measured down from the upper pole, and the angle of longitude $\phi_p \in [0, 2\pi]$ is the parameterization of the unit sphere in spherical coordinates. In the above equation, it is assumed that the signal power for $l > l_{\max}$ is insignificant, and l_{\max} is chosen such that aliasing does not occur. The expansion coefficients are calculated using

$$f_l^m = \frac{4\pi}{n} \sum_{p=0}^{n-1} f(\xi_p) Y_l^m(\xi_p), \quad (7)$$

where $Y_l^m(\xi_p)$ is the real-valued spherical harmonic defined by

$$Y_l^m(\xi_p) = \begin{cases} \sqrt{2} \kappa_l^m \cos(m\phi_p) P_l^m(x) & \text{if } m > 0 \\ \sqrt{2} \kappa_l^m \sin(|m|\phi_p) P_l^{|m|}(x) & \text{if } m < 0 \\ \kappa_l^0 P_l^0(x) & \text{if } m = 0 \end{cases} \quad (8)$$

where $P_l^0(x) = P_l(x)$ is the Legendre polynomial of degree l and $x = \cos(\theta_p)$.

IV. FAST EIGENDECOMPOSITION ALGORITHM

Our objective is to estimate the desired subspace dimension k , as well as the principal eigenimages $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k$ of X such that $\rho(X, \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k) \geq \mu$, where μ is the user specified energy recovery ratio. The first step in computing the desired subspace dimension k , as well as the principal eigenimages, is to construct the image data matrix X . As mentioned previously, the approach taken here is to consider the object placed at the center of an imaginary unit sphere, and sample S^2 using the HEALPix sampling pattern (see Fig. 1).




Fig. 1. Image acquisition using the HEALPix sampling pattern to obtain the training image data matrix X . A sample image is taken at each black dot on the sphere.

Using the HEALPix sampling pattern is based on subdividing the sphere using the parameter N_{side} , resulting in $12N_{\text{side}}^2$ images captured from various vantage points on the sphere [29]. Setting $l_{\max} = 3N_{\text{side}} - 1$, the Legendre polynomials will not alias on the sphere. Furthermore, because the sampling pattern is isolatitudinal, the computation of the Legendre polynomials (which is the most computationally expensive portion of the SHT) is minimized. In this paper, we are using CAD generated ray-traced images, examples of which are shown in Fig. 2, (the CAD models were provided by [31]).

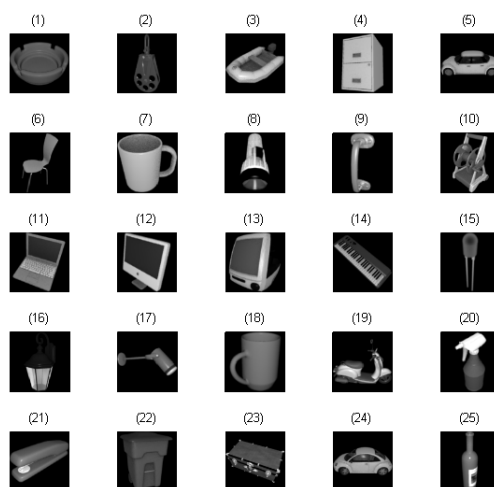


Fig. 2. Ray-traced CAD models courtesy of Kator Legaz [31]. Each object is sampled as discussed above at a resolution of 128×128 . Each of the above images are both scale and intensity normalized.

Once the image data matrix X has been constructed, we compute the matrix F whose i^{th} row is the SHT of the i^{th} row of X , denoted from this point forward as SHT(X). This can be computed quickly using the methods described in [28], however for small bandwidths the computational savings of this method are slim. Alternatively, the SHT(X) can be cast into a linear algebra problem as

$$F = XZ, \quad (9)$$

where $X \in \mathbb{R}^{m \times (12N_{\text{side}}^2)}$ is the image data matrix with the images ordered in terms of ξ_p as

$$X = \left[\mathbf{f}(\xi_0), \mathbf{f}(\xi_1), \dots, \mathbf{f}(\xi_{12N_{\text{side}}^2 - 1}) \right], \quad (10)$$

and $Z \in \mathbb{R}^{(12N_{\text{side}}^2) \times (9N_{\text{side}}^2)}$ is the matrix of spherical harmonics ordered in terms of ξ_p , l , and m as shown in (11). Note that in computing the SHT(X) the matrix Z may be pre-computed for several different bandwidths and stored for later use.

Even though the image data matrix X is correlated in S^2 (i.e., correlated in both θ and ϕ), similar to [8], most of the energy of X is contained in the low frequency harmonics of F . An example of this is illustrated in Fig. 3, which is the spherical harmonic power spectra of object (17) from Fig. 2. As a result, the principal eigenimages $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k$ of SVD(F) serve as excellent estimates to those of X . Moreover, the accuracy of the approximated eigenimages with power spectra centered around the lower frequencies will tend to be better. We now propose an algorithm for estimating the required subspace dimension k , as well as the principal eigenimages of X , based on a user specified energy recovery ratio μ .

SHT BASED EIGENDECOMPOSITION ALGORITHM

- 1) Form the matrix F which is the SHT(X).
- 2) Form the matrix H whose columns are the ordered columns of F in descending order according to their norm.
- 3) Set $q = (3N_{\text{side}})^2 [1 - (1/2)^{N+1}]$, with $N=0$ initially.




Fig. 3. This figure shows the spherical harmonic power spectra of object (17) from Fig. 2. As can be seen from the figure, most of the energy is concentrated around the low frequency harmonics.

- 4) Construct the matrix H_q which is the matrix consisting of the first q columns of H .
- 5) Compute SVD(H_q). The key observation here is that H_q contains q columns which is considerably less than the $12N_{\text{side}}^2$ columns of X .
- 6) If $\rho(X, \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_q) < \mu$. Let $N = N + 1$ and repeat Steps 3 through 6. Because the SVD of H_q is already available, the eigenspace can simply be updated by modifying the algorithm outlined in [7].
- 7) Return $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k$ such that $\rho(X, \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k) \geq \mu$. Note that $k \leq q$.

Note that for $12N_{\text{side}}^2$ samples on the sphere, only $9N_{\text{side}}^2$ harmonic images are generated, as a result, it is not possible to recover 100% of the energy. While this may be viewed as a drawback in terms of image reconstruction, this is actually an advantage in terms of computational complexity. Furthermore, because subspace methods typically work well with a much smaller subspace dimension, and as a result less recoverable energy, the proposed algorithm has shown significant computational savings with little loss in performance. The actual amount of recoverable energy is dependant on the object, however for all 25 objects in Fig. 2, over 98% of the energy was recoverable at $N_{\text{side}} = 8$.

V. EXPERIMENTAL RESULTS

The proposed algorithm detailed in Section IV was tested on the objects shown in Fig. 2. The parameter $N_{\text{side}} = 8$ was used, resulting in $12N_{\text{side}}^2 = 768$ images per object. The images were then both scale and intensity normalized to create the image data matrix X . Finally, the matrix F was computed condensing the image data set from 768 images to 576 harmonic images. The true SVD(X) was also computed for a ground truth comparison.

A. Performance and Computational Savings

Figure 4 shows a plot of the energy recovery ratio ρ as a function of the subspace dimension k averaged across all objects. As can be seen from the figure, the estimated left singular vectors $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k$ for the proposed algorithm are very good approximations to the true left singular vectors as computed by the SVD(X) in terms of being able to recover the total energy in X . Note, this data is for the first split, i.e., $N = 0$, in step 3 of the algorithm, if the entire $N = (9N_{\text{side}}^2)$ harmonic images are used, the difference becomes indistinguishable.

Table I shows the computation time required to calculate the subspace dimension k , and estimate the left singular vectors $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k$ needed to meet the user specified energy recovery ratio $\mu = 0.9$. Also depicted in the table is the result as computed by the true SVD(X). As is apparent from the table, in most cases, the left singular vectors $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k$ are very good estimates of $\mathbf{u}_1, \dots, \mathbf{u}_k$ at a significant computational savings for the proposed algorithm. Again, only the first split, i.e., $N = 0$, in step 3 of the algorithm is required to recover over 90% of the energy as depicted in the table.

$$Z = \begin{bmatrix} Y_0^0(\xi_0) & Y_1^{-1}(\xi_0) & Y_1^0(\xi_0) & Y_1^1(\xi_0) & Y_2^{-2}(\xi_0) & \cdots & Y_{3N_{\text{side}}-1}^{3N_{\text{side}}-1}(\xi_0) \\ Y_0^0(\xi_1) & & & \ddots & & & Y_{3N_{\text{side}}-1}^{3N_{\text{side}}-1}(\xi_1) \\ \vdots & & & & & & \vdots \\ Y_0^0(\xi_{12N_{\text{side}}-1}) & & \cdots & & & & Y_{3N_{\text{side}}-1}^{3N_{\text{side}}-1}(\xi_{12N_{\text{side}}-1}) \end{bmatrix} \quad (11)$$




Fig. 4. This figure shows the energy recovery ratio ρ vs. subspace dimension k averaged across all 25 objects in Fig. 2.

TABLE I

THIS TABLE SHOWS THE REQUIRED SUBSPACE DIMENSION k , AND THE TIME REQUIRED TO ESTIMATE THE FIRST k LEFT SINGULAR VECTORS FOR EACH OBJECT IN FIG. 2 TO EXCEED THE USER SPECIFIED ENERGY RECOVERY $\mu = 0.9$. THE RESULTS ARE COMPARED AGAINST THE TRUE SVD USING MATLAB.

Object no.	Dimension k		Time [sec.]	
	True	Proposed	True	Proposed
1	2	2	21.78	6.60
2	9	9	24.12	6.88
3	18	18	24.36	7.03
4	3	3	23.89	6.70
5	21	21	24.71	7.13
6	67	71	25.97	8.34
7	5	5	24.00	6.92
8	3	3	26.38	6.67
9	26	26	24.27	7.20
10	60	61	25.81	8.08
11	22	22	24.77	7.16
12	35	36	29.87	7.48
13	26	26	24.62	7.25
14	47	48	25.12	7.76
15	8	8	23.34	6.77
16	19	20	24.79	7.28
17	47	49	25.31	7.81
18	3	3	23.38	6.69
19	54	56	29.93	7.95
20	11	11	23.82	6.87
21	12	12	31.95	6.94
22	2	2	27.44	6.71
23	91	102	26.35	9.11
24	35	35	25.46	7.42
25	4	4	23.39	6.82

B. Error Analysis

Based on the data provided in Table I, it is obvious that the proposed algorithm is capable of estimating the required

subspace dimension and left singular vectors of X at a significant computational savings. Because the SHT is lossy however, it is important to try to quantify the amount of energy lost. To do this, all $9N_{\text{side}}^2$ harmonic images were used to estimate the left singular vectors of X . Using these estimates, the energy recovery ratio $\rho(X, \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_{9N_{\text{side}}^2})$ was computed for each object in Fig. 2. The top plot in Fig. 5 shows the maximum amount of energy recovered per object. The bottom plot shows the maximum difference in energy recovery per object if only the first $9N_{\text{side}}^2$ true left singular vectors are used to recover the same energy, i.e., the bottom plot shows $\max(\rho(X, \mathbf{u}_1, \dots, \mathbf{u}_k) - \rho(X, \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_k))$ for all $k \leq 9N_{\text{side}}^2$ for each object. As seen in the figure, using the proposed algorithm, over 98% of the energy in X is recoverable for all objects. Furthermore, there is less than a 2% difference between the energy recovered by the first $9N_{\text{side}}^2$ true left singular vectors and the left singular vectors computed by the proposed algorithm for all objects in Fig. 2. Finally, using

$$\frac{1}{4.5N_{\text{side}}^2} \sum_{i=2}^{4.5N_{\text{side}}^2} [\rho(X, \mathbf{u}_1, \dots, \mathbf{u}_i) - \rho(X, \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_i)] \quad (12)$$

the average error in energy recovery is computed for each object in Fig. 2. The results are depicted in the top plot of Fig. 6 which shows that the average error one could expect to see for the objects of Fig. 2 is less than 1.5%. Note that $4.5N_{\text{side}}^2$ is chosen in the summation of (12) because this is the number of harmonic images resulting in the first split (step 3) of the proposed algorithm. The bottom plot in Fig. 6 shows the error as a function of subspace dimension averaged across all objects in Fig. 2. Note that the first four splits are shown and that there is less than 1% error across all objects regardless of the number of splits used.

VI. CONCLUSIONS AND FUTURE WORK

We have illustrated a computationally efficient algorithm for estimating the eigendecomposition of images correlated in S^2 using the discrete SHT. The algorithm was tested on a variety of 3-D objects with images captured from different vantage points around the sphere. In addition to significant computational savings as compared to the direct SVD approach, we have shown that the estimated subspace dimension is typically very close to that computed by the direct SVD approach. We have also shown that the errors associated with using the proposed algorithm are minimal with respect to the computational savings. Future work will focus on validating the proposed algorithm on true 3-D objects rather than CAD models, and extending the algorithm to the 3-D rotation group $SO(3)$.




Fig. 5. This figure shows the maximum achievable energy recovery ratio ρ for each object in Fig. 2 [top], and the maximum difference in energy recovery between the first $9N_{\text{side}}^2$ true left singular vectors as computed by SVD(X) and the estimated left singular vectors $\hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_{9N_{\text{side}}^2}$ as computed by the proposed algorithm [bottom].




Fig. 6. This figure shows the average error in energy recovery using the true left singular vectors as computed by SVD(X), and the approximated left singular vectors calculated using the proposed algorithm for each object in Fig. 2 across the subspace dimension k [top]. The bottom plot shows the error in energy recovery averaged across all objects in Fig. 2 using the approximated left singular vectors calculated by the first four splits (step 3) of the proposed algorithm.

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, 1981.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] 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.

- [7] 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.
- [8] 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.
- [9] 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.
- [10] P. N. Swartztrauber, "The vector harmonic transform method for solving partial differential equations in spherical geometry," *Mon. Wea. Rev.*, vol. 121, pp. 3415–3437, 1993.
- [11] M. A. Wicczorek and F. J. Simons, "Localized spectral analysis on the sphere," *Geoph. J. Int.*, vol. 162, no. 3, pp. 655–675, 2005.
- [12] F. J. Simons, F. A. Dahlen, and M. A. Wicczorek, "Spatiospectral localization on a sphere," *SIAM Review*, vol. 48, no. 3, pp. 504–536, 2006.
- [13] D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonskii, *Quantum Theory of Angular Momentum*. New Jersey, USA: World Scientific, 1988.
- [14] G. B. Arfken, *Mathematical Methods for Physicists*. New York, USA: Academic Press, 1970.
- [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] L. Qing, S. Shan, and W. Gao, "Eigen-harmonic faces: Face recognition under generic lighting," in *IEEE Int. Conf. Auto. Face Gesture Rec.*, May 17-19, 2004, pp. 296–301.
- [17] 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.
- [18] 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.
- [19] 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.
- [20] D. Saupe and D. V. Vranic, "3-D model retrieval with spherical harmonics and moments," in *DAGM-Symp. on Patt. Rec.*, 2001, pp. 392–397.
- [21] D. V. Vranic, "An improvement of rotation invariant 3d-shape based on functions on concentric spheres," in *IEEE Conf. Image Proc.*, Sept 14-17, 2003, pp. 757–760.
- [22] A. Makadia and K. Daniilidis, "Rotation recovery from spherical images without correspondences," *IEEE Trans. PAMI*, vol. 28, no. 7, pp. 1170–1175, July 2006.
- [23] N. Sneeuw, "Global spherical harmonic analysis by least squares and numerical quadrature methods in historical perspective," <http://www.geomatics.ucalgary.ca/~sneeuw/publ/neumann.ps>, 2007.
- [24] P. N. Swartztrauber and W. F. Spitz, "Generalized discrete spherical harmonic transforms," *J. of Comp. Phys.*, vol. 159, no. 2, pp. 213–230, Apr. 2000.
- [25] P. N. Swartztrauber, "On the spectral approximation of discrete scalar and vector functions on the sphere," *SIAM J. Numer. Anal.*, vol. 16, no. 6, pp. 934–949, Dec. 1979.
- [26] —, "Computing the points and weights for Gauss-Legendre quadrature," *SIAM J. on Sci. Comp.*, vol. 24, no. 3, pp. 945–954, 2002.
- [27] J. R. Driscoll and D. M. Healy Jr., "Computing Fourier transforms and convolutions on the 2-sphere," *Adv. App. Math.*, vol. 15, no. 2, pp. 202–250, 1994.
- [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] 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.
- [30] R. C. Hoover, A. A. Maciejewski, and R. G. Roberts, "An analysis of sphere tessellations for pose estimation of 3-D objects using spherically correlated images," in *IEEE SSIAP*, accepted to appear March 24-27, 2008.
- [31] K. Legaz, "Kator Legaz," http://www.katorlegaz.com/3d_models/index.php, 2007.