

Eigendecomposition of Images Correlated on S^1 , S^2 , and $SO(3)$ Using Spectral Theory

Randy C. Hoover, *Student Member, IEEE*, Anthony A. Maciejewski, *Fellow, IEEE*, and Rodney G. Roberts, *Senior Member, IEEE*

Abstract—Eigendecomposition represents one computationally efficient approach for dealing with object detection and pose estimation, as well as other vision-based problems, and has been applied to sets of correlated images for this purpose. The major drawback in using eigendecomposition is the off line computational expense incurred by computing the desired subspace. This off line expense increases drastically as the number of correlated images becomes large (which is the case when doing fully general 3-D pose estimation). Previous work has shown that for data correlated on S^1 , Fourier analysis can help reduce the computational burden of this off line expense. This paper presents a method for extending this technique to data correlated on S^2 as well as $SO(3)$ by sampling the sphere appropriately. An algorithm is then developed for reducing the off line computational burden associated with computing the eigenspace by exploiting the spectral information of this spherical data set using spherical harmonics and Wigner- D functions. Experimental results are presented to compare the proposed algorithm to the true eigendecomposition, as well as assess the computational savings.

Index Terms—Computer vision, correlation, data compression, eigenspace, image sampling, pose estimation, singular value decomposition, spherical harmonics, Wigner- D functions.

I. INTRODUCTION

OVER the last several decades, pose detection of 3-D objects from 2-D images has become an important issue in computer vision applications. Subspace methods, also referred to as eigenspace methods, principal component analysis, or the Karhunen-Loeve transformation [1], [2], represent one computationally efficient approach that has been applied to a range of computer vision problems. Specific examples include face char-

acterization and recognition [3]–[8], lip reading [9], [10], object recognition and pose estimation [11]–[20], as well as a host of applications that arise in industrial automation [21]. These applications take advantage of the fact that a set of highly correlated images can be approximately represented by a reduced dimensional set of eigenimages [14], [22].

Once the principal eigenimages of an image data set have been determined, using these eigenimages is very computationally efficient for the online 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 [23]–[25]. A fundamentally different approach was proposed by Chang *et al.* [16] where the authors show that the fast Fourier transform (FFT) may be used to estimate the desired subspace dimension, as well as the principal eigenimages if the image data set is correlated in one dimension. Examples of one dimensionally correlated image data sets given in [16] were arbitrary video sequences, as well as sequences of objects rotated through a single axis of rotation (referred to here as S^1). Chang's eigendecomposition algorithm was extended to data correlated in higher dimensions in [19] where the image data set was generated by capturing images of objects from a spherical patch above the object. Capturing images from a spherical patch above the object allows the sampling dimensions to remain orthogonal, and, thus, an FFT may be applied to each dimension. Unfortunately, the size of the patch must be limited to maintain orthogonality. As a result, this technique is not directly applicable to full 3-D pose estimation. In this paper, we present a computationally efficient method to estimate the eigenspace for images correlated in two dimensions (S^2) as well as correlation in three dimensions ($SO(3)$) by sampling the sphere appropriately and using spherical harmonics and Wigner- D functions to extract the spectral information of these spherically correlated data sets.

The remainder of this paper is organized as follows. In Section II, we review the fundamentals of applying an eigendecomposition to a related image data set and discuss three quality measures used in this paper. In Section III, a brief introduction to spherical harmonics is provided along with a discussion of harmonic analysis on S^2 and $SO(3)$ required to compute the spectral information of a spherically correlated data set. We also propose a method for sampling S^2 as well as $SO(3)$. In Section IV, we show how Chang's algorithm [16] motivates a possible extension to correlation in higher dimensions. We then present a

Manuscript received July 11, 2008; revised June 02, 2009. First published June 30, 2009; current version published October 16, 2009. This work was supported in part by the Missile Defense Agency under Contract HQ0006-05-C-0035. Preliminary versions of portions of this work were presented at the IEEE International Conference on Robotics and Automation (ICRA), Pasadena, CA, May 19–23, 2008, the IEEE Southwest Symposium on Image Analysis and Interpretation (SSIAI), Santa Fe, NM, March 24–26, 2008, and the IEEE Conference on Automation Science and Engineering (CASE), Washington, DC, August 23–26, 2008. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Sharathchandra Pankanti.

R. C. Hoover is with the Department of Mathematics and Computer Science, South Dakota School of Mines and Technology, Rapid City, SD 57701 USA (e-mail: randy.hoover@sdsmt.edu).

A. A. Maciejewski is with the Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO 80523-1373 USA (e-mail: aam@colostate.edu).

R. G. Roberts is with the Department of Electrical and Computer Engineering, Florida A & M—Florida State University, Tallahassee, FL 32310-6046 USA (e-mail: rroberts@eng.fsu.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2009.2026622

computationally efficient algorithm for computing the eigendecomposition used for fully general 3-D pose estimation based on the harmonic analysis of Section III. Experimental results are given in Section V in which we compare the proposed algorithm to the true eigendecomposition in terms of computational savings and accuracy of estimation. Because the spherical harmonic transform (SHT) is a lossy transform, we also present an error analysis in this section. Finally, concluding remarks and future directions are presented in Section VI.

II. MATHEMATICAL PRELIMINARIES

A. Introduction

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 $\mathcal{X} \in [0, 1]^{h \times v}$. 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 $m > n$ [16]. Because we will be sampling images on the sphere, it should be noted that $n = ab$, where a is the number of samples defined on the sphere, and b is the number of planar rotations captured at each sample. The image vector is then $\mathbf{f} = \mathbf{f}(\xi_p, \gamma_r)$ where ξ_p , $p \in \{0, \dots, a-1\}$, is the unit vector pointing at the angle of co-latitude $\beta_p \in (0, \pi)$ measured down from the upper pole, and the angle of longitude $\alpha_p \in [0, 2\pi)$, which is the parameterization of the sphere in spherical coordinates. In $\mathbf{f}(\xi_p, \gamma_r)$, the value $\gamma_r \in [0, 2\pi)$ is the r th planar rotation at sample p where $r \in \{0, \dots, b-1\}$. Using this notation, sampling S^2 can be achieved by setting γ_r to a constant, and allowing α_p and β_p to vary accordingly, and $SO(3)$ can be sampled by allowing all three parameters to vary accordingly. The average image vector is then subtracted from each image in 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.

The *thin* singular value decomposition (SVD) of \hat{X} is given by $\hat{X} = \hat{U}\hat{\Sigma}\hat{V}^T$ where $\hat{U} \in \mathbb{R}^{m \times n}$ and $\hat{V} \in \mathbb{R}^{n \times n}$ are orthogonal, and $\hat{\Sigma} \in \mathbb{R}^{n \times n}$ where $\hat{\Sigma} = \text{diag}(\hat{\sigma}_1, \dots, \hat{\sigma}_n)$ with $\hat{\sigma}_1 \geq \hat{\sigma}_2 \geq \dots \geq \hat{\sigma}_n \geq 0$. The columns of \hat{U} , denoted $\hat{\mathbf{u}}_i$, $i = 1, \dots, n$, are referred to as the left singular vectors or 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 corresponding singular values measure how “aligned” the columns of \hat{X} are with the associated eigenimage. (See Table I for a list of all notations used in this paper.)

B. Quality Measures

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 the quality of the estimates. In this work, the following measures are used to quantify the quality of the estimated eigenimages.

TABLE I
NOMENCLATURE USED THROUGHOUT THIS PAPER

\mathcal{X} :	image matrix
X :	image data matrix
\hat{X} :	“unbiased” (average subtracted) image data matrix
h :	number of pixels in one column of each image
v :	number of pixels in one row of each image
m :	total number of pixels ($h \times v$) in each image
n :	total number of images in an image data set
S^1 :	the circle group
S^2 :	the surface of the sphere (2-sphere)
$SO(3)$:	the rotation group
α_p :	angle of longitude
β_p :	angle of co-latitude
γ_r :	planar rotation angle
p :	p th sample on the surface of the sphere (S^2)
r :	r th planar rotation at sample p
a :	number of images captured on the spheres surface
b :	number of planar rotations at each of the a samples
ξ_p :	unit vector parameterization of S^2
$f(\xi_p, \gamma_r)$:	function sampled on $SO(3)$
U :	matrix of left singular vectors of X
Σ :	diagonal matrix of singular values of X
V :	matrix of right singular vectors of X
\mathbf{u}_i :	the i th column of U
σ_i :	the i th diagonal element of Σ
\mathbf{v}_i :	the i th column of V
U_k :	matrix consisting of the first k columns of U
$\hat{\mathbf{u}}_i$:	estimate of \mathbf{u}_i
$\hat{\sigma}_i$:	estimate of σ_i
$\hat{\mathbf{v}}_i$:	estimate of \mathbf{v}_i
\hat{U} :	matrix of left singular vectors of \hat{X}
$\hat{\Sigma}$:	diagonal matrix of singular values of \hat{X}
\hat{V} :	matrix of right singular vectors of \hat{X}
$\hat{\mathbf{u}}_i$:	the i th column of \hat{U}
$\hat{\sigma}_i$:	the i th diagonal element of $\hat{\Sigma}$
$\hat{\mathbf{v}}_i$:	the i th column of \hat{V}
\hat{U}_k :	matrix consisting of the first k columns of \hat{U}
\tilde{U} :	an estimate for \hat{U}
$\tilde{\Sigma}$:	an estimate for $\hat{\Sigma}$
\tilde{V} :	an estimate for \hat{V}
\tilde{U}_k :	matrix consisting of the first k columns of \tilde{U}
Y_l^m :	spherical harmonic of degree l and order m
$D_{m,m'}^l$:	Wigner- D matrix of degree l and orders m, m'
$d_{m,m'}^l$:	Wigner’s (small) d -matrix of degree l and orders m, m'

1) *Residue Between Subspaces*: The possibility that the data matrix $\hat{U} \in \mathbb{R}^{m \times k}$ can be rotated into the data matrix $\tilde{U} \in \mathbb{R}^{m \times k}$ is explored in [26] by solving the problem

$$\Delta = \min_Q \|\hat{U} - \tilde{U}Q\|_F \quad (1)$$

where $\|\cdot\|_F$ represents the Frobenius norm, $Q \in \mathbb{R}^{k \times k}$ is an orthogonal matrix, and Δ is the residue. The residue can be computed as

$$\Delta = \sqrt{2 \left(k - \sum_{i=1}^k \sigma_{ci} \right)} \quad (2)$$

where σ_{ci} is the i th singular value of $\tilde{U}^T \hat{U}$ [27]. The smaller the residue Δ , the closer \hat{U} and \tilde{U} are to representing the same subspace.

2) *Energy Recovery Ratio*: True and estimated eigenimages of \hat{X} can also be compared in terms of their capability of recovering the amount energy in \hat{X} . The “energy recovery ratio” ρ is defined as

$$\rho(\hat{X}, \tilde{U}_k) = \frac{\sum_{i=1}^k \|\tilde{\mathbf{u}}_i^T \hat{X}\|^2}{\|\hat{X}\|_F^2} \quad (3)$$

where $\|\cdot\|_F$ denotes the Frobenius norm, and $\tilde{\mathbf{u}}_i$ is the i th column of \tilde{U}_k . Note that if the vectors $\tilde{\mathbf{u}}_i$ are orthonormal, then $\rho \leq 1$ [16]. In this work, we also use the change in ρ

$$\Delta\rho(\hat{X}, \tilde{U}_k) = \rho(\hat{X}, \tilde{U}_k) - \rho(\hat{X}, \tilde{U}_{k-1}) \quad (4)$$

which quantifies how much additional energy is recovered by adding the k th basis vector to the subspace [16].

3) *Subspace Criterion*: True eigenimages give an optimum energy recovery ratio; therefore, it is possible that more estimated eigenimages are required to achieve the same user specified energy recovery ratio. Hence, another measure used in this paper is the degree to which estimated eigenimages span the subspace of the first k^* true eigenimages, which will be referred to as the subspace criterion, SC, given by [27]

$$\text{SC} = \sqrt{\frac{1}{k^*} \sum_{i=1}^k \sum_{j=1}^{k^*} (\tilde{\mathbf{u}}_i^T \hat{\mathbf{u}}_j)^2}. \quad (5)$$

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 [28], geophysics [29], [30], quantum mechanics [31], [32], as well as a host of other related applications [33]. 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 [34]–[37], as well as 3-D model retrieval [38], [39], 3-D shape descriptors [40], and pose estimation [17], [18], [20]. Spherical harmonics have also been applied to rotation estimation and convolution of spherical images [41].

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

$$Y_l^m(\alpha, \beta) = \kappa_l^m P_l^m(\cos(\beta)) e^{im\alpha} \quad (6)$$

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

$$\kappa_l^m = \sqrt{\left(\frac{2l+1}{4\pi}\right) \frac{(l-|m|)!}{(l+|m|)!}}. \quad (7)$$

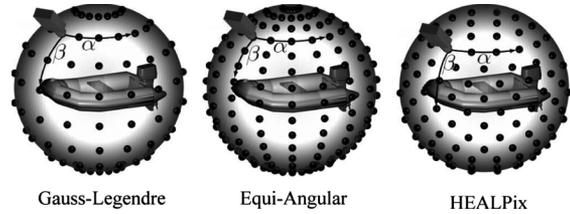


Fig. 1. Three tessellations of the sphere discussed in the section.

Using the above normalization, the spherical harmonics satisfy the condition

$$\int_0^\pi \int_0^{2\pi} Y_{l_1}^{m_1} (Y_{l_2}^{m_2})^* \sin(\beta) d\alpha d\beta = \delta_{l_1 l_2} \delta_{m_1 m_2} \quad (8)$$

where the superscript $*$ is the complex conjugate and δ_{ij} is the Kronecker delta, i.e., they form an orthonormal basis for S^2 . As a result, any square integrable function $f(\alpha, \beta) \in L^2(S^2)$, where L^2 is the Hilbert space of square integrable functions, may be projected onto this basis as

$$f(\alpha, \beta) = \sum_{l=0}^{\infty} \sum_{|m| \leq l} f_l^m Y_l^m(\alpha, \beta) \quad (9)$$

with the expansion coefficients computed as

$$f_l^m = \int_0^{2\pi} \int_0^\pi f(\alpha, \beta) Y_l^m(\alpha, \beta)^* \sin(\beta) d\beta d\alpha. \quad (10)$$

Unfortunately, computing the harmonic coefficients by evaluating the integrals in (10) is prohibitively expensive. In order to reduce this computational expense, the integrals in (10) need to be approximated by finite sums, and the development of a discrete SHT is needed. This topic has been addressed in various ways dating back to the 1800s [42] and is the subject of the next subsection.

B. Discrete Spherical Harmonic Transform

In the development of a discrete spherical harmonic transform, a first step is deciding the best tessellation of the sphere to define the sampling pattern. Three popular discretizations are commonly used when performing spectral analysis on the surface of the sphere. In [43] and [44], Swartztrauber *et al.* proposed a method for computing the discrete SHT using the Gauss-Legendre grid, as well as an efficient method for computing the quadrature weights and points [45]. Alternatively, Driscoll and Healy proposed a method for computing the discrete SHT using an equi-angular grid of Chebyshev nodes [46], [47]. In [48], 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 and as a result, they do not oversample the polar regions. An example of all three of the above mentioned tessellations is shown in Fig. 1. In [49], the authors compared the three tessellations and 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. Therefore, we use the HEALPix tessellation to define the sampling pattern over S^2 used for the construction of the discrete SHT.

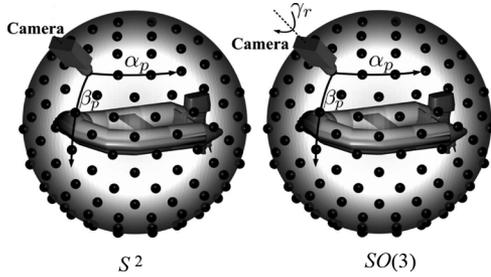


Fig. 2. Method of sampling used in this paper. The above samples are described by three parameters (α, β, γ) , and as such this method is sufficient to describe sampling over S^2 [left] and $SO(3)$ [right].

To construct the image data matrix \hat{X} , consider capturing images on the surface of a sphere where an object is placed at the sphere's center as shown in Fig. 2. If $\beta_p \in (0, \pi)$, $\alpha_p \in [0, 2\pi)$, and γ_r is set to a constant, the sampling pattern is restricted to the surface of the sphere (refer to the left image in Fig. 2). Sampling the object in this manner results in a two dimensionally correlated image data matrix that is correlated on S^2 .

A real valued band-limited function $f(\xi_p, 0)$ whose domain is $L^2(S^2)$ may be represented by its discrete spherical harmonic expansion as

$$f(\xi_p, 0) = \sum_{l=0}^{l_{\max}} \sum_{|m| \leq l} f_l^m Y_l^m(\xi_p, 0) \quad (11)$$

where $f(\xi_p, 0) \in [0, 1]$ is a single pixel of the image data vector $\mathbf{f}(\xi_p, 0)$. Recall that $\xi_p, p \in \{0, \dots, a-1\}$, is the unit vector pointing at the angle of co-latitude $\beta_p \in (0, \pi)$ measured down from the upper pole, and the angle of longitude $\alpha_p \in [0, 2\pi)$, which is the parameterization of the sphere in spherical coordinates. In the above equation, it is assumed that the signal power for $l > l_{\max}$ is insignificant and that l_{\max} is chosen to prevent aliasing. The expansion coefficients are calculated using

$$f_l^m = \frac{4\pi}{n} \sum_{p=0}^{n-1} f(\xi_p, 0) Y_l^m(\xi_p, 0) \quad (12)$$

where $Y_l^m(\xi_p, 0)$ is the real-valued spherical harmonic defined by

$$Y_l^m(\xi_p, 0) = \begin{cases} \sqrt{2} \kappa_l^m \cos(m\alpha_p) P_l^m(x), & \text{if } m > 0 \\ \sqrt{2} \kappa_l^m \sin(|m|\alpha_p) P_l^{|m|}(x), & \text{if } m < 0 \\ \kappa_l^0 P_l^0(x), & \text{if } m = 0 \end{cases} \quad (13)$$

and $P_l^0(x) = P_l(x)$ is the Legendre polynomial of degree l with $x = \cos(\beta_p)$. In the development of the S^2 SHT, real spherical harmonics are used because the functions we are dealing with (namely images) are real valued. Examples of the real spherical harmonics projected onto the sphere for $l = 8$ and three different values of m are shown in Fig. 3. With a slight modification to the above development, pose detection from an aerial perspective can be achieved as discussed in [17].

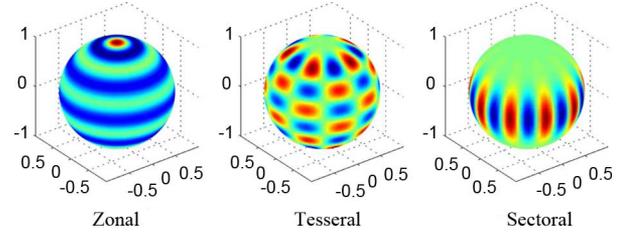


Fig. 3. Real spherical harmonics $Y_l^m(\cdot)$ for $l = 8$. The left plot is for $m = 0$ and is referred to as the zonal harmonics, the center plot is for $m = l/2$ and is referred to as the tesseral harmonics, and the right plot is for $m = l$ and is referred to as the sectoral harmonics.

C. Harmonic Analysis on $SO(3)$

In this subsection, we extend the spherical harmonic transform to the rotation group $SO(3)$ using some of the techniques developed in [50]. This extension is necessary for the development of a fully general 3-D pose estimation algorithm. For the construction of the $SO(3)$ FFT we compute the rotation of a function defined on S^2 using elements of the rotation group. The rotation group is the set of real 3×3 orthogonal matrices of determinant +1, which define proper rotations about the origin of \mathbb{R}^3 . In spectral theory, it is often the convention to define these matrices using standard $z-y-z$ Euler rotation matrices where the z -axis is the upper pole [31], [32], [51]. Therefore, any rotation $g(\alpha, \beta, \gamma) \in SO(3)$ can be written as

$$g(\alpha, \beta, \gamma) = R_z(\alpha) R_y(\beta) R_z(\gamma) \quad (14)$$

where $R_z(\alpha)$ and $R_y(\beta)$ represent a rotation about the z -axis by α radians and a rotation about the y -axis by β radians, respectively. Given any $g \in SO(3)$, we define the linear operator $\Lambda(\alpha, \beta, \gamma) : f(\alpha, \beta) \rightarrow f(\alpha', \beta')$ where (α, β) and (α', β') are the coordinates of the position vector in the original and rotated coordinate frames, respectively. The effect this has on the function in the spectral domain (i.e., the effect on the harmonic coefficient f_l^m) can be deduced from the fact that rotated versions of the spherical harmonics are simply linear combinations of harmonics of the same degree. That is

$$\begin{aligned} \Lambda(\alpha, \beta, \gamma) Y_l^m(\alpha, \beta) &= Y_l^m(\alpha', \beta') \\ &= \sum_{|m'| \leq l} Y_l^{m'}(\alpha, \beta) D_{mm'}^l(\alpha, \beta, \gamma) \end{aligned} \quad (15)$$

where $D_{mm'}^l(\alpha, \beta, \gamma)$ is the $(2l+1) \times (2l+1)$ Wigner- D matrix [31]. The D matrices satisfy the condition

$$\begin{aligned} \int_0^{2\pi} \int_0^\pi \int_0^{2\pi} D_{m_1 m_1'}^{l_1*}(\alpha, \beta, \gamma) D_{m_2 m_2'}^{l_2}(\alpha, \beta, \gamma) d\Omega \\ = \frac{8\pi^2}{2l_1 + 1} \delta_{l_1 l_2} \delta_{m_1 m_2} \delta_{m_1' m_2'} \end{aligned} \quad (16)$$

where $d\Omega = \sin(\beta) d\alpha d\beta d\gamma$, i.e., they form an orthogonal basis over $SO(3)$.

Using (15) and (16), it can be shown that a function $f \in L^2(SO(3))$ may be represented by its discrete harmonic expansion using Wigner- D matrices [41], [50]. That is

$$f(\xi_p, \gamma_r) = \sum_{l=0}^{l_{\max}} \sum_{|m| \leq l} \sum_{|m'| \leq l} f_{mm'}^l D_{mm'}^l(\xi_p, \gamma_r) \quad (17)$$

where, once again, $f(\xi_p, \gamma_r) \in [0, 1]$ is a single pixel of the image data vector $\mathbf{f}(\xi_p, \gamma_r)$. The expansion coefficients $f_{mm'}^l$ are then calculated using

$$f_{mm'}^l = \frac{4\pi}{a} \sum_{p=0}^{a-1} \sum_{r=0}^{b-1} f(\xi_p, \gamma_r) D_{mm'}^{l*}(\xi_p, \gamma_r). \quad (18)$$

In the above equations, with the unit vector ξ_p parameterized by (α_p, β_p)

$$D_{mm'}^l(\alpha_p, \beta_p, \gamma_r) = e^{-im\alpha_p} d_{mm'}^l(\beta_p) e^{-im'\gamma_r} \quad (19)$$

where $d_{mm'}^l(\beta_p)$ is known as Wigner's (small) d -matrix defined by

$$d_{mm'}^l(\beta_p) = \sqrt{\frac{(l+m')!(l-m)!}{(l+m)!(l-m)!}} \left(\sin \frac{\beta_p}{2}\right)^{m'-m} \times \left(\cos \frac{\beta_p}{2}\right)^{m+m'} P_{(m'-m, m+m')}^{(l-m')}(\cos \beta_p) \quad (20)$$

and $P_{(\cdot, \cdot)}^{(\cdot)}(x)$ is a Jacobi polynomial. For computational convenience, the d -matrices may be computed quickly using a three term recurrence relationship [31], [50], [52].

Because $f(\xi_p, \gamma_r)$ is a real-valued function, it is more convenient to work with the real-valued Wigner- D matrices denoted here as $\Delta_{mm'}^l$. The construction of rotation matrices in the basis of real spherical harmonics is discussed in [52]–[54], and can be defined as

$$\begin{aligned} \Delta_{mm'}^l &= \text{sign}(m') \Phi_m(\alpha_p) \Phi_{m'}(\gamma_r) \\ &\times \frac{d_{|m'| |m|}^l + (-1)^m d_{|m| (-|m'|)}^l}{2} \\ &- \text{sign}(m) \Phi_{-m}(\alpha_p) \Phi_{-m'}(\gamma_r) \\ &\times \frac{d_{|m'| |m|}^l - (-1)^m d_{|m| (-|m'|)}^l}{2} \end{aligned} \quad (21)$$

where

$$\Phi_m(x) = \begin{cases} \sqrt{2} \cos(mx), & \text{if } m > 0 \\ 1, & \text{if } m = 0 \\ \sqrt{2} \sin(|m|x), & \text{if } m < 0 \end{cases} \quad (22)$$

and β_p has been omitted from the Wigner- d matrices for notational convenience. The $SO(3)$ FFT may then be computed using (18) by replacing $D_{mm'}^{l*}(\xi_p, \gamma_r)$ with $\Delta_{mm'}^l(\xi_p, \gamma_r)$.

To show the harmonic extension from S^1 to S^2 and finally $SO(3)$, we use the separation of variables technique discussed in [50] and the definition of the Wigner- D matrices given in (19). Applying these to the summations of (18), the $SO(3)$ Fourier coefficients can be computed as

$$f_{mm'}^l = \frac{4\pi}{a} \sum_{\beta} d_{mm'}^l(\beta_p) \sum_{\gamma} e^{im'\gamma_r} \sum_{\alpha} e^{im\alpha_p} f(\xi_p, \gamma_r). \quad (23)$$

Note that the last summation is only concerned with data along lines of constant latitude, and is equivalent to computing the harmonic coefficients of data correlated on S^1 using the DFT. The second summation is equivalent to the special case discussed in [16], and can be computed in closed-form using FFT techniques. Finally, the first summation computes the harmonic coefficients of the co-latitudinal coordinate, thus completing the full $SO(3)$ FFT. Notice that if only the first and last summation are computed (or equivalently $m' = 0$), then the transform is equivalent to computing the discrete SHT for a function correlated on S^2 . It can be shown that by setting $m' = 0$, the Wigner- D matrices are equivalent to the spherical harmonics [31].

IV. FAST EIGENDECOMPOSITION ALGORITHM

A. Motivation

In [16], Chang *et al.* showed that for image data sets correlated in one dimension, the right singular vectors are approximately spanned by the first few low-frequency Fourier harmonics. Therefore, for image data sets correlated on S^1 , the eigendecomposition can be efficiently computed using FFT techniques [16]. Unfortunately, Chang's algorithm is not directly applicable when dealing with data sets correlated on S^2 and $SO(3)$. For example, if S^2 is sampled along lines of longitude and co-latitude, and then an FFT is performed on these dimensions independently, the samples cluster near the poles and the resulting eigendecomposition will be biased. Furthermore, for any equal-area samplings, there is no natural 2-D ordering [19]. Therefore, capturing the spectral information of a spherically correlated image data set cannot be obtained by using a 2-D FFT and instead must be done using the techniques developed in Section III [17], [18], [20].

Chang's algorithm makes use of the fact that for data correlated in one dimension, most of the energy is concentrated around the low-frequency Fourier harmonics. While this is not true for image data sets that are spherically correlated in higher dimensions, most of the energy of these image data sets is concentrated around the lower frequency spherical harmonics. An example of this is shown for object 15 from Fig. 6 where the spherical harmonic power spectra in S^2 and $SO(3)$ are plotted in Figs. 4 and 5, respectively. As can be observed in both figures, in general, as l increases, the magnitude of the power spectra decreases. As a result, the left singular vectors \tilde{U}_k of the SVD of a relatively small set of the spherically transformed harmonic images serve as excellent estimates to those of \hat{X} at a significant computational savings.

B. Algorithm Construction

Our objective is to estimate the first k principal eigenimages \tilde{U}_k of \hat{X} such that $\Delta\rho(\hat{X}, \tilde{U}_k) \leq \epsilon$, where ϵ is the user specified change in energy. The first step in computing the principal eigenimages is to construct the image data matrix \hat{X} . As mentioned in Section III, the approach taken here is to consider the object placed at the center of an imaginary unit sphere (Fig. 2) and sample $SO(3)$ appropriately using the HEALPix sampling pattern for the discretization of the sphere. The mean image is then subtracted from each sample image to form the image data matrix \hat{X} .

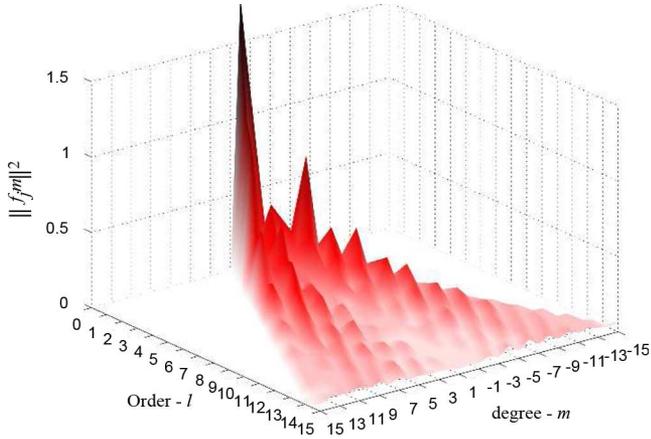


Fig. 4. Spherical harmonic power spectra $\|f_l^{mm'}\|^2$ of object 15 from Fig. 6. As can be seen from the figure, most of the energy is concentrated around the low frequency harmonics.

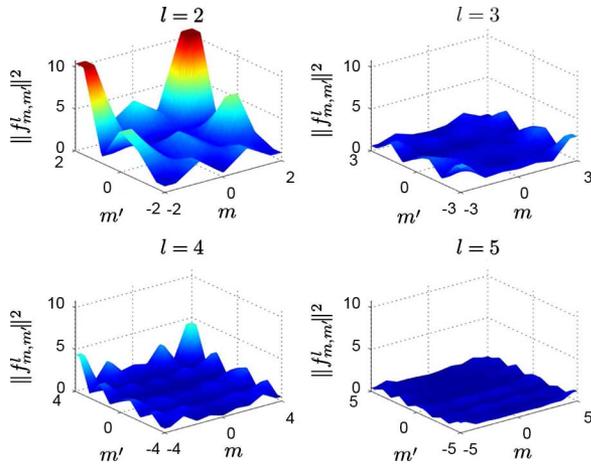


Fig. 5. Power spectra $\|f_l^{mm'}\|^2$ of object 15 from Fig. 6 for degrees $l = 2, 3, 4, 5$. As can be seen from the figure, most of the energy is concentrated around the lower frequency harmonics, i.e., those with lower values of l . (We omit $l = 0$ as this represents the mean value and is significantly higher in magnitude).

Using the HEALPix sampling pattern is based on subdividing the sphere using the parameter N_{side} , resulting in $a = 12N_{\text{side}}^2$ sample points on the sphere [48]. At each of the $12N_{\text{side}}^2$ sample points, b planar rotated images of the object are captured by rotating the camera through an angle γ_r . The number of planar rotated images captured depends on the angular resolution of the $12N_{\text{side}}^2$ samples defined on S^2 . The angular resolution is denoted θ_{pix} and calculated as $\theta_{\text{pix}} = \sqrt{3/\pi}60^\circ/N_{\text{side}}$. To maintain homogeneous sampling on $SO(3)$, $b = \lfloor 360/\theta_{\text{pix}} \rfloor$, where $\lfloor x \rfloor$ denotes the largest integer less than or equal to x . Finally, to prevent aliasing, $l_{\text{max}} = 3N_{\text{side}} - 1$ is used in the forward transform. Because the HEALPix sampling pattern is isolatitudinal, the computation of the Wigner- d matrices (which is the most computationally expensive portion of the $SO(3)$ FFT) is minimal.

Once the image data matrix \hat{X} has been constructed, we compute the matrix F whose i th row is the $SO(3)$ FFT of the i th row of \hat{X} , denoted from this point forward as $SOFT(\hat{X})$. This can be computed quickly using the methods described in

[50]; however, for small l_{max} the computational savings of this method are slim. Alternatively, $SOFT(\hat{X})$ can be cast as the matrix multiplication

$$F = \hat{X}Z \quad (24)$$

where $\hat{X} \in \mathbb{R}^{m \times n}$ is the image data matrix with the image mean removed and the images ordered in terms of (ξ_p, γ_r) as

$$\begin{aligned} \hat{X} = & [\mathbf{f}(\xi_0, \gamma_0), \mathbf{f}(\xi_0, \gamma_1), \dots, \mathbf{f}(\xi_0, \gamma_{b-1}) \\ & \mathbf{f}(\xi_1, \gamma_0), \mathbf{f}(\xi_1, \gamma_1), \dots, \mathbf{f}(\xi_1, \gamma_{b-1}), \dots \\ & \mathbf{f}(\xi_{a-1}, \gamma_0), \mathbf{f}(\xi_{a-1}, \gamma_1), \dots, \mathbf{f}(\xi_{a-1}, \gamma_{b-1})]. \end{aligned} \quad (25)$$

The matrices $\Delta_{mm'}^l(\cdot) \in \mathbb{R}^{(2l+1) \times (2l+1)}$ may be row scanned such that each row is concatenated to form the row vector $\delta^l(\cdot) = \text{vec}(\Delta_{mm'}^l(\cdot)) \in \mathbb{R}^{1 \times (2l+1)^2}$ for any given l . Using this notation, the matrix $Z \in \mathbb{R}^{n \times N_{\text{side}}(36N_{\text{side}}^2 - 1)}$ can be constructed as

$$Z = \begin{bmatrix} \delta^0(\xi_0, \gamma_0) & \dots & \delta^{3N_{\text{side}}-1}(\xi_0, \gamma_0) \\ \delta^0(\xi_0, \gamma_1) & \dots & \delta^{3N_{\text{side}}-1}(\xi_0, \gamma_1) \\ \vdots & \vdots & \vdots \\ \delta^0(\xi_0, \gamma_{b-1}) & \dots & \delta^{3N_{\text{side}}-1}(\xi_0, \gamma_{b-1}) \\ \delta^0(\xi_1, \gamma_0) & \dots & \delta^{3N_{\text{side}}-1}(\xi_1, \gamma_0) \\ \vdots & \vdots & \vdots \\ \delta^0(\xi_{a-1}, \gamma_{b-1}) & \dots & \delta^{3N_{\text{side}}-1}(\xi_{a-1}, \gamma_{b-1}) \end{bmatrix} \quad (26)$$

where the rows of Z are ordered in terms of (ξ_p, γ_r) such that the matrix product $F = \hat{X}Z$ makes sense. Note that in computing $SOFT(\hat{X})$, the matrix Z may be precomputed for several different values of l_{max} and stored for later use.

We now present an algorithm for estimating the first k principal eigenimages \tilde{U}_k of \hat{X} such that $\Delta\rho(\hat{X}, \tilde{U}_k) \leq \epsilon$, where ϵ is the user specified change in energy.

EIGENDECOMPOSITION ALGORITHM:

- 1) Form the matrix F by computing $SOFT(\hat{X})$.
- 2) Form the matrix H whose columns are the ordered columns of F in descending order according to their norm.
- 3) Set $q = \lfloor N_{\text{side}}(36N_{\text{side}}^2 - 1)[1 - (1/2)^{N+1}] \rfloor$, with $N = 0$ initially.
- 4) Construct the matrix H_q , i.e., the matrix consisting of the first q columns of H .
- 5) Compute $\text{SVD}(H_q) = \tilde{U}_q \tilde{\Sigma}_q \tilde{V}_q^T$. (The key observation here is that H_q contains q columns, which is considerably less than the n columns of \hat{X}).
- 6) If $\Delta\rho(\hat{X}, \tilde{U}_q) > \epsilon$ then set $N = N + 1$ 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 [25].
- 7) Return \tilde{U}_k such that $\Delta\rho(\hat{X}, \tilde{U}_k) \leq \epsilon$. Note that $k \leq q$.

The above algorithm takes advantage of the fact that most SVD algorithms require mn^2 flops to compute the full SVD of \hat{X} ; this is computationally prohibitive when n is large (which is the case for fully general pose estimation). Because the $SO(3)$ FFT is lossy, step one of the algorithm condenses the

energy in \hat{X} from $12bN_{\text{side}}^2$ images, where $b \approx 6N_{\text{side}}$, to $N_{\text{side}}(36N_{\text{side}}^2 - 1)$ harmonic images (roughly half) ordered in terms of the magnitude of their power spectra to form the matrix H . Furthermore, because most of the energy of \hat{X} is concentrated around the lower frequency harmonics, the matrix H_q is constructed and the SVD is performed on the first $q = \lfloor N_{\text{side}}/2(36N_{\text{side}}^2 - 1) \rfloor$ harmonic images initially, significantly reducing the computational cost. If more harmonics are required, then half of the remaining harmonic images are concatenated to H_q and the eigenspace is updated.

It is difficult to evaluate the computational complexity of the above algorithm due to the binary split used to determine k . However, if only one iteration of steps 3–6 is performed, i.e., $N = 0$, then the complexity is on the order of mq^2 flops where $q = \lfloor N_{\text{side}}/2(36N_{\text{side}}^2 - 1) \rfloor$. This is the cost of computing the SVD(H_q) in step 5 of the algorithm. It should be noted that for all of the objects tested, over 90% of the energy in \hat{X} was recovered with a single iteration. (Likewise, $\Delta\rho(\hat{X}, \hat{U}_k) < 0.01$ when k was obtained from a single iteration). Extensive simulation has shown that for fully general pose estimation, a subspace that can recover between 60–70% of the energy in \hat{X} is typically more than sufficient. A similar algorithm can be constructed for data correlated on S^2 , or the upper hemisphere of S^2 if desired, details of which can be found in [17] and [18].

V. EXPERIMENTAL RESULTS

A. Test Data

The proposed algorithm detailed in Section IV was implemented using Matlab and tested on the objects shown in Fig. 6. The parameter $N_{\text{side}} = 5$ was used, resulting in $b = 30$ and $12bN_{\text{side}}^2 = 9000$ images per object at an angular resolution $\theta_{\text{pix}} \approx 12^\circ$. The images were then both scale and intensity normalized to create the image data matrix X . The “unbiased” image data matrix \hat{X} was then constructed by subtracting the mean image from the image data matrix X . Finally, the matrix F was computed, condensing the image data set from 9000 images to 4495 harmonic images. The true SVD(\hat{X}) was computed using Matlab for a comparison.

B. Quality of Subspace Estimation

To validate the quality of the estimated eigenimages, the quality measures outlined in Section II were used. Fig. 7 shows all three of the quality measures averaged across all objects in Fig. 6. The top plot shows the energy recovery ratio for the first 50 true eigenimages \hat{U}_{50} as well as the first 50 estimated eigenimages \tilde{U}_{50} . As can be seen from the plot, the two curves are nearly indistinguishable. The second plot shows the residue Δ of the first 50 estimated eigenimages compared to the first 50 true eigenimages. The residue is normalized by $\sqrt{2k}$ resulting in a worst case bound of one. The third plot shows that based on the subspace criterion, the first 50 estimated eigenimages nearly span the same space spanned by the first 50 true eigenimages. As can be seen from all three measures, quantitatively, the estimated left singular vectors \tilde{U}_k for the proposed algorithm are very good estimates of the true left singular vectors as computed by the direct SVD.

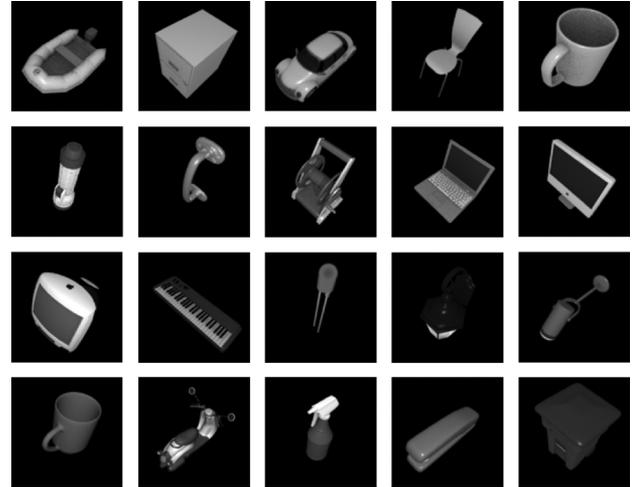


Fig. 6. Example ray-traced images of CAD models courtesy of Kator Legaz [55]. Each object is sampled using the HEALPix sampling pattern at a resolution of 128×128 . Each of the above images is then scale and intensity normalized. The objects are ordered from left to right, then top to bottom.

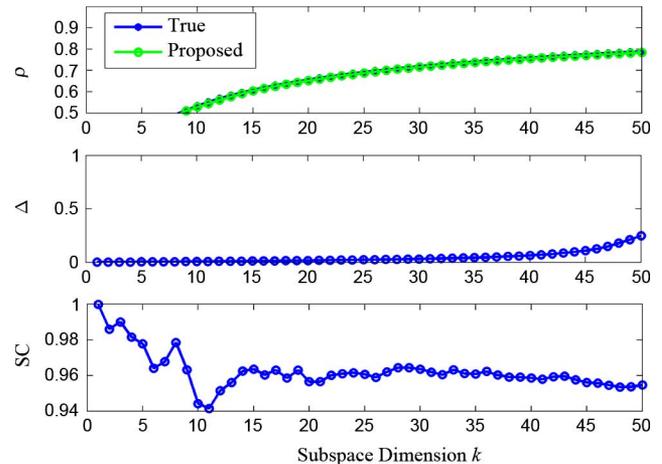


Fig. 7. Quality measures outlined in Section II averaged across all objects in Fig. 6. As is apparent from the figure, the estimated left singular vectors \tilde{U}_k for the proposed algorithm are very good approximations to the true left singular vectors as computed by the direct SVD.

To illustrate qualitatively the accuracy of the estimated eigenspace, the first seven eigenimages of object 8 from Fig. 6 were plotted and depicted in Fig. 8 as an example. The true eigenimages as computed by the direct SVD are depicted in the top row, while the bottom row depicts the eigenimages as computed by the proposed algorithm. As is apparent from a visual inspection of the figure, the first few eigenimages computed by the proposed algorithm are nearly identical to those computed by the SVD.

C. Computational Savings

Table II shows the required subspace dimension k , the amount of energy recovered at this subspace dimension, and the time required to estimate the first k left singular vectors \tilde{U}_k for each object depicted in Fig. 6 to meet the user specified change in energy $\epsilon = 0.01$. The user specified change in energy $\epsilon = 0.01$ implies that the k th eigenimage recovers less than

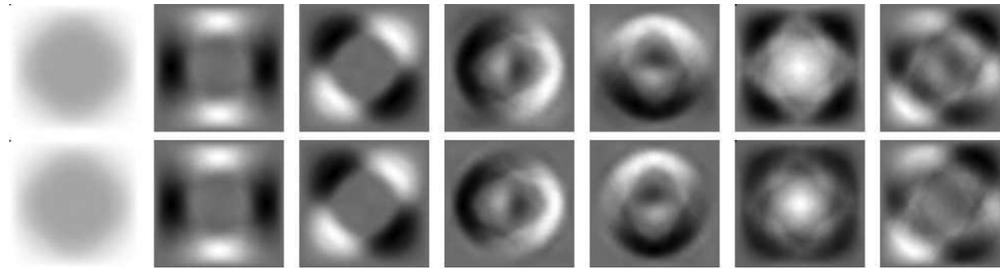


Fig. 8. First seven eigenimages of object 8 from Fig. 6. The true eigenimages as computed by the SVD using Matlab are shown in the top row, and the eigenimages as computed by the proposed algorithm are shown in the bottom row. As is apparent from the figure, the first few eigenimages computed by the proposed algorithm are nearly identical to those computed by the SVD.

TABLE II
ACCURACY AND COMPUTATIONAL EXPENSE OF THE PROPOSED ALGORITHM VERSUS THE TRUE SVD AS COMPUTED USING MATLAB (\bar{k}_{\min} IS THE AVERAGE SUBSPACE DIMENSION REQUIRED TO DETERMINE THE POSE OF 72 TEST IMAGES)

Obj.	k	\bar{k}_{\min}	Time [hours]		Energy ρ [%]		
			True	Prop.	True	Prop.	At k_{\min}
1	19	5.9	3.550	0.082	69.14	66.94	37.86
2	17	5.7	3.432	0.076	68.39	68.22	47.96
3	16	6.0	3.544	0.083	64.80	63.42	43.24
4	16	4.9	3.558	0.124	54.68	54.30	29.34
5	14	6.0	3.589	0.076	63.78	62.53	43.87
6	14	6.2	3.322	0.076	76.66	76.33	60.41
7	19	5.1	2.424	0.092	60.40	60.13	33.22
8	16	5.9	3.688	0.101	45.65	39.83	24.20
9	16	5.6	3.678	0.084	60.12	58.83	41.72
10	21	5.9	3.634	0.086	58.05	59.63	33.27
11	20	5.4	3.584	0.083	64.76	64.44	31.33
12	15	5.3	3.598	0.096	57.61	56.21	39.16
13	10	5.9	3.403	0.102	50.65	42.31	29.77
14	15	5.0	3.627	0.091	64.35	64.25	43.20
15	16	5.8	3.393	0.093	66.43	64.62	43.05
16	14	5.7	3.881	0.075	70.77	70.62	56.15
17	17	5.7	3.617	0.095	60.02	59.80	38.47
18	15	5.9	3.599	0.087	68.43	65.72	47.96
19	8	5.3	3.542	0.079	71.32	66.46	54.16
20	9	6.4	4.039	0.077	64.23	57.00	51.24
Mean			3.535	0.088	63.01	61.08	50.04
Min.			2.424	0.075	45.65	39.83	47.25
Max.			4.039	0.124	76.66	76.33	51.24

1% additional energy when added to the subspace spanned by the first $k - 1$ eigenimages. This result is then compared to the true SVD. Both the proposed algorithm and the true SVD were computed on the same machine using MATLAB. As is apparent from the table, the proposed algorithm is capable of recovering nearly the same amount of energy as the direct SVD for the given subspace dimension at a significant computational savings. Only the first split $N = 0$ (step 3) of the algorithm was used to compute this data, resulting in an average speed-up factor of 40.

D. Accuracy of Pose Estimation

To validate that the subspace dimension k (given in Table II) as determined by the proposed algorithm is sufficient to accurately determine the 3-D pose of arbitrary objects, each of the objects shown in Fig. 6 was sampled on a different $SO(3)$ grid resulting in 72 images per object that were not included in the original data matrix \hat{X} . For each of the 72 different poses, the subspace dimension k required to accurately estimate the correct pose was determined using the estimated eigenimages as

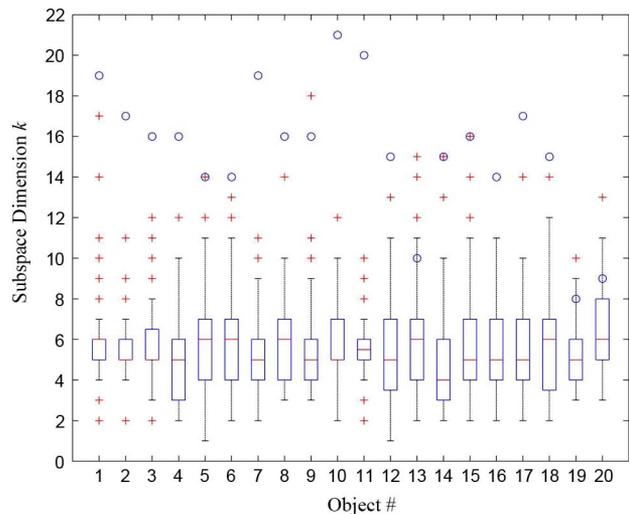


Fig. 9. Box plots showing the distribution of the subspace dimension k required to determine 72 different test poses for each object in Fig. 6. The plus signs depict outliers in the data, and the circles depict the subspace dimension calculated using $\Delta\rho(\hat{X}, \tilde{U}_k) \leq 0.01$ as shown in Table II.

computed by the proposed algorithm. Fig. 9 shows the distribution of the subspace dimension for each of the 20 object in Fig. 6. The boxes show the inner-quartile region with the horizontal bar representing the median of the data. The plus signs are the outliers in the data and represent poses that were either easier or more difficult to determine as compared to the rest of the distribution. The circles represent the subspace dimension as computed using $\Delta\rho(\hat{X}, \tilde{U}_k) \leq 0.01$ (given in Table II). As is apparent from the figure, with the exception of objects 9, 19, and 20, the subspace dimension as determined by computing $\Delta\rho(\hat{X}, \tilde{U}_k)$ is sufficient to accurately determine all 72 test poses. Furthermore, only one test pose for objects 9, 19, and 20 required a subspace dimension higher than that determined by $\Delta\rho(\hat{X}, \tilde{U}_k)$. These outliers are due to the fact that these poses lie nearly equi-distant to two adjacent poses in the original image data matrix \hat{X} . Therefore, an interpolation procedure between p of the closest matching poses would resolve this issue, even with a much smaller subspace dimension k .

Table II also shows the average subspace dimension \bar{k}_{\min} required to determine the pose of each of the 72 above mentioned test images, as well as the energy recovered at this dimension. As shown in the table, the average energy recovered at \bar{k}_{\min} is

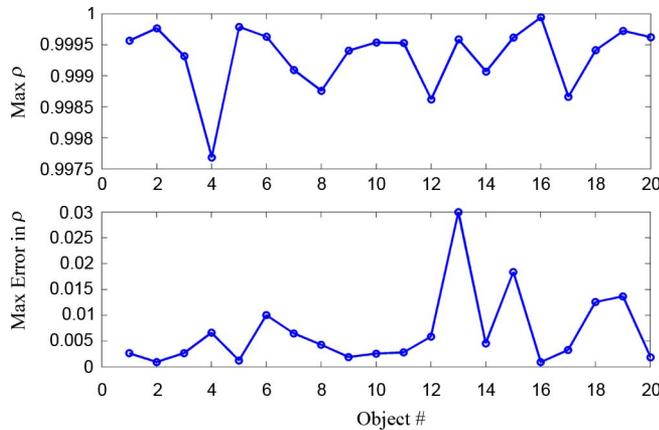


Fig. 10. Maximum achievable energy recovery ratio ρ for each object in Fig. 6 [top], and the maximum difference in energy recovery between the first $N_{\text{side}}(36N_{\text{side}}^2 - 1)$ true left singular vectors of $\text{SVD}(\hat{X})$ and the estimated left singular vectors $\hat{U}_{N_{\text{side}}(36N_{\text{side}}^2 - 1)}$ as computed by the proposed algorithm [bottom].

50.04% with the min and max being 47.25% and 51.24%, respectively. This data supports the claim in Section IV that a subspace that can recover between 60–70% of the energy in \hat{X} is typically sufficient for accurate pose estimation.

E. Maximum Energy Recovery

Based on the data provided in Table II, and Figs. 7–9, it is obvious that the proposed algorithm is capable of estimating the required subspace dimension and left singular vectors of \hat{X} at a significant computational savings for use in pose estimation. However, because the $SO(3)$ FFT is lossy, it is important to quantify the maximum amount of energy recovery possible using this approach. To do this, all $N_{\text{side}}(36N_{\text{side}}^2 - 1)$ harmonic images were used to estimate the left singular vectors of \hat{X} . Using these estimates, the energy recovery ratio $\rho(\hat{X}, \hat{U}_{N_{\text{side}}(36N_{\text{side}}^2 - 1)})$ was computed for each object in Fig. 6. The top plot in Fig. 10 shows the maximum amount of energy recovered per object. As seen in the plot, over 99.7% of the energy in \hat{X} is recoverable for all objects using the proposed algorithm. The bottom plot in Fig. 10 shows the maximum difference in energy recovery per object if only the first $N_{\text{side}}(36N_{\text{side}}^2 - 1)$ true left singular vectors are used to recover the same energy, i.e., the bottom plot shows $\max(\rho(\hat{X}, \hat{U}_k) - \rho(\hat{X}, \hat{U}_k))$ for all $k \leq N_{\text{side}}(36N_{\text{side}}^2 - 1)$ for each object. As seen in this plot, the maximum error one could expect to see is less than 3% across the entire subspace for the objects shown in Fig. 6.

VI. CONCLUSION

This paper presented a computationally efficient algorithm for estimating the subspace of images correlated in multiple dimensions for use in fully general 3-D pose estimation. This required a method of sampling objects to produce a spherical data set by sampling S^2 and $SO(3)$ appropriately. It was illustrated that for data correlated on S^2 , the spherical harmonic transform can be used to calculate the frequency information of this spherical data set. If the data is correlated on $SO(3)$,

then spherical harmonics in conjunction with Wigner- D matrices can be used to capture the frequency information of this data set. It was shown that computing the eigendecomposition on a subset of harmonic images in the transform domain can significantly reduce the computational burden. Experimental results were presented to validate both qualitatively and quantitatively the accuracy of the estimation, as well as the significant computational savings. Furthermore, the maximum energy recovery when using the spherical harmonic transform was quantified. It was shown that the errors associated with the proposed algorithm are minimal with respect to computational savings.

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 Recognit.*, vol. 14, no. 1–6, pp. 375–381, Dec. 1981.
- [3] M. Kirby and L. Sirovich, "Application of the Karhunen-Loeve procedure for the characterization of human faces," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 1, pp. 103–108, Jan. 1990.
- [4] M. Turk and A. Pentland, "Eigenfaces for recognition," *J. Cogn. Neurosci.*, vol. 3, no. 1, pp. 71–86, Mar. 1991.
- [5] P. N. Belhumeur, J. P. Hespanha, and D. J. Kriegman, "Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 19, no. 7, pp. 711–720, Jul. 1997.
- [6] R. Brunelli and T. Poggio, "Face recognition: Features versus templates," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 15, no. 10, pp. 1042–1052, Oct. 1993.
- [7] A. Pentland, B. Moghaddam, and T. Starner, "View-based and modular eigenspaces for face recognition," in *Proc. IEEE Conf. Comp. Vis. and Patt. Rec.*, Seattle, WA, Jun. 1994, pp. 84–91.
- [8] M. H. Yang, D. J. Kriegman, and N. Ahuja, "Detecting faces in images: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 1, pp. 34–58, Jan. 2002.
- [9] H. Murase and R. Sakai, "Moving object recognition in eigenspace representation: Gait analysis and lip reading," *Pattern Recognit. Lett.*, vol. 17, no. 2, pp. 155–162, Feb. 1996.
- [10] G. Chiou and J. N. Hwang, "Lipreading from color video," *IEEE Trans. Image Process.*, vol. 6, no. 8, pp. 1192–1195, Aug. 1997.
- [11] H. Murase and S. K. Nayar, "Illumination planning for object recognition using parametric eigenspaces," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 16, no. 12, pp. 1219–1227, Dec. 1994.
- [12] C. Y. Huang, O. I. Camps, and T. Kanungo, "Object recognition using appearance-based parts and relations," in *Proc. IEEE Conf. Comp. Vis. and Patt. Rec.*, San Juan, PR, Jun. 1997, pp. 877–883.
- [13] R. J. Campbell and P. J. Flynn, "Eigenshapes for 3-D object recognition in range data," in *Proc. IEEE Conf. Comp. Vis. and Patt. Rec.*, Jun. 1999, pp. 505–510.
- [14] H. Murase and S. K. Nayar, "Visual learning and recognition of 3-D objects from appearance," *Int. J. Comput. Vis.*, vol. 14, no. 1, pp. 5–24, Jan. 1995.
- [15] S. K. Nayar, S. A. Nene, and H. Murase, "Subspace methods for robot vision," *IEEE Trans. Robot. Autom.*, vol. 12, no. 10, pp. 750–758, Oct. 1996.
- [16] C. Y. Chang, A. A. Maciejewski, and V. Balakrishnan, "Fast eigenspace decomposition of correlated images," *IEEE Trans. Image Process.*, vol. 9, no. 11, pp. 1937–1949, Nov. 2000.
- [17] R. C. Hoover, A. A. Maciejewski, and R. G. Roberts, "Aerial pose detection of 3-D objects using hemispherical harmonics," in *Proc. IEEE SSIAT*, Santa Fe, NM, Mar. 2008, pp. 157–160.
- [18] 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 *Proc. IEEE Int. Conf. Robot. Automat.*, Pasadena, CA, May 2008, pp. 993–998.
- [19] K. Saitwal, A. A. Maciejewski, and R. G. Roberts, "Computationally efficient eigenspace decomposition of correlated images characterized by three parameters," *Pattern Anal. Apps.*, 2009.
- [20] 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 *Proc. IEEE Conf. Automat. Sci. and Engr.*, Washington, DC, Aug. 2008, pp. 47–52.

- [21] E. N. Malamas, E. G. Petrakisa, M. Zervakisa, L. Petit, and J.-D. Legat, "A survey on industrial vision systems, applications and tools," *Image Vis. Comput.*, vol. 21, no. 2, pp. 171–188, 2003.
- [22] H. Murakami and V. Kumar, "Efficient calculation of primary images from a set of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 4, no. 5, pp. 511–515, Sep. 1982.
- [23] 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. Speech Signal Process.*, vol. 37, no. 10, pp. 1550–1556, Oct. 1989.
- [24] 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.
- [25] S. Chandrasekaran, B. Manjunath, Y. Wang, J. Winkeler, and H. Zhang, "An eigenspace update algorithm for image analysis," *CVGIP: Graphic Models Image Process.*, vol. 59, no. 5, pp. 321–332, Sep. 1997.
- [26] G. H. Golub and C. F. V. Loan, *Matrix Computations*. Baltimore, MD: Johns Hopkins Univ. Press, 1996.
- [27] 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 Process.*, vol. 15, no. 8, pp. 2376–2387, Aug. 2006.
- [28] P. N. Swartztrauber, "The vector harmonic transform method for solving partial differential equations in spherical geometry," *Month. Weather Rev.*, vol. 121, pp. 3415–3437, May 1993.
- [29] M. A. Wiecezorek and F. J. Simons, "Localized spectral analysis on the sphere," *Geophys. J. Int.*, vol. 162, no. 3, pp. 655–675, Sep. 2005.
- [30] F. J. Simons, F. A. Dahlen, and M. A. Wiecezorek, "Spatiospectral localization on a sphere," *SIAM Rev.*, vol. 48, no. 3, pp. 504–536, 2006.
- [31] D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonskii, *Quantum Theory of Angular Momentum*. Hackensack, NJ: World Scientific, 1988.
- [32] G. B. Arfken, *Mathematical Methods for Physicists*. New York, NY: Academic, 1970.
- [33] G. S. Chirikjian and A. B. Kyatkin, *Engineering Applications of Non-commutative Harmonic Analysis: With Emphasis on Rotation and Motion Groups*. Boca Raton, FL: CRC, 2001.
- [34] R. Ramamoorthi, "Analytic PCA construction for theoretical analysis of lighting variability in images of a Lambertian object," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 10, pp. 1322–1333, Oct. 2002.
- [35] L. Qing, S. Shan, and W. Gao, "Eigen-harmonic faces: Face recognition under generic lighting," in *Proc. IEEE Int. Conf. Auto. Face Gesture Rec.*, Seoul, Korea, May 2004, pp. 296–301.
- [36] S. Romdhani, J. Ho, T. Vetter, and D. J. Kriegman, "Face recognition using 3-D models: Pose and illumination," *Proc. IEEE*, vol. 294, no. 11, pp. 1977–1999, Nov. 2006.
- [37] L. Zhang and D. Samaras, "Face recognition from a single training image under arbitrary unknown lighting using spherical harmonics," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 3, pp. 351–363, Mar. 2006.
- [38] B. Bustos, D. A. Keim, D. Saupe, T. Schreck, and D. V. Vranic, "Feature-based similarity search in 3-D object databases," *ACM Comput. Surv.*, vol. 37, no. 4, pp. 345–387, Dec. 2005.
- [39] D. Saupe and D. V. Vranic, "3-D model retrieval with spherical harmonics and moments," in *Proc. DAGM-Symp. Pattern Recognit.*, Munich, Germany, Sep. 2001, pp. 392–397.
- [40] D. V. Vranic, "An improvement of rotation invariant 3D-shape based on functions on concentric spheres," in *Proc. IEEE Conf. Image Process.*, Barcelona, Spain, Sep. 2003, pp. 757–760.
- [41] A. Makadia and K. Daniilidis, "Rotation recovery from spherical images without correspondences," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 7, pp. 1170–1175, Jul. 2006.
- [42] N. Sneeuw, Global Spherical Harmonic Analysis by Least Squares and Numerical Quadrature Methods in Historical Perspective 2007 [Online]. Available: <http://www.geomatics.ucalgary.ca/~sneeuw/publ/neumann.ps>
- [43] 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.
- [44] P. N. Swartztrauber and W. F. Spatz, "Generalized discrete spherical harmonic transforms," *J. Comput. Phys.*, vol. 159, no. 2, pp. 213–230, Apr. 2000.
- [45] P. N. Swartztrauber, "On computing the points and weights for Gauss-Legendre quadrature," *SIAM J. Sci. Comput.*, vol. 24, no. 3, pp. 945–954, 2002.
- [46] J. R. Driscoll and D. M. Healy, Jr., "Computing Fourier transforms and convolutions on the 2-sphere," *Adv. Appl. Math.*, vol. 15, no. 2, pp. 202–250, Jun. 1994.
- [47] D. M. Healy, Jr., D. Rockmore, P. Kostelec, and S. Moore, "FFTs for the 2-sphere-improvements and variations," *J. Fourier Anal. Appl.*, vol. 9, no. 4, pp. 341–385, Jul. 2003.
- [48] K. M. Górski, E. Hivon, A. L. Banday, 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," *Astrophys. J.*, vol. 622, pp. 759–771, Apr. 2005.
- [49] 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 *Proc. IEEE SSIAP*, Santa Fe, NM, Mar. 2008, pp. 41–44.
- [50] P. J. Kostelec and D. N. Rockmore, "FFTs on the rotation group," presented at the Santa Fe Institute Working Papers Series Paper #03–11-060, 2003.
- [51] E. P. Wigner, *Group Theory and its Application to the Quantum Mechanics of Atomic Spectra*. New York, NY: Academic, 1959.
- [52] M. Blanco, M. Florez, and M. Bermejo, "Evaluation of the rotation matrices in the basis of real spherical harmonics," *J. Molec. Structure: THEOCHEM*, vol. 419, pp. 19–27, Dec. 1997.
- [53] J. Ivancic and K. Ruedenberg, "Rotation matrices for real spherical harmonics. Direct determination by recursion," *J. Phys. Chem.*, vol. 100, no. 15, pp. 6342–6347, Apr. 1996.
- [54] T. Risbo, "Fourier transform summation of Legendre series and D-functions," *J. Geodesy*, vol. 70, no. 7, pp. 383–396, 1996.
- [55] K. Legaz, Kator Legaz: 3-D Model Database for Blender 2007 [Online]. Available: http://www.katorlegaz.com/3d_models/index.php

Randy C. Hoover (S'06) received the B.S. degree in electrical engineering and M.S. degree in measurement and control engineering from Idaho State University in 2002 and 2004, respectively, and the Ph.D. degree in electrical engineering from Colorado State University, Fort Collins, in 2009.

He is currently an Assistant Professor in the Department of Mathematics and Computer Science, South Dakota School of Mines and Technology, Rapid City. His research interests are in the areas of computer vision, robotics, and controls. Dr. Hoover was a National Science Foundation Fellow from 2003–2004.

Anthony A. Maciejewski (F'05) received the B.S.E.E., M.S., and Ph.D. degrees from The Ohio State University, Columbus, in 1982, 1984, and 1987, respectively.

From 1988 to 2001, he was a Professor of electrical and computer engineering at Purdue University, West Lafayette, IN. He is currently the Department Head of Electrical and Computer Engineering, Colorado State University, Fort Collins. A complete vita is available at <http://www.engr.colostate.edu/~aam>.

Rodney G. Roberts (SM'02) received the B.S. degrees in electrical engineering and mathematics from the Rose-Hulman Institute of Technology, Terre Haute, IN, in 1987, and the M.S.E.E. and Ph.D. degrees in electrical engineering from Purdue University, West Lafayette, IN, in 1988 and 1992, respectively.

From 1992 to 1994, he was a National Research Council Fellow at Wright Patterson Air Force Base, Dayton, OH. Since 1994, he has been with Florida A & M University, Florida State University College of Engineering, Tallahassee, where he is currently a Professor of electrical and computer engineering. His research interests are in the areas of robotics and image processing.