Pose Estimation of Spherically Correlated Images Using Eigenspace Decomposition in Conjunction with Spectral Theory

Randy C. Hoover

Electrical and Computer Engineering
Colorado State University

Adviser: Dr. Anthony A. Maciejewski

Committee Members:
Dr. Rodney G. Roberts, Dept. of ECE - Florida A & M - Florida State University
Dr. Edwin Chong, Depts. of ECE & Mathematics
Dr. Peter M. Young, Dept. of ECE
Dr. Chris Peterson, Dept. of Mathematics

Outline

- Motivation
- Mathematical preliminaries
- First part of my thesis (variation in pose):
  - Recap to preliminary exam
  - Introduction to spherical harmonics and Wigner-D functions
  - Eigenspace decomposition on $SO(3)$
- Second part of my thesis (variation in illumination and pose):
  - Motivation and preliminaries
  - Illumination variation (single source)
    - spherical harmonics
  - Illumination variation (multiple sources)
  - Eigenspace decomposition with variation in illumination and pose
  - Pose estimation and manifold analysis
- Conclusions

What is Computer Vision and Machine Intelligence?

- The ability for a machine to interpret $3-D$ information from a $2-D$ image
- Current focus:
  - Pose Estimation
    - Eigenspace decomposition algorithm development
    - Determining the spatial orientation of known object
- Applications:
  - Automated assembly
  - Automated part inspection
  - Human-Robot interaction
  - National security

How is this done??? - Training

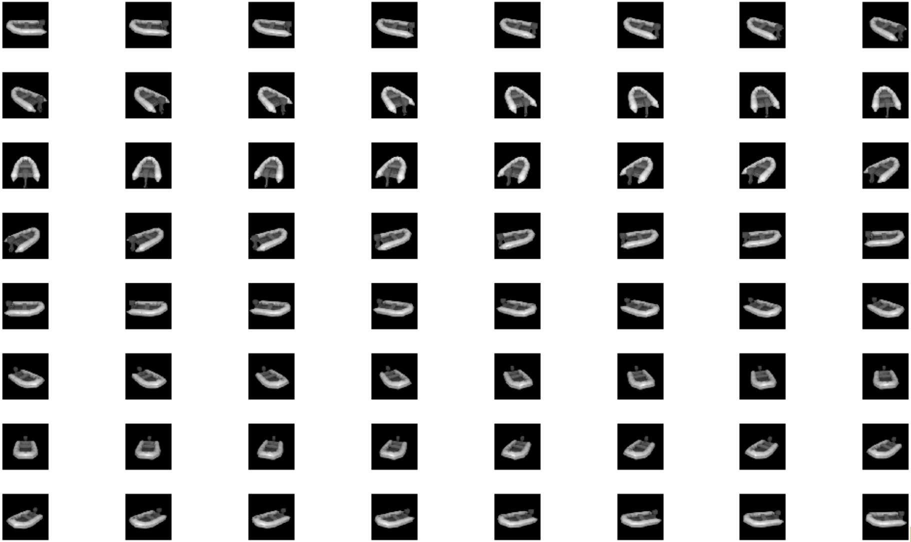

Too Much Data - Eigenspace Decomposition

- Dimensionality reduction
  - Exploit correlation between images
  - Represent by a smaller subspace
- Advantages
  - Computationally efficient on-line
  - Works well on a variety of applications
  - Strictly appearance based
    - No feature extraction
    - No edge detection
- Drawbacks
  - Background clutter/occlusion
  - Large number of training images required
  - Variation in illumination
  - Computationally expensive off-line

Preliminaries

- Gray-scale images $\mathcal{X} \in [0,1]^{h\times v}$
  - Row-scanned $\mathbf{f} = \mathrm{vec}(\mathcal{X}^T)$
  - Sets of related images $X = [\mathbf{f}_1, \cdots, \mathbf{f}_n]$
  - Subtract the mean image to get $\hat{X}$ (unbiased image data matrix)
- $\mathrm{SVD}(\hat{X}) = \hat{U}\hat{\Sigma}\hat{V}^T$
  - $\hat{U}$-left singular vectors (eigenimages)
  - $\hat{\Sigma}$-diagonal matrix of singular values
  - $\hat{V}$-right singular vectors
    - eigenimages are the eigenvectors of $\hat{X}\hat{X}^T = \hat{U}\hat{\Sigma}^2\hat{U}$
    - singular values measure how "important" each eigenimage is
    - right singular vectors measure how aligned each image is with the corresponding eigenimage

- Compute $\mathcal{M} = \hat{U}_k^T X$
  - $\mathcal{M} \rightarrow$ discrete approximation to a 1-dimensional manifold in k-dimensional space (consisting of n points)
- On-line computation consists of dot products ($\mathcal{P} = \hat{U}_k^T \mathbf{f}_{\text{new}}$) and a closest point search in high-dimensional space

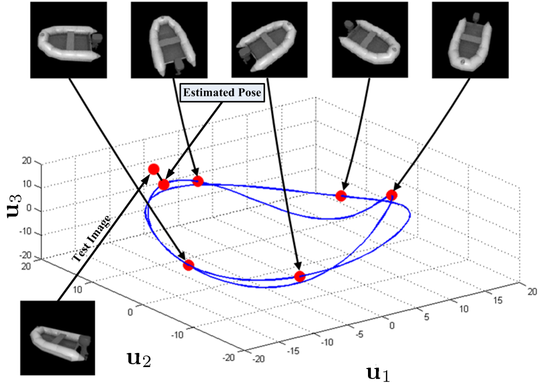

General Idea

- Compute $\mathcal{M}_k = \hat{U}_k^T \hat{X}$
- Compute $\mathcal{P} = \hat{U}_k^T \mathbf{f}_{\text{new}}$
- Search the eigenspace
- Biggest issue is computing the subspace $\hat{U}_k$
  - Computationally prohibitive when m and n are large

First Part of Dissertation - Pose Estimation (Ambient Illumination)

Objective

- Accurately estimate the first k principal eigenimages $\hat{U}_k$ of $\hat{X}$

Fully general 3-dimensional pose estimation

- Correlation in three-dimensions
  - Representative sampling ($SO(3)$)
  - Spherical harmonics in conjunction with Wigner-D matrices
  - Exploit spherical correlation
    - Eigenspace decomposition in transform domain to estimate $\hat{U}$

Quality Measures

Energy recovery

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

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

Quality Measures

Energy recovery

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

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

Two subspaces span the same space - Subspace criterion (SC)

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

Quality Measures

Energy recovery

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

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

Two subspaces span the same space - Subspace criterion (SC)

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

Rotation between subspaces - Residue

$$\Delta = \min_{Q} ||\hat{U}_k - \hat{\tilde{U}}_k Q||_F$$

- Compute $SVD(\hat{\tilde{U}}_k^T \hat{U}_k) = U_c \Sigma_c V_c^T$
- $\Delta^2 = 2\left(k - \sum_{i=1}^{k} \sigma_{ci}\right)$

Spherical Sampling

- Sampling on $SO(3)$ $\mathbf{f} = \mathbf{f}(\boldsymbol{\xi}_p, \gamma_r)$
  - $\boldsymbol{\xi}_p$: $p \in \{0, \ldots, a-1\}$ is the unit vector
  - $\beta_p \in (0, \pi)$ - angle of co-latitude
  - $\alpha_p \in [0, 2\pi)$ - angle of longitude
    - parameterization of the sphere
  - $r \in \{0, \ldots, b-1\}$ is the r^{th} planar rotation $\gamma_r \in [0, 2\pi)$ at sample p
- Hierarchical Equal Area isoLatitude Pixelization (HEALPix)

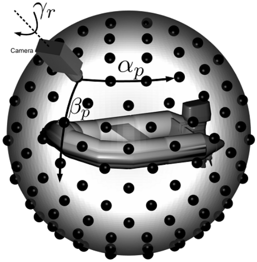

Harmonic Analysis on S^2

$f(\xi_p) \in S^2$ - Spherical harmonic transform

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

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

Spherical Harmonics

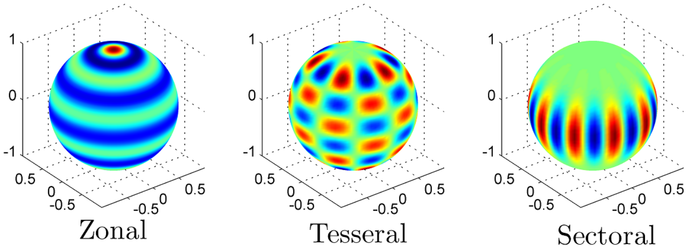

Spherical harmonics

$$Y_l^m(\xi_p) = \kappa_l^m P_l^m(\cos(\beta_p))e^{jm\alpha_p}$$

$P_l^m(\cos(\beta_p))$ associated Legendre polynomial

$$\kappa_l^m = \sqrt{\frac{2l+1}{4\pi}\frac{(l-|m|)!}{(l+|m|)!}}, \ |m| \leq l < l_{\max}$$

Harmonic Analysis on $SO(3)$

Rotation of spherical harmonics

$$\Lambda(\alpha,\beta,\gamma)Y_l^m(\alpha,\beta) = Y_l^m(\alpha',\beta') = \sum_{|m|\le l} Y_l^m(\alpha,\beta)D^l_{mm'}(\alpha,\beta,\gamma)$$

- $D^l_{mm'}(\alpha,\beta,\gamma)$ is the $(2l+1)\times(2l+1)$ Wigner-D matrix
  $D^l_{mm'}(\alpha_p,\beta_p,\gamma_r) = e^{-im\alpha_p}d^l_{mm'}(\beta_p)e^{-im'\gamma_r}$
  - $d^l_{mm'}(\beta_p)$ - Wigner's small-d matrix (related to the Jacobi polynomials)

$f(\xi_p,\gamma_r) \in SO(3)$ - $SO(3)$ harmonic transform

$$f(\xi_p,\gamma_r) = \sum_{l=0}^{l_{\max}} \sum_{|m|\le l} \sum_{|m'|\le l} f^l_{mm'} D^l_{mm'}(\xi_p,\gamma_r)$$

$$f^l_{mm'} = \frac{4\pi}{a}\sum_{p=0}^{a-1}\sum_{r=0}^{b-1} f(\xi_p,\gamma_r)D^{l*}_{mm'}(\xi_p,\gamma_r)$$

$SO(3)$ Harmonic Power Spectra

- Observation (1)
  - Low-frequency harmonics
- Consequence
  - SVD of low freq. harmonics of $\hat{X}$ is a good estimate of SVD($\hat{X}$)
- Observation (2)
  - Transform is lossy
    - harmonic images $\approx 1/2$ samples
- Consequence
  - 100% energy recovery not possible (low pass filtered)
  - Bad for compression
  - Good for computational savings

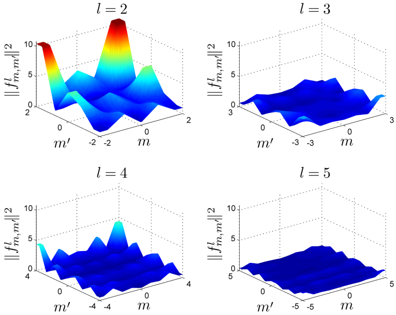

Algorithm Summary

Eigenspace Decomposition Algorithm on $SO(3)$

1. Compute the matrix F whose i^{th} row is the $SO(3)$ FFT of the i^{th} row of $\hat{X}$.
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 which is the matrix consisting of the first q columns of H.
5. Compute $\text{SVD}(H_q) = \tilde{\hat{U}}_q \tilde{\hat{\Sigma}}_q \tilde{\hat{V}}_q^T$.
6. If $\Delta\rho(\hat{X}, \tilde{\hat{U}}_q) > \epsilon$. Let $N = N + 1$ and repeat Steps 3 through 6.
7. Return $\tilde{\hat{U}}_k$ such that $\Delta\rho(\hat{X}, \tilde{\hat{U}}_k) \leq \epsilon$.

Publications Resulting From First Part

1. R. C. Hoover, A. A. Maciejewski, and R. G. Roberts, "Pose Estimation from Images Correlated on S^1, S^2, and $SO(3)$ Using Eigendecomposition in Conjunction with Spectral Theory," accepted to appear in *IEEE Transactions on Image Processing*, 2009.
2. 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," *2008 IEEE International Conference on Robotics and Automation*, pp. 993-998, Pasadena, CA, May 19-23, 2008.
3. R. C. Hoover, A. A. Maciejewski, and R. G. Roberts, "Aerial Pose Detection of 3-D Objects Using Hemispherical Harmonics," *2008 IEEE Southwest Symposium on Image Analysis and Interpretation*, pp. 157-160, Santa Fe, NM, March 24-26, 2008.
4. 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," *2008 IEEE Southwest Symposium on Image Analysis and Interpretation*, pp. 41-44, Santa Fe, NM, March 24-26, 2008.
5. 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," *4th Annual IEEE Conference on Automation Science and Engineering*, pp. 47-52, Washington, DC, August 23-26, 2008.

Second Part of Dissertation - Variation in Illumination and Pose

- Motivation:
  - Most objects are illuminated from unknown directions
  - Multiple sources of illumination can exist
  - Eigenspace decomposition is appearance based

Second Part of Dissertation - Variation in Illumination and Pose

- Motivation:
  - Most objects are illuminated from unknown directions
  - Multiple sources of illumination can exist
  - Eigenspace decomposition is appearance based
  - Objects from the same pose under different illumination can appear considerably different

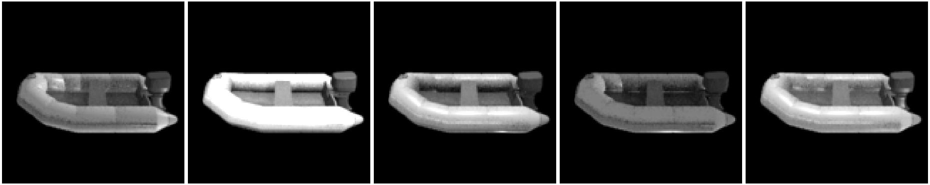

Second Part of Dissertation - Variation in Illumination and Pose

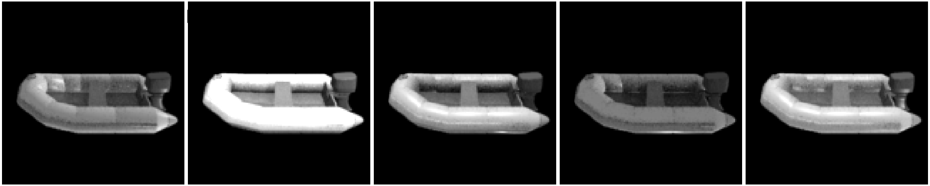

- Contributions:
  - Develop an algorithm to efficiently estimate the principle eigenimages when variations in illumination and pose exist
  - Evaluate the effects of multiple illumination sources
  - Propose a method to efficiently estimate the pose of objects when variations in pose and illumination conditions from multiple sources exist

Image Acquisition

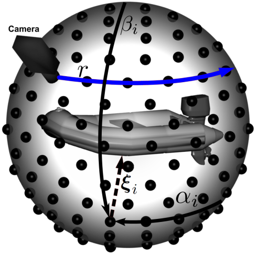

- Treat S^2 as an illumination sphere
  - $\boldsymbol{\xi}_i$: $i \in \{0, \ldots, b-1\}$
    - i^{th} illumination direction
  - $\beta_i \in (0, \pi)$ - angle of co-latitude
  - $\alpha_i \in [0, 2\pi)$ - angle of longitude
- Sampling on $S^1 \rightarrow \mathbf{f} = \mathbf{f}(\boldsymbol{\xi}_i, r)$
  - $r \in \{0, \ldots, a-1\}$ is the r^{th} pose of the object
- Image data matrix:

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

Image Acquisition

- Treat S^2 as an illumination sphere
  - $\boldsymbol{\xi}_i$: $i \in \{0, \ldots, b-1\}$
    - i^{th} illumination direction
  - $\beta_i \in (0, \pi)$ - angle of co-latitude
  - $\alpha_i \in [0, 2\pi)$ - angle of longitude
- Sampling on $S^1 \rightarrow \mathbf{f} = \mathbf{f}(\boldsymbol{\xi}_i, r)$
  - $r \in \{0, \ldots, a-1\}$ is the r^{th} pose of the object
- Image data matrix:

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

Harmonic Analysis on S^2 - Revisited

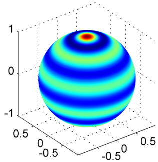

Zonal

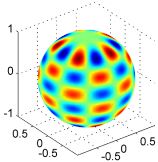

Tesseral

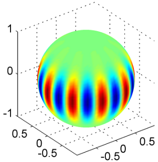

Sectoral

Spherical harmonic transform (SHT) - $f(\xi_i, r) \in S^2$ (for each r)

- $f(\xi_i, r) = \sum_{p=0}^{p_{\max}} \sum_{|q| \leq p} f^r_{p,q} Y_{p,q}(\xi_i)$
- $f^r_{p,q} = \frac{4\pi}{b} \sum_{i=0}^{b-1} f(\xi_i, r) Y_{p,q}(\xi_i)$
- $f(\xi_i, r)$ is a single pixel in $\mathbf{f}(\xi_i, r)$
- Expand all m pixels: $\mathbf{f}^r_{p,q} \in \mathbb{R}^{m \times 1}$ is a harmonic image of degree p and order q at pose r

Single Illumination Source

Goals:

1. Verify that the set of harmonic images at each pose are band-limited
2. Verify that for most objects, orthonormalizing the truncated set of harmonic images provides a good approximation to the eigenimages as computed using the SVD directly
3. Reduce the dimensionality of the image data due to a change in illumination at each of the a poses

Test Data

- Test objects
  - Each image 128×128
  - 90 different poses on S^1
  - 48 different illumination directions (HEALPix)
  - Reduce dimensionality to 9, 16, 25, and 36 harmonic images
  - SVD for ground truth

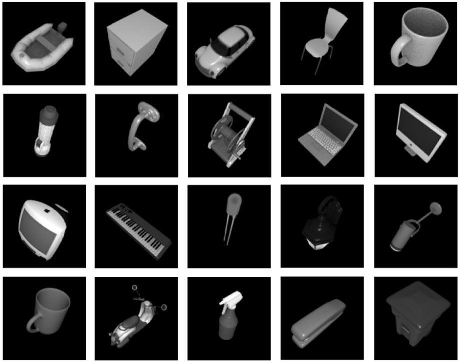

Evaluation

- 95% energy recovered by 9-D subpspace ($p = 2$)
  - low-pass filter removes specular spikes

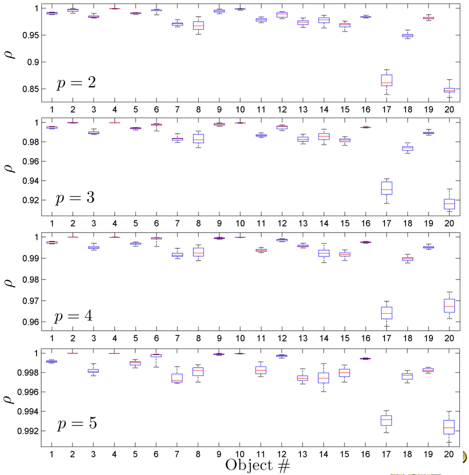

Evaluation

- 95% energy recovered by 9-D subpspace ($p = 2$)
  - low-pass filter removes specular spikes
- 9-D subspace spans over 85% of the same space

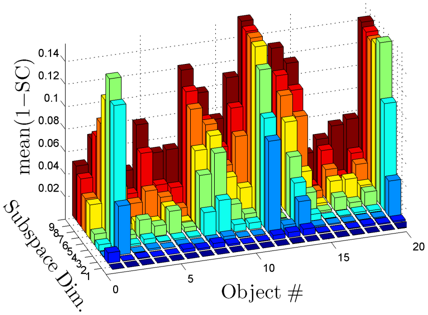

Evaluation

- 95% energy recovered by 9-D subpspace ($p = 2$)
  - low-pass filter removes specular spikes
- 9-D subspace spans over 85% of the same space
- 7-D subspace very comparable to true SVD
  - the 8th and 9th eigenimage account for large specular spikes
    - low-dimensional subspace can't recover these

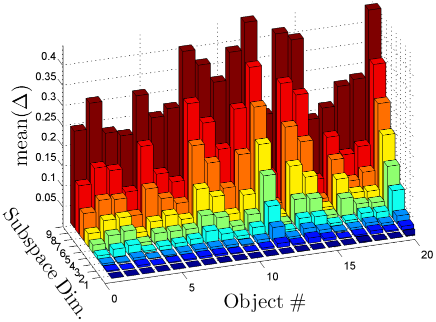

Evaluation

- 95% energy recovered by 9-D subpspace ($p = 2$)
  - low-pass filter removes specular spikes
- 9-D subspace spans over 85% of the same space
- 7-D subspace very comparable to true SVD
  - the 8th and 9th eigenimage account for large specular spikes
    - low-dimensional subspace can't recover these
- Conclusion: reducing the dimensionality of the data in the illumination dimension can be efficiently done using a truncated SHT

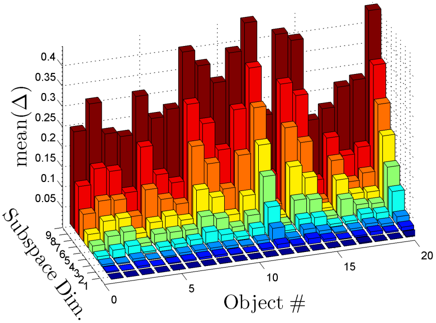

Multiple Sources

- In most real world applications illumination may exist from multiple sources and several directions
- Object properties
  - Cast shadows
  - Attached shadows
  - Surface reflections (specularities)
- It has been shown that for single sources: [Epstein *et al.* 95]
  - The first few eigenimages account for diffuse shading
  - The next few account for specular lobes
  - The higher order eigenimages account for sharp specular spikes

Multiple Sources

- In most real world applications illumination may exist from multiple sources and several directions
- Object properties
  - Cast shadows
  - Attached shadows
  - Surface reflections (specularities)
- It has been shown that for single sources: [Epstein *et al.* 95]
  - The first few eigenimages account for diffuse shading
  - The next few account for specular lobes
  - The higher order eigenimages account for sharp specular spikes

  - Evaluate how well the 9-D subspace can recover information from images of objects when multiple sources exist

Multiple Sources - Test Data

- 10 different poses on S^1
- 48 illumination directions
  - Single source (SS)
  - Two sources (DS)
  - Three source (TS)
  - All three sources were placed at random on S^2
  - 9-D subspace using a single source and HEALPix

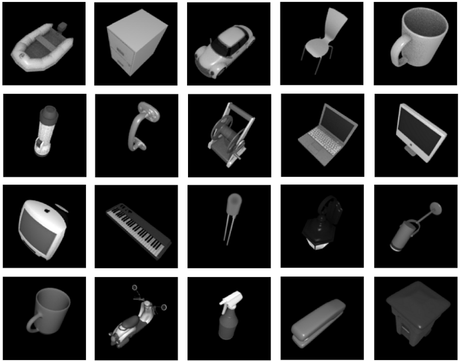

Multiple Sources - Test Data

- 10 different poses on S^1
- 48 illumination directions
  - Single source (SS)
  - Two sources (DS)
  - Three source (TS)
  - All three sources were placed at random on S^2
  - 9-D subspace using a single source and HEALPix
- Energy recovery is still high and distributions are tight for most objects

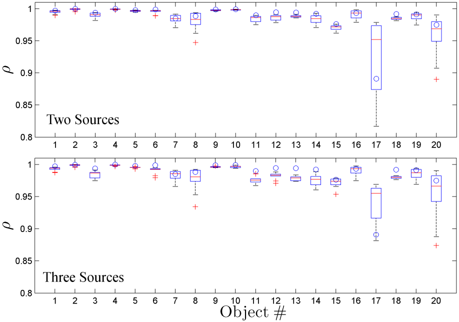

Empirical Evaluation

Reconstruction under three illumination directions and a fixed pose

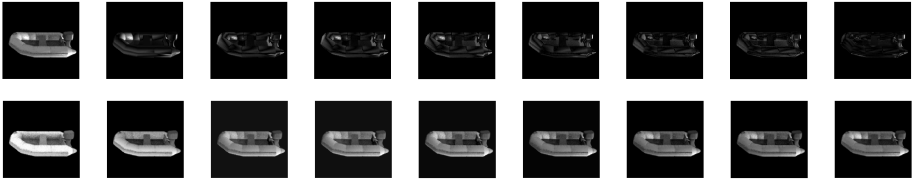

Empirical Evaluation

Reconstruction under three illumination directions and a fixed pose

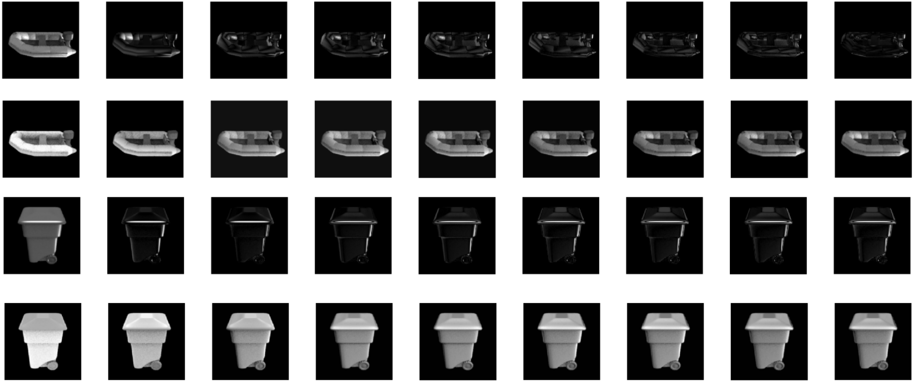

Reconstruction Error

- Treat each row-scanned image as a point in m-dimensional space
- Compute Euclidean distance between reconstruction and original image
  - Use this metric for all 480 test images (10 poses and 48 illumination conditions)

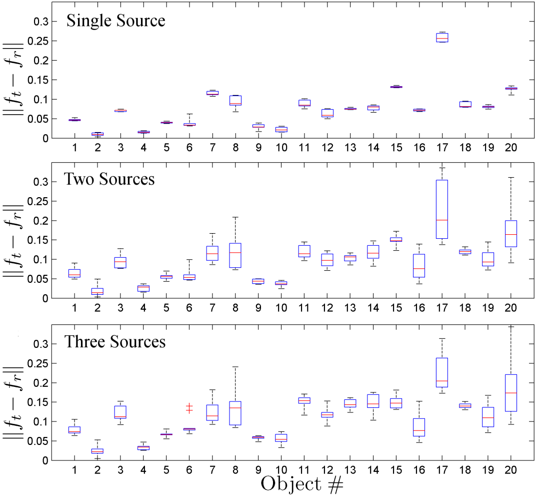

Reconstruction Error

- Treat each row-scanned image as a point in m-dimensional space
- Compute Euclidean distance between reconstruction and original image
  - Use this metric for all 480 test images (10 poses and 48 illumination conditions)
- Increased probability that local specularities will be illuminated

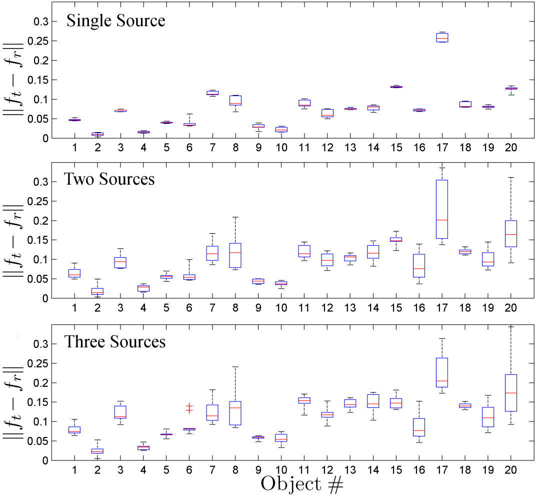

Dimensionality Reduction in Illumination Dimension (Graphical Interpretation)

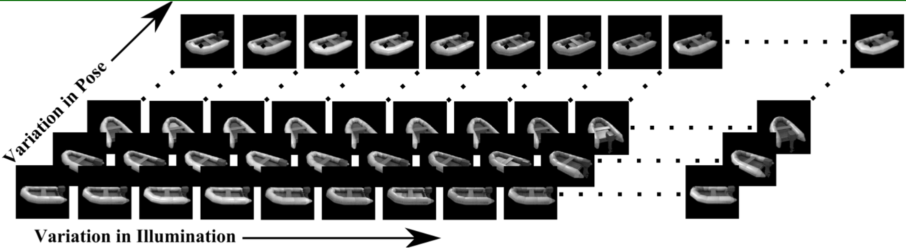

Dimensionality Reduction in Illumination Dimension (Graphical Interpretation)

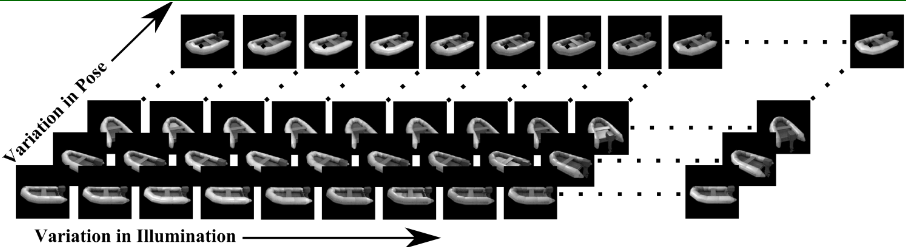

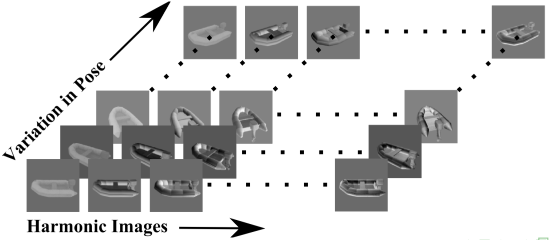

Dimensionality Reduction in Illumination Dimension

- Each harmonic image corresponds to a spherical harmonic of degree p and order q
- Construct: $\hat{X}_{p,q} = [\mathbf{f}^{0}_{p,q}, \mathbf{f}^{1}_{p,q}, \cdots, \mathbf{f}^{a-1}_{p,q}]$ for each (p, q) combination
  - There will be nine such matrices

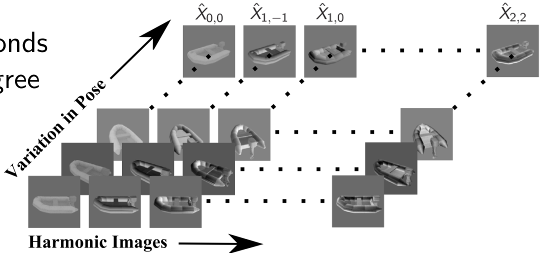

Each matrix $\hat{X}_{p,q}$ is one-dimensionally correlated for each (p, q) combination

Correlation on S^1 - Chang's Algorithm

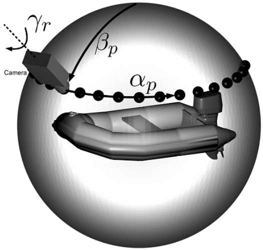

- Sample on lines of constant co-latitude
- Image data matrix $\hat{X}$ - correlated on S^1
- General idea behind Chang's algorithm
  - Majority of pixels change slowly throughout sequence
  - Right singular vectors are approx. sinusoids
  - Most of the energy in $\hat{X}$ is concentrated around the low frequency Fourier harmonics
  - SVD of the p low frequency harmonics of FFT($\hat{X}$) $\approx \hat{U}_p$

Correlation on S^1 - Example

Example Images

Eigenimages

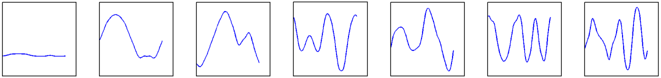

Right Singular Vectors

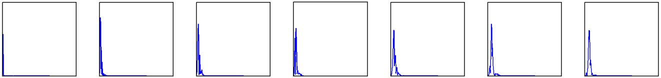

Power Spectra of Right Singular Vectors

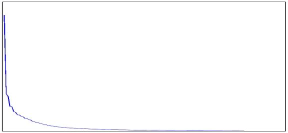

Singular Values

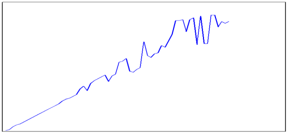

Dominant Frequencies of the Power Spectra

Dimensionality Reduction in the Pose Dimension (Chang's Algorithm)

Chang's Algorithm

- Construct the matrix ($c_k = \cos(2\pi k/n)$ and $s_k = \sin(2\pi k/n)$)

$$H = \sqrt{\frac{2}{n}} \begin{bmatrix} \frac{1}{\sqrt{2}} & c_0 & -s_0 & c_0 & -s_0 & \cdots \\ \frac{1}{\sqrt{2}} & c_1 & -s_1 & c_2 & -s_2 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \cdots \\ \frac{1}{\sqrt{2}} & c_{n-1} & -s_{n-1} & c_{2(n-1)} & -s_{2(n-1)} & \cdots \end{bmatrix}$$

- Compute $\hat{X}H$ by means of the FFT
- Find the smallest number p such that $\rho(\hat{X}, H_p) \geq \mu$
- Chang proved that: $\rho(\hat{X}, U_p) \geq \rho(\hat{X}, \tilde{\hat{U}}_p) \geq \rho(\hat{X}, H_p)$

Recombination

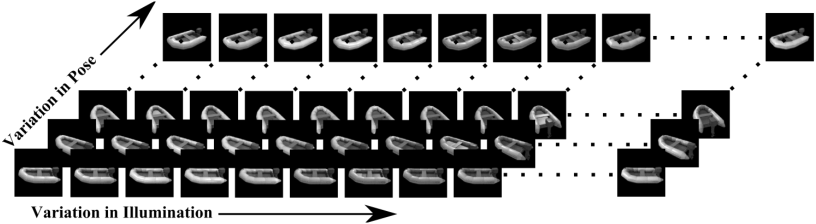

Recombination

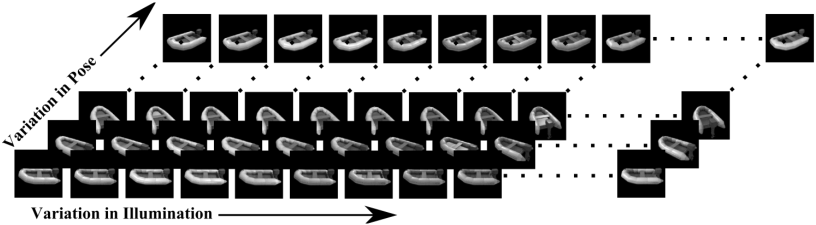

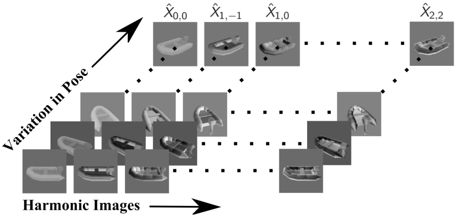

Recombination

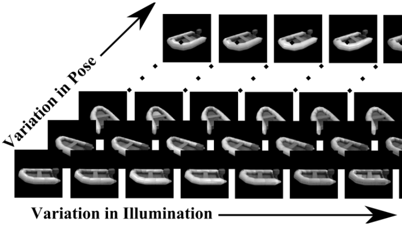

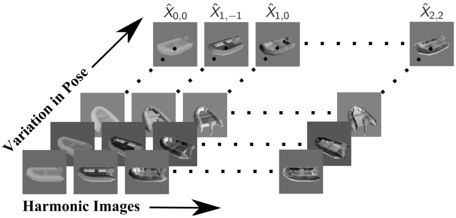

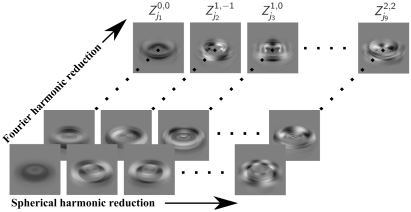

Recombination

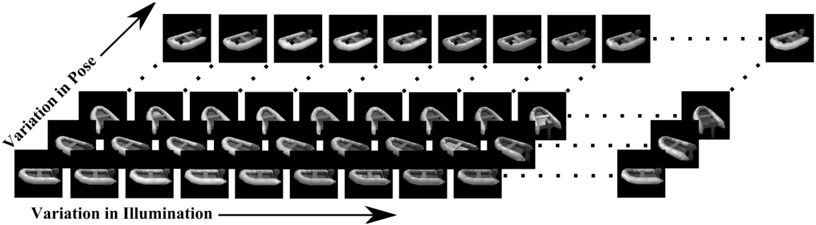

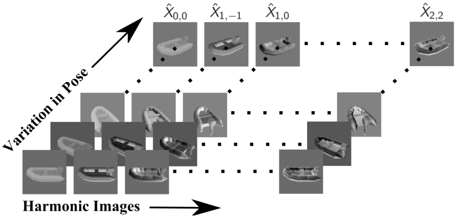

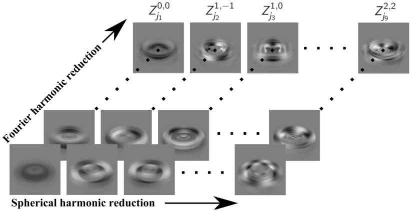

- Recombine variation due to a change in both illumination and pose: $\bar{X} = [Z^{0,0}_{j_1}, Z^{1,-1}_{j_2}, Z^{1,0}_{j_3}, Z^{1,1}_{j_4}, Z^{2,-2}_{j_5}, \cdots, Z^{2,2}_{j_9}]$
- Compute SVD$(\bar{X})$

Algorithm Summary

Eigenspace Decomposition Algorithm Summary

1. Use the SHT to compute the matrices $\hat{X}_{p,q}$ for each r.
2. For each of the nine matrices $\hat{X}_{p,q}$, determine the smallest number j_i such that $\rho(\hat{X}_{p,q}^T, H_{j_i}) \geq \mu_t$, where μ_t is the user specified energy recovery ratio in the pose dimension, and $i = 1, 2, \dots, 9$ corresponds to the i^{th} matrix $\hat{X}_{p,q}$. Def. of H
3. Let $Z_{j_i}^{p,q}$ denote the matrix $\hat{X}_{p,q} H_{j_i}$ and construct the matrix $\bar{X} = [Z_{j_1}^{0,0}, Z_{j_2}^{1,-1}, Z_{j_3}^{1,0}, Z_{j_4}^{1,1}, Z_{j_5}^{2,-2}, \cdots, Z_{j_9}^{2,2}]$. Note that the matrices $Z_{j_i}^{p,q}$ can be efficiently computed using the FFT.
4. Compute the SVD of $\bar{X} = \tilde{\hat{U}} \tilde{\hat{\Sigma}} \tilde{\hat{V}}$.
5. Return $\rho(\hat{X}, \tilde{\hat{U}}_k) \geq \mu$. Where μ is the user specified energy recovery ratio.

Test Data

- Test objects
  - Each image 128×128
  - 90 different poses on S^1
  - 48 different illumination directions (HEALPix)
  - 9-D subspace at each pose (using SHT)
  - $\mu_t = 0.95$ and $\mu = 0.8$
  - SVD for ground truth

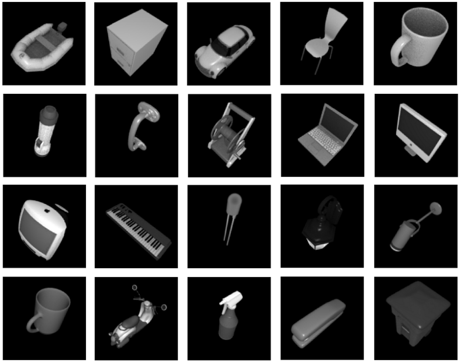

Experimental Results

- Computational Savings
  - Col. dim. of $\hat{X} = 4230$
  - Col. dim. of $\bar{X}$ never exceeds 576
  - Average speed-up = 214

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

Experimental Results

- Computational Savings
  - Col. dim. of $\hat{X} = 4230$
  - Col. dim. of $\bar{X}$ never exceeds 576
  - Average speed-up = 214
- Quality of estimates
  - Subspace dim $\approx$ SVD($\hat{X}$)

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

Experimental Results

- Computational Savings
  - Col. dim. of $\hat{X} = 4230$
  - Col. dim. of $\bar{X}$ never exceeds 576
  - Average speed-up = 214
- Quality of estimates
  - Subspace dim $\approx$ SVD($\hat{X}$)
  - Difference in energy recovery is comparable for most objects

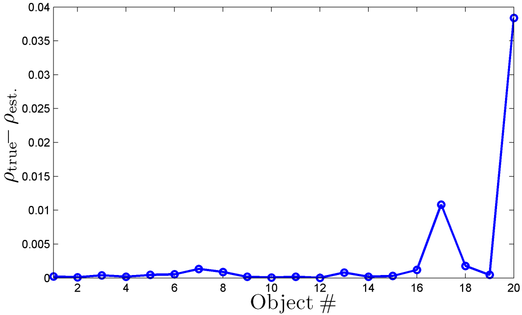

- Computational Savings
  - Col. dim. of $\hat{X} = 4230$
  - Col. dim. of $\bar{X}$ never exceeds 576
  - Average speed-up $=$ 214
- Quality of estimates
  - Subspace dim $\approx$ SVD($\hat{X}$)
  - Difference in energy recovery is comparable for most objects
  - Estimated eigenimages are very comparable to the true eigenimages

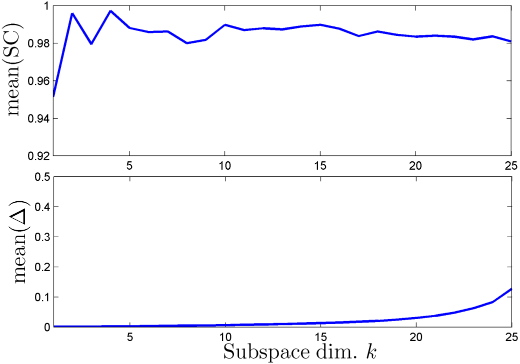

How to Estimate the Pose of Objects with Unknown Illumination Conditions?

- Recall: $\mathcal{M}_k = \hat{\tilde{U}}_k^T \hat{X}$ and $\mathcal{P} = \hat{U}_k^T \mathbf{f}_{\text{new}}$
- Denote each illumination manifold by $\mathcal{I}_r$

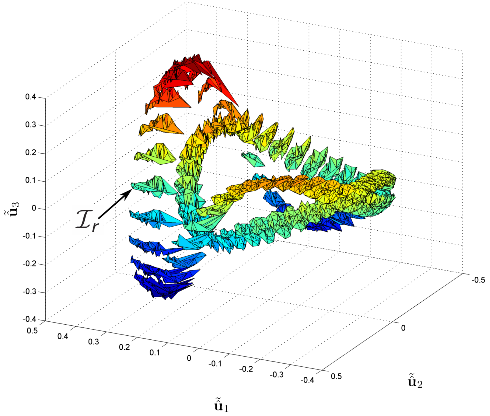

How to Estimate the Pose of Objects with Unknown Illumination Conditions?

- Recall: $\mathcal{M}_k = \hat{\tilde{U}}_k^T \hat{X}$ and $\mathcal{P} = \hat{U}_k^T \mathbf{f}_{\text{new}}$
- Denote each illumination manifold by $\mathcal{I}_r$
- Compute: $\mathcal{C} = [\mathbf{c}_0, \mathbf{c}_2, \ldots, \mathbf{c}_{a-1}]$, where

$$\mathbf{c}_r = \frac{1}{b} \sum_{i=1}^{b} \mathcal{I}_r^{(i)}$$

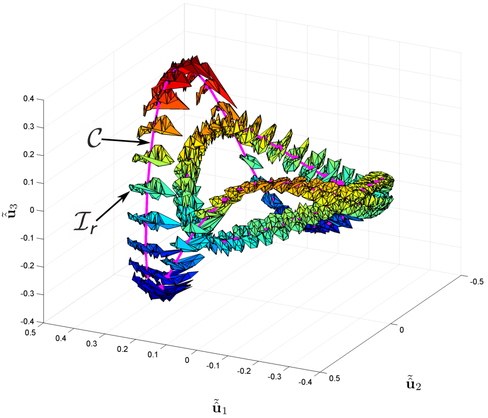

How to Estimate the Pose of Objects with Unknown Illumination Conditions?

- Recall: $\mathcal{M}_k = \tilde{\hat{U}}_k^T \hat{X}$ and $\mathcal{P} = \hat{U}_k^T \mathbf{f}_{\text{new}}$
- Denote each illumination manifold by $\mathcal{I}_r$
- Compute: $\mathcal{C} = [\mathbf{c}_0, \mathbf{c}_2, \ldots, \mathbf{c}_{a-1}]$, where

$$\mathbf{c}_r = \frac{1}{b} \sum_{i=1}^{b} \mathcal{I}_r^{(i)}$$

- $\mathcal{C}$ has significantly fewer points than $\mathcal{M}$

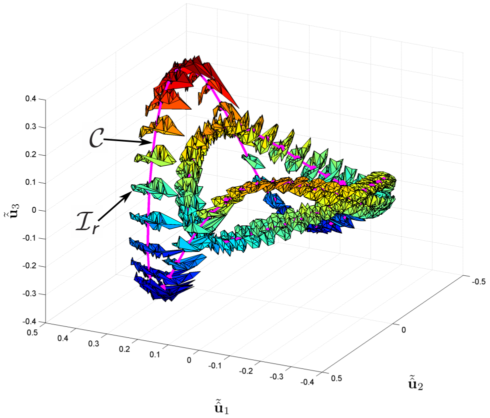

Eigenspace Partitioning

1. Issue: $\mathcal{C}$ may not be sufficient for accurate pose estimation

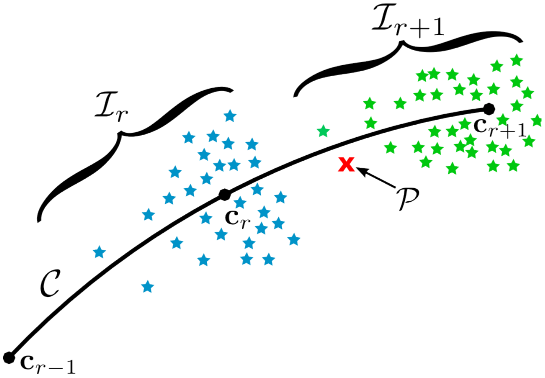

Eigenspace Partitioning

1. Issue: $\mathcal{C}$ may not be sufficient for accurate pose estimation
   - Search both $\mathcal{I}_r$ and $\mathcal{I}_{r+1}$ that bound $\mathcal{P}$

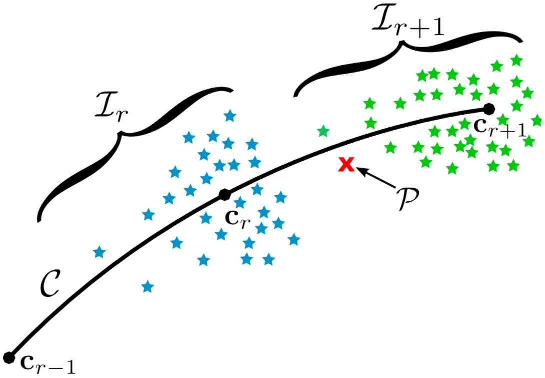

Eigenspace Partitioning

1. Issue: $\mathcal{C}$ may not be sufficient for accurate pose estimation
    - Search both $\mathcal{I}_r$ and $\mathcal{I}_{r+1}$ that bound $\mathcal{P}$
2. Issue: The illumination manifolds may intersect

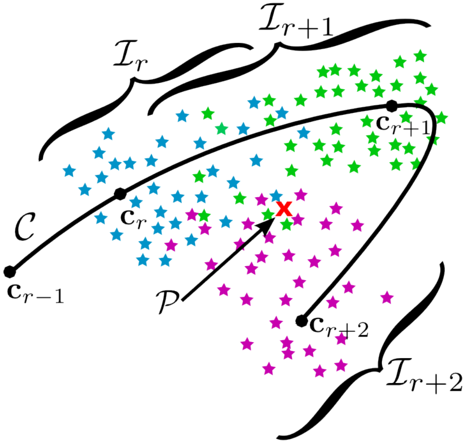

Eigenspace Partitioning

1. Issue: $\mathcal{C}$ may not be sufficient for accurate pose estimation
   - Search both $\mathcal{I}_r$ and $\mathcal{I}_{r+1}$ that bound $\mathcal{P}$
2. Issue: The illumination manifolds may intersect
   - Compare $p_r^{r+1} + p_{r+1}^r$ to the r^{th} element in

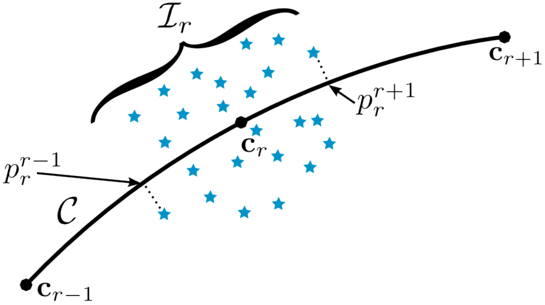

$$\Delta\mathcal{C} = [\|\mathbf{c}_2 - \mathbf{c}_1\|, \|\mathbf{c}_3 - \mathbf{c}_2\|, \cdots, \|\mathbf{c}_a - \mathbf{c}_{a-1}\|, \|\mathbf{c}_1 - \mathbf{c}_a\|]$$

- If $\Delta\mathcal{C}_r \geq (p_r^{r+1} + p_{r+1}^r)\ \forall r$, then, $\mathcal{I}_r \cap \mathcal{I}_{r+1}$ is empty
  - Fortunately, the structure of the eigenspace can be analyzed off-line

Analysis of Eigenspace Partitioning

- Object 1: No intersections $\Longrightarrow$ accurate pose estimation
  - Variations due to a change in pose are larger than illumination

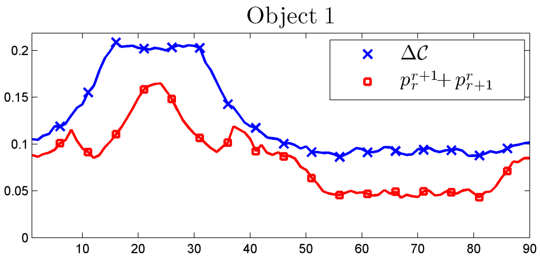

Analysis of Eigenspace Partitioning

- Object 1: No intersections $\Longrightarrow$ accurate pose estimation
  - Variations due to a change in pose are larger than illumination
- Object 13: Several intersections $\Longrightarrow$ inaccurate pose estimation
  ▸ Objects
  - Variations due to a change in pose are NOT larger than illumination

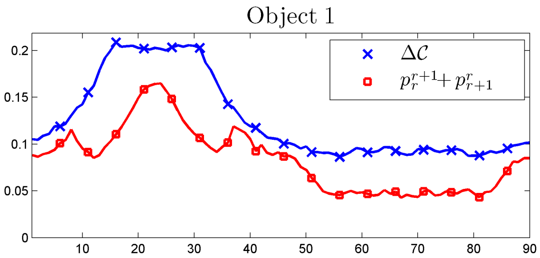

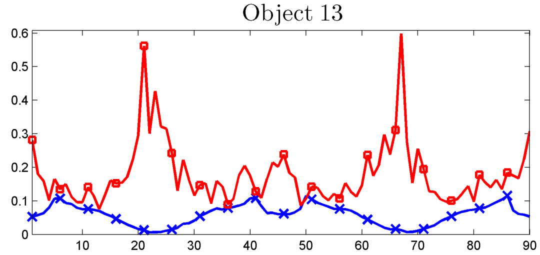

Analysis of Eigenspace Partitioning

- Object 1: No intersections $\Longrightarrow$ accurate pose estimation
  - Variations due to a change in pose are larger than illumination
- Object 13: Several intersections $\Longrightarrow$ inaccurate pose estimation
  Objects
  - Variations due to a change in pose are NOT larger than illumination
  - The problem is ill-posed

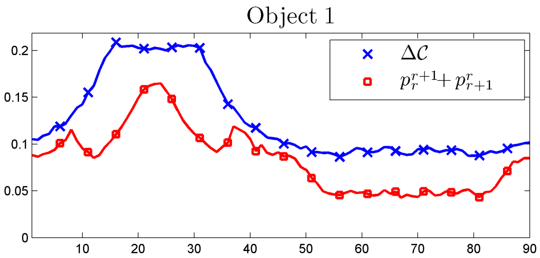

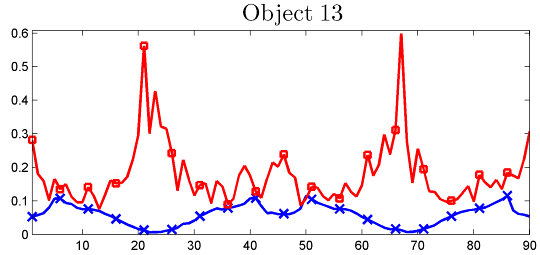

Estimation Accuracy

Evaluate how well the centroid manifold performs as compared to traditional eigenspace search techniques

Estimation Accuracy

Evaluate how well the centroid manifold performs as compared to traditional eigenspace search techniques

- 90 random but known poses on S^1
- 270 test conditions
  - Single source (SS)
  - Two sources (DS)
  - Three source (TS)
  - All sources were placed at random on S^2
  - 9-D subspace using a single source and HEALPix
- $n = 4320$ evaluations using traditional methods
- $a + 2b = 186$ evaluations using proposed method

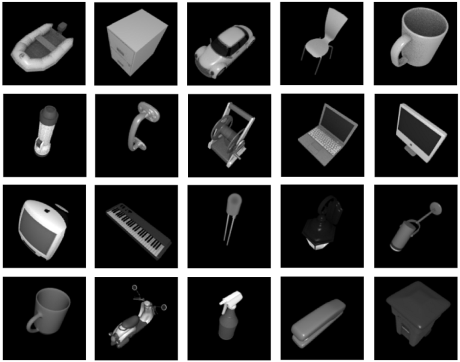

Estimation Accuracy

- Single source using exhaustive search (measure of difficulty)
- Pose est. for some objects is inherently difficult

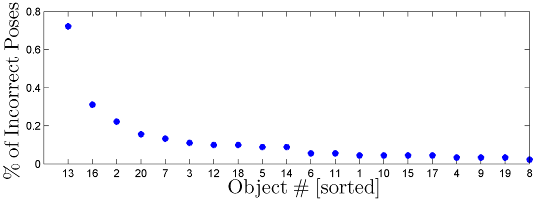

Estimation Accuracy

- Single source using exhaustive search (measure of difficulty)
- Pose est. for some objects is inherently difficult
- Proposed technique is comparable to traditional methods
- Multiple sources has little effect on accurate estimation

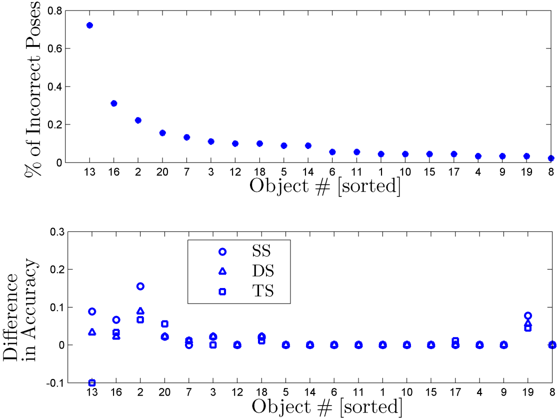

Publications Resulting From Second Part

1. R. C. Hoover, A. A. Maciejewski, and R. G. Roberts, "Fast Eigenspace Decomposition for Illumination Invariant Pose Estimation," under review in *IEEE Transactions on Systems, Man, and Cybernetics B: Cybernetics*, 2009.
2. R. C. Hoover, A. A. Maciejewski, R. G. Roberts, and R. P. Hoppal, "An Illustration of Eigenspace Decomposition for Illumination Invariant Pose Estimation," accepted to appear in *2009 IEEE International Conference on Systems, Man, and Cybernetics*, pp. -, San Antonio, TX, Oct. 11-14, 2009.
3. R. C. Hoover, A. A. Maciejewski, and R. G. Roberts, "Designing Eigenspace Manifolds: With Application to Object Identification and Pose Estimation," accepted to appear in *2009 IEEE International Conference on Systems, Man, and Cybernetics*, pp. -, San Antonio, TX, Oct. 11-14, 2009.

Conclusions

- First part: Eigenspace decomposition of spherically correlated images
  - Proposed method of sampling - spherical (HEALPix)
  - Proposed efficient algorithm for computing the eigenspace decomposition of images correlated on S^2 and $SO(3)$
    - Spherical harmonics - Wigner-D matrices
  - Shown significant computational savings
  - Good quality of estimation

Conclusions

- First part: Eigenspace decomposition of spherically correlated images
  - Proposed method of sampling - spherical (HEALPix)
  - Proposed efficient algorithm for computing the eigenspace decomposition of images correlated on S^2 and $SO(3)$
    - Spherical harmonics - Wigner-D matrices
  - Shown significant computational savings
  - Good quality of estimation
- Second part: Eigenspace decomposition of images with variations in pose as well as illumination
  - Reduce illumination dimension by using the SHT
  - Evaluated the effects of multiple illumination sources
    - Increased probability that a local specularity will be illuminated
  - Reduce the pose dimension by using Fourier harmonics (Chang's Alg.)
  - Analyze the structure of the eigenspace manifold
  - Proposed a technique to partition the eigenspace for efficient searching
  - Multiple illumination sources have little effect on pose estimation for most objects

Thank you for your attention!!!

$$H = \sqrt{\frac{2}{n}} \begin{bmatrix} \frac{1}{\sqrt{2}} & c_0 & -s_0 & c_0 & -s_0 & \cdots \\ \frac{1}{\sqrt{2}} & c_1 & -s_1 & c_2 & -s_2 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \cdots \\ \frac{1}{\sqrt{2}} & c_{n-1} & -s_{n-1} & c_{2(n-1)} & -s_{2(n-1)} & \cdots \end{bmatrix}$$

- Where $c_k = \cos(2\pi k/n)$ and $s_k = \sin(2\pi k/n)$

Quality Measures

Back

Subspace criterion (SC)

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

Residue between subspaces

$$\Delta = \min_Q ||A - BQ||_F$$

- Compute $SVD(\tilde{\hat{U}}_k^T \hat{U}_k) = U_c\Sigma_c V_c^T$
- The matrix $Q_{\min} = U_c V_c^T$
- $\Delta^2 = 2\left(k - \sum_{i=1}^{k}\sigma_{ci}\right)$
  - Normalized by $\sqrt{2k}$

Objects ◂ Back