# Biased Manifold Embedding: A Framework for Person-Independent Head Pose Estimation

Vineeth Nallure Balasubramanian
Center for Cognitive Ubiquitous Computing
Arizona State University, Tempe, AZ 85281 USA
vineeth.nb@asu.edu

Jieping Ye
Arizona State University,
Tempe, AZ 85281 USA
jieping.ye@asu.edu

Sethuraman Panchanathan
Center for Cognitive Ubiquitous Computing
Arizona State University, Tempe, AZ 85281 USA
panch@asu.edu

## Abstract

*The estimation of head pose angle from face images is an integral component of face recognition systems, human computer interfaces and other human-centered computing applications. To determine the head pose, face images with varying pose angles can be considered to be lying on a smooth low-dimensional manifold in high-dimensional feature space. While manifold learning techniques capture the geometrical relationship between data points in the high-dimensional image feature space, the pose label information of the training data samples are neglected in the computation of these embeddings. In this paper, we propose a novel supervised approach to manifold-based non-linear dimensionality reduction for head pose estimation. The Biased Manifold Embedding (BME) framework is pivoted on the ideology of using the pose angle information of the face images to compute a biased neighborhood of each point in the feature space, before determining the low-dimensional embedding. The proposed BME approach is formulated as an extensible framework, and validated with the Isomap, Locally Linear Embedding (LLE) and Laplacian Eigenmaps techniques. A Generalized Regression Neural Network (GRNN) is used to learn the non-linear mapping, and linear multi-variate regression is finally applied on the low-dimensional space to obtain the pose angle. We tested this approach on face images of 24 individuals with pose angles varying from $-90\,^\circ$ to $+90\,^\circ$ with a granularity of 2. The results showed substantial reduction in the error of pose angle estimation, and robustness to variations in feature spaces, dimensionality of embedding and other parameters.*

## 1. Introduction

Human face analysis has been growing in its importance by the day as a problem studied by several research communities, as technology assumes a human-centric approach. The estimation of head pose angle from face images is a sub-problem of human face analysis with several applications like 3D face modeling, gaze direction detection, driver monitoring safety systems, etc. With the expanding need for robust face recognition systems, realistic solutions to this problem require the ability to handle significant head pose variations. While coarse head pose estimation has been successful to a large extent [3], accurate person-independent head pose estimation, which is crucial for applications like 3D face modeling, is still in the works.

Current literature [5] [10] [14] separates the existing methods for head pose estimation into distinct categories:

- Shape-based geometric analysis, where head pose is discerned from geometric information like the configuration of facial landmarks.
- Model-based methods, where non-linear parametric models are derived before using a classifier like a neural network (Eg. Active Appearance Models (AAMs)).
- Appearance-based methods, where the pose estimation problem is viewed as a pattern classification problem on image feature spaces.
- Template matching approaches, which are largely based on nearest neighbor classification against texture templates/signatures.

1-4244-1180-7/07/$25.00 ©2007 IEEE

- Dimensionality reduction based approaches, where linear/non-linear embedding of the face images is used for pose estimation.

To obtain better representations of face images, earlier work [4] [10] [5] suggests to consider the high-dimensional feature space of face image data as a set of geometrically related points lying on a smooth manifold in the feature space.

Different poses of the head, although captured in high-dimensional image feature spaces, can be visualized as data points lying on a low-dimensional manifold in the high-dimensional space. Raytchev et al [10] stated that the dimension of this manifold is equivalent to the number of degrees of freedom in the movement during data capture. For example, images of the human face with different angles of pose rotation (yaw, tilt and roll) can intrinsically be conceptualized as a 3D manifold in image feature space. This idea underlies the family of non-linear dimensionality reduction techniques under the umbrella of manifold learning, like Isomap, Locally Linear Embedding (LLE), Laplacian Eigenmaps, Local Tangent Space Alignment (LTSA), etc, which have become popular in recent times.

In prior work in this domain, Raytchev et al [10] and Hu et al [6] employed a straight-forward approach to learn the non-linear mapping onto the low-dimensional space through manifold learning, and estimated the pose angle using a pose parameter map. In the work carried out so far, the pose information of the given face images is ignored while computing the embedding. While supervised approaches have been attempted in the past for classification problems [11] [15], the formulation of pose estimation assumes a regression flavor, thus laying the basis for this work. We propose a novel framework called Biased Manifold Embedding to extend traditional manifold learning techniques, which provides a semantic bias to the manifold-based embedding process, using pose information from the given face image data. While the proposed method is illustrated using Isomap, LLE and Laplacian Eigenmaps in this paper, it can easily be extended as a complete framework to all other manifold learning techniques with minor adaptations. As broader impact, the work proposed here is a significant contribution to a supervised approach to manifold-based non-linear dimensionality reduction techniques across all regression problems.

We discuss the background with a brief description of manifold learning techniques, followed by related work and an insight into the significance of our work in Section 2. Section 3 details the mathematical formulation of the proposed Biased Manifold Embedding framework. The experimental setup and the methodology of our experiments are briefed in Section 4. The results of the experiments are discussed in Section 5. We then discuss the advantages and limitations of the approach in the concluding section in Section 6, and provide future directions to this work.

## 2. Background

### 2.1. Non-linear Dimensionality Reduction using Manifold Learning

The computation of low-dimensional representations of high-dimensional observations (like images, spectral data, etc) is a problem that plagues every field of science and engineering. Techniques like Principal Component Analysis (PCA) [7] are categorized as linear dimensionality reduction techniques, and are often applied to obtain the low-dimensional representation. Other dimensionality reduction techniques like Multi-Dimensional Scaling (MDS) [9] use the Euclidean distance between data points in the high-dimensional space to capture the relationships between them. However, when data points lie on a manifold in the high-dimensional space, Euclidean distances do not capture the geometric relationship between the points. In such cases, it becomes necessary to consider the geodesic distances between data points to obtain a truthful representation of the data. While techniques like Isomap capture the global geometry of the surface on which the data points lie, LLE and Laplacian Eigenmaps adopt a local fitting approach based on the neighborhood of each data point.

#### 2.1.1 Isomap

To capture the global geometry of the data points, Tenenbaum et al [13] proposed Isomap to compute an isometric low-dimensional embedding of a given set of high-dimensional data points. In this method, the neighbors of a point on the manifold $M$ are determined, and the neighborhood of each point is represented as a weighted graph $G$, with each edge characterized by the distance $dx(i, j)$ between the pair of neighboring points, $x_i$ and $x_j$. The geodesic distances between all pairs of points on the manifold $M$ are estimated by computing their shortest path distance in the graph $G$. This is done using the Floyds or Djkstraas algorithm, i.e. $dM(x_i, x_j) = min_k \{dM(x_i, x_j), dM(x_i, x_k) + dM(x_k, x_j)\}$. Classical MDS is then applied to the geodesic distance matrix, deriving an embedding of the data in a low-dimensional Euclidean space that best preserves the estimated intrinsic geometry of the manifold.

#### 2.1.2 Locally Linear Embedding (LLE)

Roweis and Saul [12] proposed the LLE algorithm that embodied the think globally, fit locally paradigm. In this technique, the neighbors of a point of the manifold are deter-

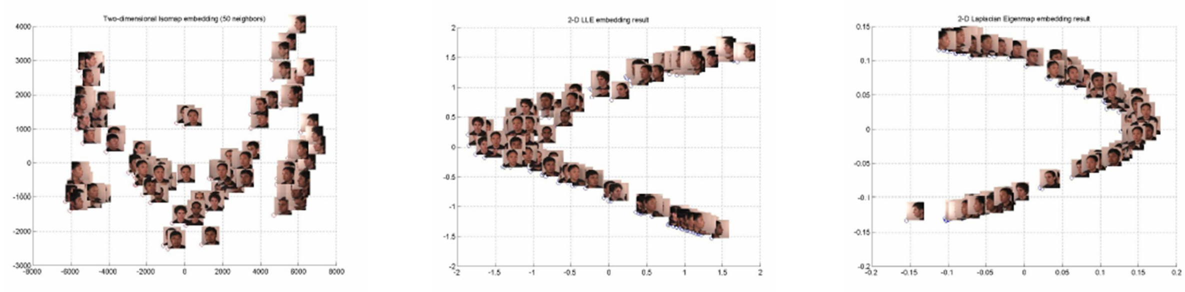


(a) Isomap embedding (b) LLE embedding (c) Laplacian Eigenmap embedding

Figure 1. Embedding of face images with varying poses onto 2 dimensions

mined as for Isomap. The data point is shifted to the origin along with its neighborhood to form a local data matrix Z, and the local covariance $C = Z'Z$ is computed. The linear system $CW = 1$ is solved for the weights $W$ in the neighborhood, which are subsequently normalized. The bottom eigenvectors of a sparse matrix $M$, constructed as $M = (I - W)'(I - W)$, are used to project the input vectors into the low-dimensional embedding space.

#### 2.1.3 Laplacian Eigenmaps

Belkin and Niyogi [1] proposed another geometrically motivated algorithm based on the Laplace-Beltrami operator on a manifold. In this approach, the Laplacian of the graph of the neighborhood of every data point in the feature space is viewed as an approximation to the Laplace-Beltrami operator. A weighted graph is constructed with weight values $W$ drawn from the heat kernel or with a simplistic version, where a weight of unit value is assigned if the nodes are neighbors. The generalized eigenvector problem $Ly = \lambda Dy$, is solved for the embedding $y$, where $D$ is the diagonal weight matrix i.e. $D_{ii} = \sum_j Wji$, and $L = D - W$ is the Laplacian matrix.

While these techniques capture the geometry of the data points in the high-dimensional space, the disadvantage of this family of manifold learning techniques is the lack of a projection matrix to embed out-of-sample data points after the training phase. This makes the method more suited for data visualization, rather than classification problems. However, the advantage of these techniques to capture the relative geometry of data points enthuses researchers to adopt this methodology to solve problems like head pose estimation, where the data is known to possess geometric relationships in a high-dimensional space.

Figure 1 shows the visualization results of using Isomap, LLE and Laplacian Eigenmaps to embed face images onto 2 dimensions. Faces of 10 individuals with 11 pose angles ($-75\,^\circ$ to $+75\,^\circ$ in increments of 15) were used to perform this embedding from the grayscale pixel intensity feature space. On close observation of the iconic images on the plots, Figure 1 illustrates that the embedding of the face images reflects an intrinsic ordering on the corresponding pose angles. The frontal view falls in the center of an elliptical trajectory, with all negative pose angles on one side, and the positive pose angles on the other. While this result supports the application of manifold learning based methods for face images with varying pose angles, the images of the same individual with different pose angles tend to cluster together forming a clutter in the ordering. This suggests that fine estimation of pose angle still remains a challenging problem.

### 2.2. Related Work

Over the last few years since the arrival of manifold learning techniques, a reasonable amount of work has been done using manifold-based dimensionality reduction techniques for head pose estimation. Chen et al [4] considered multi-view face images as lying on a manifold in high-dimensional feature space. However, they compared the effectiveness of Kernel Discriminant Analysis against Support Vector Machines in learning the manifold gradient direction in the high-dimensional feature space, and did not adopt manifold learning for non-linear dimensionality reduction. Raytchev et al [10] studied the effectiveness of Isomap for head pose estimation against other view representation approaches like the Linear Subspace model and Locality Preserving Projections (LPP). While their work established the possible gain in accuracy through use of manifold learning techniques, the face images used by them were sampled at pose angle increments of $15\,^\circ$, and relied on the robustness of the captured mapping and interpolation to obtain the precise pose angle estimate. Hu et al [6] developed a unified embedding approach for multiple individuals, where the embedding obtained from Isomap for a single individual was parametrically modeled as an ellipse. The ellipses for different individuals were subsequently normalized through scale, translation and rotation based transformations to obtain a unified embedding. In more recent work, Fu and Huang [5] presented an appearance-based strategy for head pose estimation using a supervised form of Graph Embedding, which internally used the idea of Locally Linear Embedding (LLE). This work obtained a linearization of manifold learning techniques to treat out-of-sample data points. Recent work by Ridder et al [11] and Yu et al [15] focused on obtaining a supervised approach to manifold learning techniques. However, their approaches are strictly oriented towards classification problems, and do not exploit the label information as possible for regression problems like head pose estimation.

### 2.3. Proposed Approach

Unlike class labels in classification problems, the pose labels of sample points can be viewed as an ordered single-dimensional value with an established distance metric. In the proposed Biased Manifold Embedding approach, we use the given pose information to bias the non-linear embedding to obtain accurate pose angle estimation. The significance of our contribution is realized in the fact that the proposed Biased Manifold Embedding framework, although validated in this work with Isomap, LLE and Laplacian Eigenmaps, can be extended to all manifold learning techniques with minor modifications, and in general, can be applied to any regression problem that uses manifold learning methods. In addition, while most current approaches use face images sampled with pose angles at increments of $10-15\,^\circ$ [10], we use the FacePix database [8] that includes images of faces taken at a wide range of precisely measured pose angles with a granularity of $1\,^\circ$. This reinforces the validity of our experiments with the proposed approach.

## 3. The Biased Manifold Embedding Framework

In the Biased Manifold Embedding framework, face images whose pose angles are closer to each other are maintained nearer to each other in the low-dimensional embedding, and images with farther pose angles are placed farther, irrespective of the identity of the individual. In the general unbiased case, the identity of an individual affects this process, causing different poses from the same individual to lie near each other in the low-dimensional embedding, rather than a monotonic ordering of pose angles. We achieve this goal with a modification to the computation of the neighborhood of each data point. The distances between data points in the high-dimensional feature space are biased with distances between the pose angles of the corresponding images. Since a distance metric can easily be defined on the pose angle values, the problem of finding closeness of pose angles is straight-forward.

The mathematical formulation of the Biased Manifold Embedding approach is given below. We would like the modified biased distance between a pair of data points to be of the form:

$$\tilde{D}(i,j) = f(P(i,j)) \otimes D(i,j)$$

where $D(i,j)$ is the Euclidean distance between two data points xi and xj , $\tilde{D}(i,j)$ is the modified biased Euclidean distance, $P(i,j)$ is the pose distance between $x_i$ and $x_j$, $f$ is any function of the pose distance, and $\otimes$ is a binary operator. If $\otimes$ was chosen as the multiplication operation, the function $f$ would be chosen as inversely proportional to the pose distance, $P(i,j)$. In general, the function $f$ could be picked from the family of reciprocal functions ($f \in \mathcal{F}_R$) based on the needs of an application. In this work, we choose the function as:

$$f(P(i,j)) = \frac{1}{max_{m,n} P(m,n) - P(i,j)}$$

This function could be replaced by an inverse exponential or quadratic function of the pose distance, for example. To ensure that the biased distance values are well-separated for different pose distances, we multiply this quantity by a function of the pose distance:

$$\tilde{D}(i,j) = \frac{\alpha(P(i,j))}{max_{m,n} P(m,n) - P(i,j)} * D(i,j)$$

where the function $\alpha$ is directly proportional to the pose distance, $P(i,j)$, and is defined in our work as:

$$\alpha(P(i,j)) = \beta * |P(i,j)|$$

where $\beta$ is a constant of proportionality, and allows parametric variation for performance tuning. In our current work, we used the pose distance as the one-dimensional distance i.e. $P(i,j) = |Pi - Pj|$, where $P_k$ is the pose angle of $x_k$.

In summary, the biased distance between a pair of points can be given by:

$$\tilde{D}(i,j) = \begin{cases} \frac{\alpha(P(i,j))}{max_{m,n} P(m,n) - P(i,j)} * D(i,j) & P(i,j) \neq 0, \\ 0 & P(i,j) = 0. \end{cases} \tag{1}$$

This biased distance matrix is used for Isomap, LLE and Laplacian Eigenmaps to obtain a pose-ordered low-dimensional embedding. While the geodesic distances are computed using this biased distance matrix in Isomap, LLE and Laplacian Eigenmaps have been modified to use these distance values to determine the neighborhood of each data point. Since the impact of the proposed approach

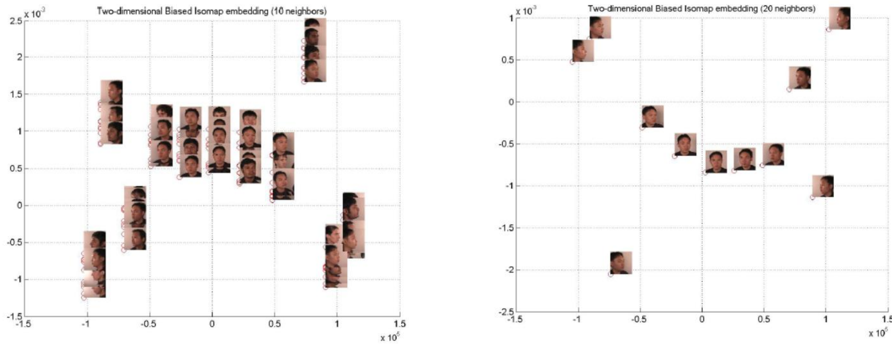


(a) Biased Isomap embedding with 10 neighbors (b) Biased Isomap embedding with 20 neighbors

Figure 2. Biased Isomap Embedding of face images with varying poses onto 2 dimensions. Note in 2(b) that all the face images with the same pose angle have merged onto the same 2D point

is restricted to the computation of the biased distances, the BME framework can easily be extended to other manifold-based dimensionality reduction techniques.

Figure 2 shows the results of using the Biased Manifold Embedding approach to embed the same facial images used in Figure 1 onto 2 dimensions. As the number of neighbors used to capture the embedding is increased, face images with the same pose merge onto the same data point in 2 dimensions (see Figure 2), irrespective of the identity of the individual. The embedded images establish the tendency of the method to elicit person-independent representations of the pose angles of the given facial images. This renders the low-dimensional embedding more conducive to deliver reliable person-independent pose angle values from the face images.

(a) Grayscale image (b) Laplacian of Gaussian (LoG) tranformed image

Figure 3. Image feature spaces used for the experiments

## 4. Head Pose Estimation: Design and Methodology

The proposed Biased Manifold Embedding framework was validated using the FacePix face database [8], which has face images with precisely measured pose variation. The results of the application of this framework to Isomap, LLE and Laplacian Eigenmaps are compared against the performance of the same manifold learning techniques without the pose bias. These three manifold learning based dimensionality reduction techniques were selected based on popular application amongst other similar techniques like spectral clustering and Local Tangent Space Alignment (LTSA). In this work, we considered a set of 2184 face images, consisting of 24 individuals with pose angles varying from $-90^\circ$ to $+90^\circ$ in increments of $2^\circ$. The images were subsampled to $32x32$ resolution, and different feature spaces of the images were considered for the experiments. The results presented here include the grayscale pixel intensity feature space and the Laplacian of Gaussian (LoG) transformed image feature space (see Figure 3). The LoG transform, which captures the edge map of the face images, was used since pose variation in face images is a direct result of geometric transformation. Preliminary experiments conducted with Gabor filters and Fourier-Mellin transformed images indicated that texture-based features may not be ideal for this problem. The images were subsequently rasterized and normalized.

Non-linear dimensionality reduction techniques like manifold learning do not provide a projection matrix to handle test data points. While different approaches have been used by earlier researchers to capture the mapping from the high-dimensional feature space to the low-dimensional embedding, we adopted a Generalized Regression Neural Network (GRNN) with Radial Basis Functions to learn the non-linear mapping. While this approach has been adopted by earlier researchers [16], the parameters involved in training the network (just one - the spread of the Radial Basis Function) are minimal, thereby facilitating better evaluation of the proposed framework. Once the low-dimensional embedding was obtained, linear multi-variate regression was used to obtain the pose angle of the test image.

The Biased Manifold Embedding approach was compared against the traditional flavors of Isomap, LLE and Laplacian Eigenmap approaches using a 8-fold cross-validation model. In this validation model, face images of 3 individuals were used for the testing phase in each fold, while all the remaining images were used in the training phase. In addition, the performance of the proposed method with different embedding dimensions and neighborhood values was studied for different image feature spaces.

# 5. Results and Discussion

The results of the experiments were evaluated by the error in the estimated pose angle against the ground truth pose angle from the FacePix database. The error values for pose angle estimation are shown in Table 1 for Isomap, LLE and Laplacian Eigenmaps with different dimensions of embedding from the grayscale pixel intensity image feature space. Table 2 presents the results when the Laplacian of Gaussian transform of the face images was used as the feature space. The number of neighbors selected for this set of experiments was uniformly fixed at 50. The improved performance of the Biased Isomap Embedding framework is unanimously reflected in the significant reduction in error values for different image feature spaces across the selected manifold learning techniques. While the results obtained for Isomap show stability across the dimensions of embedding, the estimation of pose angle using traditional LLE and Laplacian Eigenmaps is fairly good, even without the BME framework. However, when the proposed approach was used to obtain the embedding, the error values are lower, especially in case of Laplacian Eigenmaps, where the BME framework provided excellent results. The values for the error in estimation of pose angle is a substantial improvement over earlier work [10].

In addition, the performance of the Biased Manifold Embedding approach was analyzed with varying choices of the number of neighbors used for embedding. Table 3 captures these results with the embedding dimension fixed at 8. These experiments were conducted with Isomap using the grayscale pixel intensity feature space.

| Number of Neighbors | Error using traditional Isomap | Error using Biased Isomap |
|---|---|---|
| 30 | 11.56 | 5.10 |
| 50 | 12.96 | 5.06 |
| 100 | 13.83 | 5.03 |
| 200 | 12.59 | 5.06 |
| 500 | 14.36 | 5.07 |

Table 3. Analysis of performance with varying number of neighbors for embedding

As evident from the results, the significant reduction in the error of estimation of pose angle substantiates the effectivness of the proposed approach. In addition, as the results in Tables 1, 2 and 3 illustrate, the Biased Manifold Embedding method is robust to variations in feature spaces, dimensions of embedding and choice of number of neighbors. While the traditional Isomap embedding has fluctuating results for these parameters, the range of error values obtained for the Biased Manifold Embedding method across these parameter changes suggests the robustness of the method, thanks to the biasing of the embedding.

# 6. Conclusions

We have proposed the Biased Manifold Embedding method, a novel supervised approach to manifold learning techniques for regression problems. The proposed method was validated for accurate person-independent head pose estimation. The use of pose information in the manifold embedding process improved the performance of the pose estimation process significantly. The pose angle estimates obtained using this method are accurate, and can be relied upon with an error margin of $3-4\ ^\circ$, or even lower based on the manifold learning technique used. Our experiments also demonstrated that the method is robust to variations in feature spaces, dimensionality of embedding and the choice of the number of neighbors for the embedding. The proposed method can easily be extended from the current implementations to apply to the envelop of all manifold learning techniques, and has been developed as a framework to cater to all regression problems at large.

## 6.1. Limitations and Future Work

As mentioned earlier, a significant drawback of manifold learning techniques is the lack of a projection matrix to treat new data points. While we used the GRNN to learn the non-linear mapping in this work, there have been other approaches adopted by various researchers. Bengio et al [2] proposed a mathematical formulation focussed to overcome this problem. We plan to use these approaches to support the validity of our approach. In addition, we plan to study the use of different functions of pose distance used to bias the distance matrix to infer the applicability of different reciprocal functions for pose estimation.

# Acknowledgements

This work was supported by the National Science Foundation NSF-ITR grant IIS-0326544.

# References

[1] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15:1373–1396, 2003. 3

| Dimension of Embedding | Error using Traditional Isomap | Error using Biased Isomap | Error using Traditional LLE | Error using Biased LLE | Error using Traditional Laplacian Eigenmap | Error using Biased Laplacian Eigenmap |
|---|---|---|---|---|---|---|
| 3 | 16.21 | 5.66 | 8.47 | 6.59 | 11.69 | 8.21 |
| 5 | 12.57 | 5.05 | 7.34 | 5.56 | 9.35 | 4.59 |
| 8 | 12.96 | 5.07 | 6.78 | 4.18 | 7.61 | 2.52 |
| 20 | 11.35 | 5.04 | 6.04 | 3.31 | 6.32 | 1.51 |
| 50 | 10.86 | 5.04 | 4.37 | 2.56 | 4.57 | 1.47 |
| 100 | 10.41 | 5.02 | 3.27 | 2.11 | 3.93 | 1.44 |

Table 1. Performance with varying dimensions of embedding from grayscale pixel intensity feature space

| Dimension of Embedding | Error using Traditional Isomap | Error using Biased Isomap | Error using Traditional LLE | Error using Biased LLE | Error using Traditional Laplacian Eigenmap | Error using Biased Laplacian Eigenmap |
|---|---|---|---|---|---|---|
| 3 | 11.24 | 4.01 | 9.33 | 9.11 | 12.47 | 7.70 |
| 5 | 10.30 | 3.54 | 8.37 | 5.43 | 8.34 | 3.98 |
| 8 | 9.90 | 3.38 | 7.68 | 4.06 | 7.14 | 2.17 |
| 20 | 9.21 | 3.30 | 6.71 | 3.20 | 6.94 | 1.78 |
| 50 | 8.76 | 3.23 | 5.23 | 2.89 | 5.23 | 1.73 |
| 100 | 8.23 | 3.02 | 4.31 | 2.56 | 4.52 | 1.72 |

Table 2. Performance with varying dimensions of embedding from Laplacian of Gaussian feature space


[2] Y. Bengio, J. F. Paiement, P. Vincent, and O. Delalleau. Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering, 2004. 6

[3] L. Brown and Y.-L. Tian. Comparative study of coarse head pose estimation, 2002. 1

[4] L. Chen, L. Zhang, Y. X. Hu, M. J. Li, and H. J. Zhang. Head pose estimation using fisher manifold learning. In *IEEE International Workshop on Analysis and Modeling of Faces and Gestures (AMFG03)*, pages 203–207, Nice, France, 2003. 2, 3

[5] Y. Fu and T. S. Huang. Graph embedded analysis for head pose estimation. In *7th International Conference on Automatic Face and Gesture Recognition*, Southampton, UK, 2006. 1, 2, 4

[6] N. Hu, W. Huang, and S. Ranganath. Head pose estimation by non-linear embedding and mapping. In *IEEE International Conference on Image Processing*, pages 342–345, Genova, 2005. 2, 3

[7] I.T.Jolliffe. *Principal Component Analysis*. Springer-Verlag, New York, 1986. 2

[8] G. Little, S. Krishna, J. Black, and S. Panchanathan. A methodology for evaluating robustness of face recognition algorithms with respect to variations in pose and illumination angle. In *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 89–92, Philadelphia, USA, 2005. 4, 5

[9] M.W.Trosset. The formulation and solution of multidimensional scaling problems. Technical Report Technical Report TR93-55, Rice University, 1993. 2

[10] B. Raytchev, I. Yoda, and K. Sakaue. Head pose estimation by nonlinear manifold learning. In *17th International Conference on Pattern Recognition (ICPR04)*, Cambridge, UK, 2004. 1, 2, 3, 4, 6

[11] D. d. Ridder, O. Kouropteva, O. Okun, M. Pietikainen, and R. P. Duin. Supervised locally linear embedding. In *International Conference on Artificial Neural Networks and Neural Information Processing*, 333-341, 2003. 2, 4

[12] S. Roweis and L. Saul. Non-linear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000. 2

[13] J. B. Tenenbaum, V. d. Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000. 2

[14] M. T. Wenzel and W. H. Schiffmann. Head pose estimation of partially occluded faces. In *Second Canadian Conference on Computer and Robot Vision (CRV05)*, pages 353–360, Victoria, Canada, 2005. 1

[15] J. Yu and Q. Tian. Learning image manifolds by semantic subspace projection. In *ACM Multimedia*, Santa Barbara, CA, USA, 2006. 2, 4

[16] Q. Zhao, D. Zhang, and H. Lu. Supervised lle in ica space for facial expression recognition. In *International Conference on Neural Networks and Brain (ICNNB05)*, volume 3, pages 1970–1975, Beijing, China., 2005. 5