Illumination and Person-Insensitive Head Pose Estimation Using Distance Metric Learning

Xianwang Wang, Xinyu Huang, Jizhou Gao, and Ruigang Yang

Center for Visualization & Virtual Environments,
University of Kentucky,
Lexington KY 40507, USA
{xwang,xhuan4,jgao5,ryang}@cs.uky.edu

Abstract. Head pose estimation is an important task for many face analysis applications, such as face recognition systems and human computer interactions. In this paper we aim to address the pose estimation problem under some challenging conditions, e.g., from a single image, large pose variation, and un-even illumination conditions. The approach we developed combines non-linear dimension reduction techniques with a learned distance metric transformation. The learned distance metric provides better intra-class clustering, therefore preserving a smooth low-dimensional manifold in the presence of large variation in the input images due to illumination changes. Experiments show that our method improves the performance, achieving accuracy within 2-3 degrees for face images with varying poses and within 3-4 degrees error for face images with varying pose and illumination changes.

1 Introduction

Human face analysis, due to its vast application from biometric authentication to human-computer interactions, is a very active topic in computer vision research. Head pose estimation is a central component for many of these applications. For example, face recognition systems require the capability of handling the significant pose variations. Zhao *et al.* [1] shows that the pose and illumination are the major factors affecting the performance of face recognition algorithms. The difference between individuals' face images taken under the same lighting conditions is smaller than the difference between two face images of same individual taken under varying lighting conditions. That is, image variation due to lighting changes is more significant than variation due to different personal identities [2]. While person-independent head pose estimation has been studied reasonably well in recent years [3,4,5], robust illumination- and person-independent head pose estimation remains a challenging problem. The main contribution of this paper is to introduce a new approach to address this problem.

There exists many methods for pose estimation from a single image (we exclude pose estimation from a video sequence in the context of this paper). These methods can be classified into five categories [4,5]: shape-based geometric analysis methods [6], appearance-based methods [7,8], model-based methods [9],

D. Forsyth, P. Torr, and A. Zisserman (Eds.): ECCV 2008, Part II, LNCS 5303, pp. 624–637, 2008.

template-based methods [10,11], and dimensionality reduction based methods. It is beyond the scope of this paper to discuss all the related work in pose estimation, and we will focus on the fifth category: dimensionality reduction (DR) techniques, which is also what our proposed method belongs to. Every DR approach is essentially to learn a distance metric that removes the data redundancy and leads to more compact feature representation. However, the problem of illumination variation is usually conceded or treated lightly as a pre-processing step. As we shown in Figure 1, we aim to deal with some very harsh illumination conditions. The lack of DR-based pose estimation under these difficult conditions is probably due to the fact that illumination changes are usually larger than that from different persons or small pose variations.

Fig. 1. Sample input images for pose estimation, notice the large variation in rotation and the harsh un-even illumination

In this paper we present a novel approach to estimate the pose from a single input image by *combining both the unsupervised metric learning technique and the supervised metric learning techniques*. Unlike previous DR-based pose estimation method that treats illumination as a part of pre-processing step, we develop a *unified* framework that does not require any pre-processing for illumination normalization or correction. This is possible by applying a *learned* distance transformation after the use of nonlinear DR techniques. This is different from prior approaches in which the original images are modified or filtered before applying nonlinear DR techniques. To the best of our knowledge this is the first time a DR-based method is capable of producing fairly accurate pose estimation (within a few degrees) under harsh illumination conditions as shown in Figure 1.

The rest of the paper is organized as follows. Section 2 reviews the related work on head pose estimation. Section 3 describes our proposed approach of pose estimation. Section 4 presents the experimental results, and the conclusion is made in section 5.

2 Related Works

In this section we review the research work of distance metric learning and head pose estimation approaches based on dimensionality reduction. Many methods have been developed. We discuss each one of them in the following sections.

2.1 Distance Metric Learning

The main idea for methods in this category is to learn metrics that separate data points from different classes, while keeping close data points within the same class. Xing *et al* [12] posed metric learning as a convex programming problem to satisfy the constraints of the better representation. Relevant Components Analysis (RCA) [13] learns a full ranked Mahalanobis distance metric using equivalence constraints. Neighborhood Components Analysis (NCA) [14] maximizes the leave-one-out cross validation to learn a distance metric for KNN classifier. Large Margin Nearest Neighbor(LMNN) [15] extends NCA through a maximum frame work. Local Fisher Discriminant Analysis (LFDA) [16] can be viewed as localized variant of FDA.

2.2 Manifold Learning

The goal for approaches in this category is to learn a low-dimensional manifold in which most "intrinsic information" (e.g., distance) are preserved. Popular approaches include ISOMAP [17], Locally Linear Embedding (LLE) [18], and Laplacian Eigenmaps (LE) [19]. ISOMAP preserves the geodesic inter-point distances, LLE preserves the distance based on locally linear combination of neighborhood, and LE preserves the distance described by a weighted connected graph constructed from neighborhood.

Raytchev *et al* [20] apply ISOMAP-based manifold learning technique for user-independent pose estimation and evaluate their method in comparison with the Linear Subspace and Locality Preserving Projections(LPP) [21].

Chen *et al.* [22] uses the face images of two specific head poses and estimates the head poses between them through classification-based nonlinear interpolation. This approach is based on the assumption that the face images of multiple view lie on a manifold in the original image feature space.

Fu and Huang [5] presented an appearance-based strategy for head pose estimation using supervised Graph Embedding(GE) analysis. The neighborhood weighted graph is first constructed in the sense of supervised LLE. The out-of-sample data points may be treated using the projection transformation solved in closed-form based on GE linearization. The K-nearest neighbor classification is then employed to estimate the head pose. Their method is successful with low pose estimation error. They consider face images with only pose variation, but not illumination change in their experiment.

Balasubramanian *et al.* [3] proposed the Biased Manifold Embedding (BME) framework for head pose estimation. The pose information of the given face

Table 1. performance comparison of different methods for head pose estimation

Method	Interval	Increment	Best result: error	Illumination (Yes/No)
ISOMAP [20]	$[-90^\circ + 90^\circ]$	15°	11°	No
Fisher manifold learning [22]	$[-10^\circ + 10^\circ]$		3°	No
BME with ISOMAP, LLE [3]	$[-90^\circ + 90^\circ]$	2°	3°	Yes$^\diamond$
BME with LE [3]	$[-90^\circ + 90^\circ]$	2°	2°	Yes$^\diamond$
LEA [5]	$[-90^\circ + 90^\circ]$	1°	2°	No

$^\diamond$LoG filter is used.

image data is used to compute a biased neighborhood of each point in the feature space, before determining the low-dimensional embedding. BME uses a Generalized Regression Regression Neural Network (GRNN) to learn the nonlinear mapping for dealing with out-of-sample data points, and applies linear multivariate regression to estimate the pose. Essentially, BME proposes a more general model to modify the distance matrix, which unifies the other supervised manifold learning approaches [23,24]. The interested reader can refer to the discussion in [4] for details. Their experiment shows this method works well with person-independent face images with pose variation. BME uses LoG to remove the illumination effects, but LoG representation are not sufficient for pose estimation under a wide variety of lighting conditions, in particular hash lightings with shadows, as in our experiments.

The performance of representative DR-based approaches for head pose estimation are summarized in Table 1. It shows that LEA and BME are the current state-of-art techniques in pose estimation. But none of them treats illumination variations in a principled way, most of them do not discuss the effect of illumination.

3 Our Approach

In this section, we present our approach for pose estimation. We assume that there is a training face database, and each face image in the database is associated with a pose label. Our goal is to estimate the unknown pose label from an input face image that is not in the training database.

A good low-dimensional embedding for pose estimation should have the following properties: (1) Separation. The embedding from different poses are kept apart, and there is no overlap among them. Furthermore, the embeddings of different individuals with different illuminations but same pose should be close to each other, i.e., within a cluster. (2) Smoothness. The low-dimensional manifold should change smoothly according to the pose. Figure 2 shows an ideal 3-dimensional manifold embedding of 24 subjects with the same illumination while the pose angles vary from -90° to $+90^\circ$ with a granularity of 4 from our training database. In this figure, there are 46 clutters in total, each with unique color corresponding to a specific pose angle. Within each of clutters, there are 24 data points, which are embeddings of 24 faces from the same pose.

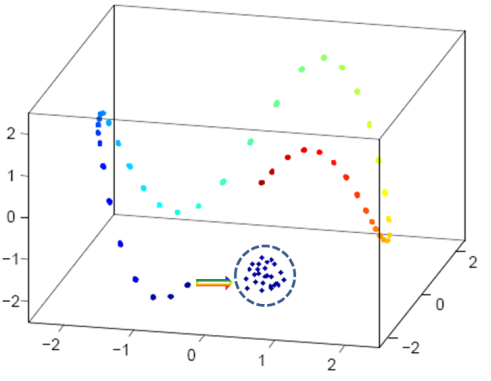

Fig. 2. The ideal 3-dimensional embedding of 24 subjects' face images with only pose variation between $[-90° + 90°]$ at 4° increments. The pose changes are represented by different colors.

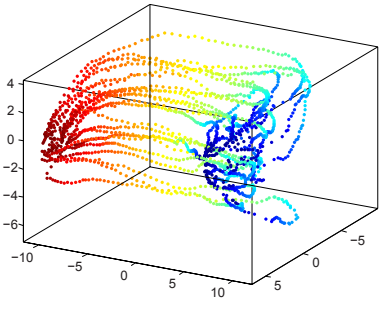

(a) ISOMAP embedding (NN = 200)

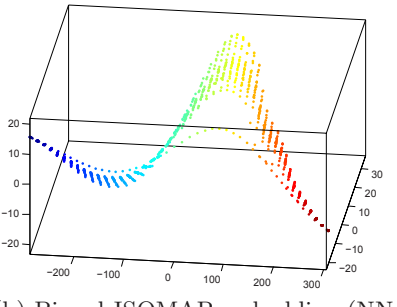

(b) Biased ISOMAP embedding (NN = 200)

Fig. 3. The 3-dimensional embedding of 24 subjects' face images with only pose variation between $[-90° + 90°]$ at 4° increments, using ISOMAP and BME. The pose changes are represented by different colors.

The 3D embedding of ISOMAP and BME with ISOMAP are shown in Figure 3. Figure 3(a) shows the ISOMAP embedding, in which 200 nearest neighbors (NN) are used. It maps face images of 24 subjects into 24 different pose manifolds. This is because ISOMAP can not find the nearest neighbors of each point accurately when there are multiple individuals in the training set. To deal with identity variations, BME finds the right nearest neighbors for each data point by the given pose labels. However, when there are illumination variations, especially large illumination variations, BME can not generate a good pose manifold either. As shown in Figure 4, we generate two manifolds with ISOMAP and biased ISOMAP from 10 subjects with pose angles varying from $-90°$ to $+90°$ with a granularity of 4 and illumination changes from 0° to $+45°$ at 5° increments. It

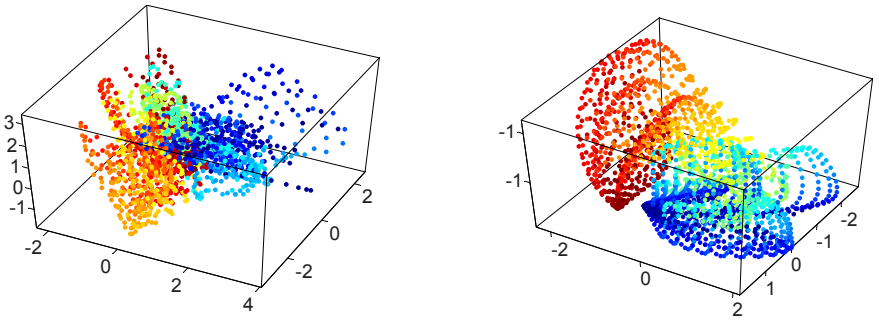

(a) ISOMAP embedding (NN = 200)

(b) Biased ISOMAP embedding (NN = 200)

Fig. 4. The 3-dimensional embedding of 10 subjects' face images with pose variation between $[-90° + 90°]$ at 4° increments and illumination changes from 0° to +45° at 5° increments, using ISOMAP and BME. The pose changes are represented by different colors.

is clear that there are many overlaps between pose angles in ISOMAP embedding (Figure 4(a)), as well as BME embedding (Figure 4(b)). Hence, it is quite difficult to estimate poses from the manifolds generated by ISOMAP and BME. This is because that 1) the computation of a nonlinear manifold relies on the distance between data points. For example, in order to compute a smooth pose manifold, the distance between face images under large illumination variation and the same pose should be small. 2) BME only uses pose labels to find the right nearest neighbors. The distances between data points remain unchanged. 3) the distortion caused by illumination variation is much larger than the distortion caused by the identity of individuals by comparing Figure 4(b) and 3(b). Therefore, the distance between data points should be modified in order to obtain a smooth manifold. This brings forth the need to develop pose estimation techniques that can work well with the face images that have large illumination changes and pose changes from many different individuals.

To obtain a good low-dimensional embedding, we propose an approach based on manifold learning techniques and supervised distance metric learning techniques for head pose estimation. We first construct the low-dimensional embedding using ISOMAP (Figure 4(a)). The low-dimensional embedding is then linearly mapped to the transformed feature space with modifying the distance between data points, using Local Fisher Discriminant Analysis (LFDA) [16] by pose labels. The combination of ISOMAP and Fisher Discriminant Analysis (FDA) was ever proposed in the work [25]. In their work, each data point is represented by a feature vector, which is its geodesic distance from other points. Then, FDA is applied to find an optimal projection direction for classification. The main difference between our approach and this extended ISOMAP is that LFDA was employed to refine the low-dimensional manifold and maintain pose class separation.

3.1 ISOMAP

The classical linear algorithms such as Principle Component Analysis (PCA) and Multidimensional Scaling (MDS) cannot always reveal the intrinsic distribution of a given complex data set. We therefore adopt ISOMAP for nonlinear dimensionality reduction [17]. The input is data matrix $X = (x_1, \cdots, x_N) \in \Re^{D \times N}$ containing N face images from the training data, where $x_i \in \Re^D (i = 1, 2, \ldots, N)$ is the D-dimensional samples of face images. In our case, D is equal to 1024, since we vectorized face images with the resolution 32×32. The output is the nonlinear embedding matrix $Y = (y_1, \cdots, y_N) \in \Re^{d \times N}$ of X, where $y_i \in \Re^d (i = 1, 2, \ldots, N)$ is the d-dimensional data points in the low-dimensional embedding. It first determines the neighbor relationship on the manifold M based on the pairwise Euclidean distance $d_X(i, j)$ between pairs of face images x_i,x_j. These neighbor relations are represented as a weighted graph G over the data points, with the edges of weights $d_X(i, j)$ between neighboring points. The pairwise geodesic distance $d_M(i, j)$ on the manifold are then estimated with the distance of the shortest path $d_G(i, j)$ in the graph G using Floyds or Djkstraas algorithm. The classical MDS is finally applied to the matrix of graph distances $D_G = d_G(i, j)$ to construct d-dimensional embedding Y.

3.2 Local Fisher Discriminant Analysis

We use LFDA to learn the matrix $P_{LFDA} \in \Re^{d \times d}$ that transforms $y_i (i = 1, 2, \ldots, N)$ to $z_i (i = 1, 2, \ldots, N)$. $z_i \in \Re^d (i = 1, 2, \ldots, N)$ is the d-dimensional data points in the transformed feature space. $z_i = P_{LFDA}^T y_i$ in the same pose angle are kept close together, while z_i from different pose angles are well separated.

LFDA evaluates within-class scatter and between-class scatter in a local manner by combining the idea of FDA and LPP [21]. The local within-class scatter matrix $\tilde{S}_B$ and the local between-class scatter matrix $\tilde{S}_W$ are defined as follows,

$$\tilde{S}_W = \frac{1}{2} \sum_{i,j=1}^{N} \tilde{W}_{i,j}^{(w)} (y_i - y_j)(y_i - y_j)^T \tag{1}$$

$$\tilde{S}_B = \frac{1}{2} \sum_{i,j=1}^{N} \tilde{W}_{i,j}^{(b)} (y_i - y_j)(y_i - y_j)^T \tag{2}$$

where

$$\tilde{W}_{i,j}^{(w)} = \begin{cases} W_{i,j}/N_c, & y_i \in c, y_j \in c \\ 0, & otherwise \end{cases} \tag{3}$$

$$\tilde{W}_{i,j}^{(b)} = \begin{cases} W_{i,j}(1/N - 1/N_c), & y_i \in c, y_j \in c \\ 1/N, & otherwise \end{cases}, \tag{4}$$

and $W_{i,j}$ is the affinity between in y_i and y_j that are ranged in $[0, 1]$. The discuss about the definition of $W_{i,j}$ can be found in [21,26]. LFDA considers

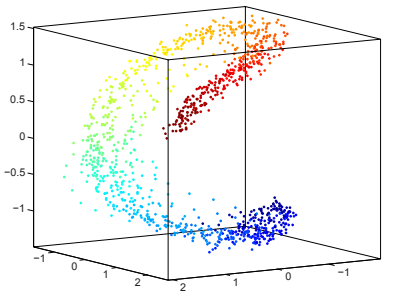

(a) The 3-dimensional embedding of 24 subjects' face images with only pose variation from -90° to $+90^\circ$ at 4° increments using our method (NN =200). The pose changes are represented by different colors.

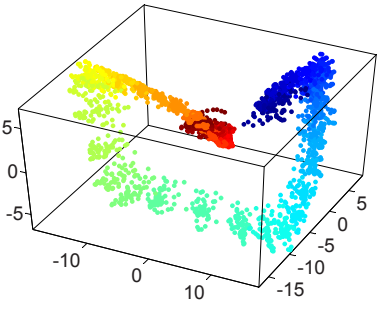

(b) The 3-dimensional embedding of 10 subjects' face images with pose variation from -90° to $+90^\circ$ at 4° and illumination changes from 0° to $+45^\circ$ at 5° increments using our method (NN =200). The pose changes are represented by different colors.

Fig. 5. The results of our method for face images with only pose variation and pose plus illumination variation

the maximization of the following objective to find the transformation matrix P_{LFDA},

$$J(P) = \frac{\|P^T \tilde{S_B} P\|}{\|P^T \tilde{S_W} P\|} \tag{5}$$

Noticing that J is invariant with respect to scale, we can formulate the objective into the constrained optimization problem as follows:

$$\begin{aligned} &\min_P \; -\tfrac{1}{2} P^T \tilde{S_B} P \\ &s.t. \; P^T \tilde{S_W} P = I \end{aligned} \tag{6}$$

The lagrangian corresponds to this optimization problem is,

$$\ell = -\frac{1}{2} P^T \tilde{S_B} P + \frac{1}{2}(P^T \tilde{S_W} P - I) \tag{7}$$

Using the KKT conditions, the problem is transformed into the following generalized eigenvalue problem,

$$\tilde{S_B}\varphi = \lambda \tilde{S_W}\varphi \tag{8}$$

Then, the LFDA transformation matrix is defined by the solution as follows,

$$P_{LFDA} = [\varphi_1, \varphi_2, \cdots, \varphi_d] \tag{9}$$

where $\{\varphi_i\}_{i=1}^d$ are generalized eigenvectors corresponding to the generalized eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_d$.

Algorithm 1. Pose Estimation Pipeline

1: Learning phase
Input: the training face images $x_i(i = 1, 2, \cdots, N)$ and their corresponding labels of pose angles ℓ_i.
Output: the nonlinear mapping between x_i and y_i, P_{LFDA}, the regression model for pose estimation.
(a): Get the low-dimensional embedding y_i using ISOMAMP.
(b): Learn the nonlinear mapping between x_i and y_i using GRNN.
(c): Find the transformation matrix P_{LFDA} from low-dimensional embedding y_i.
(d): Learn the regression model between y_i and its corresponding pose labels using RVM.

2: Testing phase
Input: the test face images $x_i(i = N + 1, N + 2, \cdots, NN)$.
Output: the pose labels for x_i.
(a): Compute its low-dimensional embedding y_i using learned nonlinear mapping.
(b): Map y_i to the transformed feature space using $z_i = P_{LFDA}^T y_i$.
(c): Estimate the pose angle by applying the learned regression model on z_i

We apply our method to the same data that we use in ISOMAP and BME methods to get the low-dimensional embedding, shown in Figure 5. Compared to the results of ISOMAP and BME in Figure 3 and Figure 4, our method is much better in clustering the face images in the same pose angle and separating the face images from different pose angles better than ISOMAP and BME methods in both pose-only variation case and pose+illumination variation case. For a new input face image, we can compute its low-dimensional embedding using the nonlinear mapping learned by Generalized Radial Neural Network (GRNN), and then project it to the transformed feature space by applying the linear transformation. Its pose angle will be finally estimated by Relevance Vector Machine (RVM) [27]. The detailed algorithm procedure using ISOMAP and LFDA for pose estimation is shown in algorithm 1.

4 Experiment and Results

4.1 Data Sets

To evaluate the performance of our approach, we employed the 3D face dataset from [28]. The pose changes horizontally from $-90°$ to $+90°$ at 2 degree increment. The illumination varies from $0°$ to $+45°$ at $1°$ increment, as shown in Figure 6. Other public face databases such as FERET, the CMU-PIE database,Yale Face database, and MIT database, are not used in our experiment, because none of them provide a precise measure for pose and illumination angles and also they do not contain face images with a wide variety of illumination

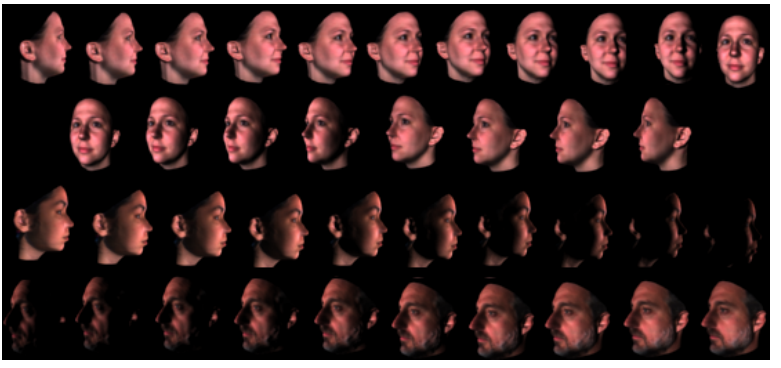

Fig. 6. Samples of face images with varying pose and illumination from 3D face scans

and pose changes [3]. To assess the robustness of our approaches, we perform the experiment in two cases: pose estimation for face images without and with illumination variation. We compared our method with the state-of-art pose estimation techniques, BME. The performance is analyzed with varying choices of the embedding dimensions (marked by "o" in the figures) and 200 neighbors.

4.2 Pose Estimation for Face Images without Illumination Variation

The experiment was performed over 24 subjects with pose angles varying from $-90°$ to $+90°$ at $2°$ increments and illumination of $22°$ using 8-fold cross-

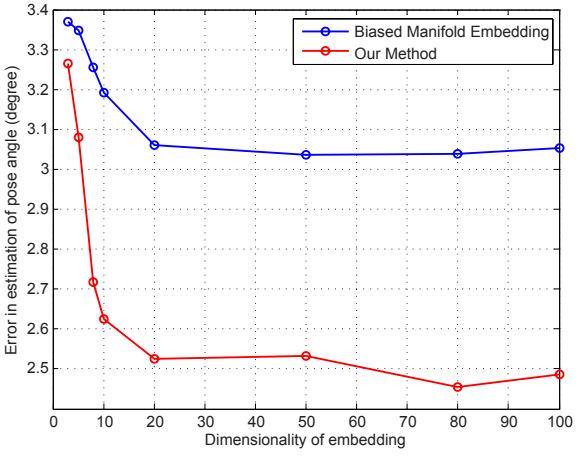

Fig. 7. Pose estimation results comparison of our method against BME with ISOMAP for the face images without illumination variation (NN=200) in different dimensionality. The red line indicates the results of our method.

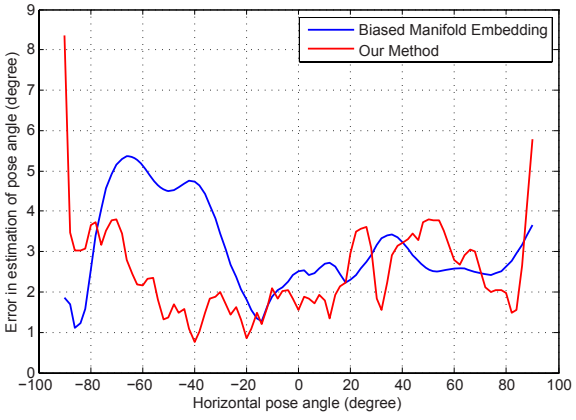

Fig. 8. Distribution of the average error of Our method against the BME framework in pose estimation for the face images without illumination variation. Each of the views is between $[-90°, +90°]$ at $2°$ increments. The red line indicates the results of our method.

validation. We use 1911 face images from 21 subjects (91 images per subject) as the training data in each fold, and then use the 273 images from the other 3 subjects as testing data. The images were down-sampled to 32×32 resolution. The results of the experiment are shown in Figure 7. The red line indicates the performance of our method, while the blue line shows the performance of the BME framework. The result shows the accuracy of our method is slightly better than that of the BME with ISOMAP. Figure 8 shows the head pose estimation error of our method against BME in each of the views in this pose angle interval.

4.3 Pose Estimation for Face Images with Illumination Variation

Due to memory limitation, this experiment was performed over 10 subjects, with pose angles varying from $-90°$ to $+90°$ at $4°$ increments and illumination variation from $0°$ to $-45°$ with a granularity of $5°$. We use leave-one-out cross-validation (LOOCV), i.e., we sequentially take out the face images of one individual and use all the remaining images of the other individuals as training data. The images are also down-sampled to 32×32 resolution. The experiment results are shown in Figure 9. The red line indicates the performance of our method, while the blue line shows the performance of the BME framework. It shows that our method significantly improves the head pose estimation performance compared to the BME framework. Figure 10 shows the head pose estimation error of our method against BME in each of the views in this pose angle interval (100-dimensional embedding) .

In both cases, our method performs well for the frontal and intermediate poses, but not for the profile (almost $-90°$ or $+90°$) views. This may be caused

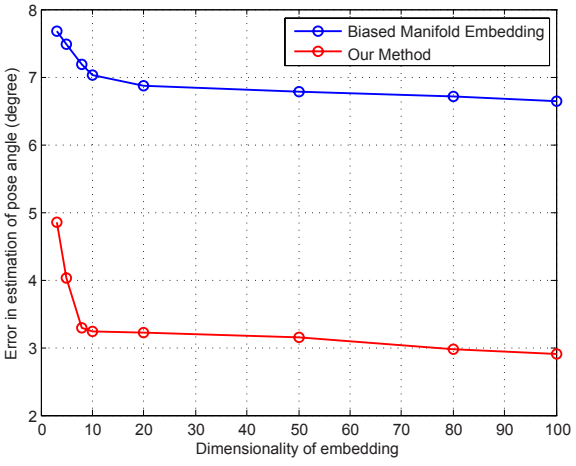

Fig. 9. Pose estimation results comparison of our method against BME with ISOMAP for the face images with illumination variation. The red line indicates the results of our method.

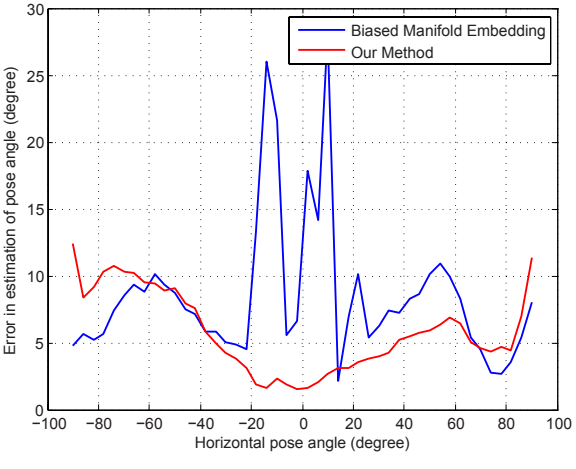

Fig. 10. Distribution of the average error of our method against the BME framework in pose estimation for the face images illumination variation. each of the views is between $[-90°, +90°]$ at $4°$ increments. The red line indicates the results of our method.

by the noisy features in the face images of these profile views, which results in some overlaps between the data points in the transformed feature space.

Comparing our results with those listed in Table 1, we can see that our method is comparable with the the state-of-art methods for face images with only pose variation. When the illumination also varies, our proposed method maintains the high accuracy for most cases and performs better that BME.

5 Conclusions

In this paper, we proposed an approach to illumination- and person-insensitive head pose estimation. We studied the limitation of related approaches in pose estimation for face images with large illumination variation, and addressed the problem by combining ISOMAP and LFDA. We conducted several experiments to evaluate our approach. The experiment results demonstrate that our method is robust to variation in dimensionality of embedding, illumination and identity of individuals. Our method can be easily extended to other manifold learning techniques, such as LLE and LE, and supervised distance metric learning techniques like RCA, NCA, and LMNN. Furthermore, our method can be used to estimate the illumination direction by using illumination labels in LFDA. Looking into the future, we will apply our approach in the application of face synthesis.

References

1. Zhao, W., Chellappa, R., Rosenfeld, A., Phillips, P.: Face recognition: A literature survey (2003)
2. Adini, Y., Moses, Y., Ullman, S.: Face recognition: The problem of compensating for changes in illumination direction. IEEE Transactions on Pattern Analysis and Machine Intelligence 19(7), 721–732 (1997)
3. Balasubramanian, V., Ye, J., Panchanathan, S.: Biased Manifold Embedding: A Framework for Person-Independent Head Pose Estimation. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–7 (2007)
4. Balasubramanian, V.N., Krishna, S., Panchanathan, S.: Person-Independent Head Pose Estimation using Biased Manifold Embedding. EURASIP Journal on Advances in Signal Processing (accepted, 2007)
5. Fu, Y., Huang, T.S.: Graph Embedded Analysis for Head Pose Estimation. In: International Conference on Automatic Face and Gesture Recognition, pp. 3–8 (2006)
6. Hu, Y., Chen, L., Zhou, Y., Zhang, H.: Estimating Face Pose by Facial Asymmetry and Geometry. In: IEEE International Conference on Automatic Face and Gesture Recognition, pp. 651–656 (2004)
7. Ba, S.O., Odobez, J.M.: A Probabilistic Framework for Joint Head Tracking and Pose Estimation. In: International Conference on Pattern Recognition, pp. 264–267 (2004)
8. Li, S.Z., Lu, X., Hou, X., Peng, X., Cheng, Q.: Learning Multiview Face Subspaces and Facial Pose Estimation using Independent Component Analysis. IEEE Transactions on Pattern Analysis and Machine Intelligence 14(6), 705–712 (2005)
9. Cootes, T.F., Edwards, G., Taylor, C.J.: Active Appearance Models. IEEE Transactions on Pattern Analysis and Machine Intelligence 23(6), 681–685 (2001)
10. McKenna, S., Gong, S.: Real-Time Face Pose Estimation. Real-Time Imaging, Special Issue on Visual Monitoring and Inspection 4(5), 333–347 (1998)
11. Bichsel, M., Pentland, A.: Automatic Interpretation of Human Head Movements. In: International Joint Conference on Artificial Intelligence, Workshop on Looking At People (1993)
12. Xing, E., Ng, A., Jordan, M., Russell, S.: Distance metric learning, with application to clustering with side-information (2003)

13. Bar-Hillel, A., Hertz, T., Shental, N., Weinshall, D.: Learning a Mahalanobis Metric from Equivalence Constraints. Journal of Machine Learning Research 6, 937–965 (2005)
14. Goldberger, J., Roweis, S.T., Hinton, G.E., Salakhutdinov, R.: Neighbourhood components analysis. In: NIPS (2004)
15. Weinberger, K., Blitzer, J., Saul, L.: Distance metric learning for large margin nearest neighbor classification. In: Advances in Neural Information Processing Systems, vol. 18, pp. 1473–1480 (2006)
16. Sugiyama, M.: Local fisher discriminant analysis for supervised dimensionality reduction. In: International Conference on Machine Learning (2006)
17. Tenenbaum, J.B., de Silva, V., Langford, J.C.: A Global Geometric Framework for Nonlinear Dimensionality Reduction. Science 290(5500), 2319–2323 (2000)
18. Roweis, S.T., Saul, L.K.: Nonlinear Dimensionality Reduction by Locally Linear Embedding. Science 290(5500), 2323–2326 (2000)
19. Belkin, M., Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation (2003)
20. Raytchev, B., Yoda, I., Sakaue, K.: Head Pose Estimation by Nonlinear Manifold Learning. In: International Conference on Pattern Recognition, pp. 462–466 (2004)
21. He, X., Niyogi, P.: Locality preserving projections. In: Advances in Neural Information Processing Systems, vol. 16 (2003)
22. Chen, L., Zhang, L., Hu, Y., Li, M., Zhang, H.: Head pose estimation using fisher manifold learning. In: Proceedings of the IEEE International Workshop on Analysis and Modeling of Faces and Gestures (AMFG 2003), pp. 203–207 (2003)
23. Ridder, D.D., Kouropteva, O., Okun, O., Pietikainen, M., Duin, R.: Supervised locally linear embedding (June 2003)
24. Li, C.G., Guo, J.: Supervised isomap with explicit mapping (2006)
25. Yang, M.H.: Extended isomap for classification. In: Proceedings of the 16th International Conference on Pattern Recognition, vol. 3, p. 30615 (2002)
26. Zelnik-Manor, L., Perona, P.: Self-tuning spectral clustering. In: NIPS (2004)
27. Tipping, M.: The relevance vector machine. In: Advances in Neural Information Processing Systems, San Mateo, CA. Morgan Kaufmann, San Francisco (2000)
28. Blanz, V., Vetter, T.: A Morphable Model for the Synthesis of 3D Faces. In: SIGGRAPH, pp. 187–194 (1999)