An Effective Approach to Pose Invariant 3D Face Recognition

Dayong WANG
Nanyang Technological University

Steven C. H. HOI
Singapore Management University, chhoi@smu.edu.sg

Ying HE
Nanyang Technological University

Citation

WANG, Dayong; HOI, Steven C. H.; and HE, Ying. An Effective Approach to Pose Invariant 3D Face Recognition. (2011). *Advances in Multimedia Modeling: 17th International Multimedia Modeling Conference, MMM 2011, Taipei, Taiwan, January 5-7, 2011, Proceedings, Part I*. 6523, 217-228.
Available at: https://ink.library.smu.edu.sg/sis_research/2357

An Effective Approach to Pose Invariant 3D Face Recognition

Dayong Wang, Steven C.H. Hoi, and Ying He

School of Computer Engineering, Nanyang Technological University, Singapore
{S090023, chhoi, yhe}@ntu.edu.sg

Abstract. One critical challenge encountered by existing face recognition techniques lies in the difficulties of handling varying poses. In this paper, we propose a novel pose invariant 3D face recognition scheme to improve regular face recognition from two aspects. Firstly, we propose an effective geometry based alignment approach, which transforms a 3D face mesh model to a well-aligned 2D image. Secondly, we propose to represent the facial images by a Locality Preserving Sparse Coding (LPSC) algorithm, which is more effective than the regular sparse coding algorithm for face representation. We conducted a set of extensive experiments on both 2D and 3D face recognition, in which the encouraging results showed that the proposed scheme is more effective than the regular face recognition solutions.

1 Introduction

Face recognition, an important biometrics technique, plays a critical role in many real-world multimedia applications. Despite being studied extensively in literature [1, 9], existing face recognition techniques still suffer from a lot of challenges when being applied in real-world applications. In particular, many 2D face recognition approaches work excellently under well-controlled conditions (well-posed and good lighting), but their recognition accuracy often decreases considerably when handling real-world face recognition tasks where variations are common for pose, illumination and expression [13, 18].

On the other hand, along with the advances of various 3D capture devices, 3D face recognition techniques are receiving more and more research attention [1, 5]. The highly detailed 3D mesh data can capture rich information, which potentially provides much more clues to tackle some unsolved challenges in 2D face recognition tasks, especially for pose and illumination variations.

Following this direction, in this paper, we investigate a novel 3D face recognition scheme that addresses the open challenge of *Pose Invariant Face Recognition*(PIFR). In particular, we propose an effective approach to tackling the pose invariant 3D face recognition task, which is equipped with a set of effective 3D parametrization, alignment, and spares feature representation techniques.

Specifically, the main contributions of this paper include:

- We propose a new 3D face recognition approach to pose invariant face recognition, which employs a state-of-the-art 3D parameterization technique to resolve the challenge of pose invariant face alignment.

– We propose a new Locality Preserving Sparse Coding (LPSC) algorithm for facial image feature representation, which is empirically more effective for face recognition than the regular sparse coding method [10].

The rest of this paper is organized as follows. Section 2 reviews related work on PIFR and sparse coding. Section 3 presents an overview of the proposed PIFR system. Section 4 discusses a geometry-based face alignment approach by applying an effective 3D mesh parameterization technique. Section 5 presents the proposed novel LPSC algorithm. Section 6 presents an extensive set of empirical studies for performance evaluation, and Section 7 concludes this work.

2 Related Work

Below we review two major groups of related work: pose invariant face recognition and sparse coding techniques.

Pose Invariant Face Recognition To attack a pose-invariant face recognition task, one possible remedy approach is to capture multi-view face images from each individual and estimate all the other possible pose positions. However, it is often not practical to collect multi-view images for each individual in real applications. As a result, the virtual view synthesis scenarios, which base on *2D pose transformation* or *3D face reconstruction*, are proposed to substitute the demand of real views from limited known views(i.e. only the frontal view in our framework) [17].

The 2D pose transformation schemes, such as active shape model (ASM) and active appearance models(AAM) [6], have been demonstrated to handle the PIFR problems effectively within small-scale pose variation. Unfortunately, they often fail for large-scale in-depth face rotation (i.e. larger than 45°) because of the image discontinuities. By utilizing the local feature instead of the whole image, some transformation algorithms could partially overcome the former limitation and further boost the performance. For example, Prince et al. [12] proposed a statistical linear model, called "tied" factor analysis model(TFA), which constructs a one-to-many mapping from the "identify" space to the observed image space with the pose-contingent linear transformation. Comparing with the state-of-the-art 3D face reconstruction approach, they achieved comparative experimental results with 14 manually-identified keypoints on each face. Their method however falls short in very intensive computational costs using local features, and their performance often highly depends on the benchmark point detection or even lots of manual labeling efforts.

The 3D face reconstruction schemes have the potentiality to overcome the pose variance challenge and achieve satisfactory results. However, the 3D reconstruction schemes are complex to implement and extremely computationally expensive because of the slow 3D face modeling process.

With the rapid improvement in 3D capture devices, face recognition techniques, based on 3D data directly, are receiving more and more research attention. It is potentially promising technique to overcome this challenge by exploiting the internally invariance of viewpoint and illumination in the 3D scanned data. Among these works, the techniques, such as extended Gaussian images(EGI), ICP matching, hausdorff distance matching and so on, are proposed for 3D shape recognition. Multi-model approaches, which combine the 2D and 3D results, are also developed to enhance the performance [5].

However, existing 3D face recognition methods rarely take advantage of sophisticated recognition algorithms in image domain.Besides, although 3D faces have the internal pose-invariance characters, incomplete partial faces often bring a lot of difficulties to most existing 3D algorithms.

Sparse Coding Sparse coding aims to represent each input instance $\mathbf{x} \in \mathbb{R}^d$ using a set of basis vectors $\{\mathbf{b}_j \in \mathbb{R}^d, j = 1, \ldots, n\}$ with a sparse coefficient vector $\mathbf{s} \in \mathbb{R}^n$, such that $\mathbf{x} \approx \sum_j \mathbf{b}_j s_j$. Let us denote by matrix $X \in \mathbb{R}^{d\times N}$ for a set of N input data instances, denoted by $B = [B_1, B_2, ..., B_n] \in \mathbb{R}^{d\times n}$ the basis matrix, and denoted by $S \in \mathbb{R}^{n\times N}$ the sparse coefficient matrix, then the sparse codes problem could be formulated as the following optimization problem [10]:

$$\min_{B,S} \frac{1}{2}\|BS - X\|_F^2 + \lambda \sum_{i,j} \phi(S_{ij}) \quad s.t. \ \|B_i\|^2 \leqslant c, \forall i = 1, ..., n. \tag{1}$$

where $\|\cdot\|_F$ denotes the Frobenius norm of a matrix, $\phi(\cdot)$ is a sparsity penalty function (a typical choice is an L_1 penalty function, i.e., $\phi(\mathbf{s}_j) = \|\mathbf{s}_j\|_1$), c is a constant, and λ is a parameter to balance tradeoff between *fitness* and *sparsity*.

In recent years, a variety of sparse coding algorithms have been proposed to improve sparse coding, such as the efficiency and optimization issues [10]. Meanwhile, spare coding is also widely applied for many applications. For example, sparse coding has been applied to face recognition tasks in literature [15]. Finally, we note that there are also a lot of emerging studies that attempt to improve the performance of sparse coding techniques by different ways [4, 7]

3 The Proposed Pose Invariant 3D Face Recognition

We first present the system architecture of the proposed Pose Invariant Face Recognition scheme. Figure 1(a) shows the system flow of the proposed PIFR solution. We discuss the details of each step below.

Capturing Faces we install a 3D camera to capture raw 3D facial mesh data and 2D facial images at the same time. The 2D facial images are mainly used by conventional 2D image based face recognition methods as baseline for empirical comparison.

Face Extraction We have developed some automated tools to detect and extract facial regions from the raw data using the state-of-the-art AAM algorithm [6];

Parameterization We adopt the state-of-the-art Inverse Curvature Map (ICM) [16] for 3D parameterization, which will be further discussed in Section 4.

Feature Extraction We propose a novel Locality Preserving Sparse Coding (LPSC) algorithm, which can extract potentially more salient facial features for the recognition tasks. The details will be discussed in Section 5.

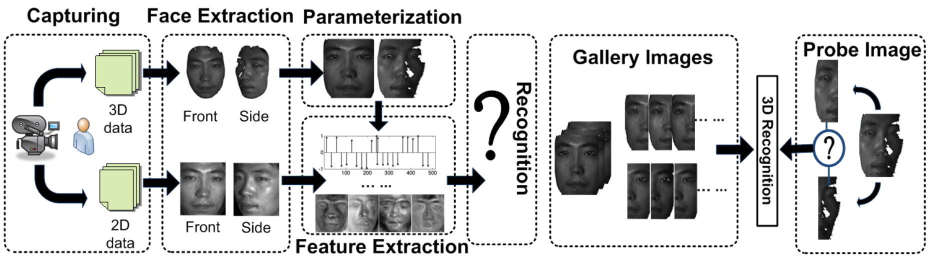

(a)Process of PIFR (b)Probe face identification

Fig. 1. The system flow of the proposed Pose Invariant Face Recognition scheme

Recognition We simply employ the linear SVM for classification task. We believe it is possible to further boost the performance of our system by kernel based SVM, but the classifier design is out of our focus in this paper.

Recognition of Unseen Faces The assumptive situation for our recognition task is : there are only front-view 3D mesh faces for each individual in gallery and the probe faces are captured from a very different pose.

In our approach, since we use only one single 3D camera, it is impossible to capture the complete 3D mesh for the side-view faces. However, we are aware a fact that the 3D camera can always capture at least half of the complete human face no matter how the person rotates his/her head (within 90°). Motivated by this practical trick, we propose a half-face recognition approach that automatically finds the complete half-face for prediction, which empirically works effectively, as shown in Figure 1(b). The left-hand side shows the gallery images where each full face is cut to two pieces. During the recognition phase, given an input novel face for prediction, our system automatically generates the complete half-face on-the-fly, and employs it for performing the recognition.

4 Geometry Alignment via 3D Mesh Parametrization

In this section, we present a geometry-based face alignment by applying an effective 3D mesh parameterization technique to face recognition applications. It is important to note that the captured 3D faces may have holes and very different boundaries due to the various poses and occlusions. Since the conventional harmonic map based parameterization highly depends on the boundary, they can hardly be used for the pose-invariant 3D face recognition.

To parameterize the 3D faces, we adopt the inverse curvature map (ICM) [16], which will minimize both the angle distortion and the area distortion by finding the best discrete conformal mapping. The reason is two-fold: firstly, ICM is a boundary free method, thus well-suited for the captured 3D faces with irregular boundaries and even holes. Second, ICM is an intrinsic curvature diffusion approach that only depends on the first fundamental form, i.e., the edge length of the input mesh.

Figure 2 shows a series of 2D faces and parameterized 3D faces for comparison, where the first row shows a series of 2D facial images, the second row shows a series of 3D facial modes, and the last row shows the corresponding aligned facial images by 3D face parameterization, each of which is overlapped with a aligned front-view image as the background image. The alignment results indicate that the geometry-based alignment can effectively align the 3D mesh into a well-posed 2D domain.

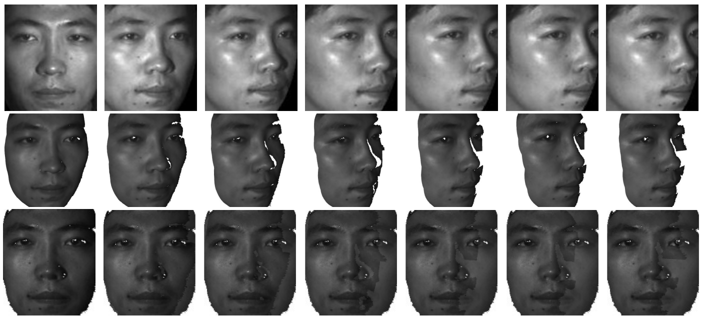

Fig. 2. Examples of facial images used in our experiment. The first row shows a series of 2D facial images; the second row shows a series of 3D facial modes; and the third row shows the corresponding aligned facial images by 3D face parameterization, each of which is overlapped with a aligned front-view image as the background image.

5 Locality Preserving Sparse Coding for Facial Images

In this section, we roughly introduce the new proposed spare coding scheme, named Locality Preserving Sparse Coding (LPSC). The **LPSC** is designed to address some limitation of the existing sparse coding technique and enhance its performance for face recognition.

Formulation One key limitation of the existing sparse coding method is that each input vector has been treated equally and independently without exploiting the input data dependency. On the other hand, face images are widely considered to reside on a non-linear submanifold space. This assumption could be used as priori knowledge and explicitly included in spare coding algorithm.

In order to capture the dependency in input instances, we introduce the following regularizer $g(S, W)$ which measures the inconsistency between the learned sparse code representation S and the weight matrix W of the input patterns, following the manifold regularization approach [3]:

$$g(S, W) = \frac{1}{2} \sum_{i,j=1}^{N} W_{ij} \|\mathbf{s}_i - \mathbf{s}_j\|^2 = tr(SLS^\top) \tag{2}$$

where $tr(\cdot)$ is the *trace* function, and $L = D - W$, $D = diag(d_1, \ldots, d_N)$ is the degree matrix with the diagonal elements defined as $d_i = \sum_{j=1}^{N} w_{ij}$.

Using the above regularizer, we can modify the original optimization problem of sparse coding as follows:

$$\min_{S,B} \frac{1}{2}\|BS - X\|_F^2 + \lambda_1 \sum_{i=1}^{N} \|\mathbf{s}_i\|_1 + \lambda_2 tr(SLS^\top) \quad s.t. \ \|B_i\|^2 \leqslant c, \forall i = 1, ..., n. \tag{3}$$

where λ_1 and λ_2 are two regularization parameters.

In order to solve the above optimization problem, we separate the learning process of Locality Preserving Sparse Coding into two optimization tasks: (1) *Coefficient Learning*, i.e., find the solution of S by fixing the dictionary B; and (2) *Dictionary Learning*, i.e., find the solution of dictionary B by fixing S.

Coefficients Learning By fixing the dictionary B, the optimization in Eq. (3) reduces to a convex optimization. In our approach, we employ a coordinate descent approach for solving the optimization iteratively. In particular, we iteratively optimize only one of the N coefficient vectors $\mathbf{s}_i$ by leaving the other coefficient vectors intact at one time, and repeat until convergence is arrived.

Specifically, consider iteration t, the solution to $\mathbf{s}_i^{(t)}$ can be found by solving the following optimization:

$$\arg\min_{\mathbf{s}_i^{(t)}} \frac{1}{2}\|B\mathbf{s}_i^{(t)} - \mathbf{x}_i\|_2^2 + \lambda_1\|\mathbf{s}_i^{(t)}\|_1 + \lambda_2 L_{ii}\|\mathbf{s}_i^{(t)}\|_2^2 + 2\lambda_2 L_i Z^\top \mathbf{s}_i^{(t)} \tag{4}$$

where $L_i^\top \in \mathbb{R}^{N-1}$ is the i-th vector of L by removing the element L_{ii}, and $Z \in \mathbb{R}^{n\times(N-1)}$ is the sub-matrix of $S^{(t)}$ by removing its i-th column vector.

The above optimization is known as a nonsmooth L1 minimization problem. In our approach, we develop an efficient optimization algorithm to solve the above problem by adapting the state-of-the-art non-smooth convex optimization technique proposed in [11], which is able to achieve a fast convergence of $O(1/t^2)$.

Dictionary Learning Once the coefficient matrix S is found, the other task is to learn the dictionary B given the matrix S. Specifically, when S is given, the optimization of (3) can be reduced to the following:

$$\min_{B} \|X - BS\|_F^2 \qquad s.t. \ \|B_i\|^2 \leqslant c, \forall i = 1, ..., n. \tag{5}$$

The above optimization is essentially the same as the Dictionary learning task of the regular sparse coding method. In our approach, we adopt the existing dictionary learning algorithm proposed in [10].

Out-of-Sample Coding Consider a set of training data points $X^{(train)} \in \mathbb{R}^{d\times N}$, by applying the previous algorithm we are able to obtain the optimal dictionary matrix $B^* \in \mathbb{R}^{d\times n}$ and the coefficients matrix $S^* \in \mathbb{R}^{n\times N}$ for the training data. Given an unseen/test data instance $x^{(test)} \in \mathbb{R}^d$, putting $x^{(test)}$ together with the collection of training data points, we can update the Laplacian matrix $L' \in \mathbb{R}^{(N+1)\times(N+1)}$ with $[X^{(train)}, x^{(test)}] \in \mathbb{R}^{n\times(N+1)}$ to capture the dependency between the test instant and training instants.

To find the coefficient vector $\hat{s}$ without modifying the existing coefficients S^*, we can simply solve the following optimization:

$$\arg\min_{\mathbf{s}} \frac{1}{2}\|B^*\mathbf{s} - \mathbf{x}^{(test)}\|_2^2 + \lambda_1\|\mathbf{s}\|_1 + \lambda_2 L_{(N+1)(N+1)}\|\mathbf{s}\|_2^2 + 2\lambda_2\mathbf{v}(S^*)^\top\mathbf{s}$$

where $\mathbf{v} = \{v_1, v_2, ..., v_N\}$ is the $(N+1)$-th row of L' by removing the last element $L'_{(N+1)(N+1)}$.

6 Experiments

6.1 Experimental Testbed

In our database, we collected totally 10 individuals with different head positions, including the front-view faces and side-view faces with rotation angle varying from 10° to 75° in-depth. Figure 2 already shows some example images of varied poses in our database.

6.2 Evaluation of LPSC for 2D Face Recognition

This section mainly aims to examine the efficacy of the proposed LPSC algorithm. To this purpose, we apply our algorithm to the benchmark face recognition task using the well-known ORL face database and YALE face database.

The ORL database contains 400 facial images from 40 individuals and the YALE face database contains 160 facial images from 15 individuals with different expression and illumination condition. For comparison, we adopt four famous baseline techniques: the Sparse Coding (SC) method [10], the LDA method [2], the LPP method [8] and the PCA method [14].

For the ORL database, we build the training sets by randomly choosing M (2,4,6,8) images for each individual, and putting the rest facial images to form the test set. For the parameter selection, for each data partition, the best feature dimension of LDA, LPP and PCA are chosen by maximizing the classification performance on the training set. For both SC and LPSC methods, we chose the number of basis vector as 1.5 times the input data dimension to balance the sparsity and computation complexity. Similar, the other parameters for SC and LPSC, and the cost parameter C for linear SVM, are found by validating their classification performance on the training set.

M	LPSC	SC	LDA	LPP	PCA
2	0.7981 ±0.015	0.7926 ±0.014	0.7770 ±0.012	0.7800 ±0.013	0.7134 ±0.009
4	0.9314 ±0.016	0.9176 ±0.021	0.9000 ±0.021	0.9030 ±0.017	0.8800 ±0.009
6	0.9680 ±0.012	0.9527 ±0.015	0.9400 ±0.012	0.9410 ±0.015	0.9225 ±0.012
8	0.9867 ±0.012	0.9589 ±0.020	0.9520 ±0.014	0.9550 ±0.020	0.9300 ± 0.008

Table 1. Result of ORL database

M	LPSC	SC	LDA	LPP	PCA
3	0.7006 ±0.018	0.6889 ±0.020	0.6828 ±0.033	0.6889 ±0.043	0.6111 ±0.048
5	0.7963 ±0.038	0.7756 ±0.036	0.7444 ±0.042	0.7502 ±0.035	0.6867 ±0.033
7	0.8478 ±0.023	0.8122 ±0.035	0.7933 ±0.043	0.8011 ±0.043	0.7733 ±0.039
9	0.8711 ±0.038	0.8370 ±0.033	0.8063 ±0.053	0.8133 ±0.054	0.7919 ±0.055

Table 2. Result of YALE database

Figure 3(a) shows the average experimental results over ORL database with varied number of training images for each individual(M). Table 1 illustrates the mean value and the standard deviation in details. For the YALE database, we conduct the similar experimental evaluation with varied $M = 3, 5, 7, 9$. The experiment results of YALE database are shown in Figure 3(b) and Table 2.

Several observations can be drawn from the results. First of all, for both of the two databases the proposed LPSC method achieved the best overall performance among all the compared methods for varied M values. For example, considering the case of $M = 8$ for the ORL database, the average accuracy achieved by LPSC is about 98.67%, which is much higher than the others, including SC (95.89%), LDA(95.20%), LPP(95.50%) and PCA(93.00%). Second, it is evident that the larger the number of training images (M), the better the recognition performance achieved by all the compared methods. Finally, the performance difference between LPSC and SC becomes more significant when increasing the value of M, which indicates that the larger the training size, the more data dependency information can be exploited by LPSC.

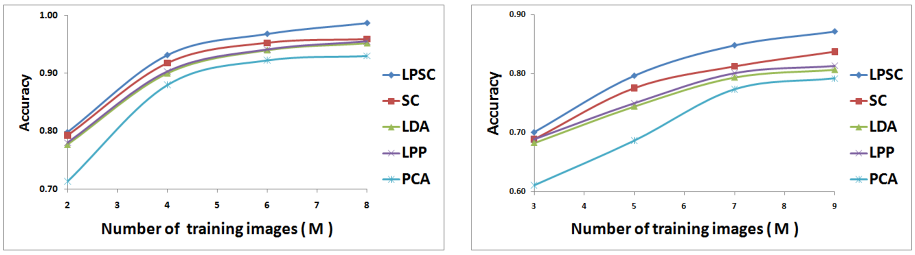

(a) (b)

Fig. 3. Experimental results of facial image recognition on the ORL database and YALE database. X-axis value M denotes the number of randomly selected training images from each person. (a) Results of ORL database (b) Results of YALE database.

As a summary, the above results showed that the proposed LPSC method can learn more effective features for improving the face recognition tasks.

6.3 Evaluation of Pose Invariance Face Recognition : 2D vs. 3D

To evaluate the performance of the proposed pose invariant 3D face recognition system, we compare our solution with regular 2D face recognition in two settings: (1) **Front-Front** recognition task, and (2) **Front-Side** recognition task.

Task I: Front-Front Recognition For the Front-Front task, our PIFR system is similar to a regular 2D face recognition system because no pose variance should be exploited. Our goal is to evaluate whether the parameterized faces have the same efficiency as the regular front-view face recognition.

In our experiments, we follow the experiment setting in the Section 6.2. The same five algorithms (LPSC, SC, LDA, LPP and PCA) are engaged for comparison following the similar experimental scheme. We perform the recognition task over 2D and 3D front-view images respectively. Table 3 shows the average recognition accuracy and their standard deviations achieved by different algorithms. Figure 4 further illustrates these results.

Two observations can be drawn from these results. First of all, for both 2D faces and 3D faces, the accuracies of the Front-Front recognition tasks are very high. There is no significant difference between 2D and 3D approaches according to statistical t-test. Second, similar to the previous observations, LPSC performs the best for both 2D cases and 3D cases among all the compared methods.

As a summary, for the Front-Front task, no significant difference exists between regular face recognition systems and our PIFR system, which shows that our solution is comparable to the regular 2D recognition systems for simple and easy face recognition tasks without pose variations.

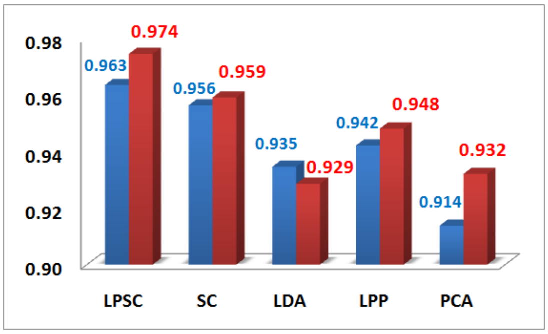

Fig. 4. Comparison of 2D and 3D Front-Front recognition results.

	2D	**3D**
LPSC	0.9634	0.9744
	±0.033	±0.036
SC	0.9562	0.9590
	±0.033	±0.029
LDA	0.9345	0.9285
	±0.043	±0.032
LPP	0.9420	0.9480
	±0.042	±0.037
PCA	0.9137	0.9320
	±0.033	±0.020

Table 3. Comparison of 2D and 3D Front-Front Face Recognition.

Task II: Front-Side Recognition This task assumes that there are only a small size of front-view 2D or 3D images in the gallery (5 images for each person). On the other hand, the test faces consist of pose-variant images with large rotation angles varying from 10°to 75°(about 5 to 10 images for each person).

In this experiment, we engage the same five algorithms for both 2D and 3D recognition and found their parameters similar to the former experiments. Moreover, for the 3D face recognition, we train two different models for the left and right half-faces, respectively. Further, our system can automatically detect the complete half-face corresponding to the probe partial 3D faces, so as to choose the corresponding model for recognition.

In order to evaluate pose tolerance ability of our system, we also compare our scheme with the state-of-the-art algorithm (TFA) [12] based on the image space for fairness, because the former five algorithms all utilize the whole images as input instances. The 2D face images are broadly divided into 7 poses (0°, ±10°, ±30° and ±75°) for model training and recognition. The further experiment follows the description of the experiment section in [12].

Figure 5 shows the comparison of 2D and 3D Front-Side recognition for $M = 4$. Several observations can be drawn. First of all, for this challenging task, the regular face recognition system performs poorly, with the best accuracy of about 50.00%. Second, the state-of-the-art pose-invariant TFA algorithm seems to be fairly effective for this challenging pose-invariant face recognition task, with an average recognition accuracy of about 80%.

Finally, the proposed PIFR system achieved a much higher accuracy than the TFA algorithm. In particular, the accuracy of our system is about 83% with PCA, which is further boosted to 90% with the proposed LPSC algorithm. Similar to the TFA study [12], we believe that the performance of our system could be further enhanced by introducing local features that are less sensitive to pose rotations.

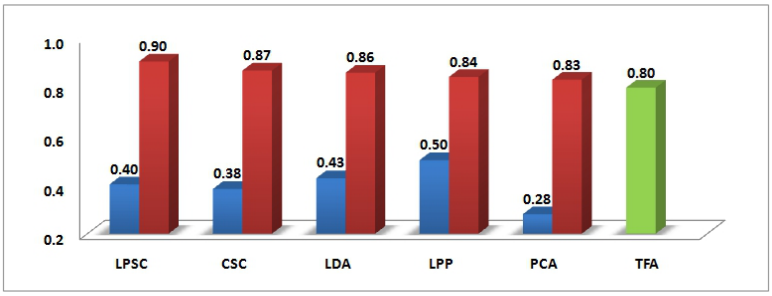

Fig. 5. Comparison of 2D and 3D Front-Side recognition results ($M = 4$). The blur and red bars correspond to 2D and 3D results, respectively, and the green bar represents the result by TFA.

To evaluate the performance of our PIFR system under different m values (from 1 to 4), Figure 6 and Table 4 give more comparison results. The similar observations further confirm the efficacy of LPSC.

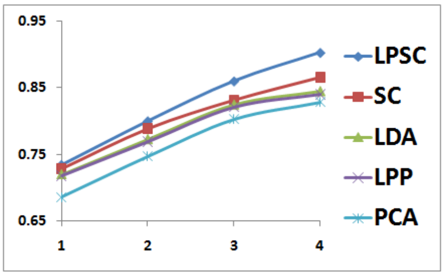

Fig. 6. Results of using different numbers of front face images per individual

M	LPSC	SC	LDA	LPP	PCA
1	0.7343	0.7286	0.7199	0.6977	0.6857
	±0.020	±0.014	±0.028	±0.018	±0.014
2	0.8000	0.7886	0.7725	0.7589	0.7468
	±0.005	±0.015	±0.016	±0.010	±0.014
3	0.8600	0.8314	0.8242	0.8257	0.8029
	±0.012	±0.006	±0.016	±0.014	±0.016
4	0.9029	0.8657	0.8443	0.8400	0.8286
	±0.006	±0.014	±0.011	±0.020	±0.006

Table 4. Results of 3D Front-Side recognition task with different front images.

As a summary, our empirical results show that the proposed PIFR system can effectively handle the challenging pose variance problem and achieve a comparative result with the state-of-the-art pose-invariance face recognition scheme. By combining the proposed LPSC feature extraction algorithm, its accuracy could be further significantly boosted.

7 Conclusions

This paper proposed a new and effective pose-invariant 3D face recognition scheme equipped with a set of effective parameterization, alignment and feature representation techniques. Our extensive experiments demonstrated the proposed scheme is effective and promising for tackling the PIFR task and achieve better results than the state-of-the-art TFA algorithm. For future work, we will investigate 3D recognition techniques to address the other challenging face recognition tasks, such as the variations of illumination and expression issues, making our 3D face recognition system practical for real-world applications.

Acknowledgement

The work was supported by the Singapore National Research Foundation Interactive Digital Media R&D Program, under research grant NRF2008IDM-IDM004-006.

References

1. Abate, A.F., Nappi, M., Riccio, D., Sabatino, G.: 2d and 3d face recognition a survey. Pattern Recognition Letters 28(14), 1885–1906 (October 2007)
2. Belhumeur, P., Hespanha, J., Kriegman, D.: Eigenfaces vs. fisherfaces: recognition using class specific linear projection. Pattern Analysis and Machine Intelligence, IEEE Transactions on 19(7), 711 –720 (jul 1997)
3. Belkin, M., Niyogi, P., Sindhwani, V.: Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. Journal of Machine Learning Research 7, 2399–2434 (2006)
4. Bengio, S., Pereira, F., Singer, Y., Strelow, D.: Group sparse coding. In: Bengio, Y., Schuurmans, D., Lafferty, J., Williams, C.K.I., Culotta, A. (eds.) Advances in Neural Information Processing Systems 22, pp. 82–89 (2009)
5. Bowyer, K., Chang, K., Flynn, P.: A survey of approaches to three-dimensional face recognition. In: Pattern Recognition. vol. 1, pp. 358 – 361 Vol.1 (23-26 2004)
6. Cootes, T., Edwards, G., Taylor, C.: Active appearance models. Pattern Analysis and Machine Intelligence, IEEE Transactions on 23(6), 681 –685 (jun 2001)
7. Gao, S., Tsang, I.W., Chia, L.T., Zhao, P.: Local features are not lonely - laplacian sparse coding for image classification. In: CVPR (2010)
8. He, X., Yan, S., Hu, Y., Niyogi, P.: Face recognition using laplacianfaces. IEEE Trans. Pattern Anal. Mach. Intell. 27(3), 328–340 (2005)
9. Hoi, C.H., Lyu, M.R.: Robust face recognition using minimax probability machine. In: Proc. IEEE Conf. on Multimedia and Expo (ICME'04). pp. 1175–1178 (2004)
10. Lee, H., Battle, A., Raina, R., Ng, A.Y.: Efficient sparse coding algorithms. In: In NIPS. pp. 801–808. NIPS (2007)
11. Nesterov, Y.: Smooth minimization of non-smooth functions. Mathematical Programming 103(1), 127–152 (2005)
12. Prince, S., Warrell, J., Elder, J., Felisberti, F.: Tied factor analysis for face recognition across large pose differences. Pattern Analysis and Machine Intelligence, IEEE Transactions on 30(6), 970 –984 (june 2008)
13. Tan, X., Chen, S., Zhou, Z.H., Zhang, F.: Face recognition from a single image per person: A survey. Pattern Recognition 39(9), 1725–1745 (2006)
14. Turk, M.A., Pentland, A.P.: Face recognition using eigenfaces. In: CVPR. pp. 586–591 (1991)
15. Wright, J., Yang, A.Y., Ganesh, A., Sastry, S.S., Ma, Y.: Robust face recognition via sparse representation. IEEE Trans. Pattern Anal. Mach. Intell. 31(2) (2009)
16. Yang, Y.L., Kim, J., Luo, F., Hu, S.M., Gu, X.: Optimal surface parameterization using inverse curvature map. IEEE Transactions on Visualization and Computer Graphics 14(5), 1054–1066 (2008)
17. Zhang, X., Gao, Y.: Face recognition across pose: A review. Pattern Recognition 42(11), 2876–2896 (2009)
18. Zhao, W., Chellappa, R., Phillips, P.J., Rosenfeld, A.: Face recognition: A literature survey. ACM Comput. Surv. 35(4), 399–458 (2003)