Automated Scene-Specific Selection of Feature Detectors for 3D Face Reconstruction

Yi Yao, Sreenivas Sukumar, Besma Abidi, David Page, Andreas Koschan, and Mongi Abidi

Imaging, Robotics, and Intelligent System Lab
The University of Tennessee, Knoxville, TN, 37996

Abstract. In comparison with 2D face images, 3D face models have the advantage of being illumination and pose invariant, which provides improved capability of handling changing environments in practical surveillance. Feature detection, as the initial process of reconstructing 3D face models from 2D uncalibrated image sequences, plays an important role and directly affects the accuracy and robustness of the resulting reconstruction. In this paper, we propose an automated scene-specific selection algorithm that adaptively chooses an optimal feature detector according to the input image sequence for the purpose of 3D face reconstruction. We compare the performance of various feature detectors in terms of accuracy and robustness of the sparse and dense reconstructions. Our experimental results demonstrate the effectiveness of the proposed selection method from the observation that the chosen feature detector produces 3D reconstructed face models with superior accuracy and robustness to image noise.

1 Introduction

The 3D reconstruction from uncalibrated video sequences has attracted increasing attention recently. Most of the proposed algorithms regarding feature matching and projective/metric reconstruction have applications in 3D reconstruction of man-made scenes [1, 2]. Recently, because of the difficulties in 2D face recognition caused by illumination and pose variations, recognition algorithms using 3D face models have emerged [3], which calls for reconstruction algorithms designed particularly for faces. Hu *et al.* used salient facial feature points to project a 2D frontal view image onto a 3D face model automatically [4] and illustrated improved face recognition rates using the 3D model despite pose and illumination variations. Chowdhury *et al.* reconstructed 3D facial feature points and obtained a 3D face model by fitting these points to a generic 3D face model [5].

Most existing 3D reconstruction algorithms start with feature selection and matching [1, 2, 5]. Based on the matched features in consecutive frames, 3D projective and metric structures are recovered. Therefore, the accuracy and robustness of feature detection and matching directly affect the overall performance of the reconstruction. Popular features for 3D reconstruction are image corners and lines. For a man-made scene, there exist well-defined corners, which facilitate the use of fast and

G. Bebis et al. (Eds.): ISVC 2007, Part I, LNCS 4841, pp. 476–487, 2007.

straightforward feature detectors such as Harris corners. However, for face images, corners are not as distinguishable as in man-made scenes. In addition, face images include smooth areas, for example cheek and forehead, where feature matching becomes more ambiguous. Therefore, it is important to find an appropriate feature detector, which can make full use of facial features and avoid smooth areas simultaneously for 3D face reconstruction.

In this paper, we propose a data driven feature detector selection algorithm, where the optimal detector is dynamically selected according to different scene structures using a cost function based on information complexity (ICOMP) [6]. Our selection framework, referred to as *MuFeSaC* [7], improves existing algorithms by using an adaptive strategy for automatic extraction of relevant features from face images that contribute to facial structure and thus lead to improved accuracy and robustness of the resulting 3D reconstruction. An example face image, detected corners, and reconstructed 3D face model are shown in Fig. 1.

(a) (b) (c)

Fig. 1. Reconstructing 3D models of human faces from uncalibrated video sequences. (a) Input face image. (b) Corners relevant for 3D reconstruction. (c) Reconstructed 3D face model.

A performance comparison of various corner detectors using repeatability and information content can be found in [8]. The comparison was conducted based on various scene structures, including man-made and natural, and intended to provide general conclusions independent of input scene structures. In comparison, our algorithm is a data driven method which dynamically chooses the optimal feature detector based on the input sequences. Such a method is particularly useful for face reconstruction from surveillance videos where faces are tracked in various backgrounds.

We apply our selection scheme to face sequences, evaluate and compare the performance of the chosen detector against various types of feature detectors in terms of accuracy and robustness of the 3D face reconstruction, and prove that the chosen feature detector produces the best performance.

The major contributions of this paper are: (1) defining a data-driven selection framework that automatically chooses an appropriate feature detector for a face and eventually produces a better 3D reconstruction and (2) investigating the importance of feature detection methods in 3D reconstruction of faces by comparing results of several widely used corner detectors. Section 2 describes the 3D face reconstruction pipeline. Our automatic feature detector selection algorithm is discussed in Section 3. Experimental results are presented in Section 4 before drawing conclusions in Section 5.

2 3D Face Reconstruction

Since 3D face reconstruction is the target application of our automated feature selection algorithm, the overall process of constructing 3D face models is presented in this section. A schematic illustration of 3D reconstruction from a 2D uncalibrated image sequence with automated feature selection mechanism is shown in Fig. 2, which includes feature detection and matching, projective/metric reconstruction, 3D deformation, and texture mapping.

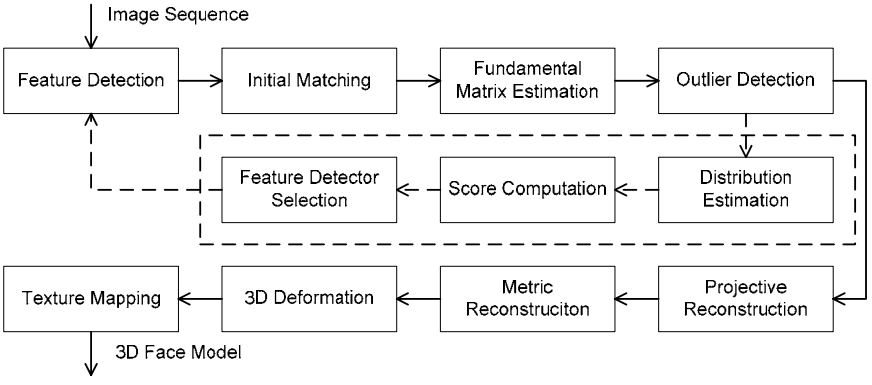

Fig. 2. Schematic illustration of 3D face reconstruction with automated selection of feature detectors

A two step feature matching process is employed. The initial matching is performed using a correlation based method. The RANSAC method is used for outlier detection and robust matching based on estimated fundamental matrices. Our automated feature selection method will be discussed in detail in Section 3. The selection of feature detectors can be performed offline for applications with consistent scene structures. Once the optimal feature detector is chosen, it can be applied directly for 3D reconstruction. However, for dynamic scenes where the characteristics of the scene structures are time varying, online selection is necessary.

We use the factorization approach for projective reconstruction, as demonstrated in [9]. Let the j^{th} point in the i^{th} frame, $\mathbf{x}_{ij}$, be projected from the scene point X_j by $\lambda_{ij}\mathbf{x}_{ij} = P_i X_j$, where λ_{ij} and P_i denote the projective depths and projection matrices, respectively. Given N_p matched points in N_f frames we have:

$$\begin{bmatrix} \lambda_{11}\mathbf{x}_{11} & \lambda_{12}\mathbf{x}_{12} & \cdots & \lambda_{1N_p}\mathbf{x}_{1N_p} \\ \lambda_{21}\mathbf{x}_{21} & \lambda_{22}\mathbf{x}_{22} & \cdots & \lambda_{2N_p}\mathbf{x}_{2N_p} \\ \vdots & \vdots & \vdots & \vdots \\ \lambda_{N_f 1}\mathbf{x}_{N_f 1} & \lambda_{N_f 2}\mathbf{x}_{N_f 2} & \cdots & \lambda_{N_f N_p}\mathbf{x}_{N_f N_p} \end{bmatrix} = \begin{bmatrix} P_1 \\ P_2 \\ \vdots \\ P_{N_f} \end{bmatrix} \begin{bmatrix} X_1 & X_2 & \cdots & X_{N_p} \end{bmatrix}, \tag{1}$$

where the matrix on the left hand side is the measurement matrix. We start with an initial estimate of the projective depths. The initial depths can be set to ones or obtained using Sturm and Triggs' method [10]. The depths are then normalized so that $\sum_{i=1}^{N_f} \lambda_{ij}^2 \mathbf{x}_{ij}^T \mathbf{x}_{ij} = 1$ and $\sum_{j=1}^{N_p} \lambda_{ij}^2 \mathbf{x}_{ij}^T \mathbf{x}_{ij} = 1$. We find the nearest rank-4 approximation of the measurement matrix using SVD, based on which 3D reconstructed points are derived. These reconstructed points are reprojected into each image to obtain new estimates of the depths. The process is repeated until the variations in the projective depths are negligible.

We then find a transformation matrix H and upgrade the projective structure by HX_j. Using the dual absolute quadric Ω^*, we have $\omega_i^* \sim P_i \Omega_1^* P_i^T$ where $\omega_i^* = K_i K_i^T$ with K_i as the camera's intrinsic matrices [11]. A linear solution of Ω^* can be obtained by imposing additional constraints on the camera's intrinsic parameters, such as zero skew, unit aspect ratio, and zero principal point, and the rank-3 property is applied for improved accuracy. The transformation matrix is then obtained by forcing $H\Omega_1^* H^T = diag(1, 1, 1, 0)$ and projective reconstruction is elevated to metric reconstruction by $P_{E,i} = P_i H^{-1}$ and $X_{E,j} = HX_j$. Finally bundle adjustment is carried out to minimize the projection errors: $\min \sum_{i,j} \left\| \mathbf{x}_{ij} - P_{E,i} X_{E,j} \right\|^2$.

The output of the previous module is a cloud of points but not a smooth surface representing a human face. By assuming that sufficient face information is embodied in the sparse point cloud, we deform a generic 3D face model. The vertices in the generic mesh are deformed using energy minimization principles similar to [5]. The procedure matches features in the generic mesh model by rotating *R*, translating **t** and scaling *s* 3D points and aligning the two models using an energy minimization function $E(s, R, \mathbf{t}) = \sum_j \| X_{E,j} - X_{M,j} \|^2$, where $X_{E,j}$ is the reconstructed point and $X_{M,j}$ refers to the point in the generic mesh. Once the initial alignment is obtained, the points in the generic mesh model are refined by weighing the distance along the surface normal of the mesh and the nearest reconstructed points and then deformed to preserve the features of the face reconstructed from the image sequence.

Once a 3D face mesh model is obtained, texture mapping is carried out by constructing the texture on a virtual cylinder enclosing the face model. A color is associated with each vertex of the deformed model by computing blending weights of each rectified image used for 3D reconstruction. The blending weights are based on the angle between the reconstructed camera's direction and the surface normal of the deformed mesh model.

3 Feature Selection

With the processing pipeline described in Section 2, our initial experiments indicated that sparse reconstruction had a major influence on the accuracy of dense reconstruction, and that the choice of the feature detection method was critical in converting the

image sequence into a meaningful and accurate sparse 3D point cloud. This observation motivated our study of different feature detectors leading to the definition of *MuFeSaC* [7], short for multiple feature sample consensus, as an extension of RANSAC used in the sparse reconstruction pipeline to include multiple feature detectors.

The contribution of *MuFeSaC* over RANSAC is an inference engine that, in addition to finding the parameters of the interest model fit based on noisy data, also evaluates the confidence in the parameter estimates. *MuFeSaC* operates towards computing confidence scores from RANSAC iterations at the same time combining the confidence from one single feature detector with the information from other competing interest points, thereby reducing the risk due to the choice of the feature detector. We list the different stages of the *MuFeSaC* procedure in Table 1 and explain the model selection criteria called information complexity that acts as the consensus scoring tool. The implementation details of single feature outlier consensus and competing feature consensus are discussed in Section 3.1 and 3.2, respectively.

Table 1. Pseudo code of *MuFeSaC*

1. ***For*** *each feature detector* FD_i **,** $i = 1,2,3\ldots N$
   a) *Extract interest points from two successive frames.*
   b) *Find the putative matches using proximity and cross correlation.*
   c) *Perform RANSAC and iterate to a convergence. Collect d-estimated parameters S of model M fitted during the iterations of RANSAC.*
   d) *Estimate probability distribution* B_i *based on n (n > 30) iterations of parameter estimates* $(S_1 \ldots S_n)$ *collected.*

   End
2. *Score Single Feature Outlier Consensus* $(SFOC_i)$ *using the model selection criterion.*
3. *Compute Competing Feature Consensus Score* $(CFCS_i)$ *by evaluating competing distributions* B_i *for different hypothesis.*
4. *Choose the optimal feature detector with minimum* $SFOC_i + CFCS_i$.

Repeat steps 1-4 every k frames. (Typically k=10 for face videos)

3.1 Single Feature Outlier Consensus

If we were to choose the optimal feature detector based on the RANSAC convergence consensus alone, we would ideally want to pick the method that is indicative of maximum likelihood of the parameters of the model fitted by RANSAC with minimum uncertainty, or in simpler words B_i with minimal variance. This can be mathematically expressed as the minimizer of criterion (2) that simultaneously considers the likelihood and also penalizes the uncertainty associated with the likelihood of the parameters of model *M*. This model selection criterion in the statistics literature [6] is known as ICOMP and derives from the Kullback-Liebler (KL) distance between estimated and unknown underlying probability densities. Without much modification, we are able to apply this criterion in evaluating the confidence in the model fit during the iterations of RANSAC with each feature. We note that Eq. (2) does not involve distributional assumptions and can be applied to even Parzen window estimates of B_i.

$$\begin{aligned} SFOC_i = & \text{ Lack of fit + Profusion of uncertainty} \\ = & \text{ } -2 \log (\textit{Likelihood of } \mu_i) + 2\, C_1 (F^{-1}(\Sigma_i)), \end{aligned} \tag{2}$$

where F^{-1} is the inverse Fisher information matrix, μ_i and Σ_i are the maximum likelihood estimates of the mean and covariance computed as the first two moments of B_i. The C_1 measure and the F^{-1} are computed using Eq. (3) and (4):

$$C_1(F^{-1}(\Sigma_i)) = \frac{s}{2}\log\left[\frac{tr(F^{-1}(\Sigma_i))}{s}\right] - \frac{1}{2}\log\left|F^{-1}(\Sigma_i)\right|, \tag{3}$$

where s denotes the rank of F^{-1}, $|\bullet|$ refers to the determinant, tr refers to the trace of the matrix, and

$$F^{-1}(\Sigma_i) = \begin{bmatrix} \Sigma_i & 0 \\ 0 & D^{+}{}_{p}(\Sigma_i \otimes \Sigma_i) D^{+}{}_{p}' \end{bmatrix}, \tag{4}$$

with D^{+}_{p} being the Moore-Penrose inverse of vectorized Σ_i and $\otimes$ representing the Kronecker product. The C_1 measure for penalizing uncertainty is obtained by maximizing mutual information in d-dimensions. We direct the reader to [6] for more implementation details on the finite sampling form of Eq. (2).

3.2 Competing Feature Consensus

The $CFCS_i$ quantifies the agreement between the competing models M fit by RANSAC from each feature detector. The score is obtained by first evaluating different hypothesis listed below and then choosing the optimal consensus combinatorial cluster among competing feature detectors:

Case 1: All B_i's maximizing the likelihood of the same parameters for model M. All μ_i*'s* equal and Σ_i*'s* equal.
Case 2: All μ_i*'s* equal but Σ_i*'s* are not.
Case 3: All μ_i*'s* and Σ_i*'s* unequal, but there exists a maximal cluster of μ_i*'s* equal.

The verification of these hypotheses is like performing multi-sample clustering based on information distances in an entropic sense as described by Bozdogan in [6]. We follow a similar approach to verify these three cases, by considering the samples that contribute to distributions B_i to have to come from the same distribution and evaluate the complexity in model-fitting as the criterion to decide which of the three cases has occurred. We use the Akaike information criterion (AIC) to score the different hypothesis based on the likelihood of feature cluster L and parameter parsimony estimation m.

$$\begin{aligned} AIC(\mu_i, \Sigma_i, \kappa) &= -(\text{Likelihood of feature cluster}) + \\ &\quad \text{Parameter parsimony after clustering} \\ &= -2\log L + 2m. \end{aligned} \tag{5}$$

The evaluation of the likelihood of feature cluster L only considers the samples that contributed to the distributions B_i's within the cluster evaluated for consensus. We evaluate the parameter parsimony factor for the 2^N different cluster combinations based on the formulae listed in Table 2. The hypothesis that has minimum AIC is the statistical decision. Initially, we only evaluate the three case hypotheses. This initial

3-case hypothesis verification can avoid the combinatorial evaluations when all methods are accurate. We assign the minimizer of the AIC for the 3-case hypothesis as $CFCS_i$ to the corresponding feature detectors. If the minimizer indicates the occurrence of Case 2 or 3, we perform the evaluation on all combinatorial "feature detector" clusters shown under Case 3 in Table 2. The minimizer of the AIC score on these sub-clusters points to the cluster with maximal κ feature detectors contributing to the same model parameters. This AIC score is assigned only to the "feature detectors" within the maximal cluster. This cluster evaluation procedure eliminates the possibility that we select a feature detector that has minimal outliers but is giving us totally different parameters after RANSAC convergence.

We just add the two definitions of information measures which is a common practice with log utility functions. Also, we note that our formulation with $SFOC_i$ minimizes the error in the model used for estimating the geometry from successive views. On the other hand, $CFCS_i$ takes care of the risk in the model itself by inferring from different model generators in feature detectors. Thus, with *MuFeSaC* automatically selecting the interest points from a face video sequence, we now present experimental results.

Table 2. An example of parameter parsimony estimation (m) for a simple d-parameter model M with $N = 3$

	κ	Clustering	m
Case 1	3	(F_1,F_2,F_3)	$d+ d(d+1)/2$
Case 2	1	$(F_1)(F_2)(F_3)$	$Nd +d(d+1)/2$
Case 3	2	$(F_1,F_2)(F_3)$ $(F_1,F_3)(F_2)$ $(F_2,F_3)(F_1)$	$\kappa d + \kappa d(d+1)/2$

4 Experimental Results

Eight face images of each subject, collected from different viewpoints, are used for 3D face reconstruction. Five types of corner detectors are implemented: curvature corner [12], Harris corner [13], STK [13], phase congruency corner (PCC) [14], and FAST [15]. These feature detectors are chosen for the different heuristics that motivates them, Harris corners being intensity gradient-based, phase congruency being spatial-frequency inspired, and curvature corners being edge-derived. We compare the 3D reconstruction using the sparse point cloud and dense 3D face model. The sparsely reconstructed point cloud is an intermediate output, the accuracy of which determines the performance of the successive dense reconstruction. The performance of the sparse reconstruction provides a direct measure of evaluating the influence of the corner detector on the final 3D face reconstruction excluding factors which may come into play from 3D deformation. Hence we compare results at both stages of reconstruction. In the interest of space, only the results from one subject are presented with similar observations applicable to other tested subjects. The curvature corner was chosen as the optimal detector for this sequence by our automated selection algorithm. Therefore, in the following discussion, its performance is compared against the performances of other tested corner detectors.

4.1 Sparse Reconstruction

Fig. 3 illustrates our experimental results and compares the performance of the chosen feature detector, curvature corner, against other corner detectors. Sample images with detected corners show that corner detectors behave in different ways for face images which include both clustered corners (eyes and mouth) and smooth areas (cheek and forehead). Curvature, FAST, and PCC detectors focus on image edges while the corners detected by Harris and STK also appear in smooth areas.

The reconstructed 3D point clouds are shown from the best viewpoint in terms of illustrating the facial structure. From visual inspection of these 3D plots, we see that the reconstructed structure using curvature corners, the chosen feature detector, yields the best visual representation. From Harris and FAST corner detectors, visible facial structures are obtained with an increased number of noisy points. As for the STK corner detector, the noise level further increases resulting in a structure barely distinguishable. In comparison, the reconstructed structure from the PCC detector has most of the points with good approximation but a few points with large errors.

We use the variations in the projective depth to describe the accuracy and stability of the reconstruction. From actual measurement, the variation in depth in the world coordinates should be within 30cm, which corresponds to 0.375 when normalized to the camera coordinates. The back projected depths in the camera coordinates are plotted for all detected corners in all frames in Fig. 3. We use the percentage of the 3D points with a projective depth exceeding the theoretical range to quantitatively evaluate and compare the performance of various feature detectors, as shown in Fig. 4. All of the back projected depths from the curvature corner detector fall within the theoretical range. The variations from other detectors exceed the theoretical range with the STK detector yielding the most variations. The frequency and magnitude of variations exceeding the theoretical range are from noisy reconstruction caused by erroneous corner matching. Based on this quantitative measure, the chosen feature detector, curvature corner, has the best performance. From both visual inspection and quantitative measure, the chosen detector produces properly spaced corners around major facial features and hence generates the most accurate and robust sparse reconstruction.

4.2 Dense Reconstruction

The final 3D dense model is obtained by deforming a generic face model using the point cloud recovered from sparse reconstruction. The reconstructed face model is compared with a 3D reference model scanned using a GenexFace Cam. The original scan is processed in RapidForm to remove holes and spikes. The deviation from the reference model describes the accuracy of the reconstruction and hence the efficiency of the corresponding corner detector. Fig. 5 illustrates the recovered 3D dense face models and their deviation from the reference model. The result from the FAST detector is not included because the sparse reconstruction does not have sufficien depths recovered on identifiable facial features. From Fig. 5, we observe that the dense reconstruction based on the chosen detector, curvature corner, yields the smallest maximum deviation (5.32) as compared with Harris (5.50), STK (5.93), and PCC

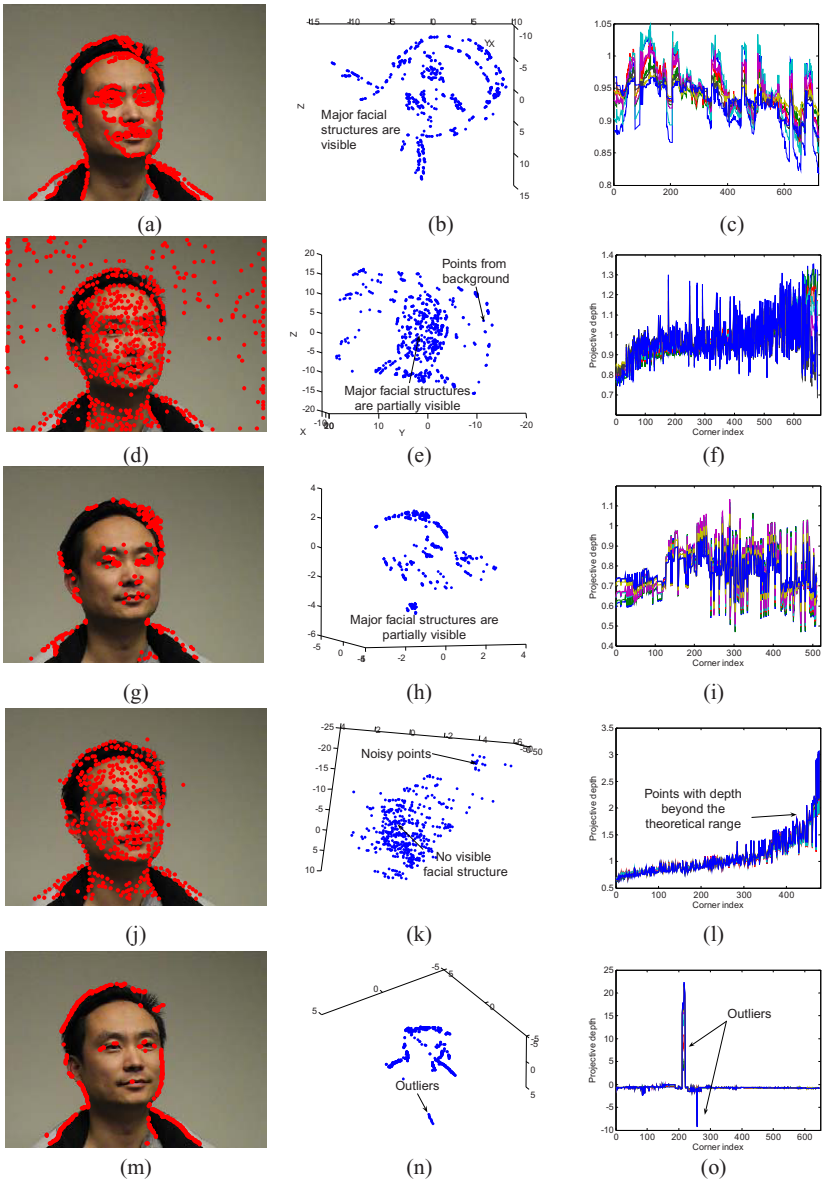

Fig. 3. Performance comparison of sparse reconstruction. Sample face images with detected corners: (a) curvature corner (the chosen feature detector by our automated selection algorithm for the illustrated sequence), (d) Harris, (g) FAST, (j) STK, and (m) PCC. 3D plots of sparsely reconstructed point clouds shown in the best view for illustrating facial structures: (b) curvature corner, (e) Harris, (h) FAST, (k) STK, and (n) PCC. Plots of projective depths: (c) curvature corner, (f) Harris, (i) FAST, (l) STK, and (o) PCC. Colored curves illustrate the projective depths from different frames. The 3D sparse reconstruction using the chosen feature detector, curvature corner, preserves major facial structures and produces accurate and robust reconstruction with all projective depths in the theoretical range.

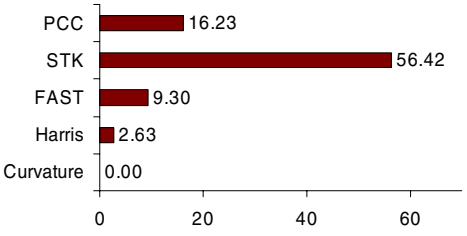

Fig. 4. Quantitative performance comparison of sparsely reconstructed 3D point clouds using various corner detectors: the percentage of reconstructed points with a projective depth exceeding the theoretical range. A smaller number suggests better accuracy. The point cloud from the chosen feature detector, curvature corner, has the most accurate variations in depth estimation.

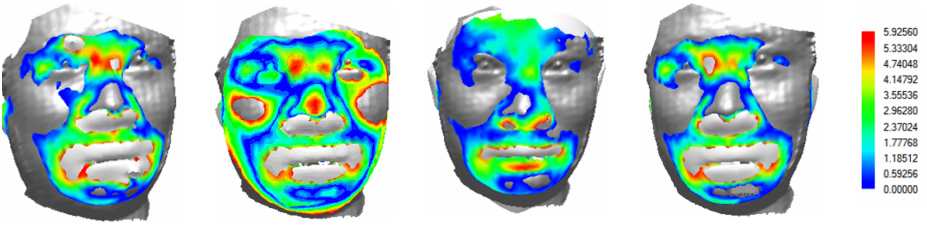

(a) (b) (c) (d) (e)

Fig. 5. Performance comparison of 3D dense reconstruction: (a) curvature corner, (b) Harris, (c) STK, and (d) PCC. The reconstructed 3D face model is compared with a reference 3D scan of the face. The reference 3D face model is shown in grey. The deviations from the reference model are shown in color code (e). The reconstruction based on the chosen detector, curvature corner, produces the smallest maximum and average deviations.

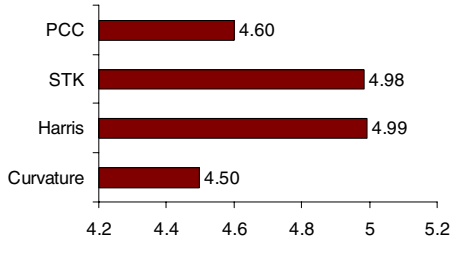

Fig. 6. Quantitative performance comparison (Hausdorff distance) of the recovered 3D face model using various corner detectors. The reconstruction using the chosen detector, curvature corner, yields the smallest deviation from the reference model.

(5.36). We also list the Hausdorff distance as a quantitative measure of the accuracy of dense reconstruction in Fig. 6. We have used the Hausdorff distance instead of the commonly used ICP metric because the point clouds are not of comparable resolution. Moreover, Hausdorff distance provides an unbiased basis for the comparison of point clouds from different feature point detectors. As expected the 3D model using the chosen detector, curvature corner, produces the smallest Hausdorff distance and thus the best accuracy. Note that to obtain a meaningful reconstruction, the outliers from PCC are excluded before deforming the generic face model. Therefore, the final 3D reconstruction using PCC presents a comparable accuracy as curvature corners
Compared with the results from sparse reconstruction, we observe that the performance gap between the curvature and other corner detectors decreases for dense reconstruction. This is due to the use of the generic face model as a constraint in the optimization process solving for the final dense reconstruction, which imposes additional bounds on the variations in the structure and reduces the influence of the errors from 3D sparse reconstruction. However, the use of the generic model may also reduce the useful discriminant structures for differentiating different subjects. In our future work, principal component analysis (PCA) will be explored for improved accuracy in 3D deformation. With the introduction of more principal components in addition to the generic face model, which is actually an averaged face model, more useful discriminant structures can be reserved. With PCA, the accuracy of sparse reconstruction plays more important role and hence more substantial performance difference can be observed by using different feature detectors.

5 Conclusions

In this paper, we proposed an automatic feature selection algorithm that adaptively chooses the optimal feature detector according to the input data for the purpose of 3D face reconstruction from uncalibrated image sequences. Several widely used corner detectors were implemented and served as competitive candidates for our selection algorithm. Experiments based on various subjects were performed to qualitatively and quantitatively compare the accuracy and robustness of the sparse and dense 3D face reconstructions. The chosen detector by our method produced a 3D face reconstruction with the best accuracy and robustness, which proves the effectiveness of the proposed selection algorithm.

References

1. Pollefeys, M., Gool, L., van Vergauwen, M., Verbiest, F., Cornelis, K., Tops, J.: Visual modeling with a hand-held camera. Int'l J. of Computer Vision 59, 207–232 (2004)
2. Kien, D.T.: A review of 3D reconstruction from video sequences. MediaMill3D Technical Report, University of Amsterdam (2005)
3. Lu, X., Jain, A.K., Colbry, D.: Matching 2.5D face scans to 3D models. IEEE Trans. on Pattern Analysis and Machine Intelligence 28, 31–43 (2006)
4. Hu, Y., Jiang, D., Yan, S., Zhang, L., Zhang, H.: Automatic 3D reconstruction for face recognition. In: IEEE Int'l Conf. on Automatic Face and Gesture Recognition, Washington D.C, pp. 843–848 (2004)

5. Chowdhury, A.R., Chellappa, R., Krishnamurthy, S., Vo, T.: 3D face reconstruction from video using a generic model. In: IEEE Int'l Conf. on Multimedia and Expo, Lausanne, Switzerland, pp. 449–452 (2002)
6. Bozdogan, H.: Akaike's information criterion and recent developments in information complexity. J. of Mathematical Psychology 44, 62–69 (2000)
7. Sukumar, S.R., Bozdogan, H., Page, D., Koschan, A., Abidi, A.: MuFeSaC: Learning when to use which feature detector. In: Accepted to the IEEE International Conference on Image Processing, San Antonio, TX (2007)
8. Schmid, C., Mohr, R., Bauckhage, C.: Evaluation of interest point detectors. Int'l J. of Computer Vision 37, 151–172 (2000)
9. Mahamud, S., Hebert, M., Omori, Y., Ponce, J.: Provably-convergent iterative methods for projective structure from motion. In: IEEE Conf. on Computer Vision and Pattern Recognition, Kauai, Hawaii, pp. 1018–1025 (2001)
10. Sturm, P., Triggs, B.: A factorization based algorithm for multi-image projective structure and motion. In: European Conf. on Computer Vision, Cambridge, England, pp. 709–720 (1996)
11. Pollefeys, M., Koch, R., van Gool, L.: Self-calibration and metric reconstruction in spite of varying and unknown intrinsic camera parameters. In: IEEE Int'l Conf. on Computer Vision, Bombay, India, pp. 90–95 (1998)
12. He, X.C., Yung, N.H.C.: Curvature scale space corner detector with adaptive threshold and dynamic region of support. In: Int'l Conf. on Pattern Recognition, Cambridge, England, pp. 791–794 (2004)
13. Mikolajczyk, K., Schmid, C.: Scale and affine invariant interest point detectors. Int'l J. of Computer Vision 1, 63–68 (2004)
14. Kovesi, P.: Phase congruency detects corners and edges. In: Australian Pattern Recognition Society Conf., Sydney, Australian, pp. 309–318 (2003)
15. Rosten, E., Drummond, T.: Machine learning for high speed corner detection. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) ECCV 2006. LNCS, vol. 3951, pp. 430–443. Springer, Heidelberg (2006)