RECONSTRUCTING 3D FACES IN CULTURAL HERITAGE APPLICATIONS

Andreas Lanitis[a] and Georgios Stylianou[b]

[a] Dept. of Multimedia and Graphic Arts, Cyprus University of Technology , P.O. Box 50329, 3036, Lemesos, Cyprus,
andreas.lanitis@cut.ac.cy
[b] Dept. of Computer Science, European University Cyprus , P.O. Box 22006, 1516, Nicosia, Cyprus,
g.stylianou@euc.ac.cy

KEY WORDS: Image-Based 3D Face Reconstruction, 3D Face Shape Models

ABSTRACT:

Image-based 3D face reconstruction techniques can be used for generating a 3D model of a person's face given a single or multiple images of the face of a person. Such techniques could be used for generating 3D models of faces appearing in cultural heritance artefacts, based on images captured with ordinary digital cameras. The ability to generate 3D reconstructions of faces appearing in cultural heritage artefacts will facilitate the development of numerous applications such as the generation of cultural heritage related 3D animations, short films, educational applications and customized educational games. In this paper we review key image-based 3D face reconstruction algorithms belonging to different categories and assess their applicability for reconstructing 3D faces appearing in cultural heritage artefacts. We also describe a dedicated 3D face reconstruction algorithm that can be used for efficient reconstruction of faces in cultural heritage applications based on a single image. The proposed algorithm requires the location of a number of landmarks on the face to be reconstructed. A deformable 3D shape model is used for generating an approximation of the 3D face geometry of the face shown on the artefact and texture mapping is used for generating the full 3D model. In the paper we present initial quantitative results and visual results obtained when the method was used for reconstructing 3D faces appearing in different types of cultural heritage artefacts such as statues, paintings, murals and mosaics.

1. INTRODUCTION

Ancient sites and artefacts such as wall paintings, paintings, statues, pots, anaglyphs, mosaics and icons contain a plethora of human faces that usually display important persons in history (see figure 1).

Figure 1: Faces appearing in different types of CH artefacts

In this paper we review methods that can be used for generating three-dimensional (3D) reconstructions of faces appearing in such artefacts. The ability to generate a complete 3D models of faces appearing in cultural heritance artefacts can be very important in numerous applications such as:

- 3D animations: Once 3D facial models are created, it is trivial to use animation techniques for generating animations that can be used in short films and other educational applications.
- Educational Games. 3D reconstructed faces can be used for generating realistic characters appearing in CH related educational games.
- Art related studies: The generation of the actual 3D face geometry of subjects appearing in CH artefacts will allow the comparison of faces of the same subject generated by different artists, so that comparison of artistic methodologies and trends will be enabled. The definition of such trends will be most important for studying the design methodologies and styles adopted in different countries and different time periods.
- Digital restoration: 3D reconstructions of damaged faces shown in artefacts can be used as the basis for digital restoration of the facial appearance of missing/damaged parts.

Although it is possible to generate a 3D model of a face using dedicated 3D scanners, this method is only applicable in applications involving the digitisation of existing 3D facial structures i.e. faces of statues or faces in anaglyphs. 3D scanners cannot be used for generating 3D models of faces appearing as 2D surfaces such as faces appearing in paintings, mosaics and murals. Even when dealing with the generation of 3D models of statues, the use of 3D scanners possesses several limitations due to the size and cost of 3D scanners, the requirement of controlled lighting during the scanning process, the considerable amount of post processing required for removing noise and joining 3D scans, and the overall processing time required during the scanning process. Such problems are minimized when on-site data acquisition is performed using 2D instead of 3D imaging equipment.

As an alternative to using 3D scanners, image-based 3D reconstruction techniques can be used for generating 3D face models of faces appearing in CH artefacts. Over the past decade the technology of image-based 3D face reconstruction has evolved dramatically. Realistically looking 3D faces can be reconstructed from a single or more photographs including video frames (Stylianou, 2008). However, most 3D face reconstruction techniques reported in the literature aim towards the reconstruction of real faces appearing in digital face images. To the best of our knowledge as of today no systematic efforts in reconstructing 3D faces appearing in CH artefacts were reported in the relevant literature.

In this paper we provide a general overview of 3D face reconstruction methods reported in the literature. In particular we describe typical example-based (Blanz, 1999), stereo-based (Akimoto, 1993; Chen 2001) and video-based (Leung, 2000) 3D reconstruction methods. For each category of 3D reconstruction algorithms we provide a discussion that assesses the applicability of typical algorithms for reconstructing faces appearing in CH artefacts. We also propose a dedicated method for 3D face reconstruction of CH artefacts and preliminary results when the method is applied for reconstructing faces appearing in different types of artefacts.

2. 3D FACE RECONSTRUCTION METHODOLOGIES

Image-based 3D face reconstruction methods reported in the literature can be classified as example-based, stereo and video-based methods. Usually example-based methods operate on a single input image, stereo methods operate on a pair of images and video-based methods operate on multiple input frames. In this section we provide a brief overview of typical 3D face reconstruction approaches for each category, outline the major advantages and disadvantages of each category and discuss the applicability of those approaches for reconstructing faces appearing in CH artefacts. The review provided in this section aims to present the basic principles pertaining to 3D reconstruction methods - a more comprehensive review appears elsewhere (Stylianou, 2008).

2.1 Example based Methods

Example-based methods rely on the analysis of a large number of training samples for accomplishing 3D reconstruction so that the structure and typical deformations associated with 3D faces are learned. The information obtained during the training process is used to guide the reconstruction process.

An example-based 3D face reconstruction method has two phases: The morphable model generation phase and the actual reconstruction phase. During the morphable model generation phase the training set, which is composed of a face database containing 3D face models, is used for building a parametric face model, capable of generating 3D faces consistent with samples appearing in the training set (Atick, 1996; Blanz, 1999). Once a parametric face model is trained, it is possible to generate different 3D faces by setting values to the model parameters. 3D face reconstruction is accomplished by defining optimum values for the face model parameters so that a 3D face generated by the model resembles the face appearing in a given image. The actual steps involved in this process are:

1. The values of the model parameters are initialized to zero so that the model generates the mean 3D face instance.
2. Based on the values of the model parameters the corresponding 3D face instance is generated.
3. The 3D face instance is projected on the face in a given image.
4. The similarity between the given face and the 3D model instance is estimated. Usually the similarity measure is defined based on differences in texture between the 3D model instance and the face to be reconstructed.
5. An optimization algorithm is used for defining improved values for the model parameters.
6. Steps 2 to 5 are repeated until the similarity between the 3D face instance and the given face is maximized.

When the similarity measure is maximised a 3D face with appearance similar to the face in the given image is created so that the task of reconstructing a 3D face from a single face image is accomplished.

The most cited example-based 3D reconstruction methodology is the method developed by Vetter and Blanz (Blanz, 1999) where a PCA based model is used. In this case 3D reconstruction is achieved by computing the weights of the shape and texture eigenvectors for a novel 2D face image. The method requires manual initialization to roughly align the 3D morphable model and the novel face image. During the reconstruction process 22 rendering parameters such as pose, focal length of the camera, light intensity, color and direction are estimated. Other typical example-based methods are reported by Hu et al (Hu, 2004), Lee et al (Lee, 2005) and Romdhani (Romdhani, 2005).

The main advantage of example-based 3D reconstruction methods is the ability to produce visually good 3D face reconstructions based on a single face image. The use of a face model allows the implicit use of constraints related to the expected structure of a human face, so that it is possible to generate full 3D reconstructions based on missing data (i.e. data seen only from a single face view). However, example-based techniques exhibit the following disadvantages:

- Limited anatomical accuracy of generated models. This is reasonable as information from a single image is not enough for obtaining a highly accurate estimate of the 3D face structure.
- Due to the high computation requirements involved, example-based techniques usually do not run in real time.

In the majority of example-based methods reported in the literature, the calculation of the similarity between a 2D projected face model instance and the face to be reconstructed mainly relies on texture differences between corresponding pixels. For this process to be successful it is assumed that faces in the given image have skin-like texture similar to facial texture encountered in digital images. For this reason this method is not applicable to reconstructing faces appearing in CH artefacts since in such cases irregular facial textures are encountered according to the type and material used for designing the artefacts (see figure 1).

2.2 Stereo Methods

Stereo 3D reconstruction methods can be classified to two main categories: (a) methods based on orthogonal images and (b) methods based on non-orthogonal images. In category (a) it is necessary to have two images each containing a face in frontal and profile views, respectively. In the second category, it is necessary to have at least two calibrated cameras (between them) in order to capture two images of the same face from different viewpoints. Both methods identify corresponding features from the input images but use them differently as described in sections 2.2.1 and 2.2.2 respectively.

2.2.1 Orthogonal images. Reconstruction algorithms based on orthogonal images rely on the location of common features on both frontal and side photos that correspond to predefined features on a generic head model. The location of common features enables the calculation of the x and y coordinates of each vertex using the frontal image and the calculation of the z coordinate of each vertex based on the profile image. Once the coordinates of a number of vertices are defined it is possible to deform a generic head model in order to approximate the geometry defined by the extracted vertices. Once the geometry of a 3D model is defined, the face texture is generated from both images and texture mapped onto the resulting 3D model. Typical orthogonal-based 3D reconstruction methods are reported by Akimoto (Akimoto, 1993), Lee (Lee, 1997) and Park (Park, 2005).

Ortogonal-based methods can produce geometrically accurate faces since information from orthogonal images provides an adequate description of the 3D geometry of a face. Even though only two photographs are used it is trivial to extend a stereo orthogonal method to use a third photograph in order to improve the reconstruction accuracy. Furthermore, these methods are invariant to the lighting conditions even though inappropriate lighting can spoil the quality of the texture in reconstructed faces. Despite the good features of stereo methods using orthogonal images, these methods create smoothed faces because generic face models used in such cases usually have a relatively small number of vertices that does not allow the geometric description of local facial structures. 3D reconstruction based on orthogonal images can only be used in cases where images are captured in controlled conditions in order to ensure that the input pair of images is indeed orthogonal.

For face reconstruction in CH applications, this method is only applicable for reconstructing faces appearing on statues, since in those cases it is feasible to obtain two orthogonal pictures of the face using ordinary digital cameras. The fact that the texture of the input faces is not utilized for computing the geometry of the face, implies that this method can be applied to statues designed using different materials, colours and tones.

2.2.2 Non-orthogonal Images. Non-orthogonal stereo-based 3D reconstruction is a technique used to 3D reconstruct very diverse scenes. However, as this technique is quite successful, it has also been applied to the 3D face reconstruction domain (Chen, 2001; Leclercq, 2004). Non-orthogonal stereo-based techniques require as input two images of the same object captured from different angles using two calibrated digital cameras. Pixel correspondences are established between the two images to create the disparity map. The disparity map and the knowledge of the relative distance between the two cameras are used to compute the depth map, which shows the geometric structure of the face.

In general, stereo non-orthogonal techniques provide geometrically accurate results with the main drawback being the necessity of structured lighting. In addition, such techniques do not require a face database or a 3D generic face model.

Stereo-based face reconstruction methods can be used for generating 3D face models of statues based on two photographs captured from different angles. This method is also applicable for reconstructing the 3D geometry of faces appearing in CH artefacts in cases where the face of the same subject is shown as seen from different viewpoints. However, reconstructing a 3D model using two unconstrained views of the same subject, possess extra difficulties, as in such cases camera parameters and the distance between the two cameras are not known.

2.3 Video-Based Methods

Video based techniques are used for generating a 3D face model based on multiple face images captured from different viewpoints. The main steps in video-based 3D face reconstruction are:

1. Capture an image sequence showing the face to be reconstructed as seen from various viewpoints.
2. On the first frame locate a number of key features.
3. Track the location of the features in the remaining frames.
4. Use the location of feature points to reconstruct the geometry of the 3D face.
5. Collect texture information from the video frames and synthesize the texture of the reconstructed 3D face.

Researchers in this area have developed a diverse number of techniques. These techniques share similarities in the initial steps, as they all require appropriate feature point detection and tracking. The 3D face reconstruction step is done using either generic model morphing (Leung, 2000; Liu, 2001), linear combinations of basis functions that do not require a generic model (Bregler, 2000; Brand, 2001), shape from motion techniques (Chowdhury, 2002) or by using PCA based techniques (Xin, 2005) similar to example-based 3D face reconstruction.

An extension to video-based methods are Silhouette-Based methods. The objective of silhouette based methods is to produce a 3D reconstruction (or optimization) of a face by using multiple face outlines extracted from several face images or a video sequence. A complete set of silhouettes captured from different angles provides details related to the geometrical structure of a face hence it is possible to generate a 3D face from silhouettes. Figure 2 shows typical face silhouettes extracted from an image sequence. Silhouette-based methods have been used either individually or in combination with other 3D reconstruction methods such as stereo (Fua, 1999) and example-based methods (Moghaddam, 2003; Lee, 2003).

Figure 2: Examples of face silhouettes extracted from an image sequence.

In applications involving 3D face reconstruction in CH applications, video-based methods are only applicable to the reconstruction of faces appearing in statues, since in such cases it is possible to obtain multiple images of the same face as seen from different angles.

2.4 Image-Based 3D Reconstruction in CH Applications

In the previous section, we have briefly discussed image-based 3D face reconstruction methods belonging to three different categories. In this section we discuss the applicability of those methods to reconstructing images in CH applications.

Among the three categories example-based methods are the most applicable for reconstructing faces from CH artefacts since in such cases only one view of the face is usually available. However, since typical example-based methods depend on comparing the texture between the model and the face to be reconstructed, these methods are not applicable for reconstructing faces with diverse texture tones and colours such as the ones encountered in CH applications (see figure 1).

Since stereo and video-based methods require two or more views of the face to be reconstructed, such methods are applicable only in the case of statues or in cases that multiple instances of the same face are available on an artefact. Techniques from this category are not so sensitive to the type of texture, hence they can be applied for reconstructing faces with varying textures.

Ideally an image-based 3D reconstruction technique suitable for reconstructing faces appearing in CH artefacts should have the following characteristics:

- Ability to generate a 3D reconstruction from a single image, since certain types of artefacts display faces as seen from a single view.
- Ability to deal with significantly different facial geometries, since artists in different time periods adopt different painting styles, resulting in deformed facial geometries. For example faces showing Saints drawn during the Byzantine period tend to be extremely thin.
- Ability to generate 3D models with textures with different tones, shadings and colours.
- Ability to deal with damaged faces, since in some cases part of the face to be reconstructed is destroyed.

3. 3D FACE RECONSTRUCTION METHOD

In this section we describe a 3D face reconstruction method that can be used for generating 3D models of faces appearing in CH artefacts. In order to deal with most of the issues raised in section 2.4 we developed a dedicated 3D reconstruction algorithm that can be used for generating 3D models of faces shown in CH artefacts. The proposed method is a model-based approach that utilizes a PCA-based 3D shape model (Blanz, 1999) generated during the training process. The main steps involved in the reconstruction process are shown in figure 3. Details about each step of the reconstruction process are provided in the following sections.

3.1 3D Shape Model

The training set consists of 60 3D models of volunteers where each model contains about 75000 vertices. 3D models used in the training set were captured using a laser 3D scanner and pre-processed in order to eliminate noise. Typical samples from the training set are shown in figure 4.

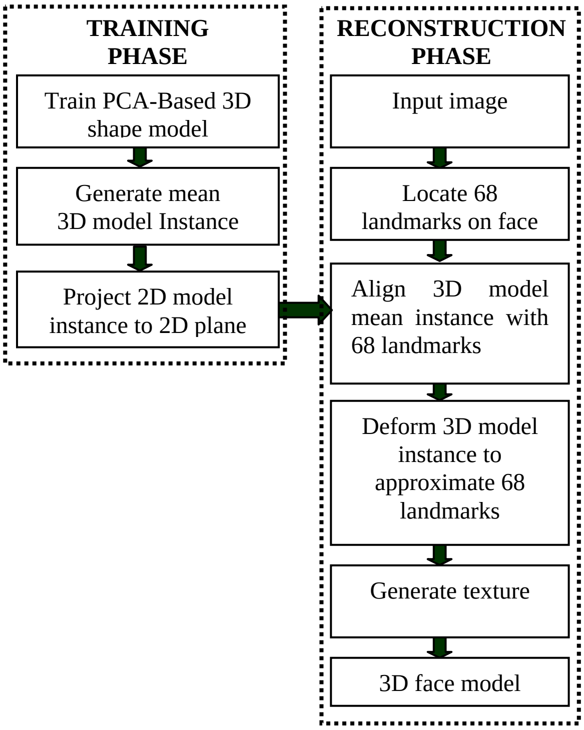

Figure 3: Block Diagram of the 3D reconstruction method

Figure 4: A typical 3D faces from the training set. Row 1: textured models, Row 2: models with triangles overlaid.

The model training procedure involves the alignment of all training samples so that there is one-to-one correspondence between vertices of all samples in the training set (Blanz, 1999), and the estimation of the mean 3D face model among the training set. A PCA-based model training procedure similar with the one used by Cootes et al (Cootes, 2001) and Blanz and Vetter (Blanz, 1999) is used for buiding a face shape model. During the process we represent training samples using the coordinates of each vertex and the covariance matrix of the deviations of vertices from the mean shape is estimated. By applying Principal Component Analysis on the covariance matrix the main modes of shape variation within the training set are defined and 3D face shapes from the training set can be represented using a small number of 3D shape model parameters. In this framework model parameters are the weights of the eigenvectors derived during the Principal Component Analysis decomposition process. Based on this framework, it is possible to generate novel 3D face shapes by setting different values to the model parameters. 3D shapes generated by the model display shape instances consistent with the training set. More details related to the process of training such models appear elsewhere (Cootes, 2001; Blanz, 1999).

3.2 3D Face Reconstruction

The 3D reconstruction algorithm involves the following steps.

3.2.1 Location of facial features. Given a face to be reconstructed we locate 68 landmarks on predefined locations as shown in figure 5.

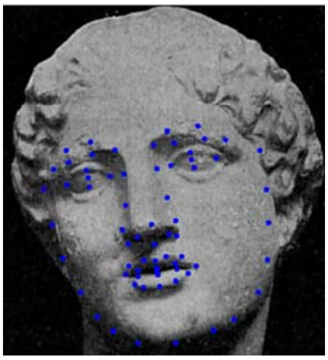

Figure 5: Locations of 68 landmarks on a face to be reconstructed.

The location of the landmarks is done based on a supervised semi-automatic method, where the users provide rough estimates of the locations of the face outline and template matching–based and edge-based algorithms are employed in order to locate the correct position for each landmark. Due to the diversity of the texture of faces to be reconstructed, a totally automatic method would lead to inaccurate landmark location. The process of locating the 68 landmarks on a face image, using the semi-automatic tool requires about five minutes to be completed.

3.2.2 3D Model Alignment. During this step we aim to align the 3D model vertices with the 68 landmarks located on the face to be reconstructed. In order to enable the alignment, we first establish the correspondences between the 68 landmarks and the 75000 3D vertices, i.e. among the 75000 vertices we select the 68 vertices that correspond to the locations of the 68 landmarks. The establishment of the correspondences is done once during the training phase using a dedicated visualization tool. Once the 68 vertices from the 3D model are selected, all further processing is carried out using only the selected 68 vertices, minimizing in that way the execution time. Only during the final step of texture mapping all 3D model vertices are used.

During the alignment phase we generate the mean 3D model instance of the 3D shape model (this is done by setting the model parameters to zero) and project the resulting 3D vertices to the 2D plane using an orthographic projection transformation. Among the 75000 vertices projected to 2D we select the 68 that correspond to the landmarks of the face to be reconstructed and calculate the overall distance (dis) between the two sets of points using:

$$dis = \sum_{i=1}^{N} \sqrt{[(\mathbf{X2D}_i - \mathbf{X3D}_i)^2 + (\mathbf{Y2D}_i - \mathbf{Y3D}_i)^2]} \qquad (1)$$

Where N is the number of selected landmarks (i.e 68),

X2D, **Y2D** are the x and y coordinates of the 68 landmarks located on the face to be reconstructed.

X3D, **Y3D** are the x and y coordinates of the selected 68 shape model vertices projected to 2D.

The alignment process involves three steps:

a. *Rigid transformations:* The 2D projected vertices are translated, scaled and rotated around the z-axis until the distance between them and the 68 2D landmarks is minimized. In the actual implementation the Generalized Procroustes Analysis alignment method (Dryden, 1998) is used. This process results in the estimation of the best rigid transformation parameters (i.e. the best x and y translation, scaling factor and z-rotation angle (roll angle)) that ensure the best alignment between the two sets of points.

b. *Definition of Yaw and Pitch angles:* 3D vertices of the shape model are rotated between angles of –60 to 60 degrees around the x and y axis. For each rotation angle the distance dis is calculated (see equation 1) and the combination of the yaw and pitch angles that minimize dis is recorded. In effect during the process the 3D pose of the face to be reconstructed is determined so that projected 3D vertices are better aligned with the 2D landmarks.

c. *Non-rigid transformations:* Up to this point the alignment process is carried out based on transformations of the mean 3D shape. However, for the reconstruction algorithm to function correctly, the mean shape must be deformed in order to approximate the shape of the face in the 2D image. The optimum shape deformation required is estimated by defining the values of the 3D shape model parameters that minimize the distance dis. In this case a minimization algorithm based on a pattern search method (Yin, 2000) is used for defining the optimum values of the model parameters.

The basic steps of the 3D model alignment process are demonstrated in figure 6.

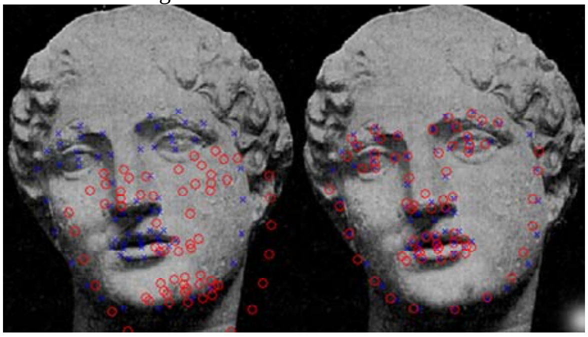

Figure 6: Demonstration of the alignment process. (Blue points are the actual 2D landmarks on the face to be reconstructed and red points are the corresponding 2D-projected vertices of the shape model instance). Left image (Before alignment) (dis=68000). Right image after rigid transformations, adjustment of yaw and pitch angles and adjustment of model parameters (dis=551).

3D Model Generation. The completion of the alignment process results in the definition of transformation and deformation parameters that can be used for aligning the 68 selected 3D vertices on the 68 landmarks located on the face to be reconstructed. The definition of the deformation parameters (i.e the parameters of the 3D shape model) can be used for generating the geometry of the face to be reconstructed. When both the transformation and deformation parameters are applied to all vertices of the 3D model (so far all work was done using the 68 selected vertices) it is trivial to generate the texture of the 3D model by sampling the 2D image at the positions where each vertex is projected. As a result both the geometry and the texture of the face are defined and the 3D reconstruction process is completed.

When the proposed algorithm was implemented in MATLAB (www.mathworks.org) the 3D reconstruction process needs about one minute to be completed. However, this time does not include the time required for locating the 68 landmarks on the face to be reconstructed. About five minutes are also required for locating the 68 landmarks on the face, using the semi-automatic tool.

3.3 Experimental Evaluation

In order to assess quantitatively the accuracy of the proposed method in reconstructing 3D faces, we staged an experimental evaluation. For our experiments we used a dataset containing 100 3D human faces (see figure 4) captured with a 3D laser scanner. We have used 60 3D faces among the dataset for training the system and the remaining 40 3D faces were used for testing the reconstruction accuracy.

During the experiments we train a 3D shape model using the 60 samples from the training set. We then project all 40 faces from the test set to a frontal 2D face image and we define the coordinates of the 68 landmarks located on each test face image. The reconstruction algorithm described in section 3.2 is then used for reconstructing the 3D geometry and texture of the 2D projected faces from the test set. Since for each test face we know the actual geometry and texture of the corresponding 3D face, we obtain quantitative measurements that assess the similarity between reconstructed and actual 3D faces by calculating the Euclidean distance between the actual and reconstructed vertices and color intensities. The results of the experiments are shown in Table 7.

	Mean	Standard Deviation
Shape Error	0.0298 mm	0.0078 mm
Intensity Error	0.15*	0.046*

* When using a 0-255 range of color intensities

Table 7: Quantitative evaluation results

The results of the experiment indicate that the accuracy of reconstructed faces is satisfactory. Bearing in mind that the precision of commercially available 3D scanners is around 0.05mm the performance of the proposed method is deemed adequate for most relevant applications.

3.4 Reconstructing Faces in CH Artefacts

Figure 9 (see page 8 of the paper) shows preliminary reconstruction results when we tested the proposed 3D reconstruction method in reconstructing faces appearing in CH artefacts of different styles and types. The results demonstrate that this method can be used for obtaining plausible reconstructions of the 3D facial appearance of the faces shown in the corresponding artefacts.

4. CONCLUSIONS

In this paper we have provided an overview of image-based 3D face reconstruction methods described in the literature. In particular we described typical example-based, stereo-based and video-based methodologies. In each case we provided a discussion pertaining to the applicability of the methods for reconstructing faces appearing in CH artefacts. According to the relevant discussion a successful method for CH applications should be able to reconstruct a face based on a single view but in order to deal with the diversity of texture variation it should generate the texture of the model through a sampling rather than a generative process.

Based on the conclusions of the review we developed a face reconstruction algorithm, which uses a 3D shape model as the basis of estimating the geometry of a face. During the process, the deformation parameters of a 3D shape model that best approximate the positions of 68 landmarks located on the face to be reconstructed are defined enabling in that way the definition of the 3D geometry of a face. A direct texture mapping approach is then used for generating the texture of the resulting 3D model. The results of a preliminary quantitative performance evaluation prove that the proposed method is able to reconstruct the 3D geometry of frontal faces appearing in images with reasonable accuracy.

Visual results indicate that the proposed method performs well in reconstructing faces appearing in CH artefacts despite the fact that the method was applied to cases involving significant shape and texture variation. In the future we plan to carry out more work in this area in order to deal with the following issues.

- Since the texture is collected directly from the image currently the method is only applicable to reconstructing faces with approximately frontal view in the given image. In the future we plan to upgrade the texture mapping process so that it will be possible to generate the texture of the missing parts based on the visible texture. For this purpose techniques that employ symmetry constraints and related face restoration techniques (Draganova, 2005) will be used. As part of this effort we also plan to augment the basic 3D shape model with 3D models describing the shape of occluding structures that may be encountered in such applications (i.e beards) so that such structures could be added on reconstructed faces. It is anticipated that our overall effort along these lines will lead to efficient and accurate digital restoration techniques that can be applied to damaged faces.
- Currently the reconstruction method relies on the use of a semi-automatic technique for locating 68 landmarks on the face to be reconstructed. In the future we plan to fully automate this process. In order to deal with this issue we plan to use an extended training set that includes faces on different types of artefacts, so that customized templates for specific artefact types will be used during the process of feature location.
- Although the initial quantitative and visual results are promising, we plan to stage a more rigorous quantitative evaluation of the reconstruction accuracy. During the process we plan to reconstruct 3D models of faces of statues and compare the reconstructed models with the actual 3D geometry obtained using 3D laser scanners. We also plan to investigate the sensitivity of the method in reconstructing non-frontal faces.

The ultimate aim of our work is to use the results of our work in different CH related applications. For example we plan to use 3D models obtained using the proposed method for generating complete models that can be used in animations, games and other educational software applications. Also we plan to use the generated 3D models for assessing trends and styles of different artists. For example in figure 8 we show two mosaics of Ktisis (a figure that symbolizes creation) and the corresponding 3D reconstructions. Differences in the face geometry can be used for obtaining conclusions related to the design style adopted in each case.

Figure 8: 3D Reconstructions of Ktisis from two different mosaics. The first column shows the original face and the remaining columns show the reconstructed face as seen from different views. Comparison of the 3D geometry of the two reconstructed models will enable the comparison of different styles and trends.

ACKNOWLEDGEMENTS

This work was supported by a Cyprus Research Promotion Foundation research grant (EPYAN/0205/08).

REFERENCES

Akimoto, T., Suenaga, Y., Wallace, R.S., 1993. Automatic creation of 3D facial models. IEEE Computer Graphics and Applications.

Atick, J.J., Griffin, P.A., Redlich, N.A., 1996. Statistical approach to shape from shading: Reconstruction of 3D face surfaces from single 2d images. Neural Computation. Vol. 8, pp. 1321-1340.

Blanz, V., Vetter, T., 1999. A morphable model for the synthesis of 3D faces. ACM Siggraph. pp. 187-194.

Brand, M., 2001. Morphable 3D models from video. Technical report, Mitsubishi Electric Research Lab.

Bregler, C., Hertzmann, A., Biermann, H., 2000. Recovering non-rigid 3d shape from image streams. Proc. of IEEE Conf. on Computer Vision and Pattern Recognition.

Chen, Q., Medioni, G., 2001. Building 3-D human face model from two photographs, Journal of VSLI Signal Processing, pp. 127-140.

Chowdhury, A.R.K., Chellappa, R., Vo, T., Krishnamurthy, S., 2002. 3D face reconstruction from video using a generic model. Proc of Int. Conf. on Multimedia. pp. 449-452.

Cootes, T.F., Edwards, G.J., Taylor, C.J., 2001. Active Appearance Models. IEEE Transactions of Pattern Analysis and Machine Intelligence, 23, pp 681-685.

Draganova, C., Lanitis, A., Christodoulou, C., 2005. Restoration of Partially Occluded Shapes of Faces Using Neural Networks. Procs. of the IV International Conference on Computer Recognition Systems.

Dryden, L. Mardia, K.V., 1998. Statistical Shape Analysis. John Wiley & Sons.

Fua, P., Miccio, P., 1999. Animated heads from ordinary images: A least squares approach. Computer Vision and Image Understanding. 75(3), pp. 247-259.

Hu, Y., Jiang, D., Yan, S., Zhang, L., Zhang, H., 2004. Automatic 3D reconstruction for face recognition. IEEE Int. Conf. on Automatic Face and Gesture Recognition.

Leclercq, P., Liu J., Woodward A., Delmas, P., 2004. Which stereo matching algorithm for accurate 3d face creation?, Int. Workshop on Combinatorial Image Analysis, Lecture Notes in Computer Science, 2004.

Lee, W., Kalra, P., Thalmann, N.M., 1997. Model based face reconstruction for animation. Int. Conf. on Multimedia Modeling.

Lee, J., Moghaddam, B., Pfister, H., Machiraju, R., 2003. Silhouette-based 3D face shape recovery. Graphics Interface.

Lee, J., Machiraju, R., Pfister, H., Moghaddam, B., 2005. Estimation of 3D faces and illumination from single photographs using a bilinear illumination model. Proc. of the Eurographics Symposium on Rendering Techniques. pp. 73-82.

Leung, W.H., Tseng, B.L., Shae, Z., Hendriks, F., Chen, T., 2000. Realistic video avatar. IEEE Multimedia and Expo.

Liu, Z., Zhang, Z., Jacobs, C., Cohen, M., 2001. Rapid modeling of animated faces from video. The Journal of Visualization and Computer Animation. 12(4), pp. 227-240.

Moghaddam, B., Lee, J., Pfister, H., Machiraju, R., 2003. Model-based 3d face capture with shape-from-silhouettes. Analysis and Modeling of Faces and Gestures.

Park, I.K., Zhang, H., Vezhnevets, V., 2005. Image-based 3d face modeling system. EURASIP Journal on Applied Signal Processing. No. 13, pp. 2072-2090.

Romdhani, S., 2005. Face Image Analysis using a Multiple Features Fitting Strategy. Ph.D. thesis, University of Basel.

Stylianou, G., Lanitis, A., 2008. Image Based 3D Face Reconstruction: A Survey. International Journal of Image and Graphics, to appear.

Xin, L., Wang, Q., Tao, J., Tang, X., Tan, T., Shum, H., 2005. Automatic 3D face modeling from video. IEEE Int. Conf. on Computer Vision.

Yin, S., Cagan, J. 2000. An extended pattern search algorithm for three-dimensional component layout. Journal of mechanical design , Vol. 122, No 1, pp. 102-108.

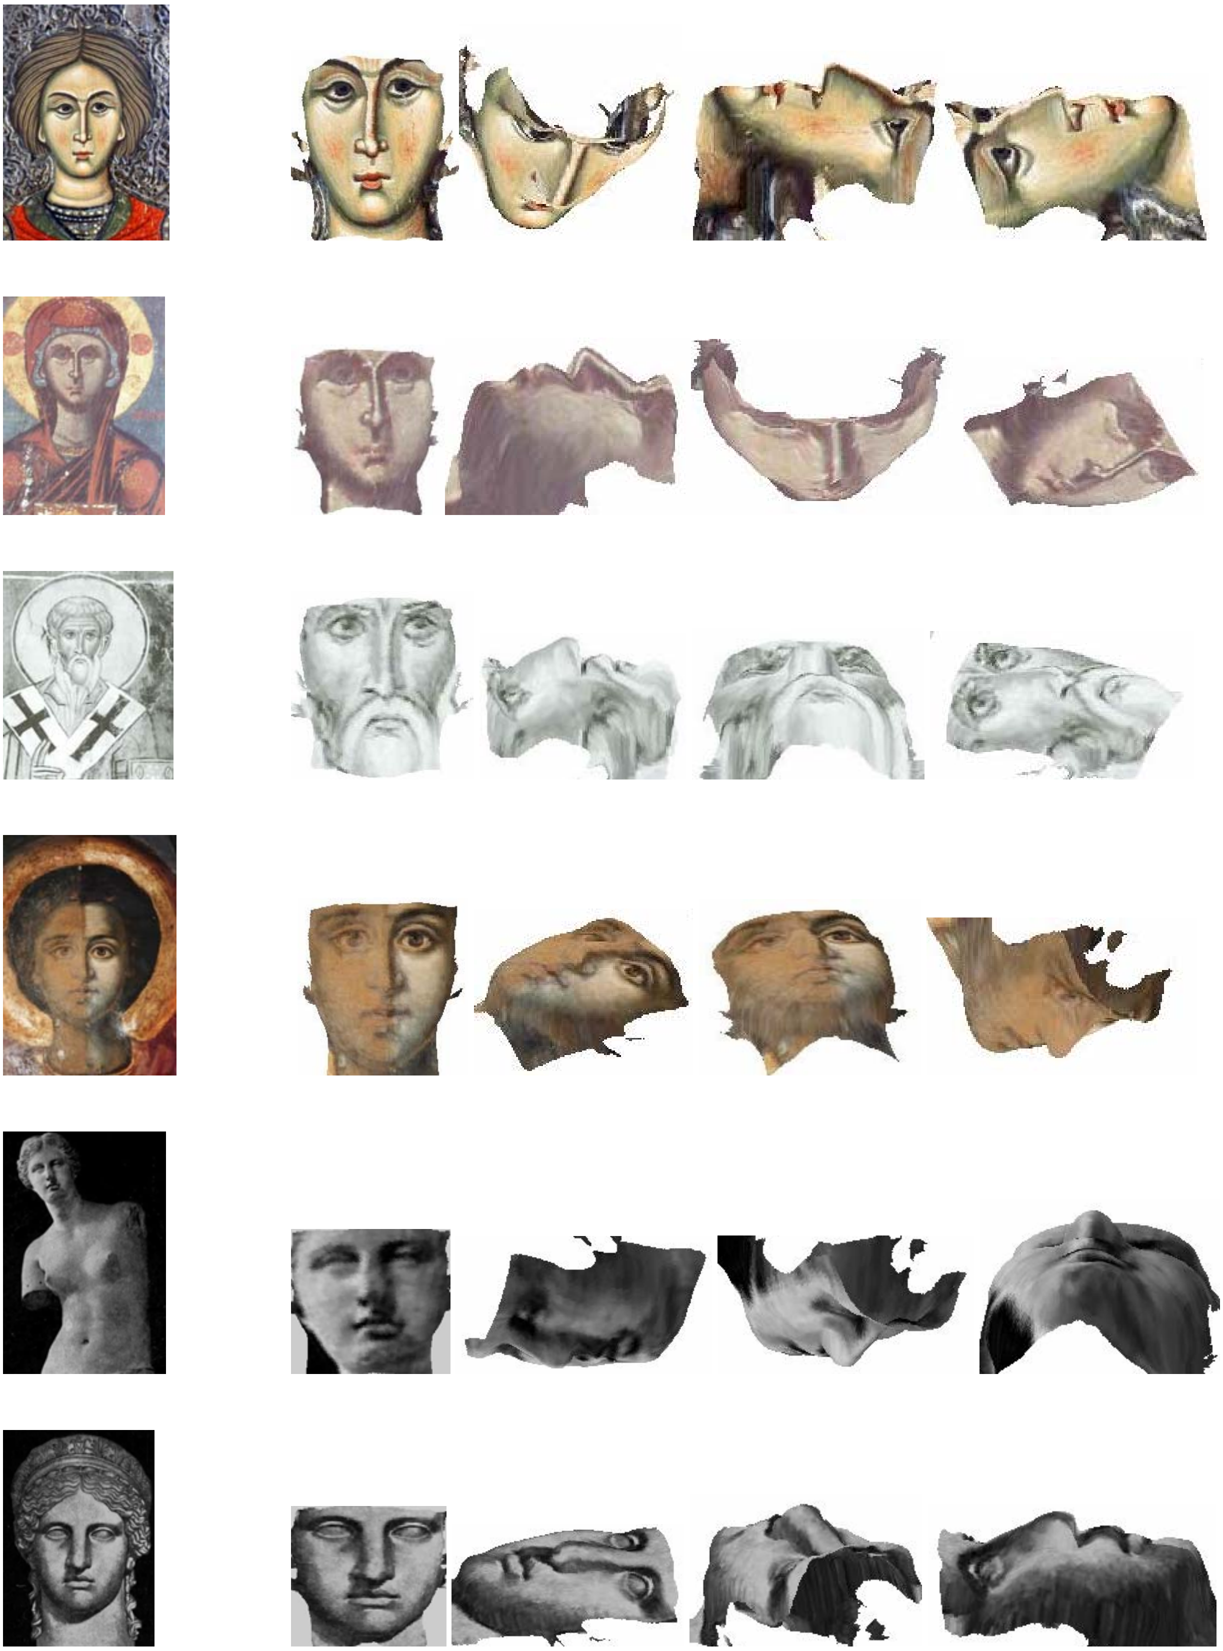

Figure 9: 3D reconstruction of faces appearing in CH artefacts. The first column shows the original face and the remaining columns show the reconstructed face as seen from different views.