

Automatic Face Recognition from Skeletal Remains

Peter Tu, Rebecca Book, Xiaoming Liu, Nils Krahnstoever
GE Global Research

Carl Adrian, Phil Williams
Federal Bureau of Investigation

Abstract

The ability to determine the identity of a skull found at a crime scene is of critical importance to the law enforcement community. Traditional clay-based methods attempt to reconstruct the face so as to enable identification of the deceased by members of the general public. However, these reconstructions lack consistency from practitioner to practitioner and it has been shown that the human recognition of these reconstructions against a photo gallery of potential victims is little better than chance. In this paper we propose the automation of the reconstruction process. For a given skull, a data-driven 3D generative model of the face is constructed using a database of CT head scans. The reconstruction can be constrained based on prior knowledge such as age and or weight. To determine whether or not these reconstructions have merit, geometric methods for comparing reconstructions against a gallery of facial images are proposed. First, Active Shape Models are used to automatically detect a set of facial landmarks on each image. These landmarks are associated with 3D points on the reconstruction. Direct comparison of the reconstruction is problematic since in general the camera geometry used for image capture is unknown and there are uncertainties associated with the reconstruction and landmark detection processes. The first method of comparison uses constrained optimization to determine the optimal projection of the reconstruction on to the image. Residuals are then analyzed resulting in a ranking of the gallery. The second method uses boosting to learn which points are both reliable and discriminating. This results in a match/no-match classifier. Experimental evidence indicating that skull recognition from facial images can be achieved is presented.

1. Introduction

When an unknown skull is found at a crime scene, law enforcement officials attempt to reconstruct the victim's face so that it can be presented to the public and/or compared against a gallery of facial images of missing persons. Traditional approaches often rely on clay based methods [23, 2]. However [22] has shown that when persons




Figure 1. The identity of a skull (top-left) is established by comparing a statistical skin-surface reconstruction (top-right) and associated landmarks with a set of images depicting missing persons (bottom) for which corresponding landmarks have been estimated using an active appearance model.

not familiar with the deceased are asked to compare clay based reconstructions against facial images, the resulting recognition rates are little better than chance. There are a number of contributing factors. Clay based reconstruction is a highly subjective process and results can vary drastically from practitioner to practitioner. Soft-tissue uncertainty associated with regions such as the tip of the nose can be considerable. Humans seem to rely on various appearance cues that cannot always be inferred directly from the bone and they are often distracted by inaccuracies incurred during reconstruction. In this work a fully automatic face reconstruction process is proposed. To validate this approach, geometric methods are used to automatically compare these reconstructions against a facial image gallery of missing persons (See Figure 1).

In previous work [24] we have proposed the collection of a CT head scan database for the purposes of creating a face space that is constrained by the shape of an unknown skull. Skin and bone surfaces are extracted from the CT head scans. Via bone-to-bone registration of the database against the unknown skull, each skin surface becomes an estimate

of the unknown face. In this work, the warped skin surfaces along with their age and weight information are represented as structurally identical vectors. This means that vector elements representing semantically identical points such as the tip of the nose and the corners of the eyes will be in correspondence. Principle Component Analysis (PCA) on these skin vectors constitutes a face-space tailored to the unknown skull. Prior knowledge such as the estimated age and/or weight of the unknown individual can be used to constrain the face space.

The automatic recognition of an unknown skull against a gallery of facial images of missing persons is a difficult task. There is uncertainty associated with the soft-tissue of the deceased and information regarding the imaging conditions of the gallery will be imperfect. To our knowledge this type of automatic recognition has never been accomplished. In this paper we seek evidence that the task of automatic recognition is achievable. Our approach starts by taking advantage of the fact that each reconstruction is structurally identical. Landmark positions in the images can then be associated with specific points on the reconstructions. Such image landmarks are found using Active Appearance Models (AAM). The task at hand is to determine whether or not a given set of reconstruction and image landmarks constitutes a match. Two approaches are considered. In the first method, constrained optimization based on the landmark positions with the least amount of soft-tissue variance is used to determine a projective transformation between the 3D reconstruction and the 2D image. Projection residuals are then calculated providing a ranking of the missing person image. In the second approach, boosting is used to construct a strong match/no-match classifier. In this way we learn which landmarks are both reliable and discriminating.

2. Face Reconstruction

A database of CT head scans of over 280 individuals across 6 demographics has been collected. The age and weight of each individual were also recorded. For purposes of validation, a frontal and profile image of each subject was taken. The marching cubes algorithm [16] was used to extract a polygonal mesh of the skin and bone surfaces of each individual. Each skin surface was manually labeled with a number of specific landmark points and using an approach similar to [17] they are represented by structurally identical skin vectors of the form

$$\mathbf{X}_i = [w_i, a_i, x_{i1}, y_{i1}, z_{i1}, \dots, x_{iM}, y_{iM}, z_{iM}] \quad (1)$$

where (x_{ij}, y_{ij}, z_{ij}) is the j^{th} point on the skin of individual i which is semantically identical to the j^{th} point of all other skin vectors. For example if the j^{th} point for a particular skin vector corresponds to the tip of the nose, then the j^{th} point for all skin vectors will correspond to the tip

of the nose. The value w_i and a_i are the age and weight of individual i . Each skin vector is based on M surface points.

When presented with an unknown skull, a registration and deformation is automatically computed between each database bone surface and the unknown skull. Details on this step can be found in [25]. Based on this registration process, thin plate splines [6] are used to warp each skin vector onto the unknown skull. This process is illustrated in Figure 2. Like the work of [4], Principle Component Analysis (PCA) can be used to create a face space. However, the morphing process restricts the face space based on the shape of the unknown skull, hence this face space is tailored to the unknown individual.

Let $\bar{\mathbf{X}}$ and \mathbf{C} equal to the mean and covariance of the warped skin vectors. The eigenvectors and diagonal matrix of eigenvalues of \mathbf{C} are \mathbf{E} and λ . Each reconstruction \mathbf{R} is a weighted sum of the eigenvectors added to the mean and can be formulated as:

$$\mathbf{R} = \bar{\mathbf{X}} + \mathbf{E}\mathbf{y} \quad (2)$$

where \mathbf{y} is a unique set of weights defining the reconstruction. Figure 3 shows an example of a slice through a face space.

For \mathbf{y} equal to zero, the reconstruction is reduced to the average face. In the absence of any *a priori* knowledge, this is the most likely and hence a reasonable estimate of the unknown face. However, if prior knowledge such as age and/or weight estimates of the unknown individual are available, this knowledge can be used to constrain the reconstruction. A vector \mathbf{c} is defined such that its i^{th} element is set to 0 unless this element corresponds to a component of the skin vector that can be estimated in which case it is set to the estimated value minus the mean value of this component. An indicator matrix \mathbf{S} is defined such that $\mathbf{S}(\mathbf{i}, \mathbf{i}) = \mathbf{1}$ iff $c(i) \neq 0$. All other elements of \mathbf{S} are set to 0.

A cost functional $\mathbf{Q}(\mathbf{y})$ that is to be minimized can now be defined as:

$$\mathbf{Q}(\mathbf{y}) = (\mathbf{S}\mathbf{E}\mathbf{y} - \mathbf{c})^T(\mathbf{S}\mathbf{E}\mathbf{y} - \mathbf{c}) + \alpha\mathbf{y}^T\lambda^{-1}\mathbf{y} \quad (3)$$

The first part of the cost function penalizes a reconstruction that deviates from the constraints. The second part of the function is a regularization term that keeps the reconstruction from deviating too far from the statistical mean of the population. The α term is user defined.

By setting $\frac{d\mathbf{Q}(\mathbf{y})}{d\mathbf{y}}$ equal to zero and solving for the optimum value $\hat{\mathbf{y}}$ we find that

$$\hat{\mathbf{y}} = (\alpha\lambda^{-1} + \mathbf{E}^t\mathbf{S}^t\mathbf{S}\mathbf{E})^{-1}\mathbf{E}^t\mathbf{S}^t\mathbf{c} \quad (4)$$

In general there will be a unique solution for a given choice of α , however if the matrix $(\alpha\lambda^{-1} + \mathbf{E}^t\mathbf{S}^t\mathbf{S}\mathbf{E})$ is rank deficient, singular value decomposition can be used and the

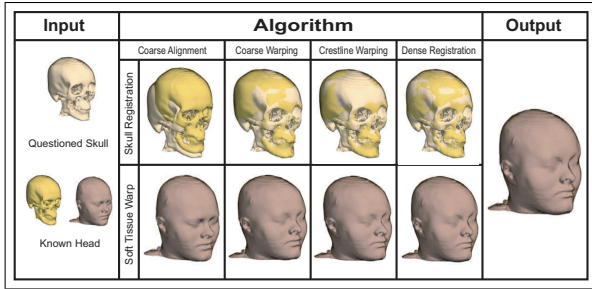


Figure 2. The registration of a database skull onto an unknown skull is achieved in four stages. The first generates a rough alignment of the skulls. The second aligns the crest-line points. The third generates a tighter alignment of crest-line points. The fourth aligns both low curvature and crest-line points. The transformations are also applied to the database skin resulting in an estimate of the unknown face. At each stage we see the unknown skull in gray and the warped database skull in yellow.

user must choose a reconstruction from the resulting null space.

Note that prior knowledge need not be restricted to age and weight, it can also be the location of specific points on the skin such as the tip of the nose. Figure 4 shows a set or reconstructions covering a range of ages and weights. This form of estimation can be viewed as a constrained solution. In contrast, approaches such as [14] perform age and weight progression by establishing trajectories in a general face space covering a large population - one major difference is that our approach does not require multiple samples of the same individual at different ages and weights.

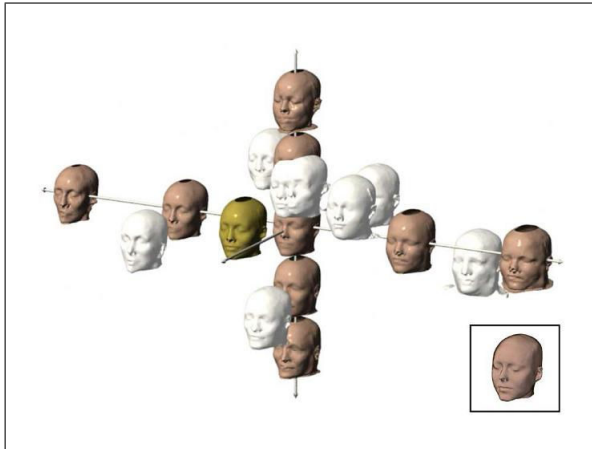


Figure 3. Cross-section of a tailored face space: The gray (light) faces are entries from the database warped onto an unknown skull. The tan (dark) faces are synthesized from the face space along the axes of the first two eigenvectors. The face in the bottom right corner is the true face to be estimated. The gold face (dark face left from center) is the projection of the true face into the face space.

Work related to automatic facial reconstruction can be

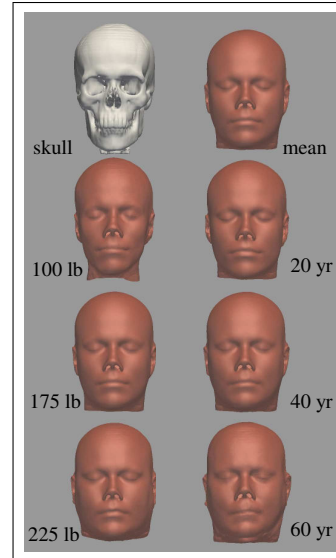


Figure 4. This figure shows an unknown skull, the mean reconstruction and a variety of reconstructions constrained by age or weight.

found in [18, 26, 12, 19, 8]. To our knowledge no other reconstruction system is based on a CT database of the magnitude presented in this work.

3. Face Recognition

Many appearance cues such as facial hair and skin texture which are important for face recognition by humans [20] cannot be accurately estimated based solely on knowledge of the skull. In order to validate the efficacy of our reconstructions we seek an automatic method of recognition that does not require such information. The problem at hand is how best to automatically compare a 3D reconstruction with a 2D facial image and to determine whether or not there is a match. Kakadiaris *et al.* [13] have shown that 3D structure plus appearance information can be used to recognize 2D facial images. However in our application, appearance information is not available. Blanz and Vetter [5] have shown that 3D models can be generated from 2D images. Iterative Closest Point algorithms [1] could then be used to perform recognition. Conversely, as shown by Viola and Wells III [27], mutual information can be used to directly compare a 3D reconstruction with 2D images. However, it is not clear that these methods would be successful under the uncertainty associated with the facial reconstruction process. Given an image that is to be compared to a reconstruction, specific landmarks such as the tip of the nose and the corners of the eyes can be defined automatically using Active Shape and Appearance Models [3]. Since each reconstruction is structurally identical, equivalent 3D landmarks on the reconstructions will be known in advance. The recognition of known 3D objects based on

landmark correspondences with their images is a well studied problem in computer vision, however there is relatively little work in the area of face recognition based on comparing 3D landmarks to 2D image landmarks [7]. Two approaches are considered: projective registration and learning via boosting.

3.1. Projective Registration

As a first approach we evaluate possible matches between unknown skulls and images of missing individuals by comparing 3D reconstruction landmarks with their corresponding 2D image landmarks. We assume that the 2D image landmarks $\mathbf{x} = \{(\mathbf{x}_i, \mathbf{y}_i), \mathbf{i} = 1, \dots, \mathbf{v}\}$ are projections of the 3D reconstruction landmarks $\mathbf{X} = \{(\mathbf{X}_i, \mathbf{Y}_i, \mathbf{Z}_i), \mathbf{i} = 1, \dots, \mathbf{v}\}$ obtained with an unknown camera \mathbf{P}

$$\mathbf{x}_i = \mathbf{P}\mathbf{X}_i. \quad (5)$$

Given a sufficient number of 2D to 3D correspondences, the projection \mathbf{P} can be estimated. A good quantitative evaluation of the match is then given as the reprojection error

$$RMS^2 = \frac{1}{v} \sum_{i=1}^v \|\mathbf{x}_i - \mathbf{P}\mathbf{X}_i\|_2^2. \quad (6)$$

The datasets under consideration make this approach challenging for multiple reasons: First, inherent errors in the skull-to-skin surface reconstruction and errors in the 2D image landmark location estimation violate the assumption that the 2D landmarks are actually a good projection of the 3D landmarks under the true but unknown camera \mathbf{P} . In practice many landmarks, especially from the jaw contour and the eyebrows correspond poorly between the 3D and 2D datasets. Hence, only certain landmarks provide good guidance for camera estimation. Secondly, the camera estimation process can sometimes obtain reasonably good estimates even for mismatched datasets where a skin reconstruction is compared to an image of an incorrect individual. The reason for this is that important geometric facial properties such as aspect ratio can be obtained from any set of 3D face points by modifying the aspect ratio of the camera intrinsics. Hence such freedoms in the camera estimation process need to be prohibited.

We use the following approach for estimating the camera $\mathbf{P} = \mathbf{KR}[\mathbf{I} - \mathbf{c}]$. First, we assume a unit (or at least known) aspect ratio and zero skew as well as an image centered principle point. This leaves only seven degrees of freedom: rotation \mathbf{R} , translation \mathbf{c} and focal length in the camera matrix \mathbf{K} . These seven parameters are estimated by performing linear estimation of the unconstrained \mathbf{P} matrix, followed by non-linear least squares minimization of the reprojection errors with penalty terms for constraint violations, followed by LS minimization with the enforced

constraints. This first process is performed using all available landmarks. We then further refine the camera estimate with only a *subset* of the facial landmarks to overcome the landmark error challenge.

We employ two approaches for choosing the landmark subset. The first approach is to manually perform the selection based on experience and intuition for what are likely to constitute stable and accurate landmarks that are good for the camera estimation and subsequent reprojection error based match evaluation. The second approach is to use training data and exhaustively search for the optimal landmark subset.

3.2. Learning via Boosting

RANSAC [10] approaches have been shown to be robust to the noisy landmark problem. In this method a minimal number of correspondences are randomly selected to generate a transformation between the 3D reconstruction space and the 2D image space. This transformation is then applied to a subset of the remaining landmarks. Residuals are measured and analyzed. Based on this approach, a classifier using specific landmarks for transformations and residual comparisons can be used to discriminate between true and false image/reconstruction pairs. Unlike the previous section which relies on a single optimal projection, it is proposed that the AdaBoost algorithm [11] be used to generate a strong classifier based on a linear combination of weak landmark classifiers each with its own projection transformation. The weak classifiers are iteratively selected by minimizing the expected error associated with a training set of true and false matches. In this way we learn which landmarks are both stable and discriminating. Thus recognition becomes a process of classification between true and false matches.

For each weak classifier a more restrictive projective transform that is simple to compute is used. It is assumed that the face plane containing the corners of the eyes and the bottom of the chin is roughly parallel to the imaging plane. Points on the reconstruction can be mapped to the image by first projecting them on to the face plane and then applying a metric transform between the face plane and the image plane. A metric transform can be defined by as little as two correspondences. This transform can be applied to a third landmark on the reconstruction and if the resulting residual is less than a threshold, the weak classifier responds positively. If the converse is true, a negative response is reported. At each round of boosting the optimal weak classifier with respect to expected classification error is selected.

3.3. Facial Landmarks

We now consider the method used for extracting the 2D image landmarks. The shape model and appearance model part of an AAM are trained with a representative set of




Figure 5. The mean and first 7 basis vectors of the shape model (a) and the appearance model (b) trained from the IMM database. The shape basis vectors are shown as arrows at the corresponding mean shape landmark locations.

facial images. The procedure for building a shape model is as follows. Given a face database, each facial image is manually labeled with a set of 2D landmarks, $[x_i, y_i]$ $i = 1, 2, \dots, v$. The collection of landmarks of one image is treated as one observation from the random process defined by the shape model, $\mathbf{s} = [x_1, y_1, x_2, y_2, \dots, x_v, y_v]^T$. Eigen-decomposition is applied to the observation set and the resulting model represents a shape as,

$$\mathbf{s}(\mathbf{P}) = \mathbf{s}_0 + \sum_{i=1}^n p_i \mathbf{s}_i \quad (7)$$

where \mathbf{s}_0 is the mean shape, \mathbf{s}_i is the shape basis, and $\mathbf{P} = [p_1, p_2, \dots, p_n]$ are the shape parameters. By design, the first four shape basis vectors represent global rotation and translation. Together with other basis vectors, a mapping function from the model coordinate system to the coordinates in the image observation is defined as $\mathbf{W}(\mathbf{x}; \mathbf{P})$, where \mathbf{x} is a pixel coordinate within the face region $F(\mathbf{s}_0)$ defined by the mean shape \mathbf{s}_0 .

Given the shape model, each facial image is warped into the mean shape via a piecewise affine transformation. These shape-normalized appearances from all training images are processed via eigen-decomposition and the resulting model represents an appearance as,

$$\mathbf{A}(\mathbf{x}; \lambda) = \mathbf{A}_0(\mathbf{x}) + \sum_{i=1}^m \lambda_i \mathbf{A}_i(\mathbf{x}) \quad (8)$$

where \mathbf{A}_0 is the mean appearance, \mathbf{A}_i is the appearance basis, and $\lambda = [\lambda_1, \lambda_2, \dots, \lambda_m]$ are the appearance parameters. Note that the resolution of the appearance model here is the same as the resolution of the training images. Figure 5 shows an AAM trained using 40 images from the IMM face database [21].

An AAM can synthesize facial images with arbitrary shape and appearance within a population. Thus, the AAM can be used to *explain* a facial image by finding the optimal shape and appearance parameters such that the synthesized image is as similar to the image observation as possible.

This leads to the cost function used for model fitting [9],

$$J(\mathbf{P}, \lambda) = \frac{1}{N} \sum_{\mathbf{x} \in F(\mathbf{s}_0)} \|I(\mathbf{W}(\mathbf{x}; \mathbf{P})) - \mathbf{A}(\mathbf{x}; \lambda)\|^2 \quad (9)$$

which is the mean-square-error (MSE) between the warped observation $I(\mathbf{W}(\mathbf{x}; \mathbf{P}))$ and the synthesized appearance instance $\mathbf{A}(\mathbf{x}; \lambda)$, and N is the total number of pixels in $F(\mathbf{s}_0)$.

Traditionally this minimization is solved by gradient-descent methods. Baker and Matthews [3] proposed the Inverse Compositional (IC) and Simultaneously Inverse Compositional (SIC) method that greatly improves the fitting performance. Their basic idea is that the role of the appearance templates and the input image is switched when computing $\Delta \mathbf{P}$. Thus the time-consuming steps of parameter estimation can be pre-computed and remain constant during each iteration. Liu *et al.* [15] introduced a face model enhancement scheme, where face modeling and model fitting are iteratively performed using the training image set, to improve the fitting convergence. In this paper, we employ the enhanced ASM/AAM and the SIC method to process the set of 2D images.

4. Experiments

For purposes of validation, a testbed of 35 reconstructions of Caucasian males with known frontal images has been assembled. No prior knowledge was assumed, hence the mean reconstructions were used. Each image was labeled with 48 landmark positions using the methods described in section 3.3. All possible image/reconstruction pairs were processed using the projective registration approach described in section 3.1. It was found that the average ranking of the true reconstruction was 11.1 based on user defined landmarks. The user selected landmarks where $\{0, 1, 19, 17, 18, 45, 3, 16, 7, 13\}$ (see Fig. 6) and resulted in the Cumulative Match Characteristic (CMC) graph is shown in Figure 7. To determine how representative the user selection was, an exhaustive search was performed on a larger landmark subset. During the search six-tuples of landmarks where selected from the subset $\{0, 1, 19, 17, 18, 45, 3, 16, 7, 13, 41, 43\}$ and the selection that obtained the best average rank was recorded. This experiment obtained the optimal subset $\{0, 19, 18, 45, 3, 16\}$, with an average rank of 9.97 and hence only slightly better than the expert selected landmarks. Both of these average rankings are significantly higher than 17.5 which would be the expected result of a totally random process.

In the next experiment, we consider the boosting approach described in section 3.2. In this experiment, the database was divided into two, one for training and one for testing. The training database consists of 31 image/reconstruction pairs. This results in 31 true matches




Figure 6. The landmark set used for registration and matching.

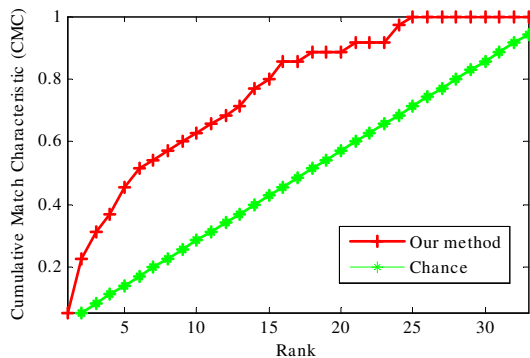


Figure 7. The Cumulative Match Characteristic Graph. For validation purposes a database of 35 Caucasian males was assembled. All possible image/reconstruction pairs were processed using the projective registration approach described in section 3.1. The horizontal axis shows the ranking of the true match. The vertical axis shows the percentage of subjects that achieve a true match ranking greater than a given value. The red curve shows the performance of our method and the green curve shows the performance that would be expected based on chance alone.

and 930 false matches. The testing database consists of 5 image/reconstruction pairs. Boosting was applied so as to generate a strong classifier and when the same classifier was used to process the testing data, a true positive rate of 0.8 and a false positive rate of 0.2 was observed. This results in a κ score of 0.49 with an 80% confidence interval of ± 0.24 . A κ score above 0.4 implies that it is reasonable to infer that the classifier is performing better than chance. Since the confidence interval is still relatively large, we believe that the collection of a larger database is warranted - it should be noted that for validation purposes, difficult to acquire CT head scans are not necessary - skulls from deceased individuals with associated photos will suffice.

5. Conclusions

In this work, we have presented a fully automatic approach to face reconstruction from skeletal remains. The system is based on a database of CT head scans, which to our knowledge is the only such database of this magnitude in existence. The ability to generate a person specific face space enables the user to incorporate prior knowledge such as estimates of age and weight directly into the reconstruction.

We have presented landmark face recognition approaches that enables the recognition of skull identity based on a gallery of facial images. To our knowledge this is the first time that such results have been reported. Initial experiments indicate that better than chance recognition rates are achievable. However, from a statistical point of view, it will be important to perform more experimentation on larger datasets.

From a law enforcement perspective, hundreds of unknown skulls are discovered each year. The ability to match these skulls against a database of missing persons would be a boon to law enforcement and victim's advocacy. The work presented here is a significant step towards this goal.

Future work will incorporate contour-based features and distance measures as well as shape from shading types of approaches for comparing face reconstructions with image galleries.

6. Acknowledgments

This work has been funded by a research grant from the Federal Bureau of Investigation.

References

- [1] B. B. Amor, M. Ardabilian, and L. Chen. New experiments on icp-based 3D face recognition and authentication. In *18th International Conference on Pattern Recognition (ICPR 2006)*, 20-24 August 2006, Hong Kong, China, pages 1195–1199, 2006.
- [2] W. A. Aulsebrook, M. Y. İşcan, J. H. Slabbert, and P. Becker. Superimposition and reconstruction in forensic facial identification: a survey. *Forensic Sci Int*, 75:101 – 120, 1995.
- [3] S. Baker and I. Matthews. Lucas-Kanade 20 years on: A unifying framework. *IJCV*, 56(3):221–255, March 2004.
- [4] V. Blanz and T. Vetter. A morphable model for the synthesis of 3D faces. In *Computer Graphics Proceeding, Annual Conference Series*, pages 187 – 194, New York, 1999. ACM SIGGRAPH.
- [5] V. Blanz and T. Vetter. Face recognition based on fitting a 3D morphable model. *PAMI*, 25(9):1063–1074, 2003.
- [6] F. Bookstein. Principle warps: Thin plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:567–585, 1989.
- [7] K. W. Bowyer, K. Chang, and P. J. Flynn. A survey of approaches and challenges in 3D and multi-modal 3D+2d face

- recognition. *Computer Vision and Image Understanding*, 101(1):1–15, 2006.
- [8] P. Claes, D. Vandermeulen, S. D. Greef, G. Willems, and P. Suetens. Statistically deformable face models for craniofacial reconstruction. *Journal of computing and information technology*, 2006.
- [9] T. Cootes, G. Edwards, and C. Taylor. Active appearance models. *PAMI*, 23(6):681–685, June 2001.
- [10] M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Comm. of the ACM*, 24:381–395, 1981.
- [11] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- [12] K. Kähler, J. Haber, and H.-P. Seidel. Reanimating the dead: reconstruction of expressive faces from skull data. In *ACM Transactions on Graphics. Proceedings of ACM SIGGRAPH 2003; 2003 July 28 – 31; San Diego (CA)*, pages 554 – 561, New York, 2003. ACM SIGGRAPH.
- [13] I. Kakadiaris, G. Passalis, G. Toderici, N. Murtuza, and T. Theoharis. 3D face recognition. In *Proc. British Machine Vision Conference, Edinburgh, United Kingdom*, pages 200–208, 2006.
- [14] A. Lanitis, C. Taylor, and T. Cootes. Modeling the process of ageing in face images. In *International Conference on Computer Vision 1999*, pages 131–136, 1999.
- [15] X. Liu, P. Tu, and F. Wheeler. Face model fitting on low resolution images. In *Proc. 17th British Machine Vision Conference, Norwich, UK*, volume 3, pages 1079–1088, 2006.
- [16] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *Computer Graphics*, volume 21, pages 163 – 169. ACM, 1987.
- [17] C. Lorenz and N. Krahnstoeber. Generation of point-based 3D statistical shape models for anatomical objects. *Computer Vision and Image Understanding*, 77:175–181, 2000.
- [18] L. A. Nelson and S. D. Michael. The application of volume deformation to three-dimensional facial reconstruction: A comparison with previous techniques. *Forensic Sci Int*, 94:167 – 181, 1998.
- [19] G. Quatrehomme, S. Cotin, G. Subsol, H. Delingette, Y. Garidel, G. Grévin, M. Fidrich, P. Baillet, and A. Ollier. A fully three-dimensional method for facial reconstruction based on deformable models. *J Forensic Sci*, 42(4):649 – 652, 1997.
- [20] P. Sinha and T. Poggio. I think I know that face... *Nature*, 384:404, Dec. 1996.
- [21] M. B. Stegmann, B. K. Ersboll, and R. Larsen. Fame - a flexible appearance modeling environment. *IEEE Transactions on Medical Imaging*, 22(10):1319–1331, 2003.
- [22] C. N. Stephan and M. Henneberg. Building faces from dry skulls: are they recognized above chance rates? *J Forensic Sci*, 46(3):432 – 440, 2001.
- [23] K. T. Taylor. *Forensic Art and Illustration*. CRC Press, New York, 2001.
- [24] P. Tu, R. I. Hartley, W. E. Lorensen, A. Alyassin, R. Gupta, and L. Heier. Face reconstruction using flesh deformation modes. In *Computer Graphic Facial Reconstruction, Chapter 8, ed J.G. Clement, M. K. Marks, Pub. Elsevier: Oxford, 2005, ISBN 0-12-473051-5*, pages 145–163, 2005.
- [25] W. D. Turner, R. E. B. Brown, T. P. Kelliher, P. H. Tu, M. A. Taister, and K. W. P. Miller. A novel method of automated skull registration for forensic facial approximation. *Forensic Science International*, 154:149 – 158, 2005.
- [26] P. Vanezis, R. W. Blowes, A. D. Linney, A. C. Tan, R. Richards, and R. Neave. Application of 3D computer graphics for facial reconstruction and comparison with sculpting techniques. *Forensic Sci Int*, 42:69 – 84, 1989.
- [27] P. A. Viola and W. M. W. III. Alignment by maximization of mutual information. In *Proceedings of the Fifth International Conference on Computer Vision (ICCV 95), June 20-23, 1995*, pages 16–23, 1995.