Face Reconstructions using Flesh Deformation Modes

Peter Tu[1], *Richard I. Hartley*[1], *William E. Lorensen*[1], *Majeid Allyassin*[1], *Rajiv Gupta*[2], *Linda Heier*[2]

[1] GE Corporate R&D Center
1 Research Circle
Niskayuna, N.Y. 12309

[2] Weill Medical College
Cornell Cornell University

Abstract

An important forensic problem is the identification of human skeletal remains. For this purpose the skull is commonly used as a means for building a face model, which may be used for recognition of the subject. This paper describes a method whereby a 3D computer model of the face and head is built using computer 3D graphics techniques. In reconstructing the shape of the face from a skull, it is essential to have data relating the shape of the face to the shape of the underlying skull. In manual reconstruction techniques this information is derived from flesh depth tables. We offer a method in which Computer Tomography (CT) scans are used to get a dense set of measurements of flesh depth. To model the variation of the facial shape, Principal Component Analysis is used as a means of exploring the space of the major facial deformations. After a subject skull and a set of database skulls have been aligned, the average face and principal deformations may be computed. Also provided are interactive tools for substituting facial features (such as nose, eyes, ears) from a catalog of such parts.

1 Introduction

In this paper, a database of CT (Xray Computer Tomography) head scans is used to reconstruct the face of an unknown subject given its skeletal remains. Skull and face surfaces can be extracted from each CT head scan. By establishing point correspondences between the CT skull and the subject's skull, the CT face can be morphed to coincide with the subject's skull. The morphing process removes differences that are caused by deviation in skull structure. Any remaining discrepancies can be attributed to variation in soft tissue.

By morphing each head scan in the CT database, a collection of estimates of the subject's face is generated. These estimates can be viewed as samples from a *face space* that is tuned to the subject's skull. Given these estimates, principal component analysis is applied to determine the main modes of deformation. Initially the operator is provided with the average estimate of the subject's face. The operator can then explore face space by applying the dominant deformations in a convenient manner.

The face space approach provides methods for estimating the overall structure of the subject's face. A set of post-processing tools allows the operator to define specific details.

Figure 1: *Method of reconstruction of a face model from a skull, using depth markers and modeling clay. The image shown here is of a 200-year old female skull found in Albany NY. Thanks to Gay Malin and the New York State Museum for permission to use this image.*

A face editing process is used to change the shape of features such as the nose, eyes and lips. Texture maps are tailored to the subject creating a lifelike appearance.

This paper is divided into three main parts. In section 2 details regarding development of the CT head scan database are provided. In section 3 the method used to create an estimate of the subject's face from a CT head scan is described. Finally, in section 4 various post-processing tools are demonstrated.

2 Collection and processing of CT database

A problem of basic forensic importance is to determine the identity of human skeletal remains. For this purpose the skull often gives the best available information. One approach to the identification problem is to superimpose images of peoples faces onto a given skull (see [1], [13] and [9]). This is essentially a method of verifying a hypothesized identification by determining the goodness of fit of a photographic image of a person with a correctly oriented image or the skull.

In the absence of a specific hypothesized identity, a commonly used practice is to build a model of the face using all available evidence, but essentially based on the skull shape. Traditional manual methods of reconstructing a face from a skull are based on tissue-depth tables ([4]). Depth markers are placed on the skull at strategic locations and are used to guide the depth of modeling clay that is used to build up a face over the top of the skull. An example of such a reconstruction process is shown in figure 1.

Instead of relying on tissue depth tables, our approach is to acquire tissue-depth information from a set of CT scans. In [7] CT scans were used to make specific flesh depth measurements. We utilize the entire dataset. Conceptually, this is a simple idea. The surface of the skull and the face are extracted from a CT data set, and from these tissue thickness may be measured. For useful statistics to be gathered from this data, it is necessary to align the scans to each other, and most of the description of the data-gathering

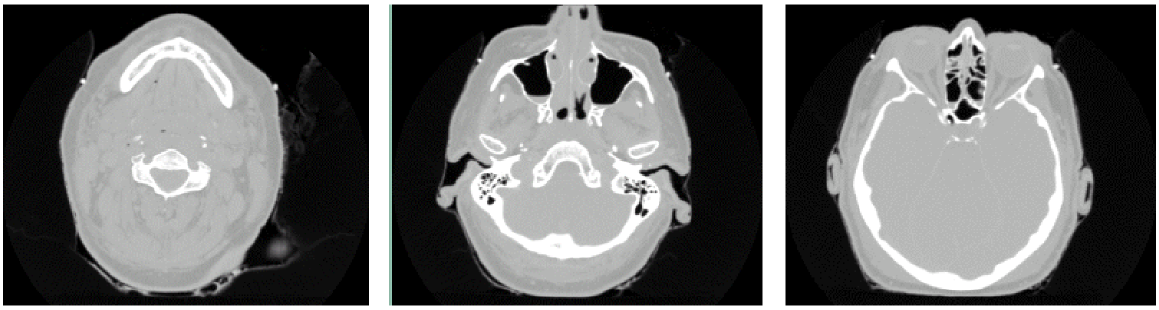

Figure 2: *Slices from a CT data set of a head.*

phase of this project will be concerned with this alignment process.

The output of a complete head-scan is a 3D block of intensity values (known as voxel data), which is usually viewed slice by slice by radiologists in clinical medical applications. Figure 2 shows some typical CT slices of a head. For the present application, however a 3D model needs to be built. Based on calibrated intensity values for flesh and bone, the surface of the bone and the air-flesh interface is extracted from the voxel data. A well known algorithm that gives excellent results is the "Marching Cubes" algorithm ([5]). Subsequently, the surface is triangulated to give a polyhedral model, and each triangle is shaded according to a lighting model and view direction to give a realistic image of the skull or skin surface. The model may be viewed from any direction. Figure 3 shows a rendered skull overlaid with the triangulated skin surface.

The complete process of data collection is made up from several steps, which will be explained more clearly later. The reader may refer back to these steps in reading the following sections.

1. Obtain a CT scan of the head
2. Apply a 3D reconstruction algorithm (Marching Cubes) to extract skull and face (to be more exact, skin) surfaces.
3. Select alignment points on the skull surface.
4. Use these alignment points to align the skull with a canonic skull via a scaled rigid transformation. Apply the same transformation to the face model. The canonic skull can be chosen arbitrarily. Its function is to make sure that all headscans are in the same relative position, orientation and scale.
5. Project both face and skull models into a cylindrical coordinate system, and produce an image-like representation of skull and face in which the radius coordinate represents image intensity. This will be referred to as the 2.5D skull or face model, namely a 2D image in which intensity represents radial depth of the surface. See figure 5.
6. Pick new fiducial points on the 2.5D skull model (see figure 9) to enable alignment with other models.

The fiducial points are used to align each database head scan to a target model, working with the 2.5D skull image model. The alignment takes place in two steps as follows.

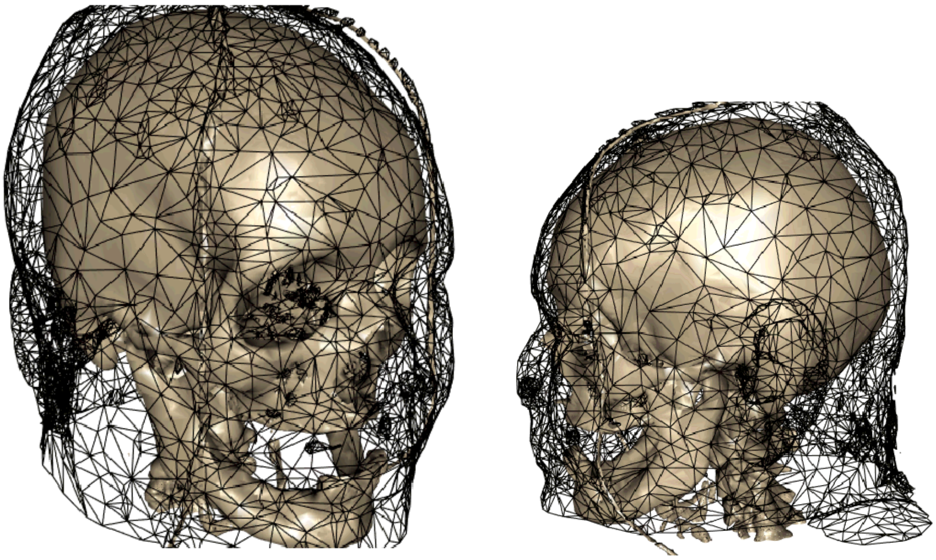

Figure 3: *Reconstructed skull surface and overlaid triangulation of the skin surface obtained from CT data set. The model is 3-dimensional and may be viewed from any direction.*

1. Fiducial points are used to align the head scan to the target model by 2D warping of the 2.5D image model. This 2D feature-based alignment is based on a concept of *saliency*, aligning sections of high variability in the skull images.
2. The second stage of morphing is depth adjustment, in which a depth offset δ_r is added at each pixel, based on fitting an interpolating surface.

This method of alignment through morphing is described in detail later.

2.1 Rigid alignment of skull

In order to standardize the data base, each set of skull and face surfaces is aligned with a canonical model. This process starts by establishing a set of point correspondences between the scanned skull and the canonical skull. Once this is done an optimal metric transform (rotation, translation, and uniform scaling) is computed ([3]) . This processing step is performed on all scanned data. Figure 4 illustrates this process. This process defines an overall alignment of the skull (and hence the face) with the canonical model, but without any deformation taking place. This alignment allows a cylindrical coordinate system to be aligned to the skull model, as will be seen later.

2.2 Non-rigid morphing of skull to model

Next the mechanism for morphing a single CT head scan onto a model skull is described. The main idea is to apply a transformation to the CT head scan so that its skull assumes the shape of the target skull. The same transformation is applied to the face of the

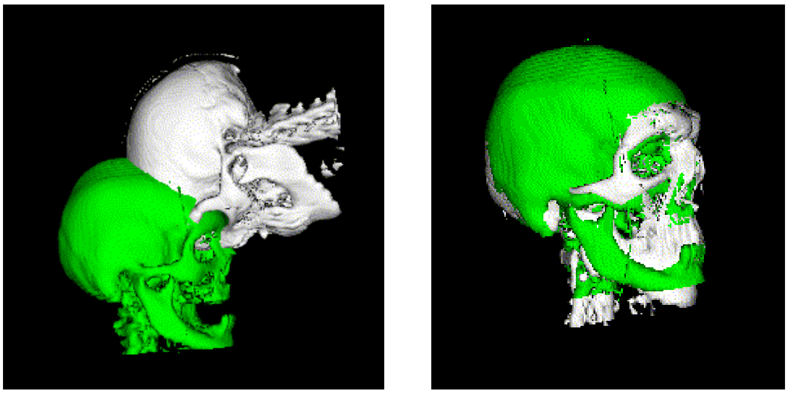

Figure 4: *The green skull is the canonical model and the white skull is the subject. The left image shows the initial misalignment. Once the operator specifies a few point correspondences, the two skulls can be aligned with respect to a scaled rigid transformation. The right image shows the two skulls after alignment.*

head scan resulting in an estimate of the model's face. This estimate will be reasonably accurate for areas with little soft-tissue variation such as around the eyes, the forehead and the cheekbones. In section 3.1 methods for modeling the remaining areas will be described.

Both 2D and 3D morphing rely on a set of point correspondences between a source (the skull of the head scan) and the target (the subject's skull). These correspondences are used to create a *warping field* where any point in the source space can be transformed into the target space (see [12]). In [11] face surfaces were added to digitized skulls using adjustments based on conventional flesh depth tables. In [6] it was shown that 3D volumetric morphing can be used to transform one skull onto another. In this paper, it is argued that the morphing mechanism should be divided into two steps: a 2D surface morph, where each point on the surface becomes registered, and a depth adjustment process. Both steps require a set of point correspondences to create a warping field. However, different criteria are used for selecting these control points. As shown in ([2]), a so-called 2.5D image-like representation of the data based on cylindrical coordinates provides a convenient framework for the morphing process. The PCA and post-processing methods described in sections 3.1 and 4 can also take advantage of this representation.

2.3 Cylindrical coordinate system

In [2] it was shown that it is useful to have an image-like representation of a 3D face surface. We have found that this is also true for a skull surface. Let the origin of a Cartesian coordinate system be placed at the center of the base of the head. The z axis points toward the top of the head. The x axis points toward the left side of the head and the y axis points toward the front of the head. A point can be defined by its Cartesian coordinate (x, y, z) or by its cylindrical coordinates

$$\theta = \tan^{-1}(\frac{y}{x})$$

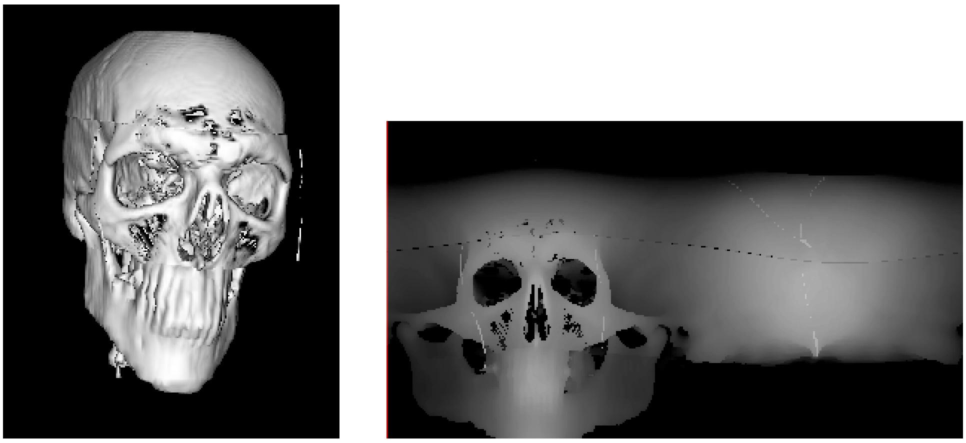

Figure 5: *The left image shows a 3D polygonal mesh of a skull. The right is a 2.5D representation of the skull. The horizontal axis, vertical axis and pixel intensity are associated with the cylindrical coordinates θ, h and r.*

$$
\begin{aligned}
h &= z \\
r &= \sqrt{x^2 + y^2}
\end{aligned}
$$

A 2D image is defined such that the horizontal axis corresponds to the angular coordinate θ, the vertical axis corresponds to the height coordinate h and the intensity of each pixel is set to the radial coordinate r. A surface defined by a polygonal mesh can be transformed into this image representation in the following way: a ray in space is defined for each pixel i with coordinates (θ_i, h_i). The point of intersection between the ray and the polygonal mesh is found. This establishes the value of r_i. The intensity of pixel i is set to r_i. In figures 5 and 6 a face and a skull are shown in this format. One advantage of this approach is that many 2D image processing operations can be performed on data in this format.

2.4 2D morphing

Given image representations of the head scan and subject skulls, a two step morphing process can be implemented. First, the image representation of the head scan skull is morphed in a two dimensional sense. That is, a warping field is used to align each pixel in the head scan image with a pixel in the subject's skull image. Second, an intensity alignment is performed so that the two images become almost indistinguishable. These operation are also performed on the face of the head scan. Figure 7 illustrates this process.

The first 2D morphing step is now described. A 2D warping field is defined so that pixels $(\hat{\theta}, \hat{h})$ in the source image are mapped to pixels (θ, h) in the target image. The warping field is defined as:

$$
\begin{aligned}
\theta &= F_\theta(\hat{\theta}, \hat{h}) \\
h &= F_h(\hat{\theta}, \hat{h})
\end{aligned}
$$

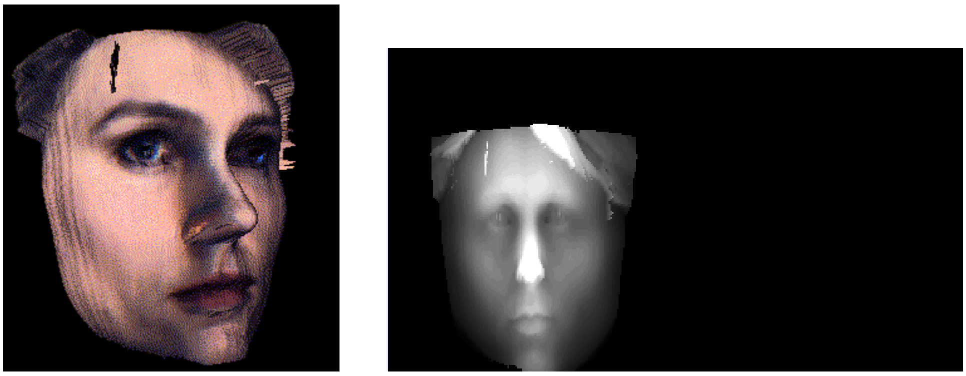

Figure 6: *The left image shows a 3D polygonal mesh of a face. The right is a 2.5D representation of the face. The vertical axis, horizontal axis and pixel intensity are associated with the cylindrical coordinates θ, h and r.*

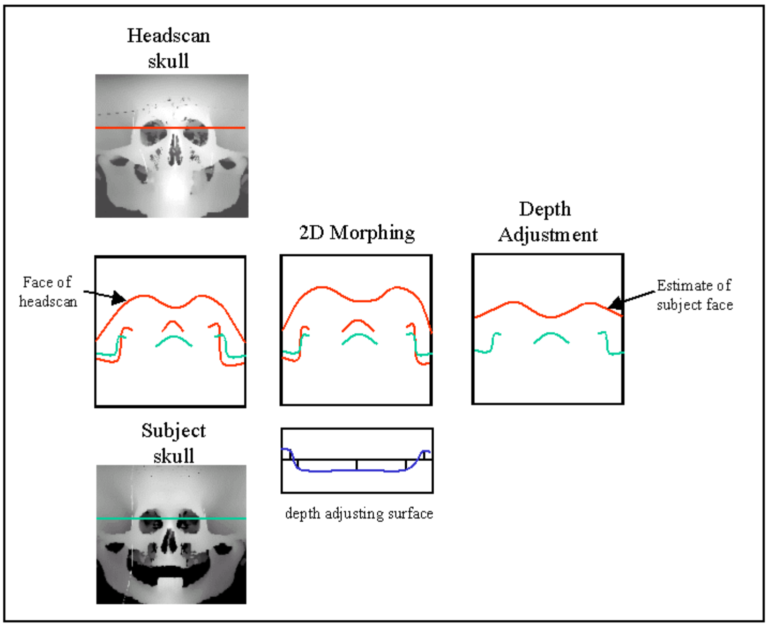

Figure 7: *This figure illustrates the two phases to the morphing process. A cross section of the two skulls is shown to illustrate the method better. The face of the head scan (from the database) is also shown in the cross section. In the first phase, the head scan undergoes a 2D morph so that each pixel becomes registered to the subject skull. This aligns the cross sections in the image plane. In the second phase, depth or intensity adjustment is performed. This aligns the cross sections in the radial direction. The processing steps applied to the head scan skull are also applied to it's face. The modified face becomes the estimate of the subject's face.*

where F_θ and F_h are 2D surfaces defined parametrically. Given a set of P point correspondences $(\hat{\theta}_i, \hat{h}_i) \rightarrow (\theta, h)$, two cost functions C_θ and C_h are defined as:

$$
\begin{aligned}
C_\theta &= \sum_1^P (\theta_i - F_\theta(\hat{\theta}_i, \hat{h}_i))^2 \\
C_h &= \sum_1^P (h_i - F_h(\hat{\theta}_i, \hat{h}_i))^2.
\end{aligned}
$$

The parameters for F_θ and F_h are determined such that C_θ and C_h are minimized is a least squares sense. There are various sorts of surfaces that can be used to define F_θ and F_h. In this system a set of continuous bicubic finite elements are used.

The criterion for selecting the initial point correspondences is based on saliency. This method favors points on the two images that are easily identifiable such as the region around the eye socked and boundaries of the cheek and jaw bone. In terms of image features, these are regions where there are large discontinuities in pixel intensity.

2.5 Skull depth alignment

Once the head scan pixels have been aligned with the subject pixels, an adjustment must be made on the r intensity values. The approach taken in this paper is to define an adjusting surface $\delta_r(\theta, h)$. Each head scan pixel i with coordinates (θ_i, h_i) will have its pixel intensity adjusted by adding the value of $\delta_r(\theta_i, h_i)$. Once again the relationship between the head scan skull and the subject skull will be used. The value of the adjusting surface cannot be measured for all points since the skull surfaces have holes. Therefore, the adjusting surface must be interpolated based on a set of Q points where skull depths can be accurately measured. Since the pixels are registered, the criterion for selecting the Q control points shifts from saliency to stability. Measurements should be taken where the bone image intensity values are changing slowly and continuously. Note that this is the exact opposite of the criterion taken for selecting points that define the 2D morph.

A least squares criterion for generating $\delta_r(\theta, h)$ is used:

$$
\sum_{i=1}^{Q} ((r_{skull}(\theta_i, h_i) - \hat{r}_{skull}(\theta_i, h_i)) - \delta_r(\theta_i, h_i))^2
$$

where $r_{skull}(\theta_i, h_i)$ and $\hat{r}_{skull}(\theta_i, h_i)$ are the pixel intensity values for the subject and the head scan skulls at the Q measurement points. As in the previous section $\delta_r(\theta, h)$ is a bicubic finite element surface defined parametrically.

This process is demonstrated by taking two head scans from the data base and using the skull of the first head scan as the subject skull. Figure 8 shows the faces of these two head scans. Figure 9 shows a tool used to perform the morph. The estimated face and the true face of the subject will be different. It can be argued that variation due to skeletal structure has been removed and that remaining differences are the result of non-uniform soft tissue structures. In Section 3.1 principle component analysis will be used to model these soft tissue or flesh depth variations.

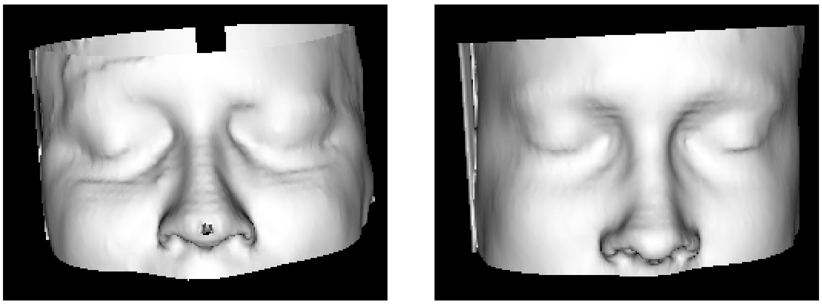

Figure 8: *The morphing process is demonstrated with two scanned heads taken from the CT data base. The left face is that of the subject. The right face will be morphed so as to approximate the left face.*

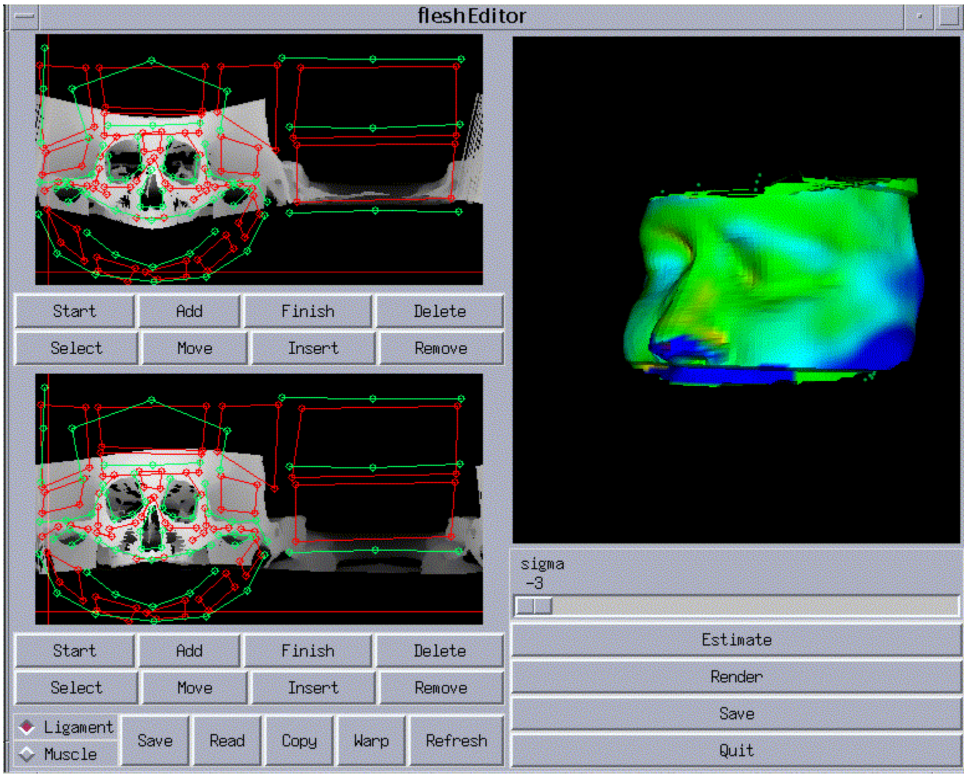

Figure 9: *This figure shows a tool used to transform the right face of figure 8 in to the left face. The left images are the cylindrical representations of the two skulls. The connected green points are used to control the 2D image morphing. The red lines show regions used to control the depth adjustment. The right figure shows the result of morphing. Color is used to depict the difference between the true subject face and the morphed face where green is small error, blue is positive error and red is negative error. As can be seen most of the error occurs where soft tissue is most prominent.*

Figure 10: *The right image shows the CT scan of the skull shown in the left image.*

Extraction of statistics

In the previous sections, we have described how the skull and face models generated from a CT scan may be aligned with a canonical skull model. Once this has been done for several data sets, it is possible to derive statistical data about the depth and variance of soft-tissue at any chosen point on the skull model. More generally it is possible to compute the correlation of tissue depths at different points on the model, and in fact at every pixel in the 2.5D cylindrical representation of the skull. By carrying out principal component analysis (to be described below) on the set of face and skull models it is possible to compute the principal modes of variation of the face or of tissue depth.

In doing this, one may choose any skull as the canonical skull model. This may be one of the skulls from the database, or even a general CAD model. In practice, we find that there are advantages to using a subject's skull (for instance, the one being considered in a specific forensic investigation) as the canonical skull model. This process will be considered in the next section.

3 Reconstruction from a subject skull

The sequence of steps that are used to process a subject skull and from it to reconstruct a face model will be described in this section. The first step in this process is to obtain a computer model of the skull. The recommended method is to use a CT scan, though other methods are possible, for instance laser scanning. Figure 10 shows a model generated from a CT scan of (a plaster cast of) the skull shown in figure 1.

Once a 3D model is obtained, each of the database skull models is aligned with it, and statistics are gathered. The steps of the complete process are as follows.

1. Scan the subject skull.

2. Carry out Marching Cubes algorithm to get a skull model.
3. As in the data-collection phase, select reference points to be used to align it with the head scan skulls from the database.
4. Create a cylindrical coordinate image model and pick fiducial points to be used for morphing the database head scans.
5. Morph each of the database skulls in turn to align them to the subject skull model, as described previously.
6. **PCA:** Carry out a Principal Components Analysis on the aligned *face* images to obtain an average face and the principal deformation modes of the face.
7. Modify the appearance of the face by varying the principal deformation modes using control-bar sliders.
8. **Post-processing:** Substitute features from a feature library to obtain the desired appearance.
9. Texture map the reconstructed face to give a more realistic appearance, add hair, etc.

Note that the average face and principal deformation modes of the face obtained by this method are computed after database models are morphed to the actual shape of the subject skull. Thus, the average face and its variants are in fact average faces **subject to** the restrictions imposed by the subject skull's shape.

The remaining details of this procedure will now be given.

3.1 Principle component analysis

In the previous section, a method for morphing a specific head scan onto the subject's skull was described. By repeating this process for all of the N heads in the CT database, a collection of estimates of the subject's face is generated. Each estimate is represented by an image structure, which can be viewed as a vector $\mathbf{X}$ of length L, where L is the number of pixels in the image. Each element of $\mathbf{X}$ is equal to a particular pixel intensity in a one to one fashion. The vector $\mathbf{X_i}$ corresponds to the estimate generated from the i^{th} head scan.

The idea of principal component analysis is to model a complex set of data in as simple a manner as possible. In this case the data is the set of subject face estimate vectors $\mathbf{X_i}$. One can compute the average data vector, and then express any particular data vector as the average vector, plus a deviation from the average. One wants to approximate any possible face vector in terms of an average, plus a linear combination of a small number of fixed representative (or basis) vectors representing deviations from the average. This represents a face space continuum. In general, it is not possible to model a complex data set precisely with a small number of basis vectors. However it is desirable to choose the basis vectors so that the expected error is as small as possible. In terms of linear algebra, one desires that the span of the basis vectors approximates the full data set as closely as possible. The solution to this problem is based on the eigenvector decomposition of a covariance matrix ([8]).

3.2 Deformation modes

The first step is to compute the mean face

$$\hat{\mathbf{X}} = \frac{1}{N}\sum_{1}^{N}\mathbf{X_i}.$$

Let the matrix

$$\mathbf{M} = [(\mathbf{X_1} - \hat{\mathbf{X}}), (\mathbf{X_2} - \hat{\mathbf{X}}), ...(\mathbf{X_N} - \hat{\mathbf{X}})].$$

The matrix $\mathbf{M}$ has L rows and N columns. The covariance matrix $\mathbf{C}$ is defined as

$$\mathbf{C} = \mathbf{M}\mathbf{M^T}$$

The eigenvectors $\mathbf{E}$ of $\mathbf{C}$ are a set of independent basis vectors and are defined by

$$\mathbf{CE} = \mathbf{E}\boldsymbol{\Lambda}$$

where $\boldsymbol{\Lambda}$ is a diagonal matrix of eigenvalues. In principle, $\mathbf{E}$ and $\boldsymbol{\Lambda}$ can be computed directly from $\mathbf{C}$. The matrix $\mathbf{C}$ is a an L by L matrix where L is the number of pixels in each image. Since L is very large, the required computational cost would be overwhelming. However, it was shown in [10] that the number of independent eigenvectors is limited to $N-1$ where N is the number of head scans in the database. This leads to the following formulation: compute $\mathbf{V}$ and $\boldsymbol{\Phi}$ such that

$$(\mathbf{M^T}\mathbf{M})\mathbf{V} = \mathbf{V}\boldsymbol{\Phi}.$$

This is a cheap computation since $\mathbf{M^T M}$ is a N by N matrix. Multiply both sides of the equation by $\mathbf{M}$ resulting in

$$\mathbf{CMV} = \mathbf{MV}\boldsymbol{\Phi}.$$

Therefore

$$\begin{aligned} \mathbf{E} &= \mathbf{MV} \\ \boldsymbol{\Lambda} &= \boldsymbol{\Phi} \end{aligned}$$

Given the mean face $\bar{\mathbf{X}}$ and the eigensystem $\mathbf{E}$ and $\boldsymbol{\Lambda}$, reconstruction based on a continuum of faces can now be performed.

3.3 Reconstruction

A particular reconstruction can be defined as the mean estimate $\bar{\mathbf{X}}$ plus a linear combination of the eigenvectors in $\mathbf{E}$. The eigenvalues define the dominance of each mode. Large eigenvalues are associated with large global deformations. Low eigenvalues correspond to specific aspects of a particular face. Since we are only interested in the overall shape of the face, only the major eigenvectors need be considered. Figures 11, 12 and 12 illustrate how the face space models are used to create an initial reconstruction. The user starts off with the average estimated face and then slider bars are used to define the contribution of the various deformation mode.

Once the operator is satisfied with the estimated face, a set of post-processing tools are used to make local alterations and to generate a texture map resulting is a lifelike reconstruction.

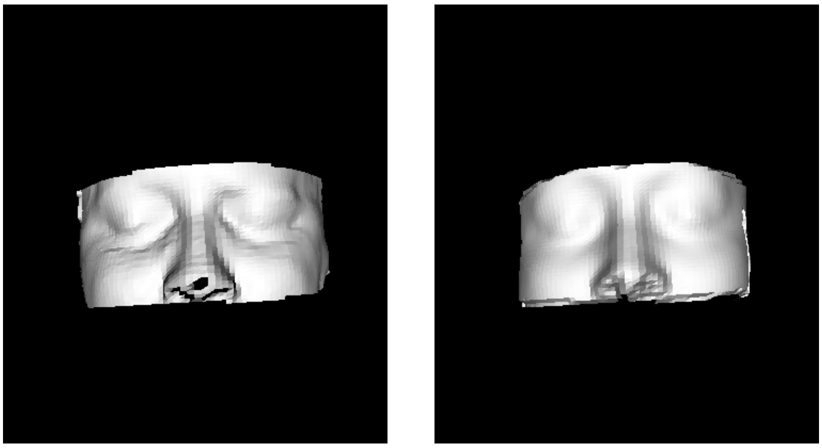

Figure 11: *The left image shows the true face of the subject. The right face shows the average morphed face*

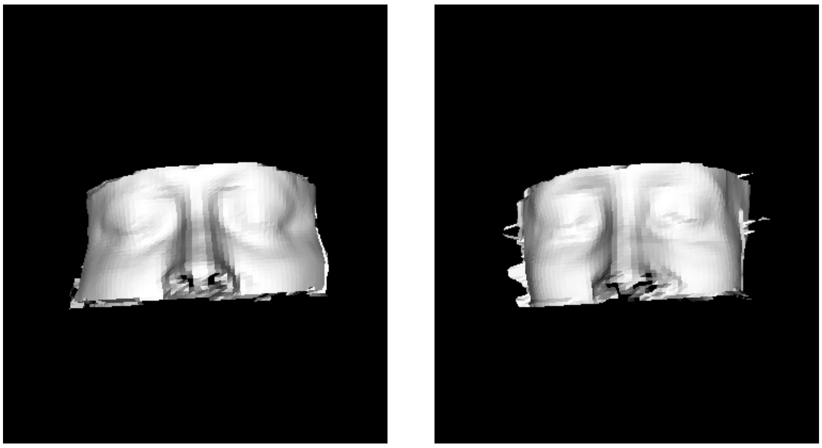

Figure 12: *These figures show the range of deformations associated with a single deformation mode. This deformation mode corresponds approximately to fatness/thinness of the face. By varying this mode, one obtains a spectrum of different facial types. Varying other principal modes will cause the shape of the face to vary in different ways.*

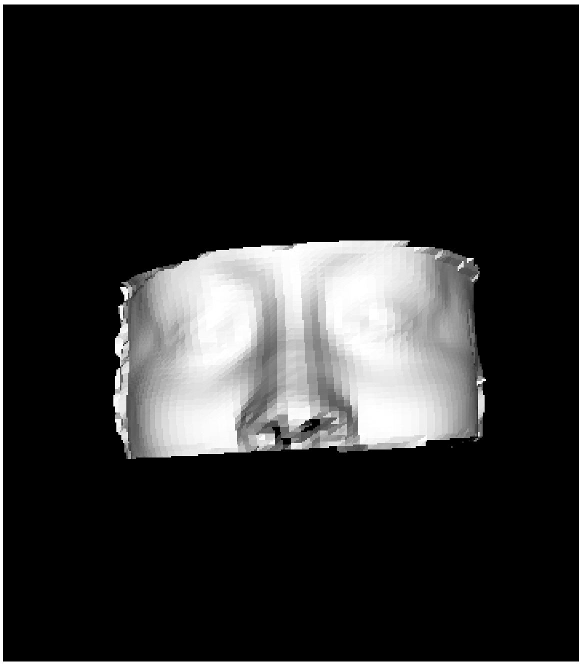

Figure 13: *This figure shows a particular reconstruction after adjusting various deformation modes.*

4 Post-processing

A reconstruction generated from the face-space continuum represents a linear combination of all the estimated faces generated from the CT head scans. The motivation of the PCA procedure was to generate an initial estimate of the basic shape of the subject's face. Based on prior evidence, heuristics or intuition, the operator may decide to change certain aspect of the face. A set of tools for this process has been provided.

4.1 Face editing

The face editing facility is based on a parts library. A collection of eye, nose and lip models have been extracted from 3D scans of various individuals. Figure 14 shows the part gathering process. By defining a few point correspondences, the part can be placed directly onto the model, replacing what was previously there. A blending operation is performed so that a smooth transition occurs between the model and the transplanted part. Many of these operations are considerably simplified by using the image representation of the face.

4.2 Texture mapping

The final geometric face model can be rendered as a polygonal mesh. The model will appear more lifelike if it is viewed with a real texture map. There are a number of devices that can be used to capture a texture map from an individual. These maps can

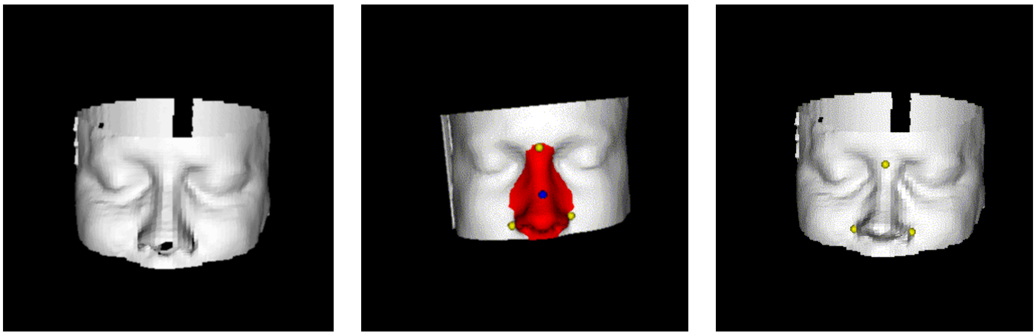

Figure 14: *This figure illustrates a face editing procedure. The left image is the original face. The middle image shows a part (the nose region shown in red) from a donor face. The right image shows the combined face. The spheres show the point correspondences. Merging algorithms are used to create a smooth border.*

Figure 15: *This figure illustrates a mechanism for transferring texture maps. The top left image is the donor face. The bottom left image is the subject face. The middle image is a texture map of the donor face. The right image is the subject face with the donor texture map.*

be tailored to the face model by specifying a few fiducial points on the source texture map and the geometric model. Figure 15 shows an example of this process. Since the source texture map comes from a particular individual, too much fine detail may be introduced. Another approach is to create an average texture map and use this to render the model. In this way specific features that can lead the eye are removed however the overall lifelike appearance still remains.

5 Discussion

In this paper we have described a system for creating a set of face estimates based on the skeletal remains of an unknown individual. These estimates are viewed as samples of a face space tailored to a specific skull. Principle component analysis is performed so that operators can explore this space in a convenient manner.

The CT data used to demonstrate our approach was derived from a set of sinus studies. The initial results have been encouraging. We are now in the process of collecting a set of complete head scans so that full reconstructions can be performed. The utility of this system will be determined by measuring human recognition rates on this new data.

It can be argued that a major obstacle in face reconstruction is the inability to predict the shape of features such as the nose and the eyes based of skeletal information. One path may be to take a more anatomical approach where muscle models are used to build up the face. An alternative solution is to partition the face space into regions that contain distinct facial features. Based on CT head scan data, it may be possible to correlate these regions with various measurable features derived directly from the skull. If successful, this would result in a major improvement in face reconstruction capability.

Acknowledgments

This work has been funded by a research grant from the Federal Bureau of Investigation. We would also like to thank Gay Malin and Chuck Fisher from the New York State Museum for their many contributions.

References

[1] Aulsebrook W.A., Iscan M.Y., Slabbert J.H., “Superimposition and reconstruction in forensic facial identification: a survey”, Forensic Science International, 75(1995) 101-120.

[2] Blanz V., Vetter T., “A Morphable model for the synthesis of 3D faces”, SIGGRAPH 1999, pp 187-194.

[3] Horn, B.K., “Closed-form solution of absolute orientation using unit quaternions”, J. Opt. Soc. Am. A, vol. 4, No. 4, pp 629-642. April 1987.

[4] Iscan M.Y., Helmer R.P., “Forensic analysis of the skull”, Wiley-Liss, 1993.

[5] Lorensen W.E., Cline H.E., “Marching cubes: a high resolution 3D surface construction algorithm”, Computer Graphics 21(3), pp. 163-169. July 1987.

[6] Nelson L.A., Michael S.D., “The application of volume deformation to three-dimensional facial reconstruction: A comparison with previous techniques”, Forensics Science International, 94(1998) 167-181.

[7] Phillips V.M., Smuts N.A., “Facial reconstruction: utilization of computerized tomography to measure facial tissue thickness in mixed racial population”, Forensic Science International, 83(1996) 51-59.

[8] Press W.H., Flannery B.P., Teukolsky S.A., Vetterling W.T., ‘Numerical Recipes in C the art of scientific computing’, University of Cambridge Press Syndicate, 1992.

[9] Shahrom A.W., Vanezis P., Chapman R.C., Gonzales A., Blenkinsop C., Rossi M.L., “Techniques in facial identification: computer-aided facial reconstruction using a laser scanner and video superimposition”, Int. Journal of Legal Medicine, 108, pp. 194-200, 1996.

[10] Turk M.A., Pentland A.P., “Face recognition using eigenfaces”, Proc. Computer Vision and Pattern Recognition, 1999, pp. 598-603.

[11] Vanezis P., Blowes R.W., Linney A.D., Tan A.C., Richards R., Neave., "Application of 3-D computer graphics for facial reconstruction and comparison with sculpting techniques", Forensic Science International 42(1989) 69-84.

[12] Wolberg G., "Digital Image Warping", IEEE Computer Society Press, 1988.

[13] Yoshino M., Matsuda H., Kubota S., Imaizumi K., Miyasaka S., Seta S., "Computer-assisted skull identification system using video superimposition", Forensic Science International, 90(1997) 231-244.