Reanimating the Dead: Reconstruction of Expressive Faces from Skull Data

Kolja Kähler[*]　　　Jörg Haber[†]　　　Hans-Peter Seidel[‡]

MPI Informatik, Saarbrücken, Germany

Abstract

Facial reconstruction for postmortem identification of humans from their skeletal remains is a challenging and fascinating part of forensic art. The former look of a face can be approximated by predicting and modeling the layers of tissue on the skull. This work is as of today carried out solely by physical sculpting with clay, where experienced artists invest up to hundreds of hours to craft a reconstructed face model. Remarkably, one of the most popular tissue reconstruction methods bears many resemblances with surface fitting techniques used in computer graphics, thus suggesting the possibility of a transfer of the manual approach to the computer. In this paper, we present a facial reconstruction approach that fits an anatomy-based virtual head model, incorporating skin and muscles, to a scanned skull using statistical data on skull / tissue relationships. The approach has many advantages over the traditional process: a reconstruction can be completed in about an hour from acquired skull data; also, variations such as a slender or a more obese build of the modeled individual are easily created. Last not least, by matching not only skin geometry but also virtual muscle layers, an animatable head model is generated that can be used to form facial expressions beyond the neutral face typically used in physical reconstructions.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation G.3 [Probability and Statistics]—Multivariate statistics G.1.2 [Numerical Analysis]: Approximation—Approximation of surfaces and contours

Keywords: facial modeling, forensic art, face reconstruction

1 Introduction

1.1 Background

For well over a hundred years, forensic art and science has been assisting law enforcement. One of the major areas of concern in this area is facial reconstruction for postmortem identification of humans from their physical remains. Manual reconstruction and identification techniques build on the tight shape relationships between the human skull and skin: for instance, the presumed identity

---

[*]e-mail: kkaehler@acm.org

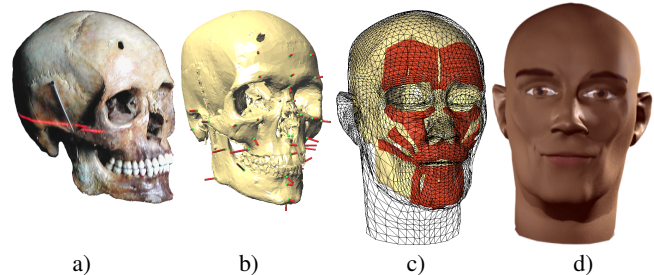[†]e-mail: haberj@acm.org

[‡]e-mail: hpseidel@mpi-sb.mpg.de

Figure 1: Reconstruction of a face from the skull: a) scanning the skull; b) skull mesh tagged with landmarks; c) skin mesh with muscles fitted to the skull; d) textured skin mesh, smiling expression.

of a murder victim can be confirmed by superimposing a facial photograph with a properly aligned and sized image of the skull. If no photograph is available, the look of the face can be reconstructed to a certain degree by modeling the missing tissue layers directly onto the skull or a plaster cast made from it.

The first documented case using three-dimensional facial reconstruction from the skull dates back to 1935 [Taylor 2001]. A key experiment was later performed by KROGMAN [1946]: given the body of a deceased person, he took a picture of the cadaver head before extracting the skull. The skull was provided to a sculptor along with information about sex, origin, and age of the late owner, plus data on the average tissue thicknesses at several positions in the face. From this material, a reconstruction sculpture was created that could be compared to the original head. Since that time, three-dimensional facial reconstruction from the skull has been much refined, but the method has essentially remained the same. Researchers have examined the skull / skin relationships for different ethnic groups [Lebedinskaya et al. 1993] and analyzed the correspondences of skull morphology and facial features [Fedosyutkin and Nainys 1993]. Others found correlations between muscle activity and skull shape [Moore and Lavelle 1974; Weijs and Hillen 1986]. In her comprehensive textbook, TAYLOR [2001] describes the craft in great detail.

Much of the fascination of the topic is due to the combined efforts of science and art, resulting in often astonishingly lifelike reconstructions, given the little available input (see Fig. 2). Many parameters of the outward appearance of an individual cannot be readily derived from the skull, though. The process is thus highly dependent on rules of thumb, the experience of the artist, and some guesswork. It is, for instance, next to impossible to reconstruct the shape of the ears based on scientific reasoning, although empirically there seems to be a relation of ear height to the length of the nose.

1.2 The Manual Reconstruction Process

The traditional work process for facial reconstruction begins with preparation of the skull. Since the skull is often evidence in a criminal case, great care needs to be taken in handling it: some parts are extremely thin and fragile, especially in the nose and the orbits. For identification, the teeth often provide a lot of useful informa-

Figure 2: Comparison of sculpted reconstructions with photographs. Left: male subject; right: female subject. (Images: Copyright ©[Helmer et al. 1993], reprinted by permission of Wiley-Liss, Inc., a subsidiary of John Wiley & Sons, Inc.)

tion, so a dental analysis is usually performed at this stage. For the reconstruction of the lower face, the mandible needs to be properly aligned and secured to the skull. In cooperation with an anthropologist, and possibly given more information from the remains of the victim, an estimation of age, ancestry, sex, and stature can now be obtained.

The actual face reconstruction proceeds with one of two available approaches: the *anatomical method* and the *tissue depth method*. The anatomical method attempts reconstruction by sculpting muscles, glands, and cartilage, fleshing out the skull layer by layer. This technique is more often used in the reconstruction of fossil faces, where no statistical population data exists [Zollikofer et al. 1998]. As TAYLOR states, this technique is very time consuming, occupying "many hundreds of hours". It also requires a great deal of detailed anatomical knowledge. Therefore, the alternative tissue depth method has become the more popular reconstruction technique in law enforcement. Here, standard sets of statistical tissue thickness measurements at specific points on the face are used. Each measurement describes the total distance from skin surface to the skull, including fat and muscle layers. The method is thus more rapid than the anatomical method and does not require as much anatomical knowledge. Such measurements have been collected for males and females of several racial groups, using needles, X-rays, or ultrasound techniques. The tissue depth data most often used by police artists today was collected primarily by RHINE *et al.* [Rhine and Campbell 1980; Rhine and Moore 1984]. The data is sorted into "slender", "normal", and "obese" groups, as well as by sex and race.

Given the set of measurements, tissue depth markers are now placed on the skull or a cast made from it, reflecting the tissue thickness at the sample points. These markers are oriented orthogonally to the skull surface, corresponding to the direction of the tissue thickness measurements. Using the markers and other features on the skull for guidance, the face is modeled on top of the skull using clay. A snapshot of the beginning stages of a reconstruction using the tissue depth method is shown in Fig. 3.

1.3   Our approach

Looking at the facial reconstruction process as described above from a computer graphics perspective, it essentially boils down to a surface interpolation problem. We thus implement the manual "dowel placement" method as an interactive procedure, obtaining position and distance constraints that define the relation between skin and skull at selected sample positions. The sculpting of the skin surface is mapped to a volume deformation applied to a head model template, satisfying these constraints. The deformation approach has the additional advantage of being applicable to additional structures attached to the template: in our system, we map a muscle structure to the fitted head model (see Fig. 1), enabling

animation on the reconstructed head in a physics-based facial animation framework.

The remainder of this paper is organized as follows: after reviewing related work in Section 2, we discuss acquisition of skull data and interactive landmark placement for setting up surface constraints in Section 3. Section 4 describes the structure of our generic head model and how it is fitted to the skull. Animation and texture generation for the resulting head model are touched upon in Section 5. We present examples in Section 6 and draw conclusions from our results in Section 7.

2   Previous and Related Work

2.1   Computer-Aided Face Reconstruction

Perhaps due to the lack of rigid taxonomies and hard rules, the use of computers and computer graphics in this forensic application is still very limited. The procedures described above cannot be cast easily into a computer program that produces good results in an automated manner—the experience and judgment of the practitioner remain a vital part of the system.

In law enforcement practice, computer-aided techniques restrict to relatively simple image and video manipulation: face photographs are used for skull superimposition [Grüner 1993; Miyasaka et al. 1995], while image warping and retouching enable a basic simulation of aging [Taylor 2001, p. 253]. This situation is unfortunate, since the traditional three-dimensional face reconstruction process is extremely time-consuming and expensive. It is hardly feasible to produce a variety of different plausible reconstructions from one skull, simply due to the effort that has to be put into the creation of each model. Also, repeated physical handling of the original skull increases the risk of damage.

One prototypical computer-based face reconstruction system, allowing fitting of a generic hierarchical B-spline head model to a skull mesh, is described by ARCHER in her Master's thesis [1997]. The user places dowels on a skull model with prescribed tissue thickness values, resulting in targets for a B-spline surface fitting process. The interpolation process is tricky and requires careful preparation of the template head model.

In the approach presented by MICHAEL and CHEN [1996], a source head model H_s that includes a skull S_s is deformed using a volume distortion function V such that the deformed source skull approximately matches the target skull S_t: $V(S_s) \approx S_t$. It is assumed that the deformed source head model $V(H_s)$ bears a good resemblance to the (unknown) target head model. The volume distortion function V is set up as a field warp using fourty pairs of disc fields, which are manually placed around the skull. No details are given about the placement of these control fields.
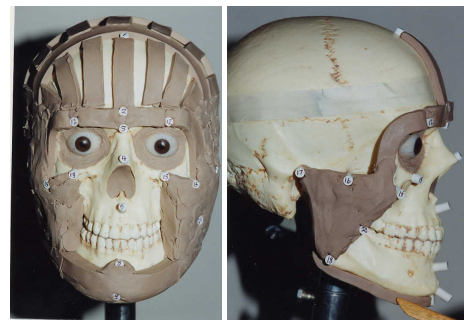

Figure 3: Modeling the face with clay on top of the skull using the tissue depth method. (Images [Taylor 2001], reprinted by permission.)

A deformation technique similar to the one used in our approach is employed by VANEZIS *et al.* [2000]. A facial template chosen from a database of scanned faces is deformed to match the position of target face landmarks, which have been derived from adding statistical tissue thickness values to the corresponding skull landmarks. The resulting reconstructed heads are not always complete (for instance, the top of the head is usually missing). The authors suggest to export an image of the reconstructed head and to apply a final image-processing step to add eyes, facial and head hair.

The above methods require a lot of manual assistance in setting up the interpolation function [Archer 1997; Michael and Chen 1996], or rely on a database of head templates [Vanezis et al. 2000]. In contrast, we develop reconstructions from one head template with relatively few markers, and use additional mechanisms to improve reconstruction results (see Section 4.3). Our approach always generates complete head models. Instead of using higher-order surfaces or point samples, the surface of our deformable head template is an arbitrary triangle mesh, simplifying later artistic modifications of the result using standard modeling tools. To the best of our knowledge, integration of expressive facial animation is not discussed by any other computer-aided facial reconstruction approach.

Other than explicit treatment of facial reconstruction, the creation of virtual head models based on human anatomy is well researched and documented in the computer graphics literature. Major developments in this area are discussed in the following section.

2.2 Human Head Modeling

A variety of techniques exists to create a face model from images or scan data. In the method presented by LEE *et al.* [1995], animatable head models are constructed semi-automatically from range scans. A generic face mesh with embedded muscle vectors is adapted to range scans of human heads. This process relies on the planar parameterization of the range scans as delivered, for instance, by the Cyberware digitizers. PIGHIN *et al.* [1998] interactively mark corresponding facial features in several photographs of an individual to deform a generic head model using radial basis functions. Animation is possible by capturing facial expressions in the process and blending between them. CARR *et al.* [2001] use radial basis functions to generate consistent meshes from incomplete scan data. Employing a large database of several hundred scanned faces, BLANZ *et al.* [1999] are able to create a geometric head model from only a single photograph. This model has the same resolution as the range scans in the database and cannot be readily animated. In the context of medical imaging, SZELISKI *et al.* [1996] minimize the distance between two surfaces obtained from volume scans of human heads by applying local free-form deformations [Sederberg and Parry 1986] and global polynomial deformations. The method does not require specification of corresponding features on the geometries.

Several facial animation systems use an approximation of the layered anatomical structure. WATERS [1987] represents skin and muscles as separate entities, where muscle vectors and radial functions derived from linear and sphincter muscles specify deformations on a skin mesh. In contrast to this purely geometric technique, physics-based approaches attempt to model the influence of muscle contraction onto the skin surface by approximating the biomechanical properties of skin. Typically, mass-spring or finite element networks are used for numerical simulation [Platt and Badler 1981; Lee et al. 1995; Koch et al. 1998]. From an initial triangle mesh, TERZOPOULOS and WATERS [1990] automatically construct a layered model of the human face. The model structure consists of three layers representing the muscle layer, dermis, and epidermis. The skull is approximated as an offset surface from the skin. Free-form deformations are employed by CHADWICK *et al.* [1989] to shape the skin in a multi-layer model, which contains bones, muscles, fat

tissue, and skin. SCHEEPERS *et al.* [1997] as well as WILHELMS and VAN GELDER [1997] introduce anatomy-based muscle models for animating humans and animals, focusing on the skeletal musculature. Skin tissue is represented only by an implicit surface with zero thickness [Wilhelms and Van Gelder 1997].

We build our system on the deformable, anatomy-based head model described by KÄHLER *et al.* [2002]. There, a generic face mesh with underlying muscle and bone layers is deformed to match scanned skin geometry. This process is adopted here to match the muscle and skin layers to given skull data instead.

3 Preparation of the Skull

Our approach uses three-dimensional skull data acquired, for instance, from volume scans and extraction of the bone layers, or by range scanning a physical skull. The test data used for the examples in Section 6 was acquired using both types of scans. To speed up processing, a triangle mesh of the skull model comprised of 50-250k polygons is produced by mesh decimation techniques [Garland and Heckbert 1997]. In general, the original data should be simplified as little as possible since minute details on the skull can give important clues for the reconstruction. The mesh resolution is chosen for adequate responsiveness of our interactive skull editor application. In practice, it is helpful to have the original data set (or the physical skull) ready as a reference during editing.

In the editor, the skull model is equipped with landmarks, as shown in Fig. 4. Points on the skull surface are simply picked to create a landmark, which can then be moved around on the surface for fine positioning. Each landmark is associated with a vector in surface normal direction, corresponding to the typical direction of thickness measurements. As can be seen on the right image in Fig. 4, some skull / skin correspondences are in fact non-orthogonal to the skull surface in the area of the lips. This is corrected for at a later step of the fitting process, as described in Section 4.3. The landmark vector is scaled to the local tissue thickness, which is looked up automatically by the landmark's assigned name in a table based on RHINE's data (see Section 1.2). The specific set of landmarks used in our system is listed in Appendix A.

4 Fitting the Deformable Head Model

4.1 Head Model Structure

When the skull is tagged with landmarks, it serves as the target for deformation of the generic head model shown in Fig. 5. Since the head model is used in a physics-based animation system, it does
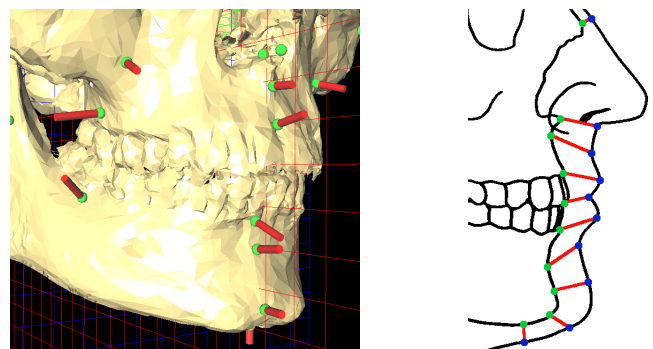

Figure 4: Skull landmark specification in the mouth area. Left: snapshot from our landmark editor; right: correspondences between skull and skin markers (Image after [y'Edynak and İşcan 1993])
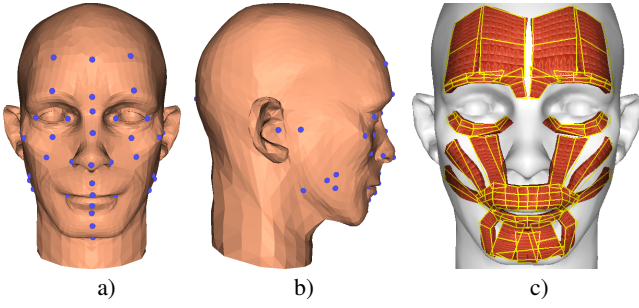
Figure 5: The deformable head model: a) head geometry with landmarks (blue dots), front view; b) side view; c) underlying muscles (red) created from layout grids (yellow).

not only consist of the visible outer geometry. The encapsulated structure includes:

the skin surface represented as a triangle mesh. The mesh resolution should be high enough to ensure good fitting results. Our template head mesh consists of 8164 triangles.

virtual muscles to control the animation. Each muscle is specified by a grid laid out on the skin, the actual muscle shape being computed automatically to fit underneath the skin surface. Each muscle consists of an array of fibers, which can contract in a linear or circular fashion. Our model includes 24 facial muscles responsible for facial expressions. Fig. 5(c) shows the muscle layout on the head template.

a mass-spring system connecting skin, muscles, and skull, built after the head model is fitted to the skull. For animation, muscles pull at spring nodes attached to their surface, in turn causing deformation of the spring mesh in the skin surface layer.

landmarks defined on the skin surface, as shown in Fig. 5(a) and (b). The majority of these landmarks corresponds to the landmarks interactively specified on the skull. These landmark pairs control the basic fitting of the head structure as described in Section 4.2. A few additional landmarks are only defined on the skin and are used for the final adjustments of the reconstructed shapes discussed in Section 4.3.

The head model is similar to the one in [Kähler et al. 2002], where detailed descriptions of the muscle model and animation approach can also be found.

4.2 Landmark-Based RBF Deformation

Given the deformable head model with n predefined skin landmark positions $\mathbf{p}_i \in \mathbb{R}^3$ and the corresponding landmarks $\mathbf{s}_i \in \mathbb{R}^3$ ($i = 1,\ldots,n$) specified on the skull, we set up a space deformation that fits the skin and the muscle layout to the skull.

The target skull landmarks have associated tissue depth vectors $\mathbf{d}_i$, so corresponding skin landmark positions $\mathbf{q}_i$ are defined as

$$\mathbf{q}_i = \mathbf{s}_i + \mathbf{d}_i.$$

The problem can now be treated as one of interpolation: we need to find a function $\mathbf{f}$ that maps the $\mathbf{p}_i$ to the $\mathbf{q}_i$:

$$\mathbf{q}_i = \mathbf{f}(\mathbf{p}_i), \quad i = 1,\ldots,n.$$

The unknown function $\mathbf{f}$ can be expressed by a radial basis function, i.e., a weighted linear combination of n basic functions ϕ_i and an additional explicit affine transformation:

$$\mathbf{f}(\mathbf{p}) = \sum_{i=1}^{n} \mathbf{c}_i \phi_i(\mathbf{p}) + \mathbf{R}\mathbf{p} + \mathbf{t}, \qquad (1)$$

where $\mathbf{p} \in \mathbb{R}^3$ is a point in the volume, $\mathbf{c}_i \in \mathbb{R}^3$ are (unknown) weights, $\mathbf{R} \in \mathbb{R}^{3 \times 3}$ adds rotation, skew, and scaling, and $\mathbf{t} \in \mathbb{R}^3$ is a translation component. The ϕ_i are defined by the source skin landmark points. According to BOOKSTEIN [1997], for deformation of biological solids an approach based on thin-plate splines is favorable. We thus use the simple biharmonic basic function $\phi_i(\mathbf{p}) := \|\mathbf{p} - \mathbf{p}_i\|_2$, which minimizes bending energy for the deformation [Duchon 1977].

To remove affine contributions from the weighted sum of the basic functions [Pighin et al. 1998; Carr et al. 2001], we include the additional constraints

$$\sum_{i=1}^{n} \mathbf{c}_i = \mathbf{0} \quad \text{and} \quad \sum_{i=1}^{n} \mathbf{c}_i^T \mathbf{p}_i = 0.$$

The resulting system of linear equations is solved for the unknowns $\mathbf{R}, \mathbf{t}$, and $\mathbf{c}_i$ using a standard LU decomposition with pivoting, to obtain the final warp function $\mathbf{f}$. This function can now be used according to Eq. (1) to transform a point $\mathbf{p}$ in the volume spanned by the landmarks. We apply $\mathbf{f}$ to the skin and muscle components of the generic model in the following ways:

- The *skin mesh* is deformed by direct application of the function to the vertices of the mesh.

- The *muscles* are transferred to the new geometry by warping their layout grid vertices, followed by recomputation of the shape to fit the deformed skin mesh.

Since our landmark set is comprised of only 40 landmarks (see Appendix A), the computed deformation doesn't properly align the skin to the skull in all places, as can be seen in Fig. 6(a). Interactive specification of more landmarks puts an undesirable additional burden onto the user, so additional landmark pairs are computed automatically by interpolation between existing ones on the upper and back part of the cranium, as well as on the mandible, as shown in Fig. 6(b). The thickness value of an interpolated skull landmark is also interpolated, where only such skull areas are chosen for landmark interpolation where the tissue thickness is near-constant. Tissue depth interpolation would be problematic, for instance, in the mid-face area, where thickness values change drastically from the cheekbone to the mid-face region below.

4.3 Additional Reconstruction Hints

The tissue depth values at the marker positions define the basic shape of the reconstructed head, assuming depth measurements being always strictly orthogonal to the skull surface. As mentioned in Section 3, this assumption is not always valid. A number of rules are thus used in traditional facial reconstruction to help locate certain features of the face based on the skull shape, employing empirical knowledge about shape relations between skin and skull [Taylor
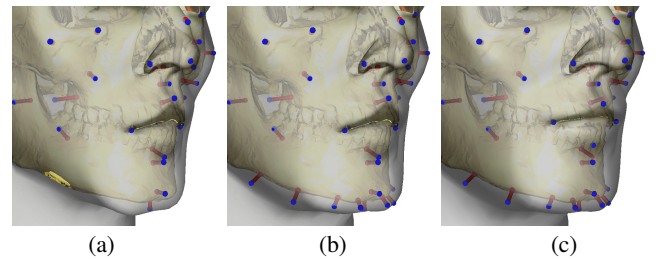

Figure 6: Fitting stages, shown on the lower face. a) Warp using only user-specified landmarks (some skull areas still intersecting the skin); b) with automatically interpolated landmarks on the mandible; c) using additional heuristics for lip and nose shaping.
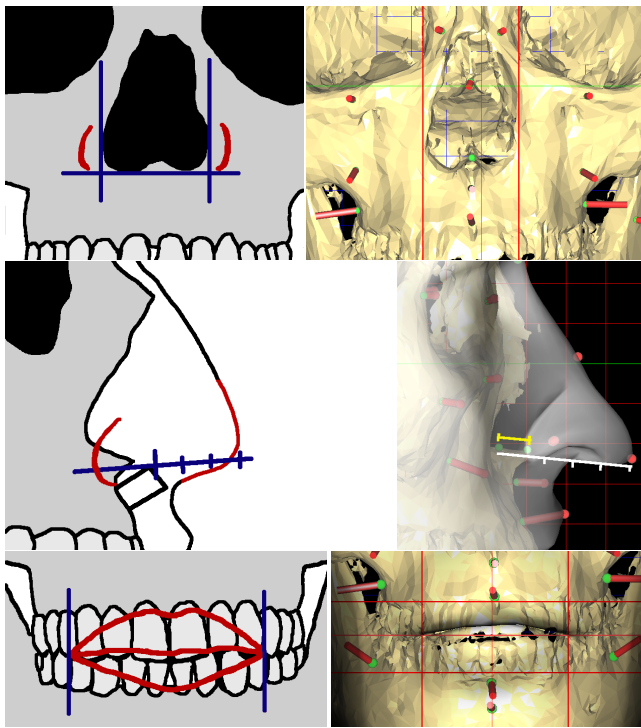
Figure 7: Comparison of heuristics used in traditional reconstruction (left) with our graphical interface (right). (Note: different skulls are used in the adjoining images.) Top: estimation of nose width; center: positioning of the nose tip; bottom: setting lip width, height, and mouth corner position.

2001]. We have translated some of these heuristics for use with the skull landmark editor: the final fitting result, as shown in Fig. 6(c), is obtained by including this additional user input.

To keep the user interface uniform, most rules are expressed by the placement of vertical and horizontal guides in a frontal view of the skull. From this user input, the placement of a few landmarks on the skin is adjusted, resulting in a new target landmark configuration. The updated landmark set is used to compute another warp function, which deforms the pre-fitted head model in the adjusted regions. Five rules influence the shape of the nose and the shape of the mouth, as shown in Fig. 7:

- The width of the nose wings corresponds to the width of the nasal aperture at its widest point, plus 5mm on either side in Caucasoids. In the editor, the user places two vertical guides to the left and right of the nasal aperture. From their position, the displacement of the two al[1] skin landmarks placed at the nose wings is computed (cf. Fig. 7, top row).

- The position of the nose tip depends on the shape of the anterior nasal spine. According to KROGMAN's formula [Taylor 2001, p. 443], the tip of the nose is in the extension of the nasal spine. Starting from the z value of the tissue depth marker directly below the nose (mid-philtrum, see Appendix A), the line is extended by three times the length of the nasal spine (cf. the white and yellow lines in the rightmost image of Fig. 7, middle row). In the editor, begin and end points of the nasal spine are marked. The prn landmark at the nose tip is then displaced according to the formula.

[1] see, e.g., [Farkas 1994] for a definition of standard facial landmarks

- The width of the mouth is determined by measuring the front six teeth, placing the mouth angles horizontally at the junction between the canine and the first premolar in a frontal view. Two vertical guides are used for positioning the ch landmarks located at the mouth angles (vertical lines in Fig. 7, bottom row).

- The thickness of the lips is determined by examining the upper and lower frontal teeth. Seen from the front, the transition between the lip and facial skin is placed at the transition between the enamel and the root part of the teeth. Two horizontal guides are placed by the user at the upper and lower transition, respectively. This determines the vertical position of the id and sd landmarks marking the lip boundary (top and bottom horizontal lines in Fig. 7, bottom row).

- The parting line between the lips is slightly above the blades of the incisors. This determines the vertical placement of the ch landmarks (middle horizontal line in Fig. 7, bottom row).

Using these heuristics, a better estimate of the mouth and nose shapes can be computed. The effect is strongest on the lip margins, since the assumption of an orthogonal connection between corresponding skin and skull landmarks is in fact not correct at these sites, as the right part of Fig. 4 shows. The initial deformation thus gives a good estimate of the tissue thickness of the lips while the second deformation using the information provided by interactive guide adjustment refines the vertical placement of the lip margins.

5 Facial Expressions and Rendering

In manual facial reconstruction, a neutral pose of the face is preferred as the most "generic" facial expression. Other expressions could be helpful for identification purposes, but the cost of modeling separate versions of the head model is prohibitive. In our virtual reconstruction approach, this does not pose a problem. Since the fitted head model has the animatable structure of skin and muscles, different facial expressions can be assumed by setting muscle contractions, as in other physics-based facial animation systems [Kähler et al. 2001; Lee et al. 1995]. Fig. 8 shows how muscles are used to form different facial expressions.

For a completely animatable head model, it is necessary to include a separately controllable mandible, a tongue, rotatable eyeballs, and eye lids into the head model. We have decidedly left them out of the reconstruction approach since these features are not particularly useful in this application: while a modest change of expression such as a smile or a frown might aid identification, rolling of eyes, blinking, and talking would probably not. It is also nearly impossible to correctly guess details such as a specific way of speaking—errors in this respect would produce rather misleading results in a real identification case. The effort of placing tongue, eye, and potentially teeth models thus does not offset the benefits.

Figure 8: Expressions on the generic head model and the corresponding muscle configurations.
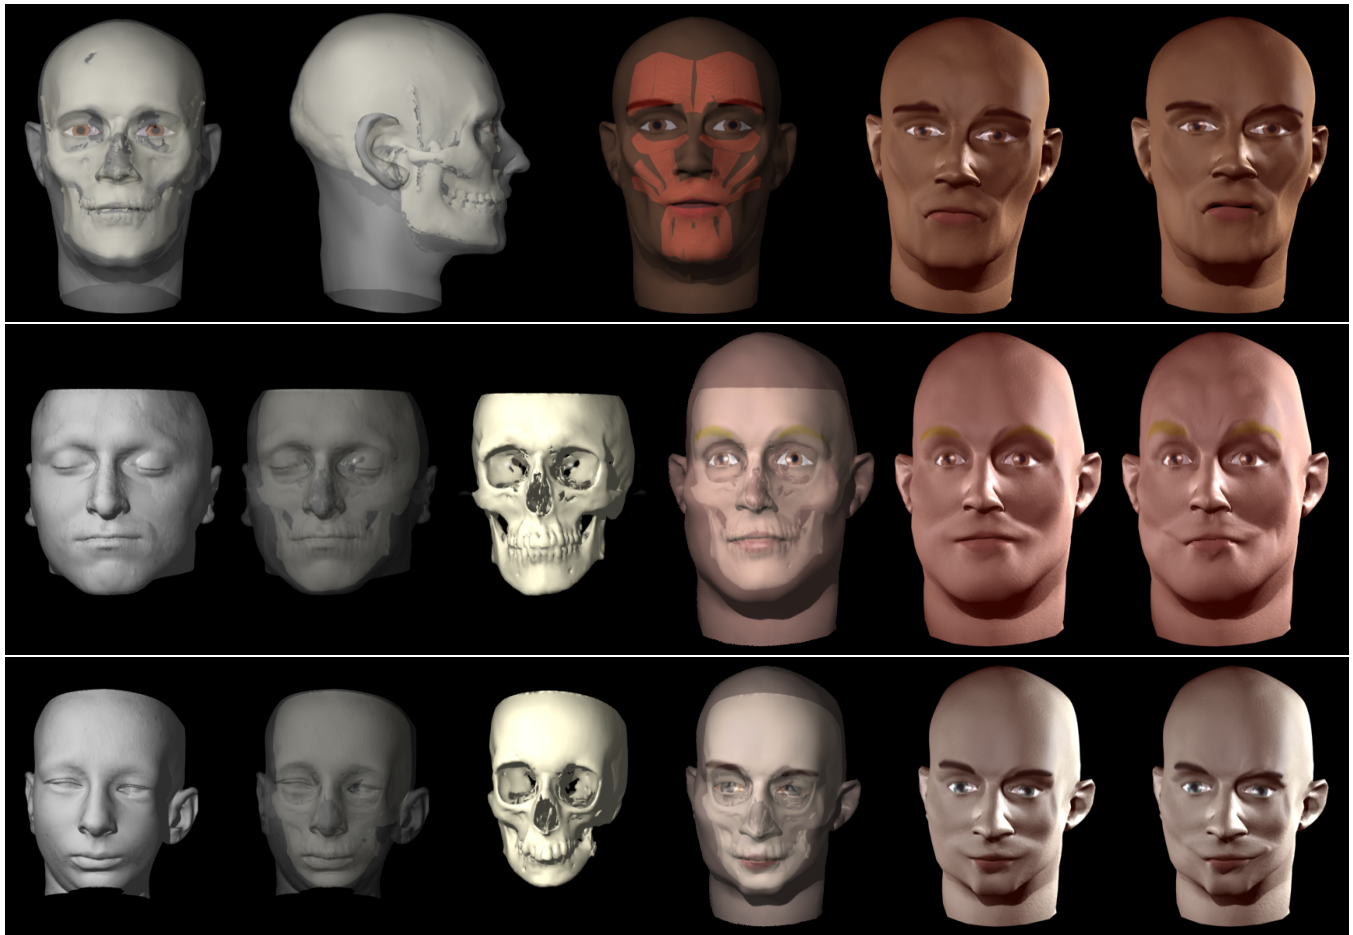
Figure 9: Examples of facial reconstructions created with our system. Top: model created from a scanned real skull, showing fit of skin to skull, transferred muscles, and two facial expressions. Middle: Reconstruction from a volume scan of a male, showing the actual face as contained in the data, superimpositions of the actual and the reconstructed face with the skull, and the reconstruction with neutral and "worried" expression. Bottom: Reconstruction from volume scan of a female with strong skull deformations. The CT data sets don't contain the top and bottom of the heads, thus the source skull and face models are cut off. The actual head height had to be guessed in these cases.

If additional information about the modeled person is available, for instance, from remnants of hair found with the skull, the resulting mesh can be colored correspondingly. Our system includes basic capabilities for coloring the parts associated with skin, lip, and eyebrows in the model's texture map. Colors can be adjusted interactively in HSV space on the reconstructed head model. Finally, the color adjustments are merged into a neutral base texture and saved as a new texture map. The fitted, texture-mapped triangle mesh can be easily imported into various rendering packages for display. The examples shown in Fig. 9 show three different skin colorations created in this way.

6 Results

We have tested our technique on a real skull that was made available to us by a forensic institute and on two medical volume scans. All data pertains to individuals of Caucasian type. Each reconstruction required approximately an hour of interactive work, excluding time for data acquisition.

The real skull, depicted on the first page of this paper, was unearthed on a construction site and belongs to an unidentified male, approximately 35 years of age. As can be seen from the hole in the frontal bone, he was killed by a head shot—the owner of this skull

probably was a war victim or a soldier. After scanning the skull, the resulting mesh was simplified to 100k triangles. Interactive placement of skull landmarks and facial feature guides was relatively easy in this case since the skull is complete and in good condition. Due to its war-time origin, we assumed the face to be rather skinny, so we selected the "slender" tissue thickness table. Fitting results can be seen in Fig. 9, top row. Since the actual appearance of the individual is unknown, the accuracy of the reconstruction can only be guessed. Nonetheless, our reconstruction seems plausible. Notably, the shape of the chin, which can be predicted from the corresponding region on the skull, has been reproduced well.

To show examples utilizing other data sources, and also for validation, we extracted skull and skin surfaces from medical volume scans. The first data set, shown in the middle row of Fig. 9, pertains to a male subject of roughly 30 years. The subject's face is rather bulky, so we chose the "obese" tissue thickness data set (in a real case, this choice would have to be made based on other available information such as the size of clothes, if present). Our first reconstruction attempts showed a consistent emphasis on prominent cheek bones and hollow cheeks: no matter which data set we picked, the face would become more bulky, but not show the expected general roundness of the face. This effect is demonstrated in Fig. 10 on variations of our first model. A closer examination
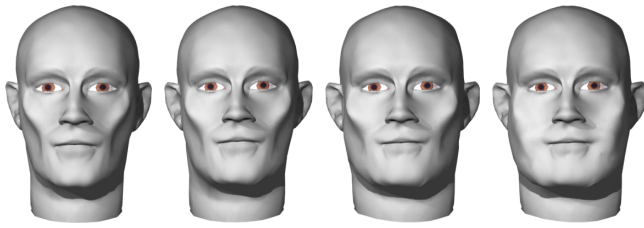
559

Figure 10: Left to right: RHINE's traditional "slender", "average", and "obese" tissue depth tables (cf. [Taylor 2001, p. 350 ff.]) often result in hollow cheeks and prominent cheekbones (see also Fig. 9). Rightmost image: the shape can be improved by "bulging out" the affected mesh areas.

revealed that the reason lies in the relatively low thickness values RHINE assigned to the landmarks defining the cheek region (*sbm2* and *spm2* in Table 1). After excluding these two landmarks, we obtained the results shown in Fig. 9. The rightmost image in Fig. 10 shows how simple mesh modeling techniques could be used at this point to improve and individualize the reconstruction.

The second volume data set shows a female patient with strong skull deformations. We produced a reconstruction of this face to test the method with a decidedly non-average skull shape. The result can be seen in the bottom row of Fig. 9. Since our automatic landmark interpolation scheme (see Section 4.2) is designed to handle the normal range of skull variations, the unusual shape of the mandible resulted in very sparse sampling of the chin area. Another prominent feature of the skull data is the protrusion of one incisor, pushing the upper lip to the front. We modeled this effect by moving the *sd* landmark a few millimeters down onto the blade of the incisor, thus pushing the associated skin landmark forward as well. This did not impair the positioning of the upper lip boundary since this is adjusted separately by the mouth guides (cf. Fig. 7).

7 Conclusion and Future Work

The face reconstruction approach presented in this paper mirrors the manual tissue depth method and thus has essentially the same prediction power. Our results show overall good reproduction of facial shape and proportions, and some surprisingly well-matched details. It should be noted that our examples were produced by computer scientists with no training in forensic reconstruction.

The advantages of the computerized solution are evident: instead of weeks, it takes less than a day to create a reconstructed face model, including scanning of the skull. Once the scan data is marked with landmarks, different varieties such as slimmer or more obese versions can be produced within seconds at the push of a button, which is practically impossible with the manual method due to the vast amount of time needed for production of a single model. Slight variations in facial expression can also be obtained quite easily by animating the muscle structure underlying the model.

Since the virtual reconstruction is based on 3D scans, which can be acquired contact-free, the risk of damage to the original skull is reduced. On the other hand, the scanning process has inherent limitations: depending on the maximum resolution of the digital scanner, much of the finer detail on the skull is lost. The delicate structure of, for instance, the nasal spine cannot be fully captured with current scanning technology. For this reason, it is necessary to consult the original skull from time to time for reference.

In our experiments, we often found that surface normals on the scanned skull geometry do not always behave the way they should, reflecting the orientation of the surface only very locally. It might be useful to consider an average of normals in a larger area around

the landmark position to solve this. Sometimes, it would be desirable to adjust the orientation manually.

The interactive system allows for an iterative reconstruction approach: a model is produced quickly from a given landmark configuration, so landmarks can be edited repeatedly until the desired result is obtained. The emphasis on the interaction component makes the speed of the fitting process an important issue. While the actual calculation of the warp function and the deformation of the mesh are performed instantaneously, about five seconds are needed in our test setting on a 1.7 GHz Pentium Xeon to examine skull and skin for potential insertion of additional landmarks. This time is for the largest part used for ray intersections of the skull and skin meshes, which are done in a brute force manner. We expect a big speed-up through the use of space partitioning techniques.

For practical use, the facial reconstruction system should provide more editing facilities for skin details and hair. Useful additions include, for instance, a choice of templates for haircuts and facial features such as eyebrow shapes, beards, and wrinkles. At this point, large-scale validation of the system would be necessary to evaluate the usability of the system.

As TAYLOR writes in her book, the tissue depth values should not be taken at face value in three-dimensional facial reconstruction, but rather act as guides for the final facial reconstruction, which still relies heavily on artistic skills and intuition. Our tests confirm that strict adherence to RHINE's data for the solution of the interpolation problem is too limiting. This indicates not a weakness in our method, but reflects the low number of samples (between 3 and 37 in each group) and the technical limitations at the time RHINE assembled his data tables. Given the current state of technology, more samples of higher precision could be acquired, resulting in much more comprehensive and usable data. Ultimately, computer-based facial reconstruction could then even become superior to the traditional approach.

8 Acknowledgements

References

ARCHER, K. M. 1997. *Craniofacial Reconstruction using hierarchical B-Spline Interpolation*. Master's thesis, University of British Columbia, Department of Electrical and Computer Engineering.

BLANZ, V., AND VETTER, T. 1999. A Morphable Model for the Synthesis of 3D Faces. In *Proc. ACM SIGGRAPH 1999*, ACM Press / ACM SIGGRAPH, Computer Graphics Proceedings, Annual Conference Series, 187–194.

BOOKSTEIN, F. L. 1997. *Morphometric Tools for Landmark Data*. Cambridge University Press.

CARR, J. C., BEATSON, R. K., CHERRIE, J. B., MITCHELL, T. J., FRIGHT, W. R., MCCALLUM, B. C., AND EVANS, T. R. 2001. Reconstruction and Representation of 3D Objects With Radial Basis Functions. ACM Press / ACM SIGGRAPH, Computer Graphics Proceedings, Annual Conference Series, 67–76.

CHADWICK, J. E., HAUMANN, D. R., AND PARENT, R. E. 1989. Layered Construction for Deformable Animated Characters. In *Computer Graphics (Proc. ACM SIGGRAPH 89)*, 243–252.

DUCHON, J. 1977. Spline minimizing rotation-invariant semi-norms in Sobolev spaces. In *Constructive Theory of Functions of Several Variables*, W. Schempp and K. Zeller, Eds., vol. 571 of *Lecture Notes in Mathematics*, 85–100.

FARKAS, L. G., Ed. 1994. *Anthropometry of the Head and Face*, 2nd ed. Raven Press.

FEDOSYUTKIN, B. A., AND NAINYS, J. V. 1993. *Forensic Analysis of the Skull*. Wiley-Liss, ch. 15: The Relationship of Skull Morphology to Facial Features, 199–213.

GARLAND, M., AND HECKBERT, P. S. 1997. Surface simplification using quadric error metrics. In *SIGGRAPH 97 Conference Proceedings*, 209–216.

GRÜNER, O. 1993. *Forensic Analysis of the Skull*. Wiley-Liss, ch. 3: Identification of Skulls: A Historical Review and Practical Applications.

HELMER, R. P., RÖHRICHT, S., PETERSEN, D., AND MÖHR, F. 1993. *Forensic Analysis of the Skull*. Wiley-Liss, ch. 17: Assessment of the Reliability of Facial Reconstruction, 229–246.

KÄHLER, K., HABER, J., AND SEIDEL, H.-P. 2001. Geometry-based Muscle Modeling for Facial Animation. In *Proc. Graphics Interface 2001*, 37–46.

KÄHLER, K., HABER, J., YAMAUCHI, H., AND SEIDEL, H.-P. 2002. Head shop: Generating animated head models with anatomical structure. In *ACM SIGGRAPH Symposium on Computer Animation*, ACM SIGGRAPH, 55–64.

KOCH, R. M., GROSS, M. H., AND BOSSHARD, A. A. 1998. Emotion Editing using Finite Elements. In *Computer Graphics Forum (Proc. Eurographics '98)*, vol. 17, C295–C302.

KROGMAN, W. M. 1946. The reconstruction of the living head from the skull. *FBI Law Enforcement Bulletin* (July).

LEBEDINSKAYA, G. V., BALUEVA, T. S., AND VESELOVSKAYA, E. V. 1993. *Forensic Analysis of the Skull*. Wiley-Liss, ch. 14: Principles of Facial Reconstruction, 183–198.

LEE, Y., TERZOPOULOS, D., AND WATERS, K. 1995. Realistic Modeling for Facial Animations. In *Proc. ACM SIGGRAPH 1995*, ACM Press / ACM SIGGRAPH, Computer Graphics Proceedings, Annual Conference Series, 55–62.

MICHAEL, S., AND CHEN, M. 1996. The 3D reconstruction of facial features using volume distortion. In *Proc. 14th Eurographics UK Conference*, 297–305.

MIYASAKA, S., YOSHINO, M., IMAIZUMI, K., AND SETA, S. 1995. The computer-aided facial reconstruction system. *Forensic Science Int. 74*, 1-2, 155–165.

MOORE, W. J., AND LAVELLE, C. L. B. 1974. *Growth of the Facial Skeleton in the Hominoidea*. Academic Press, London.

PIGHIN, F., HECKER, J., LISCHINSKI, D., SZELISKI, R., AND SALESIN, D. H. 1998. Synthesizing Realistic Facial Expressions from Photographs. In *Proc. ACM SIGGRAPH 1998*, ACM Press / ACM SIGGRAPH, Computer Graphics Proceedings, Annual Conference Series, 75–84.

PLATT, S. M., AND BADLER, N. I. 1981. Animating Facial Expressions. In *Computer Graphics (Proc. ACM SIGGRAPH 81)*, 245–252.

RHINE, J. S., AND CAMPBELL, H. R. 1980. Thickness of facial tissues in American blacks. *Journal of Forensic Sciences 25*, 4, 847–858.

RHINE, J. S., AND MOORE, C. E. 1984. Tables of facial tissue thickness of American Caucasoids in forensic anthropology. *Maxwell Museum Technical Series 1*.

SCHEEPERS, F., PARENT, R. E., CARLSON, W. E., AND MAY, S. F. 1997. Anatomy-Based Modeling of the Human Musculature. In *Proc. ACM SIGGRAPH 1997*, ACM Press / ACM SIGGRAPH, Computer Graphics Proceedings, Annual Conference Series, 163–172.

SEDERBERG, T. W., AND PARRY, S. R. 1986. Free-Form Deformation of Solid Geometric Models. *Computer Graphics (Proc. ACM SIGGRAPH 86) 20*, 4 (Aug.), 151–160.

SZELISKI, R., AND LAVALLÉE, S. 1996. Matching 3-D Anatomical Surfaces with Non-Rigid Deformations using Octree-Splines. *International Journal of Computer Vision 18*, 2, 171–186.

TAYLOR, K. T. 2001. *Forensic Art and Illustration*. CRC Press LLC.

TERZOPOULOS, D., AND WATERS, K. 1990. Physically-based Facial Modelling, Analysis, and Animation. *Journal of Visualization and Computer Animation 1*, 2 (Dec.), 73–80.

VANEZIS, P., VANEZIS, M., MCCOMBE, G., AND NIBLETT, T. 2000. Facial reconstruction using 3-D computer graphics. *Forensic Science Int. 108*, 2, 81–95.

WATERS, K. 1987. A Muscle Model for Animating Three-Dimensional Facial Expression. In *Computer Graphics (Proc. ACM SIGGRAPH 87)*, 17–24.

WEIJS, W. A., AND HILLEN, B. 1986. Correlations between the cross-sectional area of the jaw muscles and craniofacial size and shape. *Am. J. Phys. Anthropol. 70*, 423–431.

WILHELMS, J., AND VAN GELDER, A. 1997. Anatomically Based Modeling. In *Proc. ACM SIGGRAPH 1997*, ACM Press / ACM SIGGRAPH, Computer Graphics Proceedings, Annual Conference Series, 173–180.

Y'EDYNAK, G. J., AND İŞCAN, M. Y. 1993. *Forensic Analysis of the Skull*. Wiley-Liss, ch. 16: Anatomical and Artistic Guidelines for Forensic Facial Reconstruction, 215–227.

ZOLLIKOFER, C. P. E., PONCE DE LEÓN, M. S., AND MARTIN, R. D. 1998. Computer-Assisted Paleoanthropology. *Evolutionary Anthropology 6*, 2, 41–54.

A   Landmark Set used for Reconstruction

Table 1 lists the paired landmarks on skin and skull that are used for the facial reconstruction approach described in this paper. Most skull landmark names and descriptions are taken from [Taylor 2001, page 350 ff.]. Short skull landmark names are listed in the *id* column. We have tried to adhere to naming conventions used in the forensic and anthropometric literature as much as possible [Taylor 2001; y'Edynak and İşcan 1993; Farkas 1994]. For simplicity, corresponding landmarks on skull and skin have the same short name in our system, which is not generally the case in the literature. In a few cases, marked by ∗ in the table, we invented short names. Not all skull landmarks have an "official" counterpart on the skin, so we placed the corresponding skin markers using our own judgment. The *mp* landmark pair is not part of the standard set. We added it to improve the alignment of skin to skull in the region behind the ears, where the mastoid process adds a bulge to the skull.

name	id	description
		Midline
Supraglabella	tr	Above glabella, identified with the hairline
Glabella	g	The most prominent point between the supraorbital ridges in the midsagittal plane
Nasion	n	The midpoint of the suture between the frontal and the two nasal bones
End of nasals	na	The anterior tip or the farthest point out on the nasal bones
Mid-philtrum	a	The mid line of the maxilla (east and west), placed as high as possible before the curvature of the anterior nasal spine begins
Upper lip margin (Supradentale)	sd	Centered between the maxillary (upper) central incisors at the level of the Cementum Enamel Junction (CEJ)
Lower lip margin (Infradentale)	id	Centered between the mandibula (lower) central incisors at the level of the Cementum Enamel Junction (CEJ)
Chin-lip fold (Supramentale)	b	The deepest mid line point of indentation on the mandible between the teeth and the chin protrusion
Mental eminence (Pogonion)	pog	The most anterior or projecting point in the mid line on the chin
Beneath chin (Menton)	me	The lowest point on the mandible
		Bilateral
Frontal eminence	fe∗	Place on the projections at both sides of the forehead
Supraorbital	sci	Above the orbit, centered on the upper most margin or border
Suborbital	or	Below the orbit, centered on the lower most margin or border
Endocanthion	en	point at the inner commissure of the eye fissure; the landmark on the skin is slightly lateral to the one on the bone
Exocanthion	ex	point at the outer commissure of the eye fissure; the landmark on the skin is slightly medial to the one on the bone
Inferior malar	im	The lower portion of the maxilla, still on the cheekbone
Lateral orbit	lo	Drop a line from the outer margin of the orbit and place the marker about 10 mm below the orbit
Zygomatic arch, midway	zy	Halfway along the zygomatic arch (generally the most projecting point on the arch when viewed from above)
Supraglenoid	sg	above and slightly forward of the external auditory meatus
Gonion	go	The most lateral point on the mandibular angle
Supra M^2	spm2∗	Above the second maxillary molar
Occlusal line	ol	On the mandible in alignment with the line where the teeth occlude or bite
Sub M$_2$	sbm2∗	Below the second mandibular molar
Mastoid process	mp∗	Most lateral part on the mastoid process behind and below the ear canal

Table 1: Landmark set used for face reconstruction.