

# A novel method of automated skull registration for forensic facial approximation

## Abstract

Modern forensic facial reconstruction techniques are based on an understanding of skeletal variation and tissue depths. These techniques rely upon a skilled practitioner interpreting limited data. To (i) increase the amount of data available and (ii) lessen the subjective interpretation, we use medical imaging and statistical techniques. We introduce a software tool, (*anonymized*) for computer-based forensic facial reconstruction. The tool applies innovative computer-based techniques to a database of human head Computed Tomography (CT) scans in order to derive a statistical approximation of the soft tissue structure of a questioned skull. A core component of this tool is an algorithm for removing the variation in facial structure due to skeletal variation. This method uses models derived from the CT scans and does not require manual measurement or placement of landmarks. It does not require tissue depth tables, can be tailored to specific racial categories by adding CT scans, and removes much of the subjectivity of manual reconstructions.

## 1 Introduction

Forensic facial reconstruction can be segmented into three main branches based on the amount of information available about the soft tissue. Consider the taxonomy of (*anonymized*) [1] in Figure 1. The leftmost branch (known face) represents those instances where sufficient soft tissue remains such that the appearance of the individual can be restored. In these cases, the reconstructionist patches the existing soft tissue and a full reconstruction is not necessary [2, 3]. The middle branch (possible face) represents those cases where there is insufficient soft tissue to restore the individuals appearance, but where other clues can limit the identity to a few candidates. For these cases superimposition of photographs onto the questioned skull may be preferred to full reconstruction [4]. The rightmost branch (questioned face) represents the cases of interest to this paper. In these cases, there is insufficient soft tissue to restore the individuals appearance and superimposition has either failed or is impractical. This branch of the taxonomy is further segmented based on dimensionality (two- or three- dimensions), and on methodology such as tissue-depth or anatomical methods [3, 5, 6, 7]. Blocks below the dotted line represent computer-based or computer-enhanced facial techniques [8].

Facial approximation methods assume that an individual’s appearance is influenced by the structure of: (i) the underlying skull, (ii) the facial soft tissue depth, and (iii) the facial features such as the nose, ear, lips, and eyes. Variations in (i), (ii), and (iii) create different faces and allow individuals to be distinguished and recognized.

Facial approximations start with the questioned skull as a template. Facial structure due to the underlying skull, item (i), can thus be explicitly considered in the final approximation. Estimating facial structure due to facial soft tissue depth, item (ii), and facial features, item (iii), is a tougher problem and the methods generally gather soft tissue data from a population in order to predict the face of the questioned individual. This amounts to a statistical argument from the general (population) to the specific (individual). Accuracy and validity of the approximation are dependent on the accuracy of the underlying model, accuracy of the measurements, and rank of the sample population. Take as an example, the sparse tissue-depth methods. The inference model is the landmark positions and the method for placing clay on the skull. The tissue depth tables are the data, and the accuracy is affected by both positioning and measurement accuracy. The number of individuals sampled to create each division in the table comprises the rank. The greater the number of points gathered, the denser and more accurately the points are placed, and the larger the population, the more confidence can be placed in the reconstruction.

The method we describe is a dense computational approach and falls in the “dense tissue box” of Figure 1. These methods can be seen as logical extensions to the sparse tissue-depth methods. Instead of the sparse landmark and clay inference model these methods use dense placement of soft-tissue depth. The tissue-depth table data is replaced by Computed Tomography (CT) scan data, and the population size is determined by the number of CT scans used in the reconstruction. In some sense, these methods are analagous to taking a rubber mask of the soft tissue and stretching it over the questioned skull. By using more points and eliminating the subjective placement of clay, we are enhancing the accuracy of sparse tissue-depth approaches. Further, by using a CT database, we can increase the confidence in our reconstruction by adding CT scans.

## 2 Materials and Methods

The dense placement model requires that the soft tissue from a known skull or skulls be placed on the questioned skull. Conceptually, this means that each point on the questioned skull is associated with a corresponding point on the known skull. The soft tissue overlying the point on the known skull is then associated with the point on the questioned skull. This can be a direct implementation, *i.e.* find dense point-to-point mappings from the questioned skull to the known skull [9], or more frequently can be implemented by generating a deformable transform or “warping” [10, 11, 12]. The warping process changes the shape of a known skull so that it has the same shape as the questioned skull. Forcing the known soft tissue to follow the same warping as the known

skull gives the desired shape of the soft tissue without requiring a point-to-point mapping.

The idea of warping is central to both existing computer-based implementations and to our system and it is important to contrast warping with rigid transformations that simply rotate and slide two items. Consider Figure 2. The figure has a line segment (Source) that is matched to the curve (Target) using both a warping and a rigid transformation. Both methods start by identifying a set of point-to-point mappings or correspondences. The points on the source are then mapped to the target using two different transformations, a rigid transformation (upper example) and a warping (lower example). Note that the rigid transformation preserves the Source while the warping better approximates the Target. Conversely, rigid transformations are simpler to calculate than deformations because the points on the source line cannot move independently. Thus rigid transformations need only specify a new position and orientation for the Source as a complete unit. The warping needs to specify a new position for each defined point on the source and must further specify the path connecting any two points. Both rigid transformations and warpings are useful in many applications and the appropriate method should be selected based on the application. Skulls have complex shapes and matching a questioned skull to a known skull requires warpings.

Many of the computer-based approaches rely on similar theories. Each method requires a model for representing the skull and face data. The two main methods either work directly on the “volumetric” or CT images *e.g.* [9, 10, 12], or work from polygonal models, *e.g.* [11], derived from the CT data or obtained directly from a range sensor. The volumetric data is comprised of individual volume elements or voxels, three dimensional analogs to pixels. Each voxel is a cube of space having a single value and representing a portion of a 3-dimensional image. The union of all the voxels create a volume, or volumetric data set. The volumetric data typically represents the skull and skin implicitly as areas in the data having specific values in the 3-dimensional image. The interface between skin and air, skin and bone, etc. are determined as areas in the 3D image where the value changes sufficiently to break the next threshold. The area of constant value where the jump occurs is called an isosurface. Polygonal models make the surfaces explicit by representing these isosurfaces as the union of small, discrete triangles. Figure 3 illustrates the concept of a polygonal model showing the skull as it appears in our viewer, along with a magnified view showing a wire-frame model (triangular patch boundaries) of a portion of the mandible with teeth.

Given a representation of the skull and soft tissue, a set of correspondences are determined to guide the warping. These correspondences may either be discovered by the algorithm (automatic), or are input by the operator (manual). A warping that maps the known skull onto the questioned skull is generated from the correspondences, and this warping is then applied to the soft tissue to provide a representation of a face for the questioned individual.

Several implementations of this basic method exist. For example, Quatrehomme *et al.* [11] described a facial reconstruction method where the correspon-

dences that guide the warping are automatically discovered based on matching “crest lines” between the known and questioned skulls. Crest lines [13, 14] are the tops of ridges in the skull and represent lines of high curvature. This matching was iterative and used both warpings and rigid transformations. Quatrehomme *et al.* used polygonal models and chose the soft tissue model (population) to be a single representative case.

Michael and Chen [10] and Nelson and Michael [12] worked directly from the volumetric data. Their warping is based on the idea of disc fields. The basic operation is to disperse circular fields defined by a center point, normal and radius through the data around features of interest. The circular fields on the questioned skull were associated with similar fields on the known skull. The relative center points of associated disk fields, their relative radii, and the relative orientation and magnitude of the normal vectors defined a warping from the known skull to the questioned skull. Once the correspondences were determined and the discs defined, the volume was deformed as a complete entity which had the effect of dragging the soft tissue along with the skull. Michael and Chen used multiple known skull/soft tissue pairings in their known database, and composited the soft tissue together using volume summation to get a representative set of tissue measurements. Nelson and Michael appeared to work from a single best representation, claiming that the compositing operation distorted the results. While the distortion claim may be valid in the context of their compositing algorithm, there is no way to get statistically verifiable results from a single sample. Statistical validity can only be assured by using a large database of varying facial characteristics. A proper implementation can allow deviations from this average face in a statistically measurable and bounded fashion should additional facial details prove valuable.

Jones [9] also worked from volumetric data. He generated correspondences between the questioned and known skulls using correlation, a signal processing technique used to find the best match between portions of a signal. In this case, the signal was a two dimensional view of the front of the skull. Once these correspondences were determined, he defined a warping which related the correspondences. Instead of applying this warping to the known skull to get a dense mapping, this warping was used to define a probe operation. This probe operation allowed each volume element in the questioned skull volumetric data to be associated with an element in the known head volumetric data. In some sense, this is a direct dense analog to methods based on tissue depth tables.

## 2.1 (anonymized)

Forensic facial reconstruction has two conflicting constraints. A reconstruction must contain sufficient detail that the face is recognizable, but the forensic artist must not inject detail where detail does not exist. For instance, in the anatomical approach, Gerasimov [5], suggested that underlying musculature should be designed to correspond to cues in the skull structure, while those that follow the tissue depth method such as Aulsebrook [7] suggested that an objective and repeatable reconstruction can only be made when average tissue-

depth tables guide the forensic artist. Details of the state of facial reconstruction are reviewed in [1, 15, 16].

(anonymized) and (anonymized) are designing a software package, (*anonymized*), to address the problem of balancing facial detail against deviations from statistical facial norms. This package provides both a composite representation of an “average” face for a particular skull structure, and a method for exploring the likely variations in facial features from within a statistical framework. Many of the algorithms exist and are being combined together and tested against a growing database of CT scans.

For a given questioned skull, our method constructs multiple candidate faces from heads within a CT database. These candidate faces are analyzed using a technique from structural analysis, Principle Components Analysis (PCA). A PCA allows us to determine the average face and the ways in which the facial structure varies in terms of independent variation vectors. These vectors describe the face in terms of likely variation from the average face. In the average case, the contribution from each of the variation vectors is clamped at 0. A face is allowed to deviate from the average by increasing the contribution of one or more of the variation vectors. The vectors have statistical significance and a given vector weight has a statistical likelihood of occurring in the general population. These variations are termed eigenvectors. Using the terminology of [17] for a similar process in 2 dimensions, we term a reconstruction based on a combination of these variations with the average face an eigenface. Similarly, we will term the set of all eigenfaces we can construct from a set of eigenvectors and an average face the “face-space”.

The approach we discuss here relies on generating a face-space that contains a recognizable representation of the questioned face. The mechanics of the face-space construction are discussed in more detail in (anonymized) [18], but for the purposes of this paper, the generation of the face-space can be shown to require a set of candidate faces each of which has the same bony structure as the questioned skull, but which have differing soft tissue. Previous studies of facial variation from 2-dimensional images indicate that between 15 and 50 independent variation vectors are required to completely characterize the facial structure. Our process is somewhat different, but we expect to require the same complexity to describe our face space and are targeting 50 eigenvectors. Each eigenvector requires at least one sample face, however, a general rule of thumb suggests that increased accuracy in the eigenvector calculation can be obtained by increasing the number of faces to 3 per eigenvector desired. This implies that we need between 50 and 150 samples to fill the face-space in a statistically valid fashion. Rather than searching for 50 heads whose skull exactly matches a questioned skull, we maintain a database of heads of varying bone and soft tissue structures and use warping to match known skulls to the questioned skull. The same warping, applied to the soft tissue, provides one candidate face to the PCA.

The need for multiple candidates implies that *anonymized* must have an automated and objective method for generating the warpings from known to questioned skulls. The method presented in Section 2.2 extends previous com-

putational approaches, particularly, the work of Quatrehomme *et al.* [11], to give improved, automatic registrations from which the candidate faces can be derived.

## 2.2 Details of the Registration

We now discuss the generation of a warping function,  $\mathbf{T}_{(K,Q)}()$  that maps the known skull  $K$  onto the questioned skull  $Q$  as

$$\mathbf{S}_Q \approx \mathbf{T}_{(K,Q)}(\mathbf{S}_K), \quad (1)$$

for questioned skull  $\mathbf{S}_Q$  and known skull  $\mathbf{S}_K$ . The candidate face,  $\mathbf{F}_Q$  can then be written as

$$\mathbf{F}_Q = \mathbf{T}_{(K,Q)}(\mathbf{F}_K), \quad (2)$$

for the soft tissue of the the known head  $\mathbf{F}_K$ .

Our deformation robustly and autonomously registers one skull to a second, with a few assumptions and caveats. First, since we have control of the collection of the known skulls, we only accept skulls that are substantially complete. The skulls have no missing bony sections, and have a closed mouth. They have limited dental work and, in particular, we assume they have all the teeth except, perhaps, for the second and third molars. For the questioned skulls, we assume the skulls also have a closed mouth and that they are largely complete. Our algorithm is designed to tolerate holes and missing bony sections in the questioned skulls, so completeness is not as much of an issue. In particular, we expect proper operation when there are holes in the maxilla, and frontal or parietal bones. However, our efforts to date have not been able to fully categorize this behavior and an analysis of the accuracy of the algorithm with damaged skulls prior to paleoanthropological restoration remains of interest. Behavior following paleoanthropological restoration should be comparable to that of complete skulls.

Our warping algorithm is designed to merge the surface of a known skull with the surface of a questioned skull. Interior details such as the nasal passages, the interior surface of the skull and the spinal column, while important to the fidelity of the reconstruction and to visualization of the results, do not aid the calculation of the warp. This complexity is ignored by the algorithm, but remains available and is presented to the user in the final reconstruction. The remaining surfaces become the input to the registration algorithm and the deformation  $\mathbf{T}_{(K,Q)}()$  is calculated based on them.

Part of our algorithm relies upon finding and matching crest lines or lines of maximal curvature on the skull [11, 14]. These lines are formally calculated based on spatial derivatives of the skull surface [13], however, since our surfaces are composed of triangular patches, we approximate crest lines by looking at the angle of intersection of two triangular faces. If the angle is greater than a threshold defined as  $30^\circ$ , the line of intersection is considered a segment of a crest line. This is a heuristic threshold designed to give coherent crest lines

around major features such as the orbits and edges of the mandible, while avoiding spurious crestlines due to noise.

Figure 4 shows two skulls from our database. Skull (a) shown after an initial scaling operation, is designated known and skull (b) is designated questioned. We take known as the known head and use the skin surface as the known soft tissue. The second skull, questioned, is taken as the questioned skull. The results following the preprocessing stage are shown in Figure 5. Note that the skulls have been reduced to their single, outer surface, and crest lines (red and green markings) have been generated. Extraneous features including the spinal column, and interior surfaces of the skulls are not shown. Figure 6a shows the two skulls superimposed upon each other, while the initial, undeformed face is shown alongside in Figure 6b.

Following this preprocessing, the algorithm proceeds as shown in Figure 7. In this figure the symbols  $t_i$  represent the deformations calculated at the  $1 \leq i \leq 3$  passes,  $S_Q$  represents the questioned skull,  $S_K$  represents the known skull, and  $S_K^i$ ,  $1 \leq i \leq 3$  represent the 3 intermediate skulls generated by the passes.

The individual passes are variants of the same base algorithm. We apply different different inputs and parameter settings to bring specific part of the skull into alignment. The first step provides rough alignment of skull features based on matching the crest lines of the known skull to those of the questioned; the second more tightly aligns areas of high curvature; and the third tightly aligns smooth areas away from the high curvature while maintaining the alignment of the high curvature areas. The output of each pass becomes one of the inputs to the next pass, and that pass 2 is actually an inverse pass, calculating a deformation for the questioned skull onto skull  $S_K^1$ , but applying the inverse of the calculated transform to  $S_K^1$ . We can write the complete deformation as the composition of the three passes

$$\mathbf{T}_{(K,Q)}(\mathbf{S}_K) = \mathbf{t}_3(\mathbf{t}_2^{-1}(\mathbf{t}_1(\mathbf{S}_K))). \quad (3)$$

Each pass has the same construction. We start by removing outliers to aid robustness when registering to incomplete skulls. If we are trying to calculate a deformable transformation,  $\mathbf{t}_i()$  for pass  $i$ , such that

$$\mathbf{S}_X \approx \mathbf{t}_i(\mathbf{S}_Y), \quad (4)$$

the outlier removal consists of generating a rigid transform  $\mathbf{r}'_i(\mathbf{S}_X)$ . This is an Iterative Closest Point (ICP) algorithm [19], that finds points on the two skulls that lie close to one another when the skulls are aligned. Given a source  $\mathbf{S}_X$  and a target  $\mathbf{S}_Y$  the ICP algorithm first selects a set  $\mathbf{P}_X$  of  $n$  points from  $\mathbf{S}_X$ , selects a set  $\mathbf{P}_Y$  of the  $n$  closest matching points from  $\mathbf{S}_Y$ , calculates the rigid transformation that minimizes the distance between the sets  $\mathbf{P}_X$  and  $\mathbf{P}_Y$ , applies the transform to  $\mathbf{P}_X$  and starts over by finding a new set  $\mathbf{P}_Y$ . At the end of the algorithm, the set  $\mathbf{P}_X$  is denoted as the source landmarks and the set  $\mathbf{P}_Y$  as the target landmarks of  $\mathbf{r}'_i()$ .

The transformation  $\mathbf{r}'_i()$  rigidly transforms the skulls in the opposite of the desired direction and will be discarded by the process. However, this transform

serves to identify points,  $\mathbf{P}_Y \in \mathbf{S}_Y$  that have close neighbors in  $\mathbf{S}_X$ . We can be reasonably certain that any point  $p_{y_i} \in \mathbf{P}_Y$  has a close corresponding point  $p_{x_j} \in \mathbf{S}_X$ ; *i.e.*, if a point is in the target landmarks of the initial rigid transformation, then we can assume the point is not an outlier, does not lie over a missing bony structure in the source skull, and hence, must have a close correspondence to a source point.

We then calculate a new ICP transform  $\mathbf{r}_i()$  based only on this reduced set of points, such that  $\mathbf{S}_X \approx \mathbf{r}_i(\mathbf{P}_Y)$ . By limiting the source points  $\mathbf{r}_i()$  can select from, we limit the members of  $\mathbf{P}_Y$  to valid correspondence points.

The final step is the calculation of a deformation,  $\mathbf{d}_i$ , based on the thin plate spline of Bookstein [20]. This is a straightforward calculation given two sets of correspondences. It treats the correspondences as lying in a deformable medium and pushes and pulls points in the source set until they lie close to the target set. Deformations propagate outward from the source points. The intuition is that the source points are embedded in a block of rubber and are moved within that block until they align with the corresponding target points. The alignment is not absolute and a relaxation parameter  $\sigma$  allows the algorithm to trade off positional accuracy of the alignment with stresses induced by the deformation. We select the source landmarks of the second ICP registration,  $\mathbf{r}_i(\mathbf{P}_Y)$ , as the points to be deformed by the thin plate spline. The fixed points of the deformation are determined by searching the skull surface  $\mathbf{S}_X$  for the closest point on the skull surface with a surface normal orientation (line perpendicular to the surface at the point) within a threshold  $\tau$  of the surface normal at the corresponding point in  $\mathbf{r}_i(\mathbf{P}_Y)$ .

The full transformation,  $\mathbf{t}_i$  is then given as

$$\mathbf{t}_i(\mathbf{X}) = \mathbf{d}_i(\mathbf{r}_i(\mathbf{X})). \quad (5)$$

As indicated in Figure 7, the three passes calculate

$$\mathbf{S}_Q \approx \mathbf{S}_K^1 = \mathbf{t}_1(\mathbf{S}_K), \quad (6)$$

$$\mathbf{S}_K^1 \approx \mathbf{t}_2(\mathbf{S}_Q), \mathbf{S}_K^2 = \mathbf{t}_2^{-1}(\mathbf{S}_K^1), \quad (7)$$

$$\mathbf{S}_Q \approx \mathbf{S}_K^3 = \mathbf{t}_3(\mathbf{S}_K^2). \quad (8)$$

### 3 Results and Discussion

While the passes are conceptually the same, we choose different inputs and different parameters to achieve the desired results. These values are determined heuristically based on experiments and may change as our database grows. For pass 1, we are coarsely aligning skull features. The selected points are limited to points on the crest lines and we limit  $\mathbf{P}_X$  and  $\mathbf{P}_Y$ , to roughly 200 points each. The angle constraint between normals is fairly rigid requiring source and target points for the deformation to lie within  $\tau = 18^\circ$ . During pass 1 we also allow the thin plate spline more freedom to define smooth transformations by setting the relaxation parameter  $\sigma = 10$ . Figure 8a shows a superposition of the two

skulls following the initial alignment pass with the associated facial deformation in Figure 8b. This stage is designed to bring areas of high curvature on the the known skull into alignment with corresponding areas on the questioned skull. This can be seen most clearly in the chin and zygomatic arch regions of Figure 8. The chin region is shortened and pulled in toward the face for both the known skull and the soft tissue. The zygomatic arch region is aligned with that of the questioned skull.

Pass 2 is designed to work similarly to pass 1, except that the matching is from the questioned to the known skulls. This pass is an attempt to reconcile skull areas where the pass 1 algorithm was unable to correctly differentiate between multiple candidate crest lines on the questioned skull. By running the same algorithm in the opposite direction, on the same points, there is a higher likelihood of resolving these conflicts. For this pass we again limit the points available for registration to those calculated as lying on crest lines of the skull but raise the number of landmarks used in the registration, typically the number of points used lies in the range of 200 - 350. Since this pass is a reverse pass, we calculate the actual warping such that  $\mathbf{S}_K^1 \approx \mathbf{t}_2 \mathbf{S}_Q$ , and generate  $\mathbf{S}_K^2 = \mathbf{t}_2^{-1} \mathbf{S}_Q^1$ . We set  $\sigma = 1$  and  $\tau = 60^\circ$ . Figure 9 shows the skull alignment and facial deformation following the completed pass 2. The most significant changes from this pass are in the maxilla/front tooth region where the known skull and soft tissue are pulled in toward the questioned skull surface and in the frontal process region where the known skull and soft tissue are pulled out toward the questioned skull surface. Note that by aligning crest lines, we are under-constraining the surface in areas where crest lines are sparse. Changes in the surfaces toward the anterior of the skull and soft tissue are the result of this and will be corrected during pass 3.

Finally, pass 3 is intended to complete the skull deformation. We again use a large number of landmarks divided among the landmarks used in pass 2 and an additional 200-500 landmarks chosen from the smooth skull surfaces to ensure dense coverage of the skull. As mentioned previously, stiffer splines force greater alignment between the correspondences chosen. Since we now assume that points on the known skull are generally close to their correspondences on the questioned skull, we can choose a stiffer spline,  $\sigma = 1$ . We again set  $\tau = 60$ . Figure 10 shows the results of this final pass of the algorithm, with Figure 10a showing the deformed known skull superimposed on the questioned skull and Figure 10b showing the corresponding face. Most changes from pass 2 occur in the flatter regions of the skulls. In particular, the mandible, frontal bone and anterior portions of the known skull have been moved into correspondence with the questioned skull. The soft tissue shows subtle changes such as a flattening of the crown, and a broader more pronounced cheek region.

Throughout the algorithm we use the ICP algorithm both to align the different instances of deformed skulls and in the selection of landmarks for the thin plate spline deformation. The landmarks selected for the three passes are shown in Figure 11a, Figure 11b and Figure 11c, respectively. As illustrated in the figure, Pass 1 (Figure 11a) and Pass 2 (Figure 11b) use a smaller number of landmarks limited to the high curvature areas of the skull. Pass 3 (Figure 11c),

targeting overall registration, augments the points selected by the second pass with an additional set of correspondences distributed on the skull surface. Figure 11c shows only those points chosen during the third pass. All of the points 11b and 11c are used in the calculation of the third pass deformation. Figure 12 shows a closeup of the skulls following each step of the registration. The area chosen is the right zygomatic process and the ascending ramus of the mandible with a progression from initial alignment through the three passes in figures 12a, 12b, 12c, and 12d, respectively. Figure 12b shows greater alignment of high curvature areas than Figure 12a. Changes in Figure 12c are less pronounced and correspond to a slight warping. The final inset, Figure 12d shows a more highly aligned surface with lower curvature areas also matched.

## 4 Conclusions

In order to improve on the accuracy of clay-based forensic facial reconstructions, we look at various sources of imprecision: sparse depth measurements, imprecise placement of landmarks, and imprecise angle of soft tissue penetration, all of which force the forensic artist into subjective and variable interpretations of the questioned face. We present a computer-based method using dense tissue measurements designed to reduce and to eliminate these errors. Our method, implemented in the software tool *anonymized*, automatically generates a dense set of objective landmarks between the questioned skull and a set of known face/skull pairs. These landmarks and the deformation field associated with them are the basis for a mathematical description of the space of facial variation.

Our results show that a multipass ICP-based algorithm can register skulls as a part of a facial reconstruction framework. This algorithm can effectively handle not only the case of registering complete skulls one to another but also the real-world situation of registering a partial skull to a complete skull from the database.

The remaining issues for validation of this research are to apply the algorithm to a larger sample of skulls selected from the CT database and to use the resultant registration to compute a “face space” [18] that accurately depicts the variations in faces. Given this face space we will be able to reconstruct faces which have a statistically quantifiable likelihood of occurring in the target population, and can explore the space of likely variation. The final step is then to validate the reconstruction process using the ground truth database that comes from the CT data. We intend to run a series of studies that seeks to reconstruct a face from the database starting from only the skull to which the face belongs and the other members of the database. Given the reconstruction based on this data we will then capture both a quantitative measure of closeness of approximation similar to the metric used to as a figure of merit for the registration and a qualitative measure of “observer ability to recognize”.

## References

- [1] anonymized.
- [2] W. Drake, L. Lukash, Reconstruction of mutilated victims for identification, *J Forensic Sci* 23 (1978) 218 – 230.
- [3] K. T. Taylor, *Forensic Art and Illustration*, CRC Press, New York, 2001.
- [4] O. Grüner, Identification of skulls: A historical review and practical applications, in: M. Y. İşcan, R. P. Helmer (Eds.), *Forensic Analysis of the Skull*, Wiley-Liss, New York, 1993, pp. 29 – 45.
- [5] M. M. Gerasimov, *The Face Finder*, JB Lippincott Co, Philadelphia (PA), 1971.
- [6] J. Prag, R. Neave, *Making Faces*, Texas A&M University Press, College Station (TX), 1997.
- [7] W. A. Aulsebrook, Facial tissue thickness in facial reconstruction, in: J. A. Siegel, P. J. Saukko, G. C. Knupfer (Eds.), *Encyclopedia of Forensic Sciences*, Academic Press, San Diego (CA), 2000, pp. 779 – 788.
- [8] P. Vanezis, M. Vanezis, G. McCombe, T. Niblett, Facial reconstruction using 3-D computer graphics, *Forensic Sci Int* 108 (2000) 81–95.
- [9] M. W. Jones, Facial reconstruction using volumetric data, in: T. Ertl, B. Girod, G. Greiner, H. Niehann, H. P. Seidel (Eds.), *Vision, Modeling, and Visualization 2001. Proceedings of VMV 2001*; 2001 Nov 21 – 23; Germany, IOS Press, Amsterdam, 2001, pp. 135 – 142.
- [10] S. D. Michael, M. Chen, The 3d reconstruction of facial features using volume distortion, in: *14th Annual Conference of Eurographics, UK Chapter*; 1996 Mar 26 – 28; London, 1996, pp. 297 – 305.
- [11] G. Quatrehomme, S. Cotin, G. Subsol, H. Delingette, Y. Garidel, G. Grévin, M. Fidrich, P. Bailet, A. Ollier, A fully three-dimensional method for facial reconstruction based on deformable models, *J Forensic Sci* 42 (4) (1997) 649 – 652.
- [12] L. A. Nelson, S. D. Michael, The application of volume deformation to three-dimensional facial reconstruction: A comparison with previous techniques, *Forensic Sci Int* 94 (1998) 167 – 181.
- [13] N. Ayache, A. Guézic, J. P. Thirion, A. Gourdon, Evaluating 3D registration of CT-scan images using crest lines, in: *Proceedings of SPIE'93, Mathematical Methods in Medical Imaging II*, Vol. 2035-06, San Diego, California, USA, 1993, pp. 60–71.

- [14] A. Guéziec, N. Ayache, Large deformable splines, crest lines and matching, in: Proceedings of SPIE'93, Geometric Methods in Computer Vision II, Vol. 2031-28, San Diego, California, USA, 1993, pp. 316–327.
- [15] W. A. Aulsebrook, M. Y. İřcan, J. H. Slabbert, P. Becker, Superimposition and reconstruction in forensic facial identification: a survey, Forensic Sci Int 75 (1995) 101 – 120.
- [16] anonymized.
- [17] M. A. Turk, A. P. Pentland, Face recognition using eigenfaces, in: IEEE Computer Society Conference on Computer Vision and Pattern Recognition. Proceedings of CVPR '91; 1991 June 3 – 6; Maui (HI), IEEE Computer Society, Washington (DC), 1991, pp. 586 – 591.
- [18] anonymized.
- [19] P. J. Besl, N. D. McKay, A method for registration of 3d shapes, IEEE Transactions on Pattern Analysis and Machine Intelligence 14 (1992) 239–356.
- [20] F. Bookstein, Principle warps: Thin plate splines and the decomposition of deformations, IEEE Transactions on Pattern Analysis and Machine Intelligence 11 (1989) 567–585.

## List of Figures

1	A taxonomy of forensic facial approximation. The field is initially segmented based on the amount of facial information available. Adapted with permission from (anonymized) [1] . . . . .	15
2	A simple, one-dimensional illustration of transformations. . . . .	16
3	An illustration of a polygonal model. The image shows the full model of a skull rendered as a surface, while the inset shows detail of the the segments bounding the triangular faces on a portion of the mandible with teeth. . . . .	17
4	Original skulls after scaling, (a) Skull known representing the known skull $K$ , and (b) Skull questioned representing the questioned skull $Q$ . Both skulls are from the ( <i>anonymized</i> ) database. . . . .	18
5	Skulls after pre-processing to extract outer surfaces and generate crest line approximations (colored lines). Images are derived from (a) Skull known , and (b) Skull questioned . . . . .	19
6	Initial alignment including rescaling. Images are derived from (a) Skull known (yellow) and Skull questioned (white) superimposed, and (b) knownface. This shows the complete skulls from Figure 4 superimposed to highlight the differences. . . . .	20
7	Overview of the registration algorithm showing the three separate phases: (a) Rough registration and gross alignment, (b) Alignment of high curvature areas using crest lines, and (c) Final surface to surface alignment. . . . .	21
8	Position of (a) known (yellow) and questioned (white) skulls superimposed, and (b) known face after initial deformation. These images represent the completion of Pass 1. Note particularly, that the chin area in both the known skull and the flesh is shortened and pulled into the face compared to 6 and the zygomatic arch is aligned. . . . .	22
9	Position of (a) known (yellow) and questioned (white) skulls superimposed, and (b) known face after second (crest line) deformation. The most significant changes from this pass are in the anterior region of the skull and the maxilla/front tooth regions where the known skull and soft tissue are pulled in toward the questioned skull surface and in the region of the frontal process where the known skull and soft tissue are pulled out. Since there are no crest lines in the anterior region of the skull, changes in that area are side effects of crest line matches at other points of the skull. These changes will be removed by pass 3. . . . .	23

10	Position of (a) known (yellow) and questioned (white) skulls superimposed, and (b) known face after the complete deformation. Note that most changes from pass 2 occur in the flatter regions of the skulls. In particular, the mandible, frontal bone and anterior portions of the known skull have been moved into correspondence with the questioned skull. The soft tissue shows subtle changes such as a flattening of the crown, and a broader more pronounced cheek region. . . . .	24
11	Alignment points chosen by the algorithm during pass (a) 1, coarse , (b) 2, crest line, and (c) 3, final alignment. Passes 1 and 2 use different sets of correspondences, but both concentrate correspondences along high curvature areas of the skulls. Pass 3 adds points along the lower curvature surfaces. . . . .	25
12	Close up view of the posterior of the mandible and the zygomatic process after (a) scaling and rough alignment, (b) pass 1 deformation, (c) pass 2 deformation, and (d) pass 3 (final) deformation. The insets show the progression of the alignment from the initial state in (a) through the three passes in (b), (c), and (d). . . . .	26

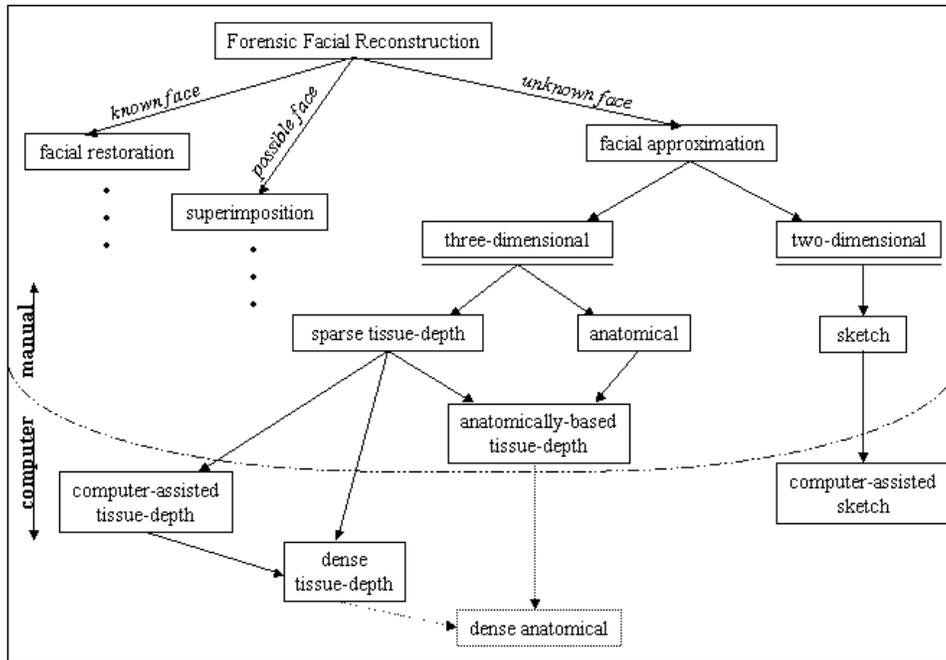


Figure 1: A taxonomy of forensic facial approximation. The field is initially segmented based on the amount of facial information available. Adapted with permission from (anonymized) [1]

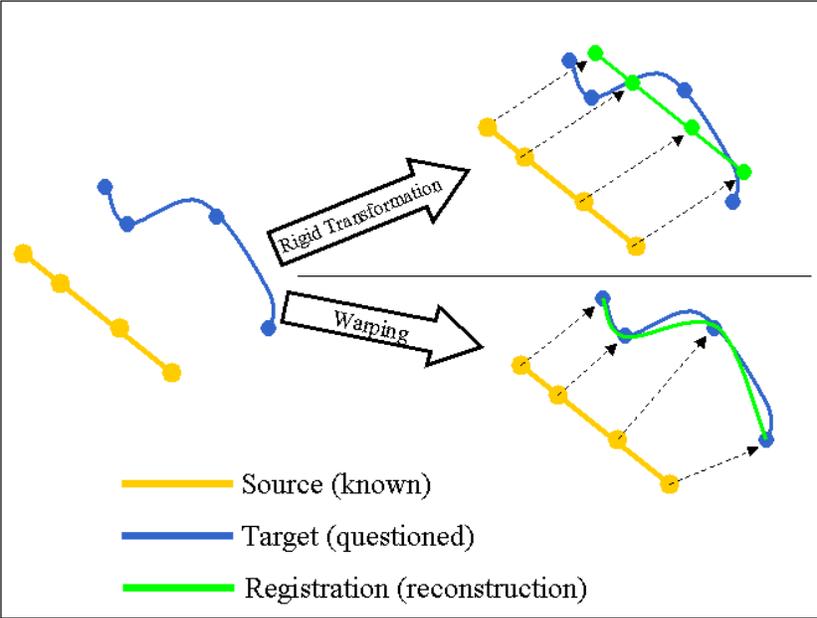


Figure 2: A simple, one-dimensional illustration of transformations.

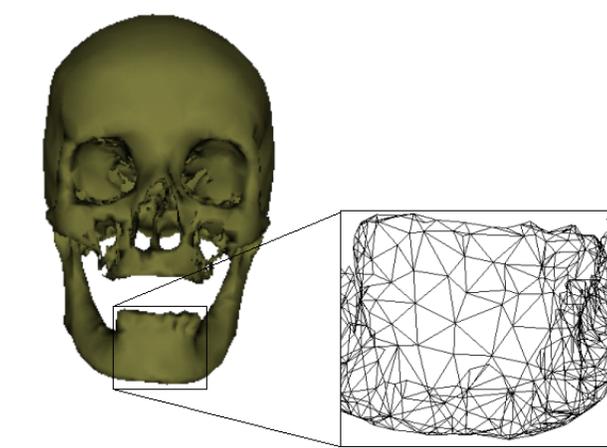


Figure 3: An illustration of a polygonal model. The image shows the full model of a skull rendered as a surface, while the inset shows detail of the the segments bounding the triangular faces on a portion of the mandible with teeth.

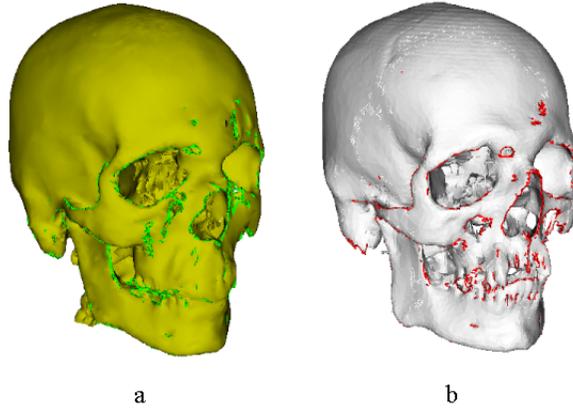


Figure 4: Original skulls after scaling, (a) Skull known representing the known skull  $K$ , and (b) Skull questioned representing the questioned skull  $Q$ . Both skulls are from the (*anonymized*) database.

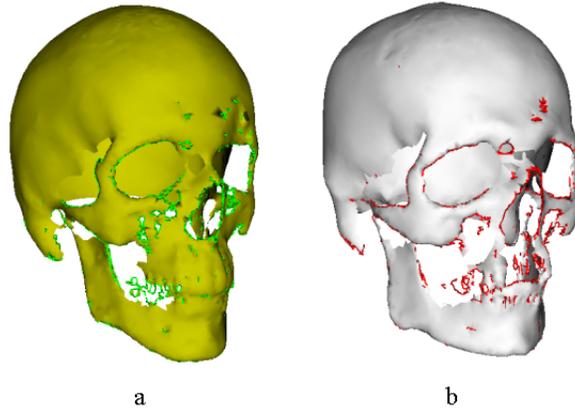


Figure 5: Skulls after pre-processing to extract outer surfaces and generate crest line approximations (colored lines). Images are derived from (a) Skull known , and (b) Skull questioned .

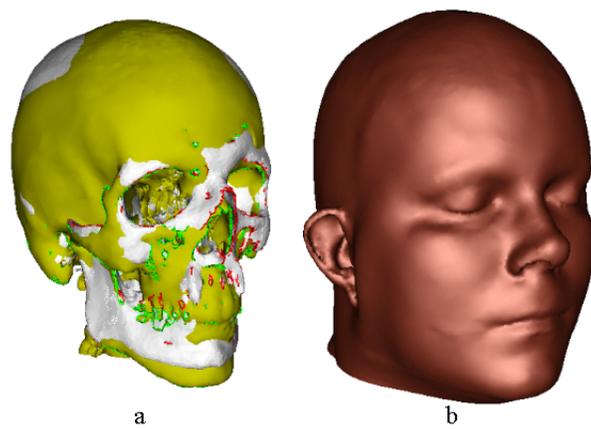


Figure 6: Initial alignment including rescaling. Images are derived from (a) Skull known (yellow) and Skull questioned (white) superimposed, and (b) known face. This shows the complete skulls from Figure 4 superimposed to highlight the differences.

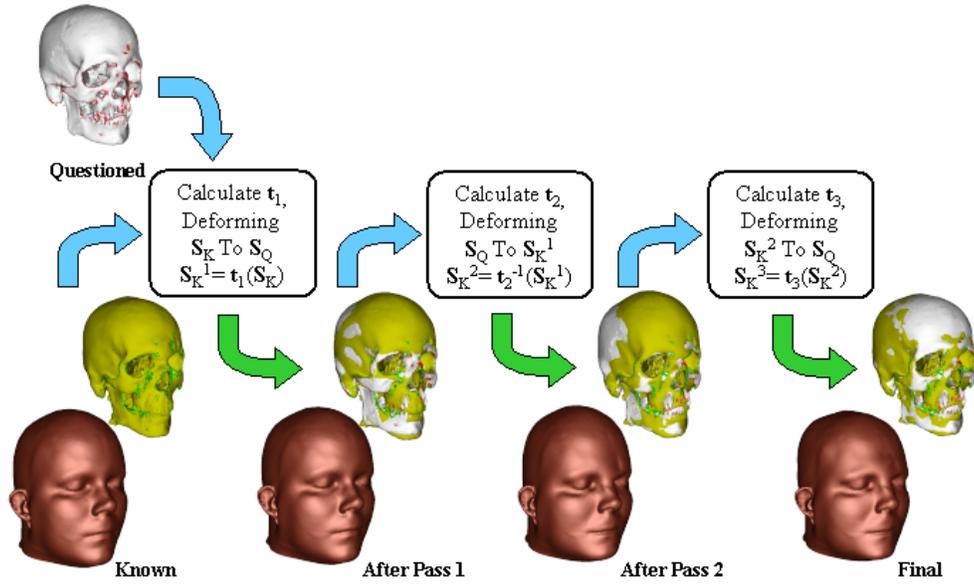


Figure 7: Overview of the registration algorithm showing the three separate phases: (a) Rough registration and gross alignment, (b) Alignment of high curvature areas using crest lines, and (c) Final surface to surface alignment.

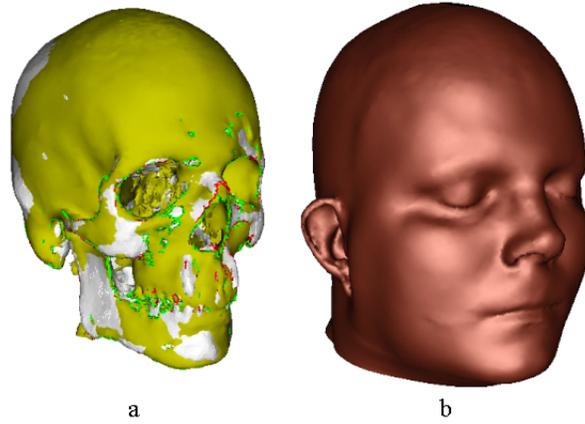


Figure 8: Position of (a) known (yellow) and questioned (white) skulls superimposed, and (b) known face after initial deformation. These images represent the completion of Pass 1. Note particularly, that the chin area in both the known skull and the flesh is shortened and pulled into the face compared to 6 and the zygomatic arch is aligned.

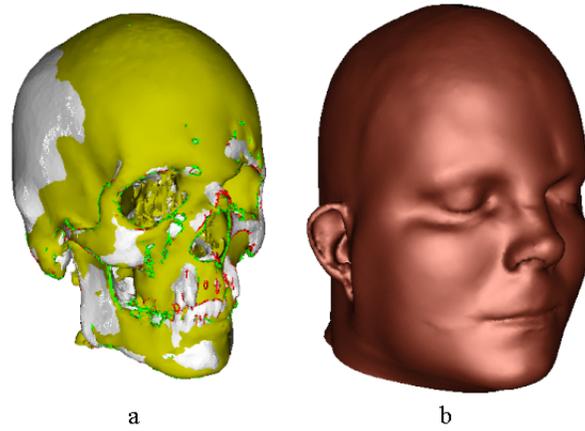


Figure 9: Position of (a) known (yellow) and questioned (white) skulls superimposed, and (b) known face after second (crest line) deformation. The most significant changes from this pass are in the anterior region of the skull and the maxilla/front tooth regions where the known skull and soft tissue are pulled in toward the questioned skull surface and in the region of the frontal process where the known skull and soft tissue are pulled out. Since there are no crest lines in the anterior region of the skull, changes in that area are side effects of crest line matches at other points of the skull. These changes will be removed by pass 3.

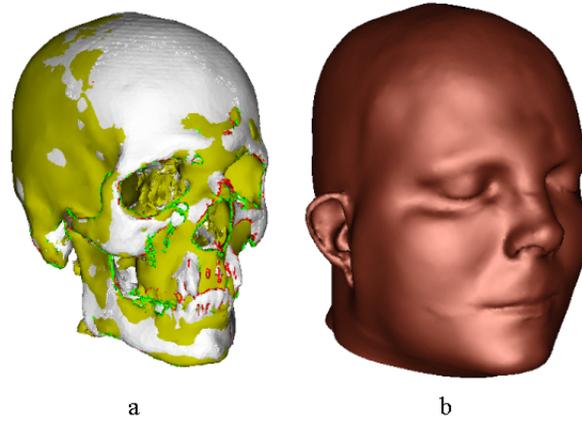


Figure 10: Position of (a) known (yellow) and questioned (white) skulls superimposed, and (b) known face after the complete deformation. Note that most changes from pass 2 occur in the flatter regions of the skulls. In particular, the mandible, frontal bone and anterior portions of the known skull have been moved into correspondence with the questioned skull. The soft tissue shows subtle changes such as a flattening of the crown, and a broader more pronounced cheek region.

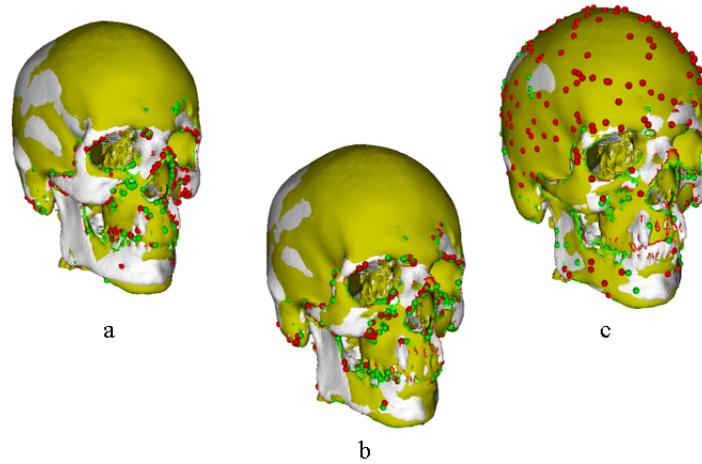


Figure 11: Alignment points chosen by the algorithm during pass (a) 1, coarse , (b) 2, crest line, and (c) 3, final alignment. Passes 1 and 2 use different sets of correspondences, but both concentrate correspondences along high curvature areas of the skulls. Pass 3 adds points along the lower curvature surfaces.

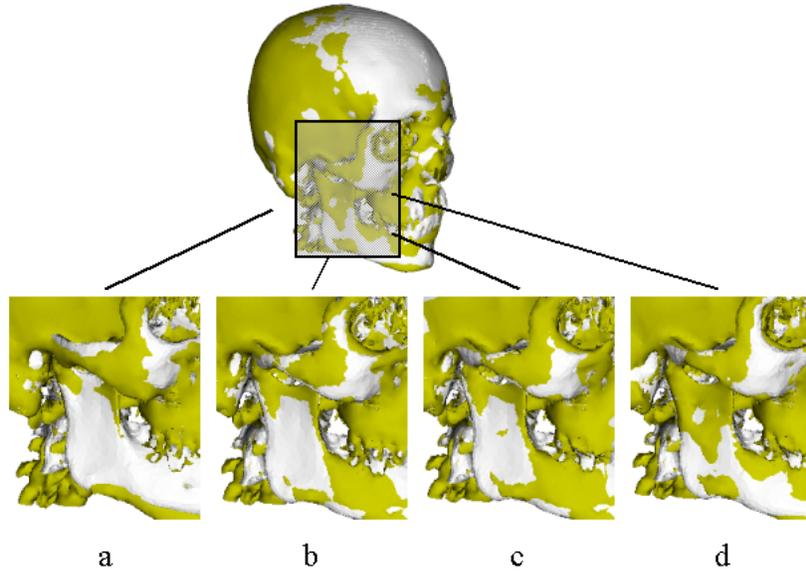


Figure 12: Close up view of the posterior of the mandible and the zygomatic process after (a) scaling and rough alignment, (b) pass 1 deformation, (c) pass 2 deformation, and (d) pass 3 (final) deformation. The insets show the progression of the alignment from the initial state in (a) through the three passes in (b), (c), and (d).