

Probabilistic Joint Face-Skull Modelling for Facial Reconstruction

Dennis Madsen Marcel Lüthi Andreas Schneider Thomas Vetter
 {dennis.madsen, marcel.luethi, andreas.schneider, thomas.vetter}@unibas.ch
 Department of Mathematics and Computer Science
 University of Basel

Abstract

We present a novel method for co-registration of two independent statistical shape models. We solve the problem of aligning a face model to a skull model with stochastic optimization based on Markov Chain Monte Carlo (MCMC). We create a probabilistic joint face-skull model and show how to obtain a distribution of plausible face shapes given a skull shape. Due to environmental and genetic factors, there exists a distribution of possible face shapes arising from the same skull. We pose facial reconstruction as a conditional distribution of plausible face shapes given a skull shape. Because it is very difficult to obtain the distribution directly from MRI or CT data, we create a dataset of artificial face-skull pairs. To do this, we propose to combine three data sources of independent origin to model the joint face-skull distribution: a face shape model, a skull shape model and tissue depth marker information. For a given skull, we compute the posterior distribution of faces matching the tissue depth distribution with Metropolis-Hastings. We estimate the joint face-skull distribution from samples of the posterior. To find faces matching to an unknown skull, we estimate the probability of the face under the joint face-skull model. To our knowledge, we are the first to provide a whole distribution of plausible faces arising from a skull instead of only a single reconstruction. We show how the face-skull model can be used to rank a face dataset and on average successfully identify the correct match in top 30%. The face ranking even works when obtaining the face shapes from 2D images. We furthermore show how the face-skull model can be useful to estimate the skull position in an MR-image.

1. Introduction

In forensics, facial reconstruction is the process of recreating the face of an unknown individual solely from their skeletal remains [27]. Forensic artists reconstruct the face surface by adding layers of muscle fibres and placing tissue depth markers on the unknown skull. However, recon-

structing the face shape from a skull is an ill-posed problem. To begin with, it is commonly known that the facial reconstruction process is highly artistic and different forensic artists might produce different face predictions [28]. In addition, many features of a shape are not due to the skull itself, but due to genetics and environmental factors such as diet and ageing. As a result, for a given skull there exists a whole distribution of possible faces. Therefore, instead of a one-to-one relation between skulls and faces, we assume a one-to-many relationship. We propose to build a joint distribution of face and skull shape to be able to reconstruct a distribution of faces contrary to a single face in facial reconstruction. It is difficult to estimate the distribution directly from Computed Tomography (CT) or Magnetic resonance imaging (MRI) data as it would be necessary to capture multiple scans of the same person exposed to different environmental influences and both the face and the skull are difficult to segment. Instead, we propose to cre-




Figure 1. Illustration of the face-skull model with the face fully visible, one side visible and fully visible with an opacity of 70%.

ate an artificial dataset of corresponding skulls and faces. To this end, we combine a statistical shape model of human faces [18] with a statistical skull shape model. No point-to-point correspondence is assumed. The two models are independent and contain scans from different individuals. We compute the posterior distribution of faces given a skull using a Markov Chain Monte Carlo (MCMC) approach. A sparse set of distance constraints between the surface of the skull and the surface of the sampled faces is everything the method needs. The distance constraints are known as tissue-depth markers within facial reconstruction and represent the

depth of the tissue in between the skull and the face surface such as fat, muscle fibres and cartilage. We generate face-skull pairs by sampling from the posterior face distribution found for each skull. Then, we model the joint distribution according to the generated data. The result is a probabilistic model of skull and face shape that combines two independent statistical shape models. With the joint statistical face-skull shape model we can reconstruct face shapes from skulls and match skulls to faces. For a given skull, we are not only obtaining a single face, but also the full distribution of likely faces. The model can even be used to predict the face distribution if only part of the skull is available or it can be constrained by additional information about parts of the face shape. The mean of the joint model is illustrated in Figure 1.

As the model provides us with a distribution of faces, we use it to evaluate the probability that a given face will fit the skull. Within forensics, craniofacial superimposition is the method of comparing a photo of a skull to a photo of a face to evaluate their similarity [26]. We show how the joint model can be used in a similar setting; to rank the likelihood of faces to match a given skull. We compute the face distribution for 9 skulls and rank a face dataset of more than 300 face shapes. The ranking allows us to find the correct matching face for a skull in the top 30%. We show as a proof of concept that this method could be beneficial in tasks such as Disaster Victim Identification.

2. Related Work

The majority of methods within automatic forensic facial reconstruction share a general framework which consists of finding the mapping between skulls and faces [2]. Independent models of the face and the skull are constructed and the parameter mapping between the two models is found from a small test set. In more recent work [22], two individual statistical shape models were created from a large dataset of CT head scans and a linear mapping is found between the two. Similar methods have previously been attempted with smaller datasets from various sources [19, 4, 13, 9, 29, 8]. The model mapping approach maps a skull to a single face and therefore, does not account for shape variance due to genetics and environmental factors such as diet and health. Several approaches, however, suggest changing the facial attributes in a final step to adjust for soft tissue variation due to age and Body Mass Index (BMI) differences [19, 9, 21].

Alternative simple methods place tissue-depth pins on a dry skull and deform a face-mesh according to the mean tissue-depth measures [3].

More recent approaches fit several local shape models to a skull and glue them together into one joint face by using partial least squares regression to find a smooth surface transition [6, 11]. In [17] they are working with images of skulls and try to identify the matching face from a larger set

of unnormalized facial photographs. The evaluation method here does also mimic craniofacial superimposition, similar to our method, by only using pure classification techniques and without taking the soft-tissue variability into account.

The evaluation method, which is most similar to ours, renders 3D landmarks and compares them to 2D landmarks from facial ground truth images [25]. They warp a series of CT images to a target dry skull and compute the mean face prediction. They are able to restrict the solution space by setting the BMI or age of the reconstruction but do not consider the full distribution of possible faces arising from a dry skull. The authors validate their method on a +200 person CT database and receive comparable results to ours (average ranking of the true identity in top 30% on average) even though we use a much smaller dataset to build the model and the data being of independent origin.

3. Background

In this section, we introduce statistical shape models which we will combine in the Method section with a Markov Chain Monte Carlo approach to integrate tissue-depth marker data to build a joint face-skull model.

3.1. Statistical Shape Models

Statistical shape models (SSMs) are linear models of shape variations. The main assumption behind SSMs is that the space of all possible shape deformations can be learned from a set of typical example shapes $\Gamma_1, \dots, \Gamma_n$. Each shape Γ_i is represented as a discrete set of points; i.e. $\Gamma_i = \{x_k^i, x_k \in \mathbb{R}^3, k = 1, \dots, N\}$, where N denotes the number of points. In order to be able to learn the shape variations from the examples, they need to be rigidly aligned to a reference shape and have to be in correspondence. This latter requirement means that the k -th point x_k^i and x_k^j of example Γ_i and Γ_j denote the same anatomical point. To build the model, a shape Γ_i is represented as a vector $\vec{s}_i \in \mathbb{R}^{3N}$, where the x, y, z -components of each point are stacked onto each other:

$$\vec{s}_i = (x_{1x}^i, x_{1y}^i, x_{1z}^i, \dots, x_{Nx}^i, x_{Ny}^i, x_{Nz}^i). \quad (1)$$

This vectorial representation makes it possible to use standard multivariate statistics to model a probability distribution over shapes. The usual assumption is that the shape variations can be modelled using a normal distribution; i.e. $\vec{s} \sim \mathcal{N}(\mu, \Sigma)$, where the mean μ and covariance matrix Σ are estimated from the example data.

As the number of points N is usually large, the covariance matrix Σ cannot be represented explicitly. Fortunately, as the covariance matrix is determined completely by the n example data-sets, it has at most rank n and can, therefore, be represented using n basis vectors. This is achieved by performing a Principal Component Analysis (PCA) [12]. In

its probabilistic interpretation, PCA leads to a model of the form

$$\vec{s} = \vec{\mu} + \sum_{i=1}^n \alpha_i \sqrt{\lambda_i} \vec{u}_i, \quad \alpha_i \sim \mathcal{N}(0, 1) \quad (2)$$

where (\vec{u}_i, λ_i) are the eigenvectors and eigenvalues of the covariance matrix Σ . Note that each shape Γ_i is uniquely determined by the coefficients $\vec{\alpha} = (\alpha_1, \dots, \alpha_n)^T$, and its probability is $p(\vec{s}) = p(\vec{\alpha}) = 2\pi^{-\frac{k}{2}} \exp(-\|\vec{\alpha}\|^2)$. Ignoring discretization and committing a slight abuse of notation, we will in the following refer the surface corresponding to the model parameters $\vec{\alpha}$ as $\Gamma[\vec{\alpha}]$.

By construction, an SSM has a fixed alignment in space. To explain a given surface, we also need to be able to adjust the pose of a surface. In addition to the shape coefficients α , we need to model a translation $\vec{t} = (t_x, t_y, t_z)^T \in \mathbb{R}^3$ and a rotation $R(\phi, \psi, \rho) \in SO(3)$. Here the rotation matrix R is parametrized by the Euler angles ϕ, ψ, ρ . We summarize all the parameter in the vector $\vec{\theta} = (\alpha_1, \dots, \alpha_n, \phi, \psi, \rho, t_x, t_y, t_z)$. Consequently, we write $\Gamma[\vec{\theta}]$ to refer to the resulting surface.

In our case, we use a statistical shape model of skull surfaces, constructed from 30 CT skull shapes. The skull model construction is described in detail in [14]. We also make use of the Basel Face Model [18] to model face shape. We refer to these models as $\Gamma_S[\vec{\theta}]$ and $\Gamma_F[\vec{\theta}]$ respectively.

4. Method

We model the joint face-skull distribution and build a joint face-skull statistical shape model. Because it is difficult or even impossible to observe the joint face-skull distribution directly, we construct it from the conditional distribution over face shapes given skull shapes. Figure 2 shows the general idea of combining the skull and the face model.




Figure 2. Combining the models. We map each skull in the statistical skull model to a distribution of face shapes in the face model. The face-skull model is a joint distribution of the individual skulls from the skull model and their individual face distributions.

4.1. Training Data

We use three sources of data; the publicly available BFM, a statistical skull model and tissue marker data as defined in the publicly available T-Table [23, 24] and shown on a skull shape in Figure 3. All three sources have to our knowledge no overlap in the individuals which the data was obtained from.

4.1.1 Tissue-depth Gaussian assumption

The Gaussian distribution is assumed for each of the individual tissue-depth markers. While the distribution of tissue-depth over age, weight and gender differences might follow a more complex distribution, tissue-depth studies do only provide mean and standard deviation and not the individual measurements. However, the nature of our method would allow easy use of more intricate distributions such as Gaussian mixture models if one has raw data measures available. As our raw skull data does not contain any metadata, we use the T-Table which uses the law of large numbers to combine measures from a variety of age, BMI and gender groups. If skull metadata would be available, alternative data tables could be used to limit the variance and update the mean, as found in [5].




Figure 3. Tissue-depth markers placed on a skull. The landmark distributions are shown at the end of the vectors where one standard deviation is visualised. The illustration on the right shows a face shape fit to the tissue-depth measures on the left.

4.2. Simulating the Joint Face-Skull Distribution

We model the joint face-skull distribution as a multivariate normal distribution over face and skull shape.

$$\Gamma_F, \Gamma_S \sim \mathcal{N}(\mu_F, \Sigma_F), \mathcal{N}(\mu_S, \Sigma_S) \quad (3)$$

As we cannot observe the full joint face-skull distribution directly, we propose to construct it from the skull shapes and tissue depth marker information. Probability theory allows us to write the joint distribution as a product of a conditional distribution and a prior. We write the joint probability of observing a specific face-skull pair as the distribution of face shapes Γ_F given the skull shape Γ_S and the prior of the skull shape.

$$P(\Gamma_F, \Gamma_S) = P(\Gamma_F|\Gamma_S)P(\Gamma_S) \quad (4)$$

The statistical skull shape model provides the prior $P(\Gamma_S)$. The posterior distribution over plausible face shapes, given a skull shape, consists of all the face shapes that fulfil the given tissue depth constraints. We approximate the posterior distribution over face shapes given an observed skull shape with Metropolis-Hastings. Having a posterior distri-




Figure 4. Illustration of the MCMC fitting procedure.

bution over faces given a skull allows us to construct the joint distribution. According to Equation 4, we first sample a random shape from the skull model. Then we approximate the posterior face distribution given that skull. To create a face-skull statistical shape model, we fit a Gaussian distribution to the skull model samples and their corresponding posterior face shape samples.




Figure 5. Pipeline describing how to compute the face-skull model by combining skull samples with the posterior face distributions found with MCMC.

Figure 4 gives an overview of how we approximate the posterior. Figure 5 shows how we build the joint face-skull statistical shape model from skull samples and posterior face shape samples.

4.2.1 Probabilistic Model

In this section, we describe how the probability distribution over possible face shapes is modelled (determined by a parameter vector $\vec{\theta}$) to fit a given skull. We model the probability distribution

$$P(\vec{\theta} | D^{tvi}, D^{sym}, c, D^{cs}), \quad (5)$$

where D^{tvi} , D^{sym} , c , D^{cs} are four different measurements that are derived from the given skull and the tissue depth markers. We will explain them in detail below. Using Bayes theorem we can write the conditional distribution as a product of its prior probability and a likelihood term. As we will

see later, we are only interested in comparing samples, so the evidence term can be ignored. By assuming independence between the individual likelihoods, we can write it as the product:

$$P(\vec{\theta} | D^{tvi}, D^{sym}, c, D^{cs}) \propto P(\vec{\theta}) P_{tvi}(D^{tvi} | \vec{\theta}) P_{tvs}(D^{sym} | \vec{\theta}) P_{fs}(c | \vec{\theta}) P_{cs}(D^{cs} | \vec{\theta}). \quad (6)$$

With the face shape prior $P(\vec{\theta})$ we penalise unlikely shapes. We enforce the tissue-depth distribution with the tissue-vector intersection likelihood P_{tvi} . Tissue-vector symmetry is encouraged by P_{tvs} and face-skull intersection is discouraged with P_{fs} . Furthermore, we enforce correspondence in a single point with the likelihood P_{cs} .

For a given skull shape Γ_S we evaluate the probability of a face shape from its face model parameters $\vec{\theta}$. We place N points $p_i^S \in \mathbb{R}^3$, $i = 1 \dots N$ on Γ_S . For each point we have a tissue-depth vector $\vec{v}_i^S \in \mathbb{R}^3$ with direction towards the face surface (see Figure 3) [24]. We define the point $p_i^F[\vec{\theta}] \in \mathbb{R}^3$ as the intersection point of \vec{v}_i^S with the face model $\Gamma_F[\vec{\theta}]$.

Tissue-vector intersection. We evaluate where the N vectors \vec{v}_i^S intersect with the face model at a given $\vec{\theta}$ by the distance $d_i \in \mathbb{R}$ for each vector. d_i is found from the skull point p_i^S to the intersection point $p_i^F(\vec{\theta})$ on the face shape as $d_i(\vec{\theta}, \Gamma_S) = \|p_i^F[\vec{\theta}] - p_i^S\|$, with D^{tvi} denoting the d_i distances for all the tissue-depth vectors. The probability over all the N points is then the product of each individual likelihood term:

$$P_{tvi}(D^{tvi} | \vec{\theta}) = \prod_{i=1}^N \mathcal{N}(d_i(\vec{\theta}, \Gamma_S), \sigma_i^2) \quad (7)$$

with σ_i^2 being the variance of the individual tissue-depth vector.

Tissue-vector symmetry. We extend the evaluation of tissue-vector intersection by evaluating how similarly mirroring vectors intersect the face surface. Each point p_i^S not placed in the centre of the head has a mirror point. There is a total of m points placed in the centre of the head. We define $p_{-i}^S \in \mathbb{R}^3$ to be the mirror point of p_i^S and $\vec{v}_{-i}^S \in \mathbb{R}^3$ to be the mirror vector of \vec{v}_i^S around the sagittal plane. Let $d_{-i}(\vec{\theta}, \Gamma_S)$ be the tissue-depth vector intersection mirroring $d_i(\vec{\theta}, \Gamma_S)$, then $d_i^{sym}(\vec{\theta}, \Gamma_S) = \|d_i(\vec{\theta}, \Gamma_S) - d_{-i}(\vec{\theta}, \Gamma_S)\|$ is the difference between a tissue-vector intersection depth and its mirror vector intersection depth. D^{sym} is the tissue-depth intersection differences. The combined likelihood for all the mirroring points is:

$$P_{tvs}(D^{sym} | \vec{\theta}) = \prod_{i=1}^{(N-m)/2} \mathcal{N}(d_i^{sym}(\vec{\theta}, \Gamma_S), \sigma^2) \quad (8)$$

with σ^2 set to one. The likelihood will return a high probability if the symmetry vectors intersect the face surface at

the same distance. Figure 6 (left) shows a sample where the tissue-vector symmetry likelihood was not used.

Face in skull detection. The likelihood

$$P_{fs}(c|\vec{\theta}) \sim \text{Exp}(\lambda) \quad (9)$$

penalizes the number of points $c(\vec{\theta})$ of the face model which intersect the skull shape Γ_S . By choosing λ large, faces intersecting with the skull region become very unlikely.

Point correspondence. In our implementation, we have made use of a single point correspondence instance. The point is placed in the centre of the face just below the nose. The correspondence point on the skull is placed just below the anterior nasal spine. The likelihood is necessary in order to avoid misalignment at the nose region as shown in Figure 6 (right). The general point correspondence likelihood




Figure 6. Face-skull alignment problem from missing likelihood terms. Left: Missing tissue-vector symmetry likelihood. The centre of the face and the skull does not align. This is especially visible at the nose which is shifted to the side. The illustrated tissue-vectors have the same length and are not intersecting the tissue at the same distance on both sides of the face. Right: Missing point-to-point correspondence likelihood. The face is rotated slightly which can be seen from the red line which should be horizontal as in the picture to the left. The nose is also not aligned at the centre of the nose, similar to the problem without the point correspondence.

consists of m points where the exact face intersection position on the face model surface has been defined as $p_j^{F_{ideal}} \in \mathbb{R}^3$, with $j = 1 \dots m$. The distance between the actual intersection point $p_i^F(\vec{\theta}, \Gamma_S)$ and the ideal intersection point $p_j^{F_{ideal}}$ is found from $d_i^{cs}(\vec{\theta}, \Gamma_S) = \|p_i^F(\vec{\theta}, \Gamma_S) - p_j^{F_{ideal}}\|$, with D^{cs} denoting all the d_i^{cs} distances. The combined likelihood is then:

$$P_{cs}(D^{cs}|\vec{\theta}) = \prod_{i=1}^m \mathcal{N}(d_i^{cs}(\vec{\theta}, \Gamma_S), \sigma^2) \quad (10)$$

with σ^2 set to one.

4.2.2 Approximating the Probabilistic Model

The distribution (Equation 5) is complicated and cannot be obtained in analytic form. However, as shown in the

previous section, we can evaluate the unnormalized distribution point-wise; i.e., we can compute the unnormalized density for any value of the parameter $\vec{\theta}$. This allows us to sample from the distribution using the well-known Metropolis-Hastings (MH) algorithm [10]. To use the MH-algorithm we need to define a proposal distribution $Q(\vec{\theta}'|\vec{\theta})$ from which we know how to draw samples. By filtering the samples in an acceptance/rejection step, the MH-algorithm yields samples that follow the desired probability distribution (Equation 5). As a proposal distribution, we choose random walk proposals; i.e., our proposals are all of the form $Q(\vec{\theta}'|\vec{\theta}) = \mathcal{N}(0, \Sigma)$. More precisely, we choose independent random walk proposals for the shape, rotation and translation parameters. These individual proposals are combined into a mixture proposal for the full parameter vector

$$Q(\vec{\theta}'|\vec{\theta}) = z_\alpha Q(\vec{\alpha}'|\vec{\alpha}) + z_t Q(\vec{t}'|\vec{t}) + z_R Q(\vec{R}'|\vec{R}) \quad (11)$$

where z_α, z_t, z_R are weighting factors of the different proposal distributions with $z_\alpha + z_t + z_R = 1$. The standard deviations for each of the proposal distributions have been scaled to have 20% acceptance rate ($\Sigma_\alpha = 0.1, \Sigma_t = 0.2, \Sigma_R = 0.005$).

5. Experiments

We evaluate the model in a recognition experiment. We show how the face distribution from the posterior face distribution varies for a single skull. We compare different sources of face shapes with respect to a verification performance and evaluate the influence of higher modes in the face shape. Furthermore, we show how the face-skull model can be used as prior knowledge to segment the skull from an MR-image. The experiments are performed with the statistical shape modelling library Scalismo¹ as introduced in [15].

5.1. Data

We perform the experiments on different types of face shapes obtained from different sources.




Figure 7. The ground-truth skull segmentation of skull # 2, 5, 7, 9.

MRI data. We use a small test set of 9 MR-images from where we segment and register the skull and the face in a semi-automatic way. The face is segmented by image intensity threshold and registered by fitting the BFM to the

¹<https://github.com/unibas-gravis/scalismo>

segmentation. The skull is segmented by multi-level thresholding according to [7] and registered by fitting a posterior skull-model (based on 20 landmarks) to the segmentation as described in [16]. A subset of the final registered skulls are shown in Figure 7.

3D surface scans. We also have high-resolution face scans available for all the skulls as well as a large face dataset of almost 300 people not belonging to any of the skulls.

2D images. We add faces reconstructed from 2d face images as an additional face shape source. This allows us to investigate the usefulness of the face-skull model in real-world applications such as identifying skulls from passport photographs. We use the MCMC based Morphable Model parameter reconstruction method [20] to obtain the face shape. Morphable Models [1] are statistical shape and texture models. Additional to a shape, they model surface texture as a linear combination of texture basis functions learned from data. They parameterize shape, texture, pose and illumination of an object such as a face. We use the BFM to reconstruct the face shape from single images. We create an artificial dataset of images that imitate passport photographs by rendering textured face scans. We reflect the passport setting by rendering them in frontal pose and a uniform illumination. The camera parameters and the distance of the face to the camera is fixed for all scans. The artificially created images allow us to test whether we can identify skulls from passport images in an optimal environment. Figure 8 shows example passport renderings, face reconstructions from 2d images, face scans, and the shape reconstruction error.




Figure 8. Face reconstruction from a photo. From left to right: ground-truth photo, face shape reconstructed from the photo, ground-truth 3D scan, and reconstruction error between the ground-truth scan and the ground truth face scan.

5.2. Identification: Matching Faces to Skulls

In the following experiments, we want to identify the face that belongs to a given skull from a set of face shapes. We will compare the identification performance of different sources of 3D face shape datasets as mentioned above.

Ranking experiment. In case of a natural disaster or a crime, unknown remains have to be matched to a list of target identities. Different types of information about the target identities can be known, such as DNA or photographs. Ranking the list according to how likely they are to match to a specific skull will speed up the process and saves resources because fewer DNA comparisons have to be done until a match has been found.

Given a skull to identify, we rank target face shapes by computing their probability under the posterior face-skull model.




Figure 9. Random samples from the posterior model conditioned on skull #5.

Alignment problem. To compute the probability of a face shape under the posterior model correctly, the posterior face model should only allow shape variability and not pose or translation. Because the face-skull model is aligned with the skull shapes and not the face, there is translation and rotation variability. We remove this unwanted variance, by constructing a new face model based on the posterior face-skull model, where the shapes are aligned to a reference face shape. The resulting statistical face model can be used to compute the correct probabilities of face shapes. To compute the probabilities of the target faces, we align them to the same reference. This allows us to rank the target face shapes. Figure 9 shows random samples of the posterior model of skull #5. We observe that the faces from the posterior face distribution look different from one another but are similar in shape.

Performance Measure. We measure the performance of the model by computing the average rank of the correct match. For every skull, we know the correct matching face shape. We compare the three face shape sources with respect to the ranking performance. The ranking results are summarised in Table 1 where we show that the face-skull model on average ranks the correct match within the high-

est 30% across the different datasets with 9 target skulls. Figure 11 and 12 visualise the top and bottom 3 ranked in the 9 face ranking experiments from face scans. The ranking result is consistent across the different data sources and dataset sizes as shown by the normalized average ranking.

Experiment	μ	μ norm	σ	Min	Max
MRI (9)	2.44	0.27	1.67	1	5
Scan (9)	2.89	0.32	1.54	1	5
Photo (9)	3.00	0.33	1.80	1	6
Scan (306)	91.89	0.30	44.23	26	159
Photo (106)	31.56	0.29	27.03	4	82

Table 1. Overall ranking statistics showing mean, normalized mean, standard deviation, min and max ranking of the 5 experiments with different face shape sources and sizes. The number of faces in the dataset is mentioned in parenthesis next to each experiment.




Figure 10. The number of principal components and the average ranking of the correct match for face shapes from different target sources.

5.3. Model evaluation

In this experiment, we answer the question if face size is sufficient to identify skulls from faces to obtain better than random rankings. The amount of details that can be expressed with a statistical shape model depends on the number of principal components used. For example, width and height of a shape are mostly contained in the first components. Therefore, we restrict the model to only the first n ranks and measure the identification performance. We restrict the model to the first 1,2,5,10,20,50, and 100 principal components and compare the average ranks of the correct matches. Figure 10 shows the results. We see that around 50 components are necessary to obtain the best-ranking performance from the model. It shows that size is not sufficient to obtain good results, as is reflected in the first two data points. Face details have to be taken into account too. We can see, that the performance deteriorates at 100 components. We hypothesise, that using too many components

introduces noise to the model which results in worse rankings.




Figure 11. Top and bottom 3 ranking for skull #5. The correct face is found as the second most likely face out of 9.




Figure 12. Top and bottom 3 ranking for skull #9. The correct face is found as the most likely face out of 9.

5.4. MRI skull segmentation

In this section we quickly outline another possible application of the model. Instead of using the face-skull model to obtain a distribution over face shapes, we obtain a distribution over plausible skull shapes by conditioning on a face shape. An immediate application for this is the segmentation of the skull from an MR-image. Segmenting the skull in an MR-image is a difficult task as the bone has the same intensity as air and consequently cannot be distinguished. Therefore, any successful segmentation method necessarily needs to include a strong shape prior. In contrast, segmentation of the face from an MR-image can be done using simple threshold segmentation. Hence, by conditioning on the known face shape, we can obtain a shape model that only assigns a high probability to skulls which fit the given face. We have illustrated this approach in Figure 13, where we




Figure 13. Mean and the first two principal components of the face-skull model conditioned on a face in an MR-image. Blue: ground-truth face, green: face-skull model conditioned on the face, red: skull from the face-skull model. Notice the skull from the face-skull model (red) in the back and on top of the head. The model is very good in describing the correct placement of the skull over the 2 components.

show the mean and the first two principal components (± 3 standard deviations). We observe that while the model still has the flexibility to adapt to the image, it already constrains the solution space well, and a model based segmentation algorithm for this task would be likely to benefit from such additional prior knowledge.

6. Discussion

The ranking results summarised in Table 1, clearly show the potential of the joint face-skull model as a tool for face identification based on the skull shape. The model is able to identify the correct face on average within the top 30%, even in a dataset of over 300 face scans. Such a dataset ranking would especially be suited for Disaster Victim Identification to make a rough initial sorting of individuals based on standardised photos. The whole pipeline from individual shape models to the ranking experiment concerns multiple steps with room for improvement, which has not been the focus of this paper.

MRI skull segmentation: the skull segmentation relies heavily on the prior skull model. The different skulls in Figure 7 all seem too smooth and might omit important details. The leftmost skull also seems unnaturally long (even though it overlaps with the skull in the MRI). Fortunately, the prediction does correctly find a longer male face from the face dataset as the most likely and has the smaller female face as the least likely. We might be able to use the face-skull model to improve the skull segmentation as shown in Figure 13 and hopefully get a better ranking.

Tissue-vector direction: the tissue-vector directions as projected from the skull in Figure 3 is another source of possible improvement. The directions are described and illustrated loosely in [24]. If for instance, the vectors around the nose point slightly inward instead of outwards, our formulation will favour faces with unnaturally narrow noses as the vectors have to intersect the face mesh within the defined tissue-vector depth. The ideal scenario would be to obtain the tissue-vector depths and directions from a large set of

CT or MR-images where it is possible to register the correct directions and measure the tissue-depth more densely.

7. Conclusion

We have presented a method for constructing a joint face-skull model by co-registering two individual statistical shape models of the skull and the face. The main novelty of our method is that it does not require a paired dataset of head images, but can construct possible face-skull pairs using readily available statistics about the soft tissue depth at individual points. From a computational point of view, our methods work by sampling possible faces from a posterior distribution using the Metropolis-Hastings algorithm. The generated face-skull pairs are then used to build the joint face-skull model. We have shown different applications of our model. First, we can simply generate plausible face shapes that fit a given skull. This is a task of great interest in forensics. Our experiments also showed that, given a skull, the space of possible face shapes is large. This implies that current approaches of facial reconstruction, which only predict a single face for a given skull, exclude a lot of important information. We have also provided a quantitative evaluation of the predictive power of our model. We have used the model to rank possible faces according to how likely the faces fit a given skull. Our results have shown that using this method we can on average exclude 2/3 of possible faces from a dataset of several hundred faces. This can have important applications in scenarios such as Disaster Victim Identification. As another application, we have also mentioned the possibility to use our model for model-based segmentation of the skull in MR-images. We believe, that our approach of combining individual statistical models is quite general and that the same principle can be used to join other shape models, in order to obtain flexible shape priors of different combined structures.

References

- [1] V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces. *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999.
- [2] P. Claes, D. Vandermeulen, S. De Greef, G. Willems, J. G. Clement, and P. Suetens. Computerized craniofacial reconstruction: conceptual framework and review. *Forensic science international*, 201(1):138–145, 2010.
- [3] P. Claes, D. Vandermeulen, S. De Greef, G. Willems, and P. Suetens. Craniofacial reconstruction using a combined statistical model of face shape and soft tissue depths: methodology and validation. *Forensic science international*, 159:S147–S158, 2006.
- [4] M. De Buhan and C. Nardoni. A mesh deformation based approach for digital facial reconstruction. *hal-01402129*, 2016.
- [5] S. De Greef, P. Claes, D. Vandermeulen, W. Mollemans, P. Suetens, and G. Willems. Large-scale in-vivo caucasian facial soft tissue thickness database for craniofacial reconstruction. *Forensic science international*, 159:S126–S146, 2006.
- [6] Q. Deng, M. Zhou, Z. Wu, W. Shui, Y. Ji, X. Wang, C. Y. J. Liu, Y. Huang, and H. Jiang. A regional method for craniofacial reconstruction based on coordinate adjustments and a new fusion strategy. *Forensic science international*, 259:19–31, 2016.
- [7] B. Dogdas, D. W. Shattuck, and R. M. Leahy. Segmentation of skull and scalp in 3-d human mri using mathematical morphology. *Human brain mapping*, 26(4):273–285, 2005.
- [8] F. Duan, D. Huang, Y. Tian, K. Lu, Z. Wu, and M. Zhou. 3d face reconstruction from skull by regression modeling in shape parameter spaces. *Neurocomputing*, 151:674–682, 2015.
- [9] F. Duan, S. Yang, D. Huang, Y. Hu, Z. Wu, and M. Zhou. Craniofacial reconstruction based on multi-linear subspace analysis. *Multimedia Tools and Applications*, 73(2):809–823, 2014.
- [10] W. R. Gilks, S. Richardson, and D. Spiegelhalter. *Markov chain Monte Carlo in practice*. CRC press, 1995.
- [11] Y. Hu, F. Duan, M. Zhou, Y. Sun, and B. Yin. Craniofacial reconstruction based on a hierarchical dense deformable model. *EURASIP Journal on Advances in Signal Processing*, 2012(1):217, 2012.
- [12] I. T. Jolliffe. Principal component analysis and factor analysis. In *Principal component analysis*, pages 115–128. Springer, 1986.
- [13] Y. Li, L. Chang, X. Qiao, R. Liu, and F. Duan. Craniofacial reconstruction based on least square support vector regression. In *Systems, Man and Cybernetics (SMC), 2014 IEEE International Conference on*, pages 1147–1151. IEEE, 2014.
- [14] M. Lüthi, A. Forster, T. Gerig, and T. Vetter. Shape modeling using gaussian process morphable models. *Statistical Shape and Deformation Analysis: Methods, Implementation and Applications*, page 165, 2017.
- [15] M. Lüthi, T. Gerig, C. Jud, and T. Vetter. Gaussian process morphable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [16] M. Lüthi, A. Lerch, T. Albrecht, Z. Krol, and T. Vetter. A hierarchical, multi-resolution approach for model-based skull-segmentation in mri volumes. In *Conference Proceedings D*, volume 3, 2009.
- [17] S. Nagpal, M. Singh, A. Jain, R. Singh, M. Vatsa, and A. Noore. On matching skulls to digital face images: A preliminary approach. *arXiv preprint arXiv:1710.02866*, 2017.
- [18] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, and T. Vetter. A 3d face model for pose and illumination invariant face recognition. In *Advanced video and signal based surveillance, 2009. AVSS'09. Sixth IEEE International Conference on*, pages 296–301. IEEE, 2009.
- [19] P. Paysan, M. Lüthi, T. Albrecht, A. Lerch, B. Amberg, F. Santini, and T. Vetter. Face reconstruction from skull shapes and physical attributes. In *Joint Pattern Recognition Symposium*, pages 232–241. Springer, 2009.
- [20] S. Schönborn, A. Forster, B. Egger, and T. Vetter. A monte carlo strategy to integrate detection and model-based face analysis. In *German Conference on Pattern Recognition*, pages 101–110. Springer, 2013.
- [21] S. Shrimpton, K. Daniels, S. De Greef, F. Tilotta, G. Willems, D. Vandermeulen, P. Suetens, and P. Claes. A spatially-dense regression study of facial form and tissue depth: towards an interactive tool for craniofacial reconstruction. *Forensic science international*, 234:103–110, 2014.
- [22] W. Shui, M. Zhou, S. Maddock, T. He, X. Wang, and Q. Deng. A pca-based method for determining craniofacial relationship and sexual dimorphism of facial shapes. *Computers in Biology and Medicine*, 90:33–49, 2017.
- [23] C. N. Stephan. The application of the central limit theorem and the law of large numbers to facial soft tissue depths: T-table robustness and trends since 2008. *Journal of forensic sciences*, 59(2):454–462, 2014.
- [24] C. N. Stephan and E. K. Simpson. Facial soft tissue depths in craniofacial identification (part i): an analytical review of the published adult data. *Journal of Forensic Sciences*, 53(6):1257–1272, 2008.
- [25] P. Tu, R. Book, X. Liu, N. Krahnstoever, C. Adrian, and P. Williams. Automatic face recognition from skeletal remains. In *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*, pages 1–7. IEEE, 2007.
- [26] D. H. Ubelaker. Craniofacial superimposition: Historical review and current issues. *Journal of forensic sciences*, 60(6):1412–1419, 2015.
- [27] C. Wilkinson. *Forensic facial reconstruction*. Cambridge University Press, 2004.
- [28] C. Wilkinson. Facial reconstruction—anatomical art or artistic anatomy? *Journal of anatomy*, 216(2):235–250, 2010.
- [29] Z. Xiao, J. Zhao, X. Qiao, and F. Duan. Craniofacial reconstruction using gaussian process latent variable models. In *International Conference on Computer Analysis of Images and Patterns*, pages 456–464. Springer, 2015.