Politecnico di Torino
Dottorato in Ingegneria Informatica e Dei Sistemi

Planning Plastic Surgery in 3D. An innovative approach and tool

Ph.D. Dissertation of:

Matteo De Simone

Advisor:

Prof. Andrea Bottino

XXV Cycle

Politecnico di Torino
Dottorato in Ingegneria Informatica e Dei Sistemi

Planning Plastic Surgery in 3D. An innovative approach and tool

Ph.D. Dissertation of:

Matteo De Simone

Advisor:

Prof. Andrea Bottino

XXV Cycle

Politecnico di Torino
Dipartimento di Automatica e Informatica
Corso Duca degli Abruzzi, 24, 10139 Torino

ACKNOWLEDGEMENTS

Thanks to Aldo and Andrea because they showed me that there is more than engineering in Engineering. Thanks to Flavia because she pushed me to start. Thanks to Lab1 because he shared lights and shadows of this adventure. Thanks to my friends and their beers because they helped me to not give up. Thanks to my parents because they taught me to never give up.

ABSTRACT

Face plastic surgery (PS) plays a major role in today medicine. Both for reconstructive and cosmetic surgery, achieving harmony of facial features is an important, if not the major goal. Several systems have been proposed for presenting to patient and surgeon possible outcomes of the surgical procedure. In this work, we present a new 3D system able to automatically suggest, for selected facial features as nose, chin, etc., shapes that aesthetically match the patient's face. The basic idea is suggesting shape changes aimed to approach similar but more harmonious faces. To this goal, our system compares the 3D scan of the patient with a database of scans of harmonious faces, excluding the feature to be corrected. Then, the corresponding features of the *k* most similar harmonious faces, as well as their average, are suitably pasted onto the patient's face, producing *k+1* aesthetically effective surgery simulations. The system has been fully implemented and tested. To demonstrate the system, a 3D database of harmonious faces has been collected and a number of PS treatments have been simulated. The ratings of the outcomes of the simulations, provided by panels of human judges, show that the system and the underlying idea are effective.

1 CONTENTS

2 INTRODUCTION

Plastic surgical procedures enjoy increasing popularity in different countries as well as among different social groups. According to recent statistics by The American Society for Aesthetic Plastic Surgery, in 2009 more than one million facial plastic surgeries were performed, with a growth of 151% in ten years [1]. According to an estimate, in the US more money is spent annually on beauty related items or services than on both education and social services [2]. The popularity of plastic surgery (PS) also triggered new studies in face recognition aimed at taking into account facial differences due to these procedures ([3],[4]). Facial PS is defined reconstructive when aimed at correcting deformities congenital or due to accidents or diseases, and cosmetic otherwise. Actually, there is not a well-defined boundary between the two fields. In any case, achieving as far as possible a harmonious or at least regular face shape is a main goal for both kinds of surgery.

2.1 Previous work

Planning PS requires taking into account a number of physiological and psychological constraints. The frequency of secondary rhinoplasties, due to unsatisfactory outcomes from previous procedures, and estimated around 8%-15% of the cases [5], shows that planning PS is far from simple. Fully planning PS would require considering several factors, as the exterior shape to obtain, its interaction with soft tissues and bones, and the evolution of tissues with time. This work deals with the problem of supporting the aesthetic judgment of surgeon and patient for choosing more harmonious facial shapes.

To this purpose, a number of computer tools have been proposed, aimed at presenting possible outcomes of the surgical procedure. Ozkul and Ozkul [6] presented a system for simulating rhinoplasty in 2D profile images. Image interaction tools emulating real surgical procedures were supplied. Rabi and Aarabi [7] emulated the outcomes of PS in 2D frontal images by suitably pasting on the patient's face facial features of other subjects. Liu et al.[8] attempted to predict the PS outcome on the basis of a number of results of former cases. Commercial 2D simulation packages are also available [9]. These packages essentially allow manipulating 2D images with morphing operators.

Since human face is a complex 3D object, several systems for manipulating 3D faces have been studied. Gao et al. presented a technique for warping the 3D B-spline surface obtained from patient's face scan [10]. Lee et al. [11] constructed a patient specific 3D model from a generic model based on CT data and a 2D image of the patient and used morphing operators for emulating surgical procedures. A 3D PS planning system specific of nose surgery was proposed by Lee, Lin and Lin in [12]. The system works on CT data of the head of the patient from which the face surface is obtained for realistic rendering purposes. Wang et al. [13] proposed a 3D simulator of augmentation nose surgery, based on a surface scan and a lateral X-ray image. Deformation tools were provided, as well as the evaluation of the size of the required prosthesis. Commercial 3D systems have also been recently proposed [14].

2.2 OUR APPROACH AND ITS RATIONALE

The quoted planning systems supply manual interfaces for changing the shape of 2D or 3D patient's images. All choices in manipulating faces and evaluating results are left to the surgeon's and patient's judgment.

Our approach is aimed at automatically suggesting, for the particular feature object of surgery and for the particular patient's face, the shapes most suitable to enhance facial harmony. To explain the rationale of our approach, we will review some results of the research on facial attractiveness. Recently this research has been a major issue in psychology, psychobiology, anthropology, evolutionary biology, behavioral and cognitive sciences, and in medical areas such as plastic surgery and orthodontics. Empirical rating studies have demonstrated high beauty rating congruence over ethnicity, social class, age, and sex (see for instance [15][16][17][18]). Ratings are in particular very similar for beautiful and ugly faces. This means that the human perception of facial attractiveness is essentially data-driven, and largely irrespective of the perceiver. Many papers aimed at analyzing and assessing beauty using Computer Vision and Image Processing techniques have been recently presented. Several of them have been surveyed in [19].

Although several results have been obtained, the question whether facial beauty can be synthesized into some relatively simple quantitative rules is far from being answered. Traditional rules, as the classic canons, have been found not to apply to many faces rated beautiful (see [20][21][22]). Several empirical results support the idea that there is not a unique beauty prototype. Although averaging facial images produces faces more attractive than each component image, attractive faces are mostly not average ([23][24]). Beauty ratings have been found largely independent on ethnicity, but faces rated beautiful can be rather different in various ethnic groups ([25][26]). Cognitive theories support the idea that humans define and use several category prototypes for face image analysis [27]. According to this idea, several beauty rating programs provide results close to those of human raters, based on classifiers trained with a set of already rated faces [19]. Using around 100 samples, the quality of the ratings increases without saturation with the number of samples in the training set [28], highlighting the needs for large DBs of rated faces.

Fig. 1. Different shapes could be more or less attractive, depending on the perceived general harmony

Concluding, simple rules based on measurable face features, able to assess beauty or suggest how to improve it, are still lacking, and we must learn attractiveness by examples. Therefore, our general idea for patient-specific beautification is to consider several harmonious face prototypes, and to suggest surgical procedures able to approach the samples closer to the patient's face. Observe that the idea of many different prototypes of harmonic faces implies that there is not a unique harmonic prototype of a particular facial feature (e.g. mouth, nose). Different shapes could be attractive, depending on the integration with the rest of the face. An example of rather different noses belonging to movie actors is shown in Fig. 1.

In more details, our approach to the face beautification works as follows:

1. consider a particular patient's face feature (target feature), and compare the face of the patient with a set of faces rated harmonic (reference database), excluding the target feature;
2. minimize some convenient metric in the face space and find the nearest k attractive faces;
3. blend the corresponding target feature of the k nearest attractive faces, as well as their average, with the original face, in order to produce k+1 surgery simulations; in the following these features will be referred to as target reference features (TRF).

We also use the average feature since, as previously mentioned, averaging faces has been found to improve attractiveness.

A first implementation of this idea is in [29] focused on 2D face profiles, and choosing as target feature the nose. In the 2D case the steps are applied as follow:

- patient's and reference faces' profiles are extracted and normalized;
- the nose is removed from all profiles. Specific techniques based on the analysis of the local curvature of the profile contour and on the local analysis of the face features had developed;
- profiles are compared by means of a suitable similarity measure in order to find the reference face most similar to the patient's face (*best candidate*). The similarity was described by the area resulting XORing the patient and reference profiles (Fig. 2);
- the nose of the best candidate is substituted to that of the patient with a suitable 2D morphing;

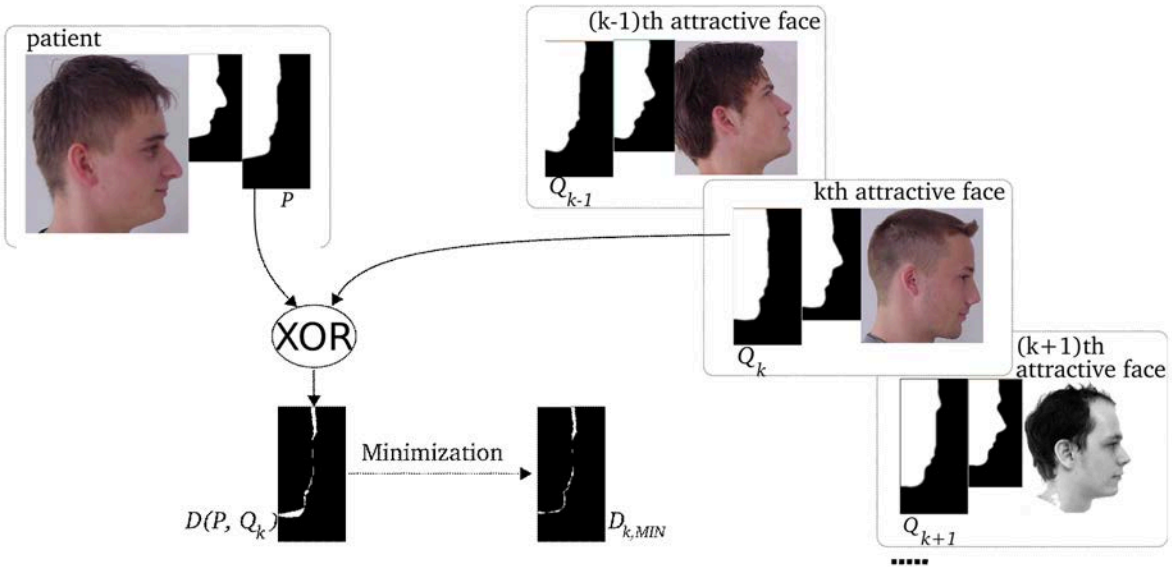

Fig. 2. 2D patient comparison with the kth subject of the database for best candidate search. The distance between the noseless patient profile P and the Q_k, *the kth* noseless profile in the reference database is $D(P, Q_k) = area(P \text{ XOR } Q_k)$

In this work, we extend the idea to 3D face scans, and to any part of the face.

The content of the thesis is as follows. In Chapter 3, we describe the 3D scans database used and we detail the various steps of the algorithm. In Chapter 4, we present examples of PS simulation. Processing times are reported in Chapter 5 and the ratings of the results of PS simulations supplied by a panel of human raters are discussed in Chapter 6. Concluding remarks are reported in Chapter 7.

In Appendix B we attach more images and results; in Appendix A we have a short description of the developed software.

3 THE PROPOSED TECHNIQUE

In this section, we describe in detail the reference database used and the various steps performed by the system, that are:

- Normalizing face scans
- Selecting and removing on the patient's face the target feature
- Determining and removing the corresponding features on the harmonious faces of the reference DB
- Retrieving the k DB faces most close to the patient's face
- Suitably pasting on the patient's face the TRF of the k closer DB faces and their average.

Underlying our technique is the assumption that facial harmony is unaffected by (slightly) non-isotropic 3D global scaling.

3.1 CREATING THE REFERENCE DB

The outlined beautification process requires a set of harmonious faces (the reference database). Unfortunately, not many 3D face databases are available, and those existing, being essentially aimed at supporting face recognition research, contain attractively average faces, and only a very few scans useful for our purposes. As mentioned before, a dense sampling of the manifold of beautiful faces in the face space is likely to require several hundreds of samples. To build such large database is beyond the scope of this paper, and we have constructed a database sufficient to demonstrate the PS simulation software implemented and the effectiveness of the general approach.

To construct the reference DB, we collected scans from two sources. First, we selected, with the help of a human raters panel, the most attractive faces (actually very few) from the available 3D databases, such as [30] and [31]. Second, we launched an initiative (named "*Faces for the Science*") aimed at acquiring high-resolution 3D attractive face scans (Fig. 3). The initiative is backed by the Human Morphology Department of the University of Milan and involves professional modeling agencies in order to exploit specific professional expertise for the attractiveness rating process. The scanning campaign collected face scans of both male and female multiethnic

subjects, age between 16 to 30 years old. Data have been acquired with the VECTRA-CR-two-pod system [32], a portable non-contact passive stereoscopic scanner that guarantees full 180-degree capture. Each 3D model has approximately 90.000 vertices with an associated 10 Megapixel resolution texture.

In total, we selected as convenient to our purposes 53 face scans of both male and female multiethnic subjects. Due to the lack of a sufficient amount of male samples we finally decided to create a reference DB containing only female scans (36 samples) and, therefore, to use merely females as subjects of the PS simulation procedure.

3.2 SCAN NORMALIZATION

The 3D face scans used for constructing the current and future reference database come from different sources, and are different for size, orientation, resolution, type of meshes, etc. (Fig. 3). Therefore, for obtaining meaningful comparisons, mesh resolution and geometric normalization are applied.

First, since many 3D scans present various imperfections, as missing polygons or spikes, mesh regularization is performed. In addition, faces have been cropped eliminating hairs and clothes (Fig. 4). Both cleaning and cropping have been performed manually with Blender [33].

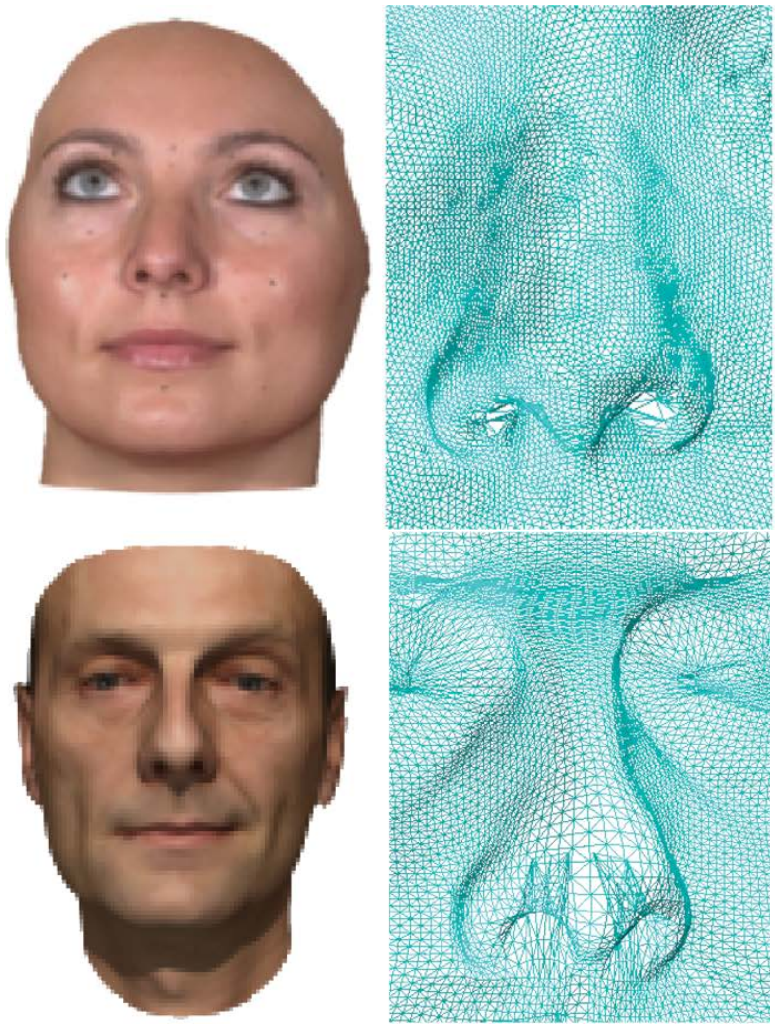

Fig. 3. Examples of face scans (rendered view and wireframe detail) from different databases

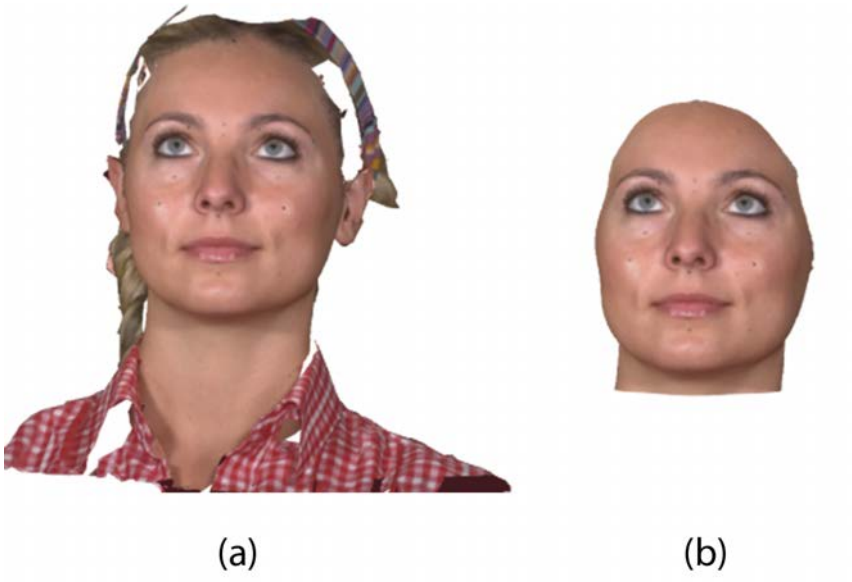

Fig. 4. An example example of scan before (a) and after (b) cleaning and cropping

Our normalization algorithm performs the following steps.

3.2.1 Orientation and size normalization

Geometric normalization of 2D face images is mostly obtained by making coincident the centers of the eyes using suitable translation, rotation and 2D isotropic scaling. In our 3D case, normalization involves isotropic 3D scaling after a first rough registration using the sagittal plane and two feature points.

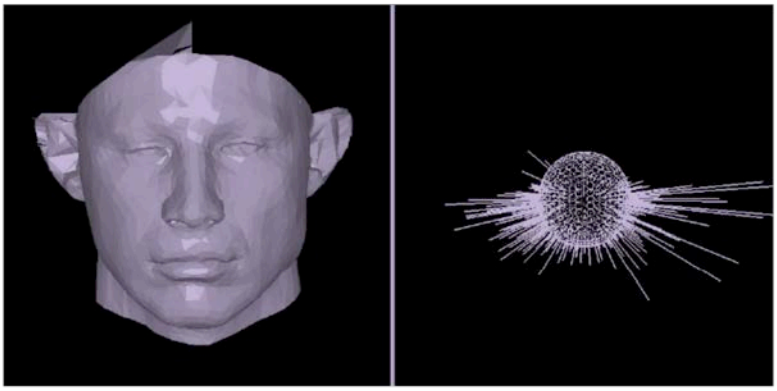

Fig. 5. (a) a face scan and (b) the correspondent EGI as in [34]

The sagittal plane is determined by the algorithm described in [34], exploiting the Extended Gaussian Image (EGI). EGI [41] is a 3D shape which is a function defined on the unit sphere (Gauss sphere [42]). Following [43], two objects with the same EGI are congruent so the problem to find the sagittal plane of an object could be transformed in finding the sagittal plane of its EGI. This method is hard to apply on objects with noise data (as 3d scan are). Anyway the approach suggested by [34] offers

a specific solution for face scans where the symmetry is not so rigorous. In Fig. 5 you can see a face scan (a) and the correspondent EGI (b)

When we have the sagittal plane, we can intersect it with the face scan in order to determine its profile, where two facial landmarks, the nasion and subnasale are identified (see Fig. 6). These landmarks are computed with the algorithm described in [35] and it is based on the local curvature of the profile contour and on the local analysis of the face features.

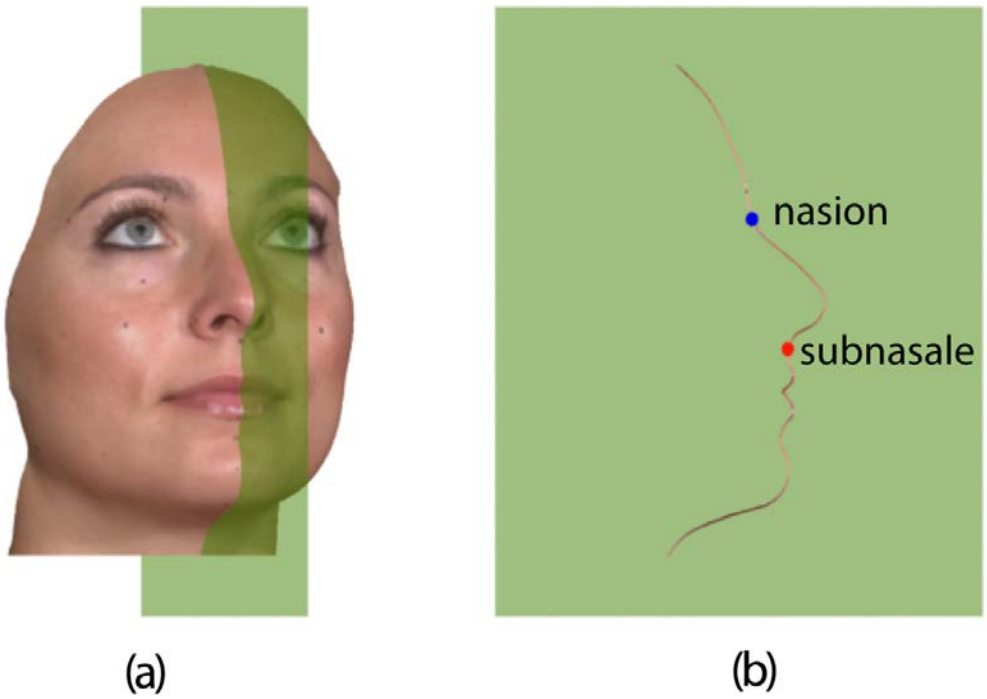

Fig. 6. Intersection of a face scan with its sagittal plane and automatic identification of landmarks on face profile

Then, the sagittal plane and the two landmarks are made coincident for all the 3D scans by translation, rotation and scaling. Finally, a further finer registration with an intermediate face (a face of the reference DB chosen at random, see section 3.4 for

details) is performed using the Iterative Closest Point algorithm (ICP, [36]) using the isotropic scaling software of the VTK library [37]. At the end of the registration process, all scans lie in a common reference system, and have common sagittal plane and reference axis on the sagittal plane, passing through the centroid of the intermediate face and oriented as described in [34].

3.2.2 Mesh resolution normalization

Reference DB samples and patient's scan could have different resolution. Mesh normalization is required by one of the following steps of the algorithm, where a dense point-to-point correspondence between different face scans is determined. Different resolutions can be used in the re-sampling process in order to balance the precision of the reconstructions and the computational costs.

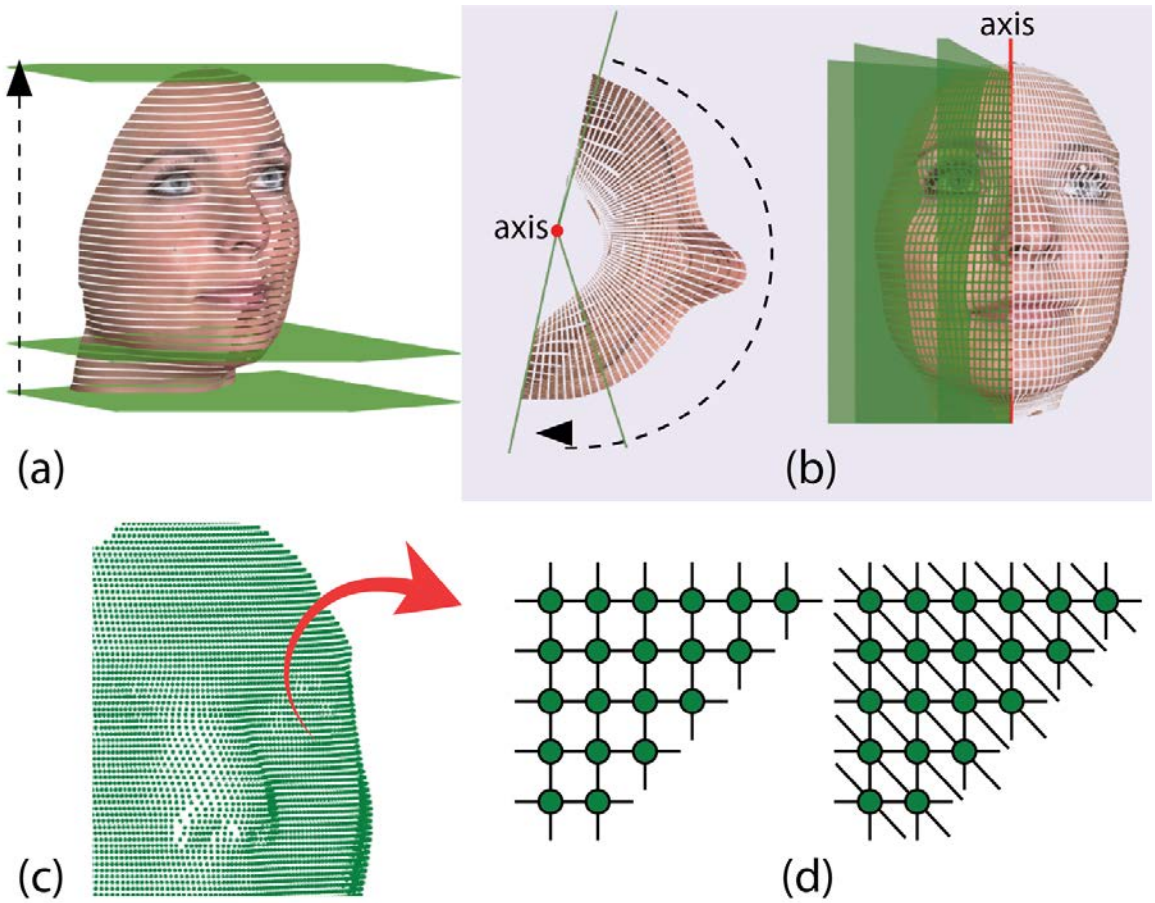

Fig. 7. Re-sampling process: (a) the mesh is sliced with horizontal planes obtaining a number of polylines; (b) the polylines are intersected with vertical planes rotated around the vertical axis. (c) The final point cloud; (d) reconstruction of the mesh connectivity

To normalize resolution, face scans are re-sampled. As shown in Fig. 7(a), each face mesh is first intersected with H horizontal planes equally spaced within a reference bounding box. Then, the obtained polylines are intersected with a set of V vertical planes (Fig. 7 (b)), passing through the reference axis, common for all aligned faces, and equally spaced angularly. So, each point, is uniquely defined by the pair (h, v) where $0 \leq h \leq H$ and $0 \leq v \leq V$. The resolution of the final meshes can be adjusted changing the H and V values. For the experiments described in chapter 4, and submitted to human raters (chapter 6), both variables were assigned the value 776,

which produced mesh resolutions near to 300K vertices. For the results presented in chapter 5 (Processing times), we also used two lower resolutions (160K and 40K).

The result of re-sampling is the point cloud shown in Fig. 7(c). Initially, we tested the 3D Delaunay triangulation algorithm [48] to derive the mesh from the point cloud but that technique introduces changes in topology. The final triangular mesh is determined, as shown in Fig. 7(d), by forming quads and splitting each of them in two triangles. This is a trivial process due the slicing process: each point p is described by the pair (h, v) so the neighbours are identified by pairs $(h-1, v)$, $(h+1, v)$, $(h, v-1)$ and $(h, v+1)$

Texture coordinates for each vertex of the new 3D mesh are computed with bilinear interpolation of those of the neighboring points in the original mesh.

3.3 SELECTING THE PATIENT'S TARGET FEATURE

Any area of the scanned patient's face can be selected by the user as target of the PS simulation. An interactive tool allows drawing on the scanned face a 3D polyline, with vertices located at the re-sampled mesh points. The polyline can be more or less detailed according to the particular target feature and to the judgment of the user. In Fig. 8 an example is shown relative to the nose area. A detail view of an example of polyline and facial triangulation is shown in Fig. 9.

The feature cutting is performed by intersecting the facial mesh with a cylinder whose generatrices pass through the polyline selected and parallel to a direction evaluated from the polyline by the Newell's Method [38].

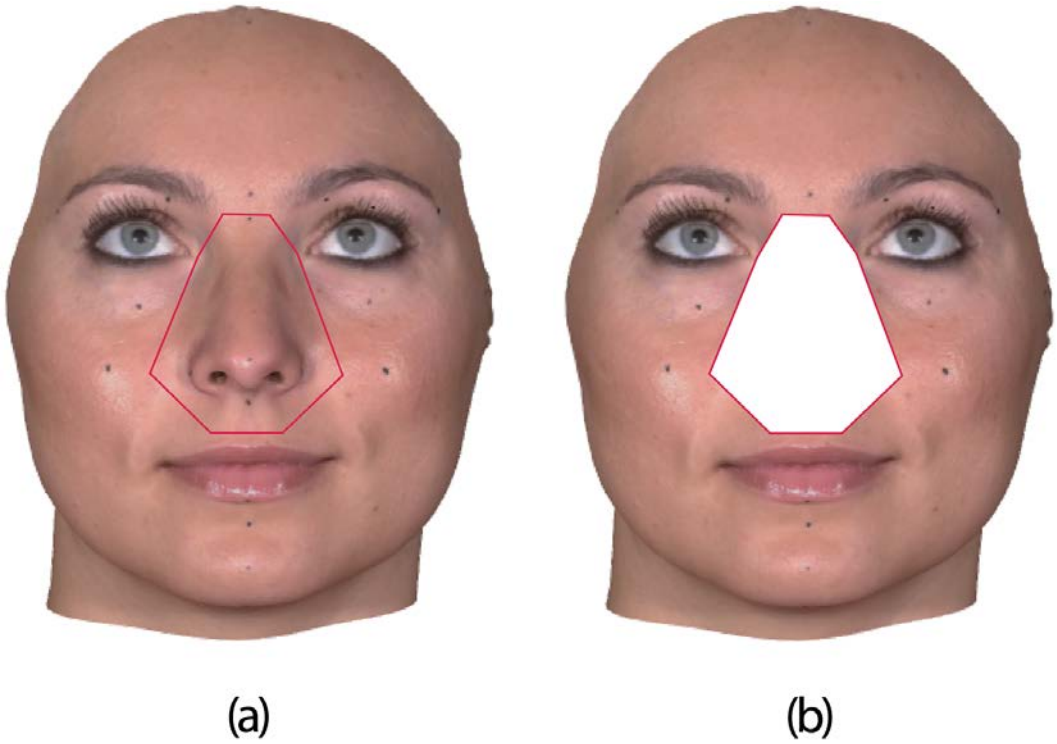

Fig. 8. Interactive feature selection: contour drawing (a); removal of the feature (b).

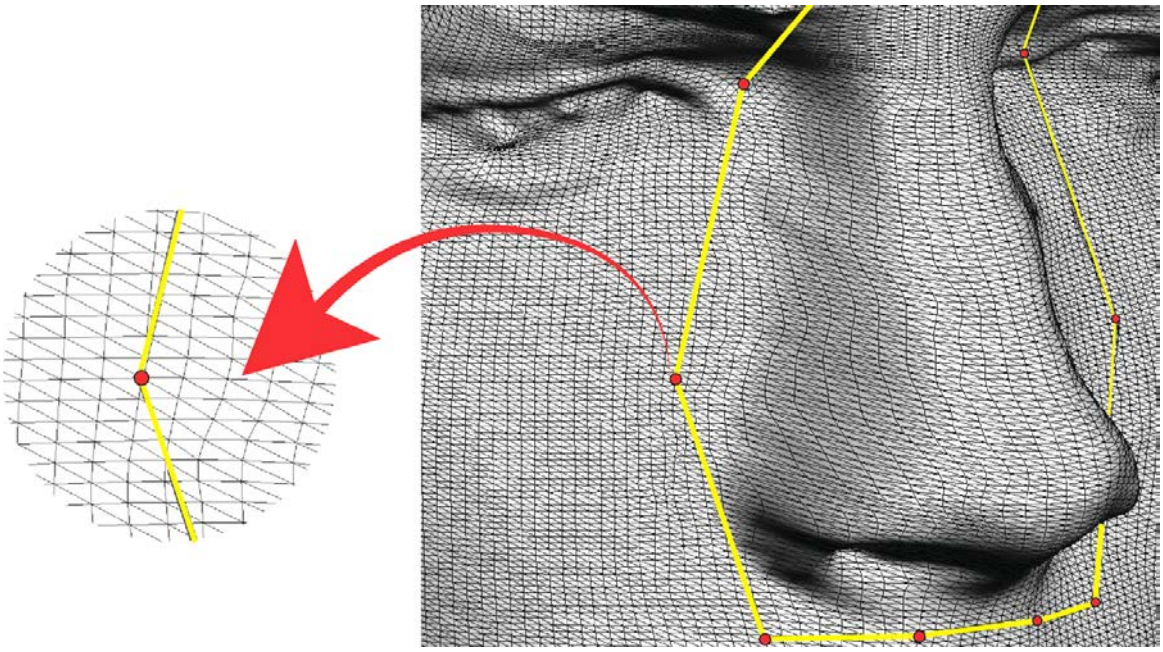

Fig. 9. A detail view of the polyline delimitating the area subject of the PS procedure (red dots are vertices of both the polyline and the face mesh)

3.4 REMOVING THE TARGET FEATURE FROM ALL THE REFERENCE FACES

Our algorithm requires removing the surface corresponding to the target feature both from the patient's face scan and from all the scans of the reference database.

The first operation is easily performed by intersecting the facial mesh with a cylinder in the same way described in 3.3

Removing the same feature from the scans of the reference database requires identifying a corresponding polyline on each face. To do this, we must first establish a dense point-to-point correspondence between the patient's face and each reference face. To this purpose, implemented two techniques.

The first approaches the problem using the method proposed by Blanz and Vetter [44] for building morphable 3D face model. First, a 3D model S is transformed into a four-channel 512x512 2D *image* $I(h,v)=\{R(h,v), G(h,v), B(h,v), d(h,v)\}$, where (h,v) are the coordinates of a surface point as resulting form the resampling process, $R(h,v)$, $G(h,v)$, $B(h,v)$ its texture values and $d(h,v)$ its distance from the vertical axis (Fig. 10). Then, the correspondences between the points of two scans S_1 and S_2 can be computed applying a modified optical flow algorithm to their transformations I_1 and I_2. The Thirion's "daemon" registration algorithm ([45],[46]) is used for optical flow evaluation.

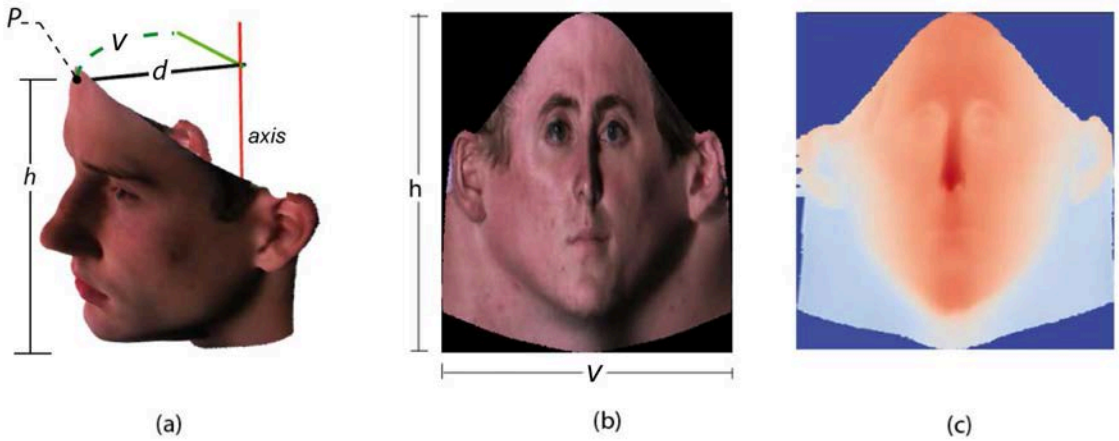

Fig. 10. The surface point P corresponding to (h,v) coordinates (a). The 512x512 RGB (b) and d channels of the 2D transformation applied to the model shown in (a).

We realized that this approach was not perfectly suitable for our algorithm because it is strictly related to textures. For example it is difficult to find a good match in presence of deformed or injured faces and the user has not an easy way to intervene to fix wrong results. Because we'd like to design a general approach suitable as for aesthetic as for reconstructive procedures, we decided to find a different way to conserve generality and freedom.

So we introduced the second technique that it is similar to that proposed in [39]. First, faces are registered with a thin plate spline transformation (TPS), a non-rigid smooth surface transformation particularly suited for morphing or animating faces, based on the correspondence of a number of keypoints. We used 37 pairs of corresponding landmarks manually selected on both the source and destination scans (see Fig. 11). These keypoints are a subset of the standard landmarks used in Anthropometry [40] with additional points in eyes area. Additional points are added to

handle regions where scanners typically generate noise (as in eye area). It is possible add more point to handle very peculiar situations as facial bone deformation or injuries.

TPS provides a smoothly interpolated mapping between source and destination scans, which makes the landmarks coincident and brings the other points of the deformed source face very close to the destination face. This allows an accurate alignment of points in smooth regions, like cheeks and forehead. Then, after TPS morphing, for each vertex of the first scan a correspondence with a vertex of the second scan is derived with a closest point searching algorithm that exploits a KD-tree data structure to ensure computational efficiency.

Consequently, a new polyline, equivalent to the polyline drawn on the patient's scan, is determined for each reference face. An example of the transfer of the feature boundary on the reference DB faces is shown in Fig. 12. Once the boundary polyline has been determined on all reference faces, the target feature is easily removed from them, as done for the patient's face.

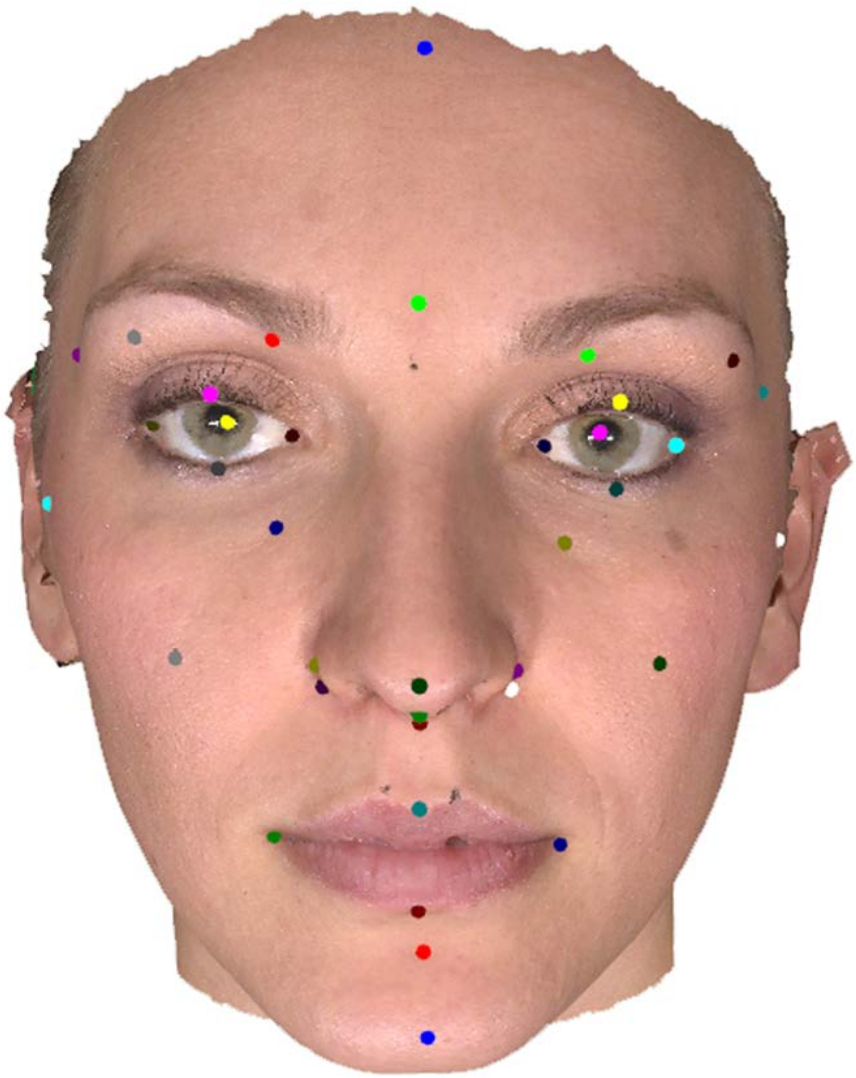

Fig. 11. The keypoints used for establishing dense point-to-point correspondence between faces

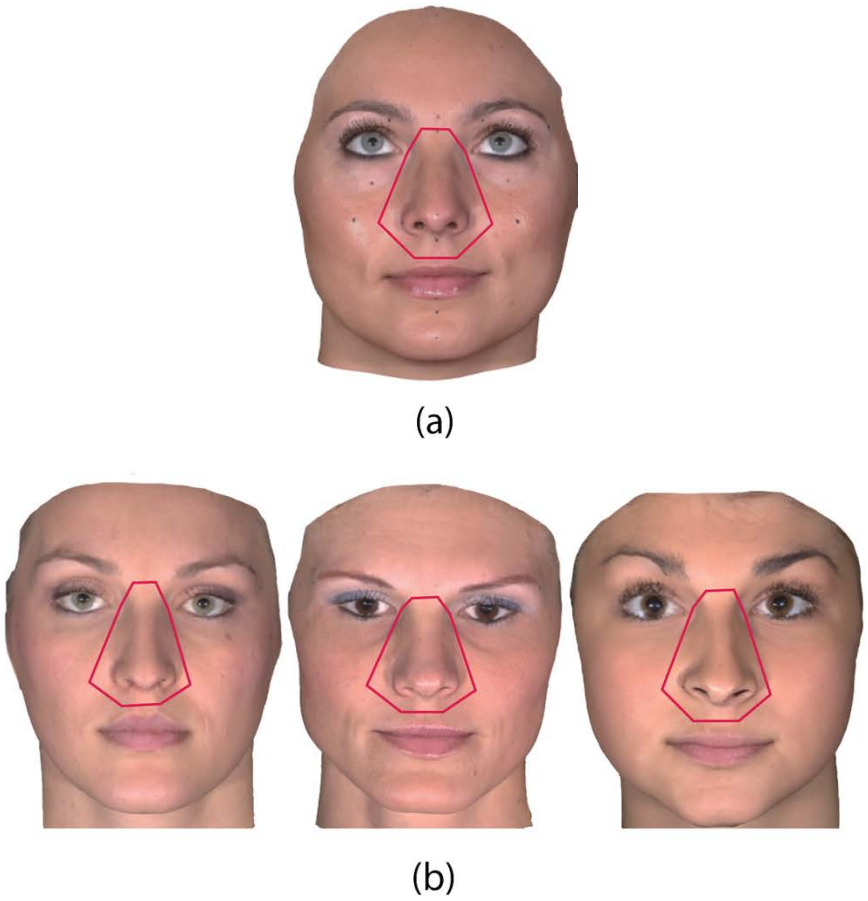

Fig. 12. An example of transferring a feature boundary from the patient's face (a) to the scans in the reference DB (b)

Finding the correspondences between two scans is a time consuming process. In principle, each new patient's scan would require computing the correspondences with all the reference faces. As mentioned before, fully sampling the attractive face manifold could require several hundreds of samples, resulting in long computation times. To avoid this, we pre-computed the correspondences between an intermediate face (a face randomly chosen) and all the reference faces. In this way, for each new patient's face, we have only to compute the correspondences with the intermediate face (see Fig. 13).

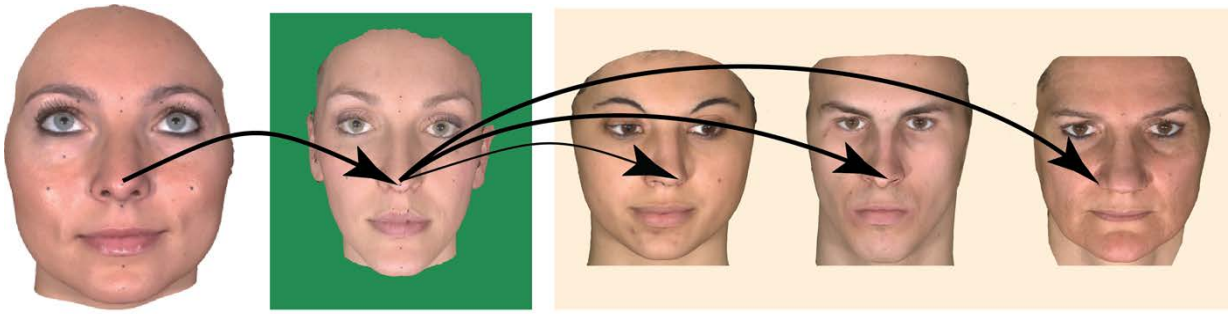

Fig. 13. A point of the input face (a) corresponds to a point on the intermediate face (b) that, in turn, corresponds to a point on each face in the reference DB (c)

Observe that, given the discrete nature of the correspondence, two vertices of the first scan could collapse in only one vertex in the second and there might be vertices of the destination scan which have no correspondences in the source scan. Therefore, the correspondence function is not bijective. However, bidirectional correspondences are necessary for our algorithm, since in the last step we also construct a mean feature, obtained from those of a set of reference faces by averaging their corresponding points. Therefore, computing the mean feature requires correspondences between all reference faces, which can be obtained evaluating the backward correspondences from each reference face to the intermediate face. These correspondences can be pre-computed as well.

3.5 RETRIEVING THE MOST SIMILAR REFERENCE FACES

After the target feature removal, reference samples are sorted in order of similarity with the patient's scan. The similarity measure between two faces is given by their RMS distance, which is the square root of the average of the sum of squares of the closest point distances, after their alignment with ICP using anisotropic scaling.

As mentioned in the introduction, average faces have often been found to be more attractive then each component face. For this reason, we also computed an average feature from the first *k*-nearest neighbors in the following way. Starting from F_1, the closest face, we evaluate the (weighted) mean of each vertex of its TRF and its corresponding points on the remaining *k*-1 faces. The correspondences between the target and the average features that are required for pasting a TRF on the patient's face, are easily obtained, being equal to those between target feature and TRF of F_1.

As for the weights, we experimented with both a uniform and a weighted mean, based on the inverse RMS distances. However, since a panel of human raters did not express significant preferences for either option, we used equal weights.

3.6 BLENDING THE TRFS

This final step produces the simulated surgery outcome, that is the 3D textured model of the patient's face where the surface of the target feature has been modified according to the surface of the TRF retrieved, and the original patient's texture has been conveniently applied on the new surface. This process should produce a smooth, realistic and seamless junction.

The process consists of three steps:

a) the surface of the TRF is registered with respect to the patient's scan;

b) the original target feature (surface and texture) is conveniently morphed toward the TRF surface. A feature of this particular morphing is that it produces a smooth blending of patient's scan and TRF surface;

c) the original texture intensity is corrected according to the different new shape.

In more detail:

a) Registration. An initial registration of the TRF on the patient's face is obtained applying the same transformations used to match patient and corresponding face of the reference DB. This registration is refined to reduce the differences between the contours of the removed target

feature and the TRF, by applying a second anisotropic ICP only to the points of their contours. The resulting transformation is applied to the entire mesh of TRF.

b) Morphing. The morphing process is schematically shown in (Fig. 14). The general idea is to use the dense correspondence computed between target feature and TRF, excluding a strip close to the contour, where the original surface and texture is morphed to an intermediate position, depending on the distance from the feature contour.

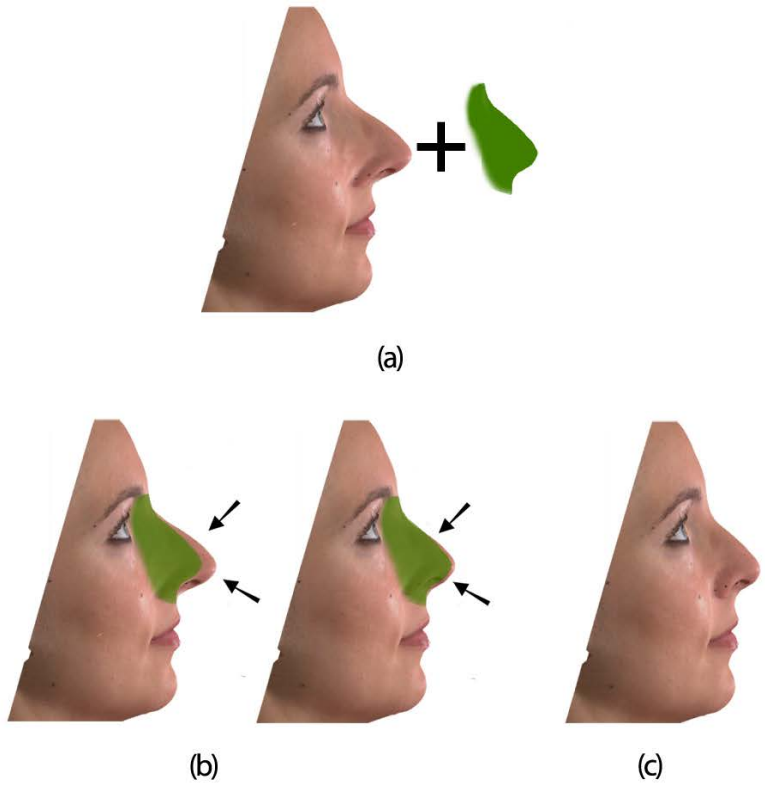

Fig. 14. The reference feature (a, in green) is used to morph the target feature (b) to obtain the final result (c)

In order to describe the details of the process, let us refer to Fig. 15, where we show in (a) the target feature (red, on the left), the TRF (gray, on the right) and in Fig. 15(b) their registration. For each point P_i of the target feature:

1. we compute c_i, its distance from the closest point on its contour, and D_i, its vector difference with its corresponding point on the transformed TRF (Fig. 15c);
2. we transform P_i (Fig. 15d) according to the following equation:

   $$P_i \leftarrow P_i + W(P_i) * D_i$$

   where $W(P_i)$ is a scalar weight in the interval [0,1] obtained as:

   $$W(P_i) = \begin{cases} c_i / (c_{\max} * t) & if\ c_i / c_{\max} < t \\ 1 & otherwise \end{cases}$$

   c_{max} is the maximal c value for all target feature points and t is a threshold (experimentally set to 0.1) that controls the width of the blending strip.

Fig. 15(e) shows, in false colors, a map of the weights applied to the target feature points, ranging from 1 (blue) to 0 (red). The green is the area of blending. As it can be seen, points closer to the patient feature boundaries are smoothly interpolated with that of the reference feature. An example of transformation of a point is shown in (c) and (d). In this case, the distance c_i is such that the point is outside the blending

strip, and the point sticks directly on the corresponding point. For computing $W(P_i)$ we used the Euclidean distance between P_i and the feature contour. A better option could have been the geodetic distance, however too expensive to compute. The geodetic distance is the length of the shortest path between two points on the mesh grid. It implies that first we have to find a shortest path and after we have to evaluate distances between each consecutive point. Furthermore, our tests show that the Euclidian distance is a good enough approximation for our purpose. This reduces the estimation of $W(P_i)$ just to evaluate a distance between two points.

c) Texture correction. The texture applied on the new shape corresponds to the original skin of the target features, but in general is not in full agreement with the new geometry, since the reflected light depends on surface normals, more or less changed. Exactly reconstructing the illumination for the new surfaces is not easy, since it would require detailed information about lights and surface materials, which is not available in general.

Actually, in most images of simulated PS the effects of this problem are not conspicuous. Anyway, we implemented an approximate algorithm providing a texture correction that appears effective in our case.

Assuming that the reflective features of the surface material (skin) is the same for all the feature points and a Lambertian reflection model, the intensity of a surface point depends approximately only on the direction of its normal.

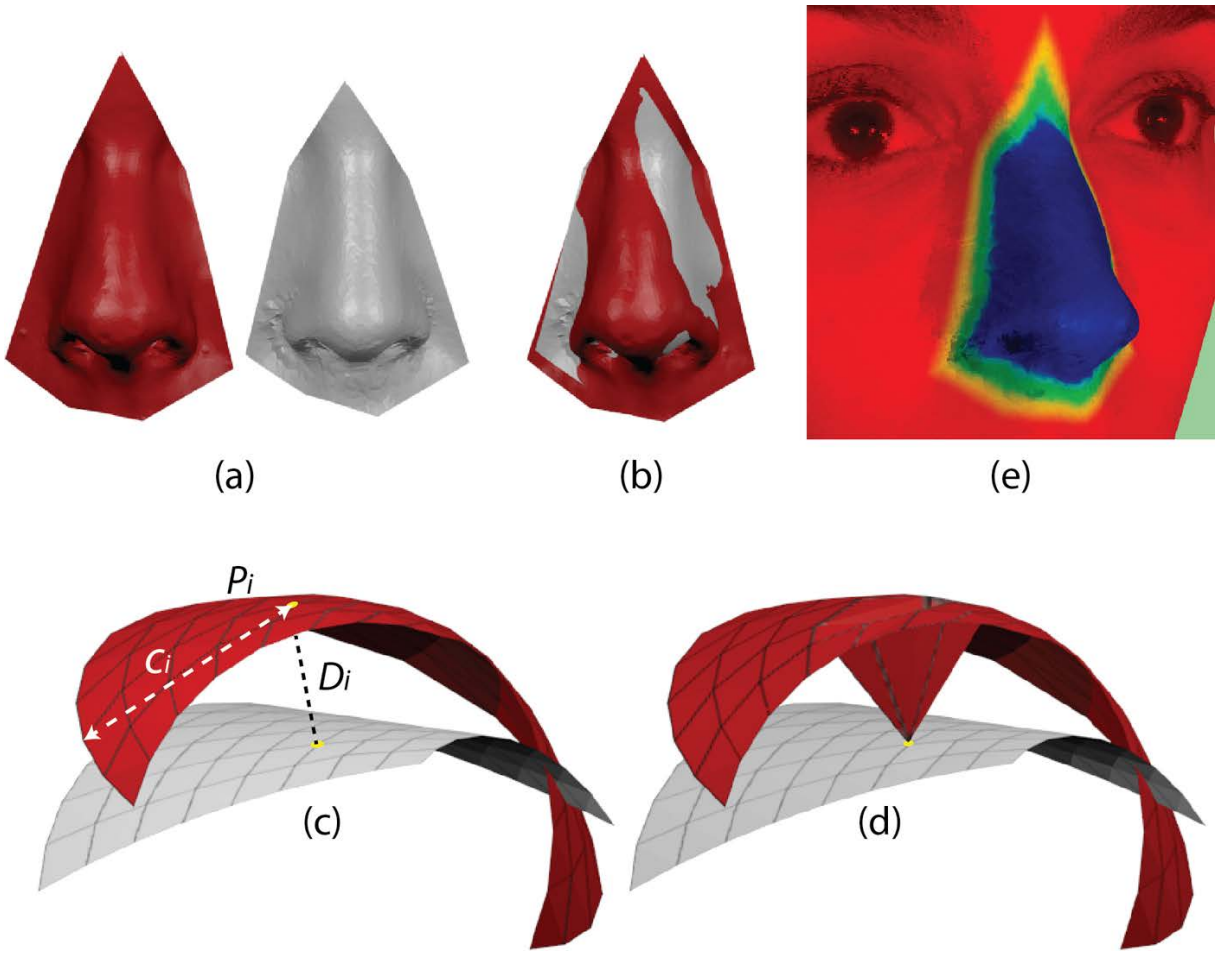

Fig. 15. a) target feature (left) and TRF (right); b) features alignment; c, d) evaluation of the weight of point P_i (with $W(P_i)=1$) and its transformation; e) a false color map of the weights $W(P_i)$ (red = 0, blue = 1)

Therefore, we assign to the texture of the morphed target feature the intensity of the texture of the original feature where the surface had the same normal. Next chapter explains in more detail how this is performed.

3.7 ADJUSTING THE TEXTURE

The previous steps described how the patient's scan can be modified to obtain a possible outcome of the surgical planning process changing the geometry of the target feature according to that of some convenient reference samples. However, this step is not always sufficient and also the texture might be adjusted, since the colors of its pixels have been obtained as a product of the interaction of incoming lights, surface materials and surface normals and, during the feature reshape process, the direction of some of normals might have changed. The main problem is that, for image recognition and understanding, the human vision system largely relies on 2D information and the original texture might contain some depth cues related to shading that, if not modified, could not reflect the true geometric changes. An interesting example, related to this remark, has been given by Bronstein at al. in [49], where the 3D face models of Osama bin Laden and George Bush have been obtained texturizing the same 3D model with their different images. Their recognition rates demonstrate the importance of image cues w.r.t. 3D geometry.

Reconstructing the correct surface illumination would require to know the position and characteristics of lights into the environment and a detailed characterization of the surface material. However this information is not available in

general (and, in specific, for our case) and obtaining it is a challenging problem. Therefore, we implemented a heuristic solution that provides reasonable results. The basic idea is that, since the surface material (skin) is the same for all the feature points1 and assuming a Lambertian model, the intensity of a surface point depends approximately only on the direction of its normal. Therefore, since we have one-to-one correspondences between the old and the new geometries, we can transfer the texture from the original feature to the morphed one correcting its intensity as a function of the change of the surface normal directions in the following way:

- we construct a look-up table, relating a normal direction with an average texture intensity on the original feature, by equally displacing a number of points along the surface of a sphere of ray one (in our experiments we used 218 points); each point represents a normal direction and its reference illumination is initialized to zero
- we convert the feature texture from RGB to HSV space; then, for each point of the original feature, we find the element of the look-up table whose direction is closer to the point normal and we add to it the V value of the texel associated with that point

[1] There are cases where the feature material is not unique, e.g. the mouth area where lips and surrounding skin can be considered as two different materials; in this case we segment these different regions on the texture plane on the basis of their chromaticity values and we process separately each material

- we compute the mean intensity for each element of the look-up table. If no values were assigned to an element, its intensity is evaluated interpolating those of the samples in a predefined neighborhood, starting with that having an higher number of neighbors and updating neighboring information at each iteration
- for each texel related to a point of the morphed feature, we keep its HS values and change the intensity (V) with that associated to the direction in the look-up table which is most similar to the point normal; the intensities of the pixels of the texture triangle corresponding to a mesh triangle are computed by interpolating its vertices intensities; a light smoothing of the intensity plane is applied in order to reduce irregularities related to the initial normal space sampling
- the transferred texture is converted back to RGB

In Fig. 16 we show an example of texture transfer, where in (a) the original patient's face, subject to a mouth reshaping, is shown; one of the planning outcomes is shown in (c) without and in (d) with the adjusted texture. Finally, in (b) we have highlighted the points where the differences of normal directions are significative, that is the area in which the texture correction takes place.

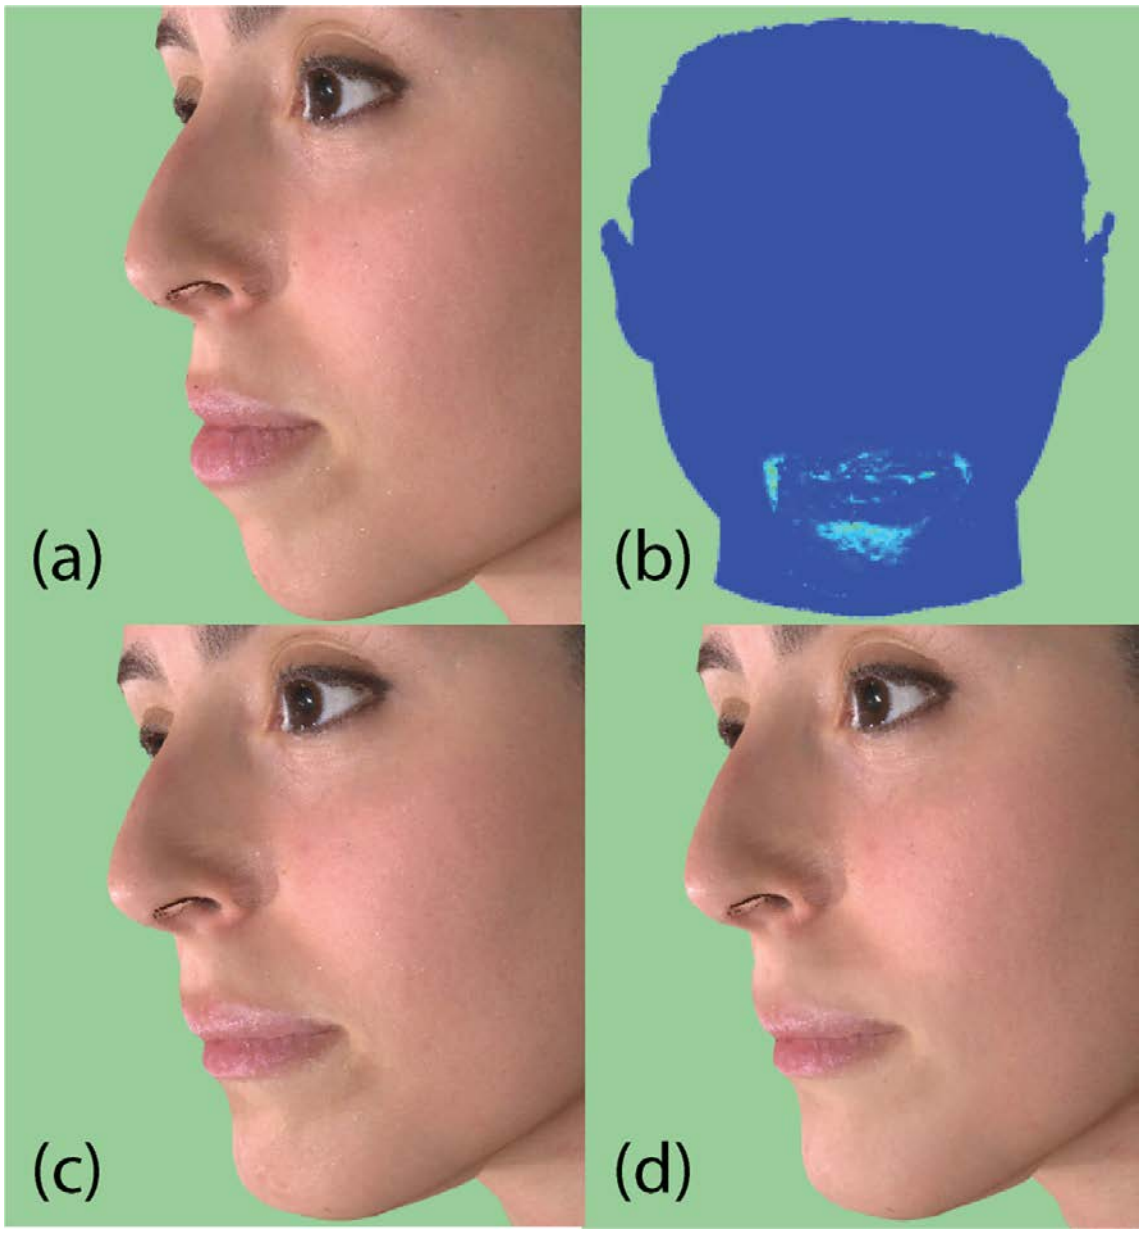

Fig. 16. Patient's face subject to mouth reshaping (a); map of the differences of the normal directions (b); planning outcome without (c) and with (d) texture adjustment

4 EXAMPLES OF PS SIMULATION

The beautification algorithm has been applied to 26 test cases, and k has been set to 4, which appears in keeping with the cardinality (36) of the reference DB. Then, for each case, we produced 4+1=5 PS simulations, including that related to the average feature, and overall 26x5=130 simulations

The test cases have been chosen from the faces discarded when constructing the reference DB. Different target features have been considered in the various cases: nose, mouth, chin and chin and mouth together. The majority of the subjects were involved into nose reshaping.

Some simulations are shown in this section. All have been used for evaluating the average computation times (Section IV) and rated by human panels (Section V).

Examples related to various face features are shown in Fig. 17 and. To make the 3D shape changes more evident, for each case we present three different views of the original patient's scan and of the PS simulation. The simulations shown are the best among 4+1 simulations for each face according to a panel of human Fig. 18 (see Section V). In Fig. 19 we also show an example where all the 4+1 outcomes of nose reshaping simulations are present. The leftmost image in the top row is the original

face. From left to right and top to bottom follow the results obtained pasting the mean feature and the feature of the closest DB samples in decreasing similarity order.

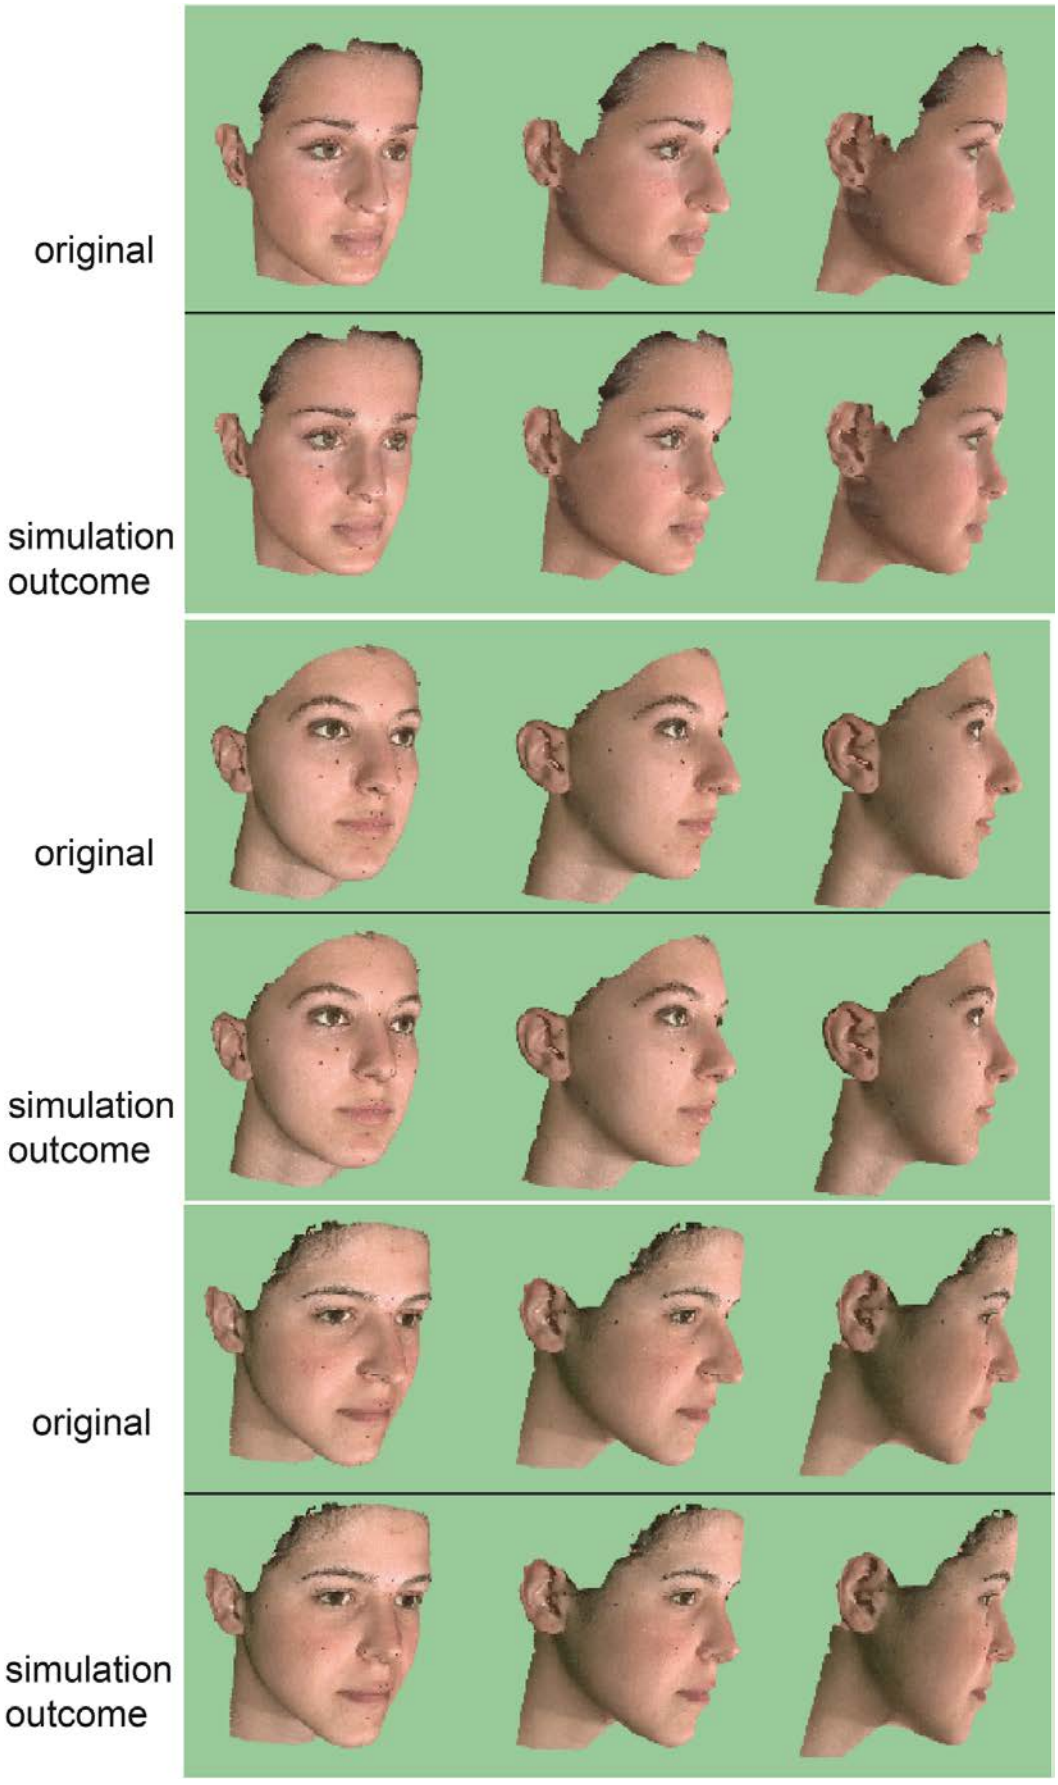

Fig. 17. Examples of simulated surgeries related nose reshaping

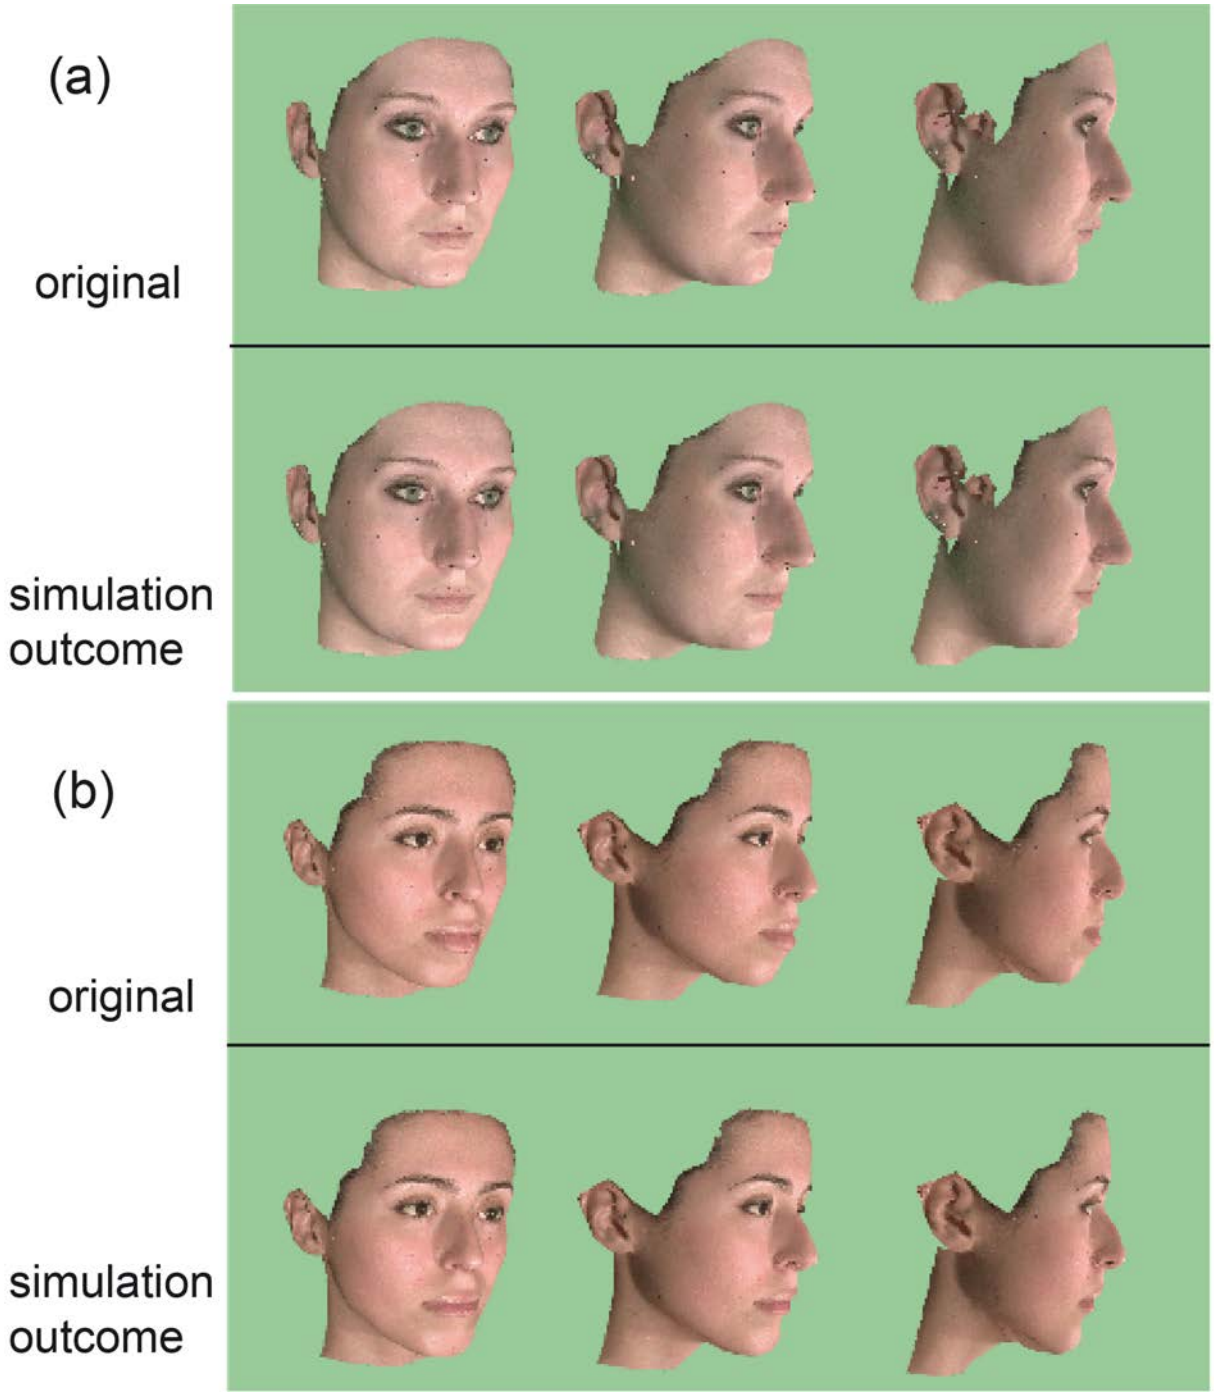

Fig. 18. Examples of simulated surgeries related to a) mouth and b) mouth and chin reshaping

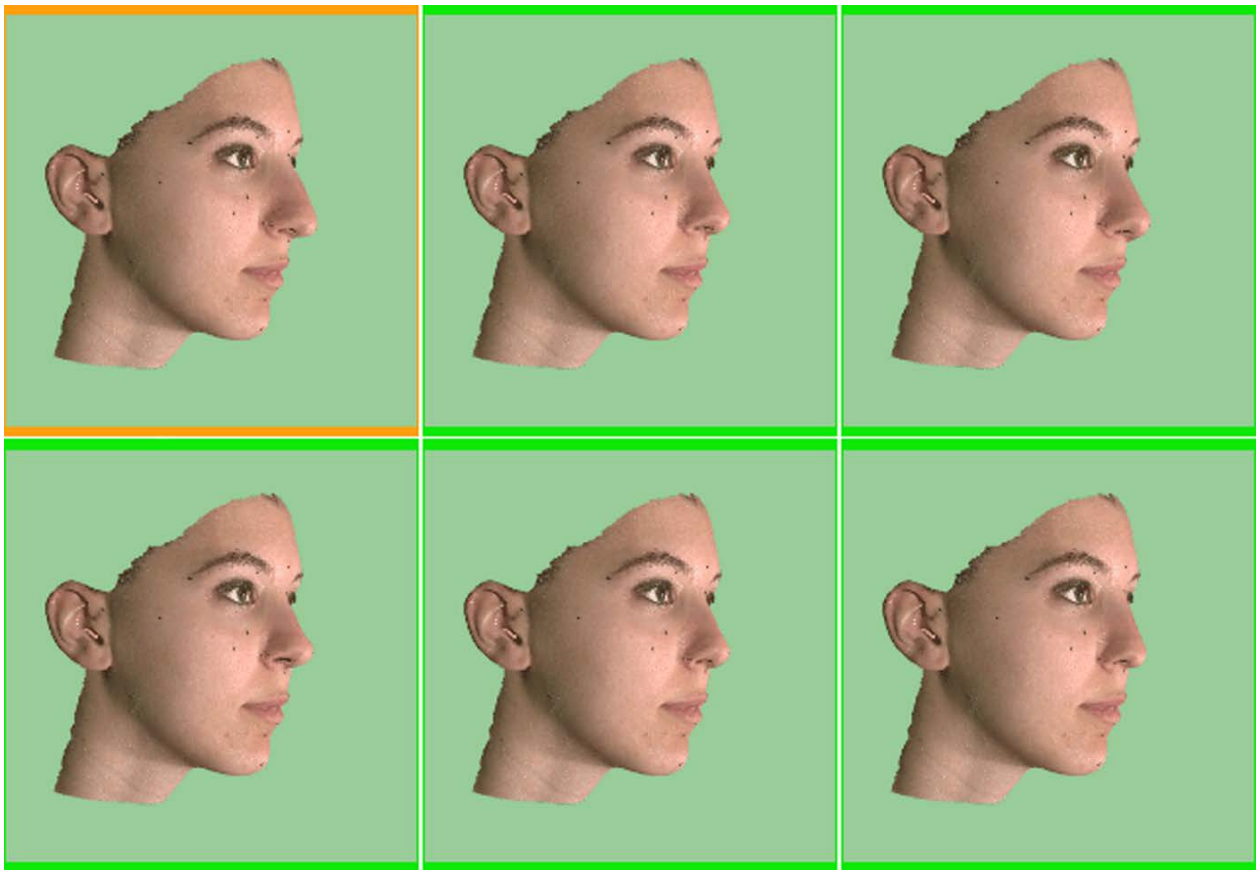

Fig. 19. In this case, all the 4+1 simulated nose reshaping surgeries are shown. Only one of them had been shown in Fig. 14-15. The leftmost image in the first row is the original patient's face

5 PROCESSING TIMES

The entire PS simulation process requires a pre-processing phase, in part manually performed by a trained operator, and a totally computer-performed computation of the possible surgery outcomes. Computation times have been determined by running our system on an Intel Core i7 3.7 GHz processor, 12GB RAM, and a 64-bit Linux operative system.

The times reported in this section are relative to processing the 3D scan of a single patient. We suppose that the construction of the reference DB and all the related operations, as finding suitable face scans, mesh re-sampling, computing correspondences with the intermediate face, have already been performed.

5.1 PRE-PROCESSING

Three manual pre-processing operations are required: 1) scan cleaning and cropping, 2) facial landmark identification and 3) patient's target feature selection. Clearly, scans' cleaning strongly depends on their quality. Anyway, a trained operator (one of the authors) was able to perform operation 1), on the average, in about 6 minutes per mesh (*Snorm*). As for facial landmarks placement, in order to speed up execution times, a rough initialization of their position, obtained from that on the intermediate face, allows the trained operator to perform the task in about 30 sec. (*Landm*). Finally, operation 3) requires to the trained operator about 10 sec. (F_{sel}).

The average times required by the remaining computer-performed phases of the pre-processing algorithm are reported in Table 1. In addition to the top resolution class for re-sampled mesh (300K), used for the images of the previous Section and presented to human raters, 2 additional classes of resolution have been used namely 160K and 40K vertices.

For each resolution class, the actual average number of vertices resulting from the re-sampling process is in column *AvgPt*, and the number of horizontal and vertical planes used for the re-sampling is indicated in column H and V. For each class, the table reports the average time required for finding the sagittal plane and normalizing

faces for size and orientation (F_{norm}), the average execution times for re-sampling the mesh (*Resamp*) and evaluating its bi-directional point-to-point correspondences with the intermediate face (*Corresp*). To compute F_{norm} and *Resamp*, actually depending also on resolution and orientation of initial scans, we used the set of faces in the reference DB, re-sampled at the three class resolutions. These faces were randomly scaled, translated and rotated before applying the normalization process.

				Pre-processing		
CLASS	AvgPt	H	V	Fnorm (sec)	Re-samp (sec)	Corresp (sec)
300K	304,962	776	776	152.64	366.85	8.27
160K	160,141	512	512	81.49	170.38	7.15
40K	40,329	256	256	48.05	126.54	6.25

Table 1. All values are in seconds, except H and V (the number of planes used during re-sampling) and AvgPt

5.2 PS SIMULATION TIMES

Table 2 reports the average execution times required by each step of the simulation process. *Similarity* is the average time for removing the target feature from a database sample and matching it with the patient's face and *Morph* the average time for morphing the target feature, whose mean resolution is in *MorphPt*. The times, depending on k, for computing the average TRF and selecting the *k* closer faces are negligible (with *k=4*, we obtained the average TRF in less than 0.003 sec. for class 300k)

PS simulation

CLASS	Similarity (sec)	Morph (sec)	MorphPt
300K	3.03	3.13	6,277
160K	2.42	2.35	3,256
40K	1.66	1.49	830

Table 2. Times required by each phase of the simulation process. MorphPt is the average number of points per category and per target feature.

5.3 TOTAL PROCESSING TIMES

Overall time T_{all} for obtaining k simulation outcomes, with a reference DB containing n faces, is given by:

$$T_{all} = S_{norm} + Landm + F_{norm} + Resamp + Corresp +$$

$$+F_{sel} + n\,Similarity + (k+1)Morph$$

For instance, with the current settings (36 samples in the reference DB, 5 different simulations), total time is 1052 sec. for resolution 300k and 647 sec. for resolution 40k.

Execution times for larger reference DBs are provided by the following examples. For *n=100* and *k=5*, execution time is 1249 sec. for resolution 300k and 755 sec. for resolution 40k. For *n=200* and *k=7* it is 1558 sec. for resolution 300k and 924 sec. for resolution 40k. These numbers shows that the approach is computationally feasible even with practical DBs. Anyway, improving and parallelizing code could shorten considerably computation times.

Observe that the most expensive part of the computation required for setting up the system, that is pre-processing all the faces of the reference DB, is executed once for all, and adding a new sample requires only pre-processing this sample.

6 HUMAN RATINGS OF PS SIMULATIONS

The performances of the proposed approach have been evaluated in order to assess the quality of the results, in terms of attractiveness improvement. We recall that a dense characterization of the manifold of harmonious faces require several hundreds samples, a number that must be compared with the reduced dimension of our reference DB. Therefore, we consider the results obtained as representative, and not definitive, in understanding the capabilities of the proposed method.

The outcomes of the algorithm were rated through a public website by a panel of colleagues of our University. For each test sample, the original face was shown on the same page together with the planning results obtained with different references (the four nearest neighbors and their average). Users were allowed to rotate interactively and simultaneously all the faces around the common reference axis in order to analyze from different views the proposed results. Raters were asked to evaluate, for each subject and for each of its planning results, if there was an improvement (positive answer), a worsening (negative answer) or no observed changes (neutral answer) in the perceived attractiveness of the face with respect to the original (see an example in Fig. 20).

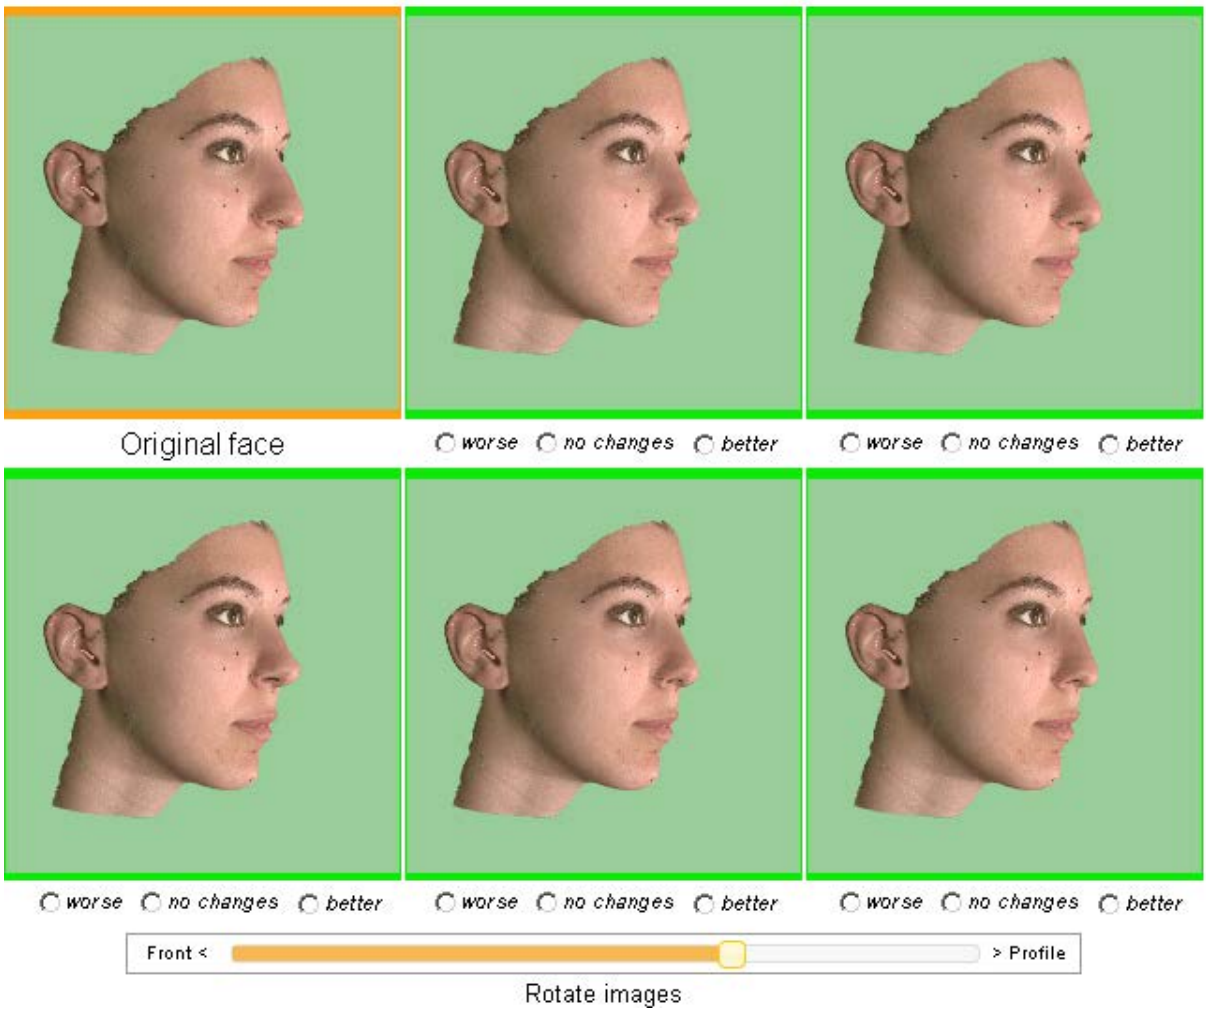

Fig. 20. Extract of a rating page relative to a nose reshaping, showing the original face and the results obtained using as reference the average feature and that of the four samples more similar to the patient; the page was also describing the part of the face subject to the planning procedure and some instructions on how to perform rating

Together with evaluation scores, a few information about the voters, that is gender and age range, were also collected. A total of 15.485 votes, with an average of 119 votes per beautification result, were obtained from 312 judges, showing a substantial rating congruence, consistent with previously reported findings ([15],[16],[17] and [18]). For instance, the difference between attractiveness

improvements expressed for the same test case by male and female raters was 0.14 on the average, with a Pearson correlation of 0.75 between the two groups of raters.

In the following, for the sake of clarity, we will summarize the results obtained from the Human Panel (HP), without reporting the detailed percentage of positive (+), neutral (0) and negative (-) answers for each subject and each reference.

If we analyze the results with respect to the distance in face space between the patient's scan and the reference faces (Table 3), we can see that, according to our initial hypothesis, the shorter is that distance the greater is the number of cases where the corresponding reference is optimal (that is, it provides the highest percentage of positive answers). As a matter of fact, the closest face is optimal in the 42% of the times, outclassing other references (optimal between 15% and 27% of the cases). Moreover, the number of times when the most frequent answer is neutral or negative is increasing with the distance. If we compare these results with those obtained with the average of the closer references as target, we have that the mean is optimal in 18 cases, but its results are very similar to those obtained with the patient's neighbor providing the highest score, since the average difference of their percentage of positive grades is only 2.3%. The number of cases where the most frequent answer is positive for the mean feature as reference is 25 other 26, the remaining case being neutral. These results were expected, and are in keeping to what we already mentioned, that is, the

result of averaging several facial images is usually rated more attractive than each component image but attractive faces are mostly not average.

Reference	Best +		Most frequent		
	nr	%	+	0	-
1st	11	42,3%	22	3	1
2nd	7	26,9%	21	0	5
3rd	6	23,1%	22	1	3
4th	4	15,4%	19	4	3

Table 3. Summary of the results related to distances in face space. Number of times (and percentage) in which the target provides the highest percentage of positive answers (Best +); number of times when the reference obtains, as most frequent, a positive (+), neutral (0) or negative (-) answer

If we choose for each subject the reference providing the highest fraction of votes, with respect to the total, indicating positive (Max +) and not negative (Max 0/+) answers, we can see that the proposed system guarantees a substantial improvement in the attractiveness of the results. As for Max +, we always obtain, excluding one case, a number of positive grades higher than 50%, with a mean value of 70%. However, the only exception, which achieves 48%, scores a minimal number of negative grades (13%) and the most frequent answer is still positive. On the contrary, the not negative

answers are always greater than 77%. The distribution of percentage of positive and not negative answers per sample in the optimal case is shown in Fig. 21.

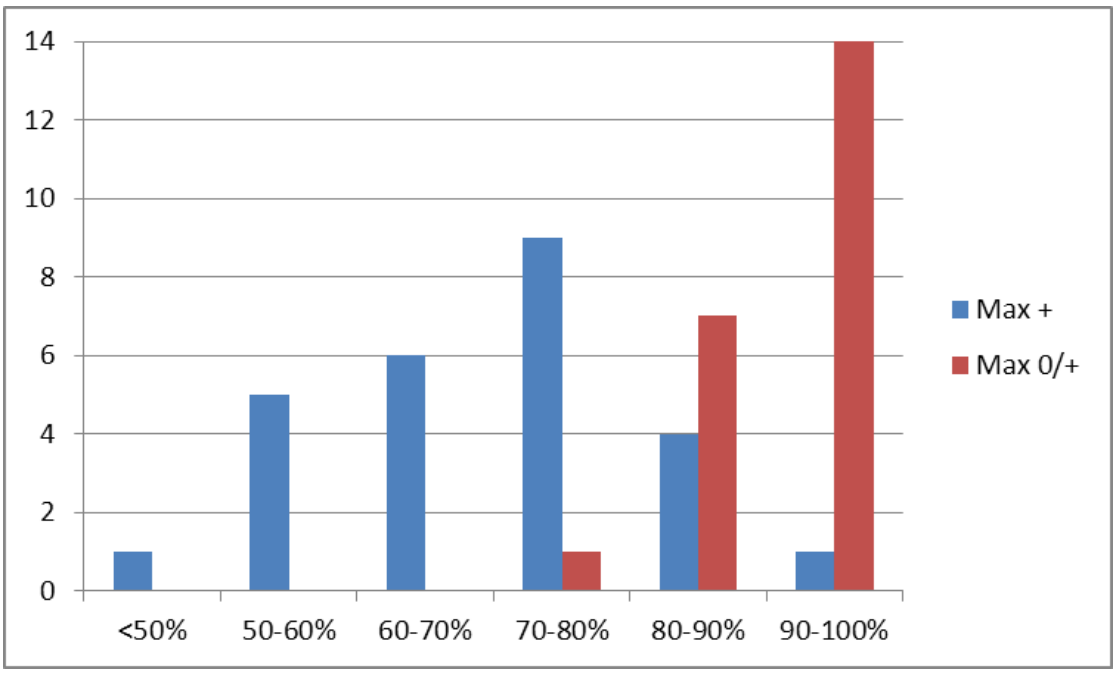

Fig. 21. distribution of percentage of positive grades

Summarizing, according to the HP, choosing the appropriate reference, almost all the beautified faces appear more attractive than the original. As for the reference to choose, experiments indicate that the closer it is in the face space, the greater is the chance that it is optimal, in terms of attractiveness improvement and compared to other references. The mean face provides a good pick as well, which is often close to the optimum. However, since the aim of the system is to provide suggestions for the

surgeon, different possible choices, which we have shown to provide better results in some cases, should not be neglected.

Summarizing, the surgical simulations proposed appear in general more attractive than the original. The simulation obtained using the most similar reference face is several times the more appropriate to select and the average TRF provides a good pick as well.

7 SUMMARY AND FUTURE WORK

We have presented a new 3D system for the aesthetic planning of face plastic surgery. The system can suggest, for any selected facial area, new 3D shapes able to improve the general harmony of the face. Since until now the research in the field has not been able to quantify rules for evaluating and improving general facial feature harmony, we proposed a "learning by examples" approach. In our case, it consists in suggesting changes able to approach the shape of similar, but more harmonious faces. For doing this, a database of harmonious faces has been constructed. The feature to be improved (target feature) is deleted from a patient's face scan, and the residual surface is compared with a set of harmonious face scans, where the same feature has been deleted. Retrieving and suitably pasting onto to patient's face the target feature of the k more similar harmonious faces, as well as the average feature, provides $k+1$ surgery suggestions.

Implementing this general idea required to face and solve a number of problems, including scan normalization, finding dense point-to-point correspondences between different scans, performing smooth and seamless blending of scans, correcting the final textures of the morphed target feature to account for differences in shape and thus in intensity of reflected light. All the corresponding sub-systems have

been implemented and tested. Several examples of surgery simulation have been provided, showing how the system works for key facial areas as nose, mouth and chin.

The tests show that the general idea at the basis of our approach is effective. Although the database of harmonious faces used is rather small, the ratings provided by a human panel demonstrate that the closer is the scan of the harmonious face, the better is the aesthetical result of pasting the correspondent target feature.

Further work for improving the system will be carried on along three directions. It has been already observed that a dense sampling of the space of harmonious faces is likely to require several hundreds of samples. In addition, an estimated 30% of today face surgeries are relative to male patients. Thus, improving both system performances and application areas requires a) increasing maybe five or ten times the cardinality of the female harmonious faces database, with particular attention to covering several ethnic groups; b) constructing a database of harmonious male scans. With regard to the last point, according to most scholars of human attractiveness in human science and medicine, the elements of human face beauty are very different in male and female subjects, and thus the harmonious female database is not likely to be effective for beautifying males.

A second important area of improvement is extending the system described, which deals with the aesthetic planning of surgery, to also deal with the surgical

constraints, combining the surface model described with a volumetric scan of soft and hard tissues. Finally, in order to deal with large reference DBs, we are planning to improve the computational performances, for instance rewriting the code, which is highly parallelizable, addressing a multi-thread architecture, or porting some of the algorithms in GPU.

8 LIST OF FIGURES

9 LIST OF TABLES

10REFERENCES

[1] American Society for Aesthetic Plastic Surgery, “Statistics 2009”. 2009. Available online at: http://www.surgery.org/sites/default/files/2009stats.pdf, last accessed Feb 2013

[2] P. Adamson and S.Doud Galli, ”Modern concepts of beauty,” Current Opinion in Otolaryngology& Head and Neck Surgery,Vol.11, pp. 295-300, 2003

[3] R.Singh, M. Vatsa, H.S. Bhatt, S. Bharadwaj, A. Noore, and S.S. Nooreyezdan,” Plastic Surgery: A New Dimension to Face Recognition,” IEEE Trans. On Info. Forensics And Security, Vol.5, No.3, pp. 441-448, 2010

[4] H. S. Bhatt, S. Bharadwaj, RSingh, M. Vatsa, and A. Noore “Evolutionary Granular Approach for Recognizing Faces Altered Due to Plastic Surgery,”IEEE Conf. on Auto. Face and Gesture Reco, 2011,pp. 720-725

[5] R. Bracaglia, R. Fortunato and S.Gentileschi, “Secondary Rhinoplasty”, Aesthetic Plastic Surgery, Vol.29, pp.230-239, 2005.

[6] T. Ozkul, M.H. Ozkul, “Computer simulation tool for rhinoplasty planning”, Comput. In Biol. And Med., Vol.34, pp.697-718, 2004

[7] S. Rabi and P. Aarabi, “Face Fusion: an automatic method for virtual plastic surgery”, Proc. Int. Conf. On Information Fusion, July 2006.

[8] J.Liu, X.Yang, T.Xi, L.Gu, and Z. Yu, ”A novel method for computer aided plastic surgery prediction,”IEEE Proc. BMEI’09, pp.1-5, 2009

[9] http://www.seattlesoftwaredesign.com/, last accessed May 2012

[10] J. Gao, M. Zhou, H. Wang and C. Zhang, “Three dimensional surface warping for plastic surgery planning”, IEEE proc.Int. Conf. On Systems, Man and Cybernetics, pp.2016-2021, 2001.

[11] T. Lee, Y. Sun, Y. Lin, L. Lin and C.Lee, “Three-dimensional facial model reconstruction and plastic surgery simulation”, IEEE Trans. On Info. Tech. In Biomed., Vol.3, No.3, pp.214-220, 1999.

[12] T. Lee, C. Lin and H. Lin, “Computer aided prototype system for nose surgery”, IEEE Trans. On Info. Tech. In Biomed., Vol.5, No.4, pp.271-278, 2001.

[13] J. Wang, S. Liao, X. Zhu, Y. Wang, C.Ling, X.Ding, Y. Fang, and X. Zhang,“ Real time 3D simulation for nose surgery and automatic individual prosthesis design,“ Computer Methods and Programs in Biomedicine 104(3), pp.472-479, 2011

[14] http://www.axisthree.com/professionals/home, last accessed May 2012

[15] Cunningham, M.R., Roberts, A.R., Barbee, A.P., Wu, C.H. and Druen, P.B.: Their ideas of beauty are, on the whole, the same as ours: consistency and variability in the crosscultural perception of female physical attractiveness. J. Pers. Soc. Psychol., Vol.68, pp. 261-279, 1995

[16] J. N. Thakera, and S. Iwawaki, “Cross-cultural comparisons in interpersonal attraction of females towards males,” Journal of Social Psychology, Vol.108, pp.121 – 122, 1979

[17] S. M. Maret, “Attractiveness ratings of photographs of blacks by Cruzans and Americans, “Journal of Psychology, Vol. 115, pp.113 – 116, 1983

[18] E. Wagatsuma and C. Kleinke,” Ratings of facial beauty by asian-american and caucasian females,” J. Of Soc. Psychology, Vol.109, pp. 299-300, 1979

[19] A.Bottino, A. Laurentini, “The analysis of facial beauty: an emerging area of research in pattern analysis”, Lecture Notes In Computer Science, pp. 11, 2010, Vol. 6111/2010,

[20] K. Schmid, D. Marx, and A.Samal, “Computation of face attractiveness index based on neoclassic canons, symmetry and golden ratio, “ Pattern Recognition, Vol.41, pp.2710-2717, 2008

[21] B.W. Baker and M.G.Woods, “The role of the divine proportions in the esthetic improvement of patients undergoing combined orthodontic/orthognathic surgical treatment,” Int. J. Adult. Orthod. Orthognath. Surg., Vol.16, pp. 108-120, 2001

[22] J.P.Moss, A.D. Linney, and M.N. Lowey,” The use of three-dimensional techniques in facial aesthetics”, Semin. Orthod., Vol.1, pp. 94-104, 1995

[23] T.R. Alley, and M.R. Cunningham, “Average faces are attractive, but very attractive faces are not average”, Psychological science, Vol.2, pp.123-125, 1991

[24] D. Perret, K. May, and S. Yoshikawa, “Facial shape and judgements of female attractiveness”, Nature, Vol.368, pp.239-242, 1994

[25] T. Le, L.Farkas, R. Ngim, L. Levin, and C. Forrest,: “Proportionality in asian and north American Caucasian faces using neoclassic facial canons as criteria”, Aesthetic Plastic Surgery, Vol.26, pp.64-69, 2002 .

[26] M. Bashour, “History and Current Concepts in the Analysis of Facial Attractiveness”, Plastic and Reconstructive Surgery, Vol.118, No.3, pp.741-756, 2006

[27] E. Rosch, “Principles of categorization”, In Rosch, E. and Lloyd B., (1978): “Eds. Cognition and Categorization”, Erlbaum, 1978.

[28] Y. Eisenthal, G. Dror, and E. Ruppin, “ Facial Attractiveness: beauty and the machine, “ Neural Computation, Vol. 18, pp. 119-142, 2006

[29] A. Bottino, A, Laurentini, and L Rosano, “A New Computer-aided Technique for Planning the Aesthetic Outcome of Plastic Surgery”, Proceedings of 16th Int. Conf. on Computer Graphic, Visualization and Computer Vision WSCG 2008, Plzen (Cz)

[30] BFM (3D Basel Face Model), available online at http://faces.cs.unibas.ch, last accessed May 2012

[31] USF Human ID, available online at http://figment.csee.usf.edu/dataandcode.php, last accessed Feb 2013

[32] Canfield Imaging System, http://www.canfieldsci.com , last accessed Feb 2013

[33] Blender 3D. Available online at: http://www.blender.org. Last accessed: Feb 2013

[34] G.Pan, Y. Wang, Y.Qi, and Z. Wu, "Finding symmetry plane of 3D face shape," IEEE Proc. ICPR 2006

[35] Bottino, A. and Cumani, S., "A fast and robust method for the identification of face landmarks in profile images", WSEAS Trans. on Computers, 7(8), 2008.

[36] Z. Zhang. "Iterative point matching for registration of free-form curves and surfaces". Int. J. Comput. Vision 13, 2. 1994

[37] The Visualization Toolkit Library. Available online at: www.vtk.org/. Last accessed: Feb 2013

[38] Filippo Tampieri. 1992. Newell's method for computing the plane equation of a polygon. In Graphics Gems III, David Kirk (Ed.). Academic Press Professional, Inc., San Diego, CA, USA 231-232.

[39] Y. Hu, M. Zhou, and Z. Wu, "A Dense Point-to-Point Alignment Method for Realistic 3D Face Morphing and Animation," International Journal of Computer Games Technology, vol. 2009

[40] Leslie G. Farkas. 1994. Anthropometry of the head and face. Raven, New York

[41] Horn B., "Extended Gaussian Image," Proc. of IEEE, Vol 72, No 12, pp. 1671-1686, 1984.

[42] Gauss K. F., " General Investigations of Curved Surfaces", Raven Press, New York, 1965.

[43] L.A.Lysternik,"ConvexFiguresandPolyhedra",DoverPublications, New York,1963.

[44] Blanz, V. and Vetter, T., "Face Recognition Based on Fitting a 3D Morphable Model", *IEEE Trans. Pattern Anal. Mach. Intell.,* vol. 25, no. 9, pp. 1063-1074, 2003

[45] Thirion, J. P., "Fast non-rigid matching of 3D medical image." Technical report, Research Report RR-2547, Epidure Project, INRIA Sophia, 1995

[46] Thirion, J.P., "Image matching as a diffusion process: an analogy with maxwell's demons." *Medical Image Analysis*, vol. 2, no.3, pp.243–260., 1998

[47] Python Programming Language Official Website. Available online at: www.python.org/. Last accessed: Feb 2013

[48] de Berg M., van Krefeld M., Overmars M., Schwarzkopf O., "Computational Geometry: Algorithms and Applications" Springer, 2 edition, 2000.

[49] A. M. Bronstein, M. M. Bronstein, R. Kimmel, "Three-dimensional face recognition", Intl. Journal of Computer Vision (IJCV), Vol. 64/1, pp. 5-30, August 2005

[50] Ghislanzoni L. H., De Simone M., Rosati R., Ugolini A., Bottino A., Sforza C., "The construction of the average adult upper dental arch: a clinical validation of a new three-dimensional method", INT. Journal of Oral Science (2013)

11 APPENDIX A

Our work produced a set of software tools and a specific code library in C++. In this appendix we briefly describe them.

The developed code is intended to be multiplatform (Linux, Windows and OSX) and it is based on the open source library VTK [37]. All the developed C++ classes follow the conventions defined by the VTK dev team.

The primly development language is C++ but we use Python [47] as supporting language for collateral tasks as experiment automation, data generation and debugging.

All tools are command line driven so we can easily concatenate and control batch process by Python. When a GUI is necessary (for example in case of setting references point for point-to-point relationship evaluation – see chapter 3.4) an interactive mode is added. This architectural approach follows a common conventional philosophy in Linux environment that is our main development operating system.

Each tool is described in the follow.

- `Cleaner`. This tool preprocesses input scans removing duplicate vertices and fixing wrong topology. Also it convert input data from .OBJ format to VTK format (see chapter 3.2);

- `Evaluator`. This tool evaluate the distance between two faces returning the RMS as described in chapter 3.5;
- `Mappoints`. It performs the evaluation of point-to-point relationship between two scans (chapter 3.4)
- `Ptsetter`. It is an interactive tool to set reference points on scan faces. These keypoints are used by `Mappoints`;
- `Remapper`. This tool re-arranges texture to be consistent (Fig. 24)

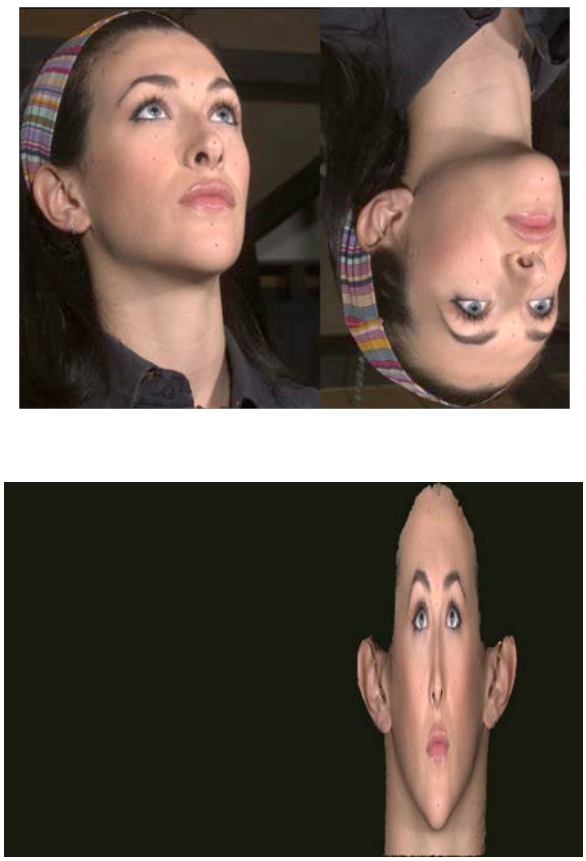

Fig. 22.Some texture needs to be rearranged. Upper picture is the original texture followed by the new one

- `Retopo`. It performs the normalization as described in chapter 3.2

- `Selector`. It is the interactive tool that allows to select the feature on a face scan;
- `Sticker`. It sticks the new feature on the initial face as described in chapter 3.6;
- `Texturefixer`. It fixes the texture to grant shading consistency with the new geometry. See chapter 3.6;
- `Morpher`. It is a debug tool. It performs a morphing between two faces to highlight resulting point-to-point relationship.
- `tools_manager.py`. It is a Python script and it calls tools described to automatize the experimental data generation. It is driven by specific configuration files that they describe set of experiments;
- `Average`. This tool generates an average mesh. If we have the point-to-point relationship between a reference model R and a set of N meshes M_j, we can use it to build a new 3D object O. Each point p_i in O is $p_i = \frac{r_i + \sum_{j=0}^{N} m_{ji}}{N}$ where r_i is a point of R, m_{ji} is the correspondent point of r_i in M_j.

All software programs perform specific tasks so it is easy to use them to face different problems. For instance, we applied our software also in a research work made in collaboration with the Università di Milano. The work was about digital construction and validation of an average adult upper dental arch and its application in the clinical environment. Laser-scanned images of the dental casts were obtained with an optical laser-scanning device. Seventy-nine landmarks (Fig. 23) were identified on each dental arch on the basis of a protocol previously validated for dental analysis. An average dental arch shell was then created by `Average` tool. After validation, the average dental arch was used as a template for comparison with other dental arches presenting some form of malocclusion. The results were published on [50].

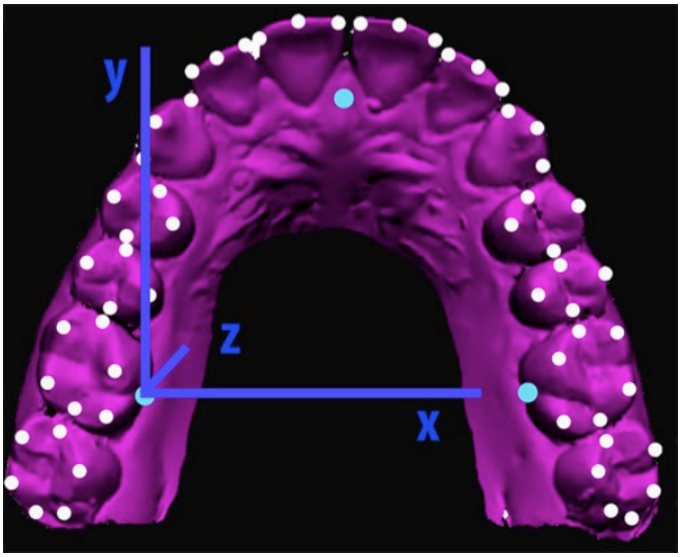

Fig. 23 Reference poinst on a dental arch scan

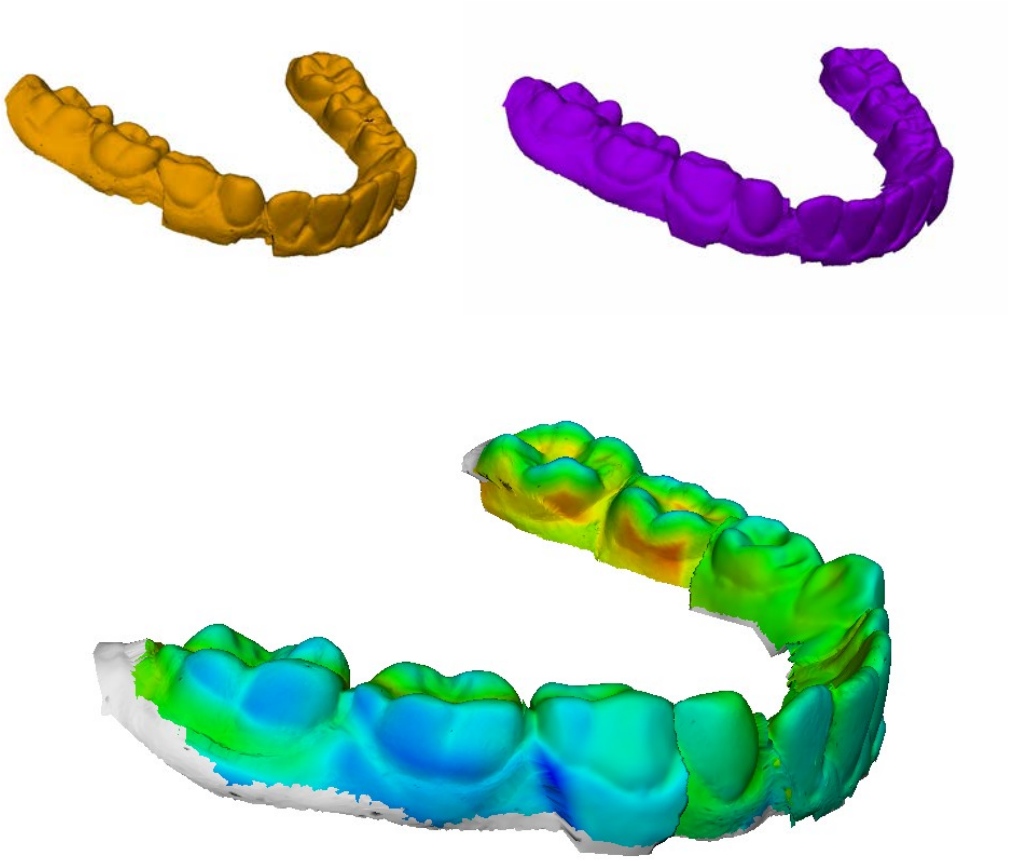

Fig. 24. The average dental arch before (upper left) and after (upper left) treatments. Lower image shows the average dental arch difference between before and after treatments by a color map.

12 APPENDIX B

Here we report more results and example generated by our algorithm as reference.

12.1 Nose

A. original face, B. average feature, C. weighted average feature, D.,E.,F. selected closest features. More details in Chapter 4.

A.
B.
C.
D.
E.
F.

A.

B.

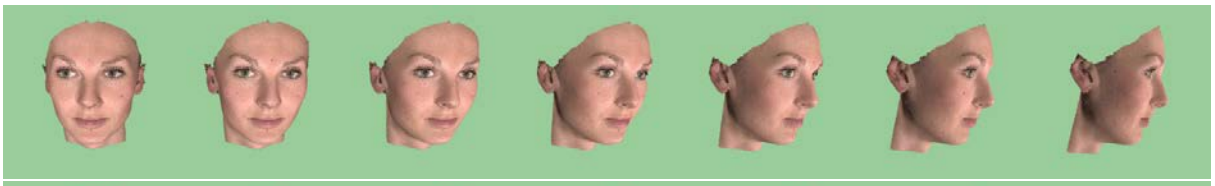

C.

D.

E.

F.

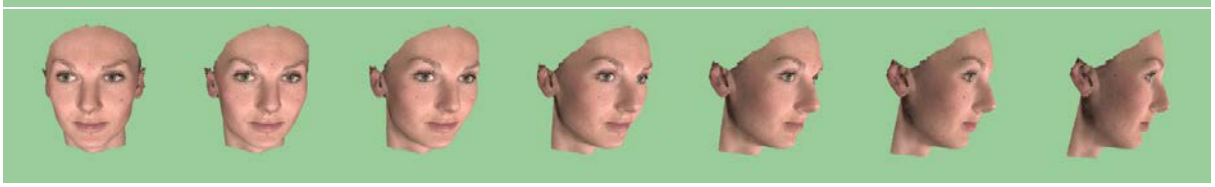

A.

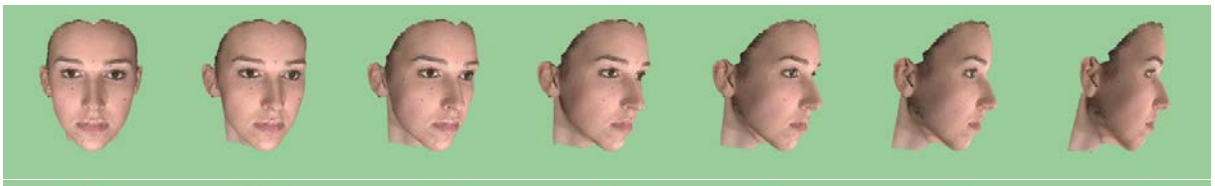

B.

C.

D.

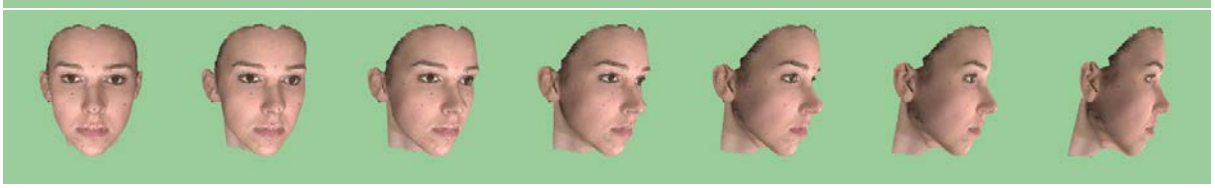

E.

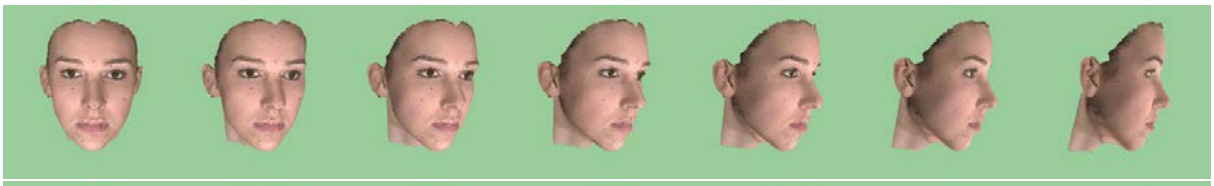

F.

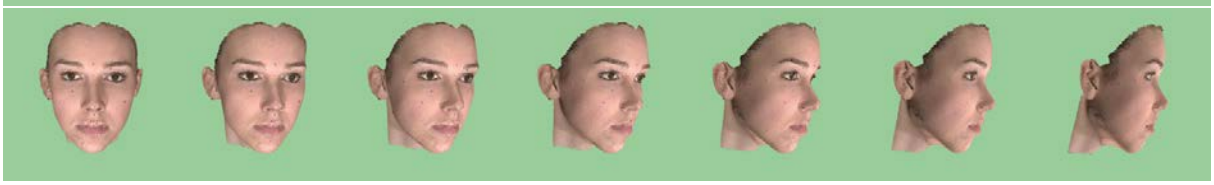

A.

B.

C.

D.

E.

F.

A.

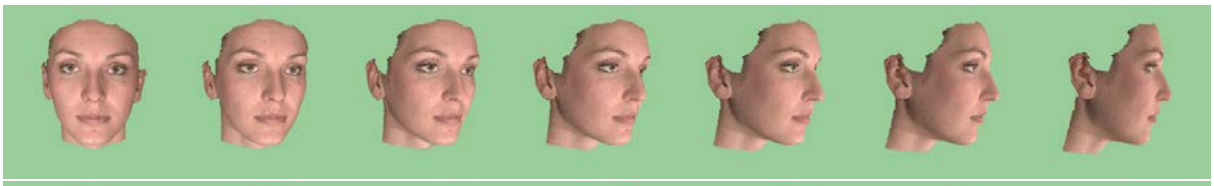

B.

C.

D.

E.

F.

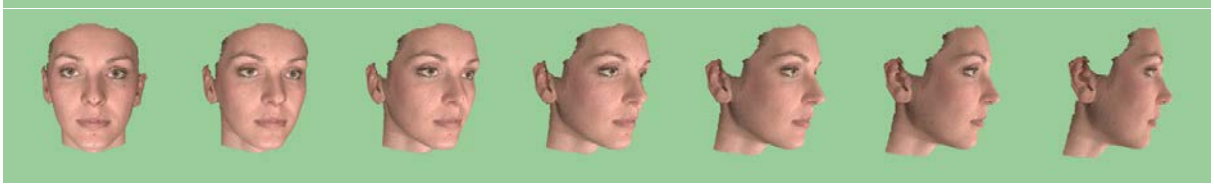

A.

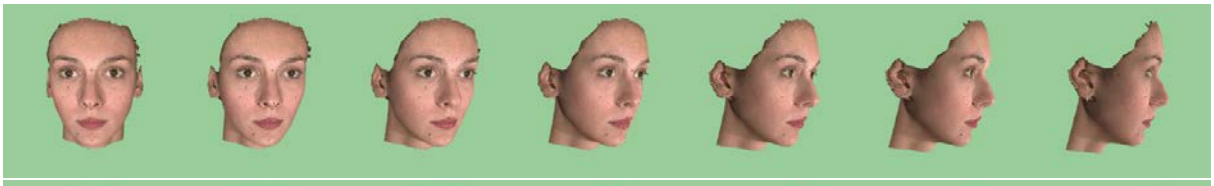

B.

C.

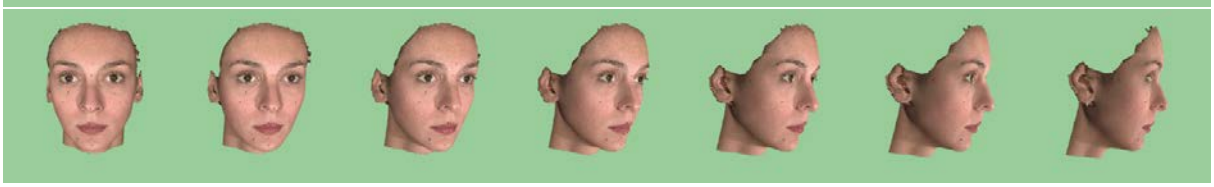

D.

E.

F.

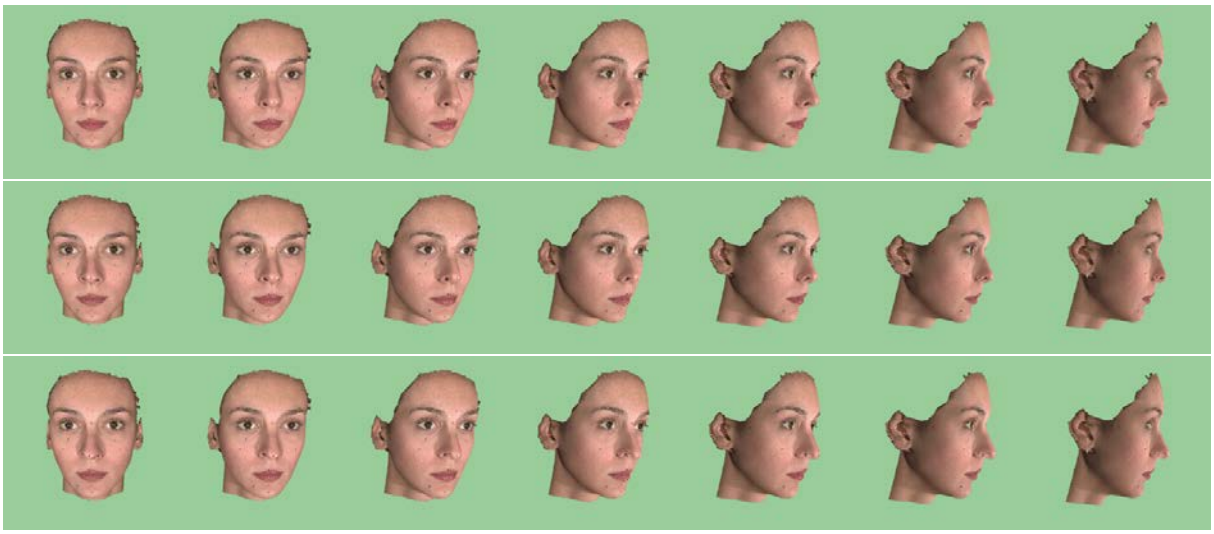

A.

B.

C.

D.

E.

F.

A.

B.

C.

D.

E.

F.

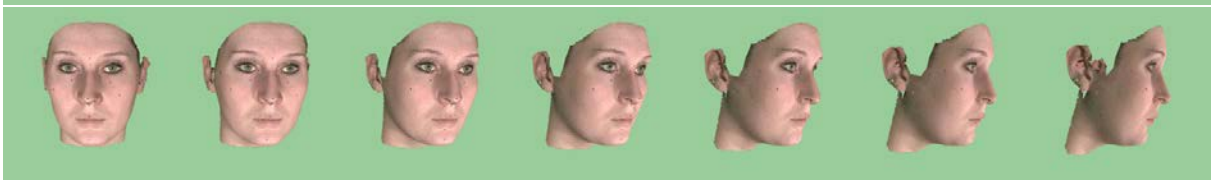

A.

B.

C.

D.

E.

F.

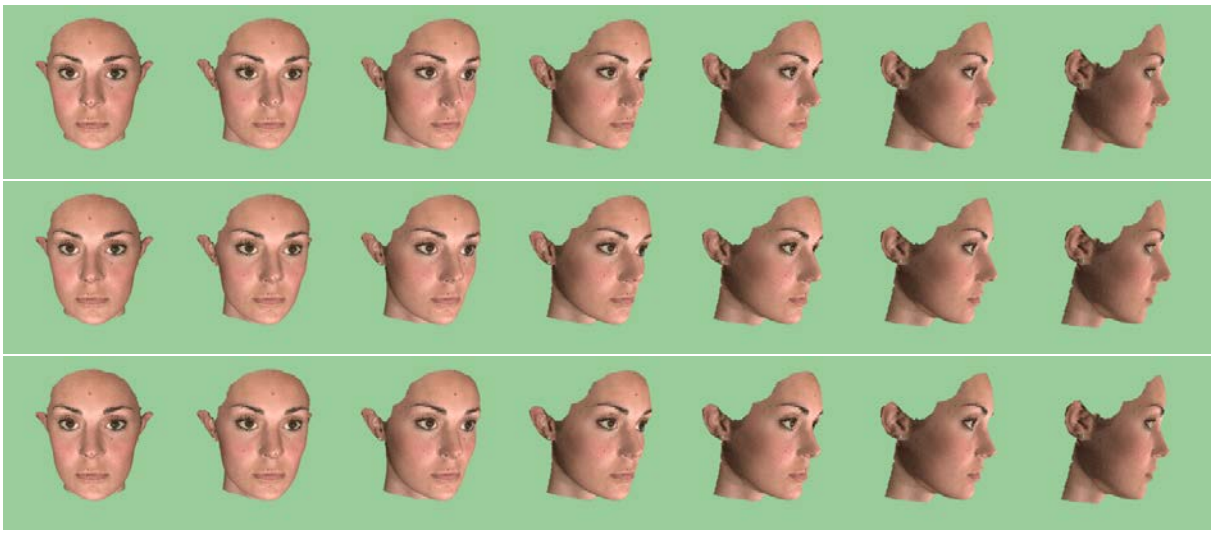

A.

B.

C.

D.

E.

F.

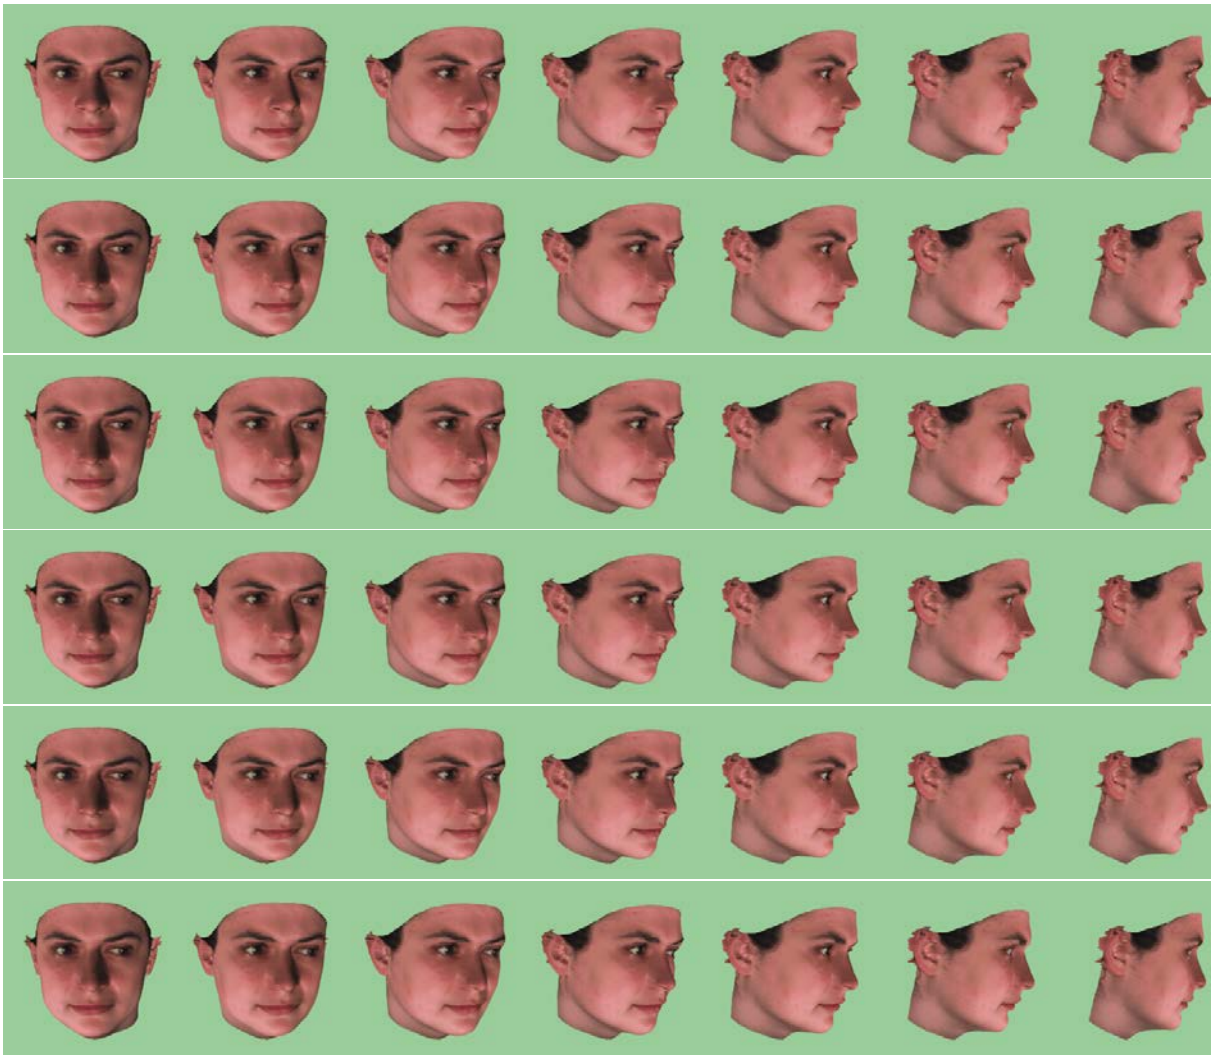

A.

B.

C.

D.

E.

F.

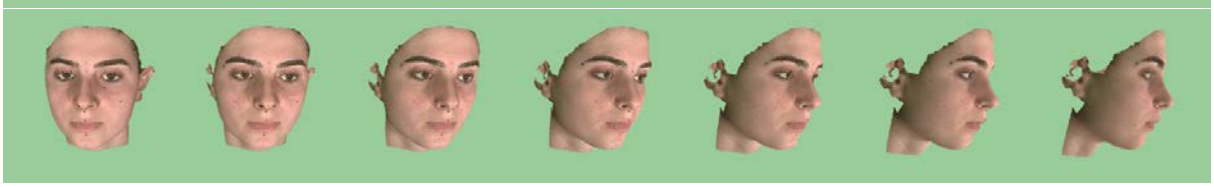

A.

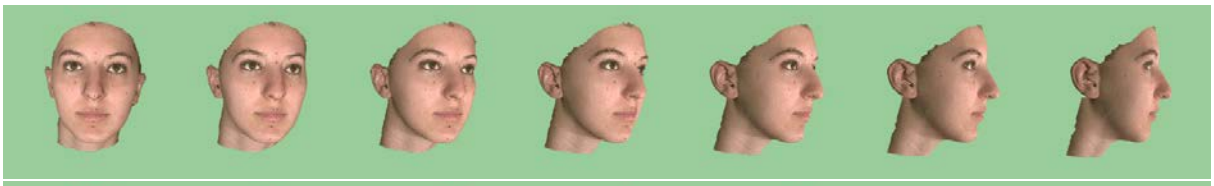

B.

C.

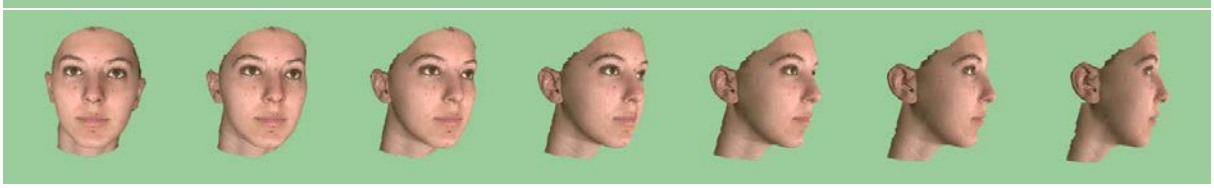

D.

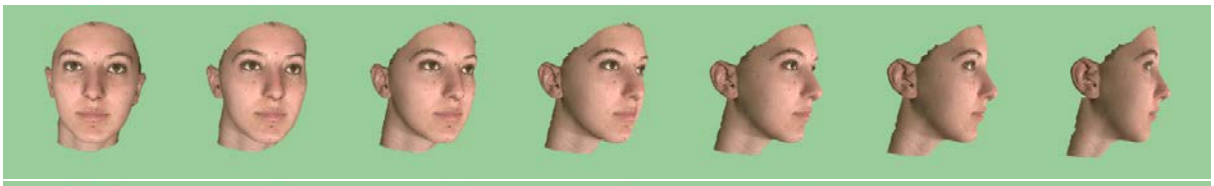

E.

F.

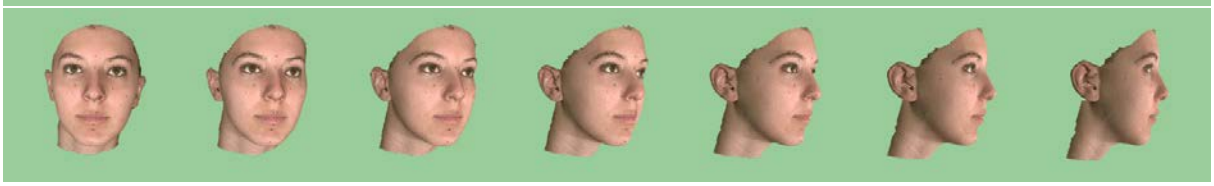

A.

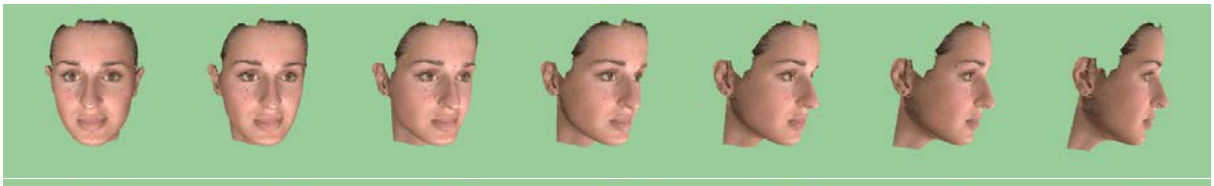

B.

C.

D.

E.

F.

A.

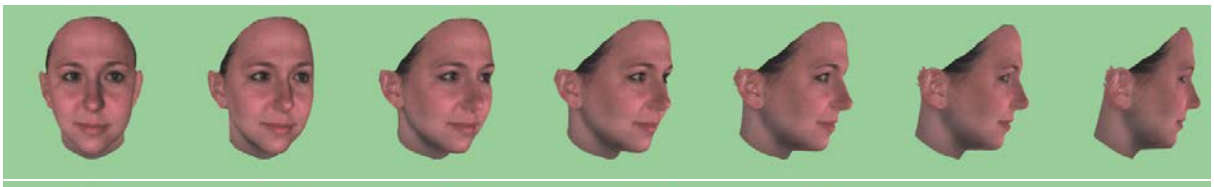

B.

C.

D.

E.

F.

A.

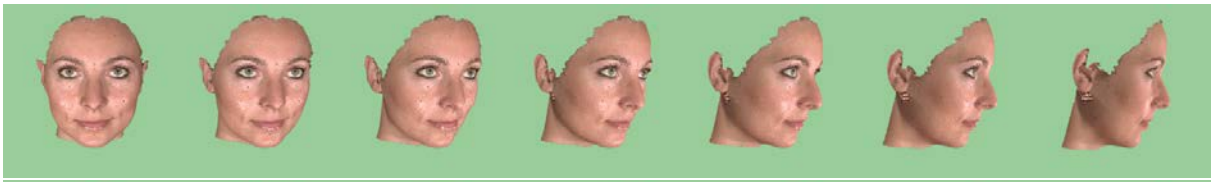

B.

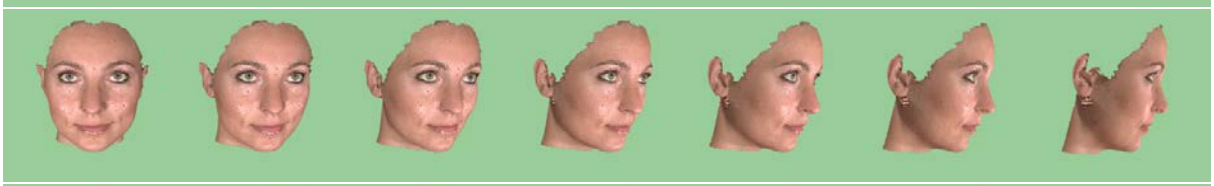

C.

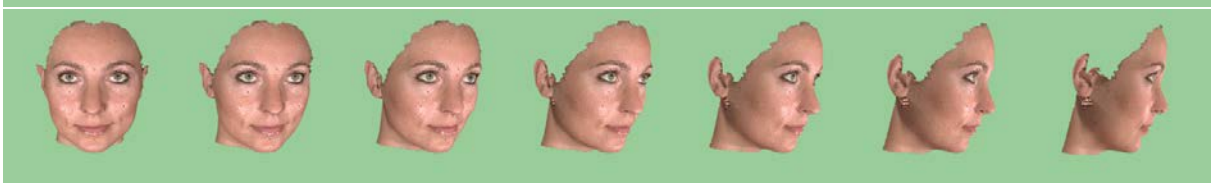

D.

E.

F.

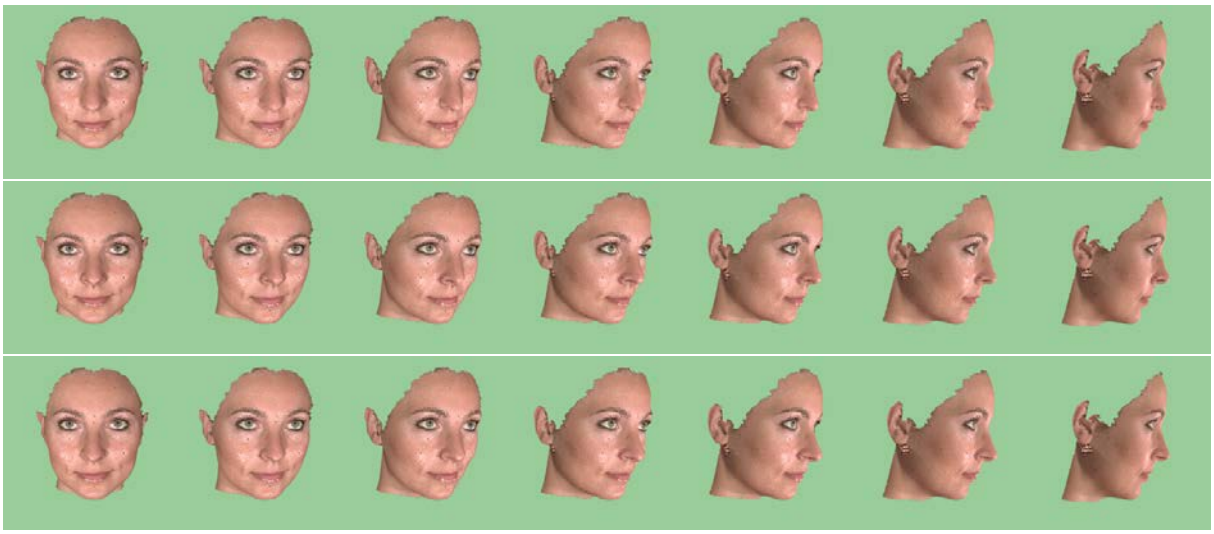

A.

B.

C.

D.

E.

F.

A.

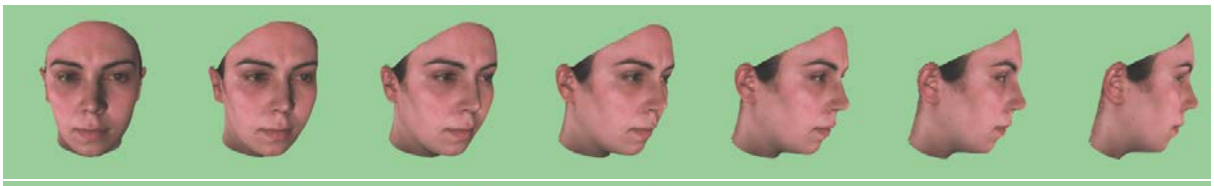

B.

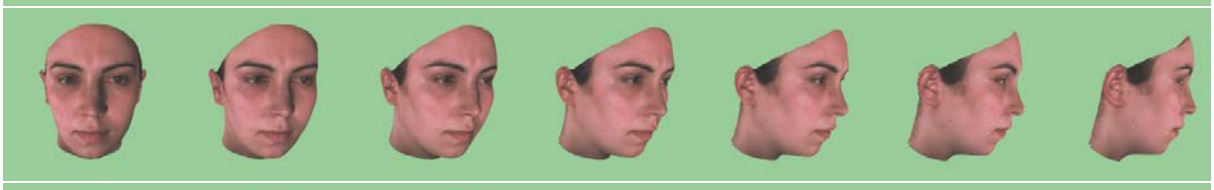

C.

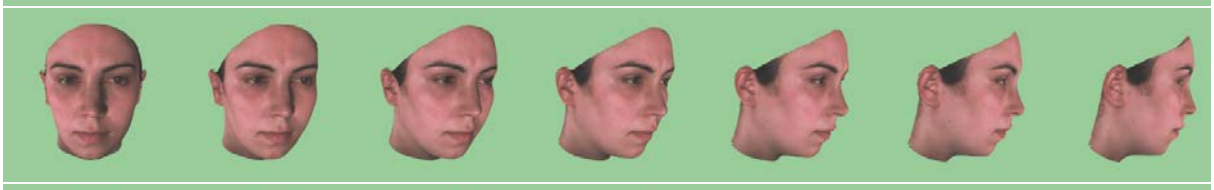

D.

E.

F.

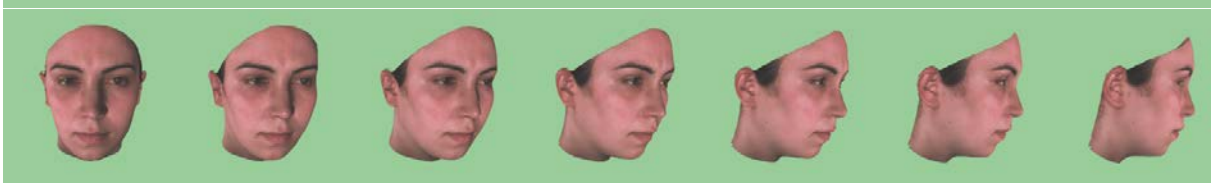

A.

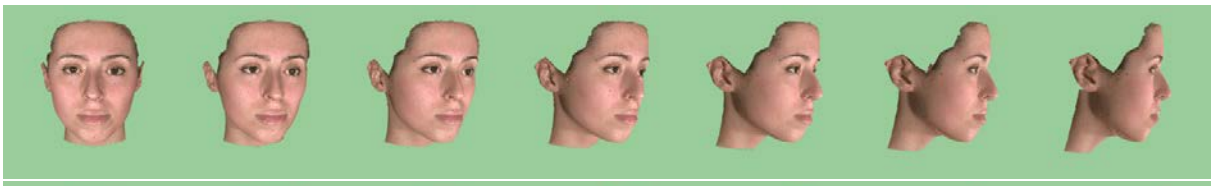

B.

C.

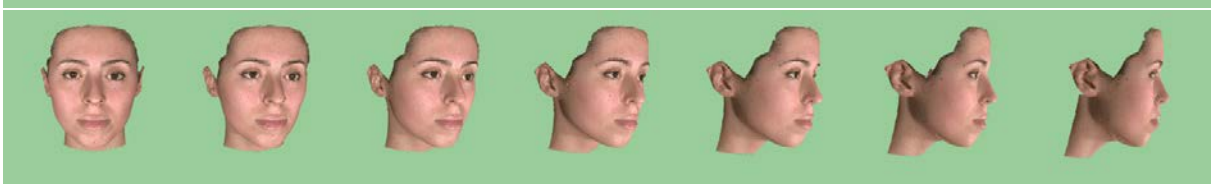

D.

E.

F.

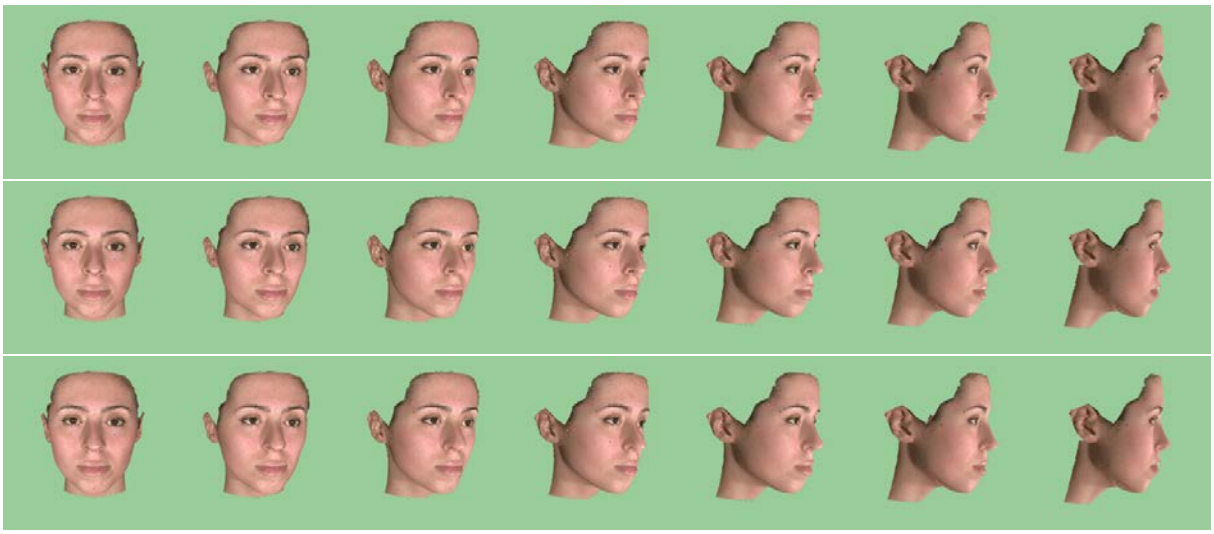

A.

B.

C.

D.

E.

F.

12.2 CHIN

A. original face, B. average feature, C. weighted average feature, D.,E.,F. selected closest features. More details in Chapter 4.

C.

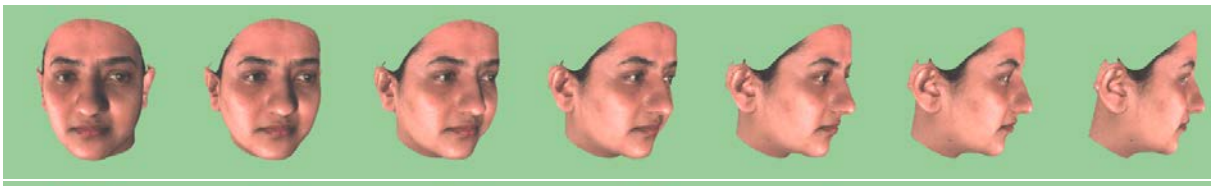

D.

E.

F.

12.3 CHIN AND MOUTH

A. original face, B. average feature, C. weighted average feature, D.,E.,F. selected closest features. More details in Chapter 4.

A. 

B. 

C. 

D. 

E. 

F. 

A. 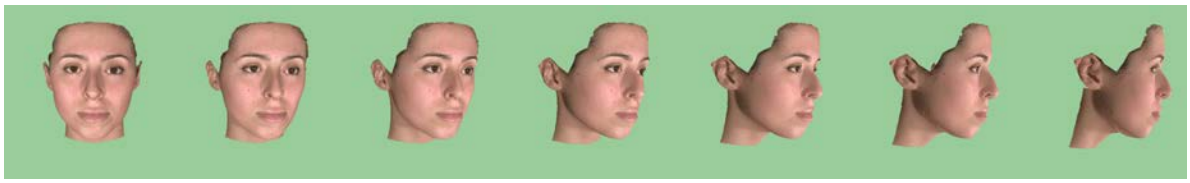

B. 

C. 

D. 

E. 

F. 

12.4 Mouth

A. original face, B. average feature, C. weighted average feature, D.,E.,F. selected closest features. More details in Chapter 4.

A.

B.

C.

D.

E.

F.

A. 

B. 

C. 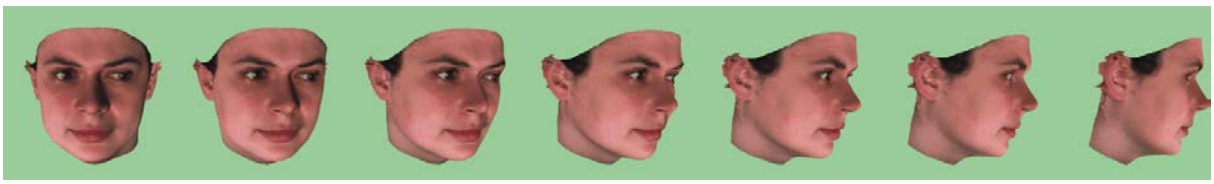

D. 

E. 

F. 

A. 

B. 

C. 

D. 

E. 

F.