POLITECNICO DI TORINO
Repository ISTITUZIONALE

A New 3D Tool for Planning Plastic Surgery

Original
A New 3D Tool for Planning Plastic Surgery / Bottino, A.G., DE SIMONE, M., Laurentini, A., C., S.. - In: IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING. - ISSN 0018-9294. - STAMPA. - 59:12(2012), pp. 3439-3449. [10.1109/TBME.2012.2217496]

Availability:
This version is available at: 11583/2502046 since:

Publisher:
IEEE

Published
DOI:10.1109/TBME.2012.2217496

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

15 September 2026

Author's version

Published in: IEEE TRANSACTIONS ON BIOMEDICAL ENGINEERING, VOL. 59, NO. 12, DECEMBER 2012, pp. 3439-3449, doi: 0.1109/TBME.2012.2217496

A New 3D Tool for Planning Plastic Surgery

A. Bottino, M. De Simone, A. Laurentini, C. Sforza

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 paper, 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.**

Index Terms— **Computer supported surgery, Plastic surgery, face scanning, planning surgery outcomes, anisotropic ICP**

I. 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 2008 more than one million facial plastic surgeries were performed, with a growth of 162% 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 an harmonious or at least regular face shape is a main goal for both kinds of surgery.

A. 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].

B. *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.

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.

Two of the authors presented a first implementation of this idea in [29] focused on 2D face profiles, and choosing as target feature the nose. In this paper, we extend the idea to 3D face scans, and to any part of the face.

The content of the paper is as follows. In Section II, we describe the 3D scans database used and we detail the various steps of the algorithm. In Section III, we present examples of PS simulation. Processing times are reported in Section IV and the ratings of the results of PS simulations supplied by a panel of human raters are discussed in Section V. Concluding remarks are reported in Section VI.

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

II. THE PROPOSED TECHNIQUE

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

- 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.

A. *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 acquired high-resolution 3D attractive face scans of professional male and female models. Data have been obtained with the VECTRA-CR-two-pod system [32]. Each 3D scan 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.

B. *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. 2). 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. 3). Both cleaning and cropping have been performed manually with Blender [33].

Fig. 2. Examples of face scans (rendered view and wireframe detail) from different databases

Fig. 3. An example of scan before (a) and after (b) cleaning and cropping

Our normalization algorithm performs the following steps.

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.

The sagittal plane is determined by the algorithm described in [34], exploiting the Extended Gaussian Image (EGI). The sagittal plane is intersected with the face scan in order to determine its profile, where two facial landmarks, the *nasion* and *subnasale* are identified (see Fig. 4). These landmarks are computed with the algorithm described in [35].

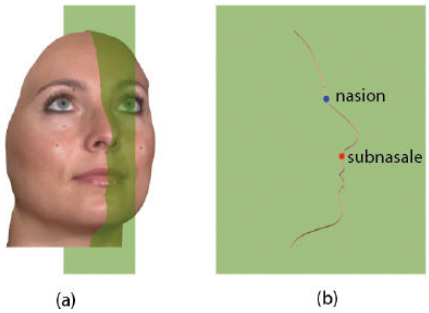

Fig. 4. Intersection of a face scan with its sagittal plane and automatic identification of landmarks on face profile

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 II.D 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].

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.

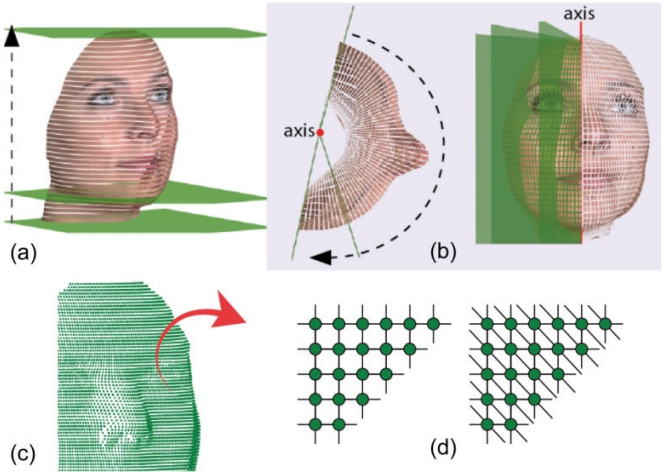

Fig. 5. Re-sampling process: (a) the mesh is sliced with horizontal planes obtaining a number of polylines; (b) the polylines are intersected with vertical planes rotated around the vertical axis. (c) The final point cloud; (d) reconstruction of the mesh connectivity

To normalize resolution, face scans are re-sampled. As shown in Fig. 5(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. 5 (b)), passing through the reference axis, common for all aligned faces, and equally spaced angularly. The resolution of the final meshes can be adjusted changing the H and V values. For the experiments described in Section III, and submitted to human raters (Section V), both variables were assigned the value 776, which produced mesh resolutions near to 300K vertices. For the results presented in Section IV (Processing times), we also used two lower resolutions (160K and 40K).

The result of re-sampling is the point cloud shown in Fig. 5(c). The final triangular mesh is determined, as shown in Fig. 5(d), by forming quads and splitting each of them in two triangles. Texture coordinates for each vertex of the new 3D mesh are computed with bilinear interpolation of those of the neighbouring points in the original mesh.

C. *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. 6 an example is shown relative to the nose area. A detail view of an example of polyline and facial triangulation is shown in Fig. 7.

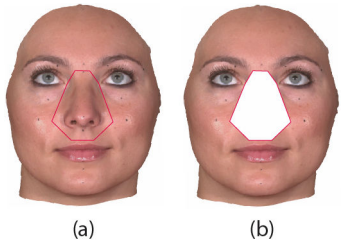

Fig. 6. Interactive feature selection: contour drawing (a); removal of the feature (b).

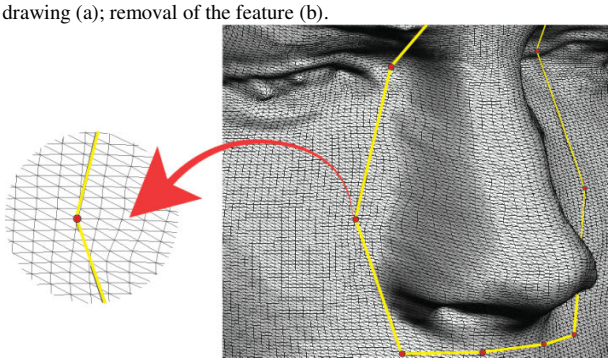

Fig. 7. A detail view of the polyline delimitating the area subject of the PS procedure (red dots are vertices of both the polyline and the face mesh)

D. 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 whose generatrices pass through the polyline selected and parallel to a direction evaluated from the polyline by the Newell's Method [38].

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, we use a technique 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. 8). These keypoints are a subset of the standard landmarks used in Anthropometry [40] with additional points in eyes area. 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. 9. 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.

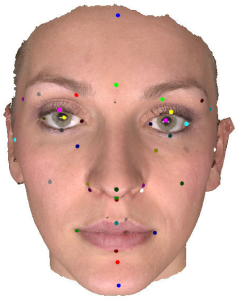

Fig. 8. The keypoints used for establishing dense point-to-point correspondence between faces

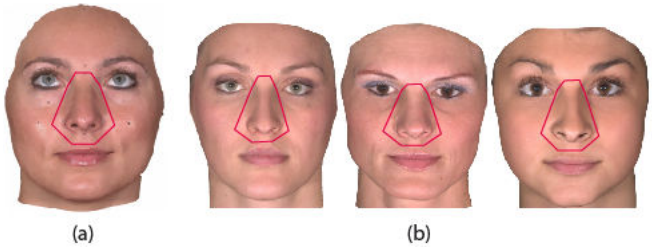

Fig. 9. An example of transferring a feature boundary from the patient's face (a) to the scans in the reference DB (b)

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. 10).

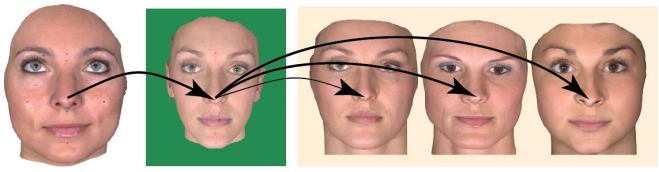

Fig. 10. A point of the input face (a) corresponds to a point on the intermediate face (b) that, in turn, corresponds to a point on each face in the reference DB (c)

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.

E. 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.

F. 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. 11). 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.

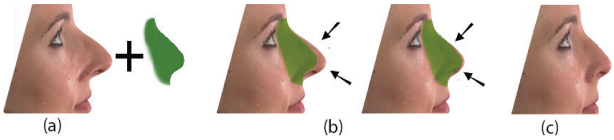

Fig. 11. The reference feature (a, in green) is used to morph the target feature (b) to obtain the final result (c)

In order to describe the details of the process, let us refer to Fig. 12, where we show in (a) the target feature (red, on the left), the TRF (gray, on the right) and in Fig. 12(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. 12c);
2. we transform P_i (Fig. 12d) 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:

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

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. 12(e) shows, in false colours, 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.

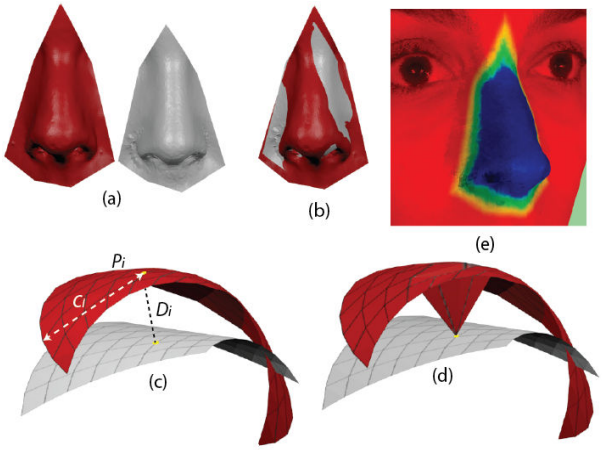

Fig. 12. a) target feature (left) and TRF (right); b) features alignment; c, d) evaluation of the weight of point P_i (with $W(P_i)$=1) and its transformation; e) a false color map of the weights $W(P_i)$ (red = 0, blue = 1)

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 the 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[1] and a Lambertian reflection model, the intensity of a surface point depends approximately only on the direction of its normal.

[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

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. In more detail, this is performed in the following way.

First, we convert the feature texture from RGB to HSV space and we construct a look-up table (LUT) relating normal direction and average texture intensity for the target feature. The LUT is constructed by averaging, for 2^{18} equally spaced directions of the normals, the V values of the mesh vertices whose normals are closer to each direction.

Then, for each vertex of the morphed mesh, we find the closest normal in the LUT and retrieve the intensity. The intensities of the pixels of the texture triangle corresponding to a mesh triangle are computed by interpolating that of its vertices. Some smoothing of the intensities is applied to reduce irregularities related to the initial normal space sampling. Finally, the adjusted intensities are combined with the original HS values and the texture is converted back to RGB.

An example of texture correction is shown in Fig. 13. A surgery simulation of a patient subject to mouth and chin reshaping (a) 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 significant, that is the area where the texture correction takes place.

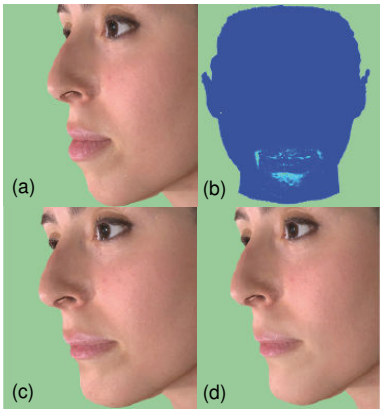

Fig. 13. Patient's face subject to mouth reshaping (a); map of the differences of the normal directions (b); planning outcome without (c) and with (d) texture adjustment

III. 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. Some of them were professional models, but presenting, in our opinion, some slight defect. 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. 14 and Fig. 15. 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 raters (see Section V). In Fig. 16 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.

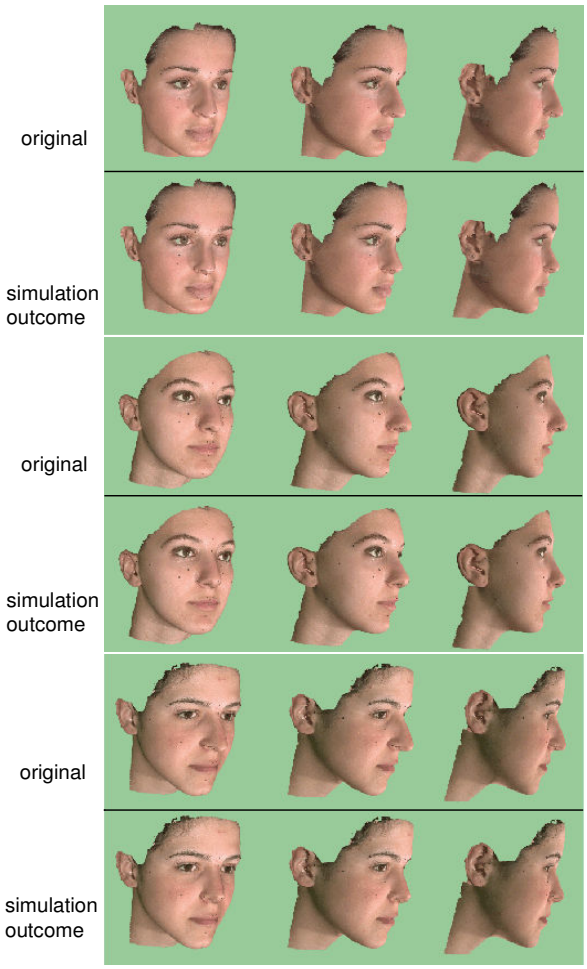

Fig. 14. Examples of simulated surgeries related nose reshaping

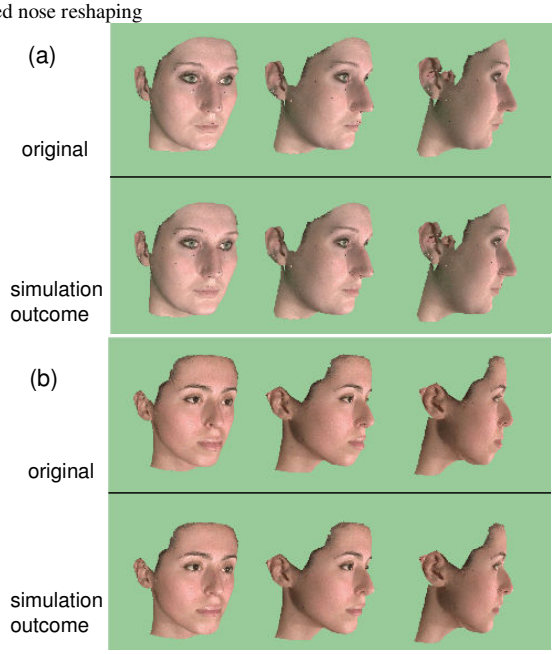

Fig. 15. Examples of simulated surgeries related to a) mouth and b) mouth and chin reshaping

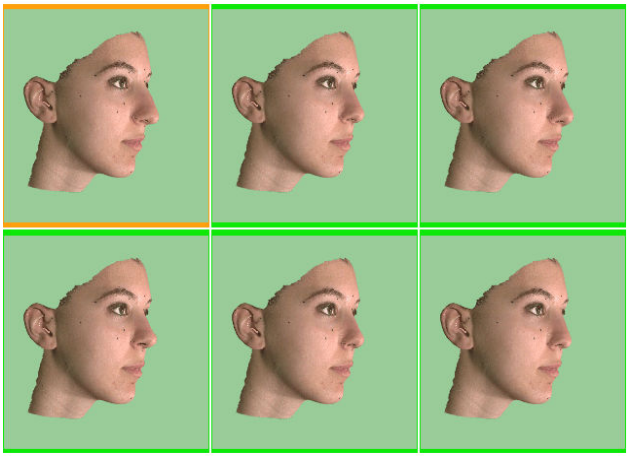

Fig. 16. In this case, all the *4+1* simulated nose reshaping surgeries are shown. Only one of them had been shown in Fig. 14-15. The leftmost image in the first row is the original patient's face

IV. 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.

A. *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. (*Fsel*).

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 (*Fnorm*), the average execution times for re-sampling the mesh (*Re-samp*) and evaluating its bi-directional point-to-point correspondences with the intermediate face (*Corresp*). To compute *Fnorm and Re-Samp*, actually depending also on resolution and orientation of the 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

B. *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 more close 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.

C. *Total processing times*

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

*Snorm + Landm + Fnorm + Re-samp + Corresp + Fsel + 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 he 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.

V. HUMAN RATINGS OF PS SIMULATIONS

We recall that a dense sampling of the manifold of harmonious (female) faces is likely to require some hundreds of samples, when our reference DB only contains 36 samples. Therefore, the effectiveness of the system could considerably improve with a larger DB, and the examples presented should essentially be considered a demonstration of the system developed.

However, we also arranged the rating of the simulation results through a public website of the Politecnico di Torino. For all the 26 test cases, the original scan and the five surgical simulations were shown on the same page in random order. Raters were allowed to rotate interactively original and modified faces. Then they were asked to evaluate, for each PS simulation, if there was an improvement (positive answer), a worsening (negative answer) or no change (neutral answer) in the perceived attractiveness. A total of 15,485 votes, with an average of 119 votes per simulated result, were obtained from 312 judges. The ratings showed a substantial congruence between gender and age range of the raters, which is consistent with congruence of human panels judgments of attractiveness reported in many papers.

A first result, actually rather obvious, is that in almost all cases the results were perceived as positive or neutral. Globally, the number of simulations scoring more than 50% of positive grades is 68.4%, while those obtaining more than 50% of not negative grades were 99.2%.

A more interesting result is the connection between closeness of scans in the face space and perceived attractiveness improvement. In Table 3, column *Best+* shows, for the four closer scans of the reference DB, the number and percentage of test patients for which that corresponding PS simulation obtained the highest number of positive grades. It can be seen that the closest scan outperforms (11 times and 42% of the cases) the others. In the table, the sum of test patients gives 28 instead of 26, since there are two cases where two different scans provide the same highest grade.

Although relative to a small DB, and thus affected by noise due to under-sampling, this is a confirmation of the general assumption at the basis of this paper.

Another interesting result is the high rating of the average TRF. Actually, it received the highest percentage of positive answers in 18 cases. However, if we compare the percentage of positive answers of the average TRF and the highest of the four closer scan, their absolute difference is very small (2.3%). Also this result is probably affected by the under-sampling. Anyway, it is in keeping with the fact already mentioned that the result of averaging several facial images is usually rated more attractive than each component image.

Closer scan	Best +	
	nr	%
1st	11	42.3%
2nd	7	26.9%
3rd	6	23.1%
4th	4	15.4%

Table 3. Number of times (and percentage) in which the TRF of the closer reference provides the highest percentage of positive answers (Best +)

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.

VI. 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.

REFERENCES

[1] American Society for Aesthetic Plastic Surgery, "Statistics 2008". 2009. Available online at: http://www.surgery.org/sites/default/files/2008stats.pdf, last accessed May 2012

[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 May 2012
[32] Canfield Imaging System, http://www.canfieldsci.com , last accessed May 2012
[33] Blender 3D. Available online at: http://www.blender.org. Last accessed: May 2012
[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: May 2012
[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

A. Bottino was born in Torino, Italy, in 1971. He received his master degree in Computer Science Engineering and his Ph.D. from Politecnico di Torino in 1995 and 2000. He is currently a researcher in Computer Science at the Dipartimento di Automatica e Informatica of the same university. He is author of several journal and conference papers. His current research interests include computer graphics, computer vision and virtual heritage.

M. De Simone was born in Lanciano, Italy, in 1975. He received his master degree in Computer Science Engineering from Politecnico di Torino in 2005. After an experience as researcher on web technologies and knowledge management in a private company, he is currently PhD student in Computer Science at the Dipartimento di Automatica e Informatica of Politecnico di Torino. He is author of several journal and conference papers. His current research interests include computer graphics, computer vision and game developing.

A. Laurentini was born in Genova, Italy, and received the degree of Ingegneria Elettronica from the Politecnico di Milano in 1963. In 1965 he joined the Politecnico di Torino, where he has been professor of Computer Science at the Dipartimento di Automatica e Informatica until November 2011. He is member of IEEE and ACM, and author of more than eighty scientific papers. His current research interests include computer vision, computer graphics and computational geometry.

C. Sforza was born in Monza, Italy, and received her M.D. degree from the Università di Milano, Italy in 1986 and her Master degree in Sports Medicine in 1989 from the same University. She is a professor of human anatomy at the Medical and Dental School of the Università di Milano, and Director of the

Laboratory of Functional Anatomy of the Stomatognathic System. She has been involved in research on the functional anatomy of the stomatognathic apparatus and on the anatomical bases of stomatognathic function and dysfunction. She has authored/coauthored more than 230 publications indexed in PubMed