

A Portable Modular System for Automatic Acquisition of 3-D Objects

N. Alberto Borghese, *Member, IEEE*, and Stefano Ferrari

Abstract—A portable and flexible system for three-dimensional (3-D) scanning is here presented. It is constituted of four main modules. The first module is devoted to the acquisition of a set of 3-D points over the surface through laser scanning (digitization). The second module reconstructs a continuous 3-D surface, filtering the measurement noise. Whenever required (e.g., in CAD applications), a third module converts the surface into a 3-D mesh which can then be colored by projecting over it a bitmap of the surface, obtained from a snapshot to obtain a highly realistic textured 3-D model. This instrument improves upon the commercially available scanners in two main aspects. The digitizer proves to be highly flexible and accurate, and it can easily accommodate objects of different dimension. The construction of the surface and the filtering of the digitization noise are performed in a single step through a fully adaptive algorithm which produces a multiscale surface and can be parallelized to work in real time. Results on the reproduction of human faces are reported and discussed.

Index Terms—Modular architectures, multiresolution, neural networks, radial basis functions, range data, real-time, scanner 3-D.

I. INTRODUCTION

VIRTUAL three-dimensional (3-D) models are required by an increasing number of applications ranging from basic image processing to video conferencing, constructive and plastic surgery, 3-D fax, reversed engineering, and 3-D CAD (virtual prototyping). A host of devices (3-D scanners), have come to the market in the last few years to provide these models. Essentially, a 3-D scanner captures the 3-D visible surface of an object as a mesh suitable to be processed by CAD and graphical systems [1]. Although ultrasound [2] or mechanical (e.g., Microscribe) devices are available, optical technology is preferred because it allows a high resolution and it does not require any contact with the surface. The gold-standard is represented by the Cyberware scanners which are suitable to most applications. However, apart from their very high cost, they have two drawbacks:

- 1) it has a complex structure, inside which the object is placed, has to be set up, and
- 2) only objects within limited size ranges can be digitized.

In this paper, a modular 3-D scanner, which is devoted to opposite criteria is described (Fig. 1). The first module is devoted

to the acquisition of a set of 3-D points over the surface through laser scanning (digitization). The second module reconstructs a 3-D continuous surface which filters out the noise introduced by the digitization process. At this stage, a 3-D model is already available. For some applications in CAD/CAM and computer graphics, the representation of the surface in the form of a mesh is more suitable; this is done in a third module, which converts the surface into a mesh. In a fourth module, a bitmap of the surface, obtained from a snapshot, is aligned and projected over the 3-D mesh to obtain a highly realistic textured 3-D model. The aim of this paper is to present the overall structure, focusing on the accuracy of the digitization process and on the reconstruction of the 3-D surface. The overall processing allows achieving a high accuracy and can be implemented onto PC boards to achieve real-time 3-D modeling. The system has been widely tested in the reconstruction of 3-D objects and of human faces in particular.

II. DATA ACQUISITION

The first step toward creating a 3-D model is the digitization of a set of 3-D points on the objects surface, which is carried out here by the Autoscan system introduced in [3]. This is constituted of a commercial laser pointer of 5 mW of power, a set of video cameras, which provide each an image of 256×256 pixels with a frame rate of 100 frames/s, a real-time image processor and a host computer. The processor computes a cross-correlation between a 6×6 mask, template of the laser spot, and each image of the stream coming from the video cameras. This allows achieving a high SNR (a laser spot can be detected also in outdoor conditions) and sub-pixel accuracy (by averaging the over-threshold pixels of the laser spot, weighted with their cross-correlation value), which has been experimentally determined in 0.1 pixel [4]. Digital cameras can therefore be used as a valid alternative to analog sensors for spot detection (cf. [5]). The pattern matching procedure allows also to automatically discard those spots which are poorly seen: in this case, cross-correlation is over-threshold only for one row and/or column of the image, losing sub-pixel accuracy in the computation of the spot center. The cross-correlation is computed in real-time through a custom board which contains essentially delay lines, adders, and multipliers [6]; it can be replaced by less expensive general-purpose hardware (e.g., on FPGA board), without degrading the performances. The 3-D position of the laser spot is computed with standard photogrammetric procedures from the spot position on two cameras. For this purpose, the cameras are calibrated by surveying a rigid bar with two spherical markers on its extremities, the same size of the laser spot, inside the volume where the object is scanned [7]. This makes the system extremely simple to

Manuscript received May 26, 1999; revised June 22, 2000.

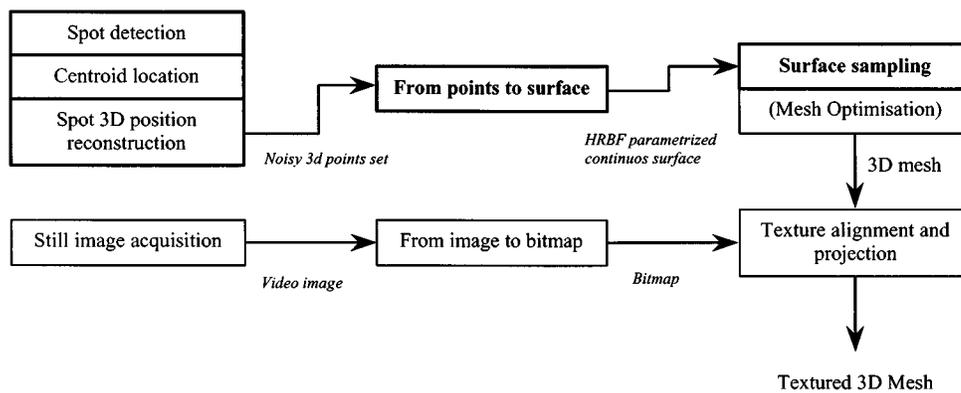
N. A. Borghese is with the Laboratory of Human Motion Analysis and Virtual Reality (MAVR), Istituto Neuroscienze e Bioimmagini—CNR, 20090 Segrate, Milano, Italy.

S. Ferrari is with the Center for Bioengineering, Fondazione ProJuventute, Department of Electronics and Information, Politecnico di Milano, 20133 Milano, Italy.

Publisher Item Identifier S 0018-9456(00)07953-5.



(a)



(b)

Fig. 1. System: scanning a human face. The beam of a commercial laser pointer is directed over the subject's face. Scanning is performed moving the laser manually with the aim of collecting more data points in the most difficult regions: lips, eyes... At the same time, a commercial CCD video camera acquires (a) a snapshot of the face. A schematic diagram of the processing blocks is reported in (b).

set-up. We explicitly remark that there are no constraints on the set-up and on the size of the object to be scanned: this depends only on the lens used and the relative position of the cameras. Flexibility, scalability, and portability therefore characterize the acquisition module.

Surface digitization is carried out by moving the laser pointer manually [Fig. 1(a)]. To help in directing the laser beam, a real-time feedback is provided on the host PC monitor. This scanning procedure offers the great advantage to increase the number of measured points in those regions which are richer of spatial details (where the surface has the highest spatial frequency content), achieving a denser sampling there. The digitization accuracy has been experimentally assessed similarly to [5], and the results are reported here (Fig. 2). The position of 5000 laser spots projected over a plane posi-

tioned in the center of the working volume is analyzed. For a medium-size volume (0.8 m \times 0.6 m \times 1 m), the accuracy is of 0.27 mm \pm 0.0063 mm (rms error), averaged over 100 different experiments, with the two cameras at an angle close to 90° (86.9°). This figure scales linearly with the dimension of the working volume, and it is consistent with the accuracy reported in the motion analysis field [7]. The distribution of the measurement error is zero-mean and normally distributed [cf. Fig. 2(b)] as it is the outcome of several error sources (quantization error in the video cameras, numerical approximation in the cross-correlators, nonuniform spot appearance, etc.).

At the end of the digitization process, a set of N 3-D points is obtained $\{P_j = (x_i, y_j, z_j)\}$ [cf. Fig. 3(a)]. At the same time, through a standard CCD camera a color snapshot of the surface is acquired [cf. Fig. 3(b)], which is converted into a bitmap of

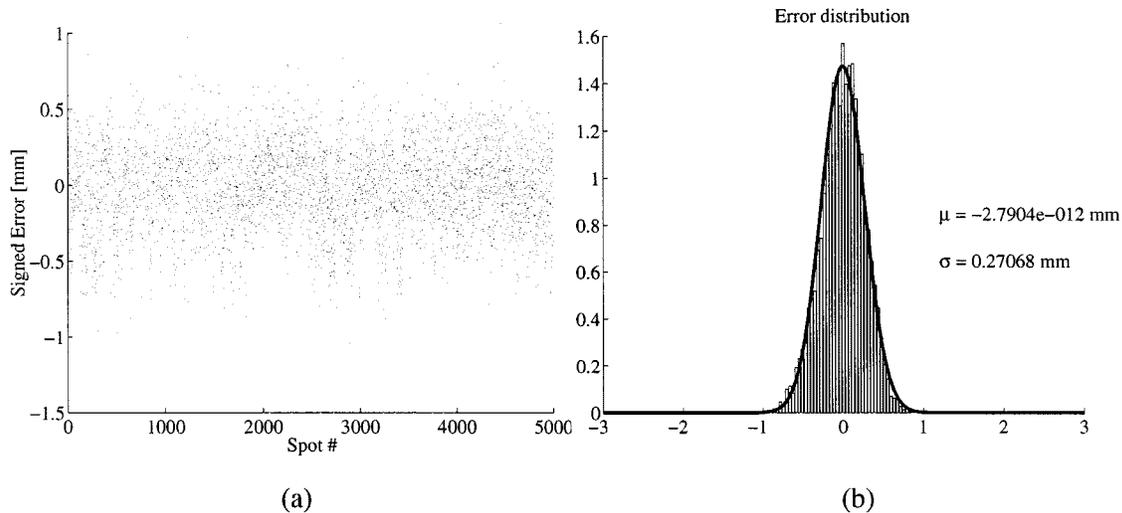


Fig. 2. (a) Distance of 5000 spots with respect to the best fitting plane. (b) The distribution of the error can be considered Gaussian; its standard deviation is assumed as a measure of the digitization rms error.

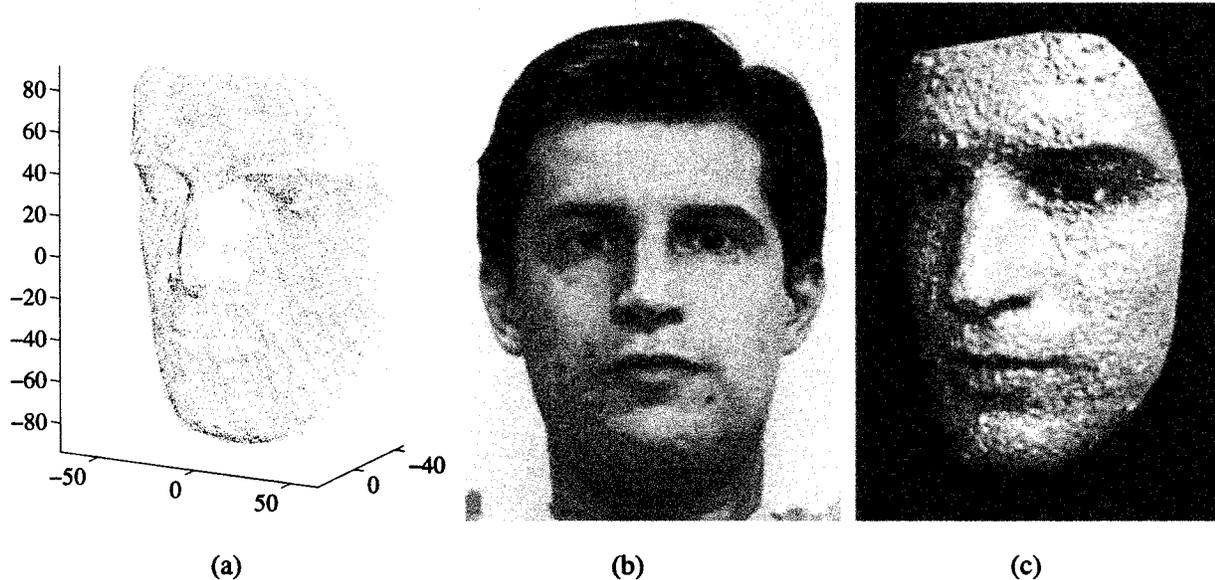


Fig. 3. (a) Ensemble of the $N = 17\,080$ 3-D points and (b) bitmap of Stefano's face. (c) A direct tessellation of the data points (e.g., through Delaunay triangulation) [33] would produce a jerky surface which is plotted with texture.

576×768 (PAL standard), with 24 bits/pixel. The sustained peak acquisition rate is of 100 spots/s. This rate decreases when the spot is not visible to the cameras. In real acquisitions, an average rate of 90–95 spots/s is obtained. By using arrays of laser pointers in place of a single pointer, this rate could be increased.

III. SURFACE RECONSTRUCTION THROUGH THE HRBF MODEL

The conversion of the cloud of 3-D points into a 3-D geometrical model is the critical operation of any 3-D scanner. Due to measurement noise a direct tessellation of the data points obtained by simply connecting them, would produce an undesirable wobbling surface [Fig. 3(c)]. Such a surface is useless to graphics, CAD or any other application, and the need of some sort of filtering is evident.

When the shape of the object surface is known, model fitting can be a valuable solution. In this approach a parameterized

surface model is progressively adapted to the range data by minimizing their distance from the surface (e.g., [8] for human faces). This approach is improved when features can be extracted from the range data [9] (cf. also [10], [11]). However these approaches are heavily based on predefined models and lack in generality. Generic parametric shapes (semi-parametric fitting), which are warped to fit the data topology, offer a more general solution [12]. The simplest of these is the two-dimensional lattice, proposed in the seventies in the connectionist domain, as a tool to represent multivariate data distributions. This model, called Self-Organizing Maps (SOMs) [13], has been more recently applied to reconstruct 3-D surfaces from range data [14]. Its limitation lies in the requirement that the surface topology is homomorphic to the lattice structure (usually a plane). When this is not the case, twisting of the plane and poor approximation arise. Superquadrics have been introduced as more powerful semi-parametric geometrical models [9].

Introducing a B-spline surface on the top of superquadric [9] or meshes [15] can further refine the obtained surface. This approach, termed “surface on surface,” gives rise to a weak hierarchical structure: the superquadric is adapted first to capture the overall shape of the data, and the spline patch is used to deform the superquadrics surface to better fit the data locally, by reproducing the surface details. To reconstruct complex surfaces, more than one superquadric can be employed [16].

All these approaches are iterative and cannot operate in real-time on the data points. A novel different approach is presented here aimed to produce 3-D models in real-time. The reconstruction is achieved with a sparse approximation with adaptive variable resolution spatial filtering, where the correlation between neighbor data points is exploited to eliminate noise introduced by the sampling process. It is based on the hierarchical radial basis function network (HRBF) model, proposed originally in the connectionist domain as a tool for multiscale signal processing [17]; it is extended here to surface approximation from scattered data. In the following, the overview of the HRBF model is presented first along with the robust determination of its parameters. The hierarchical, multiscale, structure is then introduced to make the approach fully adaptive both in resolution and scale.

A. Construction of a 3-D Surface Through Approximation

We assume that the surface can be represented as a function $S(P): R^2 \rightarrow R$. This assumption is motivated by the shape of the human face which can be described as a height field (2 1/2-D). However this does not lack in generality as full 3-D surfaces can be recovered by zipping multiple 2 1/2 D models [18]. Under this hypothesis, it is more convenient to reframe the data points collected with the scanning procedure as: $\{(x_j, y_j, z_j) | j = 1, \dots, N\}$ as the set $\{(P_j, S(P_j)) | P_j = (x_j, y_j) \in R^2; S(P_j) = z_j \in R; S: R^2 \rightarrow R\}$.

Our final goal is the reconstruction of the true surface up to the measurement error. That is, the distance between the points reconstructed by the HRBF model, and those sampled on the surface, should be distributed according to the measurement error. The reconstructed surface $S_r(P)$ should satisfy the following conditions:

$$\frac{1}{N} \sum_{j=1}^N (S_r(P_j) - S(P_j)) = E(v) = 0 \quad (1a)$$

as the measurement noise in zero mean (Section II), and

$$\frac{1}{N} \sum_{j=1}^N (S_r(P_j) - S(P_j))^2 = \sigma_v^2 \quad (1b)$$

which can be evaluated, for a given set-up, through the experiment reported in Section II. As the measurement error is uniform, (1) should be satisfied not only on the entire input domain, but also locally in sub-regions of it (cf. Section III-D).

B. HRBF Network

The HRBF model combines the output of many simple units to achieve the reconstruction of a complex surface. In particular,

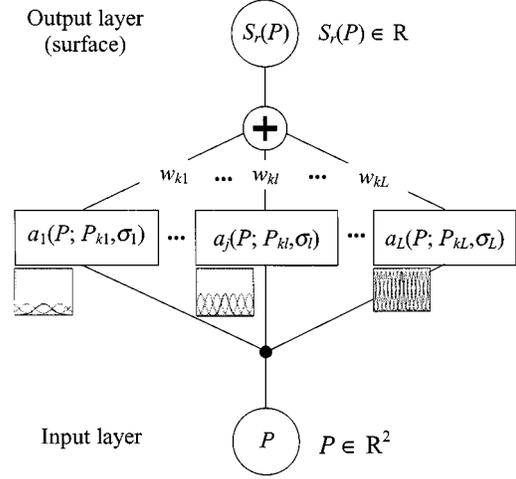


Fig. 4. HRBF network structure.

the HRBF network is composed of radially symmetric Gaussian units

$$g(P; \mu | \sigma) = \frac{1}{\sqrt{\pi}\sigma} e^{-((P-\mu)^2/\sigma^2)} \quad (2)$$

where $P, \mu \in R^2$ and $\sigma \in R$. In the HRBF model, the units are organized in layers (subnetworks), where the k th layer is composed of equally spaced Gaussians, which have the same standard deviation, σ_k (cf. Figs. 4 and 6). Therefore the surface, $S_r(P)$, is constructed by adding the contribution of a few grids of Gaussian functions, where, as shown in Fig. 5, each grid operates at a certain scale (or cut-off frequency).

Given a set, α , of parameters which characterize the HRBF network, the actual shape of the reconstructed surface (i.e., the output of the net), $S(P|\alpha)$, is

$$S(P|\alpha) = \sum_{l=1}^L \sum_{k=1}^{M_l} w_{kl} g(P; P_{kl} | \sigma_l) \quad (3)$$

where

- L number of grids;
- M_l number of Gaussian units in the l th grid;
- w_{kl} weight associated to the k th Gaussian in the l th grid;
- P_{kl} its position in the grid;
- σ_l standard deviation of all the Gaussians in the l th grid which determines the filtering scale of that grid.

The parameters $\alpha = \{L, M, \{w_{kl}\}, \{P_{kl}\}, \{\sigma_l\}\}$ determine the actual shape of $S(P|\alpha)$. Although these parameters could be determined by global optimization, e.g., [19], the time required to obtain a reasonable solution has suggested exploring alternative solution schemes. HRBF networks, in particular, offer a very fast solution as the determination of the parameters in (3) does not require any iteration, and it is performed with only local operations on the data points. This makes this approach particularly suitable to 3-D scanning application.

C. Determination of the Parameters in Each Layer

Each grid, l , of the HRBF model realizes a low-pass filter, which is able to reconstruct the surface up to a certain scale, determined by σ_l . It can be shown that σ_l and the spacing between

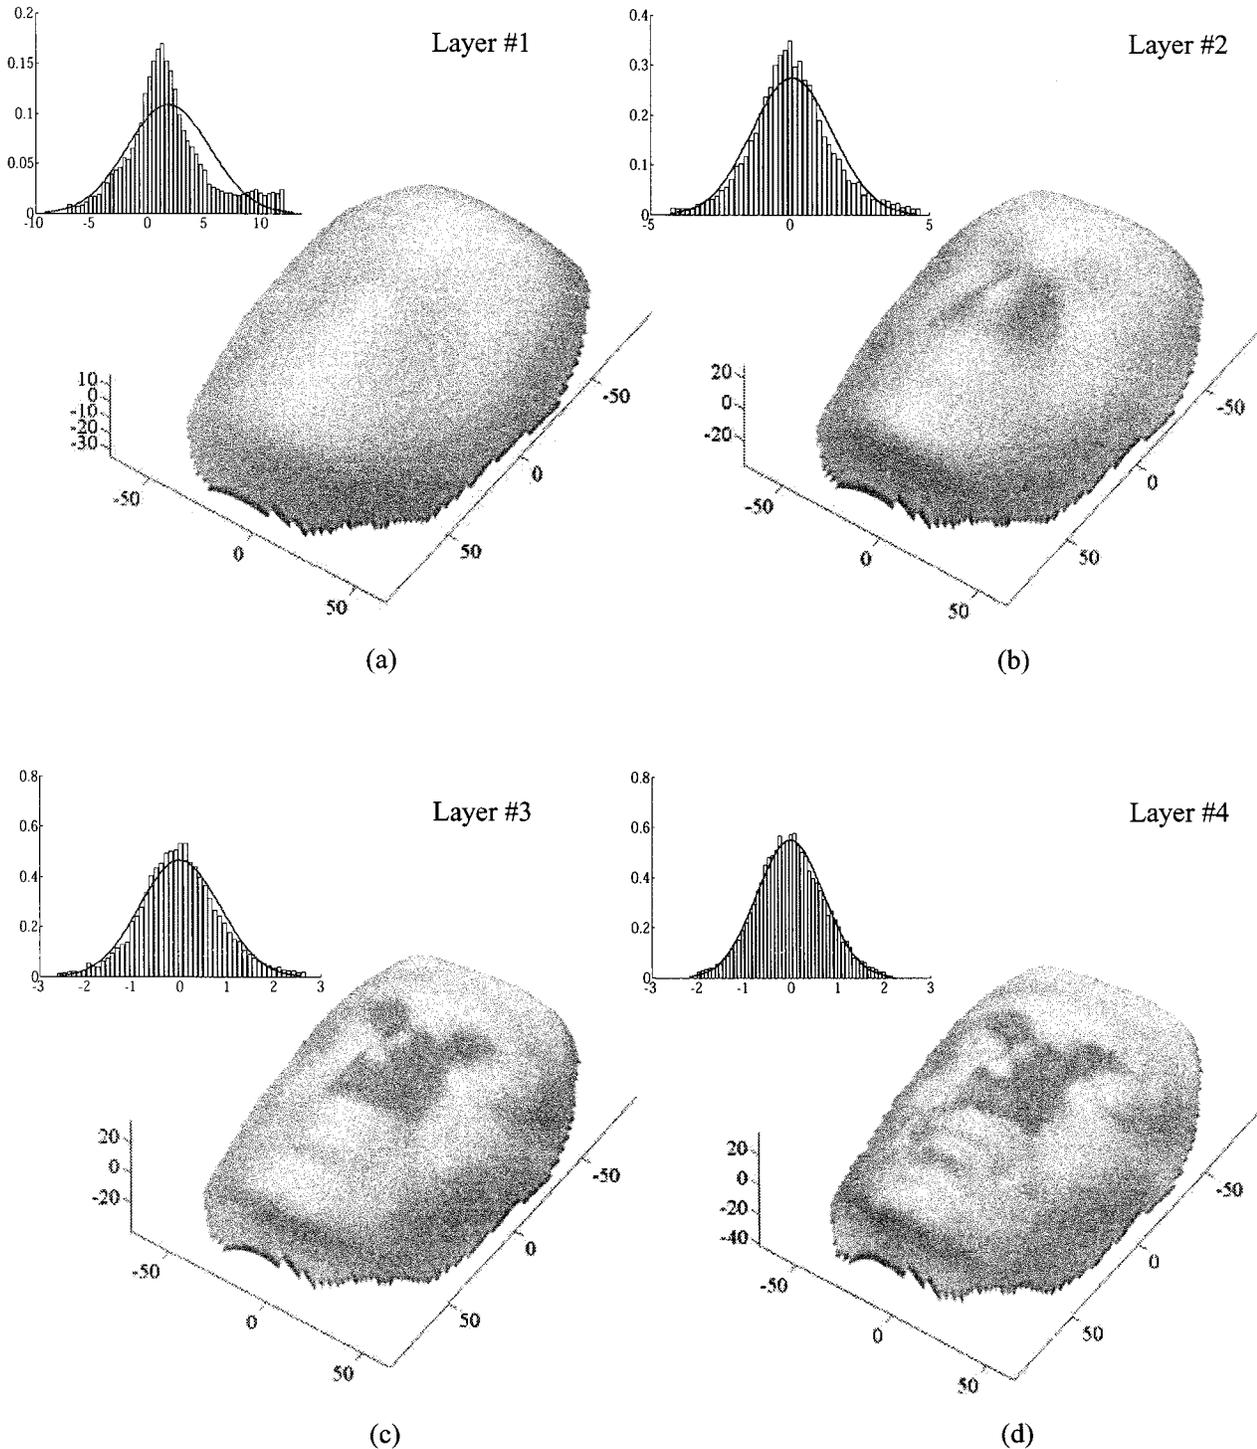


Fig. 5. Multiscale reconstruction of Stefano's face is reported here. It is obtained adding the contributions (a) of all the four layers; (b) of the first three layers; (c) of the first two layers. In (d), the output of the first layer is plotted. The residual distribution is plotted besides each layer in the form of a histogram whose area is normalized to one.

two consecutive Gaussians on the same grid, ΔP_l , are related with

$$\sigma_l = 1.465 \Delta P_l. \quad (4)$$

This relationship is obtained by accepting a maximum attenuation in the pass band of -3 dB and a minimum attenuation

in the stop band of -40 dB [20].¹ Different attenuation values lead to a different proportionality constant between ΔP_l and σ_l . To apply the Gaussian filter, the data should be equally sampled in correspondence to the grid crossings $\{k\}$. That is, the set $\{z_{k,l}\} = \{S(P_{kl})\}$ should be available. Unfortunately this

¹These values produce a cut-off frequency of $v_{\text{cut-off}} = 0.1874/\sigma_l$ and a transition band of $[0.1874/\sigma_l, 0.7327/\sigma_l]$.

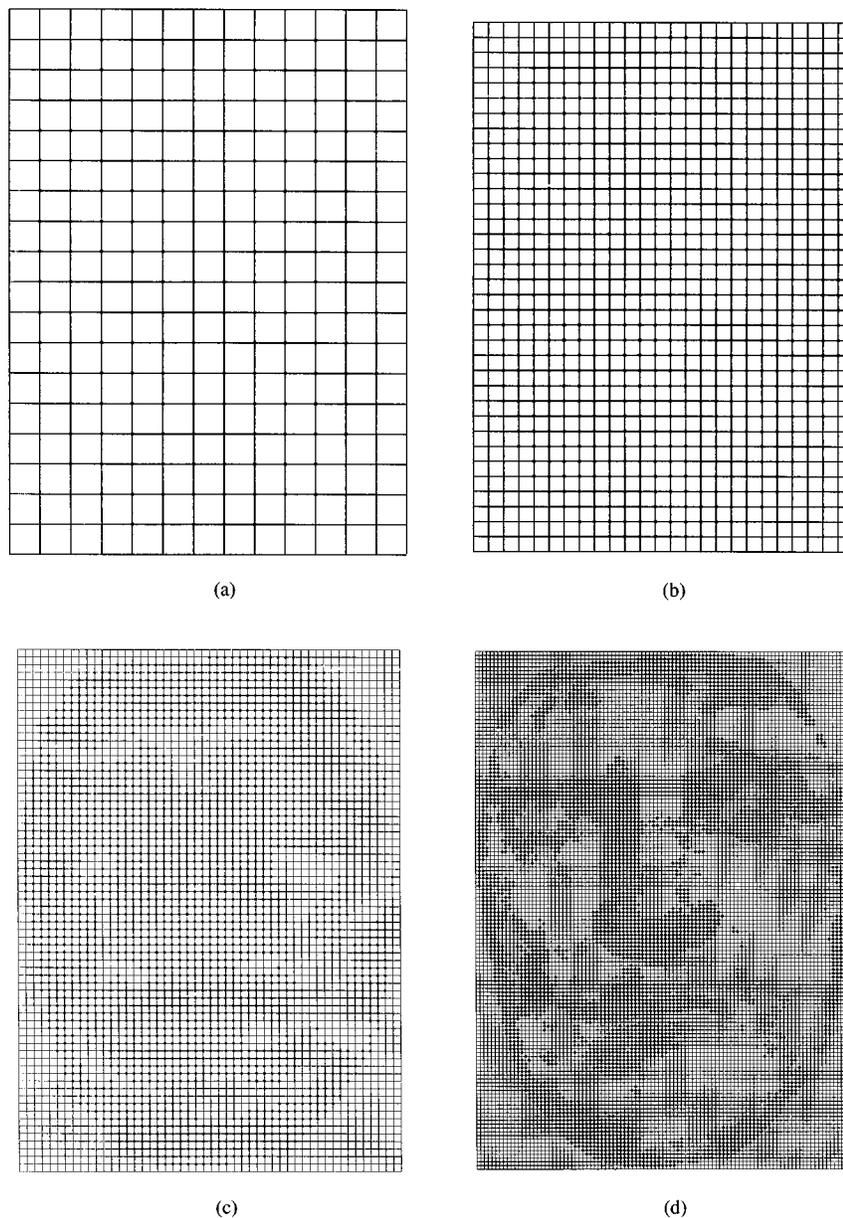


Fig. 6. Position of the Gaussian units in the four layers for Stefano reconstruction. Cluster of Gaussians are inserted in the higher layers determined an increase in the local frequency of the reconstruction.

is not the case here where the data are unevenly sampled. However, since many points are usually digitized (surface oversampling), a reliable estimate of $S(P_{kl})$ can be obtained through the following weighted mean estimate (cf. also [21])

$$\tilde{S}(P_{kl}) = \frac{\sum_{P_r \in A(P_{kl})} S(P_r) e^{-\|P_{kl} - P_r\|^2 / \sigma_l^2}}{\sum_{P_r \in A(P_{kl})} e^{-\|P_{kl} - P_r\|^2 / \sigma_l^2}} \quad (5)$$

where $A(P_{kl})$ is the Receptive field associated to the k th Gaussian in the l th grid. It is set, somewhat arbitrarily, as the square region centered P_{kl} of side ΔP_l . The estimate in (5) is carried out locally on the input space, and it can be parallelized to achieve quasi-real-time processing.

The grid filter can be now written as

$$S_r(P) = \sum_{k=1}^{M_l} \tilde{S}(P_{kl}) g(P; P_{kl} | \sigma_l) \Delta P_l^2. \quad (6)$$

Comparing (6) with (3), it can be demonstrated that the parameters $\{w_{kl}\}$ can be obtained simply as $\tilde{z}(P_{kl}) \Delta P_l^2$ [20].

D. Hierarchical Multiscale Structure

If only one grid was adopted, a serious drawback is introduced: the Gaussian scale should be small enough to resolve the finest details. This requires very dense packing of the Gaussians also in those regions where the details can be resolved at a coarser scale, causing a waste of resources and overfitting in

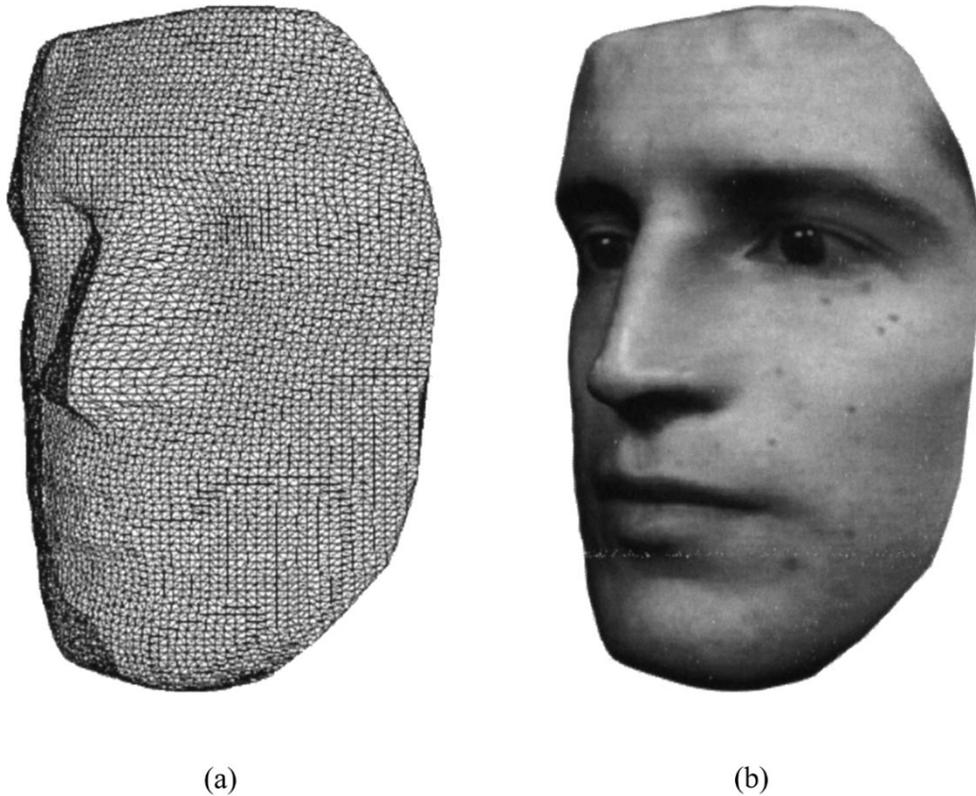


Fig. 7. (a) Tessellation resulting from sampling the HRBF surface at 1mm step. (b) The final result where texture has been applied.

those space regions. Moreover, as the sampled point local density reflects the local richness of details, those regions which lack high-frequency details will contain points spaced apart far, and receptive fields $A(P_{kl})$ may come out empty there. A better solution would be to adaptively allocate the Gaussian units, with an adequate scale and, consequently, adequate receptive fields, in the different space regions. The novel idea is to analyze the local residual and to stack noncomplete grids over a first grid at a coarse scale, obtaining what is termed a sparse approximation (cf. Fig. 6). The first grid will output a rough estimate of the surface, $a_1(P)$ [Fig. 5(a)] as

$$a_1(P) = \sum_{k=1}^{M_l} \tilde{S}(P_{kl})g(P; P_{kl}|\sigma_l) \Delta P_l^2. \quad (7)$$

The residual $\{r_1(P_j)\}$ is computed for each sampled data point, $(P_j, S(P_j))$, as

$$r_1(P_j) = S(P_j) - a_1(P_j). \quad (8)$$

This residual will be the input to a second grid which features a smaller scale than the first one; somewhat arbitrarily we have chosen $\sigma_2 = \sigma_1/2$ which is the same choice made in wavelet decomposition [22]. This second grid does not need to reconstruct the original surface but only the residual one. Its output, $a_2(P)$, will be a reconstruction of the residual $r_1(\cdot)$ at the scale σ_2 , and it will provide a second residual

$$r_2(P_j) = r_1(P_j) - a_2(P_j) = S(P_j) - a_1(P_j) - a_2(P_j). \quad (9)$$

This grid will not be full, but Gaussians will be inserted only when a poor approximation is given. This is evaluated through the residual itself: a Gaussian is inserted in the grid crossing P_{kl} only if

$$\frac{\sum_{P_r \in A(P_{kl})} r_1^2(P_r)}{N_k} < \sigma_v^2 \quad (10)$$

where σ_v^2 is the noise variance and N_k is the number of sampled points which belong to $A(P_{kl})$. A mathematical proof of HRBF convergence can be found in [23] and goes beyond the scope of this paper. Grids are created one after the other, until the condition in (10) is not true over the entire input domain [Figs. 5(a)–(d)].

IV. CONVERSION INTO A MESH AND COLOR APPLICATION

When required, the reconstructed continuous surface can be sampled as densely as desired [at 1 mm in Fig. 7(a)] to obtain a mesh of triangular elements. The sampling could be optimized according to topological or geometrical criteria to obtain a lighter mesh [24]. When the color appearance is important (e.g., in computer graphics applications), a bitmap obtained from the snapshot of the surface is applied. This is achieved by using standard texture mapping procedures after having aligned it to the mesh through a semi-automatic graphical tool [25], [26]. The final result is a 3-D very realistic reconstruction of human face [Figs. 2(b) and 7(b)]. A detailed description of this procedure goes beyond the scope of this paper.

TABLE I
PERFORMANCE INDEXES AND PARAMETERS OF EACH LAYER OF THE
HRBF NETWORK

#grid	Error mean [mm]	Error std [mm]	#gauss	$v_{\text{cut-off}}$ [Hz]	Grid dimension
1	1.89	3.66	282/315	0.015	15x21
2	0.07	1.45	1124/1333	0.03	31x43
3	-0.005	0.96	3020/5394	0.06	62x87
4	-0.002	0.72	5929/21750	0.12	125x174

V. RESULTS

The system has been intensively used in scanning human body parts and faces in particular. This is a particularly difficult task because the face's spatial frequency content is highly variable. Three to four layers are usually sufficient. A typical result is reported in Figs. 5 and 7 for Stefano's face. In the set-up used the angle between the cameras was of $\approx 50^\circ$ with an estimated rms measurement error of 0.65 mm. The surface is reconstructed starting from $N = 17080$ data points collected in 3 min. As can be seen in Fig. 6, in the higher layers the units are inserted in clusters, only in those regions where the face has the highest frequency content: the HRBF network automatically adapts its local cut-off frequency and resolution to the local frequency of the data. Only 10 355 Gaussian units have been used versus the 21 750 units, which were required by a complete Gaussian filter at the smallest scale. The quantitative results are reported in Table I, where the residuals and the parameters of each layer are shown. As it can be seen, the mean error on the reconstructed surface (bias) is already very close to zero in the second layer, while the standard deviation decreases down to 0.72 mm in the fourth layer. The distribution of the residuals becomes progressively Gaussian shaped and approaches that of the digitization error as can be seen in Fig. 5.

VI. DISCUSSION

The 3-D scanner presented here is based on a digitization module characterized by scalability, flexibility and accuracy. In fact, acquisition of objects of complex topology can be easily accomplished by zipping multiple acquisitions carried out with a single pair of cameras placed in different positions and orientations [27] or increasing the number of video cameras used simultaneously. The use of pattern recognition to compute the spot's centroid gives a spot accuracy higher than that of commercial 3-D scanners [1]. This figure decreases when the angle between the pair of video cameras, used to reconstruct the 3-D position of the spot, differs from 90° with approximately a quadratic curve [28], [29]. In real situations, where the angle between the cameras has to be negotiated with the amplitude of the common field of view of the two cameras, an angle between 40° and 60° is usually adopted, with an accuracy decrease between 1.5 and 3 times. When higher accuracy is required, smaller working volumes can be set-up or the angle between the cameras increased.

Given the relatively long acquisition time, such a high accuracy can be reached only when still objects are scanned. Therefore, when scanning human faces, maximum comfort of the

subject has to be provided to avoid nonintentional movements, which do decrease the accuracy. These small movements could be eliminated by using at least three circular markers solid with the head, of the same dimension of the laser spot; they would be recognized by the pattern matching processor, and used to compensate head motion. Given the high accuracy achieved, this procedure has not been adopted, but it could suggest a possible implementation in which moving objects can be scanned.

This high accuracy is preserved in the reconstruction of the 3-D surface, which is achieved by a novel hierarchical model: the HRBF network. The requirement of only operations carried out locally on the data points with no iterations is the main characteristic, which makes it suitable to real-time implementations. Few algorithms aimed to real-time, can be found in the literature. The closest approach is the hierarchical reconstruction based on B-splines [21] where the reconstructed surface is the sum of a set of intermediate surfaces generated through a hierarchy of control grids. The main difference is that, in the HRBF model, the grids in the superior layers are not complete, but Gaussian units are inserted in clusters where the residual is over threshold, with a large saving in resources. The same adaptive unit allocation schema offers to HRBF an advantage versus multiresolution analysis (MRA) carried out through wavelet decomposition [22]. The MRA works in a direction opposite to HRBF as it produces a hierarchy of approximating surfaces from fine to coarse. Each surface is the sum of a very low-frequency surface and a set of surfaces called details, where each of them is obtained as a linear combination of basis functions. Small coefficients can be set to zero, giving a sparse approximation of the surface. The MRA structure closely resembles the reconstruction in the HRBF model (cf. Fig. 4), although the coefficients are computed differently. In fact, the k th detail surface in the HRBF is computed from the difference between the surface approximated at the lower level and the surface samples. This allows recovering errors in the computation, e.g., due to finite precision, or the effect of zeroing the coefficients in the previous layers. In MRA instead, the detail is computed directly by projecting the residual at the higher layer onto an adequate basis. This procedure, although faster as it does not require computing the residual for each data point, does not allow correcting for any error in the computation.

Another advantage of the HRBF model is its minimalism, as it requires a single basis, the Gaussian, versus the four bases required by MRA, at least in the most common biorthogonal setting [30].

Finally, the scale of each layer in HRBF does not have to be necessarily halved as in MRA, but can be adapted to the data frequency content. Moreover, if a norm different from L^2 were required, it can be introduced by simply changing the norm in which the error and the residuals are evaluated. For example, if outliers are often collected, a more suitable norm to measure the residual would be L^1 [20].

The limitation of HRBF networks, and of basis function approximation in general, lies in the use of filtering, which makes the reconstruction of structured objects, with sharp edges and corners, difficult. For these, procedures which reconstruct the surface through an intermediate stage based on extraction of objects' geometrical primitives can be more powerful [31], [32].

VII. CONCLUSION

Overall, the 3-D scanner system presented here is able to furnish a detailed 3-D reconstruction of surfaces by using simple hardware components in conjunction with adaptive data processing. The operations involved both in spot detection and mesh construction are performed locally on the data set. They can be easily parallelized and implemented on general purpose, low-cost, processing boards (e.g., FPGA) to obtain a low cost system which shall be seen as a standard measurement device of next generation graphical workstations.

REFERENCES

- [1] M. Petrov, A. Talapov, T. Robertson, A. Lebedev, A. Zhilayev, and L. Polonskiy, "Optical 3-D digitisers: Bringing life to the virtual world," *IEEE Comput. Graph. Appl.*, vol. 18, no. 3, pp. 28–37, 1998.
- [2] H. Urban, "Ultrasonic imaging for industrial scene analysis," in *Sensor Devices and System for Robotics*. ser. NATO ASI Series, Computer and System Sciences, A. Casals, Ed. New York: Springer-Verlag, 1989, vol. 52, pp. 187–194.
- [3] N. A. Borghese, G. Ferrigno, G. Baroni, R. Saverè, S. Ferrari, and A. Pedotti, "AUTOSCAN: A flexible and portable scanner of 3-D surfaces," *IEEE Comput. Graph. Appl.*, vol. 18, no. 3, pp. 2–5, May/June 1998.
- [4] G. Ferrigno and A. Pedotti, "Modularly expansible system for real-time processing of a TV display, useful in particular for the acquisition of coordinates of known shapes objects," U.S. Patent 4 706 296, 1990.
- [5] E. Gagnon and J. F. Rivest, "Laser range imaging using the self-mixing effect in a laser diode," *IEEE Trans. Instrum. Meas.*, vol. 48, no. 3, pp. 693–699, 1999.
- [6] N. A. Borghese, M. Di Rienzo, G. Ferrigno, and A. Pedotti, "Elite: A goal-oriented vision system for moving objects detection," *Robotica*, vol. 9, pp. 275–282, 1990.
- [7] N. A. Borghese and P. Cerveri, "Calibrating a video camera pair with a rigid bar," *Pattern Recognit.*, vol. 33, no. 1, pp. 81–95, 2000.
- [8] K. Waters and D. Terzopoulos, "Modeling and animating faces using scanned data," *J. Vis. Comput. Animat.*, vol. 2, pp. 123–128, 1991.
- [9] D. Terzopoulos and D. Metaxas, "Dynamic 3-D models with local and global deformations: Deformable superquadrics," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 13, pp. 703–714, July 1991.
- [10] Y. Lee, D. Terzopoulos, and K. Waters, "Realistic modeling for facial animation," in *Proc. SIGGRAPH 95*, Los Angeles, CA, August 1995, pp. 55–62.
- [11] F. Pighin, J. Hecker, D. Lischinski, R. Szeliski, and D. H. Salesin, "Synthesizing realistic facial expressions from photographs," in *Computer Graphics Proceedings*, ser. Annual Conf. Series: ACM SIGGRAPH, 1998, pp. 151–160.
- [12] R. Mencl and H. Müller, "Interpolation and approximation of surfaces from three-dimensional scattered data points," in *Eurographics '98 State of the Art Reports*, Lisbon, Portugal, 1998, pp. 51–67.
- [13] T. Kohonen, *Self-Organizing Maps*. New York: Springer-Verlag, 1995.
- [14] A. Baader and G. Hirzinger, "Three-dimensional surface reconstruction based on self-organizing feature map," in *Proc. 6th Int. Conf. Advan. Robotics*, Tokio, 1993, pp. 273–278.
- [15] V. K. Krishnamurthy and M. Levoy, "Fitting smooth surfaces to dense polygonal meshes," in *Computer Graphics Proceedings*, ser. Annual Conf. Series: ACM SIGGRAPH, 1995, pp. 151–160.
- [16] D. De Carlo and D. Metaxas, "Shape evolution with structural and topological changes using blending," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 20, pp. 312–324, Nov. 1998.
- [17] N. A. Borghese and S. Ferrari, "Hierarchical RBF networks and local parameters estimate," *Neurocomputing*, vol. 19, pp. 259–283, 1998.
- [18] C. Dorai, G. Wang, and A. K. Jain, "Registration and integration of multiple object views for 3-D model construction," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 20, no. 1, pp. 83–89, Jan. 1998.
- [19] T. Poggio and F. Girosi, "Regularization algorithms for learning that are equivalent to multilayer networks," *Science*, vol. 247, pp. 978–981, 1990.
- [20] N. A. Borghese and S. Ferrari, "Automatic acquisition of 3-D objects through portable modular system," *Proc. IMTC*, vol. 3, pp. 1823–1827, 1999.
- [21] S. Lee, G. Wolberg, and S. Y. Shin, "Scattered data interpolation with multilevel B-splines," *IEEE Trans. Visualization and Computer Graphics*, vol. 33, no. 3, pp. 228–244, July–September 1997.
- [22] S. G. Mallat, "A theory for multiresolution signal decomposition: The wavelet representation," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 11, pp. 674–493, July 1989.
- [23] N. A. Borghese, M. Maggioni, and S. Ferrari, "Multi-scale approximation with hierarchical radial basis functions networks," Dept. Electronics and Information, Politecnico of Milano, Tech. Rep. 99-67, 1999.
- [24] P. S. Heckbert and M. Garland, "Survey of polygonal surface simplification algorithms," Carnegie Mellon Univ., Pittsburgh, PA, Tech. Rep., 1997.
- [25] F. Weinhaus and V. Devarajan, "Texture mapping 3-D models of real-world scenes," *ACM Comput. Surv.*, vol. 29, no. 4, pp. 325–365, 1997.
- [26] P. Rigioli, "Sistema per acquisizione di modelli 3-D di volti umani e loro animazione ad alto grado di realismo mediante motion-capture," Laurea thesis, Dept. Elect. Inform., Politec. Milano, Italy, 2000.
- [27] F. Bernardini, J. Mittleman, C. Silva, and G. Taubin, "The ball-pivoting algorithm for surface reconstruction," *IEEE Trans. Comput. Graph.*, vol. 5, no. 4, pp. 349–362, 1999.
- [28] N. A. Borghese and G. Ferrigno, "An algorithm for 3-D automatic movement analysis by means of standard tvcameras," *IEEE Trans. Biomed. Eng.*, vol. 37, pp. 1221–1225, 1990.
- [29] J. Oliensis, "Rigorous bounds for two-frame structure from motion," in *Proc. ECCV '96, LNCS-Series*. Berlin, Germany: Springer-Verlag, 1996, vol. 1065, pp. 184–195.
- [30] A. Cohen, "Biorthogonal wavelets," in *Wavelets: A Tutorial in Theory and Applications*, C. K. Chu, Ed. New York: Academic, 1992.
- [31] H. Hoppe, "Surface reconstruction from unorganized points," Ph.D. dissertation, Dept. Comput. Sci. Eng., Univ. Washington, Seattle, June 1994.
- [32] P. E. Debevec, C. J. Taylor, and J. Malik, "Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach," in *Computer Graphics Proceedings*, ser. Annual Conf. Series: ACM SIGGRAPH, 1995, pp. 151–160.
- [33] R. A. Dwyer, "A faster divide-and-conquer algorithm for constructing the Delaunay triangulation," *Int. J. Comput. Inform. Sci.*, vol. 9, no. 3, pp. 219–242, 1997.

N. Alberto Borghese (M'97) received the "laurea" degree in electrical engineering (with 100/100 cum laude) from Milan Politecnico, Milan, Italy, in 1984. He joined the Institute of Neuroscience and Bioimaging—CNR, Milano, in 1988 where he is currently Director of the Laboratory for Human Motion Analysis and Virtual Reality (MAVR). He has been Visiting Scientist at Center for Neural Engineering of USC, Los Angeles in 1991–1992; and at the Department of Electrical Engineering of Caltech, Pasadena, in 1992. His research interests include quantitative human motion analysis, modeling and synthesis in virtual reality, and artificial learning systems.

Stefano Ferrari received the "laurea" degree in computer science from Università degli Studi di Milano, Italy in 1995. Currently, he is pursuing the Ph.D. degree in the Department of Electronics and Information, Politecnico di Milano.

He collaborated with Bioengineering Center, Fondazione ProJuventute, Milano and MAVR, INB in 1996–1997. His research interests are soft-computing algorithms and their high-level architectural specification; and their application in the areas of image processing and virtual reality.