PERGAMON

Pattern Recognition 34 (2001) 2497–2513

PATTERN RECOGNITION
THE JOURNAL OF THE PATTERN RECOGNITION SOCIETY

www.elsevier.com/locate/patcog

3D wavelet-based multiresolution object representation

Luis Pastor[a,*], Angel Rodríguez[b], J. Miguel Espadero[a], Luis Rincón[a]

[a]*Department of Ciencias Experimentales y Tecnológicas, Universidad Rey Juan Carlos, Móstoles, Spain*
[b]*Department of Tecnología Fotónica, Universidad Politécnica de Madrid, Boadilla del Monte, Spain*

Received 17 March 2000; accepted 25 September 2000

Abstract

This paper presents a technique for computing multiresolution shape models of 3D objects acquired as clouds of 3D points. The procedure is fully automated and is able to compute approximations for any object, overcoming sampling irregularity if present (sampling irregularity is a common feature of most 3D acquisition techniques; a typical example is stereo vision). The method described here starts by computing an intermediate mesh that meets the subdivision connectivity requirement needed to allow the computation of the wavelet transform. The mesh is then adjusted to the 3D input data using an iterative deformation process. Finally, a spherical wavelet transform is computed to obtain the object's 3D multiresolution model. This paper shows a number of real objects acquired with different techniques, including hand-held 3D digitizers. The paper also gives some examples of how multiresolution representations can be used in tasks such as acquisition noise filtering, mesh simplification and shape labelling. © 2001 Pattern Recognition Society. Published by Elsevier Science Ltd. All rights reserved.

Keywords: 3D Modeling; Multiresolution representations; Wavelet-based representations; 3D vision, Automatic shape extraction

1. Introduction

Representing and recognizing objects are two of the main goals of computer vision systems, and are essential stages for understanding the environment's structure. Object representation or modeling aims at creating thorough descriptions of the entities that integrate the perceived scene. Sometimes modeling is an objective in itself, as in reverse engineering, where data acquired from real objects is used as an input to a CAD system. Usually, in computer vision, computing the representation of an item is an intermediate stage of the vision system, yielding results used by other processes that perform more abstract operations on the data acquired from the scene objects. An illustrative example is object recognition, an important application itself, and which may also be a key stage in many areas such as robotics, quality control, tracking, etc.

Interest in object representation can be tracked to the beginnings of computer vision research [1]. Early work concentrated mainly on representing 2D entities. Nevertheless, the development of 3D sensors and sensing techniques has stimulated the interest in representing 3D objects. The large amount of information involved and the complexity and speed requirements of the processing techniques demand the development of efficient and powerful object representation methods. Many of the 3D sensing techniques (particularly in computer vision) provide range data without much structural information, which makes the modeling problem more challenging.

3D geometric models are composed of a large number of primitive elements, particularly when accuracy is a requirement. Model complexity strongly affects either the hardware needed for a particular application, the kinds of operations that the system can perform or the maximum complexity of the objects that can be considered while keeping processing time reasonable. In fact, this is a situation which is often found; typically, image

* Corresponding author. Tel.: +34-91-664-7442; fax: +34-91-664-7490.

E-mail addresses: lpastor@escet.urjc.es (L. Pastor), arodri@dtf.fi.upm.es (A. Rodríguez), jespa@dtf.fi.upm.es (J.M. Espadero), lrincon@escet.urjc.es (L. Rincón).

PII: S0031-3203(00)00170-9

processing and computer vision applications combine large data sets and small processing times. It is well known that starting by processing reduced resolution versions of the input image may result in large computational savings. There are also specific tasks, such as edge detection or texture classification, that benefit from analyzing input data at different scales. It is therefore not surprising that hierarchical and multiresolution techniques have always raised interest in the computer vision community [2,3].

The development of the wavelet transforms theory has spurred new interest in multiresolution methods, and has provided a more rigorous mathematical framework. Wavelets give the possibility of computing compact representations of functions or data. Additionally, they allow variable degrees of detail or resolution to be achieved, and they are attractive from the computational point of view [4–6]. All these features make them appear as an interesting tool to be used for representing 3D objects.

Several 2D computer vision applications using wavelets have recently been developed [7,8]. Also, some 3D wavelet-based object modeling techniques have appeared lately in the computer graphics environment [9,10] Nevertheless, similar methods have not been used in 3D computer vision environments, possibly for two main reasons:

- Wavelet representations are not translation invariant. This is a significant problem, and different approaches have been described in the literature to cope with it [11,12].
- The sensors used in 3D vision provide data in a way which is difficult to analyze with standard wavelet decompositions: many 3D sensing techniques provide sparse measurements which are irregularly spread over the object's external surface. This is also important, because sampling irregularity prevents the straightforward extension of 1D or 2D wavelet techniques.

Even though the lack of translation invariance is an important drawback, we believe that multiresolution object representations have a bright future in 3D computer vision for several reasons, such as:

- Bottom-up scene analysis methods essentially attempt to create hierarchical symbolic representations. Wavelets are excellent for creating hierarchical geometric representations, which can be useful in the image data analysis process.
- Going to 3D implies an important increase in complexity. Wavelet decompositions can provide alternative domains in which many operations can be performed effectively.

This paper presents a fully automated method for computing spherical wavelet-based multiresolution descriptions of objects acquired as clouds of 3D points. Two important features of the method are that it does not impose any restriction on the input data, which can be as sparse and irregularly sampled as desired, and that it is able to establish structural relationships between data points. Within this paper, Section 2 gives an overview of object representation techniques in computer vision. Section 3 describes spherical wavelets. Section 4 describes a method to overcome the samples' possible lack of neighborhood relationships as well as the sampling irregularity. Section 5 presents some multiresolution representations computed from real 3D data. Section 6 explores some application areas, such as digitization noise removal, mesh simplification and global shape extraction. Last, Section 7 summarizes the paper's main conclusions.

2. Representing 3D objects

Depending on the system's purpose, the objective of the modeling stage can either be to collect a set of descriptors characterizing the acquired object with respect to the system's universe of discourse, or to generate an exhaustive description that allows the object representation to be manipulated afterwards.

It is generally accepted that a good modeling system should have the following properties [13,14]:

- Expressive richness.
- Stability in presence of errors or noise in the input data.
- Ability to deal with occlusions, which implies some sort of representation locality.
- Capability to reflect physical measures of the object's shape.
- Efficiency.

Many approaches have been followed to model 2D or 3D objects. For planar shapes, moments and Fourier descriptors have been used frequently [15]. For 3D systems, the following approaches are classical:

Volumetric representations: Spatial occupancy has been one of the first techniques used for representing 3D objects [16]. The main drawback of these techniques is that they are highly inefficient in the cases where only the objects' external appearence is relevant, since the ratio between "occupied" and "empty" cells is very low. Nevertheless, some novel developments using this approach have appeared recently [17,18].

Facet-based meshes: The surface of the objects is described by a set of facets,[1] edges and vertices, where the

[1] Sometimes referred to as tiles or patches.

vertices correspond to the points of the original description. In general, the order of the polynomials representing the facets varies in different approaches [19,20]. For example, it is possible to find planar, cubic, bicubic, quadric or superquadric facets. Generally, methods using higher-order facets achieve better mesh compression rates.

Deformable models: This group of techniques use meshes similar to those used by facet-based methods. The main difference between both is how the mesh is computed: deformable models start generating a mesh by covering a reference surface of variable degree, such as a tessellated sphere [21] or a hyperquadric [22]. Then, an iterative minimization process is applied in order to adjust the mesh's nodes to the 3D input data points, deforming the original reference surface's geometry. Most of the methods propose using similar deformation processes, taking the input shape's features into account, while trying to maintain some local or regular properties of the initial mesh in order to achieve a stable solution [23,24].

Grammatical or structural models: These techniques assume that it is possible to generate object descriptions using a hierarchy of geometric primitives at different abstraction levels, mimicking some theories about object recognition processes in the human vision system [14,25,26]. The hierarchy's coarser levels can be used as indexes to accelerate discriminations between very different objects, whereas the finer levels can be used for discriminating between similar objects [27–29]. Several methods have been suggested following this approach, differing in the primitives used and the way to connect them at the different hierarchy's levels.

Aspect graphs: Storing several views (2D projections) of the 3D shapes connected by a graph to preserve relationships among them allows 3D object recognition to be performed by matching the input 2D image with the available views [30,31].

3*D Invariants*: Extracting 3D invariants allows 3D shapes to be described by feature vectors in a similar way to classical 2D approaches [32–34]. Like some of the previous methods, this approach is only adequate for recognition purposes.

During the last years, some wavelet-based 3D modeling techniques have been applied to the computer graphics arena.

For example, Lounsbery et al. [35,36] present a new class of wavelets, based on subdivision surfaces, strongly extending the class of representable functions: whereas previous two-dimensional methods were restricted to functions defined in $\mathbb{R}^2$, the subdivision schemes presented in Refs. [35,36] may be applied to functions defined on compact surfaces of arbitrary topological type. They present several applications of their work, including smooth level-of-detail control for graphics rendering, compression of geometric models and animation previewing. Following Lounsbery et al.'s work, Schröder and Sweldens [9], have presented a mathematical framework for applying wavelet constructions for scalar functions defined on the sphere. They use the *lifting scheme* [37,38], to construct biorthogonal wavelets with custom properties, which they use in computer graphics environments for applications such as compressing topographic data and BDRF[2] [9], mapping synthetic textures over spherical surfaces [39], interactive mesh editing [40], etc.

Lounsbery et al.'s and Schröder and Sweldens' contributions share a common drawback: they depart from 3D meshes with data uniformly distributed over the whole object's surface and which meets the criterion of subdivision connectivity.[3] Eck et al. [41] propose a method for computing multiresolution representations of irregularly sampled meshes without subdivision connectivity, such as those provided by laser 3D digitizers. Eck et al.'s method is based on the approximation of an arbitrary initial mesh M to another mesh M^J that has subdivision connectivity and is guaranteed to be within a specified tolerance from M. They start by partitioning the original triangular mesh into a small number of regions covering the whole object following the concepts of Voronoi diagrams. Then, a Delaunay-like triangulation is constructed by computing the harmonic map h_i of each Voronoi tile τ_i into an appropriate planar polygon P_i. The inverse of h_i provides a parameterization of τ_i over P_i which is used to construct a low resolution Delaunay-like triangulation K^0 of the original mesh. Next, they define a local parameterization of the triangular regions over triangular facets of a base mesh K^0, again using harmonic maps. The local parameterizations are made to fit together continuously, in the sense that they define a globally continuous parameterization $\rho{:}K^0 \to M$. Finally, making a series of J recursive 4-to-1 splits on each of the faces of K^0 they get a triangulation K^J of K^0 with subdivision connectivity. The remesh M^J is obtained by mapping the vertices of K^J into $\mathbb{R}^3$ using the parameterization ρ, and constructing an interpolating mesh.

Lounsbery et al.'s and Schröder and Sweldens' techniques address the modeling problems found in 3D computer graphics environments. The application of similar ideas to computer vision environments needs to consider a number of aspects:

(1) Many 3D vision techniques yield sparse, irregularly sampled measurements where most of the structural information is lost. Typically, stereo vision techniques return 3D low-level primitives, such as isolated measurements within a point cloud or segments

[2] Bidirectional reflectance distribution functions.

[3] A mesh meets the *subdivision connectivity* condition if it can be generated from a base mesh through a 1–4 subdivision process.

belonging to the different object's contours rather than triangular meshes, either regularly or irregularly sampled.

(2) In computer vision applications, there are other sources of information in addition to *geometry* alone: color or other spectral information, surface texture, relative motion, etc. All of them give valuable information that can be used to make deductions about the perceived scene's structure. Consequently, each sample from the object's surface may contain, besides geometric information, additional data corresponding to the surface patch surrounding the sample point.

The method presented here specifically addresses these requirements, overcoming sampling irregularity with a fully automated procedure and representing object features in an efficient way. Since the scattered points' surface approximation is carried out by deforming a surface with *subdivision connectivity*, a subdivided icosahedron, there is no need to regularize the mesh. With respect to Eck et al.'s, our approach adds the possibility of processing 3D data clouds which do not have any additional structural information.

The following two sections describe spherical wavelets basis, as introduced by Schröder and Sweldens [9,39] and the multiresolution representation technique described in this paper.

3. Spherical wavelets

3.1. Multiresolution analysis

Following Ref. [9], we take the space $L_2 = L_2(S^2, \mathrm{d}w)$, of all square integrable functions defined on S^2, the sphere of unity radius. A multiresolution analysis is a sequence of closed subspaces $\{V_j\}\ j \geqslant 0, V_j \subset L_2$ which meets:

(1) $V_j \subset V_{j+1}$[4] The spaces are nested.
(2) $L_2(S^2, \mathrm{d}w)$ is the closure of the union of all V_j ($\bigcup_{j\geqslant 0} V_j$ is dense in $L_2(S^2, \mathrm{d}w)$).
(3) For every resolution level j, there are scaling functions $\varphi_{j,k}$ which are a Riesz basis of V_j,[5] being $k \in K(j)$ an index set where $K(j) \subset K(j+1)$.

From the first property, we have

$$\varphi_{j,k} = \sum_l h_{j,k,l}\varphi_{j+1,l}, \tag{1}$$

being $l \in K(j+1)$, $k \in K(j)$. If we have the spaces W_j so that $V_j \oplus W_j = V_{j+1}$ and we have a set of functions $\{\psi_{j,m} \,|\, j \geqslant 0, m \in M(j), M(j) \subset K(j+1)\}$ so that

(1) $\{\psi_{j,m} \,|\, j \geqslant 0,\ m \in M(j)\}$ is a Riesz basis for $L_2(S^2)$,
(2) $\{\psi_{j,m} \,|\, m \in M(j)\}$ is a Riesz basis of W_j,

then the $\psi_{j,m}$ define a spherical wavelet basis. Given that $W_j \subset V_{j+1}$, we also have

$$\psi_{j,m} = \sum_l g_{j,m,l}\varphi_{j+1,l} \tag{2}$$

for $m \in M(j)$ and $l \in K(j+1)$. In a biorthogonal setting, we have the dual spaces $\tilde{V}_j$ and $\tilde{W}_j$, spanned by the basis $\tilde{\varphi}_{j,k}$ and $\tilde{\psi}_{j,m}$, which are orthogonal to $\varphi_{j,k}$ and $\psi_{j,m}$, respectively. A function $f \in L_2$, can therefore be represented as

$$f = \sum_{j,m} \langle \tilde{\psi}_{j,m}, f \rangle \psi_{j,m} = \sum_{j,m} \gamma_{j,m}\psi_{j,m}. \tag{3}$$

It is also possible to express $\varphi_{j,l}$, the scaling functions at resolution level j, as a combination of scaling functions and wavelets at a coarser resolution level $j-1$:

$$\varphi_{j,l} = \sum_k \tilde{h}_{j-1,k,l}\varphi_{j-1,k} + \sum_m \tilde{g}_{j-1,m,l}\psi_{j-1,m}, \tag{4}$$

where $k \in K(j-1)$, $l \in K(j)$ and $m \in M(j-1)$.

Let $\lambda_{j,k}$ be the scaling coefficients of a function f, computed at resolution level j. The fast wavelet transform computes recursively $\{\lambda_{j-1,k} \,|\, k \in K(j-1)\}$, the coarse approximations, and $\{\gamma_{j-1,m} \,|\, m \in M(j-1)\}$, the detail coefficients, both at the following $j-1$ lower-resolution level:

$$\lambda_{j-1,k} = \sum_l \tilde{h}_{j-1,k,l}\lambda_{j,l} \tag{5}$$

and

$$\gamma_{j-1,m} = \sum_l \tilde{g}_{j-1,m,l}\lambda_{j,l}. \tag{6}$$

The inverse wavelet transform reconstructs the coefficients of the following higher-resolution level $j+1$ from the j lower resolution level's coefficients:

$$\lambda_{j+1,l} = \sum_k h_{j,k,l}\lambda_{j,k} + \sum_m g_{j,m,l}\gamma_{j,m}. \tag{7}$$

3.2. Extension to a triangular facet cover

Without loss of generality, we suppose that the spheric surface is partitioned by a tiling of triangular facets. This is the case when we approximate a sphere S^2 by the regular subdivision of an icosahedron with tiles $T_{j,k} \subset S^2$, $k \in K(j)$ (see Fig. 1(a)) that satisfy:

(1) $S^2 = \bigcup_{k\in K(j)} T_{j,k}$ and their union is disjoint.
(2) For each j and k, $T_{j,k}$ is the union of 4 *child* triangles $T_{j+1,l}$ (see Fig. 2).

[4] Higher values of the index j mean higher resolution.

[5] Every element f in V_j can be written uniquely as $f = \sum_k c_k f_k$, existing positive constants C_1 and C_2 so that

$$C_1\|f\|^2 \leqslant \sum_k |c_k|^2 \leqslant C_2\|f\|^2.$$

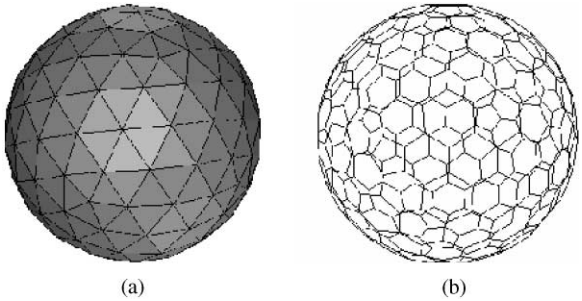

Fig. 1. Basic model structures. (a) Sphere discretization by subdividing an icosahedron j times. (b) Geodesic dome, dual of the icosahedron.

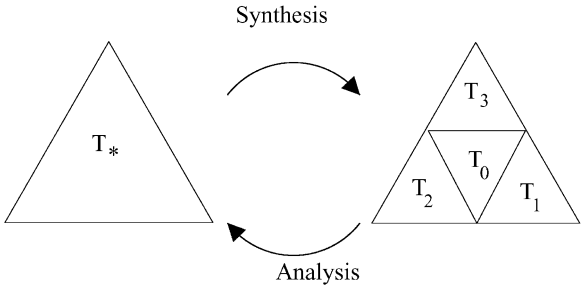

Fig. 2. *Parent* ⇄ *child* relationships between consecutive level facets. In the synthesis step, the source is triangle T_* which is divided in order to generate four new triangles. The reverse step, analysis, merges four facets into one.

Let $\alpha(T_{j,k})$ be the spherical area of a triangular patch on S^2 at resolution level j, and define the functions

$$\varphi_{j,k} = \chi T_{j,k} \quad \text{and} \quad \tilde{\varphi}_{j,k} = \alpha(T_{j,k})^{-1}\chi T_{j,k}, \tag{8}$$

where $\chi T_{j,k}$ is equal to 1 inside the patch $T_{j,k}$, and zero outside $T_{j,k}$. Then, the spaces $V_j = span\{\varphi_{j,k} \mid k \in K(j)\}$ define a multiresolution analysis of $L_2(S^2, \mathrm{d}w)$, with $\varphi_{j,k}$ and $\tilde{\varphi}_{j,k}$ as scaling and dual scaling functions. From a patch $T_{j,*}$, and its children $T_{j+1,l=0,1,2,3}$, obtained through subdivision as in Fig. 2, we can construct the bio-Haar [39] wavelets as

$$\psi_{j,m} = 2(\varphi_{j+1,m} - (\alpha(T_{j+1,m})/\alpha(T_{j+1,0}))\,\varphi_{j+1,0}), \tag{9}$$

in order to have their integral vanishing. The set of dual functions is defined as

$$\tilde{\psi}_{j,m} = 1/2(\tilde{\varphi}_{j+1,m} - \tilde{\varphi}_{j,*}). \tag{10}$$

Schröder and Sweldens present in Ref. [9] a number of vertex-based wavelet bases which are more efficient for performing certain geometric operations. In consequence, vertex basis seem better fit for computer graphics environments than facet basis such as the bio-Haar described above. In this paper, facet rather than vertex basis have been used, since they are better adapted to computer vision environments: in addition to geometrical data, each node should be able to gather other sources of information such as motion, texture, color, etc. which could be computed from the area associated to each facet. If vertex basis are used, the vertex' support regions at different resolution levels are not nested. On the other hand, the price to be paid while using facet basis is a slight increase in complexity for performing operations based on the object's geometry.

4. Overcoming irregularly sampled data

Generality is the first aspect that has been taken into account in this paper. The main objective was to develop a fully automated procedure for achieving 3D multiresolution representations that could be used in any computer vision environment, and which could take advantage of any 3D data measuring technique. Previous methods for 3D modeling in computer graphics environments departed from hypothesis that hold only for objects which have either been synthesized or acquired with dense samplings. In these cases, therefore, it is reasonable to assume sampling regularity, or at least, it is easier to compute additional structural information such as edges and facets lists.

The technique described here, presented first in Ref. [42], does not need any additional assumption. In consequence, it is able to compute and efficiently handle multiresolution superficial representations of objects acquired as irregularly sampled 3D clouds of points without any operator assistance.

In this method, we must first define a representation domain in which we can map the original input data — the 3D coordinates of each acquired point — and where we can have the possibility of handling the representation at different resolution levels. We achieve this by creating an icosahedron and recursively subdividing each of its facets in order to get a regular polyhedron with 20×4^n triangular faces. This polyhedron can then be approximated to a sphere of unity radius (see Fig. 1(a)). Using the multiresolution analysis described in the previous section, we can propagate the geometry of the object at different resolution levels. Higher values of n correspond to finer levels of detail on the object. On the other hand, lower values of n represent a coarser description of the object.

The dual of the icosahedron is a geodesic dome, a wireframe mesh of hexagonal cells (see Fig. 1(b)), where each node of the mesh has valence[6] three (like the facets

[6] Number of neighbors of a geometric primitive, i.e., nodes or facets.

in the icosahedron). An advantage of the dual is the lower dimensionality of the primitives that compose the mesh (3D points and segments), which allows higher degrees of freedom in the process of adjusting it over the scattered data.

Once the representation domain has been defined, it is necessary to adapt this geodesic dome to the input data's shape. The deformation process can be broken down into the following steps, following the solution proposed by Ikeuchi et al. [43]:

(1) Determine for each node (*face*) of the geodesic dome (*tessellated sphere*), the closest 3D point in the data file corresponding to the object being modeled.
(2) Perform an iterative deformation of the wireframe mesh while the average sum of local errors exceeds a fixed threshold. Every local error is defined by the distance between a node of the geodesic dome and its corresponding closest point, as determined in the previous step. Each node is deformed to match the object, according to an approximation force F_o and a curvature force F_g controlling the local regularity curvature, both defined for the actual position P_t of each node at time t. The new position of the node at time P_{t+1} is given by: $P_{t+1} = P_t + F_o + F_g + d(P_t - P_{t-1})$, where d represents a damping coefficient affecting the rate of convergence.

In the work of Ikeuchi et al., the first stage of the defor mation process is performed manually, establishing a connection between the main features of the object's shape and the nodes of the lower resolution mesh. Successive mesh subdivisions only take into account an Euclidean metric.

In this paper, the process has been fully automated, proceeding in the following steps:

(1) Determine an initial number of mesh nodes depending on the number of points acquired from the objects' surface. Experimental tests have proven that an initial number around 1/4 to 1/16 of the input object points yield good approximations.
(2) The relationship between nodes and input data points is established considering not only Euclidean distance, as suggested by Ikeuchi et al., but including also a term that takes into consideration the deformation direction. This directional distance can be expressed as follows:

$$\begin{aligned} &\text{If} \quad |\mathbf{l}|^2 < |\mathbf{N}_1|^2 \\ &\quad \text{then } d_d = \frac{\mathbf{n}\cdot\mathbf{l}}{|\mathbf{n}|} + 0'2\cdot|\mathbf{l}|, \\ &\quad \text{else } d_d = \text{MAXFLOAT}, \end{aligned} \tag{11}$$

where d_d is the directional distance between the input data point P_1 and the mesh node N_1, d is the Euclidean distance, $\mathbf{N}_1$ is the vector passing through the center of mass of the cloud of points and the node N_1, $\mathbf{l}$ is the vector linking the point P_1 and the node N_1, and $\mathbf{n}$ is the normal to the plane defined by the node's neighbours V_j, $j = 1, 2, 3$ (see Fig. 3).

With these modifications, it is possible to compute the shapes shown in this work via a fully automated process. Using the directional distance produces more accurate approximations, as shown in Fig. 4.

When the deformation process is finished, we map the positions of each node of the deformed geodesic dome to the corresponding facet of the icosahedron, which may also gather other information in addition to geometry. At that moment, we can apply the analysis process to the icosahedron in order to reduce the representation's level of detail. The geometry of the object is recovered at resolution level j if we compute the dual of the icosahedron, which corresponds to a wireframe mesh topologically equivalent to the geodesic dome.

The positions of the wireframe mesh's nodes at level j have been computed from the points adjusted to the input data through the recursive application of equations 12 or 13, which compute new values after one analysis step, moving from level $j+1$ to level j:

$$\lambda_{j,*} = \sum_{l=0}^{3} \frac{\lambda_{j+1,l}\cdot\alpha(T_{j+1,l})}{\alpha(T_{j,*})}, \tag{12}$$

$$\lambda_{j,*} = \left(\sum_{l=0}^{3} \lambda_{j+1,l}\right)/4, \tag{13}$$

$$\gamma_{j,l} = \lambda_{j+1,l} - \lambda_{j,*} \quad \text{for } l = 1, 2, 3. \tag{14}$$

Eq. (13), which avoids computing the weights associated with irregularly sampled distributions, can be used

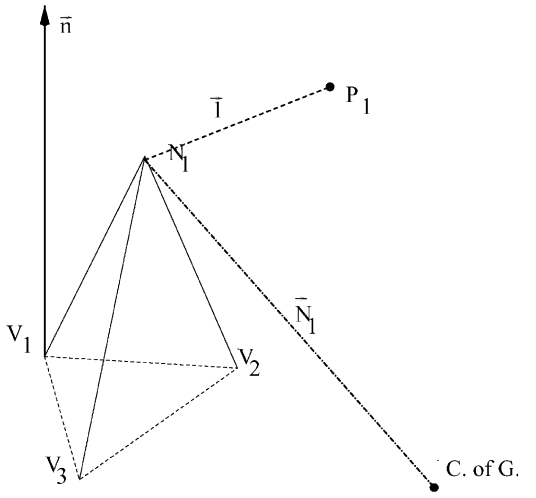

Fig. 3. Representation of the deformation orientation process in order to compute the objective function d_d.

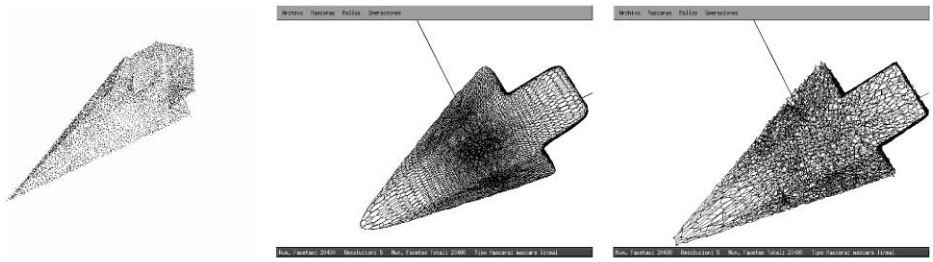

(a) (b) (c)

Fig. 4. Influence of the directional distance d_d in the deformation process. (a) Clouds of points from the tip of an arrow. (b) Mesh computed using the Euclidean distance. (c) Mesh computed using the directional distance d_d.

whenever geometric accuracy is not a requirement, or when there is an appreciable amount of sampling regularity.

In the synthesis process, we must recover the original values going from coarser level j to finer level $j+1$:

$$\lambda_{j+1,l} = \gamma_{j,l} + \lambda_{j,*} \quad \text{for } l = 1,2,3, \tag{15}$$

$$\lambda_{j+1,0} = \frac{\lambda_{j,*} \cdot \alpha(T_{j,*}) - \sum_{l=0}^{3} \lambda_{j+1,l} \cdot \alpha(T_{j+1,l})}{\alpha(T_{j+1,0})}, \tag{16}$$

$$\lambda_{j+1,0} = 4\lambda_{j,*} - \sum_{l=1}^{3} \lambda_{j+1,l}. \tag{17}$$

Eq. (17) must be used instead of Eq. (16) during synthesis whenever Eq. (13) is used during the analysis stages.

The same multiresolution analysis and synthesis processes can be applied to any other kind of information associated to each of the input sample points, such as color, surface texture, etc.

The method described above can generate multiresolution representations of objects acquired with any 3D sensing technique. The method's only restriction concerns the object's shape: only objects with a *genus* equivalent to that of a sphere can be modeled exactly. Higher genus objects can also be approximated by the representation technique described here, but it has to be noted that holes will disappear from the representation whenever the resolution falls below certain threshold. In any case, this kind of behavior seems reasonable in a multiresolution representation.

5. Multiresolution representation examples

This section presents five examples acquired with different sensors that can be used to assess the possibilities offered by multiresolution representations. In some of them (Figs. 5(a) and 7(a)), the input data is a cloud of 3D points acquired from the object's surface with the help of a hand-held 3D digitizer. This device is a tactile low-price sensor manipulated by a human operator, which allows the simulation of any other kind of sensor that provides scarce and irregularly sampled measurements. The acquired data in all of the examples provides exclusively geometric information (only the points' coordinates are kept).

Fig. 5 shows different resolution versions of a closed surface approximation of a human skull. The data sets displayed in Figs. 5(b)–(h) are a set of triangular facet models, derived from the dual polyhedron of the geodesic dome (topologically equivalent to the recursively divided icosahedron), by computing the geometric center of each hexagonal cell of the wireframe mesh at level j. The reduced resolution figures show that the overall object shape is well-kept up to resolution level 1 (5th superficial approximation), although the number of object points has been greatly reduced.

Fig. 6 shows different resolution versions of a vase, acquired with an inaccurate sensor. The data sets displayed in Figs. 6(b)–(f) are a set of wireframe models, using the adjusted 3D points as nodes. Figs. 6(b) and (c) show a rough surface, due to acquisition errors. Those errors fade away with other details when the resolution level is decreased; this smoothing effect should not be confused with that obtained through low-pass filtering operations such as those performed in Figs. 10 and 11, where the original resolution level is kept.

Fig. 7 presents again closed surface representations computed from the dual polyhedra of the geodesic dome. In this case the input data has been acquired from a composition of geometric figures with a hand-held 3D digitizer. The input cloud has been preprocessed in order to considerably reduce the number of acquired points in the object's right-half. Figs. 7(b)–(f) show that the method copes well with this strong sampling irregularity, producing only a less-accurate approximation in the object's sub-sampled right-half.

Fig. 8 shows different views of a human vertebra. The input cloud of points has been computed from volumetric data obtained through 3D tomography, which allows the selection of sample points at the desired spatial locations. The result is a dense cloud of uniformly distributed

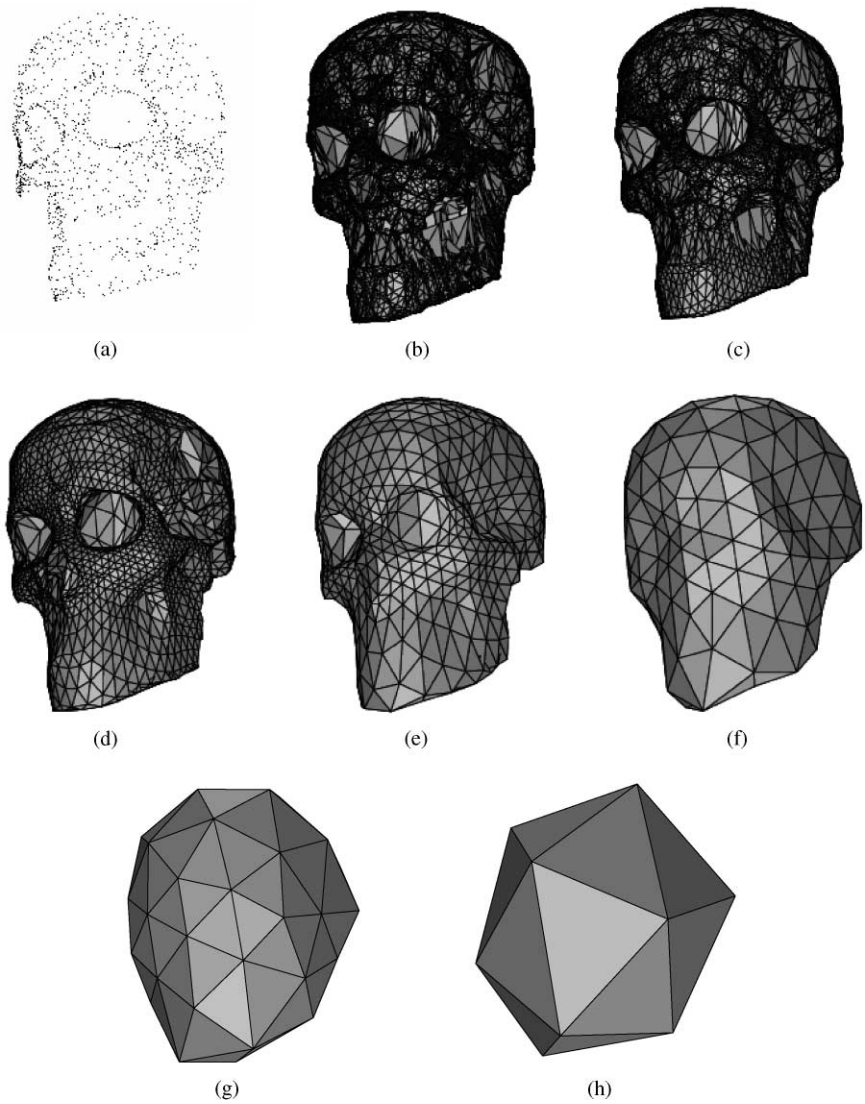

Fig. 5. Representation of a human skull. (a) Original cloud of points from a human skull. (b) Closed surface obtained from the adjusted mesh over the input data (20×4^6 facets). (c) 1st Superficial approximation (20×4^5 facets). (d) 2nd Superficial approximation (20×4^4 facets). (e) 3rd Superficial approximation (20×4^3 facets). (f) 4th Superficial approximation (20×4^2 facets). (g) 5th Superficial approximation (20×4^1 facets). (h) 6th Superficial approximation (20×4^0 facets).

points. Figs. 8(a) and (b) show a rendered view and the computed input cloud. Figs. 8(c)–(e) present different resolution wireframe approximations.

Last, Fig. 9 shows different resolutions wireframe and surface approximations of an amphora. In this case, a laser-based sensor coupled to a rotating table have been used for the 3D digitization, resulting in a more dense and regular sampling.

From a computational point of view, the most complex operation is the computation of the intermediate representation (the wireframe model), since selecting the closest input point to each mesh node has a complexity $O(m \log n)$. Computing the direct and inverse wavelet transforms is quite fast ($O(n)$), although visualizing objects which have resolution levels beyond 5 is slow without specialized graphics hardware. Most images presented in this paper have been obtained with the Geomview [44] and Xwave [45] programs; Geomview has been also used to remove occluded edges from the wireframe model (note that Geomview displays hexagonal cells with uniform grey levels, even though the cells are not planar).

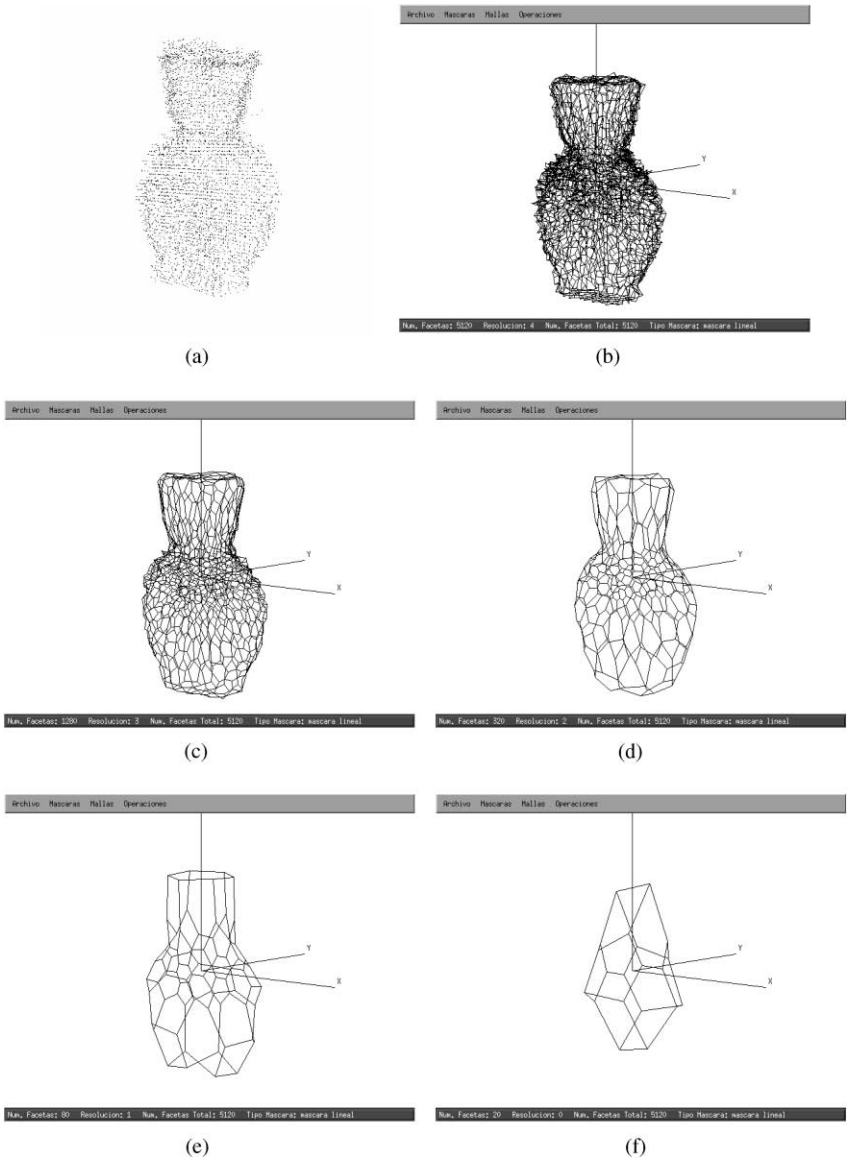

Fig. 6. Multiresolution representation of a vase. (a) Input cloud of points. (b) Adjusted mesh over the input date (20×4^4 nodes). (c) 1st Wireframe approximation (20×4^3 analysis coefficients). (d) 2nd Wireframe approximation (20×4^2 analysis coeff.). (e) 3rd Wireframe approximation (20×4^1 coeff.). (f) 4th Wireframe approximation (20×4^0 coeff.).

6. Processing multiresolution representations

3D computer vision, the process of analyzing 2D images or 3D data to make inferences about the nature of the perceived scene, is a challenging task. On one hand, sensors provide limited information about the object's features (at least, because of self-occlusions). On the other hand, real objects show a very high variability; even alterations in the environmental conditions may create strong changes in the way the scene is perceived.

Multiresolution representations which keep both the coarse approximation and detail coefficients at each intermediate level, such as those computed here, are redundant: every approximation can be computed from any of the higher-resolution coarse approximations (direct transform), or from the lower-resolution's detail and approximation coefficients (inverse transform).

These data redundancy is very useful, because it allows the resolution level which is most appropriate for fulfilling a particular task to be selected [46,47]. Also, different

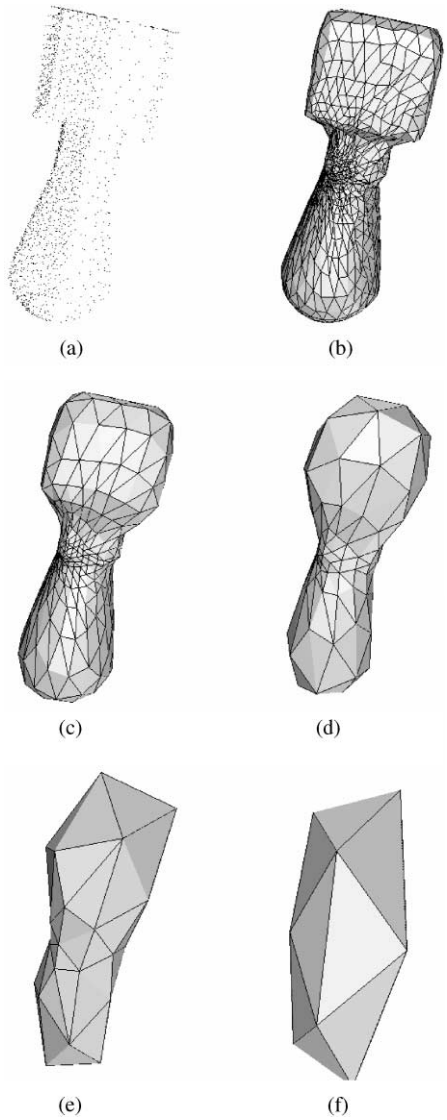

Fig. 7. Multiresolution representation of a composition of geometric blocks. (a) Input cloud of points. (b) Closed surface obtained from the adjusted mesh over the input data (20×4^4 facets). (c) 1st Surface approximation (20×4^3 facets). (d) 2nd Surface approximation (20×4^2 facets). (e) 3rd Surface approximation (20×4^1 facets). (f) 4th Surface approximation (20×4^0 facets).

levels can be processed in parallel, and their output can be combined for increased robustness [7,48]. Additionally, detail coefficients from different resolution levels can give shape information at different scales [49]. From a computational viewpoint, multiresolution representations trade memory for complexity, since processing first low-resolution versions of the perceived data first can produce large computational savings [50,51].

As stated before, the main objective of this paper is to describe a technique for computing 3D multiresolution representations without any operator intervention in the belief that these kinds of representations can help simplify the analysis process of 3D objects. For instance, this section presents some examples of operations performed effectively and efficiently using the multiresolution representations described in the previous sections. In particular, the following processes are shown:

- Preprocessing: acquisition noise removal.
- Data reduction: simplifying the input 3D point cloud without loosing significant detail.
- Object shape analysis: extracting the acquired object's global shape, to be fed to a shape-tagging process.

6.1. Acquisition noise removal

In 3D environments, the sensors' miscalibrations or lack of accuracy result in noisy depth estimates. Figs. 10(a) and 12(a) show two examples of objects' surfaces presenting spikes due to acquisition errors. These errors fade away gradually as the model resolution is decreased, as shown in Fig. 6.

Wavelet transforms' detail coefficients can be seen as the differences between "predicted" and "actual" values, (the predictions are made from lower resolution coarse approximations) [37,52], or as the output of different bandpass filters. Detail coefficients can be used for filtering, since details at specific resolution levels might have low signal-to-noise ratios. In the representations used in this paper, detail coefficients represent local shape variations, and are therefore strongly affected by high-frequency digitization noise such as that present in Figs. 10(a) and 12(a).

Figs. 10(b)–11(d) and Fig. 12(b) present different results achieved by filtering the objects in Figs. 10(a) and 12(a), respectively. The vase's case is very interesting, because it shows the effects of filtering at different resolution levels.

Figs. 10(a) and (b) show the two highest resolution approximations to the original data. Figs. 10(c) and (d) show the same two highest-resolution approximations after a filtering stage performed using the highest-resolution detail coefficients. This filter has removed the high-resolution digitization noise quite effectively.

Figs. 10(b)–(d) exhibit some lower-resolution noise which has not been removed by this first filter. Figs. 11(a) and (b) show the two highest-resolution approximations again, but this time after performing a filtering stage using only the second highest resolution coefficients. The lower frequency noise appearing in Figs. 10(b)–(d) has been removed, although the high-resolution noise is present again.

Last, Figs. 11(c) and (d) show the results of filtering using the two highest-resolution detail coefficients. The vase's surface now presents a smooth surface.

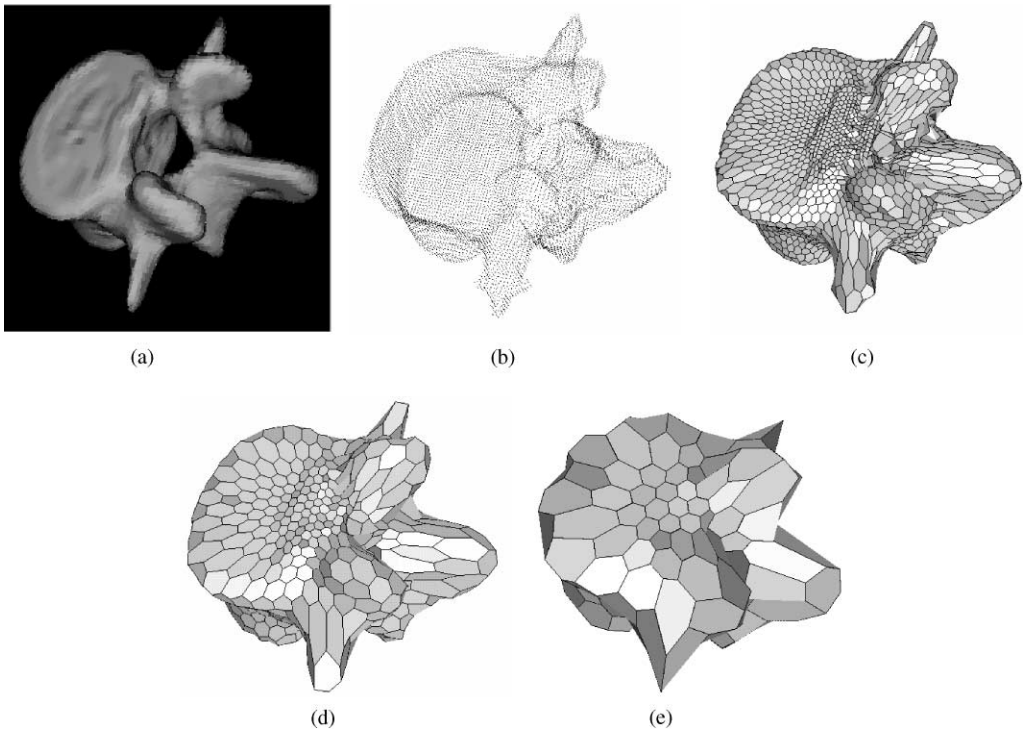

Fig. 8. Multiresolution representation of a human vertebra (8(a)). (a) Rendered view of a human vertebra obtained through 3D to morphology. (b) Input cloud of points. (c) Adjusted mesh over the input data (20×4^4 nodes). (d) 1st Wireframe approximation (20×4^3 analysis coefficients). (e) 2nd Wireframe approximation (20×4^2 analysis coeff.).

Fig. 12(b) shows the results of filtering Fig. 12(a) using the highest-resolution level detail coefficients.

6.2. Data reduction

3D meshes, such as those displayed in this paper, show a large redundancy degree, just like 2D images do. For example, if we want to represent an object preserving some areas of fine detail, a high-resolution scanning is mandatory. However, other areas of the object might not need such a high-resolution representation. The problems created by this redundancy are possibly more acute in the 3D case, since the data sets are more complex and the processing algorithms more computationally demanding.

It is evident that large reductions in mesh complexity can be achieved by eliminating nodes in smooth or flat areas, where the nodes' positions can be accurately predicted from neighboring points. Meshes can be simplified this way without loosing significant geometrical information.

Wavelet coefficients can be used for assessing potential candidates to be eliminated in the mesh-simplification process. For example, Fig. 13 presents the results of performing one and two steps in a procedure for fusing nodes at resolution level 4. Figs. 13(a) and (b) show levels 4 and 3 from the vertebra's original multiresolution representation, respectively. Data at level 4 (Fig. 13(a)) is analyzed for a selective fusion procedure. Level 3 (Fig. 13(b)) is an upper limit for the fusion, since each lower resolution approximation is achieved by fusing *every* node at the immediately higher resolution level.

The selective fusion proceeds by analyzing a surface rugosity measure based on detail coefficients: only those nodes with a rugosity below the specified threshold will be fused. The results achieved after one fusing step are presented in Fig. 13(c); it can be seen that nodes in flat areas follow a pattern identical to that at resolution level 3, while nodes in higher curvature areas such as the vertebra's protuberances or central plateau's boundary follow a pattern identical to resolution level 4. Fig. 13(d) shows further node elimination after a second fusing step.

6.3. Object shape analysis

Shape variations due to sensor noise or intra-class pattern variability produce changes in object appearance which hamper shape analysis or recognition processes.

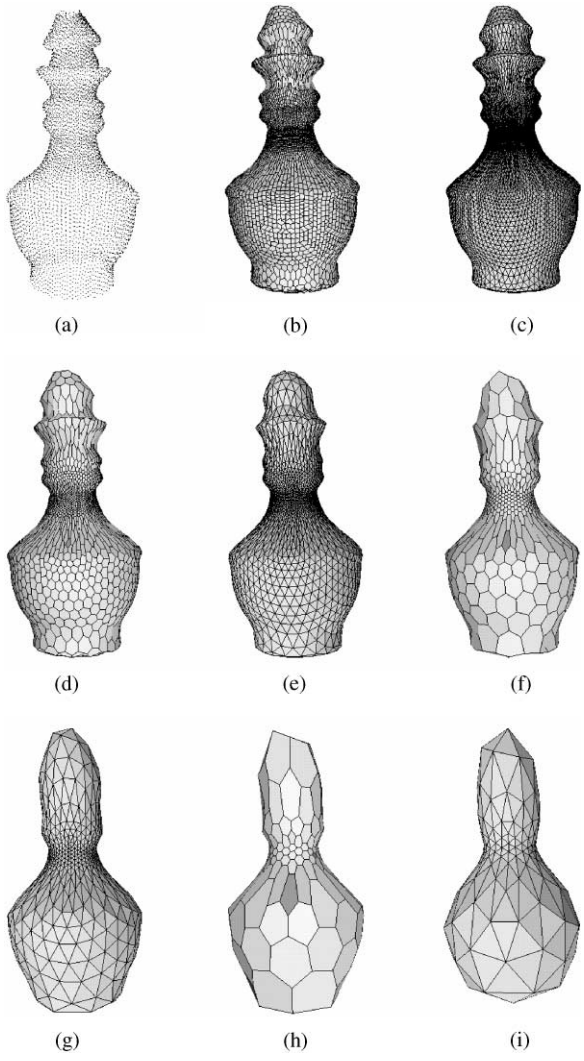

Fig. 9. Multiresolution representation of an amphora. (a) Input cloud of points (30.731 points). (b) Adjusted mesh over the input data (20×4^5nodes). (c) Closed surface obtained from the adjusted mesh over the input data (20×4^5 facets). (d) 1st Wireframe approximation (20×4^4 analysis coefficients). (e) 1st Surface approximation (5.120 facets). (f) 2nd Wireframe approximation (20×4^3 analysis coefficients). (g) 2nd Surface approximation (1.280 facets). (h) 3rd Wireframe approximation (20×4^2 analysis coefficients). (i) 3rd Surface approximation (20×4^2 facets).

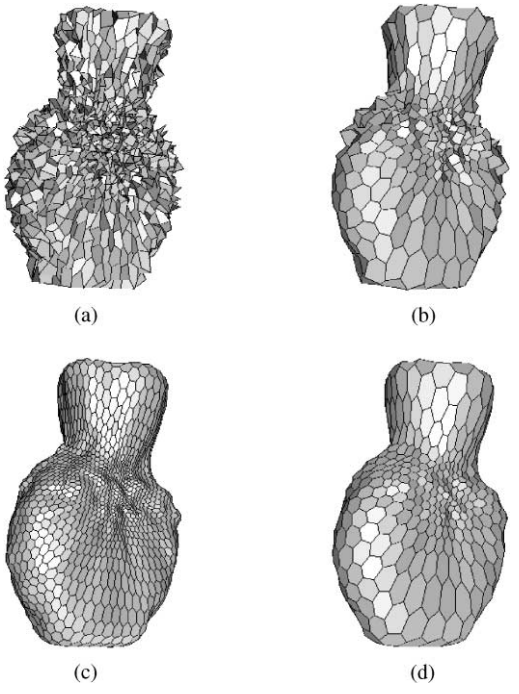

Fig. 10. Filtering results for a vase acquired with an inaccurate sensor. (a) Adjusted mesh over the input data. Initial resolution level: 4 (20×4^4 nodes). (b) 1st mesh approximation for the mesh in fiure 10(a). Resolution level: 3 (20×4^3 analysis coefficients). (c) Adjusted mesh from Fig. 10(a) filtered at resolution level 3, after one synthesis step to compute the original resolution level (4). (d) Adjusted mesh from Fig. 10(a) filtered at resolution level 3.

Coarse shape approximations extracted from multiresolution representations can be used to make the process more robust, since small or intermediate shape variations are lost whenever resolution falls below a certain threshold.

Additionally, the analysis procedures speed up, since the input data sets are significantly reduced. As an illustrative example, data points from intermediate representation levels for 3D objects, such as those presented in this paper, have been fed to a shape labelling process presented in Ref. [53]. This process computes qualitative descriptions of the global shape of 3D objects, departing from clouds of points in $\mathbb{R}^3$, such as those produced by any of the representation or processing techniques described in the previous sections. As output, the labelling process associates the input data with a volumetric label that qualitatively describes the cloud's global shape. The system is also able to decompose the input object into different volumes with a homogeneous shape. First, the system computes a hierarchy of 2D and 3D primitives that are subsequently grouped from the lower levels up in order to create a 3D shape description. The method can be decomposed in the following steps [53]:

(a) *Realignment.* Once the cloud of points has been acquired, a realignment stage is performed in order to reorient the cloud along the maximum elongation axis (MEA).
(b) *Cloud decomposition.* The cloud of points is decomposed into a continuous sequence of bins or slices

taken along the MEA. A plane which is perpendicular to the MEA is assigned to each slice, and all of the points within each bin are orthographically projected onto its corresponding bin's plane.

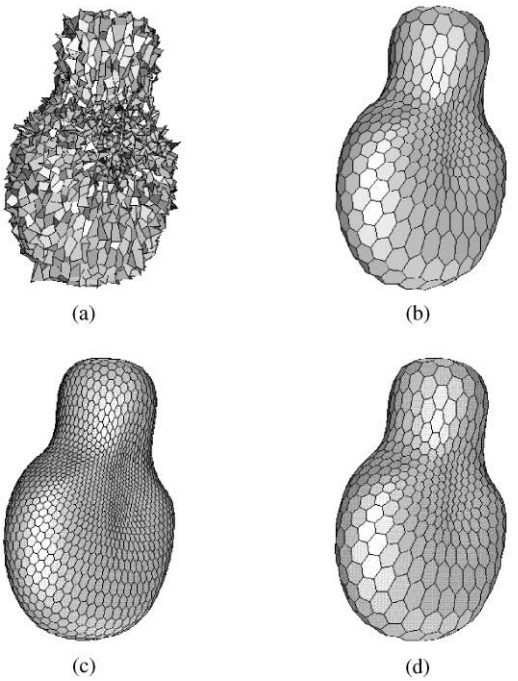

Fig. 11. Filtering results for the vase in Fig. 10(a). (a) Adjusted mesh from Fig. 10(a) filtered at resolution level 2 and reconstructed until level 4. (b) Adjusted mesh from Fig. 10(a) filtered at resolution level 2 and reconstructed until level 3. (c) Adjusted mesh form Fig. 10(a) filtered at resolution levels 2 and 3 and reconstructed until level 4. (d) Adjusted mesh from Fig. 10(a) filtered at resolution levels 2 and 3 and reconstructed until level 3.

(c) *2D Qualitative primitive extraction.* The original cloud of points is discretized over the MEA, and the 3D analysis problem can be decomposed into a series of 2D shape extraction cases. This process produces a geometrical shape tag which is assigned to each bin.

(d) *2D Shape tag association and global shape computation.* Grouping successive 2D shape tags permits the local recovery of the third dimension. In consequence, the output of the grouping process is a collection of consecutive local 3D descriptions that can be

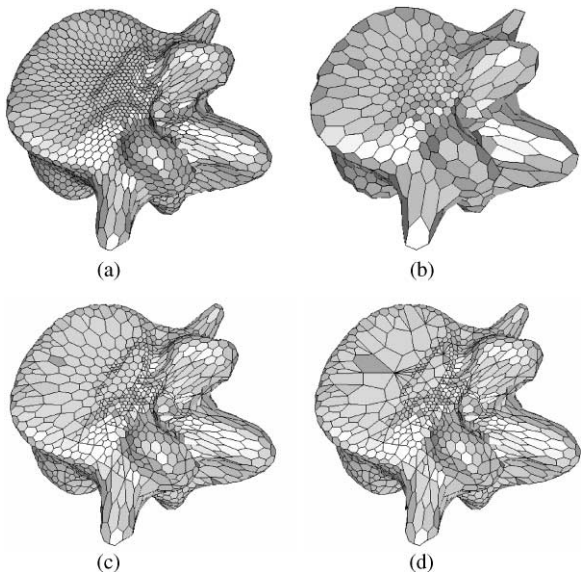

Fig. 13. Data reduction of a human vertebra starting from Fig. 13(a). (a) Adjusted mesh over the input data (20×4^4 nodes). (b) 1st mesh approximation. (c) Mesh in 13(a) after one selective fusion step. (d) Mesh in 13(a) after two selective fusion steps.

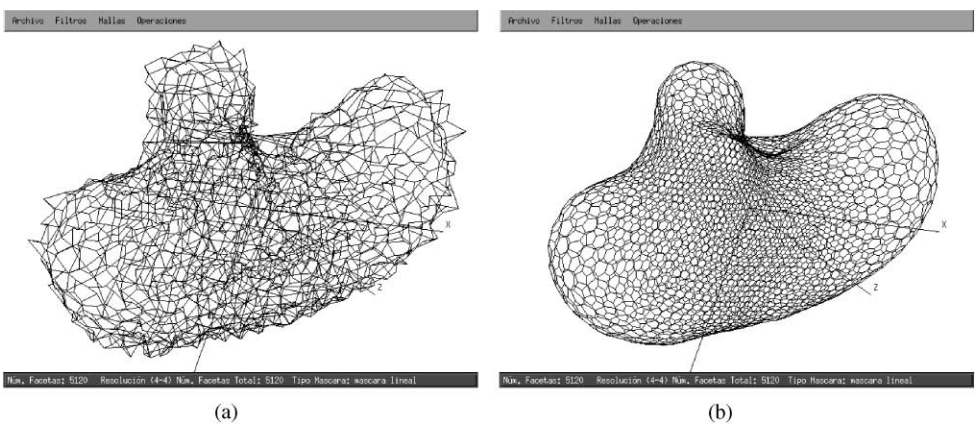

Fig. 12. Filtering results for a golf club's head acquired with an inaccurate sensor. (a) Adjusted mesh over noisy input data. Initial resolution level: 4 (20×4^4 nodes). (b) Filtering the mesh in Fig. 12(a) at resolution level 3 and reconstructing the original resolution level (4).

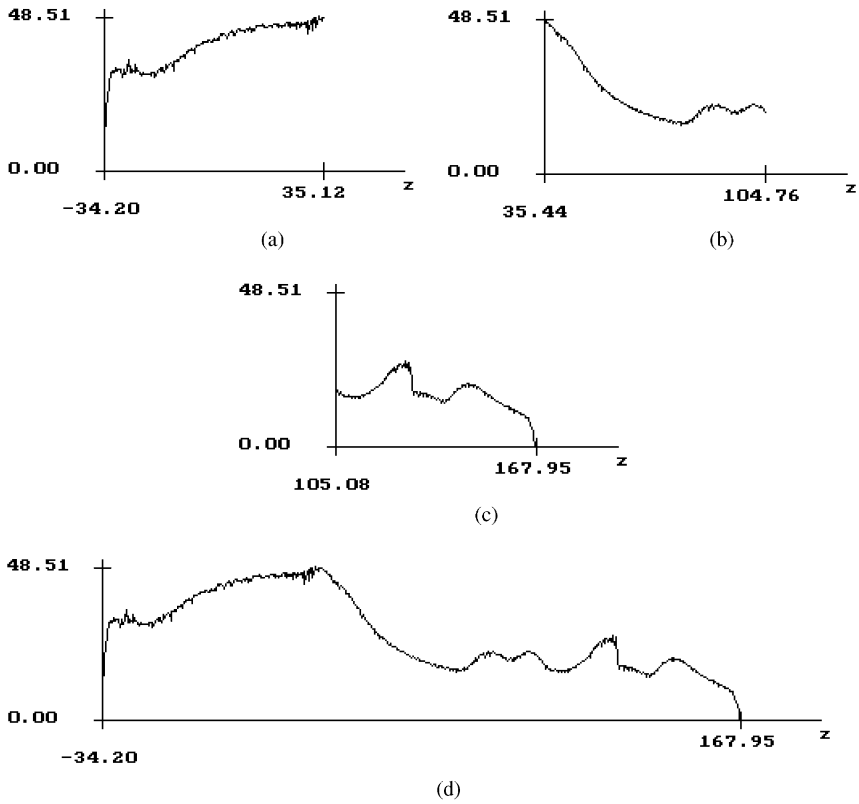

Fig. 14. Overall qualitative shape extraction process for the amphora in Fig. 9. (a) Planes: 1..216 Shape: Increasing circumference. (b) Planes: 217..432 Shape: Decreasing circumference. (c) Planes: 433..628 Shape: Convex circumference. (d) Planes: 1..103 Shape: Decreasing circumference.

recursively associated, reducing the number of local descriptions and the level of description detail along the MEA. Finally the label that best describes the object's global shape is returned.

Fig. 14 shows the results achieved after applying this labelling scheme to an input cloud of points extracted from one of the amphora's intermediate resolution levels (Fig. 9). After selecting the MEA axis, the amphora is decomposed into a set of regions, according to their shape. Finally the global tag *decreasing circumference* is selected for the overall object's shape. Note that approximately two-thirds of the amphora's profile corresponds to that label.

7. Conclusions and future work

The method described here bridges the gap between wavelet-based multiresolution representations developed for computer graphics environments and the specific problems found in computer vision environments, where there is a stronger need for fully automated procedures.

To the field of 3D multiresolution object representation, wavelets offer a sound theoretical framework. Most important, using a wavelet-based approach for computing 3D multiresolution representations allows the transference of research work performed in other application areas to computer vision environments. This is a strong advantage, because there is a very large research effort devoted nowadays to wavelets in a wide diversity of application areas. The present paper gives two examples of typical wavelet applications ported to 3D object representations: compression and filtering.

Working with multiresolution representations offers interesting possibilities. Processing wavelet and coarse approximation coefficients can facilitate certain tasks during the preprocessing or analysis stages. Some examples, such as noise removal and mesh simplification, have already been mentioned in this paper. Other examples can be found in the bibliography: FOR instance, Reissell [50] suggests using wavelet coefficients for identifying smooth surface sections to plan mobile robot paths through natural terrains. Also, within CAD or graphics environments, Stollnitz et al. [5] point out the benefits offered by multiresolution editing. Last, an advantage

that is evident from the examples given here is that multiresolution models offer a dramatic reduction in model complexity and allow the user to choose the fidelity with which the original data is modeled, letting him trade fidelity for complexity.

From the 3D object recognition point of view, multiresolution representations have the advantage of providing a set of simplified versions of the object's shape, which can be used for extracting different information to be used in the recognition process. This paper gives an example, using coarse approximations for computing the overall global shape. Further work is needed in this area.

Perhaps the method's main drawback is that, to be precisely modeled, the input object must have a genus equivalent to that of a sphere, since holes within the object's surface are lost whenever resolution decreases below a certain threshold. Somehow this behavior is reasonable, given that details such as superficial texture, protuberances or holes should disappear as the representation becomes coarser. In any case, when analyzing 3D shapes (for example, in recognition environments), holes can be dealt with as independent entities, thereby preserving the method's usefulness.

Future research lines include:

- Using other basis functions that might achieve better approximations with fewer coefficients than the bio-Haar basis used in this paper.
- Performing shape analysis and other processing stages, such as those presented in Section 6, in the wavelet transform domain.
- Dealing with objects with higher genus.

Acknowledgements

This work has been partially funded by the Spanish Commission for Science and Technology (grants CICYT TAP94-0305-C03-02 and CICYT TIC98-0272-C02-01) and Madrid's Autonomous Community (grant 06T/020/96). The authors gratefully acknowledge the help provided by Belén Moreno (3D skull data) and Jaime Gómez (amphora).

References

[1] D. Ballard, C. Brown, Computer Vision, Prentice-Hall, Englewood Cliffs, NJ, 1982.
[2] D. Marr, E. Hildreth, Theory of edge detection, Proceedings of the Royal Society, Ser. B, London, Vol. 207, 1980, pp. 187–217.
[3] A. Rosenfeld (Ed.), Multiresolution Image Processing and Analysis, Springer Series in Information Sciences, Vol. 12, Springer, Berlin, 1984.
[4] G. Strang, T. Nguyen, Wavelets and Filter Banks, Wellesley-Cambridge Press, Wellesley, MA, 1997.
[5] E.J. Stollnitz, T.D. DeRose, D.H. Salesin, Wavelets for Computer Graphics, Morgan Kauffman Publishers, Los Altos, CA, 1996.
[6] J.-L. Starck, F. Murtagh, A. Bijaoul, Image Processing and Data Analysis. The Multiscale Approach, Cambridge University Press, Cambridge, 1998.
[7] P. Wunsch, A.F. Laine, Wavelet descriptors for multiresolution recognition of handprinted characters, Pattern Recognition 28 (8) (1995) 1237–1249.
[8] Q.M. Tieng, W.W. Boles, Recognition of 2D object contours using the wavelet transform zero-crossing representation, IEEE Trans. Pattern Anal. Mach. Intell. 19 (8) (1997) 910–916.
[9] P. Schröeder, W. Sweldens, Spherical wavelets: Efficiently representing functions on a sphere. Computer Graphics Proceedings (SIGGRAPH 95), 1995, pp. 161–172.
[10] S. Campagna, H.-P. Seidel, Parameterizing meshes with arbitrary topology, in: H. Niemann, H.-P. Seidel, B. Girod, (Eds.), Image and Multidimensional Signal Processing'98, 1998, pp. 287–290.
[11] S.G. Mallat, Zero-crossing of a wavelet transform, IEEE Trans. Inform. Theory 37 (14) (1991) 1019–1033.
[12] E.P. Simoncelli, W.T. Freeman, E.H. Adelson, D.J. Heeger, Shiftable multiscale transforms, IEEE Trans Inform. Theory 38 (2) (1992) 587–607.
[13] E.L. Grimson, Introduction of the special issue on interpretation of 3D scenes, IEEE Trans. Pattern Anal. Mach. Intell. 14 (2) (1992) 97–98.
[14] S.J. Dickinson, R. Bergevin, I. Biederman, J. Eklundh, R. Munck-Fairwood, A. Pentland, The use of geons for generic 3D object recognition, Proceedings of the 13th International Conference on Artificial Intelligence, Chambery, Francia, 1993, pp. 1693–1699.
[15] J. Wood, Invariant pattern recognition: a review, Pattern Recognition 29 (1) (1996) 1–17.
[16] H. Samet, The quadtree and related hierarchical data structures, ACM Comput. Survey 16 (2) (1984) 187–260.
[17] K.N. Kutulakos, S.M. Seitz, A theory of shape by space carving, Proceedings of the Seventh IEEE International Conference on Computer Vision, IEEE, New York, September 1999, pp. 307–314.
[18] F.C.M. Martins, J.M.F. Moura, Video representation with three-dimensional entities, IEEE J. Selected Areas Commun. 16 (1) (1998) 71–85.
[19] J.D. Foley, A. van Dam, S.K. Feiner, J.F. Hughes, Computer Graphics: Principles and Practice, Addison-Wesley, Reading, MA, 1996.
[20] G. Farin, Curves and Surfaces for CAGD, Academic Press, New York, 1997.
[21] K. Ikeuchi, M. Hebert, Spherical representations: from EGI to SAI, NSF-ARPA Workshop in Object Representation in Computer Vision, 1995, pp. 327–345.
[22] S. Kumar, S. Han, D. Goldgof, K. Bowyer, On recovering hyperquadrics from range data, IEEE Trans. Pattern Anal. Mach. Intell. 17 (11) (1995) 1079–1083.
[23] D. Terzopoulos, D. Metaxas, Dynamic 3D models with local and global deformations: Deformable superquadrics, IEEE Trans. Pattern Anal. Mach. Intell. 13 (7) (1991) 703–714.

[24] K.F. Lai, R.T. Chin, Deformable contours: modeling and extraction, IEEE Trans. Pattern Anal. Mach. Intell. 17 (11) (1995) 1084–1090.
[25] M. Levine, Vision in Man and Machine, McGraw-Hill, New York, 1985.
[26] J.P. Frisby, Seeing. Illusion, Brain and Mind, Roxby Press Ltd, London, 1979.
[27] K. Wu, M.D. Levine, 3-D shape approximation using parametric geons, Image Vision Comput. 15 (2) (1997) 143–158.
[28] S.J. Dickinson, A.P. Pentland, A. Rosenfeld, 3-D shape recovery using distributed aspect matching, IEEE Trans. Pattern Anal. Mach. Intell. 14 (2) (1992) 174–198.
[29] C. Dorai, A.K. Jain, Cosmos: a representation scheme for 3D free-form objects, IEEE Trans. Pattern Anal. Mach. Intell. 19 (10) (1997) 1115–1130.
[30] S.C. Hwang, H.S. Yang, A high-quality 3D multiview model-building technique based on CAD, Mach. Vision Appl. 7 (1994) 209–216.
[31] S. Ullman, High-level Vision, Object Recognition and Visual Cognition, MIT Press, Cambridge, MA, 1996.
[32] D. Forsyth, J.L. Mundy, A. Zisserman, C. Coelho, A. Heller, C. Rothwell, Invariant descriptors for 3D object recognition and pose, IEEE Trans. Pattern Anal. Mach. Intell. 13 (10) (1991) 971–991.
[33] F.S. Cohen, J. Wang, Part I: modeling image curves using invariant 3D object curve models. a path to 3D recognition and shape estimation from image contours, IEEE Trans. Pattern Anal. Mach. Intell. 16 (1) (1994) 1–12.
[34] M. Zerroug, R. Nevatia, Three-dimensional descriptions based on the analysis of the invariant and quasi-invariant properties of some curved-axis generalized cylinders, IEEE Trans. Pattern Anal. Mach. Intell. 18 (3) (1996) 237–253.
[35] M. Lounsbery, T. DeRose, J. Warren, Multiresolution analysis for surfaces of arbitrary topological type, ACM Trans. Graphics 16 (1) (1997) 34–73.
[36] M. Lounsbery, Multiresolution analysis for surfaces of arbitrary topological type, Ph.D. Thesis, Universidad de Washington, 1994.
[37] W. Sweldens, The lifting scheme: a construction of second generation wavelets, SIAM J. Math. Anal. 29 (2) (1997) 511–546.
[38] W. Sweldens, The lifting scheme: a custom-design construction of biorthogonal wavelets, J. Appl. Comput. Harmonic Anal. 3 (2) (1996) 186–200.
[39] P. Schröder, W. Sweldens, Spherical wavelets: texture processing, in: P. Hanrahan, W. Purgathofer (Eds.), Rendering Techniques'95, Springer, Berlin, August 1995.
[40] D. Zorin, P. Schröeder, W. Sweldens, Interactive multiresolution mesh editing, Computer Graphics, SIGGRAPH 97 Proceedings, 1997, pp. 259–269.
[41] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, W. Stuetzle, Multiresolution analysis of arbitrary meshes, Computer Graphics Proceedings (SIGGRAPH 95), 1995, pp. 173–182.
[42] L. Pastor, A. Rodríguez, Surface approximation of 3D objects from irregularly sampled clouds of 3D points using spherical wavelets, Proceedings of ICIAP, Venice, Italy, September 1999, pp. 70–75.
[43] H. Delingette, H. Hebert, K. Ikeuchi, A spherical representation for the recognition of curved objects, Proceedings of the ICCV, Berlin, 1993, pp. 103–112.
[44] The Geometry Center, Geomview. http://www.geom.umn.edu/
[45] J.M. Espadero, A. Rodríguez, Visualización multirresolución de objetos 3D utilizando wavelets, Proc. de los VIII Encuentros de Geometría Computacional, Castellón, España, July 1999, pp. 259–268.
[46] B. Jawerth, W. Sweldens, An overview of wavelet based multiresolution analysis, J. SIAM 36 (3) (1994) 377–412.
[47] A. Rosenfeld (Ed.), Multiresolution Image Processing and Analysis, chapter 1, Some useful properties of pyramids, pages 2–5. Volume 12 of Rosenfeld [3], 1984.
[48] G. Van de Wouwer, S. Livens, D. van Dyck, Color texture classification by wavelet energy-correlation signatures, Proceedings of the International Conference on Computer Analysis and Image Processing, Lectures Notes in Computer Science, Vol. 1310, Firenze, Italy, 1997, pp. 327–334.
[49] L. Pastor, A. Rodríguez, D. Ríos, Wavelets for object representation and recognition in computer vision, in: B. Vidaković, P. Müller (Eds.), Bayesian Inference in Wavelet Based Models, Lectures Notes in Statistics, Vol. 141, Springer, Berlin, 1999, pp. 267–290.
[50] L.M. Reissell, Wavelet multiresolution representation of curves and surfaces, Graphical Models Image Process. 58 (3) (1996) 198–217.
[51] L. Cinque, L. Lombardi, Shape description and recognition by a multiresolution approach, Image Vision Comput. 13 (8) (1995) 599–607.
[52] W. Sweldens, P. Schröder, Building your own wavelets at home, Wavelets in Computer Graphics, ACM SIGGRAPH Course Notes, ACM, New York, 1996, pp. 15–87.
[53] L. Pastor, J. Durán, A. Rodríguez, Qualitative shape modeling of 3D data, Proceedings of Second International Conference on Mathematics & Design 98, San Sebastián, Spain, June 1998, pp. 459–466.

About the Author—LUIS PASTOR received the BSEE degree from the Universidad Polit'eecnica de Madrid in 1981, the MSEE degree from Drexel University in 1983, and the Ph.D. degree from the Universidad Politécnica de Madrid in 1985. Currently he is Professor in the University Rey Juan Carlos. His research interests include computer vision (3D modeling and recognition) and parallel computing.

About the Author—ANGEL RODRÍGUEZ received his degree in Computer Science and Engineering and the Ph.D. degree from the Universidad Politécnica de Madrid in 1991 and 1999 respectively. His Ph.D. is centered on the tasks of modeling and recognizing 3D objects in parallel architectures. He is an Associate Professor in the Photonics Technology Department, Universidad Politécnica de Madrid, Spain and has published works in the fields of parallel computer systems and computer vision. He is an IEEE and an ACM member.

About the Author—J. MIGUEL ESPADERO received his degree in Computer Science and Engineering from the Universidad Politécnica of Madrid in 1998. Since then he has been a Ph.D. student at the Photonics Technology Department in the same university,

with the research theme "Data Modeling and Automatic Recognition using Wavelets". Nowdays he is a Teaching Assistant in the Universidad Rey Juan Carlos of Madrid. His research interest include 3D objects representation and recognition, computer vision, image processing and wavelets.

About the Author—LUIS RINCÓN received his degree in Computer Science and Engineering from the Universidad Politécnica de Madrid in 1992. From 1993 to 1997 he has been an Assistant Professor in the Electrical, Electronic and Control Engineering Department in the Universidad Nacional de Educación a Distancia (UNED). Since 1997, he is an Assistant Professor in the Experimental Sciences and Technology Department in the Universidad Rey Juan Carlos of Madrid and teaches Computer Structure and Technology and Computer Architecture. He has published works in the field of computer vision and quality control. His research interests include computer vision, pattern recognition, quality control and parallel architectures.