

Blowing Bubbles for Multi-Scale Analysis and Decomposition of Triangle Meshes¹

Michela Mortara,² Giuseppe Patanè,² Michela Spagnuolo,²
Bianca Falcidieno,² and Jarek Rossignac³

Abstract. Tools for the automatic decomposition of a surface into shape features will facilitate the editing, matching, texturing, morphing, compression and simplification of three-dimensional shapes. Different features, such as flats, limbs, tips, pits and various blending shapes that transition between them, may be characterized in terms of local curvature and other differential properties of the surface or in terms of a global skeletal organization of the volume it encloses. Unfortunately, both solutions are extremely sensitive to small perturbations in surface smoothness and to quantization effects when they operate on triangulated surfaces. Thus, we propose a multi-resolution approach, which not only estimates the curvature of a vertex over neighborhoods of variable size, but also takes into account the topology of the surface in that neighborhood. Our approach is based on blowing a spherical bubble at each vertex and studying how the intersection of that bubble with the surface evolves. We describe an efficient approach for computing these characteristics for a sampled set of bubble radii and for using them to identify features, based on easily formulated filters, that may capture the needs of a particular application.

Key Words. Shape description, Shape decomposition, Multi-scale shape feature extraction, Shape indexing.

1. Introduction. Shape analysis and coding are challenging problems in Computer Vision and Graphics. An ideal shape description should be able to capture and compute the main features of a given shape and organize them into an abstract representation which can be used to automate processes such as matching, retrieval or comparison of shapes. We have tackled the problem in the context of three-dimensional objects represented by triangular meshes, bearing in mind that a good shape description should be able to distinguish between global and local features and should be based on geometric properties of the shape which are invariant under rotation, translation and scaling [5]. To characterize a shape we have used the paradigm of *blowing bubbles*: a set of spheres of increasing radius R_i , $i = 1, \dots, n$, is drawn, whose centers are at each vertex of the mesh, and whose radius represents the scale at which the shape is analyzed. The number of connected components of the intersection curve between each bubble and the surface gives a first qualitative characterization of the shape in a three-dimensional neighborhood of each vertex. Then the evolution of the ratio of the length of these components to the radius of the spheres is used to refine the classification and detect specific features such as sharp protrusions or wells, mounts or dips, blends or branching parts.

¹ This work has been supported by the Bilateral Project between GVU/Gatech (USA) and IMATI-GE/CNR, and the National Science Foundation under CARGO Award 0138420, “Multi-scale Topological Analysis of Time-Evolving Shapes”.

² Istituto di Matematica Applicata e Tecnologie Informatiche, Consiglio Nazionale delle Ricerche, Genova, Italy. {michela.patane,spagnuolo,falcidieno}@ima.ge.cnr.it.

³ College of Computing and GVU Center, Georgia Institute of Technology, Atlanta, GA 30332, USA. jarek@cc.gatech.edu.




Fig. 1. Shape decomposition using blowing bubbles.

For example, a point on a thin limb will generate an intersection curve made of only one component, which will rapidly split into two connected components as soon as the radius of the sphere exceeds the limb size. Conversely, the tip of a protrusion will always generate only one connected component, and the ratio between the intersection curve length and the sphere radius can be used to evaluate how curved the surface is at the protrusion tip. An example of the resulting decomposition is given in Figure 1.

The description achieved provides an insight on the presence of features together with their morphological type, persistence at scale variation, amplitude and/or size. The algorithm for the decomposition is concise and simple, and the decomposition is independent on the orientation of the object in space and equally distributed in all directions. The number of radii and the interval given by the minimum and maximum radius is related to the spectrum of sizes of the features of interest. Moreover, the multi-scale approach and the chosen descriptors reduce the influence of noise on shape evaluation.

In this paper we focus on the method adopted for segmentation, while a possible application of the results can be found in [12], where a skeleton describing a shape from the point of view of its sharpest protrusions is presented.

The paper is organized as follows: in Section 2 previous work relevant for the described method is briefly reviewed. Basic concepts on differential geometry delineate, in Section 3, the theoretical background of the geometric descriptors used. The approach to shape classification is presented in Section 4.

Details of the algorithm are described in Section 5, while the mesh decomposition strategy is discussed in Section 6. Finally, Section 7 includes critical considerations and remarks.

2. Previous Work. An abstract description of a shape usually combines a set of primitives that are relevant to the specific context, and is defined in terms of their type and intrinsic shape parameters.

As suggested in [11], methods for shape description can be classified into two broad categories: those considering only the local properties of the boundary of the shape, and those measuring properties of the enclosed volume. Typically, boundary-based methods evaluate accurate and mathematically well-defined local characteristics, such as critical points or curvature. They may also identify specific loci on the surface, such as curvature extrema or ridges, but they generally lack in providing a global view of the shape. Furthermore, they typically work at a single resolution and thus do not organize features into a hierarchy of global and local details.

Conversely, interior-based methods, which assume that the surface is the boundary of a solid, generally provide descriptions which highlight better the global structure of the shape. Skeletons, such as the medial axis or the Reeb graph, belong to this class of descriptors [1], [15], [17]. The great advantage of skeletons is that they provide an abstract representation by idealized lines that retain the connectivity of the original shape, thus reducing the complexity of the representation. Usually, each arc is associated with a portion of the original shape that corresponds to a feature. For example, in two dimensions the medial axis is constructed using the paradigm of the maximal enclosed disks, whose centers define a locus of points which describes, together with the associated radius, the width variation of the shape. The medial axis induces a decomposition of the shape into protrusion-like features, while concavities of the shape are not directly identified by the medial axis of the interior. Unfortunately, the medial axis of a three-dimensional shape is no longer a one-dimensional graph, but is made of surface pieces as well. Moreover, the instability of the medial axis with respect to noise has prevented its use in many application areas. Approaches to construct and store the medial axis at different scales have also been proposed, which implicitly address the problem of noise reduction as well [15], [3], [14].

Another notable example of a topology-driven skeleton is given by the Reeb graph [17], [1]. The Reeb graph is a topological structure which codes a given shape by storing the evolution of criticalities of a mapping function defined on the boundary surface. In particular, when the height function with respect to a predefined direction is chosen, the Reeb graph describes the evolution of the contours obtained by intersecting the shape with constant planes. The decomposition induced by the Reeb graph corresponds to a segmentation of the solid into slices and the corresponding branches of the Reeb graph identify the connected components of the surface. The description obtained using a Reeb graph approach is suitable for matching purposes especially if the mapping function is chosen in order to provide invariance under affine transformations. Such orientation-independent approaches have been proposed in [9] and [12]; however, they are computationally intensive and offer little control over the scale at which the shape is analyzed. We propose here an alternative and more efficient approach that gives us more flexibility to formulate the filters for shape analysis, and captures the more representative properties in a more detailed description.

3. Theoretical background. This section provides definitions and concepts [7], [10], [13] useful for describing our approach. Let $x: D \subseteq \mathbf{R}^2 \mapsto \mathbf{R}^3$ be a C^2 -parameterization

of the surface

$$\Sigma := \{x(u, v) : (u, v) \in D\}.$$

The classification of local properties of Σ is traditionally based on the study of the *mean* and the *Gaussian* curvature, which can be respectively defined as the average and the product of the maximum and minimum principal curvatures [10].

We consider the normal n to the surface Σ at a point p , and the *normal sections* of the surface around the normal vector, that is, the set of curves originated by intersecting the surface with planes containing the normal n . For each of these planar curves the curvature is classically defined as the inverse of the curvature radius. If we call κ_1 the maximum curvature of the normal sections, and κ_2 the minimum, then the mean curvature $\bar{\kappa}$ is defined as $\bar{\kappa} := (\kappa_1 + \kappa_2)/2$ and the Gaussian curvature as $K := \kappa_1\kappa_2$. The directions along which the extrema of curvature are assumed are called the principal directions. This definition formalizes the relation between the surface shape and its position with respect to the tangent plane. For example, for elliptic-shaped surfaces, the centers of curvature of all the normal sections will lie on the same side of the tangent plane, with positive values for the minimum and maximum of curvature. For hyperbolic-shaped surfaces, the centers of curvature will move from one side of the surface to the other, with a negative minimum value and a positive maximum value assumed at opposite sides with respect to the tangent plane. Finally, for parabolic-shaped surfaces, one of the principal directions will have curvature equal to zero, that is, along that direction the normal section will be a straight line. This is the case, in general, of ruled surfaces which are also said to have no double curvature. The planar case is obvious.

The Gaussian curvature represents a measurement at any point p of Σ which is the excess per unit area of a small patch of the surface, i.e. how *curved* it is. An interesting result is due to the Gauss-Bonnet theorem, which is introduced as follows. First, given a closed curve γ on a surface Σ , let T_γ be the total turning that the unit tangent t undergoes when it is carried along γ , defined as the sum of the local turnings, i.e. *exterior angles* [10] (see Figure 2). Then the quantity $I_\gamma = 2\pi - T_\gamma$ is called the *angle excess* of the curve γ and it is related to the curvature of Σ within γ , as described by the Gauss-Bonnet formula.

GAUSS-BONNET THEOREM 1. *Let γ be a curvilinear polygon of class C^2 on a surface patch of class C^k , $k \geq 3$. Suppose γ has a positive orientation and its interior on the patch is simply connected. Then*

$$(1) \quad \int_{\gamma} \kappa_g ds + \iint_{\Omega} K dS = 2\pi - \sum_i \alpha_i = I_\gamma,$$

where κ_g is the geodesic curvature along γ , Ω is the union of γ and its interior, K is the Gaussian curvature, α_i is the exterior angles of γ and ds and dS are the curve and line elements, respectively.

Among the properties of the angle excess the following ones have a particular interest for our approach:

- I_γ is independent of the chosen starting point on γ ;

- I_γ is additive;
- for any topological disk on an arbitrary surface at p , the angle excess around the boundary is equal to the total curvature of the interior.

The definition of the curvature at each point of a triangulation is not trivial because a triangular mesh is parameterized by a piecewise continuous function whose second derivatives are, almost everywhere, null. In other words, the curvature on a triangulation is concentrated along edges and at vertices, since every other point has a neighborhood homeomorphic to a planar Euclidean domain whose Gaussian curvature is null. The methods proposed in the literature for curvature evaluation can be classified in different ways but a global comparison among them is still lacking as underlined in [4], [6], and [19]. These methods can be divided into two main groups: *continuity-based* and *property-based* algorithms.

The first ones are developed, transforming the discrete case to the continuous one by using a local fitting of the surface which enables us to apply standard definitions. For example, in [8] an approximation is derived at each vertex by applying the continuous definition to a least-square paraboloid fitting its neighboring vertices, while in [18] it is evaluated by estimating its tensor curvature. The second class of algorithms defines equivalent descriptors starting from basic properties of continuous operators but directly applied to the discrete settings. The methods proposed in [2] and [16] are based on the Laplace–Beltrami operator, the Gauss map and the Gauss–Bonnet theorem, guaranteeing the validity of differential properties such as area minimization and mean curvature flow [7].

For example, the angle excess can be used to evaluate the Gaussian curvature at mesh vertices [4], [16]. Let us consider the region $Star(p)$ on the mesh defined by the triangles incident in a vertex p (see Figure 2). The boundary of $Star(p)$ defines a closed path on the mesh, to which we may apply the Gauss–Bonnet formula (1). Since the geodesic curvature along the boundary is obviously zero (edges are straight), the total curvature at p is simply quantified by the sum of the exterior angles. To understand better the geometry of the situation, we can imagine locally cutting $Star(p)$ along any of the edges incident in p , and developing $Star(p)$ onto the plane without shrinking the surface. The sum of the exterior angles corresponds to the sum of the angles at p in $Star(p)$. This result is consistent with the intrinsic nature of the Gaussian curvature since the angle excess only depends on the angles, that is, this value does not change if the mesh is deformed preserving the distance between points. Also, computation of the




Fig. 2. $Star(p)$ is enclosed by a curve γ . The exterior angles $\{a_i\}$ and interior angles $\{\alpha_i\}$ are shown.




Fig. 3. Discrete Gaussian curvature and sensibility to local noise: red and blue vertices represent elliptic and hyperbolic points (see electronic version for color figures).

angle excess can be performed without resorting to any coordinate system, as the angles may be obtained using only the edge length and not the vertex coordinates. The discrete Gaussian curvature at a vertex p of the mesh can be therefore evaluated by

$$(2) \quad K_G = \frac{2\pi - \sum_{i=1}^{\text{num. faces}} \alpha_i}{A},$$

that is, the local angle excess in p weighted by the area A of a small patch of surface around p given by $Star(p)$ or by some subregion of it containing p (see Figure 2).

The approaches mentioned are usually sensitive to noise and, even if a multi-resolution setting is chosen, they usually require smoothness conditions on the input mesh. Furthermore, the smoothing process used to get stable and uniform curvature estimations introduces a deficiency in magnitude evaluation and, consequently, difficulties in accurate distinction between planar patches and curved surfaces with low curvature. In Figure 3 evaluation of the curvature, using the area of $Star(p)$ as weights, is shown. It can be seen how the results are affected by small local undulations.

4. Geometric and Topological Classification. The approach proposed here for describing a three-dimensional shape integrates boundary and interior information of the shape finalized at defining a complete multi-scale vertex classification. The link between closed paths and curvature has suggested specializing its study to the family of closed paths built by intersecting the surface with spheres centered in each of its points. Study of the evolution of these curves and the geometric characterization of the mesh areas intersected by the spheres are the core of the proposed method. The topology of the intersection curves changes according to the object shape: in Figure 4(a) the highlighted sphere intersects the surface only at one curve, while in Figure 4(b) the boundary of the intersection area splits into two connected components. This is likely to happen, for example, near handles and branches, or around deep pits. Therefore, the variation in the boundary suggests that the vertex is located on a feature, which becomes relevant at the scale, or radius, at which the change occurs.

Given a set of radii $R_i, i = 1, \dots, n$, each vertex of the mesh is classified with an n -dimensional vector of morphological labels, each corresponding to its type at the related scale. Shape features of the mesh are then identified by connected regions of




Fig. 4. The evolution of intersection for increasing radii.

vertices with the same label at a given scale, and the geometric parameters computed to assign the label will characterize the feature. For example, a tip and a mount are both characterized by one intersection curve, but they can be distinguished by measuring the curvature induced by the intersection curve on the surface (see Figure 5).

Features which are identified by two intersection curves are further characterized by measuring the relative curve length and by checking if they define a volume which is inside or outside the shape (see Figure 6). The vector of labels and corresponding geometric parameters, together with the persistence of type through the scale values, are used to distinguish between global and local features with respect to the scale range.

Given a point p , a scale R_i and the corresponding intersection curve $\gamma(p, R_i)$, the point p is classified according to the number of connected components of $\gamma(p, R_i)$, to the curvature value if $\gamma(p, R_i)$ has only one component, or to the relative length if $\gamma(p, R_i)$ has two components, and according to a concavity/convexity check in all cases.

The classification criteria lead to a complete characterization of vertices, which expresses both geometric and morphological properties of the surface. As far as this paper




Fig. 5. A tip detected at a small radius (a) will be characterized as a mount at a larger radius (b).

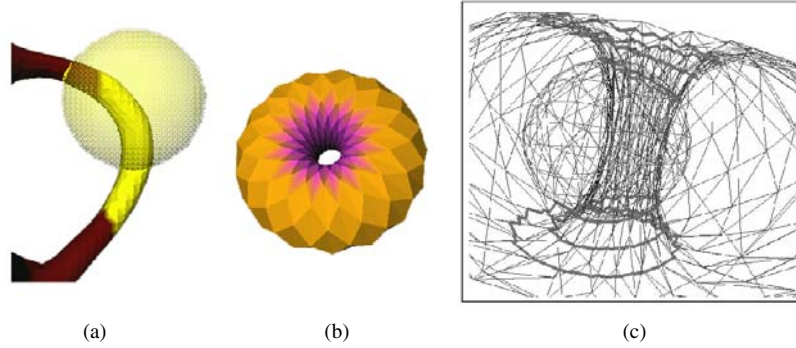


Fig. 6. A handle (a) is distinguished from a through hole (b), (c).

is concerned, only the decomposition is fully described with less emphasis on the construction of the region adjacency graph which encodes the segmentation. In the following the main steps of the classification procedure are detailed and the descriptors used to refine each class are introduced. We distinguish between geometric descriptors, which are the surface curvature and the relative length of the intersection curves, and the so-called *status* descriptors which distinguish between concave/convex or empty/full features.

4.1. Classification Based on Intersections. Given a three-dimensional mesh Σ and a set of radii $R_i, i = 1, \dots, n$, let $S(p, R_i)$ be the sphere of radius R_i and center p , and let $\gamma(p, R_i)$ be the boundary of the region of Σ containing p delimited by the intersection curves between the mesh and $S(p, R_i)$. Other regions of intersection might occur, but only the one containing p is taken into account. The first morphological characterization of the surface at a vertex p at scale R_i is given by the number of connected components of $\gamma(p, R_i)$.

We consider the following cases:

- *one component*: the surface around p can be considered topologically equivalent to a plane (see Figure 7(a)),
- *two components*: the surface around p is tubular-shaped (see Figure 7(b)),
- *three or more components*: in a neighborhood of p a branching of the surface occurs (see Figure 7(c)).

In topological terms, two components identify a handle in the object, three or more components highlight a split. If $\gamma(p, R_i)$ is made by one component, the angles excess is computed and the vertex is classified as sharp, rounded or blend (see Section 4.2). If $\gamma(p, R_i)$ is made by two components, their lengths are used to distinguish between conic and cylindrical shapes. For branching parts, no further geometric parameters are computed (see Table 1 on p. 242).

To run the process, a set of radii must be selected for the computation of the intersections. The maximum and minimum radii (R_{\max} and R_{\min} , respectively) determine an interval which is uniformly sampled according to the number of radii the user wants to

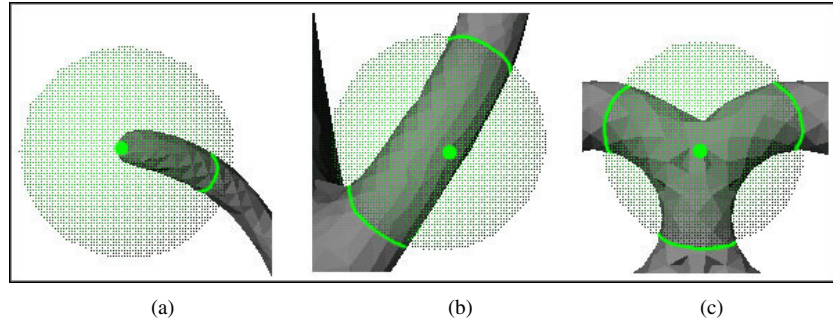


Fig. 7. Different number of connected components in the intersection boundary.

use. This step produces the values of R_i , $i = 2, \dots, n-1$ ($R_1 := R_{\min}$ and $R_n := R_{\max}$). Both R_{\min} and R_{\max} can be defined by the user by means of a slider in the Graphical User Interface; otherwise, they are automatically set proportionally to the size of the object. In this latter case, R_{\min} is set to the minimum edge length and R_{\max} to half of the diagonal bounding box of the object. The choice of the radius range obviously influences the classification results because it defines the number and sizes of neighborhoods used to analyze the surface shape around each point. In this sense, a priori knowledge about the feature size of interest surely improves the decomposition.

4.2. Curvature Characterization. As described in Section 3, when $\gamma(p, R_i)$ has only one boundary component, the curvature at the point p , at scale R_i , is the angle excess of $\gamma(p, R_i)$. Instead of using the angle excess, we use the length of $\gamma(p, R_i)$ divided by the radius R_i , i.e. $L_{\gamma(p, R_i)} = \text{length}(\gamma(p, R_i))/R_i$. Note that this value has the dimension of an angle and it always assumes a positive value. Since we want to characterize the curvature of a surface, vertices will be labeled as sharp, rounded or blend points according to their approximated curvature values by establishing some thresholds on the interval $[0, +\infty)$. We can distinguish the following cases:

- *Sharp vertices:* we consider a conic surface, with spike point p and $\alpha \in (0, \pi/2]$ the half amplitude of the cone. Intersecting the cone with a sphere centered in p and with radius R_i generates a circular curve of length $2\pi R_i \sin(\alpha)$ with $L_{\gamma(p, R_i)} = 2\pi \sin(\alpha)$, which is an increasing function of $\alpha \in (0, \pi/2]$: the lower the value of α , the more the surface around p tends to a cone-shaped point. Intuitively, we consider p a sharp vertex if $\alpha \leq \pi/4$ and consequently the curvature threshold is set to $T_s = \sqrt{2}\pi$.
- *Rounded/blend:* to distinguish between these two situations we observe that the surface is *rounded* in a neighborhood of a point if its curvature is positive while a *blend* occurs when the surface is hyperbolic, that is, its curvature is negative. The limit case between the two is obviously a *flat* point and, considering that we use $L_{\gamma(p, R_i)}$ to evaluate the curvature, it is clear that setting the threshold to $T_b = 2\pi$ discriminates between *rounded* and *blend*. The intersection between the sphere and a plane, indeed, results in an intersection curve whose length is equal to $2\pi R_i$.

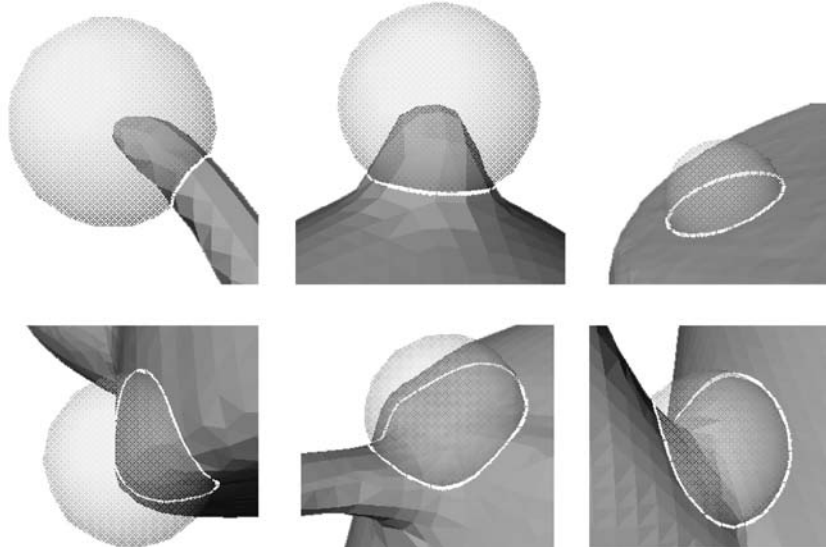


Fig. 8. Several cases of one intersection curve: note the relation between the intersection curve length and the curvature of the surface in the neighborhood of the center of the sphere.

Summarizing, the characterization of a point p at scale R_i is set as follows:

- $0 \leq L_\gamma(p, R_i) \leq \sqrt{2}\pi$: p is *sharp*,
- $\sqrt{2}\pi < L_\gamma(p, R_i) \leq 2\pi$: p is *rounded*,
- $L_\gamma(p, R_i) > 2\pi$: p is a *blend*.

For example, see Figure 8.

4.3. Relative Length Characterization. Now consider the case of two connected components in the intersection curve $\gamma(p, R_i)$. As mentioned above, this means that p lies on a region of the surface which appears as an elongated shape at that scale, like a tubular protrusion or a handle in the object. We can refine this remark as follows: if the length of the two intersection components is nearly the same, the shape at the scale R_i can be approximately considered cylindrical; if one is much longer than the other, it means that the shape may be seen as a conic part (see Figure 9). Let γ_1 and γ_2 be the two intersection components, and let l_1 and l_2 be their lengths with $l_1 \geq l_2$. The shape is considered conical if $l_1 \geq 2l_2$, cylindrical otherwise. The related threshold is $T_c = \frac{1}{2}$ which guarantees that the amount l_2/l_1 (belonging to $[0, 1]$) uniquely determines whether the local shape of the surface around p is cylindrical or conic.

4.4. Status Characterization. The extraction of morphological features on a surface is based on different operators, each related to a different aspect or property of the related feature. Up to now, we have seen how to distinguish between sharp or rounded, tubular or branching and so on, while the *status characterization* allows us to distinguish between convex/concave or, more generally, full/empty features. For instance, to discriminate




Fig. 9. Example of cylindrical (a) and conical (b) parts of a triangular mesh.

between convex and concave vertices would lead to refining the classification of a *sharp* point as a tip or a pit, a *rounded* point as a mount or a dip. Obviously, the distinction between convex and concave does not make sense for blend points. For vertices on tubular features, that is, with two or more connected components in $\gamma(p, R_i)$, it is checked if the surface intersected by the sphere encloses a volume which is inside or outside the object (see Figure 6). In this case the status characterization allows us to distinguish a handle from a deep tubular depression of the object.

We consider the status characterization in the case of one component first. As for curvature computation, concavity/convexity evaluation at a vertex p of a triangular mesh is highly affected by local noise in $Star(p)$. The local geometry is depicted in Figure 10: a given edge e shared by triangles t_1, t_2 of a mesh is convex (resp. concave) if the angle formed by t_1, t_2 , inside the object, is less (resp. more) than π .

Consequently, a given vertex p is defined strictly convex (resp. strictly concave) if all the incident edges in p are convex (resp. concave). Because in most cases the incident edges in p are both convex and concave, the previous classification cannot be applied. Note that even a slight modification of the point coordinates can produce a complete different classification. For these reasons, the method adopted for assigning a *convex* or *concave* label to a vertex p at scale R_i again uses the intersection between the mesh and




Fig. 10. Edge concavity or convexity criterion.




Fig. 11. Configuration of the intersection curve normals around a concave point.

the sphere. In the case of one connected component of the intersection curve, the center of mass b of γ and the average normal N of the intersected triangles are computed. The vertex p is considered concave (resp. convex) at scale R_i , if p lies below (resp. above) γ , that is, $N \bullet (b - p) > 0$ (resp. < 0), where “ \bullet ” denotes the inner product. We refer to Figure 11 for an easier understanding.

Suppose now that $\gamma(p, R_i)$ has two intersection components. Again we can distinguish between the case in which the local shape is a tubular protrusion or a tubular well of the surface (see Figure 6(a),(b)), in analogy with the property of convex/concave mentioned above for points generating one intersection curve.

If the number of connected components of the intersection curve is two, as in Figure 12(a), we consider the orientation of each component of $\gamma(p, R_i)$ as naturally induced by the surface orientation (see Figure 12(b)). If p lies on a tubular protrusion of the surface, it happens that the normal vector of the average plane related to each connected component of $\gamma(p, R_i)$ is directed towards p (see Figure 12(c)) according to the right-hand rule; if p appears on a tubular depression of the object, the vectors have opposite directions. This statement holds for three or more connected components too, thus it is possible to discriminate between a branch on the outer surface or a splitting cavity.




Fig. 12. Situation in which two orientation curves are generated: (a), (b) curve orientation derived by triangle orientation, and (c) average normal of intersecting curves.

5. Algorithm and Implementation Details. This section deals with the implementation details of the algorithm used for the sphere–mesh intersection, which is the core of the shape characterization. Throughout the discussion we refer to the pseudo-code given at the end of the section.

The mesh is encoded by means of a triangle-based data structure which stores:

- for each triangle, its three vertices and its three adjacent triangles, which represent the Triangle–Vertex (TV) and Triangle–Triangle (TT) relations, respectively,
- for each vertex, its coordinates and one (arbitrary) of its incident triangles (VT^*).

Actually the Vertex–Triangle relation (VT) associating to a vertex all its incident triangles is necessary to navigate the mesh, but the memory space required can be strongly optimized by coding just one of those triangles per vertex (partial Vertex–Triangle relation or VT^*). The total VT relation can be retrieved in linear time by iteratively applying the TT relation starting from the stored triangle. A scheme of the data structure is given in Figure 13. The storage of this data structure requires $3n_V * sizeof(float) + 6 * n_T * sizeof(int) + n_V * sizeof(int)$, since vertices and triangles occur in the TT , TV and VT^* structures as integers.

For each vertex v the computation of the intersection curves between the mesh and a set of n spheres centered in v with increasing radii R_1, \dots, R_n is computed as follows:

- One of the triangles incident in v is inserted in the queue Q (this operation takes constant time if we have the VT^* data structure) and it is marked as visited.
- A triangle t is extracted from Q and the main loop is repeated until Q is empty.
- The algorithm checks if the spheres intersect t : for each radius R_i , if at least two vertices p, q of t , distinct from v , satisfy the conditions $\|p-v\|_2 \leq R_i, \|q-v\|_2 \geq R_i$, t is intersected by the sphere $S(v, R_i)$. This operation takes constant time. In this case, t is considered as a seed triangle for tracing the whole line of intersection whose continuation

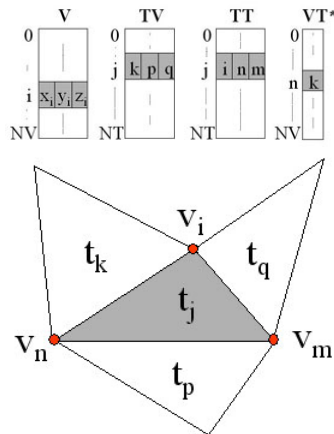


Fig. 13. Data structure organization: the information for the vertex v_i and the triangle t_j are shown.




Fig. 14. Approximated length of the intersection paths.

is searched in the triangles adjacent to t . The function *intersection* is thus invoked to complete the intersection curve starting from t and moving to its neighbors intersected by $S(v, R_i)$. The curve $t \cap S(v, R_i)$ is calculated considering the intersection points between the sphere and each edge of t . More precisely, given an edge $[a, b]$ of t , we can parameterize it as $u(s) := sa + (1-s)b$, $0 \leq s \leq 1$; the intersection points, if any, are located as $u(s_0)$, $u(s_1)$ with s_0, s_1 solutions in $[0, 1]$ of the equation $\|u(s) - v\|_2^2 = R^2$ of degree two in the unknown s . As shown in Figure 14, the length of $\gamma(v, R_i)$ is the sum of the composing arc lengths, each one belonging to an intersecting triangle, and given by $R_i\theta$ where θ is the angle \widehat{apb} , a, b being the intersection points. The call of this function increases the number of connected components of the intersection line for a given radius. Moreover, the neighbors of the intersection triangles traversed but not marked, that is, those which lay outside $S(v, R_i)$ but inside $S(v, R_n)$, are inserted in Q .

- If t is not an intersection triangle and it lies inside the sphere of maximum radius $S(v, R_n)$, its neighbor triangles (if not marked) are inserted in Q . otherwise, it is simply discarded.

The construction of a connected component of the intersection line may take as many constant operations as the number of intersection triangles, i.e. $O(n_T)$ in the worst case. However, in this implementation each triangle is visited only once: marking triangles when they are inserted in Q avoids considering them more than once, and the intersection triangles traversed during the execution of *intersection* are not stored in Q . Therefore, the main loop takes $O(n_T)$, that is, $O(n_V)$; doing this operation on the whole mesh takes $O(n_V^2)$ time. Note that if the step between the radii is small with respect to the average edge length, a triangle can be easily intersected by more than one sphere, and the function *intersection* could be invoked on the same triangle as many times as the number of radii. Anyway, once the number n of radii is chosen this is a constant value, so that the loop (L) does not increase the complexity.

```

main() {
   $Q = \emptyset$ ;
  (# connected comp. for  $R_i$ ) = 0;
  for all  $v \in V$  {
     $Q \leftarrow TV^*(v)$ ;
    // main loop
    while( $Q \neq \emptyset$ ) {
       $t =$  first element removed from  $Q$ ;
      (L) for ( $R_i = R_1, \dots, R_n$ )
        if ( $\exists v_l, v_m \in TV(t) : \|v - v_l\|_2 \leq R_i$ 
          &  $\|v - v_m\|_2 \geq R_i$ ) {
           $intersection(t, R_i, v)$ ;
          (# connected comp. for  $R_i$ ) ++ ; }
    }
  }

  for all  $t_i \in TT(t)$ 
    if( $t_i$  is not marked & ( $\exists v_l \in TV(t_i) : \|v - v_l\|_2 \leq R_n$ )
       $Q \leftarrow t_i$ ;
  }

  intersection( $t, R_i, v$ ) {
     $L_i = intersectionLength(t, R_i, v)$ ;
     $t_{next} = t$ ;
    do {
       $t_{next} = t_i \in TT[t_{next}] : t_i \cap S(v, R_i) = \emptyset$ 
       $L_i += intersectionLength(t_{next}, R_i)$ ;
      if( $t_i$  is not marked & ( $\exists v_l \in TV(j, i) : \|v - v_l\| \leq R_{max}$ ))
         $Q \leftarrow t_i$ ;
    } while( $t_{next} \neq t$ );
  }
}

```


6. Mesh Decomposition. The focus of this section is the integration of the different characterizations, described in Section 4, to achieve a unique segmentation of the input mesh into morphological features represented by closed regions with uniform properties. In Table 1 a summary of the labels assigned to vertices, for a given scale R_i , is shown. After the label assignment, the input mesh is decomposed into patches using a region-growing procedure, and each patch corresponds to a shape feature relevant at scale R_i .

The morphological classification associates a vector of feature labels to each vertex, and each label describes the vertex at the corresponding scale. Selecting the scale of interest, the surface can be rendered using a color-coding of the feature labels. The achieved decomposition is an affine-invariant segmentation into disjoint, non-empty subsets which code the geometry and shape evolution through scale changes. In Figure 15 an example is shown; the different views display the mesh decomposition at different scales, and the colors are those related to Table 1 (color figures are available in the electronic version of the paper).

Table 1. Morphological feature characterization.

Label	Feature	Color	# \cap	Geometric	Status
T	TIP	Red	1	$L/R_i \leq T_s$	Convex
P	PIT	Blue	1	$L/R_i \leq T_s$	Concave
M	MOUNT	Orange	1	$T_s < L/R_i \leq T_b$	Convex
D	DIP	Cyan	1	$T_s < L/R_i \leq T_b$	Concave
B	BLEND	Pink	1	$L/R_i > T_b$	—
L	LIMB	Yellow	2	$L_2/L_1 \geq T_c$	Full
W	WELL	Violet	2	$L_2/L_1 \geq T_c$	Empty
J	JOINT	Brown	2	$L_2/L_1 < T_c$	Full
F	FUNNEL	Gray	2	$L_2/L_1 < T_c$	Empty
S	SPLIT	Green	≥ 3	—	Full
H	HOLLOW-Y	Black	≥ 3	—	Empty

The tools defined are used to analyze a shape at different scales, but they can also be used to derive information about the *persistence* of a shape feature across the scale range. It is also possible to define a basic query language which allows us to extract features defined by the user as relations between morphological labels at different scales. For example, there has been defined a coarse feature-based query language, which allows the user to submit a query like “which are the vertices whose feature type is TIP at scale R_3 and MOUNT at scale R_6 ?” To this purpose, a query vector with wild card v_q , where $v_q[i]$ specifies the requested feature type at scale R_i , and the AND/OR Boolean operators are used. For instance, suppose we used a set of ten radii; the vector $[*,*,T,*,*,M,*,*,*,*]$ with the AND operator specify the previous query. Here, feature la-

**Fig. 15.** Shape segmentation on the pot at different scales.




Fig. 16. Queries with matched points are depicted in red. (See electronic version.) The use of AND and OR operators among the scales is specified before the label vector. The round parentheses between labels work as OR between features at the same scale.

bels are those in Table 1, and the symbol “*” means that, at that scale, each feature type is allowed.

The combination of morphological labels with logical operators enables the construction of a high-level language for shape interrogation guaranteeing a multi-task model. In fact, the user is able to extract a single shape element using a single query, to combine them and, after future improvements, to modify the geometry locally, by using other surface patches, or modifying the topology by changing the structure of the adjacency graph which codes the shape decomposition.

The results obtained (see Figure 16) suggest that further improvements of the query language will allow the extraction of higher-level features, like handles or main body of a given object. Tubular components can be extracted choosing LIMB OR JOINT vertices; among them, handles correspond to cycles in the region adjacency graph, and protrusions are adjacent to TIP or MOUNT zones. Selecting points which assume TIP OR MOUNT OR DIP OR BLEND features at most levels of detail identify the main body of the object, and so on. Using the described language, a mesh can be analyzed in a rather flexible way.




Fig. 17. Persistence analysis of the number of intersection curves (a), and its refinement using geometric information (b).

Instead, if we are willing to extract a global shape description which takes into account the whole range of scales into a single decomposition of the mesh, the following voting algorithm for persistence can be used. First, the points are classified according to the intersection connectivity, that is, according to the number of single, double or multiple components in the intersection boundary, considering the whole range of scales. The mesh is therefore segmented in parts which are characterized by having either almost always one intersection or two or more. This step provides a first insight on features which are persistently protrusion-like, handle-like or branch-like, without distinguishing if they are convex/concave or full/empty. In Figure 17(a) the result of this segmentation is shown, where the blue parts are composed by vertices having only one intersection for more than 75% of the scales, the red are those having two intersections for the same threshold. Finally the gray areas are those corresponding to shape transitions where both one and two intersections occur approximately in the same percentage.

Within each of the resulting parts, a further classification can be done considering the related geometric characterization. For example, in Figure 17(b) the shape vertices are colored with different blue saturations depending on the curvature; a darker blue corresponds to a higher curvature. An analogous criterion is applied to vertices with two intersections using their relative length where a lighter red is related to cylindrical-like features. With regard to the choice of the threshold for the persistence analysis, the value has to be tuned according to the number of radii used. In our experiments we used an average of ten radii and the value 75% provided good results. A threshold bigger than 75% has also been used, and it results in a stricter selection of patches with one or two intersections; this choice generally enlarges transition areas.

7. Method Insights. In this section results and properties of the shape characterization are discussed, focusing on the comparison between the multi-scale and the persistence analysis. Examples of an application of this method for skeleton extraction are also given.

First, we discuss the behavior of the curvature evaluation using different scales. The curvature is analyzed in a neighborhood whose size depends on R_i : for small values of R_i , such as the average length of the edges incident in p , the curvature approximation resembles the discrete curvature estimation proposed in [2] and shows similar problems




Fig. 18. The point classification corresponding to R_i chosen as the minimum edge (a): red and blue vertices locate elliptic and hyperbolic points while the green line visualizes the theoretical parabolic line (see electronic version); the results obtained with radius $2R_i$ are shown in (b), (c).

while, for increasing values, it becomes more stable to noise. We have tested the behavior of the curvature on smooth and rough surfaces, as follows.

The first test surface used is a torus and the multi-resolution curvature has been compared using two radii: R_i equal to the minimum edge of the input mesh and $2R_i$. Curvature values achieved with R_i corresponds to the theoretical point classification into parabolic, hyperbolic and elliptic regions on the torus (see Figure 18(a)). Increasing the radius to $2R_i$ results in a similar classification where parabolic points, identified as the boundary between hyperbolic and elliptic ones, are shifted with respect to their theoretical position represented by the green line (see Figure 18(b),(c)).

The shape decomposition has been tested for stability to noise also on a rough surface. In Figure 19(a),(c) the results obtained on the rabbit and on the same data set with added noise are compared. The main features, like tips and limbs, are preserved and the influence of the thresholds involved is lessened by the persistence analysis (see Figure 19(b),(d)).

The multi-scale decomposition depends on the chosen set of radii, and, on one hand, a too small radius leads to noise problems, but, on the other hand, a too large radius can give a meaningless result. Small radii can be used to determine detail features while bigger ones are able to capture the global characteristics of the surface. From these considerations, it follows that the choice of R_i is related to the scale of the features which have to be extracted, and the use of a set of increasing radii is suitable for performing a multi-scale analysis of the shape.

Multi-scale analysis can also be used for a generic segmentation of the surface according to the curvature evaluation, by grouping those points which share a similar curvature value with respect to a given threshold ε and a radius R_i , i.e.

$$p, q \text{ belong to the same patch} \iff |L_\gamma(p, R_i) - L_\gamma(q, R_i)| \leq \varepsilon.$$

This approach is commonly applied by segmentation methods based on curvature. Depending on the type and complexity of the shape to be analyzed, the simple curvature evaluation may give better results than the refined segmentation. The segmentation described by the labels in Table 1 corresponds, indeed, to a decomposition into protrusion-like and handle-like features which might not be relevant for certain shapes. As shown in Figure 20, shapes composed by detailed and unstructured features are better analyzed by multi-scale curvature evaluation, because in such case a decomposition into features, such as limbs and tips, is not meaningful.

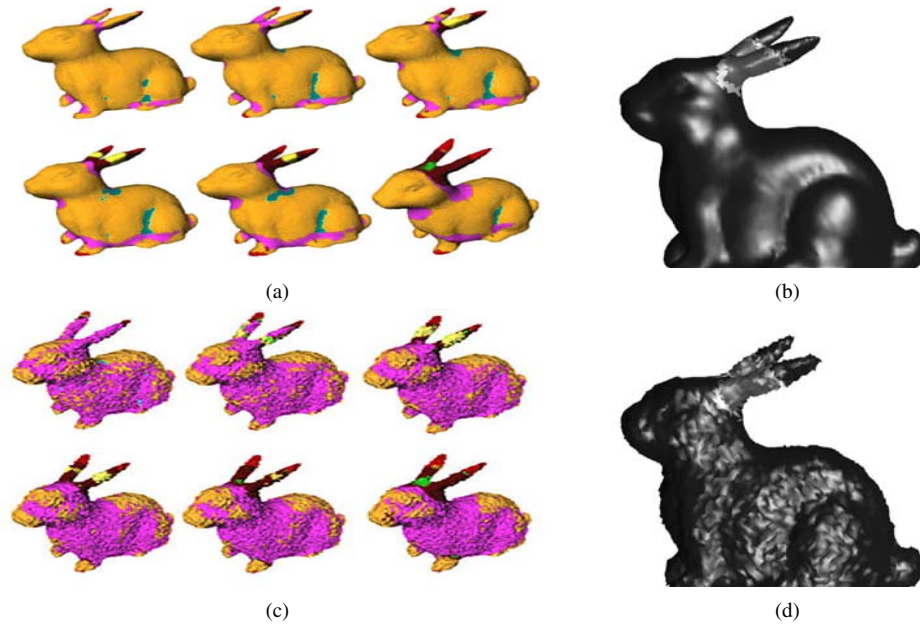


Fig. 19. (a), (b) Shape segmentation and persistence analysis on the original rabbit, and (c), (d) on the model with added noise. Segmentations achieved based on persistence analysis are nearly identical.

8. Conclusions and Future Work. The evolution of intersection curves produced by blowing bubbles at mesh vertices has been proved to be a good approach to characterize a shape using meaningful shape features. Increasing the radius of the bubbles produces an easy and efficient multi-scale analysis of the shape, which can be effectively used to produce a set of specific and flexible tools for shape analysis. The resulting description is an affine-invariant segmentation of disjoint, non-empty subsets and equally distributed in all directions, which codes the geometry and shape evolution through scale changes. Finally, the decomposition is computationally affordable and consistent with previous work on these topics.




Fig. 20. Feature decomposition on the dragon at three different scales.

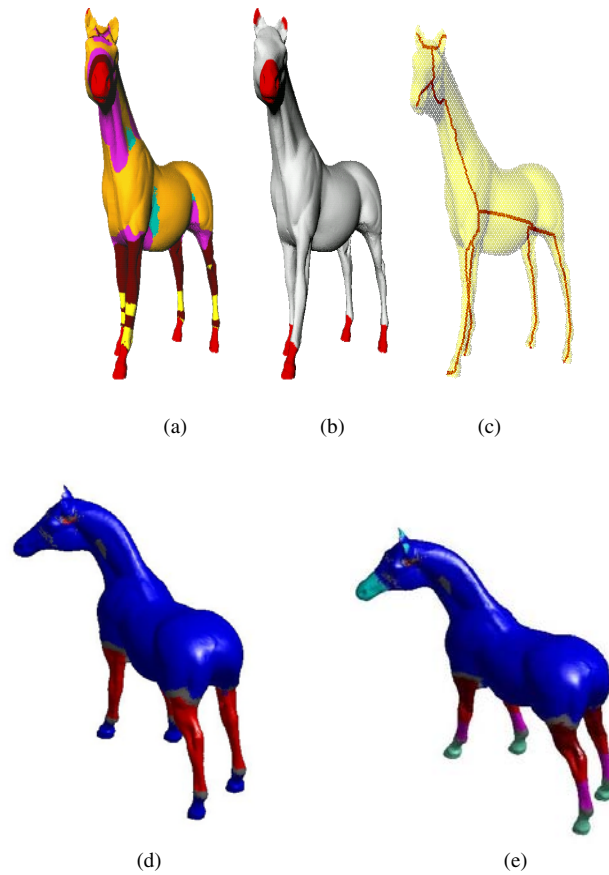


Fig. 21. Global framework: (a) curvature estimation on the horse with different radii, (b) peak regions are extracted with a query, (c) regions selected in (b) are used as seed points for extracting the skeleton, (d), coarse persistence analysis, (e) refined persistence analysis.

The usefulness of the description obtained goes beyond the problem of shape decomposition and we believe it will serve as a basis for building effective tools for shape processing and editing. For example, it has already been used to find the seed points for the construction of an affine-invariant skeleton [12] which represents the input for a wide class of applications, such as matching and topological characterization of three-dimensional shapes. Figure 21(a)–(c) shows all the steps of the construction.

Future developments will focus mainly on the specification of the feature adjacency graph and on the study of its evolution within the scale range required.

Acknowledgements. Special thanks are given to the Computer Graphics Group of IMATI-GE/CNR.

References

- [1] S. Biasotti, B. Falcidieno and M. Spagnuolo. Extended Reeb graphs for surface understanding and description. *Proceedings of the 9th Discrete Geometry for Computer Imagery Conference*. LNCS. Springer-Verlag, Uppsala, 2000, pp. 185–197.
- [2] M. Desbrun, M. Meyer, P. Schroeder and A.H. Barr. Discrete differential operators for triangulated 2-manifolds. *VisMath 2002, Proceedings*. Visualization and Mathematics III. Springer-Verlag, Heidelberg, 2003, pp. 35–57.
- [3] A.R. Dill, M.D. Levine and P.B. Noble. Multiple resolution skeletons. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **9**(4) (1987), 495–503.
- [4] B. Falcidieno and M. Spagnuolo. Geometric reasoning for the extraction of surface shape properties. In N. Magnenat Thalmann and D. Thalmann (eds.), *Communicating with the Virtual World*. Springer-Verlag, New York, 1993, pp. 166–175.
- [5] B. Falcidieno and M. Spagnuolo. A shape-abstraction paradigm for modelling geometry and semantics. *Proceedings of Computer Graphics International 1998*. IEEE Computer Society Press, Los Alamitos, CA, 1998, pp. 646–656.
- [6] R.C. Gonzales and R.E. Woods. *Digital Image Processing*. Addison-Wesley, Reading, MA, 1992.
- [7] V. Guillemin and A. Pollack. *Differential Topology*, Prentice-Hall, Englewood Cliffs, NJ, 1974.
- [8] B. Hamman. Curvature approximation for triangulated surfaces. *Computing Suppl.*, **8** (1993), 139–153.
- [9] M. Hilaga, Y. Shinagawa, T. Kohmura, and T.L. Kunii. Topology matching for fully automatic similarity estimation of 3D shapes. *Computer & Graphics, Proceeding of Siggraph 2001*. Los Angeles, 2001, pp. 203–212.
- [10] M.M. Lipschutz. *Theory and Problems of Differential Geometry*. Schaum's Outline Series. McGraw-Hill, New York, 1969.
- [11] S. Loncaric. A survey of shape analysis techniques. *Pattern Recognition*, **31**(8) (1998), 983–1001.
- [12] M. Mortara and G. Patané. Affine-invariant skeleton of 3D-shapes. *IEEE Proceeding of International Conference on Shape Modelling and Applications 2002*, Alberta, pp. 245–252.
- [13] M.E. Mortenson. *Geometric Modeling*. Wiley, New York, 1985.
- [14] R.L. Ogniewicz. Skeleton-space: a multiscale shape description combining region and boundary information. *Proceedings of Computer Vision and Pattern Recognition*, pp. 746–751, 1994.
- [15] S.M. Pizer, W.E. Oliver and S.H. Bloomberg. Hierarchical shape description via the multiresolution symmetric axis transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **9**(4) (1987), 505–511.
- [16] K. Polthier and M. Schmies. Straightest geodesic on polyhedral surfaces. In H.C. Hege and K. Polthier (eds.), *Mathematical Visualization*. Springer-Verlag, New York, 1998, pp. 135–150.
- [17] Y. Shinagawa, T.L. Kunii and Y.L. Kergosien. Surface coding based on morse theory. *IEEE Computer Graphics & Applications*, **11**(5) (1991), 66–78.
- [18] G. Taubin. Estimating the tensor curvature of a surface from a polyhedral approximation. *Proceedings of the Fifth International Conference on Computer Vision (ICCV '95)*, pp. 902–907.
- [19] E. Trucco and A. Verri. *Introductory to Techniques for 3-D Computer Vision*. Prentice-Hall, Englewood Cliffs, NJ, 1998.