Geometric models with weigthed topology

M. Attene[a], S. Biasotti[a]

[a] CNR-IMATI, Via De Marini, 6, 16149, Genova, Italy

Abstract

This paper introduces the concept of weighted topology to model a 3D object whose connectivity and metric depend on a novel notion of *weighted arc-length*. The weighted arc-length between any two points of the shape takes into account the fact that a part of the object may be either *weakly* or *strongly* connected to another part. The new model is useful to treat problems which are intrinsically not robust to small topological changes. We describe an example implementation of the model and show how it can be exploited to (1) extend the applicability domain of existing segmentation algorithms, and (2) improve the performances of a shape descriptor in a 3D object retrieval scenario.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Computational Geometry and Object Modeling—*Additional keywords:* Shape modeling, robustness, shape description.

1. Introduction

Thanks to recent technological advances on sensors and scanners, current 3D datasets are becoming larger and larger. As a consequence, precise 3D models are now made of millions of elements, and understanding the structure of such shapes in order to take decisions requires an abstraction process. On the other hand, each 3D digitization technique has some physical limitations (precision, resolution, occlusions, reflections, ...) that in some applications can lead to poor performances of the model (noise, surface holes, ...). Consequently, there is a need to abstract complex data coupled with the necessity of making a model robust to data imprecisions. In this sense, removing the unessential and keeping the relevant is the core of any good model representation.

Standard models used to represent solid objects are strongly based on how points are connected to each other. The connectivity graph of a triangle mesh, for instance, is an explicit model of the topology of the surface being represented. For some applications it has been shown that the connectivity alone can even provide a rather precise description of a shape [24]; unfortunately, they are not enough to face relatively small variations of the geometry that induce abrupt changes in topology. As an example, some shape descriptors are particularly sensitive to such small changes [31, 22], and several geometry processing and analysis algorithms are severely limited by the constraints imposed by a hard use of the encoded topology [5, 23]. Therefore, it would be desirable to have the possibility to robustly work with solid models without being spoiled by such small variations.

Figure 1: Schematic difference between standard topology (top row, abrupt change of connectivity) and weighted topology (bottom row, smooth transition from strong to weak connectivity).

The objective of this work is to propose a new perspective to study the connectedness of shapes, so that the aforementioned abrupt topological changes become smooth transitions. The approach chosen to reach this goal, which is the main contribution of this paper, is the definition of a model for solid objects where the arc-length between any two points is *weighted* to take into account that a part of the object may be either *weakly* or *strongly* connected to another part. Roughly speaking, let us imagine to have two balls in the Euclidean space and to define a solid as their union (see Figure 1). Using standard Euclidean topology, the solid is either

Email addresses: attene@ge.imati.cnr.it (M. Attene), biasotti@ge.imati.cnr.it (S. Biasotti)

connected or disconnected, depending on the relative position and size of the two spheres. Conversely, using weighted topology the solid may be either connected, disconnected, or *weakly* connected. More precisely, the connectedness between any pair of points of the solid is a number in the range $[0,1]$, where the extremes 0 and 1 indicate that the two points are disconnected or connected respectively.

In the remainder we show how weighted topology is related to fuzzy connectedness [45], and how it can be used to effectively characterize the space occupied by *blurred* solids. In turn, a blurred solid can be used as an approximated model to describe an object along with a set of possible small modifications of its shape that may also induce changes in topology. Finally, we show that our new model can be effectively employed to improve the performances of state-of-the-art shape descriptors and shape segmentation algorithms.

2. Previous work

Traditional geometric modeling techniques [33, 34] are based on a discrete representation of a geometric model which can be possibly continuous and smooth. Several models are the result of a reconstruction from a sufficiently dense point cloud [28, 1, 3, 15]. Most of these algorithms are based on the study of the shape of Voronoi cells for points sampled on smooth surfaces. For instance, it has been shown that surfaces smoothly embedded in $\mathbb{R}^3$ can be reconstructed homeomorphically from any 0.06-sampling using the Cocone algorithm [2]. For smooth submanifolds in higher dimensional spaces, similar ideas lead to provable local dimension estimation algorithms [14]. It has also been shown that appropriate offsets, or equivalently α-shapes [17], provide reconstructions with the correct homotopy type [37] under opportunely defined sampling conditions. Moreover, it has been shown that persistent homology can be used to reliably estimate Betti numbers of a wide class of compact sets (for instance it can deal also with sharp edges) under an even weaker sampling condition that depend on the weak feature size of the underlying manifold [11, 13, 9]. Finally, persistent homology has been shown to be able to code the connected components of the sublevel sets of an image also when it is partially occluded or noisily sampled [19].

Nevertheless, the definition of automatic and topology-aware methods able to reconstruct a complete watertight model from arbitrarily, poorly-sampled data scans is still a challenging task. For instance, the approach proposed in [42] requires some user input to disambiguate in the areas where the topology of the model cannot be automatically inferred with a good degree of confidence. Other methods propose a-posteriori corrections of the reconstructed topology that include skeleton-based editing [26] or topology simplification based on Reeb graphs [46].

Recently, geometry inference has been proposed as a tool to handle nonsmooth sampled objects embedded in arbitrary dimensional spaces, even in case of noise and outliers [10]. These results lead to the design of robust approximate unsigned distance functions: [35] has experimentally shown that starting from a noisily-sampled object it is possible to generate sequences of thicker bands able to capture the correct shape topology.

Despite all the efforts spent so far to deal with uncertain topology in Computer Graphics (e.g. management of outliers in reconstruction [35], imprecise annotation of sub shapes [25], weak topology characterization [42]), existing techniques are mainly based on a post-processing to *fix* the model [46] to make it suitable for downstream applications. On the contrary, we aim at defining a model which explicitly embeds the variability of its geometry and topology, so that each specific downstream application can reason on more comprehensive shape information. Therefore, our approach differs from methods that define either a surface metric, for instance based on geodesic and volume visibility such as [30], or spin images [21], because we are going to define a global representation that involves both the shape and its neighbour in the embedding space: in this sense we could configure our representation as a fuzzyfication technique.

Some attempts to embed uncertainty in the model representation have been done in Computer Vision. Besides the notion of fuzzy set and topographic connectivity [39, 40] the notions of relative and iterative relative fuzzy connectedness [45] and more recently fuzzy connectivity measures [36] have been proposed. These works introduce the concept of degree of connectedness that depends on the membership of points on paths between pairs of points of the image. The connectivity measure introduced in [36] allows the definition of connected components of an image even if two objects partially overlap and has been used to build a tree-based object description. These methods have been mainly adopted to analyze 2D and 3D images (e.g. brain, tumors and vessels segmentations). It is worth noticing that in Computer Vision there have been some attempts to perform shape analysis directly on the grey-level image or on its corresponding fuzzy-segmented image instead of using a mere crisp segmentation (eg. [8] for star-shaped images). In particular, in [8] it is stated that two different discretizations of the same continuous fuzzy descriptor generate significantly different shape descriptions: in our approach, besides the definition of shape descriptions that are available for a more general class of objects, we will overcome this kind of limitation by considering only exact Euclidean distances.

3. Topology-robust 3D object representation

After reviewing the basic concepts of arc-wise connectedness of topological spaces (Sec. 3.1), we introduce the notions of blurred solid (Sec. 3.2) and weighted connectivity (Sec. 3.3) that generalize the notions in Sec. 3.1 to obtain a model representation that robustly handles both geometric and topological noise and outliers. Its properties are discussed in Sec. 3.4.

3.1. Basic concepts

Let X be a topological space. A **path**, or **arc**, from a point $\mathbf{a}$ to a point $\mathbf{b}$ in X is a continuous function $\gamma : [0,1] \to X$ such that $\gamma(0) = \mathbf{a}$ and $\gamma(1) = \mathbf{b}$. A **path-component** of X is an equivalence class of X under the equivalence relation $\backsim$ defined by $\mathbf{a} \backsim \mathbf{b}$ if there exists a path from $\mathbf{a}$ to $\mathbf{b}$. X is said to be **path-connected** if and only if for each pair of points $(\mathbf{a},\mathbf{b})$ in X there exists a path γ such that $\gamma(0) = \mathbf{a}$ and $\gamma(1) = \mathbf{b}$, that is, if there exists only one path-component of X. A path-connected space X is also **connected**, meaning that X is not the union of two disjoint nonempty open sets.

When (X,d) is a metric space equipped with a distance d, the **arc-length** of the path γ is defined as:

$$L(\gamma) = \sup_{0=\alpha_0<\alpha_1<\ldots<\alpha_n=1} \sum_{i=0}^{n-1} d(\gamma(\alpha_i),\gamma(\alpha_{i+1})), \qquad (1)$$

where the supremum is taken over all possible partitions of $[0,1]$ and n is unbounded. In case the space X is a Riemannian manifold a more operative definition of the arc-length can be obtained integrating:

$$L(\gamma) = \int_0^1 \sqrt{g(\gamma'(t),\gamma'(t))}dt, \qquad (2)$$

where g is the metric tensor and $\gamma'(t) \in T_{\gamma(t)}X$ is the tangent vector of γ at t. Moreover, the arc-length does not change while changing the parameterization, and therefore it is an intrinsic property of the path γ.

3.2. Blurred solids

An **r-set** is a bounded and regular set [33] and we loosely call a **solid** any r-set being a subset of the three-dimensional Euclidean metric space $\mathbb{E}^3$. In this terminology a solid is not necessarily arc-connected.

Let S be a solid, and let ∂S be its boundary. We denote with $d_u(x,S)$ the Euclidean distance of a generic point x of the Euclidean space $\mathbb{E}^3$ from the closest point of ∂S. The **signed distance function** of S is the function $d_S : \mathbb{E}^3 \to \mathbb{R}$ such that:

$$d_S(x) = \begin{cases} d_u(x,S), & \text{if } x \in S \\ -d_u(x,S), & \text{otherwise} \end{cases} \qquad (3)$$

We call a **blur** of radius r of S, $r > 0$, the function $b_S^r : \mathbb{E}^3 \to [0,1]$ defined as follows:

$$b_S^r(x) = \begin{cases} \frac{d_S(x)+r}{2r}, & \text{if } -r \le d_S(x) \le r \\ 0, & \text{if } d_S(x) < -r \\ 1, & \text{if } d_S(x) > r \end{cases} \qquad (4)$$

The union of all the points $\mathbf{p} \in \mathbb{E}^3$ for which $b_S^r(\mathbf{p}) \neq 0$ defines a **blurred solid**. Thus, for any blur b_S^r we may define a corresponding blurred solid that will be denoted as $B_S^r = \bigcup_{\mathbf{p}\in\mathbb{E}^3, b_S^r(\mathbf{p})\neq 0} \mathbf{p}$.

3.3. Weighted connectivity

The notion of weighted connectivity is the key concept we will use to characterize the relationships among the points of B_S^r (i.e. in a region of the space around the solid) and it is the basis to define a point neighbor. First, let $\mathcal{A} \subset \mathbb{E}^3$ be a convex and compact superset of B_S^r (e.g. $\mathcal{A}$ can be the convex hull of B_S^r). This space plays the role of the ambient space of the solid B_S^r. Let b_S^r be the blur function, for each path $\gamma : [0,1] \to \mathcal{A}$ we define its weight $w^r(\gamma)$ as follows:

$$w^r(\gamma) = \begin{cases} \frac{1}{\int_0^1 \frac{1}{b_S^r(\gamma(t))}dt}, & \text{if } \forall t \in [0,1], b_S^r(\gamma(t)) \neq 0 \\ 0, & \text{otherwise,} \end{cases} \qquad (5)$$

where $\gamma(t)$, $t \in [0,1]$ is a parameterization of γ. Also, we define the weighted arc-length $Y^r(\gamma)$ of γ as follows:

$$Y^r(\gamma) = \begin{cases} \frac{L(\gamma)}{w^r(\gamma)}, & \text{if } \forall t \in [0,1], b_S^r(\gamma(t)) \neq 0 \\ +\infty, & \text{otherwise,} \end{cases} \qquad (6)$$

Finally, for any pair of points $\mathbf{p},\mathbf{q} \in \mathcal{A}$, we define their **weighted connectivity**, $\mathcal{W}_c^r$, as the supremum of the weights over all the paths $\gamma \in \Gamma$, i. e.,

$$\mathcal{W}_c^r = \sup_{\gamma\in\Gamma} w^r(\gamma), \qquad (7)$$

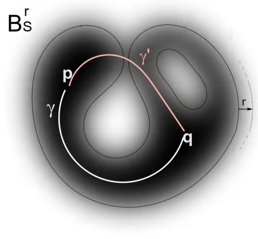

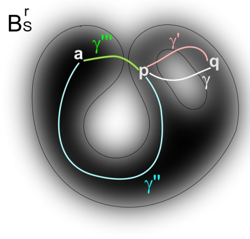

Figure 2: In the blurred solid B^r_S a path can cross ∂S: in (a) only γ is inside S, whereas in (b) only γ' is inside S.

where Γ represents the set of all paths $\gamma: [0,1] \to \mathcal{A}$ such that $\gamma(0) = \mathbf{p}$ and $\gamma(1) = \mathbf{q}$.

In Figure 2(a) a path γ is shown for which all the points are farther than r from the boundary. Therefore, its weight $w^r_\gamma = 1$. In contrast, γ' crosses the original boundary ∂S, therefore $0 \le w^r_{\gamma'} \le 1$. In figure 2(b), the weighted arc-length between p and q varies depending on the path considered; in this case $L(\gamma) = L(\gamma')$, nonetheless $Y^r(\gamma) \ge Y^r(\gamma')$ because γ crosses a more *blurred* area.

Intuitively, $\mathcal{W}^r_c$ measures the strenght of the connection between the points $\mathbf{p}, \mathbf{q}$: if there exists at least a path between $\mathbf{p}$ and $\mathbf{q}$ which is entirely contained in B^r_S, their weighted connectivity is positive, i.e., $0 < \mathcal{W}^r_c \le 1$. The points $\mathbf{p}$ and $\mathbf{q}$ are *strongly* connected if $\mathcal{W}^r_c = 1$, *weakly* connected if $0 < \mathcal{W}^r_c < 1$, *disconnected* otherwise.

3.4. Properties

The main properties of the weighted connectivity $\mathcal{W}^r_c$ and b^r_S are discussed emphasizing how they can be used to deal with complex objects possibly made of multiple components and having inner cavities.

Relation with dilation operators. The blurred solid B^r_S can be seen as the support of the blur function b^r_S and, once $r > 0$ is fixed, it is a compact space. In particular, it can be seen as a dilation of radius r of the original solid S or, better, as the r-offset of S. Anyway, B^r_S is not enough to define the weighted connectivity $\mathcal{W}^r_c$ which is generally defined over an ambient space $\mathcal{A} \subset \mathbb{E}^3$; this to include at least the whole interior part of the solid B^r_S.

Continuity and boundedness. Once r is fixed, b^r_S is a continuous function in $\mathbb{E}^3$. Based on the fact that the signed distance function d_S is 1-Lipschitz, also b^r_S is a Lipschitz continuous function[1]. To sketch the proof of this fact we consider three cases (the other ones follow by similar reasonings) where $d(x,y)$ denotes the Euclidean distance between x and y:

1. $|d_S(x)| < r$ and $d_S(y) < -r$. Then, using the triangular inequality and the hypothesis on $d_S(y)$:
$$|b^r_S(x) - b^r_S(y)| = |\frac{d_S(x)+r}{2r} - 0| \le$$
$$|\frac{d_U(x,y)+d_S(y)+r}{2r}| \le |\frac{d_U(x,y)-r+r}{2r}| = \frac{d(x,y)}{|2r|};$$
2. $|d_S(x)| < r$ and $d_S(y) > r$. Since $|d_S(x)| = d_u(x,S)$, and the $d_u(x,S)$ is the distance from the **closest** point of S, the following relations hold:
$$|b^r_S(x) - b^r_S(y)| = |\frac{d_S(x)+r}{2r} - 1| = |\frac{d_S(x)-r}{2r}| \le$$
$$\frac{|d_S(x)|}{|2r|} \le \frac{d(x,y)}{|2r|};$$
3. $|d_S(x)| < r$ and $|d_S(y)| < r$. Being the signed distance function 1-Lipschitz the following relations hold:
$$|b^r_S(x) - b^r_S(y)| = |\frac{d_S(x)+r-(d_S(y)+r)}{2r}| =$$
$$\frac{|d_S(x)-d_S(y)|}{|2r|} \le \frac{d(x,y)}{|2r|}.$$

In particular the infimum of $\frac{L(\gamma)}{w^r(\gamma)}$ over B^r_S is 0 and $\frac{1}{w^r(\gamma)}$ plays the role of a multiplier of the arc-length $L(\gamma)$. For each $\gamma \in \Gamma$, the relation $Y^r(\gamma) \ge L(\gamma)$ holds; the equality $Y^r(\gamma) = L(\gamma)$ is verified along the paths that span only the interior part of S and $d_u(x,S) \ge r$, $\forall x \in Im(\gamma)$. Finally, we notice that $w^r(\gamma)$ is bounded by 0 and 1 and, consequently, $0 \le \mathcal{W}^r_c \le 1$.

Expressiveness. The weighted arc-length defined in Eq. (6) incorporates information about the space between the two points that can be used to (partially) describe the blurred solid. We will discuss an implementation of the inner distance concept [29] driven by this notion of weighted arc-length in Section 4.3.

Relation with fuzzy connectedness. Since $0 \le b^r_S(x) \le 1$, $\forall x \in \mathbb{E}^3$, the blur b^r_S of S defines a fuzzy set on $\mathbb{E}^3$ with membership function $\mu(x) = b^r_S(x)$. Differently from the discrete notion of fuzzy connectedness defined in [45], the weighted connectivity introduced in this paper induces a smooth transition between different topologies as the geometry smoothly changes.

[1] A function $f: X \to Y$ is Lipschitz continuous if there exists a real constant $K \ge 0$ such that $\forall x,y \in X$, $d_Y(f(x),f(y)) \le K d_X(x,y)$.

Applicability. Our model representation can easily deal with objects having either multiple components or many cavities. Multi-component objects can be modeled thanks to the fact that both the blur function b_S^r and the weight connectivity do not require S to be a connected space. The weight $w^r(\gamma)$ stretches the length of $L(\gamma)$ permitting the definition of an arc-length between two points, even if they are disconnected in the traditional sense. If the object has inner cavities, the radius r can control their impact on the overall connectivity (i.e. if r is larger than the cavity, paths may cross it without the need to turn around), thus acting as the basis of a multi-level representation.

Noise treatment. Since the notion of $\mathcal{W}_c^r$ allows the definition of a weak-connectivity band in the region close to ∂S, it is possible to use this model representation to deal with noisy data, thus making it suitable for shape reconstruction and description.

Adaptivity. In principle the notions of weighted arc-length and connectivity can be defined based on non-constant values of the radius r. This fact would imply a different value of the blur radius $r_{\mathbf{p}}$ for each point $\mathbf{p}$ of ∂S, and for each point $x \in \mathbb{E}^3$ the condition $-r_{\mathbf{p}} \leq d_S(x) \leq r_{\mathbf{p}}$ is evaluated with respect to the value of $r_{\mathbf{p}}$ associated to the closest point $\mathbf{p}$ of x on ∂S. Non-constant values of r over the model can support the representation of variable entities such as, for example, magnetic fields or levels of reliability of the data.

4. Examples and preliminary results

Weighted topology is potentially useful in a number of applications, including shape segmentation, description, and processing. This section outlines our preliminary experiments and results.

4.1. A possible implementation

Despite the generality of the notions proposed in Section 3 that apply to a generic solid object, we show how the model can be effectively implemented using a rasterization of a triangle mesh that generates a 3D image.

After having ensured that the input mesh properly encloses a solid [4], we create a binary 3D image B where each voxel is *foreground* if it intersects at least one of the input triangles. By filling the interior voxels starting from such a *skin* we produce another binary 3D image S representing the solid. Then, S is converted to its signed distance function as described in [41] and, in its turn, the signed distance funtion is converted to a blurred shape R by applying Eqn. 4 to each voxel. Throughout the remainder, we will refer to S, B and R in the aforementioned meaning.

Fig. 3 depicts the main steps of the conversion process. For clarity the figure shows a single section of the blurred shape. Notice that in each section the blurred voxels depend on their 3D neighborhood, so we may observe some shape features which are not visible in the sharp version of the same section.

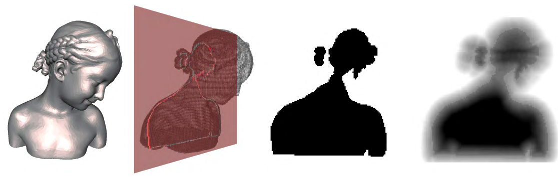

Figure 3: From left to right: the *bimba* mesh (courtesy of AIM@SHAPE's Repository); its rasterized version with a cut-plane; the section of the 3D image; the section of the final blurred shape.

Note that all the steps of this implementation proposal can be performed rather efficiently through state-of-the-art algorithms. Namely, if an input mesh has defects such as surface holes or self-intersections that prevent it to be the boundary of a solid, the algorithm described in [4] can fix it in a few seconds on a standard PC (3.2 seconds for a chair model made of 100K triangles). Furthermore, the rasterization process can be accomplished through the algorithm introduced in [16], which is able to voxelize the Blade model (1.7 million triangles) within a 256^3 grid in 75 milliseconds on the GPU of a standard PC equipped with an ATI Radeon 9800 Pro graphics card. Finally, to compute the signed distance field, the algorithm introduced in [41] guarantees a linear-time complexity, which corresponds to approximately one second for a 256^3 grid.

4.2. Shape segmentation

To show how the weighted connectivity can support existing segmentation algorithms, we have re-implemented the algorithm described in [5]. In particular, we have modified the dual graph that represents the triangle adjacencies, by assigning to each pair of triangles their weighted connectivity (the weight is 1 if the two triangles are actually adjacent).

First, we compute the weight among any pair of vertices v_i and v_j as follows: within R, let J be the segment $v_i - v_j$; each voxel traversed by J is multiplied by the length of its intersection with J. The sum of such multiplied voxels is then divided by the Euclidean distance between v_i and v_j. If all the traversed voxels are non-zero, the result is the weight between v_i and v_j which are said to be *non-disconnected*. Based on this, for each

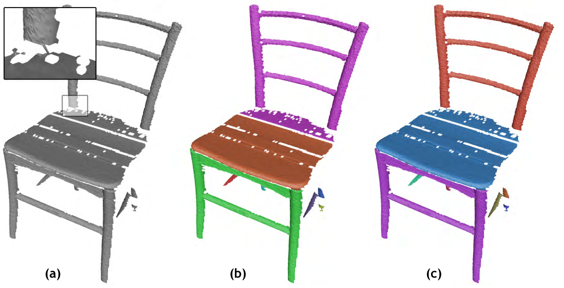

Figure 4: The scanned chair model (a) used in this example is made of a single range image which, in turn, is made of several connected components. By directly using the method in [5] on the range image, the seat and the back are not properly separated (b). Conversely, by using the blurred version of the model ($r = 10$ voxels) we could successfully segment these two features of the shape (c). Model courtesy of Aim@Shape's repository.

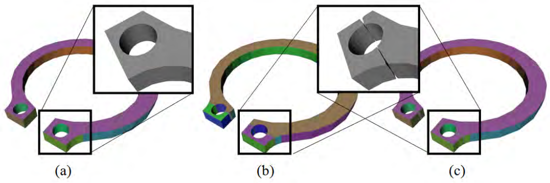

Figure 5: Mesh segmentations obtained using the method [5] on the original model (a), inserting topological modifications (b) and using the blurred solid ($r = 10$ voxels) of the modified model (c).

triangle t, we check if it has at least two vertices non-disconnected with two vertices of another triangle y. If this is the case, we create an additional arc in the dual graph whose weight is the average weight of the two triangles' non-disconnected vertices, which can be either two or three pairs. When two clusters are merged (i.e. a dual edge is contracted into a single node), to recompute the weight of the arcs incident at the resulting node we compute the average of all the contributing arcs. Finally, the cost of each contraction is recomputed by dividing the original L_2 fitting residual by the weight of the dual edge.

In other words, even if two clusters are disconnected (in standard terms), if they are sufficiently close they have still the chance to be merged, though with a lower priority. Figure 5 depicts an example that shows how a slight change in topology prevents the original algorithm to calculate a correct segmentation, while an appropriate result could be computed by using the weighted connectivity. A further, more practical example is shown if Figure 4, where one of the range images of a scanned chair model is better segmented thanks to the use of the weighted connectivity. Note that although in this case the model does not enclose a solid, we could exploit our notion of weighted connectivity by simply creating a one-voxel thick solid from the input surface mesh (i.e. each voxel which is intersected by one of the triangles is a foreground voxel).

4.3. Shape description using inner distances

The notion of weighted arc-length may be adopted to extend every shape descriptor rooted on the notion of arc-length. As an example, we consider the *inner distance* that has been recently used to describe generic 2D shapes in [29], and to compare 3D models of molecules [31]. Given a solid model bounded by a closed surface, the inner distance between two points $\mathbf{p}$ and $\mathbf{q}$, both on its boundary surface, is the length of the shortest path connecting them. Herewith, to build a shape descriptor robust to both isometric deformations and *small* topological changes, we extend the inner distance histograms to blurred solids.

Herewith, we loosely say that a deformation is *isometric* if all the distances computed on the surface are preserved after the deformation. Also, we say that a topological change is *small* if it is induced by a small deformation of the solid (i.e. no point of the solid is moved too far away from its original position).

Weighted inner distances. For each skin voxel v_i (i.e. v_i is foreground in B) we create a corresponding node $g(v_i)$ in a graph g. Then, for each pair of skin voxels v_i and v_j, we create the arc $< g(v_i), g(v_j) >$ in g if the segment J connecting the centers of v_i and v_j intersects only non zero voxels in R. Then, we multiply the value of each voxel by the length of its intersection with J, and the sum of such multiplied voxels weights $< g(v_i), g(v_j) >$ and is used as an approximation of the weighted arc-length of v_i and v_j. After g is completely built, the weighted inner distance for any node pair is computed by Dijkstra's algorithm.

To obtain a coherent description of two shapes and control the size of g, we subsample their skins using the same number of voxels and compute all the distances between these subsamples only. To actually perform this subsampling we use an adaptation of [18] on the dual graph of the skin.

Histograms. The extraction of the shape descriptor as the histogram of the inner distances between couples of points has been described in [31]. Here, it has been shown that this shape descriptor performs well when comparing molecular models. Anyhow there is a large literature in the use of histograms as signatures of object properties, see for instance [43, 32, 44]. In our implementation, we stack all the weighted inner distances into an array and create the corresponding histogram by subdividing an interval $[0, d_{max}]$, where d_{max} is an input parameter, into a fixed number N_b of equally-sized bins. The value of the bin bounded by (d_i, d_{i+1}) is equal to the number of distances belonging to range (d_i, d_{i+1}).

Given two shapes, their dissimilarity is evaluated as the L_2 distance between their inner distance histograms. Moreover, to have scale invariance, we have normalized the models with respect to their volumes. The use of the weighted connectivity significantly influences the shape classification. For instance, the six objects in Figure 6 can be clustered either into two classes made of objects with similar overall shape (i.e. the two rows in Fig. 6) or three classes of topologically equivalent objects (i.e. with respect to the number of handles). The dissimilarity matrices in Figure 7 show how the weighted version is more robust to topological changes.

Figure 6: These six objects can be classified in two different (and both meaningful) ways, depending on the focus on either the overall shape or the number of handles.

As a further example we consider a dataset made of 120 models subdivided into five classes of objects (20 human bodies, 30 cups, 30 eyeglasses, 20 aircrafts and 20 ants). In any of the five classes, each model may differ from the others for either a small topological change, for an isometric deformation, or for details of the geometry. See for instance the varying human models in Figure 8. On this dataset we have performed two types of experiments: (1) the comparison with other state-of-the-art methods and (2) the analysis of the performances when the parameter r varies. Table 1 shows some statistics on shape retrieval over our 120 model dataset comparing the performances of the weighted inner distance

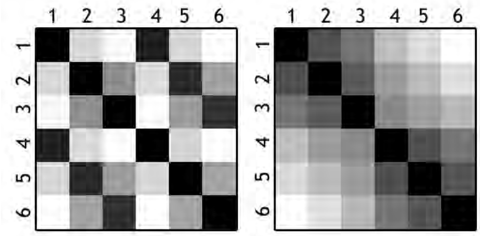

Figure 7: Dissimilarity matrices obtained using the inner distance histograms and the weighted extension on the models in Figure 6. Dark pixels correspond to a good matches while bright ones denote bad ones.

Table 1: Statistics over our 120 model dataset.

Method	NN	FT	ADR
WID ($r = 60$)	98.333%	74.208%	70.677%
Inner distance	99.167%	65.583%	66.355%
SH [27]	95.833%	68%	67.432%
LF [12]	96.5%	73.319%	72.458%

(WID) against the classical inner distance, the spherical harmonics [27] and lightfield [12] descriptors. The performance indicators reported in Table 1 are rather standard in information retrieval [38, 20]; in particular, NN represents the rate of shapes whose nearest neighbor is in the same class, FT is the first tier, and ADR represents the average dynamic recall.

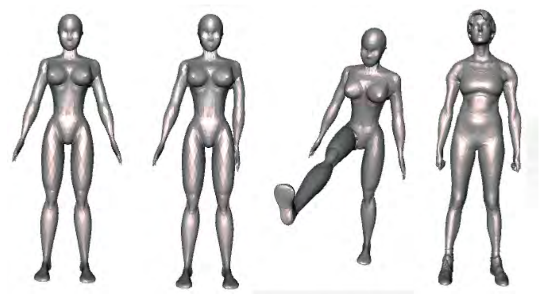

Figure 8: Four models of the class *humans* in our generic dataset. Models in a class may differ either by small topological changes, by isometric deformations, or by details of the geometry.

When considering r as a variable of the method, we observed that the retrieval performances smoothly increase as r increase up to a local maximum, around $r = 60$ voxels. For larger values of r, the performances smoothly degradate. This fact is not surprinsing bacause too small values of the blur are not able to overcome the topological changes, whereas too big values produce an indistict model.

In a different experiment, we observed that the weighted inner distance outperforms the classical one on the SHREC dataset made of 1500 models [6] affected by noise and topological artefacts, where the success rate of nearest neighbor classifier of the weighted inner distance computed with $r = 20$ passes from 78% to 81%.

5. Discussion

Clearly, the choice of the blur radius r influences both the quality of the results and the performances of the algorithms. On one extreme, when r is zero the blur is actually a "normal" solid thus, in principle, this value is not worth an actual use unless it is employed for comparison purposes.

Note that the objective of this paper is to introduce an original concept for solid models which provides advantages with respect to the state of the art within a domain of particularly topology-sensitive problems. For this reason, the task of searching an optimal value of r for each specific implementation and problem tackled, though being definitely interesting and important, is still to be undertaken and is part of our future developments. In this article, we only show that there exists a value of r that makes the use of weighted topology advantageous in two key case studies, and for a specific shape retrieval scenario we could calculate an optimal value of $r = 60$ voxels.

We have experimented that weighted topology can effectively enlarge the applicability spectrum of a classical segmentation algorithm [5]. Differently from more recent approaches such as [30], the adapted segmentation described in Section 4.2 can correctly capture even disconnected segments whose "noisy" topology would corrupt the computation of the part-aware metric. The application to topologically-robust shape description introduced in Section 4.3 can be compared to the approach proposed in [7]. Also in this case, the use of a blurred solid makes it possible to tolerate a larger set of small topological changes; for example, we can detect the similarity of a single object with a version of it which is broken into several pieces, provided that such pieces are placed sufficiently close and in the correct position. Clearly, such a modification would spoil the results of any method based solely on a surface metric, even if it is diffused such as in [7].

The conversion of an input mesh into a blurred solid based on the implementation proposed in Section 4.1 does not require a significant amount of time. Specifically, the time spent to produce the results shown in this paper is reported in Table 2. Furthermore, the 120 signatures for the models in the database used in Section 4.3 were computed in about 14 minutes, which includes the whole processing pipeline starting from the mesh loading, ranging through its voxelization and blurring, and ending to the saving of the resulting signature. Under the same circumstances, the SHREC dataset made of 1500 models described in [6] was processed in 2.5 hours.

6. Conclusion and future work

Weighted topology allows a robust analysis of abstract and qualitative aspects of the shapes, and is useful to improve the performances of shape descriptors based on distances, and in particular their robustness when the objects are subject to small topological changes. Furthermore, the new model extends the application domain of classical mesh segmentation algorithms. The new model is useful in numerous applications, ranging from mesh processing to advanced shape analysis, and we plan to investigate these areas in the future. On the theoretical side, we plan to study how the concept of Betti numbers can be extended to blurred solids with weighted topology. We expect that further studies allow to derive important properties that can be exploited to devise new robust algorithms with strong guarantees.

Acknowledgements

References

[1] Adamson, A., Alexa, M., 2003. Approximating and intersecting surfaces from points. In: Symp. on Geometry Processing. pp. 230–239.

[2] Amenta, N., Bern, M., 1999. Surface reconstruction by voronoi filtering. Discrete and Computational Geometry 22, 481–504.

[3] Amenta, N., Choi, S., Kolluri, R. K., 2000. The power crust, unions of balls, and the medial axis transform. Comp. Geom.: Theory&Applicat. 19, 127–153.

[4] Attene, M., 2010. A lightweight approach to repair polygon meshes. The Visual Computer.

[5] Attene, M., Falcidieno, B., Spagnuolo, M., 2006. Hierarchical mesh segmentation based on fitting primitives. The Visual Computer 22 (3), 181–193.

[6] Biasotti, S., Attene, M., 2008. Shape retrieval contest 2008: Stability of watertight models. In: SMI. pp. 217–218.

[7] Bronstein, A. M., Bronstein, M. M., Kimmel, R., Mahmoudi, M., Sapiro, G., September 2010. A gromov-hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. Int. J. Comput. Vision 89, 266–286.

[8] Chanussot, J., Nyström, I., Sladoje, N., 2005. Shape signatures of fuzzy star-shaped sets based on distance from the centroid. Pattern Recognition Letters 26 (6), 735 – 746.

[9] Chazal, F., Cohen-Steiner, D., Lieutier, A., March 2009. A sampling theory for compact sets in euclidean space. Discrete Comput. Geom. 41, 461–479.

[10] Chazal, F., Cohen-Steiner, D., Mérigot, Q., 2009. Geometric inference for measures based on distance functions. Tech. Rep. 00383685, INRIA.

Table 2: *Timing results*. For each model, this table reports the number of triangles in the input, the time spent to convert the mesh into a 256^3 voxelization, the time to create the blur, the time to perform the segmentation and to compute the inner distance histogram. All the times are expressed in seconds, with the exception of the voxelization which is expressed in milliseconds. The results for the *VaseCub* and *VaseRnd* models are referred to the first model on the left in the figure.

Model Name	**Fig.**	**Input Triangles**	**Voxelization time (ms)**	**Blurring Time (s)**	**Segmentation time (s)**	**Histogram Computation (s)**
bimba	3	1005382	143	1.74	-	-
chair	4	86902	46	0.97	5.81	-
CAD	5	30588	38	0.91	3.76	-
VaseCub	6(top)	95776	68	1.64	-	3.21
VaseRnd	6(bottom)	64232	64	1.61	-	3.74

[11] Chazal, F., Lieutier, A., 2007. Stability and computation of topological invariants of solids in $\mathbb{R}^n$. Discrete and Computational Geometry 37, 601–617.

[12] Chen, D., Ouhyoung, M., Tian, X., Shen, Y., 2003. On visual similarity based 3D model retrieval. Computer Graphics Forum 22, 223–232.

[13] Cohen-Steiner, D., Edelsbrunner, H., Harer, J., 2007. Stability of persistence diagrams. Discrete Comput. Geom. 37 (1), 103–120.

[14] Dey, T. K., Giesen, J., Goswami, S., Zhao, W., 2003. Shape dimension and approximation from samples. Discrete Comput. Geom. 29 (3), 419–434.

[15] Dey, T. K., Goswami, S., 2006. Provable surface reconstruction from noisy samples. Comput. Geom. Theory Appl. 35 (1), 124–141.

[16] Dong, Z., Chen, W., Bao, H., Zhang, H., Peng, Q., 2004. Real-time voxelization for complex polygonal models. In: PG '04: Proceedings of the Computer Graphics and Applications, 12th Pacific Conference. IEEE, Washington, DC, USA, pp. 43–50.

[17] Edelsbrunner, H., 1995. The union of balls and its dual shape. Discrete Comput. Geom. 13, 415–440.

[18] Eldar, Y., Lindenbaum, M., Porat, M., Zeevi, Y., 2002. The farthest point strategy for progressive image sampling. Image Processing, IEEE Transactions on 6 (9), 1305–1315.

[19] Frosini, P., Landi, C., Jan. 2010. Stability of multidimensional persistent homology with respect to domain perturbations. ArXiv e-prints.

[20] Giorgi, D., Biasotti, S., Paraboschi, L., 2007. Watertight models track. Tech. Rep. 09, IMATI, Genova, Italy.

[21] Halma, A., ter Haar, F., Bovenkamp, E., Eendebak, P., van Eekeren, A., 2010. Single spin image-icp matching for efficient 3d object recognition. In: Proceedings of the ACM workshop on 3D object retrieval. 3DOR '10. ACM, New York, NY, USA, pp. 21–26.

[22] Hilaga, M., Shinagawa, Y., Kohmura, T., Kunii, T., 2001. Topology matching for fully automatic similarity estimation of 3D shapes. In: SIGGRAPH'01. pp. 203–212.

[23] Hoppe, H., 1996. Progressive meshes. In: SIGGRAPH. ACM, pp. 99–108.

[24] Isenburg, M., Gumhold, S., Gotsman, C., 2001. Connectivity shapes. In: Procs. VIS'01. pp. 135–552.

[25] Itskovich, A., Tal, A., November 2010. Surface partial matching & application to archaeology. Computers & Graphics.

[26] Ju, T., Zhou, Q.-Y., Hu, S.-M., July 2007. Editing the topology of 3d models by sketching. ACM Trans. Graph. 26.

[27] Kazhdan, M., Funkhouser, T., Rusinkiewicz, S., 2003. Rotation invariant spherical harmonic representation of 3d shape descriptors. In: Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing. SGP '03. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, pp. 156–164.

[28] Levin, D., 2003. Mesh-independent surface interpolation. Geometric Modeling for Scientific Visualization 3, 37–49.

[29] Ling, H., Jacobs, D., 2007. Shape classification using the inner-distance. IEEE PAMI, 286–299.

[30] Liu, R., Zhang, H., Shamir, A., Cohen-Or, D., 2009. A part-aware surface metric for shape analysis. Computer Graphics Forum 28 (2), 397–406.

[31] Liu, Y., Fang, Y., Ramani, K., 2009. Idss: deformation invariant signatures for molecular shape comparison. BMC Bioinformatics 10 (1), 157.

[32] Mahmoudi, M., Sapiro, G., January 2009. Three-dimensional point cloud recognition via distributions of geometric distances. Graph. Models 71, 22–31.

[33] Mäntylä, M., 1988. Introduction to Solid Modeling. WH Freeman & Co. New York, NY, USA.

[34] Mortenson, M. E., 1986. Geometric Modeling. John Wiley & Sons.

[35] Mullen, P., De Goes, F., Desbrun, M., Cohen-Steiner, D., Alliez, P., 2010. Signing the unsigned: Robust surface reconstruction from raw pointsets. Computer Graphics Forum 29 (5), 1733–1741.

[36] Nempont, O., Atif, J., Angelini, E., Bloch, I., 2009. A new fuzzy connectivity measure for fuzzy sets. J. Math. Imaging Vis. 34 (2), 107–136.

[37] Niyogi, P., Smale, S., Weinberger, S., 2008. Finding the homology of submanifolds with high confidence from random samples. Discrete Comput. Geom. 39 (1-3), 419–441.

[38] Rijsbergen, C. J. V., 1979. Information Retrieval. Butterworth.

[39] Rosenfeld, A., 1979. Fuzzy digital topology. Information and Control 40 (1), 76 – 87.

[40] Rosenfeld, A., 1984. The fuzzy geometry of image subsets. Pattern Recognition Letters 2 (5), 311 – 317.

[41] Saito, T., Toriwaki, J., 1994. New algorithms for Euclidean distance transformation of an n-dimensional digitized picture with applications. Pattern recognition 27 (11), 1551–1565.

[42] Sharf, A., Lewiner, T., Shklarski, G., Toledo, S., Cohen-Or, D., July 2007. Interactive topology-aware surface reconstruction. ACM Trans. Graph. 26.

[43] Shilane, P., Min, P., Kazhdan, M., Funkhouser, T., 2004. The princeton shape benchmark. In: Proceedings of the Shape Modeling International 2004. IEEE Computer Society, Washington, DC, USA, pp. 167–178.

[44] Tangelder, J. W., Veltkamp, R. C., September 2008. A survey of content based 3D shape retrieval methods. Multimedia Tools Appl. 39, 441–471.

[45] Udupa, J., Saha, P., Lotufo, R., 2002. Relative fuzzy connectedness and object definition: Theory, algorithms, and applications

in image segmentation. IEEE PAMI 24 (11), 1485–1500.
[46] Wood, Z., Hoppe, H., Desbrun, M., Schröder, P., April 2004. Removing excess topology from isosurfaces. ACM Trans. Graph. 23, 190–208.