Stable volumetric features in deformable shapes

R. Litman[a], A. M. Bronstein[a], M. M. Bronstein[b]

[a] *School of Electrical Engineering, Tel Aviv University, Israel*
[b] *Institute of Computational Science, Faculty of Informatics, Università della Svizzera Italiana (USI), Lugano, Switzerland*

Abstract

Region feature detectors and descriptors have become a successful and popular alternative to point descriptors in image analysis due to their high robustness and repeatability, leading to a significant interest in the shape analysis community in finding analogous approaches in the 3D world. Recent works have successfully extended the maximally stable extremal region (MSER) detection algorithm to surfaces. In many applications, however, a volumetric shape model is more appropriate, and modeling shape deformations as approximate isometries of the volume of an object, rather than its boundary, better captures natural behavior of non-rigid deformations. In this paper, we formulate a diffusion-geometric framework for volumetric stable component detection and description in deformable shapes. An evaluation of our method on the SHREC'11 feature detection benchmark and SCAPE human body scans shows its potential as a source of high-quality features. Examples demonstrating the drawbacks of surface stable components and the advantage of their volumetric counterparts are also presented.

1. Introduction

The success of feature-based methods in image analysis and computer vision has recently driven the increased interest in developing similar methods for the analysis of 3D shapes, where the ability to detect and describe stable features is pivotal, for example, in shape correspondence, matching [TK10, HFG*06] and retrieval [MGGP06, OBBG09, TCF09].

Research in this domain has roughly followed one of the two paths: On the one hand, it is possible to apply almost straightforwardly some successful image analysis methods to 3D shapes. Notable examples of feature detectors and descriptors that follow analogous methods in image analysis include corners [SB10] and edges [KST09], histograms of gradients [ZBVH09] (similar in principle to the scale invariant feature transform [Low04] used in images), and 3D integral invariants [MGP05] (first proposed in images in [MHYS04]). On the other hand, it is possible to use 3D shape specific geometric structures for designing feature detectors and descriptors that have no direct analogs in the image domain. Spin image descriptors [JH99] and the recently introduced heat kernel signatures [SOG09] are just a few well-known examples.

1.1. Related work

A somewhat different class of feature detection methods explored in this paper try to find *regions* rather than points in the shape or image. In the image analysis community, the maximally stable extremal regions (MSER) algorithm [MCUP04] and its precursors [CB97, VS02] attempt to find level sets with the smallest area variation across intensity, resulting in regions which are affine-invariant. Furthermore, region-based features are often more robust and more informative and have become a popular alternative to their point-based counterparts in image analysis [MTS*05]. Donoser and Bischof [DB06] proposed the maximally stable volumes (MSV) detector as a direct extension of the original MSER algorithm [MCUP04] and applied it to 3D medical images. Similarly to MSER, the MSV method supports only scalar volumetric data, and was not tested quantitatively as a feature detector.

In the 3D shape analysis domain, stable regions have been explored in the works on shape decomposition, topological persistence [ELZ02], persistence-based clustering and segmentation [CGOS09], and detection of extrinsic [MPS*03] and intrinsic [TBW*11] tube-like structures. Digne *et al.* [DMAMS10] extended the notion of vertex-weighted component trees to meshes and proposed to detect MSER regions using the mean curvature. In a followup work, Litman *et al.* [LBB11] generalized the MSER framework using edge-weighted representations.

The main purpose of this paper is to develop region-based detectors and descriptors for 3D shapes represented as volumes, which are often more advantageous compared to modeling shapes as 2D surfaces [GSCO07, RBBK10]. All of the stable-regions-detectors mentioned in the previous paragraph are computed on the 2D surface of the shape, and therefore less sensitive to some "unnatural" volume-changing transformations. Volumetric representations have been employed in shape retrieval applications by Reuter *et al.* [RWP05] for the construction of global shape descriptors using Laplacian spectra, pose-oblivious signatures [GSCO07], and heat kernel descriptors [RBBK10].

1.2. Main contribution

In this paper, we expand the framework introduced in [LBB11] to stable volumetric components, with an emphasis on their use for shape matching. Our construction is based on volumetric heat diffusion [RBBK10], offering a physically meaningful analysis of 3D shapes. The method allows using either

Figure 1: Stable volumetric regions detected on the SCAPE data-set [ASK*05]. Shown are volumetric regions (left column) and their projections onto the boundary surface (right column) - note that large projected regions may be partially hidden by smaller ones that are nested in them. Corresponding regions are denoted with like colors (black indicates parts of the shape with no region detected on). The detected components are invariant to isometric deformations of the volume.

vertex- or edge weighting, a more generic construction compared to MSER or MSV, which allows for the use of vector-valued data. Our volumetric component detection approach is thoroughly evaluated on the SHREC'11 feature detection and description benchmark [BBB*11] and compares favorably to the mesh-based component detection. We also introduce a scale-invariant version of the vHKS descriptor proposed by Raviv *et al.* [RBBK10], and show that it has potential for matching regions originating from 3D human scan with synthetic human shapes.

2. Diffusion geometry

Let us consider the shape of a 3D physical object, modeled as a connected and compact region $X \subset \mathbb{R}^3$. The boundary of the region ∂X is a closed connected two-dimensional Riemannian manifold. In many application in graphics, geometry processing, and pattern recognition, one seeks geometric quantities that are invariant to inelastic deformations of the object X [RWP05, Lev06, SOG09]. Traditionally in the computer graphics community, 3D shapes are modeled by considering their 2D boundary surface ∂X, and deformations as isometries of ∂X preserving its Riemannian metric structure. In the following, we refer to such deformations as *boundary isometries*, as opposed to a smaller class of *volume isometries* preserving the metric structure inside the volume X (volume isometries are necessarily boundary isometries, but not vice versa – see Figure 2 for an illustration). Raviv *et al.* [RBBK10] argued that the latter are more suitable for modeling realistic shape deformations than boundary isometries, which preserve the area of ∂X, but not necessarily the volume of X.

2.1. Diffusion on surfaces

Recent work [RWP05, CL06, Lev06, Rus07, OSG08, SOG09, BK10] studied intrinsic description of shapes by analyzing heat diffusion processes on ∂X, governed by the *heat equation*

$$\left(\frac{\partial}{\partial t} + \Delta_{\partial X}\right) u(t,x) = 0, \tag{1}$$

where $u(t,x) : [0,\infty) \times \partial X \to [0,\infty]$ is the heat value at a point x in time t, and $\Delta_{\partial X}$ is the positive-semidefinite Laplace-Beltrami operator associated with the Riemannian metric of ∂X. The solution of (1) corresponding to a point initial condition $u(0,x) = \delta(x,y)$, is called the *heat kernel* and represents the amount of heat transferred on ∂X from x to y in time t due to the diffusion process. In particular, the diagonal of the heat kernel or the *auto-diffusivity function* $h_t(x,x)$ describes the amount of heat remaining at point x after time t. Its value is related to the Gaussian curvature $K(x)$ through $h_t(x,x) \approx \frac{1}{4\pi t}\left(1 + \frac{1}{6}K(x)t + O(t^2)\right)$, which describes the well-known fact that heat tends to diffuse slower at points with positive curvature, and faster at points with negative curvature. Due to this relation, the auto-diffusivity function was used by Sun *et al.* [SOG09] and Gebal *et al.* [GBAL09] as a local surface descriptor referred to as *heat kernel signature* (HKS). Being intrinsic, the HKS is invariant to boundary isometries of ∂X.

The heat kernel is easily computed using the spectral decomposition of the Laplace-Beltrami operator [Lev06],

$$h_t(x,y) \;=\; \sum_{i\geq 0} e^{-\lambda_i t}\phi_i(x)\phi_i(y), \tag{2}$$

where $\phi_0 = \text{const}, \phi_1, \phi_2, ...$ and $\lambda_0 = 0 \leq \lambda_1 \leq \lambda_2 ...$ denote, respectively, the eigenfunctions and eigenvalues of $\Delta_{\partial X}$ operator satisfying $\Delta_{\partial X}\phi_i = \lambda_i \phi_i$.

2.2. Volumetric diffusion

Instead of considering diffusion processes on the boundary surface ∂X, Raviv *et al.* [RBBK10] considered diffusion *inside* the volume X, arising from the Euclidean volumetric heat equation with Neumann boundary conditions,

$$\begin{aligned} \left(\frac{\partial}{\partial t} + \Delta\right) U(t,x) = 0 \qquad & x \in \text{int}(X); \\ \langle \nabla U(t,x), n(x)\rangle = 0 \qquad & x \in \partial X. \end{aligned} \tag{3}$$

Here, $U(t,x) : [0,\infty) \times \mathbb{R}^3 \to [0,\infty]$ is the volumetric heat distribution, Δ is the Euclidean positive-semidefinite Laplacian,

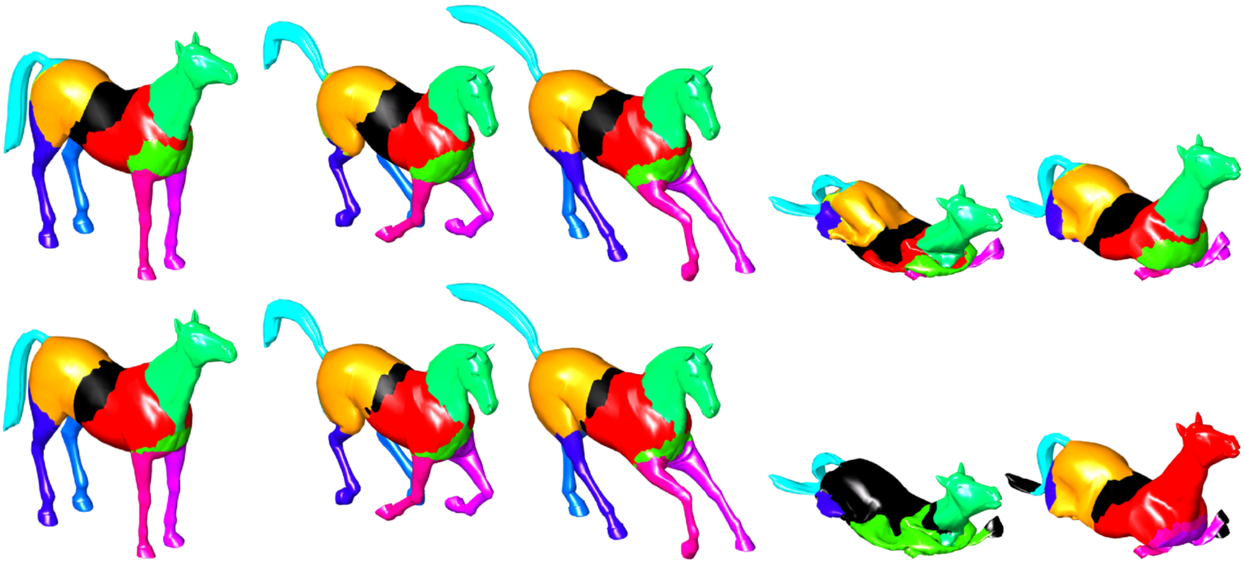

Figure 2: Maximally stable components detected on two approximate volume isometries (second and third columns) and two volume-changing approximate boundary surface isometries (two right-most columns) of the horse shape (left column). Stable regions detected on the boundary surface (2D MSER, first row) remain invariant to all deformations, while the proposed volumetric stable regions (3D MSER, second row) maintain invariance to the volume-preserving deformations only. Corresponding regions are denoted with like colors. For easiness of comparison, volumetric regions are projected onto the boundary surface.

and $n(x)$ is the normal to the surface ∂X at point x. The heat kernel of the volumetric heat equation (3) is given, similarly to (2) by

$$H_t(x,y) \quad = \quad \sum_{i \geq 0} e^{-\Lambda_i t} \Phi_i(x) \Phi_i(y), \tag{4}$$

where Φ_i and Λ_i are the eigenfunctions and eigenvalues of Δ satisfying $\Delta\Phi_i = \Lambda_i\phi_i$ and the boundary conditions $\langle \nabla\Phi_i(x), n(x)\rangle = 0$. The diagonal of the heat kernel $H_t(x,x)$ gives rise to the *volumetric HKS* (vHKS) descriptor [RBBK10], which is invariant to volume isometries of X. Compared to the 2D HKS, such descriptors were shown to be less sensitive to geometric and topological noise [RBBK10].

3. Computational aspects

Both the boundary of an object discretized as a mesh and the volume enclosed by it discretized on a regular Cartesian grid can be represented in the form of an undirected graph. In the former case, the vertices of the mesh form the vertex set V while the edges of the triangles constitute the edge set E. In the latter case, the vertices are the grid point belonging to the solid, and the edge set is constructed using the standard 6- or 26-neighbor connectivity of the grid (for points belonging to the boundary, some of the neighbors do not exist). With some abuse of notation, we will denote the graph by $X = (V, E)$ treating, whenever possible, both cases in the same way. Due to the possibility to express all the diffusion-geometric constructions in the spectral domain, their practical computation boils down to the ability to discretize the Laplacian.

3.1. Surface Laplace-Beltrami operator

In the case of 2D surfaces, the discretization of the Laplace-Beltrami operator of the surface ∂X and can be written in the generic matrix-vector form as $\Delta_{\partial X}\mathbf{f} = \mathbf{A}^{-1}\mathbf{W}\mathbf{f}$, where $\mathbf{f} = (\mathbf{f}(\mathbf{v_i}))$ is a vector of values of a scalar function $f : \partial X \to \mathbb{R}$ sampled on $V = \{v_1, ..., v_N\} \subset \partial X$, $\mathbf{W} = \mathrm{diag}\left(\sum_{\mathbf{l}\neq\mathbf{i}} \mathbf{w_{il}}\right) - (\mathbf{w_{ij}})$ is a zero-mean $N \times N$ matrix of weights, and $\mathbf{A} = \mathrm{diag}(\mathbf{a_i})$ is a diagonal matrix of normalization coefficients [FH05, WMKG07]. We used the popular *cotangent weight* scheme [MDSB03] to set the values of W and A. The eigenfunctions and eigenvalues of $\Delta_{\partial X}$ are found by solving the generalized eigendecomposition problem $\mathbf{W}\phi_\mathbf{i} = \mathbf{A}\phi_\mathbf{i}\lambda_\mathbf{i}$ [Lev06]. Heat kernels are approximated by taking a finite number of eigenpairs in the spectral expansion.

3.2. Volumetric Laplacian

In the 3D case, we used a ray shooting method to create rasterized volumetric shapes, i.e. every shape was represented as arrays of voxels on a regular Cartesian grid, allowing us to use the standard Euclidean Laplacian. The Laplacian was discretized using a 6-neighborhood stencil. We use the finite difference scheme to evaluate the second derivative in each direction in the volume, and enforced boundary conditions by zeroing the derivative outside the shape.

The construction of the Laplacian matrix under these conditions boils down to this element-wise formula (up to multiplicative factor):

$$(\Delta)_{ij} = \begin{cases} -1 & \text{if } i \neq j \text{ and } (v_i, v_j) \in E \\ -\sum_{k\neq j} (\Delta)_{kj} & \text{if } i = j \end{cases}$$

4. Stable region detection

4.1. Component trees

Let $X = (V, E)$ be a graph representing the 2D or 3D shape with the vertex set V containing the vertices of the mesh discretizing the boundary surface or the grid points discretizing the volume, and the edge set E containing the edges of the mesh in case of a 2D mesh, or capturing the standard 6-neighbor connectivity in the discretized volume). We call the graph *connected* if there exists a path between every pair of vertices in it. A graph $Y = (V' \subseteq V, E' \subseteq E)$ is called a *subgraph* of X and denoted by $Y \subseteq X$. We say that Y is a (connected) *component* of X if for any connected subgraph Z, $Y \subseteq Z \subseteq X$ implies $Y = Z$ (i.e., it is maximal). Given $E' \subseteq E$, the graph induced by E' is the graph $Y = (V', E')$ whose vertex set is made of all vertices belonging to an edge in E', i.e., $V' = \{v \in V : \exists v' \in V, (v, v') \in E'\}$.

Non-negative scalar functions $f : V \to \mathbb{R}$ and $d : E \to \mathbb{R}$ defined on the vertex and edge sets are called *vertex weight* and *edge-weight*, respectively. In the following, we will show different constructions of vertex and edge weighting functions.

Given a vertex-weighted graph (X, f) with a weighting function f, the *ℓ-cross-section* of X is defined as the graph induced by $E_\ell = \{(v_1, v_2) \in E : f(v_1), f(v_2) \le \ell\}$ for some $\ell \ge 0$. Similarly, a cross-section of an edge-weighted graph (X, d) is induced by the edge subset $E_\ell = \{e \in E : d(e) \le \ell\}$. A connected component of the cross-section is called an ℓ-level set of the weighted graph.

For any component C of X, we define the *altitude* $\ell(C)$ as the minimal ℓ for which C is a component of the ℓ-cross-section of X. Altitudes establish a partial order relation on the connected components of X as any component C is contained in a component with higher altitude. The set of all such pairs $(\ell(C), C)$ therefore forms a *component tree*.

4.2. Maximally stable components

Let us further assign to each vertex v in the graph a local element of *volume* (area in the case of 2D surfaces), denoted by $d\mathrm{vol}(v)$. The volume $\mathrm{vol}(C)$ of a component C is given as the sum of the elements $d\mathrm{vol}(v)$ over all vertices in C.

Given a sequence $\{(\ell, C_\ell)\}$ of nested components forming a branch in the component tree, the *stability* of C_ℓ is defined as

$$s(\ell) \quad = \quad \frac{\mathrm{vol}(C_\ell)}{\frac{d}{d\ell}\mathrm{vol}(C_\ell)}. \tag{5}$$

In other words, the more the relative volume of a component changes with the change of ℓ, the less stable it is. A component C_{ℓ^*} is called *maximally stable* if the stability function has a local maximum at ℓ^*. Maximally stable components are widely known in the computer vision literature under the name of *maximally stable extremal regions* (*MSER*) [MCUP04], with $s(\ell^*)$ usually referred to as the region *score*.

The construction of weighted component trees is based on the observation that the vertex set V can be partitioned into disjoint sets which are merged together going up in the tree. Maintaining and updating such a partition is done using the *union-find* algorithm and related data structures with $O(N \log\log N)$ complexity [NC06]. Such an approach is used to implement single-link agglomerative clustering which is adopted here for the construction of the component tree.

The derivative (5) of the relative component volume with respect to ℓ constituting the stability function is computed using finite differences in each branch of the tree. The function evaluation and its local maxima are detected in a single pass over the branches of the component tree starting from the leaf nodes. Furthermore, maximally stable regions with too small values of s are removed. If overlapping (over some threshold) maximally stable components are detected, only the bigger one is kept.

4.3. Vertex- and edge weighting functions

Using vertex-weighting, any scalar function that distinguishes between vertices and captures the local geometrical properties such as mean curvature [DMAMS10] can be used. For non-rigid shape analysis, diffusion-geometric weights have a clear advantage being deformation-invariant [SOCG10], and easily computed through the heat kernel. In the following discussion, we assume 3D solid objects equipped with the volumetric heat kernel H_t; the same construction applies to 2D surfaces with the corresponding heat kernel h_t.

The simplest vertex weight is obtained as the diagonal of the heat kernel $f(v) = H_t(v, v)$, which, up to a monotonic transformation, can be thought of as an approximation of the Gaussian curvature at t. The choice of the parameter t defines the scale of such an approximation [SOG09]. A scale-invariant version of this weight (the *commute-time kernel*) is obtained by integrating H_t over all time scales in the range $[0, \infty)$,

$$f(v) \quad = \quad \sum_{i=0}^{\infty} \Lambda_i^{-3/2} \Phi_i^2(v). \tag{6}$$

Edge weights offer more flexibility being able to express dissimilarity relations between adjacent vertices. Since the heat kernel $H_t(v_1, v_2)$ represents the proximity or "connectivity" of two vertices v_1, v_2, any function of the form $d(v_1, v_2) = \eta(H^{-1}(v_1, v_2))$ can define an edge weight inversely proportional to the heat kernel value (here η denotes a non-negative monotonic function).

5. Region descriptors

The diagonal of the volumetric heat kernel $H_t(x, x)$ gives rise to the vHKS descriptor defined at each point $x \in X$ for $t \in \{t_1, \ldots, t_q\}$ some fixed scale values. If one wishes to obtain a boundary surface descriptor, it is possible either to use the HKS descriptor using the surface heat kernel $h_t(x, x)$ for $x \in \partial X$, or sample the vHKS values on the boundary surface, $H_t|_{\partial X}$.

5.1. Scale-invariant volumetric descriptors

A notable disadvantage of heat kernel signatures in general is their sensitivity to scale. Given a shape X and its version X' uniformly scaled by the factor of a, it is easy to establish that eigenfunctions and eigenvalues of the the volumetric Laplacian Δ are scaled inversely proportionally to the volume of X ($\propto a^3$), $\Lambda' = a^{-2}\Lambda$, and $\Phi' = a^{-3/2}\Phi$, from which it follows that the corresponding heat kernel satisfies

$$H'_t(x, y) = \sum_{i \ge 0} e^{-\Lambda_i a^{-2} t} \Phi_i(x)\Phi_i(y) a^{-3} = a^{-3} H_{a^{-2}t}(x, y). \tag{7}$$

A scale-invariant version of the vHKS descriptor (referred to hereinafter as SI-vHKS) can be obtained in a few possible ways. Here, we use the the Fourier modulus transformation [BK10]. For this purpose, the scale parameter of the heat kernel signature is sampled logarithmically, $H(\tau) = H_{\alpha^\tau}(x,x)$. In this scale-space, the heat kernel of the scaled shape becomes $H'(\tau) = a^{-3}H(\tau + 2\log_\alpha a)$. Second, in order to remove the dependence on the multiplicative constant a^{-3}, we take the logarithm followed by a derivative w.r.t. the scale,

$$\begin{aligned} \frac{d}{d\tau}\log H'(\tau) &= \frac{d}{d\tau}\left(-3\log a + \log H(\tau + 2\log_\alpha a)\right) \\ &= \frac{\frac{d}{d\tau}H(\tau + 2\log_\alpha a)}{H(\tau + 2\log_\alpha a)}. \end{aligned} \tag{8}$$

Denoting

$$\tilde{H}(\tau) = \frac{\frac{d}{d\tau}H(\tau)}{H(\tau)} = \frac{-\sum_{i\geq 0}\Lambda_i\alpha^\tau \log\alpha e^{-\Lambda_i\alpha^\tau}\Phi_i^2(x)}{\sum_{i\geq 0}e^{-\Lambda_i\alpha^\tau}\Phi_i^2(x)}, \tag{9}$$

one thus has a new function $\tilde{H}$ which transforms as $\tilde{H}'(\tau) = \tilde{H}(\tau + 2\log_\alpha a)$ as a result of scaling. Finally, by applying the Fourier transform to $\tilde{H}$, the shift becomes a complex phase,

$$\mathcal{F}_\tau[\tilde{H}'](\omega) = \mathcal{F}_\tau[\tilde{H}](\omega)e^{-j\omega 2\log_\alpha a}, \tag{10}$$

and after taking the absolute value in the Fourier domain, we obtain scale invariance, $|\mathcal{F}_\tau[\tilde{H}'](\omega)| = |\mathcal{F}_\tau[\tilde{H}](\omega)|$.

5.2. Region descriptors

Let us now assume that we are given a point-wise descriptor at each vertex v, represented as a vector-valued function $\alpha : V \to \mathbb{R}^q$ (for example, α can be the discretized vHKS or the SI-vHKS descriptor described above). The simplest way to define a *region descriptor* of a component $C \subset V$ is by computing the average of α in C,

$$\beta(C) = \sum_{v\in C}\alpha(v)da(v). \tag{11}$$

The resulting region descriptor $\beta(C)$ is a vector of the same dimensionality q as the point descriptor α.

In [LBB11] a region descriptor was created by measuring the proximity of α to words in a pre-made "geometric vocabulary" in the descriptor space. This method showed some improvement in the performance of the HKS descriptor, but not in its scale-invariant counterpart, SI-HKS.

6. Results

To evaluate the proposed feature detector, we performed one evaluation of volumetric (3D) MSER invariance, and two experiments comparing between 3D MSER and 2D MSER: a visual comparison of the invariance of the two methods to boundary and volume isometric deformations, and a quantitative comparison evaluating the sensitivity of the two methods to shape transformations and artifacts on the SHREC'11 benchmark [BBB*11]. In addition, we also evaluated the discriminativity of region-based descriptors and performed a qualitative region-matching experiment.

The shapes in all the datasets used in our experiments were originally represented as triangular meshes. For the computation of the volumetric regions, meshes were rasterized in a cube with variable number of voxels per dimension (usually around 100-130) in order for the resulting shapes to contain approximately 45K voxels.

6.1. Detector Results

Typical times for component tree construction were 2 sec for the 2D case and 6 sec for the 3D case.Stable component detection was less than 1 sec.

A remark about visualization: Even though detected regions also reside in the interior of the shape, we chose to display only their projection back onto the original 2D surface (with the exception of Figure 1). The latter projection was done by assigning each mesh vertex the label of the (Euclidean) nearest-voxel. While such a mapping does not reveal the whole information contained in the volumetric regions, it does allow comparison of volume MSER with its surface counterpart. In order to keep images compact we also display detected regions on top of each other, resulting in parts of bigger regions being hidden by smaller regions nested in them.

In the first experiment, we applied the proposed approach to the SCAPE dataset [ASK*05], containing a scanned human figure containing around 12.5K vertices in various poses. Figure 1 shows that the detected components are consistent and remain invariant under pose variations.

In the second experiment, we used the data from [SP04]. The dataset contained an animatic sequence of a horse shape represented as a triangular mesh with approximately 8.5K vertices and included both boundary isometries (collapsing and gallop) and volume isometries (gallop). Figure 2 shows that while the surface MSERs are invariant to both types of transformations, the proposed volumetric MSERs remain invariant only under volume isometries, changing quite dramatically if the volume is not preserved – a behavior consistent with the physical intuition.

In the third experiment, we used the SHREC'11 feature detection and description benchmark [BBB*11] to evaluate the performance of the 2D and 3D region detectors and descriptors under synthetic transformations, including: *isometry*, *micro holes*, *scaling*, affine transformation, *Gaussian-* and *shot noise*, *downsampling*, and *rasterization*. Each transformation appeared in five different strengths. Vertex-wise correspondence between the transformed and the null shape was given and used as the ground truth in the evaluation of region detection *repeatability*, computed as the percentage of regions that are detected in both shapes at corresponding locations and have overlap above some threshold. Figure 4 shows the repeatability of the 3D and 2D MSERs. We conclude that volumetric regions exhibit similar or slightly superior repeatability compared to boundary regions, especially for large overlaps (above 80%). We attribute the slightly lower repeatability in the presence of articulation transformations ("isometry") to the fact that these transformations are almost exact isometries of the boundary, while being only approximate isometries of the volume. Another reason for the latter degradation may be local topology changes that were manifested in the rasterization of the shapes in certain isometries. These topology changes affected the quality of detected regions in their vicinity. Although the construction of the MSER feature detector is not affine-invariant, excellent repeatability under affine transformation is observed. We believe that this and other invariance properties are related to the properties of the component trees (which are stronger than those of the weighting functions) and intend to investigate this phenomenon in future studies.

6.2. Descriptor Results

In the fourth experiment, we computed the SI-vHKS descriptor for every volumetric stable region detected in the third experiment. Settings from [RBBK10] were used. For comparison, we also computed the SI-HKS descriptors on the boundary for the detected 2D regions using the settings from [LBB11]. For each of the descriptors, the *matching score*, defined as the percentage of regions whose nearest neighbor (in the descriptor space) in the other shape is overlapping with the corresponding groundtruth region above a fixed threshold, was measured (Figure 5). We note that results reported on the SHREC11 benchmark can be compared to other results reported in [BBB*11]. The combination of volumetric regions with volumetric descriptors show highest performance over the entire range of deformations. While conducting this experiment we also benchmarked other descriptors: vHKS (shown in [RBBK10]), vocabulary-dependant descriptors (shown in [LBB11]) and 3D integral invariants (shown in [MGP05]). The two former descriptors have shown performance inferior to those of SI-vHKS, a result consistent with [LBB11]. Similarly, the 3D integral invariants (which lack invariance to non-rigid deformation) have also shown results worse than SI-vHKS. Details about the other descriptors' performance are beyond the scope of this paper.

In the fifth experiment, we attempted matching of volumetric edge-weighted MSERs with SI-vHKS descriptors computed on the SCAPE dataset to region descriptor computed on the TOSCA dataset. Figure 3 demonstrates that while the two datasets differ drastically from each other (SCAPE shapes are 3D scans of human figures, while TOSCA contains synthetic shapes, half of which are not human), consistent correct matches were found. A quantitative experiment would have required ground truth information which we did not have, and is beyond the scope of this paper.

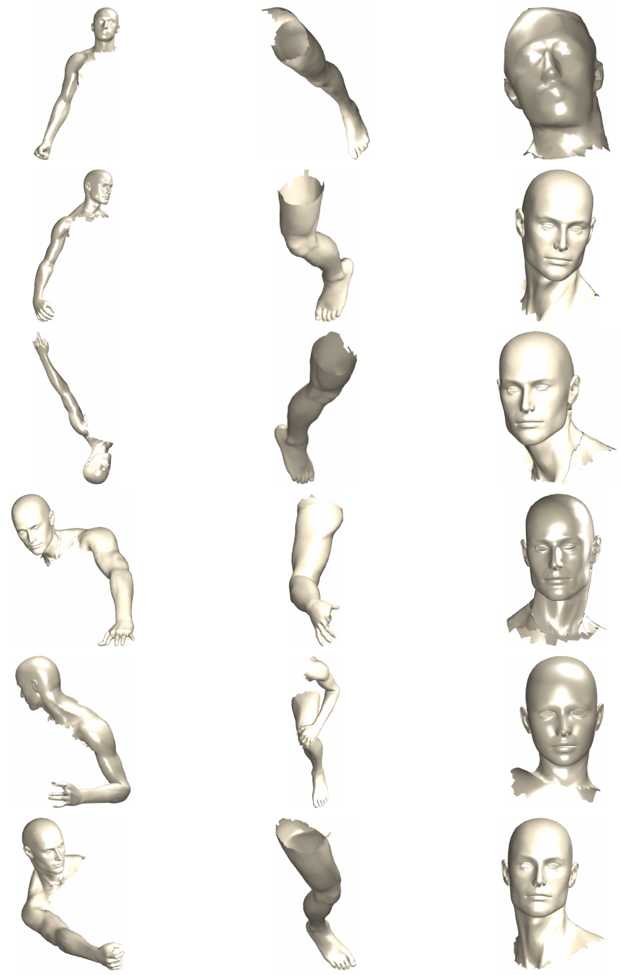

Figure 3: Examples of closest matches found for different query regions from the SCAPE dataset on the TOSCA dataset. Shown in each row from top to bottom are: query, 1st, 2nd, 4th, 10th, and 15th matches. Edge-weight $1/H_t(v_1, v_2)$ was used as the detector; average SI-vHKS was used as the descriptor.

7. Conclusions

We showed a construction of semi-local volumetric features in deformable shapes. Our approach is based on the maximization of a stability criterion in a component tree representation of the shape with vertex or edge weights derived from the properties of diffusion processed in volumes. The proposed features are invariant to non-rigid deformations, insensitive to local topological changes of the shape, and can be made invariant to global scaling. We also argue and exemplify that unlike features constructed from the boundary surface of the shape, our volumetric features are *not* invariant to volume-changing deformations of the solid object. We believe that this is the desired behavior in many applications, as volume isometries better model natural deformations of objects than boundary isometries.

We showed experimentally the high repeatability and discriminativity of the proposed volumetric features, which makes them a good candidate for a wide range of shape representation and retrieval tasks. In all experiments, our volumetric features exhibited higher robustness to deformation compared to similar features computed on the two-dimensional boundary of the shape.

Acknowledgment

We are grateful to Dan Raviv for providing us his volume rasterization and Laplacian disretization code. M. M. Bronstein is partially supported by the Swiss High-Performance and High-Productivity Computing (HP2C) grant. A. M. Bronstein is partially supported by the Israeli Science Foundation and the German-Israeli Foundation.

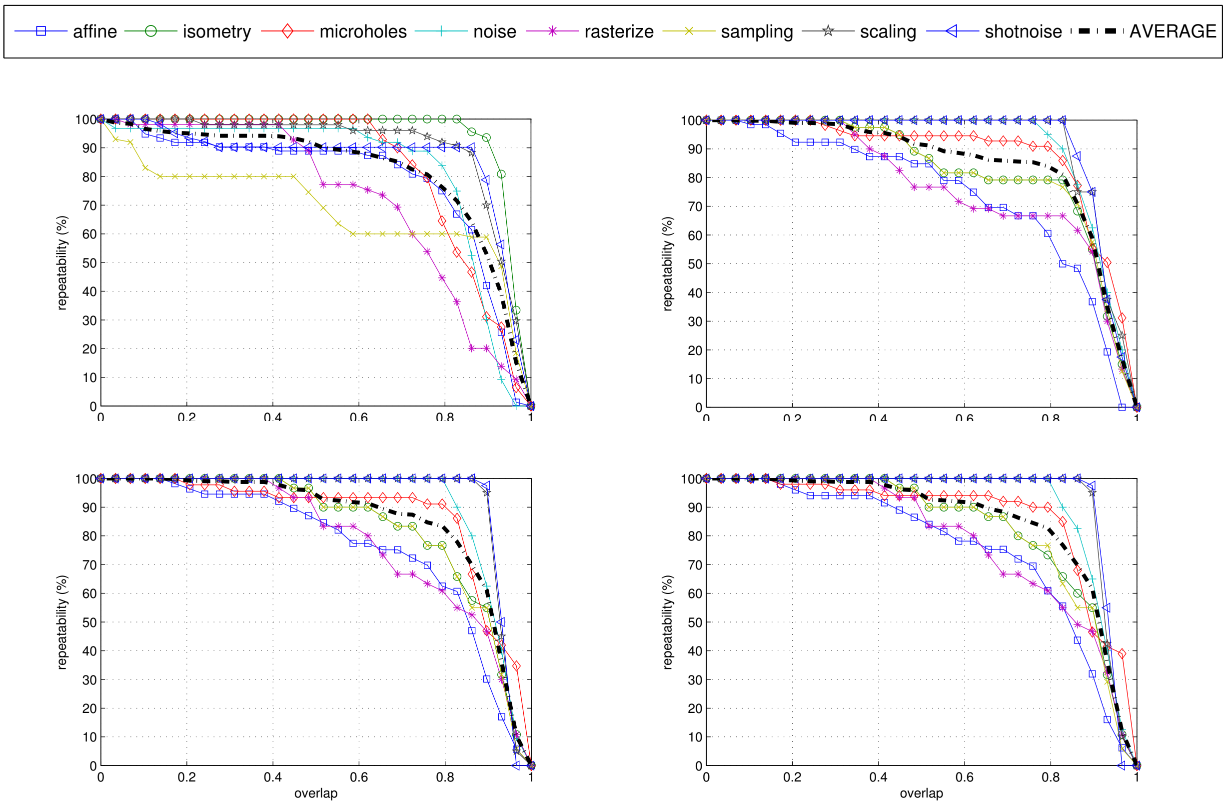

Figure 4: Repeatability of region detectors on the SHREC'11 dataset, based on different weighting functions. Upper left: 2D MSER using the edge weight $1/h_t(v_1, v_2)$, $t = 2048$. Upper right: 3D MSER using the commute-time vertex-weight. Lower left: 3D MSER using the edge weight $1/H_t(v_1, v_2)$, $t = 2048$. Lower right: 3D MSER using the vertex-weight $H_t(v, v)$, $t = 2048$.

References

[ASK*05] Anguelov D., Srinivasan P., Koller D., Thrun S., Rodgers J., Davis J.: SCAPE: shape completion and animation of people. *TOG 24*, 3 (2005), 408–416.

[BBB*11] Boyer E., Bronstein A. M., Bronstein M. M., Bustos B., Darom T., Horaud R., Hotz I., Keller Y., Keustermans J., Kovnatsky A., Litman R., Reininghaus J., Sipiran I., Smeets D., Suetens P., Vandermeulen D., Zaharescu A., Zobel V.: Shrec '11: Robust feature detection and description benchmark. In *Proc. 3DOR* (2011), pp. 71–78.

[BK10] Bronstein M., Kokkinos I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Computer Vision and Pattern Recognition* (june 2010), pp. 1704 –1711.

[CB97] Couprie M., Bertrand G.: Topological grayscale watershed transformation. In *SPIE Vision Geometry V Proceedings* (1997), vol. 3168, pp. 136–146.

[CGOS09] Chazal F., Guibas L. J., Oudot Y. S., Skraba P.: *Persistence-based clustering in riemannian manifolds*. Research Report RR-6968, INRIA, 2009.

[CL06] Coifman R., Lafon S.: Diffusion maps. *Applied and Computational Harmonic Analysis 21*, 1 (2006), 5–30.

[DB06] Donoser M., Bischof H.: 3d segmentation by maximally stable volumes (msvs). In *proc. of the 18th International Conference on Pattern Recognition* (2006), vol. 1, IEEE Computer Society, pp. 63–66.

[DMAMS10] Digne J., Morel J.-M., Audfray N., Mehdi-Souzani C.: The level set tree on meshes. In *Proceedings of the Fifth International Symposium on. 3D Data Processing, Visualization and Transmission* (May 2010).

[ELZ02] Edelsbrunner H., Letscher D., Zomorodian A.: Topological persistence and simplification. *Discrete and Computational Geometry 28*, 4 (2002), 511–533.

[FH05] Floater M. S., Hormann K.: Surface parameterization: a tutorial and survey. In *In Advances in Multiresolution for Geometric Modelling* (2005), vol. 1, Springer, pp. 157–186.

[GBAL09] Gebal K., Bærentzen J., Aanæs H., Larsen R.: Shape analysis using the auto diffusion function. In *Computer Graphics Forum* (2009), vol. 28, pp. 1405–1413.

[GSCO07] Gal R., Shamir A., Cohen-Or D.: Pose-oblivious shape signature. *Trans. Visualization and Computer Graphics 13*, 2 (2007), 261–271.

[HFG*06] Huang Q., Flöry S., Gelfand N., Hofer M., Pottmann H.: Reassembling fractured objects by geometric matching. *ACM Trans. Graphics 25*, 3 (2006), 569–578.

[JH99] Johnson A. E., Hebert M.: Using spin images for efficient object recognition in cluttered 3D scenes. *Trans. PAMI 21*, 5 (1999), 433–449.

[KST09] Kolomenkin M., Shimshoni I., Tal A.: On edge detection on surfaces. In *Computer Vision and Pattern Recognition* (2009), vol. 0, IEEE Computer Society, pp. 2767–2774.

[LBB11] Litman R., Bronstein A. M., Bronstein M. M.: Diffusion-geometric maximally stable component detection in deformable shapes. *Comput. Graph. 35* (June 2011), 549–560.

[Lev06] Levy B.: Laplace-beltrami eigenfunctions towards an algorithm that understands geometry. In *proc. of the IEEE International Conference on Shape Modeling and Applications 2006* (2006), IEEE Computer Society, p. 13.

[Low04] Lowe D.: Distinctive image features from scale-invariant keypoint. *IJCV 60*, 2 (2004), 91–110.

[MCUP04] Matas J., Chum O., Urban M., Pajdla T.: Robust wide-baseline stereo from maximally stable extremal regions. *Image and Vision Computing 22*, 10 (2004), 761–767.

[MDSB03] Meyer M., Desbrun M., Schroder P., Barr A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization and Mathematics III* (2003), 35–57.

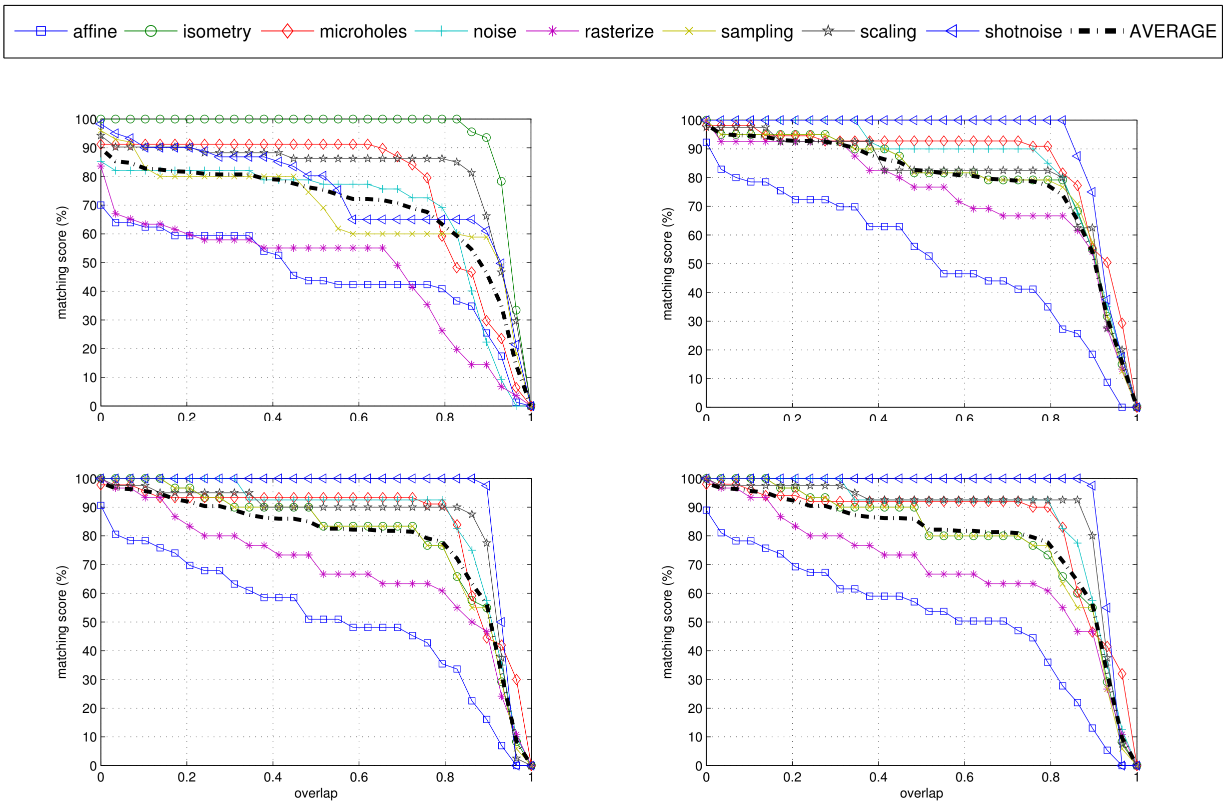

Figure 5: Matching score of descriptors based on the regions detected with the four detectors shown in Figure 4. Left-to right top-to-bottom are: 2D SI-HKS based on 2D MSER using the edge weight $1/h_t(v_1, v_2)$ with $t = 2048$, 3D SI-vHKS based on 3D MSER using the commute-time vertex-weight, 3D SI-vHKS based on 3D MSER using the edge weight $1/H_t(v_1, v_2)$ with $t = 2048$ and 3D SI-vHKS based on 3D MSER using the vertex-weight $H_t(v, v)$ with $t = 2048$.

[MGGP06] Mitra N. J., Guibas L. J., Giesen J., Pauly M.: Probabilistic fingerprints for shapes. In *proc. of the fourth Eurographics symposium on Geometry processing* (2006), pp. 121–130.

[MGP05] Mitra N. G. N. J., Guibas L. J., Pottmann H.: Robust Global Registration. In *proc. of the third Eurographics symposium on Geometry processing* (2005), pp. 197–206.

[MHYS04] Manay S., Hong B., Yezzi A., Soatto S.: Integral invariant signatures. *Lecture Notes in Computer Science* (2004), 87–99.

[MPS*03] Mortara M., Patane G., Spagnuolo M., Falcidieno B., Rossignac J.: Blowing bubbles for multi-scale analysis and decomposition of triangle meshes. *Algorithmica 38*, 1 (2003), 227–248.

[MTS*05] Mikolajczyk K., Tuytelaars T., Schmid C., Zisserman A., Matas J., Schaffalitzky F., Kadir T., Gool L.: A comparison of affine region detectors. *IJCV 65*, 1 (2005), 43–72.

[NC06] Najman L., Couprie M.: Building the component tree in quasi-linear time. *IEEE Trans. Image Proc. 15*, 11 (2006), 3531–3539.

[OBBG09] Ovsjanikov M., Bronstein A., Bronstein M., Guibas L.: Shape google: a computer vision approach to isometry invariant shape retrieval. In *Computer Vision Workshops (ICCV Workshops)* (2009), pp. 320 –327.

[OSG08] Ovsjanikov M., Sun J., Guibas L.: Global intrinsic symmetries of shapes. *Computer Graphics Forum 27*, 5 (2008), 1341–1348.

[RBB*11] Raviv D., Bronstein A. M., Bronstein M. M., Kimmel R., Sochen N.: Affine-invariant geodesic geometry of deformable 3d shapes. *Computers and Graphics 35*, 3 (2011), 692 – 697.

[RBBK10] Raviv D., Bronstein M. M., Bronstein A. M., Kimmel R.: Volumetric heat kernel signatures. In *Proc. ACM Workshop on 3D Object Retrieval* (2010), pp. 39–44.

[Rus07] Rustamov R. M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. SGP* (2007), pp. 225–233.

[RWP05] Reuter M., Wolter F.-E., Peinecke N.: Laplace-spectra as fingerprints for shape matching. In *Proc. ACM Symp. Solid and Physical Modeling* (2005), pp. 101–106.

[SB10] Sipiran I., Bustos B.: A robust 3D interest points detector based on Harris operator. In *Proc. 3DOR* (2010), Eurographics, pp. 7–14.

[SOCG10] Skraba P., Ovsjanikov M., Chazal F., Guibas L.: Persistence-based segmentation of deformable shapes. In *Proc. NORDIA* (2010), pp. 45–52.

[SOG09] Sun J., Ovsjanikov M., Guibas L.: A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. In *Computer Graphics Forum* (2009), pp. 1383–1392.

[SP04] Sumner R. W., Popović J.: Deformation transfer for triangle meshes. vol. 23, ACM Transactions on Graphics (Proc. SIGGRAPH), pp. 399–405.

[TBW*11] Tevs A., Berner A., Wand M., Ihrke I., Seidel H.-P.: Intrinsic shape matching by planned landmark sampling. *Computer Graphics Forum 30*, 2 (2011), 543–552.

[TCF09] Toldo R., Castellani U., Fusiello A.: Visual vocabulary signature for 3d object retrieval and partial matching. In *Proc. 3DOR* (2009), pp. 21–28.

[TK10] Thorstensen N., Keriven R.: Non-rigid shape matching using geometry and photometry. In *Computer Vision ACCV 2009* (2010), vol. 5996, pp. 644–654.

[VS02] Vincent L., Soille P.: Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE Trans. PAMI 13*, 6 (2002), 583–598.

[WMKG07] Wardetzky M., Mathur S., Kaelberer F., Grinspun E.: Discrete laplace operators: no free lunch. In *proc. of Eurographics Symposium on Geometry Processing* (2007), pp. 33–37.

[ZBVH09] Zaharescu A., Boyer E., Varanasi K., Horaud R.: Surface feature detection and description with applications to mesh matching. In *Proc. CVPR* (june 2009), pp. 373 –380.