

Radial Symmetry Detection and Shape Characterization with the Multiscale Area Projection Transform

A. Giachetti¹ and C. Lovato²

¹Department of Computer Science, University of Verona, Italy

²Department of Neurological, Neuropsychological, Morphological and Movement Sciences, University of Verona, Italy




Figure 1: A,B: volume rendering of the area projection transform computed at two different scales on an ant model is high near centers of radial symmetry at the selected scales. The joint multiscale map computed on a wide radius range can simultaneously detect centers of symmetry at variable scales (C). The behavior of these maps is qualitatively unchanged removing randomly 50% of the faces (D) or, thanks to adaptive smoothing (see text), adding relevant noise to vertex positions (E).

Abstract

We present a novel method to characterize 3D surfaces through the computation of a function called (multiscale) area projection transform, measuring the likelihood of points in the 3D space to be center of radial symmetry at selected scales (radii). The function is derived through a simple geometric framework based on parallel surfaces and can be easily computed on triangulated meshes. It measures locally the area of the surface well approximated by a sphere of radius R centered in the point and can be normalized in order to obtain a scale invariant radial symmetry enhancement transform. This transform can therefore be used to detect and characterize salient regions like approximately spherical and approximately cylindrical surface parts and, being robust against holes and missing parts, it is suitable for real world applications e.g. anatomical features detection. Furthermore, its histograms can be effectively used to build a global shape descriptor that provides very good results in shape retrieval experiments.

Categories and Subject Descriptors (according to ACM CCS): I.2.10 [Computer Graphics]: Vision and Scene Understanding—Shape

1. Introduction

The ability of modern technologies like 3D laser or fringe-based scanners, dense stereo, TOF cameras to provide surface data of 3D objects, usually stored as triangulated meshes, creates the necessity of finding good methods to process this kind of data, capturing and compactly encoding the relevant information about the modelled objects. A lot of research work has therefore been performed on mesh seg-

mentation, skeletonization, shape recognition, salient points detection. The search for locally symmetric parts in a shape plays a relevant role in this context. Biological and architectural objects are, in fact, characterized by planar and radially symmetrical parts (e.g. spheres, cylinders) and their detection and description is therefore useful for a variety of shape analysis tasks. In this paper we present a new spatial transform called *multiscale area projection transform*

(MAPT) measuring the local likelihood of points in 3D space to be centers of radial symmetry at selected scales. The transform can be easily estimated on triangulated meshes and can be used not only to detect salient points and to reconstruct centerlines of tubular shape parts, but also to characterize globally the shape: its histograms are able to capture invariant features of rigid and articulated bodies and can be used for shape retrieval, comparing very well with state of the art techniques.

2. Related work

Several approaches have been proposed to detect and describe salient points or geometrical location characterizing 3D shapes. Most of them are designed to find surface features through curvature analysis at different scales and to describe them with feature vectors depending on the local neighborhood. A curvature-based notion of mesh saliency has been proposed in [LVJ05] and many papers presented related methods to detect and describe salient points on surfaces [SF06, CCFM08]. Shape descriptors can be purely local [JH99, MFK*10] or add global information to obtain a more reliable point characterization [KPNK03]. Local and global descriptors can then be used for shape retrieval applications with approaches like bag of features or similar ones [BBC*10, LGB*11]. Global and local information is intrinsically encoded in features derived by heat diffusion analysis, like the *heat kernel signature* [SOG09] that has shown effective performances for 3D shape matching [BBOG11, BBC*10]. The heat kernel approach has been applied also to volumetric analysis [RBBK10], obtaining signatures related to a geometrical parameter called scalar curvature, computed on a voxelized grid.

Another popular method to characterize salient locations of 3D shapes is using the 1D curve skeleton [CSM07], a subset of the medial axis [Blu67] that, thank to its 1D structure, captures well the tree-like structure of many real objects and has been used in shape recognition and matching tasks [SSGD03]. The curve skeleton, however, is not well defined where the shape is not approximately tubular and different algorithms used to compute it give different results and are often sensitive to noise and topological changes.

The use of symmetry to characterize locally and globally shapes is also a relevant research topic. In [TZCO09], rotational symmetry is applied for a robust skeletonization approach and also the *shape diameter function* proposed in [SSCO08] is in some sense related to radial symmetry. Solomon et al. [SBCBG11] employed discrete *approximate killing vector fields* (AKVFs) to discover intrinsic primitives, e.g. shape regions with approximate intrinsic symmetry. Lipman et al. [LCDF10] introduced the *symmetry factored embedding* (SFE) and the *symmetry factored distance* (SFD) to analyze and represent symmetries in point sets. Mitra et al [MGP06] proposed a nice framework to discover discrete symmetries at different scales through a clustering

procedure in transform space. Pauly et al. [PMW*08] presented an algorithm to search for repeated structures in complex shapes using generative models. The use of symmetry for shape matching has been proposed by Kazhdan et al. [KFR04], using for this task rotational symmetry around axes passing through the center of mass and by Podolak et al. [PSG*06] that introduced a *planar-reflective symmetry transform* capturing the degree of symmetry with respect to reflection through all planes in space.

The method proposed here is rather different from the cited ones, even if linked to some of them (it detects rotational symmetry, can be used to evaluate curve skeletons, can be used to create a symmetry based descriptor for shape matching). It is based on the computation of a spatial transform encoding the local degree of radial symmetry (in a selected scale range) and on the use of a saliency notion related to high symmetry. This saliency notion is often used in image analysis, being radial symmetry related with fixation points of the human visual system even better than high contrast [KS09]. It is not surprising, therefore, that approaches sharing the basic idea of our method have been proposed in image processing to detect salient points and circular objects. The one closer to our technique is the *fast radial symmetry* [LZ03], also creating a multiscale transform projecting signal perpendicularly from edges at a set of distances. Apart from the changed domain and dimensionality, our approach is, however, different for two important reasons:

- the transform definition is purely geometric and derives from an original theoretical formulation, so that it does not rely on particular discretizations or data structures. This formulation allowed also us link the APT to the medial axis transform (Sec.4.2);
- it computes a meaningful quantitative value, i.e. the amount of converging symmetric surface area scaled by the radius. This aspect makes the map useful not only for the detection of centers of symmetry, but also for recognition and matching, characterizing with peculiar quantitative values also non symmetric shapes.

The following sections will define the transform and discuss its properties and implementation issues.

3. Area Projection Transform

Our basic idea is to compute from input surfaces a spatial function that is maximal near centers of approximately spherical surfaces with a selected radius R and along centerlines of tubular structures with approximately circular section of radius R , so that it can be used to locate salient points and regions characterized by radial symmetry. This result can be obtained by projecting surface points along the local normal vector \vec{n} at distance R . The likelihood of a point \vec{x} to be center of a radial symmetry can then be assumed to be proportional to the number of points of the original surface S that are projected in a neighborhood of \vec{x} or, in other words, to the area of the original surface projected in that neighborhood. To compute this point density for a generic surface S in




Figure 2: Basic idea of the area projection transform: we compute the parallel surface at distance R (only in a selected direction or both directions) and we compute the transform in a point \vec{x} as the area of the original surface generating the part of the parallel one falling inside a sphere of radius σ centered in \vec{x} .

the 3D space, we consider the transform $T_R(S, \vec{n})$, creating a correspondence between S and its "parallel surface" [Win52] at the distance R on a selected side (moving each point of S along the normal direction \vec{n} , Fig.2). For each point \vec{x} in \mathbb{R}^3 , a sphere with radius σ centered in \vec{x} includes a subset of the surface $T_R(S, \vec{n})$, that is $k_\sigma(\vec{x}) \subset T_R(S, \vec{n})$; we define "area projection transform" $APT(\vec{x}, S, R, \sigma)$ the area of the subset of S having its parallel surface inside the sphere:

$$APT(\vec{x}, S, R, \sigma) = \text{Area}(T_R^{-1}(k_\sigma(\vec{x}) \subset T_R(S, \vec{n}))) \quad (1)$$

If we want to characterize both the sides of the surface we can compute also the contribution of the other parallel surface at distance R . i.e. $T_{-R}(S, \vec{n})$, possibly with opposite sign to separate the different kinds of features. T_R is well defined for continuous surfaces (properties of parallel surfaces can be found in [Win52]), and this is not always the case of data captured by scanners or reconstructed from volumetric grids. However we can easily compute the APT as the sum of the contributions of all the continuous patches acquired. If we can assign a common orientation to the patches we can compute the transform only for a selected orientation (e.g. inside the object), otherwise we can compute the APT as a sum of the contributions of the two sides.

It is easy to understand why this transform captures radial symmetry: a spherical surface with radius R projects all its surface on a single point and the APT is non negligible only within a distance σ from this point (see Fig. 3). A similar reasoning can be done for cylindrical surfaces. Radial symmetries will determine maximal regions for the APT and the value of σ determines the deviation from a perfect symmetry we want to tolerate when detecting it. Furthermore, exact values of the expected maxima of the APT in the case




Figure 3: Left: The internal parallel surface at distance R of a sphere S with radius R collapses in the central point. The maximal value of the APT computed exactly at scale R is the area of the entire spherical surface, i.e. $4\pi R^2$. Right: for a cylindrical surface of radius R the internal parallel surface at distance R collapses into a line. The APT computed at scale R has a maximal value along that line equal to the area of the green part of the cylinder, $4\pi R\sigma$.

of spheres or cylinders can be computed: if the surface S is a sphere with radius R and we compute the APT with the same R , the expected value near the center is the area of the spherical surface $4\pi R^2$, independently on the choice of σ (Fig. 3). If we make the value of σ tend to zero, we can consider

$$\rho_{AP}(\vec{x}, S, R) = \lim_{\sigma \rightarrow 0} \frac{APT(\vec{x}, S, R, \sigma)}{(4/3)\pi\sigma^3} \quad (2)$$

a density function in \mathbb{R}^3 measuring the amount of points of the original surface mapped in each infinitesimal region of the parallel surface. We can, in this way, reinterpret the APT as a kernel density estimate of ρ_{AP} with a unitary spherical kernel of size σ that could, in principle replaced by an arbitrary kernel in a generalized APT definition.

4. Handling multiple scales

When we search for radial symmetry we do not usually want to do it for a fixed R , but we want to find features at different scales. If we consider the radius R as a fourth variable in previous definitions, we obtain immediately multiscale versions of APT and ρ_{AP} . However, in order to compare the values at different radii and correctly characterize the scale of the dominating symmetries, a good idea is to define a multiscale area projection transform (MAPT) with the following two properties:

- i) a perfect sphere creates a signal with values in its center independent on its radius;
- ii) two surfaces differing for a scaling transform have the same maximum.

These properties are obtained with the following definition:

$$MAPT(x, y, z, R, S) = \alpha(R) APT(x, y, z, S, R, \sigma(R)) \quad (3)$$

where $\alpha(R) = 1/4\pi R^2$ and $\sigma(R) = c \cdot R$ ($0 < c < 1$).

The coefficient α makes, in fact, a perfect sphere create

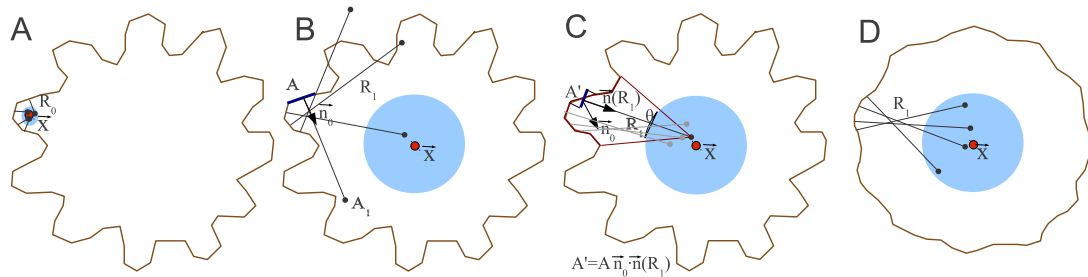


Figure 4: While the presence of symmetric surface parts with high curvature is detected by $APT(\bar{x}, R_0)$ with R_0 close to the radius of the locally tangent sphere (A), small structures may create noise in APT maps computed at larger radii (e.g. R_1), as many surface elements are projected far from the expected symmetry center (B). If, however, we compute the map at larger radii using a radially dependent smoothed normal $\bar{n}(R)$ (obtained, for example, averaging vectors in a region proportional to R^2), and then weight the projected area elements by the scalar product of local and averaged normal, the problem is removed and $APT(R)$ measures the surface of the sphere(circle) approximating the original data (C). Another simple method to reduce the problem consists of iteratively smoothing the surface at increasing R to remove structures with high curvature (D).

a central value equal to 1 independently on the radius. The choice for σ makes a shape scaled by a factor k produce a MAPT maximum at scale kR equal to the maximum of the original one at scale R . The resulting MAPT characterizes therefore the shape with values depending on the degree of radial symmetry independently on the scale. We know also that maximal value in center of perfect spheres should be equal to 1, while cylindrical shapes, following the reasoning of Fig. 3, should be characterized by maxima equal to $2\pi R \cdot 2cR / 4\pi R^2 = c$, still scale independent, but depending on the kernel size. A search for similar salient symmetric locations on input surfaces can be then performed on the 4-dimensional MAPT estimated for a reasonable range of radii, and can be also performed with a joint multiscale approach, combining the maps computed at different scales in a unique saliency function (joint multiscale APT). This map can be defined as:

$$JMPT(x, y, z) = \max_R(MAPT(R)) \quad (4)$$

and should still encode the main features related to the shape symmetry, even if it is clearly less informative than the 4D function. However, its use reduces the amount of memory required, that can be huge if the discretization is fine. Note that if we search for symmetric points on the JMPT, in order to be able to recover their scale, it is useful to create an additional matrix $SMPT(x, y, z)$ storing the related scale information, e.g. the radius corresponding to the maximum of $MAPT(x, y, z, R)$ at each spatial location.

4.1. Optimizing the behavior at different levels of detail

The above MAPT definition has a problem in the case we search for symmetry with a large radius and the surface presents a "texture" made of smaller structures. Consider the 2D sketch in Fig 4: small circular (spherical) structures are detected by projecting each segment length (triangle area) at distance R_0 along edge (face) normals (Fig 4 A), but the same procedure is less effective for $R = R_1$, even if the ker-

nel area is proportionally scaled (Fig 4 B), due to the fact that a lot of surface elements are projected along quite different normal vectors. However, it is possible to modify the MAPT definition in order to make it detect the symmetry of structures like that in Fig 4 as approximately spherical. An accurate solution consists of computing differently "smoothed normals" for sampled surface points at different values of radius, averaging local normals in a neighborhood with area proportional to the squared radius (Fig 4 C). We have just to consider the normal vector in the formula $T_R(S, n(\bar{R}))$ as a function of the radius R . To avoid then overestimation of the projected spherical part area we can then weight each projected surface patch with the scalar product between this smoothed normal and the local normal (Fig 4 C). A computationally faster solution for the discretized case (see Sec.5) consists of computing the transform sequentially at increasing radii and at each increment of the scale smoothing the surface (e.g. with a spatially controlled Laplacian operator) in order to remove high curvature detail, so that when we project at distance R_1 we don't use the original surface, but something like the smoothed one in Fig 4 D.

4.2. MAPT and the medial axis transform

Our transform is also related to the medial axis of a surface [Blu67], e.g. the locus of centers of spheres tangent to it in at least two distinct points, widely applied in geometrical processing. Actually the joint multiscale APT, defined in Sec.4, when the kernel radius tends to zero, vanishes with the exception of the points where there is an exact spherical symmetry of a surface part. These points are clearly a subset of the medial axis (centers of spheres tangent to the surface in an infinity of points), where the limit values of the function give further information related to the spherical parts area values. If we count the area elements projected into the kernel volume and consider the set of points where in the limit $\sigma \rightarrow 0$ the number is greater than 1, we can obtain the rest of the medial axis as a limit case from the JMPT. Be-




Figure 5: The medial axis of a cylinder includes also cones at the extrema (A) and when a small spherical perturbation is produced on it a complete line appears (C). This makes its computation on discretized surfaces not stable (B,D, computed with public code from [DS06]) The Joint MAPT is related to the medial axis (see text), but, weighting the area approximately tangent to spheres, it is not high in non-symmetrical regions near the bases of the cylinder (E) and, adding the small spherical perturbation, only the center of the sphere presents relevant values in addition to the axis.

ing possible to enrich the JMAPT with the associated scale, we can actually obtain the complete MAT from it.

In some sense we can consider the JMAPT an approximated and volumetric version of the medial axis transform. It is also a robust representation being not much changed by small perturbations of the surface (while the medial axis is completely changed by small bumps even far from them). This fact can be understood looking at Fig.5: the ideal medial axis estimated on a cylinder is a segment only far from the bases, where it is transformed in two conic surfaces (A). If a small perturbation is added, a new complete branch appears even if no tubular parts are present (C). Due to this low stability, algorithms for its numerical estimation on meshes provide noisy results (B,D). The ideal JMAPT for a cylinder should be maximal on the axis and decreasing around it with a behavior depending on the selected parameters. The JMAPT computed with our code on an approximately cylindrical mesh has the expected appearance (E). After the "spherical" perturbation, the JMAPT is relevant still on the axis in the center of the spherical bump, but is low elsewhere due to the negligible projected surface (F). The robustness provided by the local weighting of the symmetric surfaces is different from that provided by robust variants of the medial axis. For example the λ -medial axis [CL05] is obtained removing centers of spheres with radius lower than a threshold and the *scale axis transform* [MGP10] creates medial surfaces at different scales by extracting medial axes on shapes simplified by "multiplicative scaling" where radii of medial balls are changed. Using the APT maps we can as well select




Figure 6: 2D sketch showing the methods implemented to compute the the area projection transform with radius R . A: Points on the surface are sampled regularly and projected at distance R , with an associated original area contribution. If this point cloud is stored in a k -D tree, $APT(\vec{x}, R, \sigma)$ can be computed by querying for the points within a distance R from \vec{x} and summing the area contributions. B,C: a "counting" approach to speedup the estimation of APT on a regular grid: first each voxel is incremented of all the area contributions of the close projected points, then the obtained map is convolved with a spherical kernel with radius σ .

"medial representations" at different scales, but we have also the relevance given by the projecting surface area. Our volumetric density cannot replace medial surface representations having different nature and applications. However, the use of the MAPT as importance and scale-selective weight for medial axis branches could be investigated.

5. Discretization and implementation

In order to estimate APT for a polygonal mesh M , given a radius R and a tolerance σ we proceed as follows. First we sample points on each face and move the points of the value R along the normal direction. This can be done in a selected orientation (e.g. internal) if it is available, or both. We associate to each shifted point a contribution $A(F)/N$, where $A(F)$ is the area of the face and N is the number of sampled points. If the cloud of the shifted sampled points with associated area density is stored in a k -D tree structure it is possible to compute efficiently the APT at arbitrary location \vec{x} by retrieving all the points within a distance σ from \vec{x} and summing their density contributions.

For many applications we are interested in creating a discrete map sampling the transform in a regular three dimensional grid. Obviously we can do this by creating this grid containing the mesh of interest with the desired isotropic point spacing (voxel size) vs , and computing the APT values as described before at voxel centers. However, if we are interested only in computing the discrete map, we can skip the point cloud storage, directly accumulating the area density inside voxels. In this way the procedure is actually quite similar to a "voting" method to search for sphere/cylinder centers: each moved point increments (of $A(F)/N$) the values of the voxels in a neighborhood of its target location. In our implementation we add $A(F)/N$ for each point in the nearest voxel and finally perform the convolution of the ob-

tained map with a spherical kernel of radius σ . The resulting algorithm can be therefore summarized as follows:

- Create a matrix $APT(i,j,k)$ to store the discrete values of the transform sampled at positions $(i \cdot vs, j \cdot vs, k \cdot vs)$, initialized with all zeros.
- For each face F in M
 - Pick N sample points on it with N proportional to $A(F)$ and such as point spacing is close to vs
 - For each of these point \vec{p}_k , update the value of the voxel at distance R_i along the inward(outward) normal vector \vec{n} :

$$APT_0(\text{int}((\vec{p}_k + R_i \vec{n})/vs)) += A(F)/N$$
- Perform a convolution of the map with a spherical unitary kernel S of radius σ or an arbitrarily chosen kernel.

$$APT(M, R, \sigma) = S(\sigma) \cdot APT_0(M, R)$$

To compute a discrete multiscale $MAPT(i, j, k, R_i)$, it is then sufficient to compute the APT at different sampled radii R_i , with the corresponding $\sigma(R_i)$, proportional to them, and multiply the results by the corresponding scale factor $\alpha(R_i)$. The ratio σ_i/R_i is critical: if the size of the neighborhood where the area is integrated is small, we can have a precise localization of maxima, but only if the symmetry is perfect. If we want to detect approximate symmetry we should use sufficiently large values of the ratio, searching for a trade-off between detection and localization. In our experiments we used typically σ_i equal to $0.5R_i$. Another critical parameter is the sampling of R_i in the range considered: we usually took uniformly spaced values, with sampling step lower than the minimum spherical kernel radius applied, choice making in our test the transform able to detect spheres (and cylinders) of any radius in the range considered.

We implemented a simple C++ library with functions able to compute MAPT and JMAPT at selected scale ranges, both using the kD-tree approach or accumulating values on a grid and filtering with selectable kernel and size. Example codes are available at the web site <http://www.andreagiachetti.it/apt>. Simple volume renderings of the estimated APT/JMAPT maps, visualized together with semitransparent outer geometries show the ability of transform to enhance centers of radial symmetry at the selected scales (Fig. 1 A,B) or multiple scales (Fig. 1 C). The spatial behavior of the maps does not change if large parts of the surface are randomly removed (D) and, thanks to progressive normal smoothing, also to noise on vertices' location (E). These properties are still verified if MAPT is computed on a raw, non watertight mesh captured by a whole body scanner device (Fig. 7). Furthermore, capturing "approximate" symmetries, the map characterizes also non symmetric shapes creating peculiar structures at different scales, as shown in Fig. 8.

6. Applications and experimental results

Defined and implemented our transforms and analyzed their basic properties and critical parameters, we give some examples to show their practical useful applications.




Figure 7: The ability of MAPT to enhance centers of radially symmetrical parts is evident also in the case of noisy captured meshes with large holes, as in the case of this whole body human scan. Volume rendering of APT at three different scales (A,B,C) and of JMAPT (D) are shown.




Figure 8: APT at different scales creates peculiar structures also on flat and non spherical or tubular objects.

6.1. Salient points detection and characterization

Salient points related to high radial symmetry can be extracted as local maxima of APT (if we consider a specific radius) or of MAPT/JMAPT (if we want to search for all radial symmetries in a specific range). A simple salient point detector can be realized with the following procedure: computation of the MAPT with evenly sampled radii R_n (with a sampling step sufficiently dense compared to σ_0); Gaussian smoothing; selection of local maxima of the 4-dimensional function with intensity higher than a chosen threshold. Note that this threshold has a well defined meaning: we know that the expected value for a perfect sphere is 1, independently of the scale. We can choose the threshold as the value t corresponding to the fraction of an approximately spherical surface part that we want to detect in the specific application. Salient points can be characterized with vector descriptors, including, for example, intensity and scale. Further descriptor components can be, for example, blobness, flatness and vesselness of the map in a neighborhood of the point that can be estimated from the eigenvalues of the inertia matrix.

We implemented this salient points detector and tested it on different shapes. The program computes from the input surface a reference length s equal to a fixed fraction of the square root of the total surface area ($\sqrt{A_{tot}/80\pi}$). It then uses this value as minimum radius R_0 , as step between the tested radii and as voxel dimension in the MAPT computation. The kernel size used at the different scales is given by




Figure 9: Toy example showing the correct localization and characterization of salient points on spherical/cylindrical meshes. Detected points are represented as spheres with radius equal to half the detected scale and color encoding intensity (brightness) and sphericity (red component).




Figure 11: Salient points are reasonably extracted also in largely incomplete and noisy meshes as those here (large holes on visible sides highlighted in green), acquired with a whole body scanner on a subject in varying poses. Only maxima larger than a threshold ($t = 0.2$) are shown.

$\sigma_i = 0.5R_i$. The threshold used to consider the detected local maxima salient points is 0.2.

Fig 9 shows the results obtained on a simple synthetic mesh with three spheres of radii 0.25, 0.5 and 1 and a cylinder with radius 0.4. Detected salient points are represented as spheres with radius equal to half the detected one, brightness proportional to the intensity and hue reflecting the "sphericity" computed (red if nearly spherical, blue if elongated/flat). Maxima in the cylindrical part are detected (location and spacing depending on discretization effects), but their sphericity is low as expected.

On non perfectly symmetrical objects, salient points are still detected due to the tolerance given by σ and their sphericity values characterize regions around differently shaped surfaces. Figure 10 shows results obtained on meshes representing animals and humans (from the SHREC 2011 nonrigid 3D watertight data), where the extracted salient points correspond to relevant anatomical features that are detected almost independently on pose. Fig. 11 shows results obtained on three models of a girl acquired with a commer-

cial whole body scanner device. Salient points corresponding to major anthropometric locations are detected and are rather similar in different poses even if the quality of the models is low and large holes and noise are present. We tested also salient points extraction as maxima of the 3D JMAPT: results are rather similar to those obtained on the 4D map, only the characterization with the inertia tensor is less reliable due to the mixture of signals at different scales.

6.2. Lines extraction

Tubular parts of a shape should be characterized by a region of MAPT or JMAPT with high values, but also high vesselness with continuous orientation. Local vesselness and direction can be obtained in several ways analyzing local structure of the map or using area projection directional information. We plan to perform further investigations in order to define the best APT-based skeletonization algorithm, meanwhile we used a method commonly applied in medical images for vascular segmentation based on Frangi's hessian-based vesselness estimation [FNVV98]. Local Hessian at the selected scale is also able to provide vessels direction and lines can be obtained tracking lines of maximal vesselness in the direction of the local vessel. To simply test the feasibility of the centerline extraction based on APT, we implemented the following simple procedure to compute skeletal paths:

- compute the multiscale vesselness map from the JMAPT;
- segment connected regions of non negligible vesselness and continuous direction;
- find maxima of the map in each region and select points with the maximum distance from this point in opposite directions
- find the shortest path joining these points using fast marching with vesselness dependent speed.

This simple procedure provided visually reasonable results in tubular parts of the tested shapes. The algorithm does not provide connected paths as other methods (it is possible, however, to design specific heuristics to obtain such paths if needed). The method has, however, other interesting properties: it follows approximately tubular parts independently on local topology: in fact lines extracted do not correspond to centerlines of connected components, but to centerlines of cylindrical parts, so that, for example thumb skeleton is continued within the hand model in Fig 12 A. This can be a relevant advantage for applications, as shown in Fig. 12 B. Most skeletonization methods commonly used fail to follow tubular parts partially attached on one side (as required for correct anthropometric evaluation), as it happens, for example in the case of limbs in scanned human bodies. Lines extracted with publicly available implementations of methods based on medial geodesic function [DS06] and voxel coding [LCG09] do not follow the complete arm centerline, but the APT based method does. Furthermore, the method works on non-watertight surfaces and it is robust against the presence of relevant holes. We show this in the examples of Fig. 13 where many of the lines extracted are not changed even if a relevant percentage of the meshes triangles are removed.



Figure 10: Examples of salient points detected on models of animals and humans. They are represented as spheres with radius equal to half the detected scale and color encoding intensity (brightness) and sphericity (red component).




Figure 12: Lines extracted from the JMAPT are robust against topological noise. A: Skeletonization of a hand model. B: Lines extracted on a human body scan. Arms of the overweight woman are partially attached to the body. Curve skeleton algorithms often fail to follow the arms centerlines as in the case of publicly available codes based on medial geodesic function (blue) and voxel coding (green). Lines extracted with the APT method (red) follow the complete structures.




Figure 13: Lines extracted with the JMAPT method are robust against random triangle removal. Left: lines extracted on two models from the SHREC 2011 database. Right: lines extracted from the same models removing 50% of the original faces.

6.3. Shape retrieval with APT histograms

MAPT captures relevant properties of the shapes at different scales even in the case they are not spherical or tubular (Fig 8). This suggests the possible use of its histograms for model categorization and retrieval. The idea is to compute histograms of $MAPT(i, j, k, M, R_n, v_s, \sigma_n)$ at different scales

with a small number of bins and to concatenate them creating an unique feature vector. Of course, to characterize uniformly different shapes it is necessary to select reasonable scales for the analysis and, if we want to classify objects independently on their size, these scales should be varied with the object size.

Histograms are also expected to be approximately invariant to rigid transformations of the shapes and to present limited variations in the case of articulated motion, especially if we remove from the map voxels outside the analyzed object. We tested therefore the performances of this shape retrieval approach on the SHREC 2011 Nonrigid 3D Watertight Meshes database [LGB*11]. It is composed by 600 triangle meshes, 20 models for 30 categories, each one representing the same object undergoing an articulated deformation.

For all these models we computed a corresponding MAPT, using 8 different radii and applying the progressive smoothing method for increasing R . As in the salient point test we defined an adaptive procedure to compute the map parameters, selecting first a reference length step s proportional to the cube root of the volume. The minimum radius in the MAPT was then chosen as $2s$ and the step between tested radii s . The value of σ_i was chosen, for each scale R_i equal to $R_i/2$. We used the value s also as voxel size for discretization. Note that these choices make the retrieval method scale-independent. At each scale we computed the histogram of $MAPT(R_i)$ discretized in 12 bins equally spaced, and the concatenation of all of them gave a feature vector with 96 elements. These vectors were used to estimate the shape distance matrix required by the contest framework using the Jeffrey divergence [PBRT99], an information theory based distance measure. The SHREC'11 evaluation code was then used to compute, from the matrix, the five quality measures used in the contest: nearest neighbor (NN), first tier (FT), second tier (ST), e-measure (E) and discounted cumulative gain (DCG) [SMKF04]. Table 1 compares the values obtained using the HAPT based method with top results presented in the contest. The performances of the HAPT method resulted better than those of the con-

| | NN | 1-Tier | 2-Tier | e-Meas. | DCG |
|----------------|--------------|--------------|--------------|--------------|--------------|
| HAPT | 1.000 | 0.986 | 0.993 | 0.739 | 0.997 |
| SDGDM MeshSIFT | 1.000 | 0.972 | 0.990 | 0.736 | 0.996 |
| ShapeDNA | 0.992 | 0.915 | 0.957 | 0.705 | 0.978 |
| MDS-CM-BOF | 0.995 | 0.913 | 0.969 | 0.717 | 0.982 |
| MLSF | 0.987 | 0.809 | 0.879 | 0.643 | 0.948 |
| FOG+MRR | 0.960 | 0.881 | 0.946 | 0.696 | 0.959 |
| HKS | 0.837 | 0.406 | 0.497 | 0.353 | 0.730 |

Table 1: Our method based on histograms of the area projection transform (HAPT) performs very well for shape retrieval even if compared with best results in the SHREC'11 nonrigid watertight contest. Note that our method is robust against holes, topological noise and is scale invariant.

test winner [LGB*11]. Different parameters' choices could result in different scores for all the methods, so it does not mean that our method is always the best one, but certainly it is suitable for shape retrieval applications and it is sufficiently fast (our CPU implementation took about 10 sec. to compute each histogram). It is worth noting that we obtained perfect results for 26 of the 30 classes and clearly better than any other method for 28 classes. Not surprisingly object classes where the HAPT method resulted worse than other ones are ants and spiders, probably due to similar scales of tubular parts and similarity of articulation pattern obscuring shape similarity. Mostly flat objects (e.g. laptops, birds) are well recognized (Fig. 15).

Really appealing features of our method are the fact that it is perfectly scale-invariant and robust against topological changes and holes, at least if relatively small (they change signal intensity). Adding adaptive normal smoothing proportional to R the method is also rather robust against noise on points and normals. To test the robustness of the HAPT descriptor for shape retrieval we created three modified datasets from SHREC models: the first by removing approximately 10% of the original triangles from half of the original meshes; the second adding to the vertices of half of the original meshes a random displacement proportional to the local triangle size; the last using one third of original meshes, one third with holes and one third with random displacements. On these dataset we repeated the same classification procedure and the NN classification accuracy values were still good: 98.3%, 99,8% and 94.0% respectively.

7. Discussion

We presented a novel method for detecting radial symmetries of 3D shapes by computing a spatial function called multiscale area projection transform, defined for generic surfaces and implemented for triangulated meshes. We have shown its possible use for relevant tasks such as salient points detection and centerline tracking. The results obtained in a complex shape retrieval task with our method based on MAPT histograms revealed its ability to capture both global and local object features and are particularly remarkable




Figure 15: The precision vs recall plot for the whole set of models is rather close to the ideal one, and, if done for individual classes it is exactly the ideal one for 26 of the 30 shape labels. Only for two classes (ants and spiders) the result is worse than that obtained with other techniques.

considering that method is scale invariant, can be applied also to non watertight meshes, it is robust against topological noise and could be used as well on dense point clouds instead of meshes, projecting points along associated normals and assuming that their density is approximately constant.

Acknowledgements: thanks to Emanuele Trucco and to anonymous reviewers for useful suggestions.

References

- [BBC*10] BRONSTEIN A. M., BRONSTEIN M. M., CASTELLANI U., FALCIDIENO B., FUSIELLO A., GODIL A., GUIBAS L. J., KOKKINOS I., LIAN Z., OVSIANIKOV M., PATANE G., SPAGNUOLO M., TOLDO R.: Shrec 2010: robust large-scale shape retrieval benchmark. In *Proc. EUROGRAPHICS Workshop on 3D Object Retrieval* (2010). 2
- [BBOG11] BRONSTEIN A. M., BRONSTEIN M. M., OVSIANIKOV M., GUIBAS L. J.: Shape Google: geometric words and expressions for invariant shape retrieval. *ACM Transaction on Graphics* 30, 1 (Feb. 2011), 1:1–1:20. 2
- [Blu67] BLUM H.: A Transformation for Extracting New Descriptors of Shape. In *Models for the Perception of Speech and Visual Form*, Wathen-Dunn W., (Ed.). MIT Press, Cambridge, 1967, pp. 362–380. 2, 4
- [CCFM08] CASTELLANI U., CRISTANI M., FANTONI S., MURINO V.: Sparse points matching by combining 3D mesh saliency with statistical descriptors. *Computer Graphics Forum* 27 (2008), 643–652. 2
- [CL05] CHAZAL F., LIEUTIER A.: The " λ -medial axis". *Graph. Models* 67, 4 (July 2005), 304–331. 5
- [CSM07] CORNEA N. D., SILVER D., MIN P.: Curve-skeleton properties, applications, and algorithms. *IEEE Transactions on Visualization and Computer Graphics* 13, 3 (2007), 530–548. 2
- [DS06] DEY T. K., SUN J.: Defining and computing curve-skeletons with medial geodesic function. In *SGP '06: Proceedings of the fourth Eurographics symposium on Geometry processing* (Aire-la-Ville, Switzerland, 2006), Eurographics Association, pp. 143–152. 5, 7
- [FNVV98] FRANGI A., NIESSEN W., VINCKEN K., VIERGEVER M.: Multiscale vessel enhancement filtering. In *Medical Image Computing and Computer-Assisted Intervention-MICCAI'98*, Wells W., Colchester A., Delp S.,

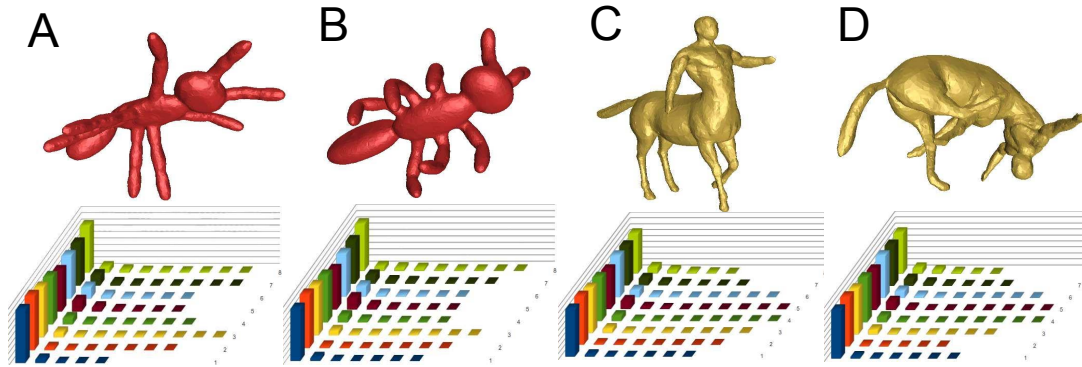


Figure 14: Histograms area projection transform computed inside objects are rather similar for the same after an articulated deformation (A-B,C-D) and rather different for shapes of different nature.

- (Eds.), vol. 1496 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 1998, pp. 130–137. [7](#)
- [JH99] JOHNSON A., HEBERT M.: Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21, 1 (May 1999), 433–449. [2](#)
- [KFR04] KAZHDAN M., FUNKHOUSER T., RUSINKIEWICZ S.: Symmetry descriptors and 3D shape matching. In *Eurographics/ACM SIGGRAPH Symposium on Geometry processing* (New York, NY, USA, 2004), ACM, pp. 115–123. [2](#)
- [KPNK03] KÖRTGEN M., PARK G., NOVOTNI M., KLEIN R.: 3d shape matching with 3d shape contexts. In *The 7th Central European Seminar on Computer Graphics* (Apr. 2003). [2](#)
- [KS09] KOOTSTRA G., SCHOMAKER L.: Prediction of human eye fixations using symmetry. In *N. A. Taatgen and H. van Rijn (Eds.), Proc. CogSci09, July 29 - Aug 1, Amsterdam, The Netherlands*. (2009). [2](#)
- [LCDF10] LIPMAN Y., CHEN X., DAUBECHIES I., FUNKHOUSER T.: Symmetry factored embedding and distance. *ACM Trans. Graph.* 29 (July 2010), 103:1–12. [2](#)
- [LCG09] LOVATO C., CASTELLANI U., GIACHETTI A.: Automatic segmentation of scanned human body using curve skeleton analysis. In *Proceedings of Mirage* (2009), LNCS, pp. 34–45. [7](#)
- [LGB*11] LIAN Z., GODIL A., BUSTOS B., DAOUDI M., HERMANS J., KAWAMURA S., KURITA Y., LAVOUÉ G., NGUYEN H., OHBUCHI R., OHKITA Y., OHISHI Y., PORIKLI F., REUTER M., SPIRAN I., SMEETS D., SUTENS P., TABIA H., VANDERMEULEN D.: Shrec'11 track: Shape retrieval on non-rigid 3d watertight meshes. In *Proc. Eurographics/ACM SIGGRAPH Symposium on 3D Object Retrieval* (2011), pp. 79–88. [2, 8, 9](#)
- [LVJ05] LEE C., VARSHNEY A., JACOBS D.: Mesh saliency. *ACM Trans. Graph.* 24 (July 2005), 659–666. [2](#)
- [LZ03] LOY G., ZELINSKY A.: Fast radial symmetry for detecting points of interest. *IEEE Trans. Pattern Anal. Mach. Intell.* 25 (August 2003), 959–973. [2](#)
- [MFK*10] MAES C., FABRY T., KEUSTERMANS J., SMEETS D., SUTENS P., VANDERMEULEN D.: Feature detection on 3d face surfaces for pose normalisation and recognition. In *BTAS '10, IEEE Int. Conf. on Biometrics: Theory, Applications and Systems* (2010), pp. 1–6. [2](#)
- [MGP06] MITRA N. J., GUIBAS L., PAULY M.: Partial and approximate symmetry detection for 3d geometry. *ACM Transactions on Graphics (SIGGRAPH)* 25, 3 (2006), 560–568. [2](#)
- [MGP10] MIKLOS B., GIESEN J., PAULY M.: Discrete scale axis representations for 3d geometry. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 101. [5](#)
- [PBRT99] PUZICHA J., BUHMANN J., RUBNER Y., TOMASI C.: Empirical evaluation of dissimilarity measures for color and texture. In *Proceedings of the Seventh IEEE International Conference on Computer Vision* (1999), vol. 2, pp. 1165–1172. [8](#)
- [PMW*08] PAULY M., MITRA N. J., WALLNER J., POTTMANN H., GUIBAS L.: Discovering structural regularity in 3d geometry. *ACM Trans. Graph.* 27, 3 (2008), 1–11. [2](#)
- [PSG*06] PODOLAK J., SHILANE P., GOLOVINSKIY A., RUSINKIEWICZ S., FUNKHOUSER T.: A planar-reflective symmetry transform for 3d shapes. *ACM Transactions on Graphics* 25, 3 (July 2006). [2](#)
- [RBBK10] RAVIV D., BRONSTEIN M., BRONSTEIN A., KIMMEL R.: Volumetric heat kernel signatures. In *Proceedings of the ACM workshop on 3D object retrieval* (New York, NY, USA, 2010), 3DOR '10, ACM, pp. 39–44. [2](#)
- [SBCBG11] SOLOMON J., BEN-CHEN M., BUTSCHER A., GUIBAS L.: Discovery of intrinsic primitives on triangle meshes. *Computer Graphics Forum* 30, 2 (2011), 365–374. [2](#)
- [SF06] SHILANE P., FUNKHOUSER T.: Selecting distinctive 3D shape descriptors for similarity retrieval. In *International Conference on Shape Modelling and Applications* (2006), IEEE Computer Society, p. 18. [2](#)
- [SMKF04] SHILANE P., MIN P., KAZHDAN M., FUNKHOUSER T.: The princeton shape benchmark. In *Shape Modeling International* (2004), vol. 105, p. 179. [8](#)
- [SOG09] SUN J., OVSIANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. In *Proceedings of the Symposium on Geometry Processing* (2009), pp. 1383–1392. [2](#)
- [SSCO08] SHAPIRA L., SHAMIR A., COHEN-OR D.: Consistent mesh partitioning and skeletonisation using the shape diameter function. *The Visual Computer* 24 (2008), 249–259. 10.1007/s00371-007-0197-5. [2](#)
- [SSGD03] SUNDAR H., SILVER D., GAGVANI N., DICKINSON S.: Skeleton based shape matching and retrieval. In *Proc. Shape Modeling International* (Washington, DC, USA, 2003), IEEE Computer Society, p. 130. [2](#)
- [TZCO09] TAGLIASACCHI A., ZHANG H., COHEN-OR D.: Curve skeleton extraction from incomplete point cloud. *ACM Trans. Graph.* 28 (July 2009), 71:1–71:9. [2](#)
- [Win52] WINTNER A.: On parallel surfaces. *American Journal of Mathematics* 74, 2 (1952), 365–376. [3](#)