Analysis of Partial Axial Symmetry on 3D Surfaces and its Application in the Restoration of Cultural Heritage Objects

Ivan Sipiran
Grupo de Reconocimiento de Patrones e Inteligencia Artificial Aplicada, GRPIAA
Departamento de Ingeniería, Pontificia Universidad Católica del Perú, PUCP
isipiran@pucp.edu.pe

Abstract

Symmetry is a ubiquitous concept that can help to understand the structure of real objects. One of the main challenging problems in the analysis of symmetry is the robustness against high partiality, i.e., when the support of the symmetry in the input geometry is small. In this paper, we address the problem of finding the partial axial symmetry of 3D objects through the analysis of surface descriptors with invariance to partiality. These descriptors are used to reduce the search space of axial symmetric correspondences, and allows us to design an effective and efficient algorithm to detect the generator axis of the symmetry. Our algorithm collects enough evidence of the presence of the axial symmetry in a consensus-based approach. Our algorithm can also identify the support of the axial symmetry. Our experiments show the robustness of our method in challenging scenarios. We show that our method is good to generate plausible restorations of damaged cultural heritage objects.

1. Introduction

Most objects around us exhibit some kind of symmetry. It can be attributed to many reasons, but it is mainly due to the fact that symmetry is an underlying characteristic that provides structure. This structure can in turn be used to understand high-level information about an object such as functionality, the process of construction, or even the beauty on it. Since one of the ultimate goal of computer vision is to let computer understand their environment through a visual stimulus, the symmetry analysis seems to be a helpful task to support other high-level tasks such as reconstruction or recognition.

We are interested in the analysis of axial symmetry. This interest arises from the concern of analyzing archaeological objects. We have noticed a predisposition of people in ancient cultures to make objects of daily use that show axial symmetry. The shape is due to the use of a pottery wheel in the construction of the objects. In general, the objects show a axially symmetric structure (the main body of the object) plus additional features (handles and ornaments). Unfortunately, many objects have been buried for a long time, which has caused severe damage to the pieces. Therefore, the problem consists of revealing the structure of objects with (possibly large) missing geometry. We call this problem *analysis of partial axial symmetry*.

The analysis of symmetry in 3D surfaces has attracted recent attention of computer vision and computer graphics communities. For readers interested in the topic of symmetry analysis, we refer to the comprehensive survey presented by Mitra et al. [14]. According to the amount of information we use to define the symmetry, we can find global symmetries and partial symmetries. The former is related to transformations that apply to the entire object, while the latter is related to transformations that apply only to a part of the object. There are many studies that focus the attention in global symmetries [5, 20, 10, 15, 22, 8, 6]. These methods assume that the object is complete and the center of mass of the object is a fixed point of the symmetric transformation. Likewise, many methods have been proposed to deal with partial symmetries [13, 3, 24, 9, 23]. In general, these methods try to match distinguishable local features in order to find the transformations between parts.

On the other hand, object repair is a challenging topic that has captured attention recently because of the increasing popularity of 3D scanning devices. If the object exhibits some sort of symmetry it is also possible to use this fact to repair the object. In a recent article, Sipiran et al. [17] used a matching algorithm to find symmetric correspondences. They proposed a variation of the Heat Kernel Signature [19] which preserves symmetry even in the presence of missing parts. Similarly, Mavridis et al. [11] proposed a registration-based technique which is able to repair 3D objects by finding general symmetries using an optimization approach. Also, Son et al. [18] presented a method to reconstruct a 3D pottery from fragments by using information

of the symmetry axis and the profile curve of the pottery.

In this paper, we present a method to tackle the detection the partial axial symmetry in 3D shapes. Our method uses the theory of heat diffusion over 3D surfaces to find evidence of regularity and structure in the input shape. In particular, we propose an algorithm that analyzes the invariance of a function defined over the surface of a partial object, to subsequently detect its circular structure with a good approximation. We also propose a method to determine the part of the geometry that is affected by the axial symmetry, disregarding additional features that are not generated by the symmetry. Finally, the detected symmetry is used to synthesize missing geometry of input objects.

The contribution of our paper is three-fold. First, we present a detailed formulation of the problem. It helps us to understand the problem and find the elements that leads to the solution. Second, we propose a robust method to compute the approximate symmetry axis of objects with missing geometry. Finally, we show the application of the proposed method to the robust generation of plausible geometry to repair cultural heritage objects.

The paper is organized as follows. Section 2 presents a formal description of the problem. Section 3 describes our method to detect the partial axial symmetry. Section 4 shows our experiments and results. Section 5 presents the application of our method to the virtual repair of cultural heritage objects. Finally, Section 6 concludes our work.

2. Problem formulation

Let $\mathcal{M}$ be a 2D manifold embedded in a 3D space. By definition, a symmetry is a non-trivial transformation T such that the transformation leave objects unchanged. Formally, T is a symmetry if $T(\mathcal{M}) = \mathcal{M}$. A relaxed version of this definition permits a degree of error in the transformation and introduces the idea of approximate symmetry. Let d be a function that measures the congruence of two manifolds. A transformation T is said to be α-approximate if $d(T(\mathcal{M}), \mathcal{M}) \leq \alpha$.

In the case of axial symmetries, the transformation T is any rotation around the *generator axis* of the object. Without loss of generality, we use the term *generator axis* instead of symmetry axis because, as we are dealing with the problem of partial symmetry, we want to denote the axis of the object as if it would be complete. With this in mind, we denote the transformations as $T = R_{\mathcal{E},\theta}$, $0 \leq \theta \leq 2\pi$, the rotation around the generator axis $\mathcal{E}$ with angle θ. The generator axis is defined as a tuple $\mathcal{E} = (\vec{N}, C)$, where $\vec{N}$ is a normal vector and C is a point that lies in $\mathcal{E}$. Therefore, if a manifold $\mathcal{M}$ exhibits axial symmetry then $\mathcal{M} = R_{\mathcal{E},\theta}(\mathcal{M})$, for any angle θ. The approximate counterpart can be formulated similarly.

The aforementioned definition works perfectly when the expected symmetry is global. That is, when the complete object is transformed and it remains unchanged. However, there are two cases in which the definitions do not hold anymore. The first case is when only a partial region of the object is invariant to the transformation $R_{\mathcal{E},\theta}$. This situation can be observed in many real objects that contain an axial symmetry in some of their sub-components. The second case is when an object with global axial symmetry is broken. This situation can be observed in cultural heritage applications for example.

Hence, the *partial axial symmetry* can be defined as follows. A manifold $\mathcal{M}$ exhibits partial axial symmetry if there exists a generator axis $\mathcal{E}$ and two regions $\mathcal{M}_1, \mathcal{M}_2 \subseteq \mathcal{M}$ such that $R_{\mathcal{E},\theta}(\mathcal{M}_1) = \mathcal{M}_2$. The problem of finding the axial symmetry $R_{\mathcal{E},\theta}$ can be interpreted as the problem of finding the generator axis $\mathcal{E}$. Nevertheless, computing the generator axis is a complex task itself due to two reasons. First, we do not know regions $\mathcal{M}_1$ and $\mathcal{M}_2$ in advance. A procedure to search two regions that hold with the definition can be a time-consuming task. Second, the generator axis do not necessarily goes through the center of mass of the input object. This makes the problem even harder since the search space needs to consider orientation and position as well.

2.1. Theory of heat diffusion

Heat diffusion studies the behavior of the heat propagation over the surface of an object as a function of the elapsed time. The analysis of the propagation is closely related to the geometrical characteristics of objects, and therefore, the analysis is helpful to understand 3D shapes and their composition. In general, the analysis of heat diffusion on a manifold $\mathcal{M}$ consists of defining a vector-valued function $f : \mathcal{M} \rightarrow \mathbb{R}^n$ over the surface of 3D objects such that the function must be preserved under isometric transformations. That is, if T is a isometry between two Riemannian manifolds $\mathcal{M}$ and $\mathcal{P}$, then it holds that $f_{\mathcal{M}}(x) = f_{\mathcal{P}}(T(x))$. If function f holds with the previous property, then we say that f is intrinsic.

This definition can be easily extended to symmetries. Formally, a symmetry is a self-isometry that maps points on a 3D shape to itself. In our problem, $T = R_{\mathcal{E},\theta}$ and, by definition, for any intrinsic surface function f, the equality $f_{\mathcal{M}}(x) = f_{\mathcal{M}}(R_{\mathcal{E},\theta}(x))$ should hold. Moreover, given a point $x \in \mathcal{M}$, there are an infinite number of points $R_{\mathcal{E},\theta}(x) \in \mathcal{M}$ such that the surface function is preserved (for any angle θ). As consequence, we plan to use the analysis of heat diffusion in the detection of partial axial symmetries. To accomplish that, we must take care of the design of an effective method to find the transformation $R_{\mathcal{E},\theta}$ from the evidence delivered by the surface function.

3. Symmetry detection

The core of our approach is the detection of the axial symmetry in objects with missing geometry. Figure 1 depicts an overview of our method to detect the partial axial symmetry of an input object. The first step of our method analyzes the surface of the input object to find evidence of circular structures. Instead of performing an analysis on the complete input mesh, we show that points sampled with a good distribution is a good choice (see Fig. 1(a)). Thereby, our method finds circles formed by similar points in the space of a surface function (see Fig. 1(b)). The second step of our method determines the dominant circular structure through a cluster analysis performed on the circles detected in the previous step (see Fig. 1(c)). Finally, our algorithm can be applied to replicate the existing object and generate missing geometry (see Fig. 1(d)).

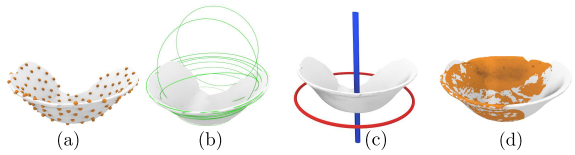

Figure 1: Overview of our method. (a) The analysis is performed over a set of well-distributed sampled points on the surface of the object. (b) Our method computes a circle for each sampled point according to the similarity in the surface function space. (c) A clustering algorithm determines the best set of circles from which we can obtain the generator axis. The circle in the figure is the average circle in the best set of circles. (d) The generator axis can finally be used to repair damaged cultural heritage objects. Best viewed in color.

3.1. Finding supporting circles

The problem can be summarized as: given a triangle mesh M, which exhibits partial axial symmetry, we want to find the transformation $R_{\mathcal{E},\theta}$. Note that it is only necessary to compute the generator axis $\mathcal{E}$ since every rotation is obtained from the generator axis and a given angle θ.

The starting point to solve the problem is the observation that a intrinsic function f_M, defined over the surface of the object M, is invariant to the symmetry transformation $R_{\mathcal{E},\theta}$. More specifically, if there exist two points $x \in M$ and $y \in M$ and $y = R_{\mathcal{E},\theta}(x)$ for some θ, then it is likely that $\|f_M(x) - f_M(y)\|_2 = 0$. Therefore, we can use this criterion to search points in the feature space of the surface function. All the points with the same values of function f_M have to be generated around the generator axis $\mathcal{E}$, so we can use the points to recover the structure of the symmetry.

Nevertheless, the previous criterion is too strict and it is useless in practice since the input shape is a discrete representation of a continuous manifold. We thus need to relax the criterion to search points with similar values in the function f_M instead of looking for exact coincidence of the function. Additionally, recall that the symmetry of the input object is only partial, so it is very likely that points on the surface do not hold the diffusion similarity criterion at all. To face this problem, our approach is to get as much relevant information as possible to recover the generator axis in a reliable way.

Our algorithm starts by computing the function f_M for every point on the surface of the input 3D shape M. To obtain as much information as possible on the existence of the axial symmetry, we take into account a set of points to perform the analysis of similarity with respect to the function f_M. A good set of points for the analysis is such that it covers the entire geometry, giving the opportunity to the algorithm to choose point that are axially symmetric. Let $P = \{p_i\}_{i=1}^{n}$ be a set of sampled points on the surface of M, obtained with a farthest point sampling (FPS) method. We chose this sampling algorithm because the resulting points are well distributed in the surface and it is simple to implement.

For every point $p_i \in P$, our algorithm searches for similar points with respect to the function f_M. In practice, we use the k-nearest neighbors associated to each point p_i. That is, the search obtains a set of points $Q_i = \{q_j \in M\}_{j=1}^{m}$ such that $\forall k \in M, k \notin Q_i$, it holds that $\|f_M(p_i) - f_M(q_j)\| < \|f_M(p_i) - f_M(k)\|$. The number of points in Q_i with a very similar value of f_M will be high if point p_i belongs to the region of the object with partial axial symmetry. In addition, the set of points Q_i should be arranged around the generator axis, forming a circle. However in practice, points are not located in a perfect circle, so we have to find the circular structure from Q_i.

Obviously, Q_i could contain outliers, i.e. points lying far from the circular structure defined by p_i and the axial symmetric correspondences. Likewise, if the input object is not complete, Q_i only contains the structure of a circular segment. To tackle these problems, we apply the RANSAC method to find the best circle that approximates the structure of points in Q_i. Briefly, we take three random points q_a, q_b and q_c in Q_i and compute the circle going through the three points. The center c of the circle is the intersection of the perpendicular lines that pass through the middle point of segments $\overline{q_a q_b}$ and $\overline{q_b q_c}$. The radius of the circle is $r = \|q_a - c\|_2$. Finally, the orientation of the 3D circle is determined by the vector $\overrightarrow{N} = \overrightarrow{q_a q_b} \times \overrightarrow{q_b q_c}$. Note that a 3D circle can now be defined as the tuple $(C, r, \overrightarrow{N})$.

Once we compute the circle defined by the three random points, we need to verify if the circle is a good candidate. For all of the remaining points in Q_i, we compute the distance to the current circle. If the distance is below a given threshold, we count a vote for the circle. The goal

of RANSAC is to find the circle with the highest number of votes. As consequence, the previous process has to be executed a number of times. Finally, the best circle is selected as evidence of the presence of the generator axis. In all our experiments, the RANSAC threshold is set to 0.5% of the diagonal of the input object and the number of random circles evaluated is 500. Furthermore, as result, we obtain a set of circles $\mathcal{C} = \{(c_i, r_i, \overrightarrow{N_i})\}_{i=1}^{n}$, one circle for each sampled point in the analysis.

3.2. Finding the generator axis $\mathcal{E}$

The set of circles $\mathcal{C}$ could contain spurious elements that do not belong to circles arranged around the generator axis. We thus need to identify the subset of circles that approximates the real axis. We present a two-step method to detect the circles around a candidate generator axis.

Clustering by orientation

In this stage, we only group the set of circles according to their orientation. Our approach focuses on finding groups of circles with similar un-oriented normals. Let $\mathcal{C}_i = (c_i, r_i, \overrightarrow{N_i})$ and $\mathcal{C}_j = (c_j, r_j, \overrightarrow{N_j})$ two circles, we define the angular distance between $\mathcal{C}_i$ and $\mathcal{C}_j$ as

$$d_{\measuredangle}(\mathcal{C}_i, \mathcal{C}_j) = 1 - \frac{|\overrightarrow{N_i} \bullet \overrightarrow{N_j}|}{\|\overrightarrow{N_i}\|\|\overrightarrow{N_j}\|} \tag{1}$$

where $\bullet$ is the dot product between two vectors. We take the absolute value of the dot product because we are only interested in the orientation and not in the direction itself. Note that the minimum angular distance is zero and it occurs when $\overrightarrow{N_i}$ is parallel to $\overrightarrow{N_j}$ or when $\overrightarrow{N_i}$ is parallel to $-\overrightarrow{N_j}$. Otherwise, the maximum angular distance is one and it occurs when $\overrightarrow{N_i}$ and $\overrightarrow{N_j}$ are orthogonal.

The previous distance is useful to find groups of normals with similar orientation. Nevertheless, we do not know how much groups there are in advance, so we propose an alternative clustering method that automatically finds good partitions. Our method is the medoid-based adaptation of a histogram-based clustering proposed in [7]. The algorithm is controlled by a few number of parameters: the maximum spread of a cluster λ, the minimum distance between clusters β and the minimum number of elements for a group to be considered a cluster K. The algorithm 1 details the steps to perform the clustering. The input to the algorithm is the pair-wise distance matrix of the input elements D. Briefly, the algorithm iterates over the set of elements, deciding in each step whether a element must form a new cluster or it must be assigned to an existing cluster. To perform the clustering by orientation, the input elements for the clustering are the circles in the set $\mathcal{C}$ and the distance used is $d_{\measuredangle}$. The output is a set of clusters, where the circles within a cluster share their orientation.

Algorithm 1 Clustering algorithm

procedure CLUSTERING(D, P, α, β, K)
  $Clusters = \emptyset$
  while $Clusters$ does not converge **do**
    for each element $p_i \in P$ **do**
      if $|Clusters| = 0$ **then**
        $min\_dist = \infty$
      else
        Find nearest cluster to p_i
        Let min_dist be the distance to the nearest medoid
      end if
      if $min\_dist \geq \beta$ **then**
        Create new cluster and add p_i
        Assign p_i as the medoid of new cluster
      else if $min\_dist \leq \alpha$ **then**
        Assign p_i to nearest cluster
        Update medoid of nearest cluster
      end if
    end for
    Keep clusters with size greater than K
  end while
  return $Clusters$
end procedure

Identification of generator axis

A cluster of circles with coherent orientation could contain elements that do not correspond to the generator axis. This can happen due to bad localizations of the matches in the surface function space. We thus need to discard the circles that do not belong to a good candidate axis. We propose to use the previously defined clustering method, but with a different distance measure. Given two circles $\mathcal{C}_i = (c_i, r_i, \overrightarrow{N_i})$ and $\mathcal{C}_j = (c_j, r_j, \overrightarrow{N_j})$, we define the axial distance

$$d_{\mathcal{E}}(\mathcal{C}_i, \mathcal{C}_j) = \|\overrightarrow{N_i} \times \overrightarrow{(c_i - c_j)}\| + \|\overrightarrow{N_j} \times \overrightarrow{(c_i - c_j)}\| \tag{2}$$

If the centers of the circles $\mathcal{C}_i$ and $\mathcal{C}_j$ are collinear to a generator axis, then the vector formed by both centers $\overrightarrow{(c_i - c_j)}$ will have a similar orientation to the circle normals $\overrightarrow{N_i}$ and $\overrightarrow{N_j}$, and therefore the distance will be small. Otherwise, if the centers are not collinear to a generator axis, then the distance is dominated by the magnitude of vector $\overrightarrow{(c_i - c_j)}$. Another way to see the behavior of this distance is using the equivalent definition of the magnitude of vector product

$$\|\overrightarrow{N_i} \times \overrightarrow{(c_i - c_j)}\| = \|\overrightarrow{N_i}\|\|\overrightarrow{(c_i - c_j)}\| \sin\theta \qquad (3)$$

where θ is the angle formed by the two vectors in the product. Note that the centers of circles that belong to a generator axis have to be aligned to the normals, and thereby θ is small. In our implementation, this distance is used to identify groups of circles aligned around candidate generator axes. As result, our algorithm computes groups of circles with similar orientation and arranged around a common axis. We consider the group with the largest number of circles as the supporting set to compute the generator axis of the symmetry. The last step is to compute the approximate generator axis which is defined by a normal vector and a point lying in the axis. We take the simple average of the normals and the centers of every circle in the final group as generator axis.

3.3. Estimation of symmetric support

It can happen that only a part of the complete geometry is affected by the axial symmetry. In fact, our algorithm is able to find a good symmetric axis without explicitly computing the supporting geometry for the symmetry. Here we show a simple and fast procedure to estimate the region that is affected by the axial symmetry.

The first step is to transform the input shape such that the generator axis coincides with the Y axis of the 3D coordinate system. It can be easily performed using the information of the normal found in the previous stage. The goal of the transformation is to facilitate the search of symmetric correspondences around the axis. The second step is to sort the points of the input shape in ascending order according to the Y-coordinate of each point. Note that each point $p_i = (x_i, y_i, z_i)$ in the shape defines an implicit circle around the Y axis, where the circle's radius is $\sqrt{x_i^2 + z_i^2}$. The third step is to identify the amount of geometry that is close to the circle defined by the point p_i. Our algorithm takes advantage of the sorted set of points to reduce the number of distance computations during the verification of the surface.

We use a threshold λ to denote the maximum permitted distance for a point to be included in the supporting set for a given circle. A first set of candidate points is computed by selecting the points on the surface with Y-coordinates in the range $[y_i - \lambda, y_i + \lambda]$. Note that since the points are sorted with respect to their Y-coordinate, the search of points in the given range can be performed in $O(\log n)$ using a binary search. The resulting set of points from the previous search is then used to compute the real distance to the circle. Points with a distance below the threshold λ are finally assigned to the circle. An easy way of computing the support for an analyzed point p_i is to assign the number of supporting points to it. A high number of supporting points means that the point p_i belongs to a axially symmetric region of the surface.

4. Experiments and Results

In this section, we present the experiments conducted in our research. Section 4.2 is dedicated to evaluate the best choice of surface function in our approach, and Section 4.3 is devoted to show the results of detecting the partial axial symmetries in real scanned 3D objects.

4.1. Experimental Setting

Here we describe the setting used for all the experiments in the paper. We use the cotangent scheme [12] to approximate the Laplace-Beltrami operator of a triangular mesh. We compute 300 eigenvalues and eigenvectors. The number of nearest neighbors per each sampled point in Section 3.1 is set to 200. The clustering parameters for the orientation estimation are $\alpha = 0.005$, $\beta = 0.01$ and $K = 10$. The clustering parameters for the estimation of the axis are $\alpha = 0.2$, $\beta = 0.4$ and $K = 5$.

4.2. The search of a robust surface function

Our proposal is based on the observation that a surface function defined on the object must be preserved around the generator axis of the symmetry. Nevertheless, when objects exhibit partial axial symmetry, the surface functions could not work anymore. The goal in this section is to evaluate several surface functions under a controlled situation where we know the symmetry transformation. We conducted our evaluation over two state-of-the-art functions: Heat Kernel Signatures (HKS) [19] and the Integral Kernel Signature(IKS) [17]. For the sake of completeness, we also include other functions to our analysis. We tested the Wave Kernel Signature [2], the SHOT descriptor [21] and the FPFH descriptor [16]. SHOT and FPFH are not intrinsic, but they provide robustness to partiality.

We built a first 3D object with a perfect axial symmetry. The object was created from a 2D profile curve lying in the XZ plane, which was rotated around the Y axis. For each point in the curve, we therefore can know the set of axially symmetric points (the points obtained by rotating the point on the curve around the Y axis). With this procedure, we built a ground-truth, where we know which points are mutually symmetric. In fact, each point p in the 2D curve generates its ground-truth set G_p composed of its symmetric correspondences. We then built three partial versions of the complete object by cutting the object against a 3D cube. The volume of the three new objects is approximately 75%, 50%, and 25% of the original 3D object. The ground-truth of these partial objects is easily computed by removing the elements in the original ground-truth which do not belong to the partial object anymore.

The evaluation determines the ability of the surface functions to detect symmetric correspondences in partial shapes and to measure the degradation against the partiality. For our evaluation, we chose 100 points on the surface of an object using the farthest point sampling. The first point in the FPS sampling was chosen randomly.

Let q be a point to be analyzed and let f be the function to be evaluated. We want to quantify the matching accuracy of f to detect the axially symmetric points to q. A good alternative is to retrieve points with a similar values of the function and to measure how much points belong to the ground-truth G_q. Formally, we compute the overlap

$$O(K_q, G_q) = \frac{|G_q \bigcap K_q|}{|G_q|}, \tag{4}$$

where K_q is the set of points with smaller distance to q with respect to f. For all our experiments, we set $|K_q| = |G_q|$, thereby the overlap measures the fraction of symmetric correspondences retrieved by the evaluated function.

We also define an approximated measure of overlap which introduces a degree of error in the localization of the correspondences. This measure is defined as follows

$$O_t(K_q, G_q) = \frac{|\{g \in G_q / \exists k \in K_q \text{ and } d_{geod}(g,k) \leq t\}|}{|G_q|} \tag{5}$$

This new measure is more general since $O_{t=0}(K_q, G_q) = O(K_q, G_q)$. In addition, it is interesting to evaluate the accuracy in presence of some error because our method to detect the symmetry involves the use of RANSAC, which is able to deal with error in the localizations. The presented results are obtained by averaging the overlaps over all the analyzed points.

Figure 2a shows the values of overlap for each function where the error t varies in the range $[\epsilon, 10\epsilon]$ where ϵ denotes the median length of edges in the triangular surface of the object. The input object has a uniform triangulation, and therefore ϵ is a good indicator of geodesic error. This first experiment is useful to show that every function is good to characterize the axial symmetry when the object is complete. Note that HKS is better than the other functions; the reason is that HKS tends to give a good description of global geometry which is not affected by partiality.

Nevertheless, the scenario is different when we compute the accuracy in partial objects. The results are presented in figures 2b, 2c and 2d. It seems that when we change the global structure of a 3D object, HKS and WKS exhibit a dramatic drop in their accuracy to detect symmetric correspondences. This would be mainly due to the predominance of low frequencies in the computation of the diffusion. It is well known that the low frequencies in the surface of an object are related to global information. This would be the reason why HKS and WKS can only retrieve 20% of the

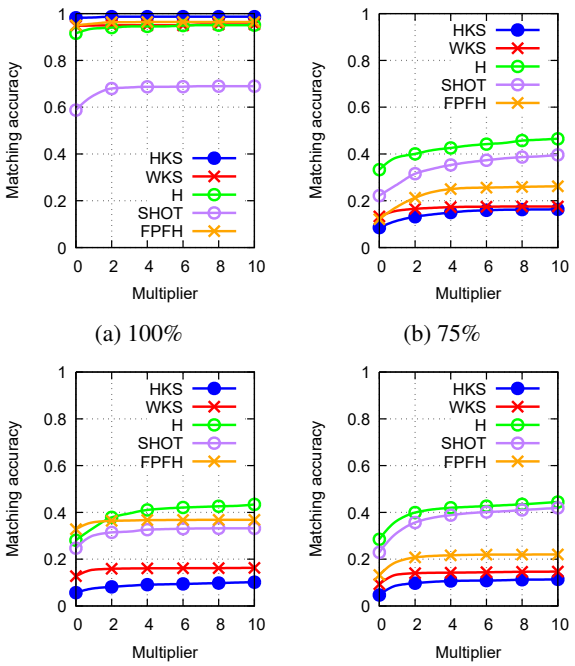

(a) 100% (b) 75%

(c) 50% (d) 25%

Figure 2: Average matching accuracy of surface functions. The plots show the impact of partial shapes in the ability of several functions to give good matches. The percentage means the amount of remaining geometry after cutting the original object.

ground-truth on average, even with a high tolerance error. Interestingly, SHOT and FPFH have a good performance to locate symmetric correspondences in partial shapes. It is mainly because they are designed to be robust to partial data. However, as they are not intrinsic, we believe that the number of false positives during the search of symmetric correspondences could lead to a bad location of the circular structures. On the other hand, IKS is less sensitive to global changes while maintaining a sustained behavior even in 25% of the object. In general, the accuracy of IKS to retrieve points in the ground-truth is close to the 50%. This evaluation shows that IKS is a good alternative for our analysis due to two reasons: i) it is robust to partiality and ii) there is a guarantee to have a low number of false positives in the matching process.

4.3. Results

We had access to a set of real damaged pottery objects from the Larco Museum in Lima, Perú. We scanned the objects using a laser scanner. The scanned objects contain approximately half a million of points and related connectivity information (triangles). In our experiments, we simplify the models to approximately 50,000 vertices to facilitate the

construction of the Laplace-Beltrami operator and the subsequent computation of its eigen-decomposition, which is in turn required to compute the intrinsic function. Figure 3 shows some results of our algorithm (the results over the complete scanned dataset will be provided in supplemental material). The six examples show the robustness of our method to detect the approximate symmetry, even in very challenging situations. Our algorithm is able to deal with missing geometry and partial symmetry support at the same time, thanks to the careful analysis of the surface function and the search of partial circular structures.

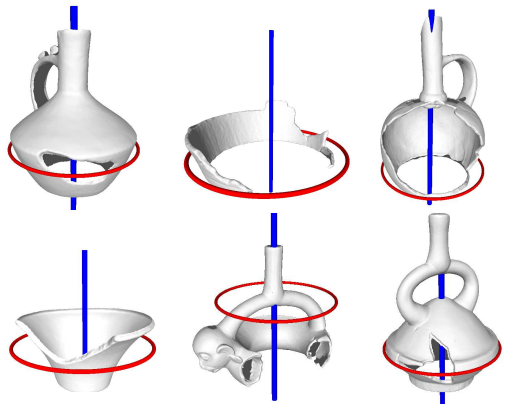

Figure 3: Partial objects and their detected symmetry axes. Objects from Museo Larco, Lima - Perú.

The approximation of the generator axis is good, despite of the existence of missing large portions. The degree of approximation is strongly related to the amount of available geometry in the analysis. Observe that there can be a visible misalignment of the detected axis and the real symmetric axis. This result is expected because we know that the ability of the surface function to detect axial correspondences degrades with partiality. Even so, the approximation is good, and therefore we can use this result to trying to repair the object.

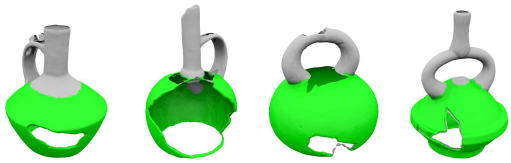

Figure 4: Partial objects with the identified symmetric support. Best viewed in color. Objects from Museo Larco, Lima - Perú

Moreover, our method can compute the geometric support of the partial axial symmetry in a reliable way. Some examples of the identification of the symmetric region can be observed in Figure 4. We used the method described in Section 3.3 to find the region that supports the detected symmetry. We set the threshold λ to 1% of the diagonal of the object.

Figure 5: Repair of damaged objects. First row: input objects. Second row: overlap between input object and the rotated region of support. Third row: final repair. Best viewed in color. Objects from Museo Larco, Lima - Perú

5. Restoration of cultural heritage objects

Our method is able to find a very good approximation of the axial symmetry and the region of support. A natural application of our method is the generation of plausible geometry to complete damaged archaeological objects. Once we have computed the symmetric region of the input shape, we can apply a rotation around the detected generator axis to generate the missing part of the object. Nevertheless, since the detected symmetry is only approximated, the overlap between the original and the rotated shapes is not perfect. Therefore, we apply a last step of non-rigid alignment to improve the blend between both surfaces and enhance the smoothness in the transitions. To tackle this problem we use the ICP-based method proposed by Amberg et al. [1]. This method iterates between two steps: the determination of point correspondences and the optimization of per-point transformations For the first step, we assign a weight zero for points with correspondences farther than 1% of the object diagonal. It avoids that the new geometry collapses to the original shape. For the second step, we constraint the stiffness parameters, which are important to keep the smoothness of the transformations, to low values. Typically in our application, we define a set of stiffness parameters equally distributed between the values two and one. As consequence, we obtain new geometry with smooth transitions.

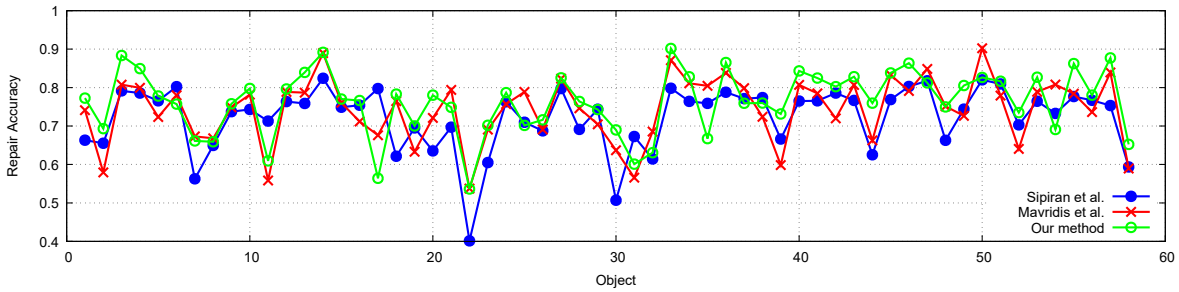

Figure 6: Values of $E_{completion}$ for the comparison of our method and the methods proposed in [17] and [11].

Figure 5 shows some examples of the repair application. Note that our method can generate plausible geometry. For objects with severe damage (forth column in Fig. 5), we apply our algorithm in incremental steps until completing as much geometry as possible. However, our algorithm has the limitation to only replicate existing geometry. In this sense, the algorithm is not able to repair the base in the shown objects.

5.1. A quantitative comparison

We use the benchmark proposed in [4] which provides data and a methodology to evaluate how robust a repair algorithm is. The benchmark comprises 76 synthetically fractured objects, which were originally collected from the Hampson dataset[1]. We selected only 58 out of the 76 objects, which contain axial symmetry. Each object O_i is associated to a set of fragments F_i. For evaluation, one fragment is discarded and the repair task must be conducted in the remaining fragments. The idea is to measure whether the repaired object is similar to the original object. To evaluate the method proposed in this paper, we used the same evaluation protocol proposed in [4].

The evaluation criterion is the congruence between the completed object and the ground-truth. Let O_i be an object in the dataset and let C_i be the object repaired with our algorithm after removing one fragment to the object O_i. The evaluation measure is then defined as

$$E_{completion}(O_i) = \frac{\text{vol}(O_i \bigcap C_i)}{\text{vol}(O_i \bigcup C_i)} \qquad (6)$$

where this measure ranges from zero (no congruence) to one (full congruence).

We compare our method with the methods proposed by Sipiran et al. [17] and Mavridis et al. [11]. The first method finds a few symmetric correspondences which are validated through a vote scheme. The second method formulates the process of finding symmetries as a sparsity-based optimization and it was designed to deal with partial data as well.

[1] http://hampson.cast.uark.edu/

Figure 6 plots the $E_{completion}$ values for each analyzed object. In general, our method shows a better repair rate in most of the objects (40 out of 58). In many objects, the improvement is even close to 10%, which means that our method recovers more significant geometry than the compared methods.

6. Conclusions

In this paper, we addressed the problem of detecting the axial symmetry in objects with partial geometry. We also showed that our method is robust and allows us to use the detected symmetry in the generation of missing geometry to repair cultural heritage objects. An important aspect in our presentation is that we showed the importance of symmetry to recover structure from raw data. In the particular case of axial symmetry, the formulation of the symmetry analysis through the search of similarity in surface functions has proven to be effective. The appropriate surface function can convey rich information about the object in different scales, which, in conjunction with a RANSAC-based matching refinement, can guarantee a detection with a good resilience to partiality.

Our method still have room for improvements. In a future work, we plan to include mid-level information in the process of detection. For example, we can add constraints related to the rim or distinguishable features of objects to improve the determination of the symmetry. On the other hand, in the application level, we plan to investigate the possibility of predicting and exporting the missing geometry to a 3D printer. The 3D printed piece can help to archaeologists to facilitate the conservation process of damages objects.

Acknowledgments

This work has been supported by Programa Nacional de Innovación para la Competitividad y Productividad, INNOVATE Perú, Grant Nr. 280-PNICP-BRI-2015.

References

[1] B. Amberg, S. Romdhani, and T. Vetter. Optimal step nonrigid icp algorithms for surface registration. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, June 2007. 7

[2] M. Aubry, U. Schlickewei, and D. Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *IEEE Int. Conf. in Computer Vision Workshops*, pages 1626–1633, 2011. 5

[3] M. Bokeloh, A. Berner, M. Wand, H.-P. Seidel, and A. Schilling. Symmetry Detection Using Feature Lines. *Computer Graphics Forum*, 28(2):697–706, 2009. 1

[4] R. Gregor, D. Bauer, I. Sipiran, P. Perakis, and T. Schreck. Automatic 3D Object Fracturing for Evaluation of Partial Retrieval and Object Restoration Tasks - Benchmark and Application to 3D Cultural Heritage Data. In I. Pratikakis, M. Spagnuolo, T. Theoharis, L. V. Gool, and R. Veltkamp, editors, *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2015. 8

[5] M. Kazhdan, B. Chazelle, D. Dobkin, T. Funkhouser, and S. Rusinkiewicz. A Reflective Symmetry Descriptor for 3D Models. *Algorithmica*, 38(1):201–225, Oct. 2003. 1

[6] S. Korman, R. Litman, S. Avidan, and A. Bronstein. Probably Approximately Symmetric: Fast Rigid Symmetry Detection With Global Guarantees. *Comput. Graph. Forum*, 34(1):2–13, Feb. 2015. 1

[7] W. K. Leow and R. Li. The analysis and applications of adaptive-binning color histograms. *Comput. Vis. Image Underst.*, 94:67–91, April 2004. 4

[8] B. Li, H. Johan, Y. Ye, and Y. Lu. Efficient View-based 3D Reflection Symmetry Detection. In *SIGGRAPH Asia 2014 Creative Shape Modeling and Design*, SA '14, pages 2:1–2:8, New York, NY, USA, 2014. ACM. 1

[9] Y. Lipman, X. Chen, I. Daubechies, and T. Funkhouser. Symmetry Factored Embedding and Distance. *ACM Trans. Graph.*, 29(4):103:1–103:12, July 2010. 1

[10] A. Martinet, C. Soler, N. Holzschuch, and F. X. Sillion. Accurate Detection of Symmetries in 3D Shapes. *ACM Trans. Graph.*, 25(2):439–464, Apr. 2006. 1

[11] P. Mavridis, I. Sipiran, A. Andreadis, and G. Papaioannou. Object Completion using k-Sparse Optimization. *Computer Graphics Forum*, 34(7):13–21, 2015. 1, 8

[12] M. Meyer, M. Desbrun, P. Schröder, and A. Barr. Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. In H.-C. Hege and K. Polthier, editors, *Visualization and Mathematics III*, Mathematics and Visualization, pages 35–57. Springer Berlin Heidelberg, 2003. 5

[13] N. J. Mitra, L. J. Guibas, and M. Pauly. Partial and Approximate Symmetry Detection for 3D Geometry. *ACM Trans. Graph.*, 25(3):560–568, July 2006. 1

[14] N. J. Mitra, M. Pauly, M. Wand, and D. Ceylan. Symmetry in 3D Geometry: Extraction and Applications. *Comput. Graph. Forum*, 32(6):1–23, 2013. 1

[15] M. Ovsjanikov, J. Sun, and L. Guibas. Global Intrinsic Symmetries of Shapes. In *Proceedings of the Symposium on Geometry Processing*, SGP '08, pages 1341–1348, Aire-la-Ville, Switzerland, Switzerland, 2008. Eurographics Association. 1

[16] R. B. Rusu, N. Blodow, and M. Beetz. Fast point feature histograms (fpfh) for 3d registration. In *2009 IEEE International Conference on Robotics and Automation*, pages 3212–3217, May 2009. 5

[17] I. Sipiran, R. Gregor, and T. Schreck. Approximate symmetry detection in partial 3d meshes. *Computer Graphics Forum (proc. Pacific Graphics)*, 33:131–140, 2014. 1, 5, 8

[18] K. Son, E. Almeida, and D. Cooper. Axially Symmetric 3D Pots Configuration System Using Axis of Symmetry and Break Curve. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 257–264, June 2013. 1

[19] J. Sun, M. Ovsjanikov, and L. J. Guibas. A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. *Comput. Graph. Forum*, 28(5), 2009. 1, 5

[20] S. Thrun and B. Wegbreit. Shape from symmetry. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, volume 2, pages 1824–1831 Vol. 2, Oct 2005. 1

[21] F. Tombari, S. Salti, and L. Di Stefano. Unique signatures of histograms for local surface description. In K. Daniilidis, P. Maragos, and N. Paragios, editors, *11th European Conference on Computer Vision, ECCV 2010*, pages 356–369, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. 5

[22] H. Wang, P. Simari, Z. Su, and H. Zhang. Spectral Global Intrinsic Symmetry Invariant Functions. In *Proceedings of Graphics Interface 2014*, GI '14, pages 209–215, Toronto, Ont., Canada, Canada, 2014. Canadian Information Processing Society. 1

[23] K. Xu, H. Zhang, W. Jiang, R. Dyer, Z. Cheng, L. Liu, and B. Chen. Multi-scale Partial Intrinsic Symmetry Detection. *ACM Trans. Graph.*, 31(6):181:1–181:11, Nov. 2012. 1

[24] K. Xu, H. Zhang, A. Tagliasacchi, L. Liu, G. Li, M. Meng, and Y. Xiong. Partial Intrinsic Reflectional Symmetry of 3D Shapes. *ACM Trans. Graph.*, 28(5):138:1–138:10, Dec. 2009. 1