Eurographics Workshop on 3D Object Retrieval (2018)
A. Telea, T. Theoharis, and R. C. Veltkamp (Editors)

Completion of Cultural Heritage Objects with Rotational Symmetry

Ivan Sipiran[1]

[1]Departamento de Ingeniería, Pontificia Universidad Católica del Perú, PUCP

Abstract

Archaeological artifacts are an important part of our cultural heritage. They help us understand how our ancestors used to live. Unfortunately, many of these objects are badly damaged by the passage of time and need repair. If the object exhibits some form of symmetry, it is possible to complete the missing regions by replicating existing parts of the object. In this paper, we present a framework to complete 3D objects that exhibit rotational symmetry. Our approach combines a number of algorithms from the computer vision community that have had good performance at solving similar problems. In order to complete an archaeological artifact, we begin by scanning the object to produce a 3D mesh of triangles. We then preprocess the mesh to remove fissures and smoothen the surface. We continue by detecting the most salient vertices of the mesh (the key-points). If the object exhibits rotational symmetry, the key-points should form a circular structure which the Random Sample Consensus (RANSAC) algorithm should be able to detect. The axis of symmetry of the circle found should correspond to the axis of symmetry of the object. Thus, by rotating the mesh around the axis of the circle we should be able to complete a large portion of the missing regions. We alleviate any misalignment caused during the rotations via a non-rigid alignment procedure. In the evaluation, we compare the performance of our approach with other state of the art algorithms for 3D object completion. The benchmark proves that our algorithm is effective at completing damaged archaeological objects.

CCS Concepts

•***Computing methodologies*** → ***Parametric curve and surface models; Shape analysis;*** *Mesh geometry models;*

1. Introduction

Object completion is an area of computer vision which seeks to repair incomplete 3D meshes. There are several scenarios where object completion is important. For instance, optic scanners are often unable to scan objects from some angles. As a result, they frequently produce incomplete datasets. Completion algorithms play a key role in acquiring 3D data because they help fill the missing regions that commonly exist in 3D scans. In the cultural heritage context, completion algorithms are valuable because they can be used to show the archaeologist how incomplete artifacts were originally. The archaeologist could even use the digital model to fabricate the missing pieces and actually complete the physical object. In this paper, we are particularly interested in this application of object completion.

Completing damaged cultural heritage objects is not a trivial task. An artifact may be badly worn out and bear no resemblance with its original form. Nevertheless, it is often possible to obtain information from the object to produce a suitable completion. In this paper, we exploit the presence of rotational symmetry to complete the artifacts. We chose this particular kind of symmetry because it is common to find archaeological objects that are hand-crafted using pottery wheels and thus have approximate rotational symmetry. Moreover, we realized that even if the object is severely damaged, it is often possible to identify the axis of symmetry and use it to complete a significant portion of the missing geometry.

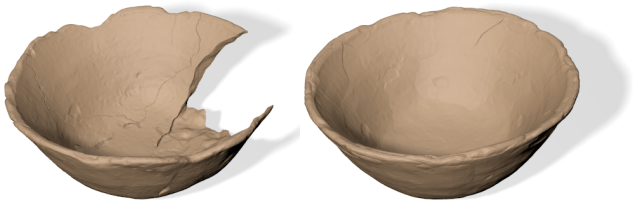

Figure 1: *Object repaired with our algorithm.*

In this paper, we propose a simple yet effective algorithm to detect the axis of symmetry of an object and subsequently complete it. It is important to mention that our approach is a combination of our own algorithms and existing techniques in the computer vision community. We claim that our main contribution is how we put together these tools to solve the completion problem in the cultural heritage context.

The first step of our method involves using key-points to detect evidence of rotational symmetry in the input object. This step is important because we are dealing with broken artifacts, therefore

DOI: 10.2312/3dor.20181057

there is no guarantee that the circular structure of the object is intact. Our method then looks for the circular edge of the object. Using the complete set of key-points, our algorithm extracts the circular structure through a RANSAC-based method. The RANSAC algorithm has proven to be very effective at discarding the outliers of the edge of the object. We continue by computing the approximate axis of symmetry and rotating the mesh about the axis found. Finally, our method performs a non-rigid alignment to merge the original shape with the rotated replicas. Figure 1 depicts a completion obtained using our method.

The paper is organized as follows. Section 2 presents the literature related to our work. We put particular emphasis on symmetry-based methods. Section 3 describes our method to detect approximate rotational symmetry in partially complete objects. Section 4 illustrates our completion algorithm. Section 5 shows the results obtained using our algorithm and compares the performance of our approach with other well-known completion algorithms. Section 6 discusses the limitations of our method and proposes future improvements. Finally, Section 7 concludes the paper.

2. Related Work

Object completion is a challenging topic that has started to capture the attention of the computer vision community thanks to the proliferation of 3D scanning devices. The existing completion methods can be roughly divided in two categories: those which require external references and those which only rely on the incomplete object. The methods of the first category complete objects by using other items as guidelines. For example, Pauly et al. [PMG*05] completed 3D objects by extracting alike artifacts from a repository and merging them with the partially complete objects. Similarly, Huang et al. [HGCO*12] presented a completion algorithm that received two inputs: the incomplete object and a similar complete object. The broken object was completed via extrapolations determined by the complete object.

The methods of the second category try to obtain the information necessary to complete an object from the incomplete object itself. They can be further classified into two subcategories: correspondence algorithms and symmetry algorithms. The correspondence algorithms search for alike local structures within an object. The mesh is completed by filling missing regions from their corresponding parts. For example, Harary et al. [HTG14a] presented a variant of the Heat Kernel Signature to find correspondences within a mesh. Once corresponding regions were found, Harary used an algorithm similar to ICP [RL01] to achieve the completion. Similarly, Zheng et al. [ZSW*10] devised an algorithm that identified repeated structures within an object and used them to predict how its empty regions should be completed.

If the artifact is symmetric, it is also possible to exploit the symmetries to complete the object. In a recent article, Sipiran et al. [SGS14] used a matching algorithm to find symmetric correspondences within a mesh. They proposed a variation of the Heat Kernel Signature [SOG09] which was capable of detecting symmetries even when the mesh was incomplete. Similarly, Mavridis et al. [MSAP15] proposed a registration-based technique which was able to reconstruct 3D objects by finding general symmetries using an optimization approach. Additionally, Son et al. [SAC13] presented a method to repair a fragmented 3D object by using the axis of symmetry and the profile curve of the object.

There are several additional approaches in the literature to computationally complete broken cultural heritage objects. In [IT11], partial shape matchings were used to detect repeated structures which were subsequently used to complete the object. Moreover, in [HT12] Harary and Tal proposed a method to repair incomplete 3D objects by introducing 3D Euler spirals. Harary's method was later improved by adding feature-based knowledge [HTG14b]. In [KLST11, KLST13], Kolomenkin et al. demonstrated that broken 3D objects can be completed from their line drawings.

Our work is closely related to symmetry detection in 3D shapes. Regarding this topic, Thrun and Wegbreit [TW05] defined a probabilistic measurement to assess the existence of symmetries in a 3D object. Xu et al. [XZT*09] proposed an algorithm to find the plane of symmetry of a figure via a voting system. To detect global symmetries, many methods have been studied such as the search of rotation spaces [ZPA95] or spherical harmonic coefficients [KCD*03, KKP13]. More recently, Korman et al. [KLAB15] designed an algorithm that could detect all of the symmetries of a 3D object. The problem with all of these approaches is the fact that they assume that the object is complete and that its centroid lies in a fixed point. It is evident that when dealing with incomplete objects the real centroid is unknown and thus symmetry detection becomes more challenging. For interested readers, a comprehensive study of symmetry detection has been presented by Mitra et al. [MPWC13].

3. Rotational Symmetry Detection

The detection of the rotational symmetry of an object is performed in three steps. The first step process the input shape to simplify and smoothen the surface. The simplification is important to maintain the subsequent tasks efficient enough. The smoothing is important to remove fissures and to merge objects that are composed from several fragments. The second step detects local features in the surface of the object. Our premise is that the rim of rotationally symmetric object should be detected in an local analysis of point neighborhoods. Finally, the third step uses the RANSAC algorithm [FB81] to determine which features form a circular structure. If a circular structure is found, the object exhibit rotational symmetry and therefore its approximate rotation axis can be computed from the geometry.

3.1. Preprocessing

The typical pipeline in the restoration of cultural heritage objects includes a reassembly stage, where many fragments are placed together to build an object. Therefore, a completion method should deal with several fragments as input. The first step before doing the completion is to merge the input to discard the fractured surfaces. We perform a volumetric merging using the OpenVDB library [MLJ*13]. This merging step first transforms out input object from vertices and triangles to a representation made of grids and level sets. The calibration of the number of grids of the level set allows us to remove the fractured surface and to smoothen the

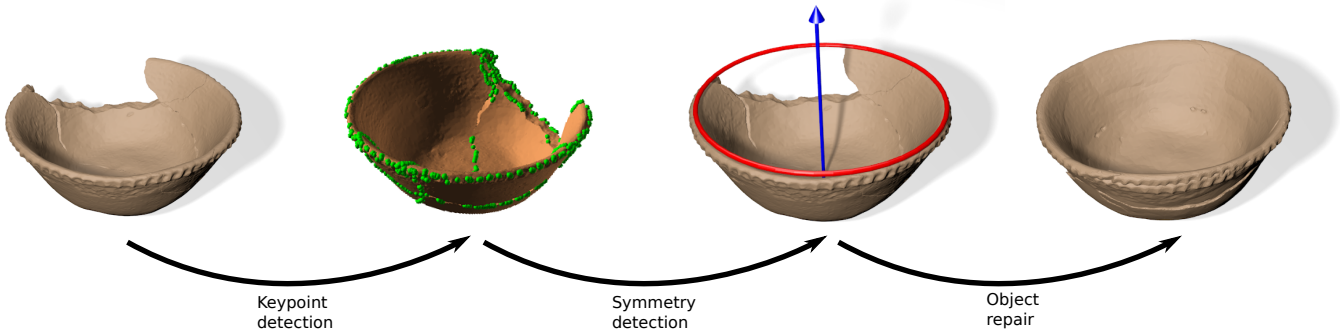

Figure 2: *The pipeline of our method. It starts with the computation of key-points. Next, we detect the rotational symmetry of the input object. Finally, our algorithm repairs the object using the rotation around the detected symmetry axis.*

fissures created by the union of two surfaces from different fragments. Finally, the volumetric output is transformed back to a triangle mesh. The output of this step is a smooth manifold mesh.

Depending on the size of the grids, the output mesh be a high resolution model. We applied the quadric error mesh simplification algorithm proposed by Garland et al. [GH97] to simplify the surface of the object. This step guarantees that the subsequent process can be performed in an efficient way.

3.2. Local Analysis

The main goal of our method is to detect the rim of the object efficiently and effectively. Despite of the preprocessing where we smooth the surface of the object, we must be careful to provide robustness to noise and local changes in the surface. We perform a fast multi-scale analysis of the local surface around a point. Our method approximates the notion of curvature in different scales by using only simple computations.

Let $\mathcal{S}$ be a two-dimensional Manifold surface embedded in $\mathbb{R}^3$. Let $\mathbf{v} \in \mathcal{S}$ be a point on the surface to be analyzed. Our method analyses the local surface around $\mathbf{v}$ to determine whether the point is outstanding in several scales. To facilitate the presentation, we describe the analysis in one scale and thus we generalize to multiple scales. Let σ be a scale parameter. In our application, σ is the radius of a 3D sphere centered in the point $\mathbf{v}$, which we denote as $\mathbb{S}^3_\sigma(\mathbf{v})$. Let $\vec{N}(\mathbf{v})$ be the normal vector associated to the analyzed point $\mathbf{v}$.

The point $\mathbf{v}$ and its normal vector $\vec{N}(\mathbf{v})$ define a tangent plane which we denote as $T(\mathbf{v})$. The intersection $T(\mathbf{v}) \bigcup \mathbb{S}^3_\sigma(\mathbf{v})$ define a 3D circle that is tangent to the point $\mathbf{v}$ and whose radius is σ. We denote this circle as $C_\sigma(v)$ This circle can be used as reference to analyze the degree of protrusion of the local surface patch centered in point $\mathbf{v}$. Let us illustrate the motivation of our analysis with two examples. Consider a flat surface around point $\mathbf{v}$. As the surface is flat around the analyzed point, it is expected that the surface patch contained in $\mathbb{S}^3_\sigma(\mathbf{v})$ is very similar to the tangent circle. As contrast, now consider a protruded surface around the point $\mathbf{v}$. The surface patch contained in $\mathbb{S}^3_\sigma(\mathbf{v})$ is likely to be far from the tangent circle. The similarity between the local patch and the tangent circle depends on how protruded the surface is around $\mathbf{v}$. Therefore, our approach is to measure the similarity between the tangent circle and the local patch surrounding the point $\mathbf{v}$.

We propose a fast approximation to measure the similarity between the tangent circle and the local surface around an analyzed vertex. Given a point $\mathbf{v}$, we extract the local patch around $\mathbf{v}$ contained in the sphere $\mathbb{S}^3_\sigma(\mathbf{v})$. The local patch is denoted as $P_\sigma(v)$. Our method consists of projecting $P_\sigma(v)$ onto the tangent circle $C_\sigma(v)$. An alternative similarity measure is to compare the areas of $C_\sigma(v)$ and the projected patch. However, we take a more efficient approach that has been effective in practice. Our method requires to identify the outer boundary of the patch $P_\sigma(v)$ and to take only vertices on this boundary in the analysis. The identification of boundary vertices can be easily performed while extracting the patch. We compute the projection of every outer vertex onto the circle $C_\sigma(v)$. We thus average the distance between point $\mathbf{v}$ and the projected points. Our criterion to compute the level of protrusion of analyzed point $\mathbf{v}$ is to compute the ratio $r_\sigma(v)$ between the average distance of projected point and the scale σ.

To add the multi-scale characteristic to our analysis, we perform the aforementioned computation in three different scales $\sigma_1 = 0.1 \times D$, $\sigma_2 = 0.2 \times D$ and $\sigma_3 = 0.3 \times D$, where D is the diagonal of the bounding box of the input object. An analyzed point $\mathbf{v}$ is identified as keypoint if the ratios $r_{\sigma_1}(v)$, $r_{\sigma_2}(v)$ and $r_{\sigma_3}(v)$ are not greater than 0.5.

3.3. Detection of Circular Structure

The set of detected key-points can be considered as a 3D point cloud. Our goal was to detect the hidden partial circle in this point cloud. As we can have many outliers, we decided to use the well known method RANSAC [FB81]. Roughly speaking, given the set of key-points and a model for a 3D circle (parameterized as a center and a normal vector), we applied the RANSAC method in order to find the better model which covers as much as possible the point cloud. In every iteration of the algorithm, we chose three random points and found the circle that fitted the three points. Let P_1, P_2, and P_3 be the selected points and let P_1P_2 and P_2P_3 be two line segments formed by the selected three points. The center C of the circle passing through P_1, P_2 and P_3 is the intersection of the perpendicular lines which pass through the midpoint of P_1P_2 and P_2P_3. In

addition, the radius of the circle is the magnitude $R = \|P_1 - C\|_2$. Finally, the orientation of the circle can be defined by the vector $N = \vec{P_1P_2} \times \vec{P_2P_3}$.

The 3D circle is then defined by the tuple (C,R,N). Once we have computed the circle for the three random points, the next step is to verify whether the circle has consensus or not. For every remaining point P in the point cloud, we compute the distance of the point to the circle. If the distance is smaller than a threshold δ, we assign the point P to the circle (we also can call the point an inlier). The goal of RANSAC is to find the circle with the greater number of assignments. In this sense, the procedure of selecting a random circle and computing the assignments is performed a number of iterations. The circle with the greater number of assignments at the end of the process is selected as the generator of the symmetry axis. In all our experiments, we used a threshold $\delta = 0.05$ and 500 iterations.

4. Object Reconstruction

Using the axis found as described in 3, it is possible to reconstruct the object by rotating it around the axis. For the examples presented in this paper, we perform three consecutive rotations of 90 degrees around the symmetry axis. Nevertheless, rotations will not always align with the original object. To address this issue we implemented the non-rigid alignment presented by Pauli et al. [PMG*05]. For every vertex v_j of the rotated mesh, a displacement vector t_j is calculated and used to translate each point. The transformation should be such that the deformed rotated model matches the original object as much as possible. Also, the distortion induced in the rotated mesh should be as small as possible. A shape penalty function is defined which takes into account the amount of distortion induced by the transformation and also the geometric error. The transformation is computed iteratively by minimizing the shape penalty function with respect to the unknown displacement vectors t_j.

In order to keep our paper self-contained, we briefly describe the non-rigid alignment formulation as proposed in [PMG*05]. The search of displacement vectors t_j can be stated as an optimization problem. The formulation includes two terms: the geometric error and the distortion. The geometric error measures the fitness between the original model and the transformed one through the application of the displacement transformation T. Let O be the original model and P be the transformed model, the geometric error can be defined as

$$E_{geometric}(O,P,T) = \sum_{j \in O} \|o_j + t_j - p_j\| \tag{1}$$

where $o_j + t_j$ is the displaced, original $j-$th point in shape O and q_j is the closest point to o_j in P.

Additionally, Pauly et al. [PMG*05] introduced a regularization term that promotes the smoothness of the displacements. The term can be defined as

$$E_{distortion}(O,T) = \sum_{j \in O} \sum_{k \in N_1(j)} A_{jk} \left(\frac{t_j - t_k}{|e_{jk}|} \right)^2 \tag{2}$$

where $N_1(j)$ is the 1-ring neighborhood of vertex j, A_{jk} is the Voronoi area related to the edge between vertex j and vertex j, and e_{jk} is the edge between vertex j and vertex k.

Both error are combined in a single optimization problem as follows

$$\arg\min_T \alpha \cdot E_{distortion}(O,T) + (1-\alpha) \cdot E_{geometric(O,P,T)} \tag{3}$$

where α weights the contribution of each measure in the optimization. Note that the aforementioned optimization problem can be solved by a conjugate gradient solver. It is important to mention that, in our work, the non-rigid alignment is performed after every rotation, so the input for the next alignment is the merged object and the new rotated version of the original input object. Furthermore, the parameter α is iteratively relaxed to avoid premature skewing and attraction of the geometries by false near surfaces.

5. Results

In this section we show the results of applying our strategy to restore archaeological objects. For our evaluation, we used the benchmark proposed by Gregor et al. [GBS*15] (available online on http://fracture-benchmark.dbvis.de). The benchmark consists of artificially fractured objects which can be used to evaluate repair algorithms. The benchmark is based on the Hampson dataset (available on http://hampson.cast.uark.edu/). It is worth mentioning that we only select the objects that exhibit approximate rotational symmetry (58 out of 76 objects, the list can be found in Table 1).

Figure 3 presents some results obtained with our method. The five examples show the ability of our method to complete objects. The common characteristic of these examples is that the circular edge has not been lost considerably. Still we want to emphasize the capacity of our method of detecting a good approximation of the symmetry axis and subsequently complete the object.

Figure 4 shows more challenging examples. The first example (Fig. 4 at left) shows the ability of our method to complete an object with a large missing part. In this case, our algorithm takes advantage of the presence of the circular edge to detect the symmetry axis and perform the subsequent reconstruction. However this example is particularly difficult because there is a long broken edge (the interior edge) which complicates the task of detecting the correct symmetry axis. Nevertheless our method had a good performance. The same scenario can be observed in the second and third examples (Fig 4 at center and right), where the circular rim is only partially present and the size of the fracture is significant. Still our method is effective at completing these two examples as well.

5.1. Quantitative Evaluation

We used the benchmark proposed by Gregor et al. [GBS*15] to do our evaluation. The benchmark consists of 76 fractured objects which were originally collected from the Hamspson dataset. We selected 58 out of the 76 objects that contain rotational 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

Figure 3: *Reconstruction of objects with our proposed method. In the first row are the input objects. In the second row are the automatically completed objects.*

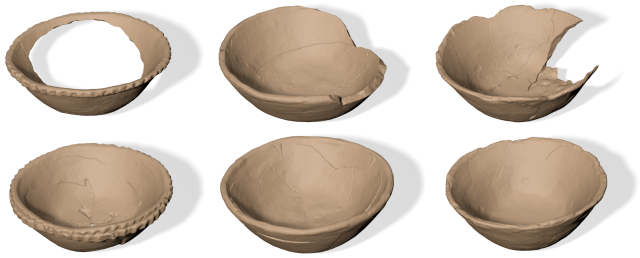

Figure 4: *Completion in challenging objects. First row: input objects. Second row: objects repaired with our algorithm.*

the remaining fragments. The idea is to measure whether the reconstruction is similar to the original object. To evaluate the method proposed in this paper, we used the same evaluation protocol proposed in [GBS*15].

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 completed with our algorithm after removing one fragment to the object O_i. The evaluation measure is then defined as

$$E_{completion}(O_i) = \frac{vol(O_i \bigcap C_i)}{vol(O_i \bigcup C_i)} \quad (4)$$

This measure ranges from zero (no congruence) to one (full congruence). We compare our method with three methods in the state of the art: the vote-based technique proposed by Sipiran et al. [SGS14], the optimization-based method proposed by Mavridis et al. [MSAP15] and the recent heat diffusion approach proposed by Sipiran [Sip17].

Table 1 shows the $E_{completion}$ measure for all the evaluated algorithms (numbers in bold are the best performance per model). Our method outperforms the reconstruction of most of the models in the benchmark collection with respect to the state-of-the-art methods. From the compared methods, only our method and the method proposed by Sipiran [Sip17] were proposed specifically for rotational symmetric shapes. Nevertheless, the advantage of our method is that it is considerably simpler and more efficient. Most of the time of our algorithm is spent in the key-point detection step, which is performed only using simple local operations. After this step, we filter a lot of information in the analysis, and subsequently the analysis time is negligible. It is still interesting how our method can perform better than other methods using simple and efficient operations.

6. Limitations and Improvements

Our method has some limitations and we would like to discuss them in order to propose future improvements. We have identified two limitations and we devote the following sections to present them.

6.1. Axis of symmetry that passes through an incomplete region

Once we have computed the symmetry axis, the completion is performed by rotating the object around this axis. However, if the symmetry axis goes through an incomplete part of the object, it is not possible to complete the geometry around the axis. This behavior is depicted in Figure 5. Although our algorithm can complete a large portion of the input shape, there is still a hole in the middle. Obviously, to cope with this problem we require a post-processing to complete the missing geometry. Depending on the amount of missing geometry, we would like to explore the analysis of shape profiles to complete such holes. When the resulting hole is big, a good choice would be the use of external references to complete the missing part.

6.2. Objects with partial rotational symmetry

As can be observed in the examples we presented in the whole article, our method perform very well when the input shape is an approximate solid of revolution. Nevertheless, there are some real cases where the object contains additional features such as handles. Some examples of such objects can be seen in Figure 6. It is interesting to note that these objects can be considered as a solid of revolution plus some additional features. Therefore, we can still

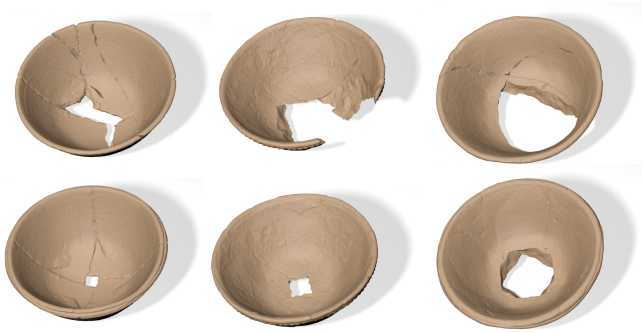

Figure 5: *First row: Input shapes with a missing part around the symmetric axis. Second row: completed shapes obtained with our algorithm.*

use our algorithm to reconstruct the part of the object that is rotationally symmetric. We call this problem a reconstruction of partial rotationally symmetric objects. To deal this problem, we would like to explore the incorporation of a previous step to perform a mesh segmentation of the input shape before applying our algorithm.

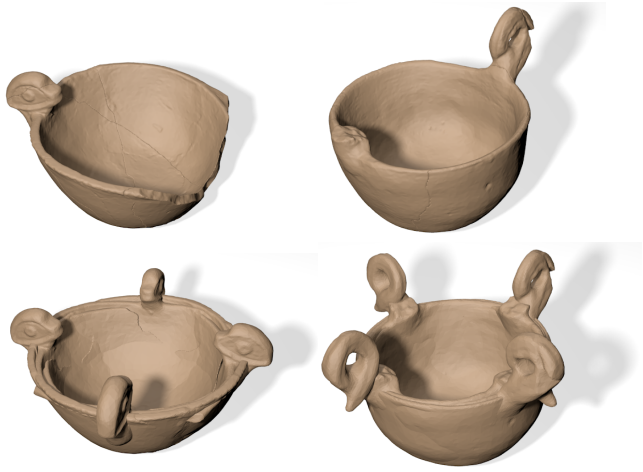

Figure 6: *First row: Input shapes with additional shape features. Second row: completed objects with our algorithm.*

7. Conclusions

In this paper, we presented a method to complete 3D archaeological objects that exhibit rotational symmetry. Our method is a conjunction of of our own key-point detection algorithm and known techniques from the computer vision and shape matching communities. The resulting framework is very effective and efficient at completing symmetrical broken objects. Even though our method is specific, the results show that our algorithm could be used in real-world scenarios. Additionally, as discussed in this paper, we would like to extend our work in several directions in the future. First, our method requires a post-processing procedure to complete missing geometry around the symmetry axis. Second, we plan to introduce the use of external reference objects to help the completion process. Finally, we also plan to include a mesh segmentation procedure in order to deal with handles and additional features in archaeological objects.

References

[FB81] FISCHLER M. A., BOLLES R. C.: Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM 24*, 6 (June 1981), 381–395. 2, 3

[GBS*15] GREGOR R., BAUER D., SIPIRAN I., PERAKIS P., SCHRECK T.: Automatic 3D Object Fracturing for Evaluation of Partial Retrieval and Object Restoration Tasks - Benchmark and Application to 3D Cultural Heritage Data. In *Eurographics Workshop on 3D Object Retrieval* (2015), Pratikakis I., Spagnuolo M., Theoharis T., Gool L. V., Veltkamp R., (Eds.), The Eurographics Association. 4, 5

[GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1997), SIGGRAPH '97, ACM Press/Addison-Wesley Publishing Co., pp. 209–216. 3

[HGCO*12] HUANG H., GONG M., COHEN-OR D., OUYANG Y., TAN F., ZHANG H.: Field-guided registration for feature-conforming shape composition. *ACM Trans. Graph. 31*, 6 (Nov. 2012), 179:1–179:11. 2

[HT12] HARARY G., TAL A.: 3d euler spirals for 3d curve completion. *Computational Geometry 45*, 3 (2012), 115 – 126. 2

[HTG14a] HARARY G., TAL A., GRINSPUN E.: Context-based Coherent Surface Completion. *ACM Trans. Graph. 33*, 1 (Feb. 2014), 5:1–5:12. 2

[HTG14b] HARARY G., TAL A., GRINSPUN E.: Feature-preserving surface completion using four points. *Comput. Graph. Forum 33*, 5 (Aug. 2014), 45–54. 2

[IT11] ITSKOVICH A., TAL A.: Surface partial matching and application to archaeology. *Computers & Graphics 35*, 2 (2011), 334 – 341. Virtual Reality in BrazilVisual Computing in Biology and MedicineSemantic 3D media and contentCultural Heritage. 2

[KCD*03] KAZHDAN M., CHAZELLE B., DOBKIN D., FUNKHOUSER T., RUSINKIEWICZ S.: A Reflective Symmetry Descriptor for 3D Models. *Algorithmica 38*, 1 (Oct. 2003), 201–225. 2

[KKP13] KAKARALA R., KALIAMOORTHI P., PREMACHANDRAN V.: Three-Dimensional Bilateral Symmetry Plane Estimation in the Phase Domain. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on* (June 2013), pp. 249–256. 2

[KLAB15] KORMAN S., LITMAN R., AVIDAN S., BRONSTEIN A.: Probably Approximately Symmetric: Fast Rigid Symmetry Detection With Global Guarantees. *Comput. Graph. Forum 34*, 1 (Feb. 2015), 2–13. 2

[KLST11] KOLOMENKIN M., LEIFMAN G., SHIMSHONI I., TAL A.: Reconstruction of relief objects from line drawings. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on* (June 2011), pp. 993–1000. 2

[KLST13] KOLOMENKIN M., LEIFMAN G., SHIMSHONI I., TAL A.: Reconstruction of relief objects from archeological line drawings. *J. Comput. Cult. Herit. 6*, 1 (Apr. 2013), 3:1–3:19. 2

[MLJ*13] MUSETH K., LAIT J., JOHANSON J., BUDSBERG J., HENDERSON R., ALDEN M., CUCKA P., HILL D., PEARCE A.: Openvdb: An open-source data structure and toolkit for high-resolution volumes. In *ACM SIGGRAPH 2013 Courses* (New York, NY, USA, 2013), SIGGRAPH '13, ACM, pp. 19:1–19:1. 2

[MPWC13] MITRA N. J., PAULY M., WAND M., CEYLAN D.: Symmetry in 3D Geometry: Extraction and Applications. *Comput. Graph. Forum 32*, 6 (2013), 1–23. 2

[MSAP15] MAVRIDIS P., SIPIRAN I., ANDREADIS A., PAPAIOANNOU G.: Object Completion using k-Sparse Optimization. *Computer Graphics Forum 34*, 7 (2015), 13–21. 2, 5, 7

[PMG*05] PAULY M., MITRA N. J., GIESEN J., GROSS M., GUIBAS L. J.: Example-based 3d scan completion. In *Proceedings of the Third Eurographics Symposium on Geometry Processing* (Aire-la-Ville, Switzerland, Switzerland, 2005), SGP '05, Eurographics Association. 2, 4

[RL01] RUSINKIEWICZ S., LEVOY M.: Efficient variants of the ICP algorithm. In *Third International Conference on 3D Digital Imaging and Modeling (3DIM)* (June 2001). 2

[SAC13] SON K., ALMEIDA E. B., COOPER D. B.: Axially symmetric 3d pots configuration system using axis of symmetry and break curve. In *Proceedings of the 2013 IEEE Conference on Computer Vision and Pattern Recognition* (Washington, DC, USA, 2013), CVPR '13, IEEE Computer Society, pp. 257–264. 2

[SGS14] SIPIRAN I., GREGOR R., SCHRECK T.: Approximate Symmetry Detection in Partial 3D Meshes. *Comput. Graph. Forum 33*, 7 (Oct. 2014), 131–140. 2, 5, 7

[Sip17] SIPIRAN I.: Analysis of partial axial symmetry on 3d surfaces and its application in the restoration of cultural heritage objects. In *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)* (Oct 2017), pp. 2925–2933. 5, 7

[SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. In *Proceedings of the Symposium on Geometry Processing* (Aire-la-Ville, Switzerland, Switzerland, 2009), SGP '09, Eurographics Association, pp. 1383–1392. 2

[TW05] THRUN S., WEGBREIT B.: Shape from symmetry. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on* (Oct 2005), vol. 2, pp. 1824–1831 Vol. 2. 2

[XZT*09] XU K., ZHANG H., TAGLIASACCHI A., LIU L., LI G., MENG M., XIONG Y.: Partial Intrinsic Reflectional Symmetry of 3D Shapes. *ACM Trans. Graph. 28*, 5 (Dec. 2009), 138:1–138:10. 2

[ZPA95] ZABRODSKY H., PELEG S., AVNIR D.: Symmetry as a continuous feature. *IEEE Trans. Pattern Anal. Mach. Intell. 17*, 12 (Dec. 1995), 1154–1166. 2

[ZSW*10] ZHENG Q., SHARF A., WAN G., LI Y., MITRA N. J., COHEN-OR D., CHEN B.: Non-local Scan Consolidation for 3D Urban Scenes. *ACM Trans. Graph. 29*, 4 (July 2010), 94:1–94:9. 2

Table 1: *Results per object*

ID	[SGS14]	[MSAP15]	[Sip17]	Ours
6	0.820	**0.902**	0.826	0.874
10	0.663	0.741	**0.772**	0.754
31	0.697	**0.793**	0.748	0.630
38	0.687	0.691	0.716	**0.813**
45	0.614	0.685	0.630	**0.743**
46	0.798	0.871	**0.901**	0.678
48	0.787	0.838	0.865	**0.877**
73	0.776	0.783	0.862	**0.943**
77	0.766	0.736	**0.780**	0.509
81	0.593	0.589	0.652	**0.759**
112	0.655	0.579	0.693	**0.900**
122	0.791	0.807	0.883	**0.971**
194	0.562	**0.673**	0.661	0.663
202	0.649	0.668	0.659	**0.672**
208	0.737	0.748	0.757	**0.987**
216	0.743	0.781	0.797	**0.820**
217	0.712	0.558	0.609	**0.846**
218	0.763	0.788	**0.796**	0.721
223	0.758	0.786	**0.839**	0.626
234	0.824	0.889	**0.891**	0.830
240	0.748	0.760	**0.770**	0.719
263	0.753	0.711	0.766	**0.769**
267	0.797	**0.675**	0.564	0.569
270	0.621	0.764	0.782	**0.879**
285	0.695	0.633	**0.700**	0.516
288	0.635	0.721	**0.780**	0.594
338	0.401	0.537	0.536	**0.638**
340	0.605	0.690	**0.702**	0.455
341	0.762	0.757	0.786	**0.987**
375	0.709	0.788	0.701	**0.985**
418	0.797	0.823	0.825	**0.899**
419	0.691	0.746	**0.763**	0.506
424	0.743	0.704	**0.742**	0.529
435	0.506	0.637	**0.689**	0.657
447	0.672	0.565	**0.601**	0.572
480	0.764	0.811	0.828	**0.978**
489	0.758	0.804	0.667	**0.864**
521	0.771	**0.798**	0.759	0.783
538	0.773	0.724	0.759	**0.810**
546	0.666	0.598	**0.731**	0.719
554	0.764	0.807	**0.843**	0.589
586	0.765	0.784	**0.824**	0.524
595	0.785	0.719	**0.802**	0.620
597	0.766	0.806	**0.828**	0.599
612	0.625	0.662	0.759	**0.820**
641	0.768	0.832	0.838	**0.921**
653	0.803	0.791	**0.863**	0.809
661	0.817	**0.848**	0.812	0.827
670	0.662	**0.751**	0.749	0.482
682	0.743	0.726	0.805	**0.998**
704	0.809	0.779	**0.816**	0.579
720	0.702	0.640	0.735	**0.840**
721	0.764	0.789	0.827	**0.837**
728	0.732	0.807	0.691	**0.972**
786	0.753	0.839	**0.877**	0.817
1300	0.785	0.799	**0.849**	0.540
1301	0.765	0.723	0.778	**0.974**
1302	0.802	0.779	0.757	**0.998**