Pacific Graphics 2015
N. J. Mitra, J. Stam, and K. Xu
(Guest Editors)

Volume 34 (2015), Number 7

Object Completion using *k*-Sparse Optimization

P. Mavridis[1], I. Sipiran[2], A. Andreadis[1] and G. Papaioannou[1]

[1]Department of Informatics, Athens University of Economics and Business, Athens, Greece
[2] Sección de Ingeniería Informática, Pontificia Universidad Católica del Perú, Lima, Perú

Abstract

We present a new method for the completion of partial globally-symmetric 3D objects, based on the detection of partial and approximate symmetries in the incomplete input dataset. In our approach, symmetry detection is formulated as a constrained sparsity maximization problem, which is solved efficiently using a robust RANSAC-based optimizer. The detected partial symmetries are then reused iteratively, in order to complete the missing parts of the object. A global error relaxation method minimizes the accumulated alignment errors and a non-rigid registration approach applies local deformations in order to properly handle approximate symmetry. Unlike previous approaches, our method does not rely on the computation of features, it uniformly handles translational, rotational and reflectional symmetries and can provide plausible object completion results, even on challenging cases, where more than half of the target object is missing. We demonstrate our algorithm in the completion of 3D scans with varying levels of partiality and we show the applicability of our approach in the repair and completion of heavily eroded or incomplete cultural heritage objects.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations

1. Introduction

Digital acquisition of 3D objects is an important task in several domains, including computer graphics, computer vision and cultural heritage preservation. However, optical 3D acquisition systems often produce incomplete datasets, due to occlusions, unfavourable surface reflectance properties or geometric restrictions that prevent capturing the target object from certain angles. Therefore, it is clear that object completion algorithms play a major role during the 3D acquisition process, in order to properly fill any missing parts from the digitized dataset.

Furthermore, in the context of cultural heritage, object completion algorithms can be used to suggest plausible completions for objects with missing parts, which is one of the main motivations behind our work. When used in conjunction with fractured object reassembly algorithms [HFG*06,MAP15], this allows for the creation of automated repair systems that can present to the archaeologist a number of plausible reassemblies and completions, in order to provide visual aids for restoration tasks. The digital expansion of the input object can even be exploited for the fabrication of the missing parts to physically complete the artefact.

Many methods in the literature focus on the completion of relatively small missing parts, thus mostly addressing small defects (holes) in the first use case that we have described. In this paper, we focus on the completion of partially scanned or incomplete 3D objects, where a significant portion of the target shape is missing from the input. Our approach exploits the fact that many natural or man-made objects exhibit some form of symmetries. By detecting such symmetries, missing geometry can be filled by replicating parts of the existing input object. However, this process can be very challenging, especially when the symmetries involved are not perfect. We successfully address these issues, by introducing an object completion pipeline that is based on a new method for partial and approximate symmetry detection.

1.1. Contributions

Our contributions can be summarized as follows:

- We present a mathematical formulation of the partial and approximate symmetry detection problem that is based on the maximization of the sparsity of a residual distance vector, subject to the constraint that *at least k elements should not be zero* (Sec. 3.1).

- We propose a RANSAC-based registration method to efficiently solve the underlying global optimization problem and detect the available symmetries (Sec. 3.2).
- We introduce a pipeline that uses the detected partial symmetries to complete the missing parts of a globally-symmetric object. This process includes a global relaxation of the accumulated error and a non-rigid deformation of the replicated parts, to better handle approximate symmetries (Sec. 4).

The key strengths of our approach is that it can directly operate on point clouds, it does not rely on the computation of features or rich local descriptors and uniformly exploits translational, rotational and reflectional partial symmetries. An additional advantage of our method is that it can be implemented with few additions to an existing surface registration system, facilitating the adoption of our approach.

We demonstrate our method in the completion of partial 3D scans of approximately symmetrical objects, where missing parts are reliably predicted with our approach, even when more than half of the object is missing from the input. We also present an application of our method in the context of cultural heritage, where our completion method is used to synthesize the missing parts of partially reassembled objects.

2. Related Work

Object completion is a challenging topic that has gained attention recently with the proliferation of scanning devices. The methods that deal with this problem can be roughly divided into two general categories. The methods in the first category use external reference objects for the completion. For example, Pauly et al. [PMG*05] proposed to complete a 3D scan using similar objects from a repository. After retrieving the most similar object to the query, a non-rigid registration step was applied in order to fit both geometries as much as possible. This approach has been later improved by Li et al. [LAGP09]. These completion methods are orthogonal and can be combined with our approach, as we demonstrate later in this paper. Also, Huang et al. [HGCO*12] proposed an object completion algorithm based on the registration with feature-conforming fields. Missing regions are completed by extrapolating the existing shape, which is not always correct, and at least two input shapes are required to do the completion, whereas our method needs only one.

The second category of methods try to obtain enough information for completion from the input object itself, using self-similarity. We can identify two sub-types of self-similarities that have been exploited so far: self-correspondences and self-symmetries. The use of self-correspondences is exploited in the sense of finding a good set of local correspondences that can be used to transfer geometry to the unknown geometry. Harary et al. [HTG14a] proposed the use of a variation of the *Heat Kernel Signature* [SOG09] to find candidate matches. With the set of candidates, the method used an ICP-like algorithm to copy the corresponding geometry. Similarly in [HTG14b], the authors included user interaction to help the completion process. In this case, the user provides four points as hint for the completion of a feature curve. Then, the algorithm performs an automatic filling of the missing geometry using the self-correspondences obtained with diffusion descriptors.

On the other hand, if the input object exhibits some form of symmetry, it is also possible to use this characteristic to tackle the problem of completion. Thrun and Wegbreit [TW05] defined a probabilistic measure to score the existence of symmetries in range scans. The method was applied for the completion of partial views of scanned objects. Similarly, Zheng et al. [ZSW*10] proposed finding repeated structures (translational symmetries) in LiDAR data. The algorithm was applied to complete 3D urban scenes from sparse point clouds. Xu et al. [XZT*09] defined an algorithm to find the intrinsic reflectional symmetry axis of a shape. The algorithm was able to deal with missing parts mainly due to the use a voting scheme of local features. The method was applied to roughly determine the missing parts of objects with non-rigid transformations. Likewise, Jiang et al. [JXCZ13] devised an algorithm to detect intrinsic symmetries using the curve skeleton in point clouds. The algorithm was also shown to produce approximate completion in objects with non-rigid transformations. More recently, Sipiran et al. [SGS14] proposed an algorithm to detect reflectional symmetries in objects with large missing geometry. The algorithm finds symmetric correspondences, which are subsequently used to produce candidate reflectional planes. The limitation of all the aforementioned approaches lies in the use of features, which are not guaranteed to exist in many practical object examples. The method proposed in this paper does not rely on the presence of features in the input shape and can be applied directly on unstructured point clouds, making our algorithm more general.

Our work is also related to more general methods that aim to detect symmetries on 3D data. A comprehensive exposition of this subject can be found in the survey by Mitra et al. [MPWC13]. In the case of global symmetry detection, many approaches have been studied such as exhaustive search in rotational space [ZPA95], spherical harmonics coefficients [KFR04, KKP13] and 3D Pseudo-polar Fourier Transform [BAK10]. More recently, Korman et al. [KLAB15] devised a global algorithm to approximately find most of the symmetries of a 3D object with theoretical guarantees. The problem with these approaches is the assumption that the object is complete and the centroid of the object is a fixed point of the symmetries. It is clear that with incomplete objects, the centroid of the real object is not known, and therefore the problem of detecting symmetries is harder. In this direction, methods that use local information have been devised. Some interesting approaches include voting scheme of local features [MGP06, LCDF10, XZJ*12] and graph matching [THW*14].

3. Partial and Approximate Symmetry Detection

At the core of our completion pipeline is a new approach for the detection of partial symmetries. In this section we first present a mathematical formulation of the problem, based on sparsity-inducing norms, and then we present a method to solve it efficiently, based on a robust *RANdom SAmple Consensus* (RANSAC) algorithm.

3.1. Mathematical Formulation

We say that a geometric object M is *symmetric* with respect to a transformation T, if $M = T(M)$. However, very few natural or man-made objects exhibit a perfect symmetry that follows this very strict definition; most of the time, there are deviations from a precise symmetry transformation. To enable a more general definition, that encompasses these cases of *approximate* symmetry, we introduce a distance function $d(X,Y)$ that measures the distance between object X and Y. We can now state that two objects are approximately symmetrical if $d(M,T(M)) \leq \varepsilon$.

There are many choices for the distance function d. As discussed in the recent survey by Mitra et al. [MPWC13], most methods use variants of squared or Hausdorff distances. In our approach, since we rely on the symmetry for the completion of a partial object, we use a distance function that measures the amount of overlap between the surface M and $T(M)$. To this end, we define the vector of *residual distances* $\mathbf{z} \in \mathbb{R}^n$, whose elements are defined as:

$$z_i(T) = \phi(T(x_i), M), \quad x_i \in M$$

The overlap can then be measured as the ℓ_0-norm of the residual vector $\mathbf{z}$ between the two surfaces, giving the following distance definition:

$$d(M,T(M)) = \|\mathbf{z}\|_0, \tag{1}$$

where the ℓ_0-norm $\|.\|_0$ counts the number of non-zero elements in a vector and the function $\phi(\mathbf{x},M)$ measures the distance of an arbitrary point $\mathbf{x} \in \mathbb{R}^3$ to the surface M. In particular, $\phi(\mathbf{x},M)$ is defined as a binary metric, which is one for the points in space whose distance to the surface M is lower than a threshold ε and zero otherwise. Such a metric is commonly used in *Largest Common Pointset* (LCP) problems, when the ℓ_0-norm is used instead of the ℓ_2 one.

In the context of object completion, we are only interested in detecting partial symmetries. In particular, the objects M and $T(M)$ should not have a complete overlap, in order to use the non-overlapping points for the completion of the final shape. We can formulate this requirement as the following minimization problem:

$$\begin{aligned} &\arg\min_T d(M,T(M)), \\ &\text{subject to} \quad d(M,T(M))/n \geq k, \end{aligned} \tag{2}$$

where n is the total number of points in surface M and the parameter k controls the minimum percentage (or fraction) of non-overlapping points in the detected symmetry. In practice, this equation maximizes the sparsity (number of zero elements) of the residual vector $\mathbf{z}$, but also keeps at least $k-$percent non-zero elements. We use the term *k-sparse* to refer to a vector with these properties.

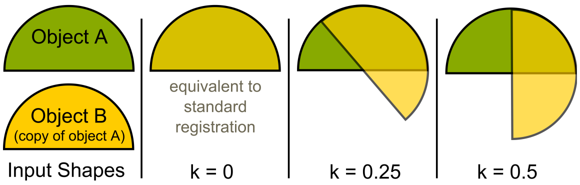

Figure 1: *k-sparse registration maximizes the sparsity of the vector of residual distances between two surfaces, but introduces the constraint that at least k% of the elements should be non-zero. For $k = 0$, the result is equivalent to standard registration. For $k > 0$, the method enforces partial overlap, even for identical surfaces.*

3.2. k-Sparse Optimization

The key observation of our approach is that Equation 2 describes the alignment of the input surface with a transformed copy of itself. By keeping the alignment transformation rigid (rotation and translation), we can detect symmetries by using well-known global rigid registration algorithms, with the additional modification that at least k-percent of the surface points should not be overlapping. Reflectional symmetries can be handled similarly, by registering the input object with the reflection of itself.

Based on this simple observation, we solve Equation 2 using the well-known Super4PCS method [MAM14], a fast global registration algorithm for point sets, which runs in optimal linear time and is robust to outliers. Super4PCS first selects a random set B of 4-points, called a *basis*, in the source surface and then efficiently extracts a list of all 4-point sets in the target surface that are approximately congruent with B. For each set of congruent 4-points in this list, a candidate transformation which aligns these points with the basis B is computed analytically [Ume91, ELF97] and then tested for validity using a voting scheme. In this validation procedure, the points of the input surface are transformed according to the candidate transformation and then the distance is computed according to Equation 1.

Parameter Sensitivity. In order to take our constraint into account (keep at least k-percent non-zero elements), the most trivial approach is to simply discard all candidate transformations for which the distance $d(M,T(M))$ is less than the parameter k. This approach effectively enforces partial overlap in the registered surfaces, as shown in Figure 1, and the computed partial symmetry transformation can be used for the completion of the input object. However, in our experiments we have found that this approach leads to very high

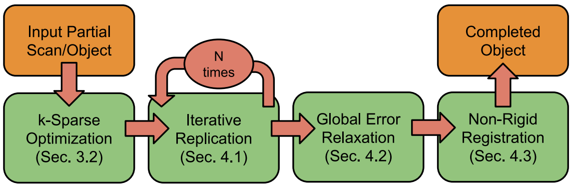

Figure 2: *An overview of our object completion pipeline. An optimization method first detects the partial symmetry with the highest score, the input object is iteratively replicated based on this symmetry, global error relaxation corrects the accumulated errors and non-rigid registration applies local deformations in order to handle approximate symmetry.*

parameter sensitivity, where k needs to be manually tweaked depending on the input dataset, in order to compute the desired symmetry.

A much more robust solution is to discard the candidate transformations T_c which are close to the optimal alignment transformation T_{opt} and satisfy the following inequality:

$$\sigma(T_c, T_{opt}) \leq \varepsilon_t,$$

where the function $\sigma(T_c, T_{opt})$ is a distance metric between two transformation matrices and ε_t is a user defined parameter. In this approach, k is not directly specified, but it is implied by the parameter ε_t, which enforces the amount of partial overlap. Various metrics can be used for the function $\sigma(T_c, T_{opt})$. The metric that we have used in our experiments computes the Euclidian norm between the two vectors that correspond to the elements of the two transformation matrices, when the size of the input object is normalized. In our experiments we have found that this approach is not sensitive to the parameter ε_t, especially when working with feature-rich objects with prominent geometric features. In fact, most examples in this paper use a value of $\varepsilon_t = 0.15$.

4. Completion Pipeline

The registration-based partial symmetry detection approach that we have described in the previous section is used as part of a larger pipeline that predicts the missing parts of an object. An overview of this pipeline is shown in Figure 2. In the remainder of this section we describe each one of the processing steps in more detail.

4.1. Iterative Replication

After detecting the partial symmetry transformation T, the input object M_0 and $T(M_0)$ are merged, in order to form a new object ($M' = M_0 \cup T(M_0)$). The main idea of our approach is that some of the missing regions in M_0 will be completed with information from $T(M_0)$. However, M' might still have some parts missing. In this case, we use an iterative completion scheme, defined by the following formula:

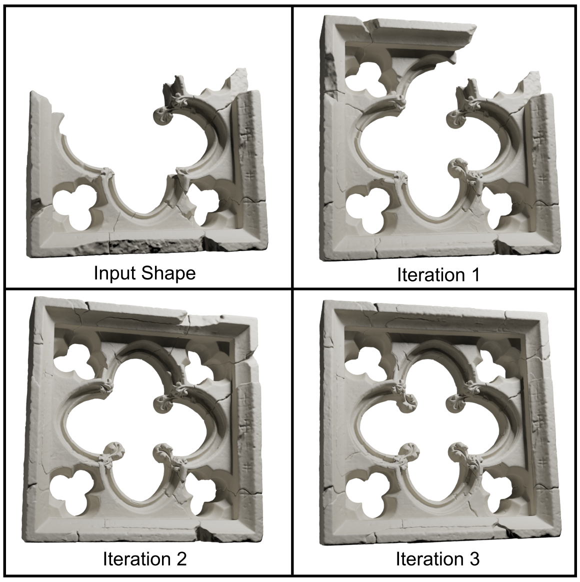

Figure 3: *Iterative completion of a partially reassembled cultural heritage object using our algorithm. Our method has properly detected the underlying discrete rotational symmetry of the embrasure shape and progressively fills the missing parts, until the method converges.*

$$M_k = M_{k-1} \cup T^k(M_0), \quad k \geq 1, \tag{3}$$

where M_k denotes the completed object after iteration k. An example of this iterative procedure is shown in Figure 3.

The algorithm stops when the prediction converges ($d(M_k, M_{k-1}) < \varepsilon$, where d is defined in Equation 1) or after a user-defined amount of iterations. Convergence can be expected only for objects that exhibit rotational or reflectional symmetries. For objects with translational symmetries, such as a segment of a fence (see Fig. 9-left), a maximum number of iterations should be set.

Using this iterative completion scheme, our algorithm needs to detect the desired partial symmetry transformation T only once, and this information is reused in all subsequent iterations. Therefore, the cost of each additional iteration is minimal.

4.2. Global Error Relaxation

The partial symmetry transformation T, which is computed with the algorithm in Section 3.2, aligns the input object M_0 with one part of the replicated object $T(M_0)$. However, this transformation might include an alignment error, which is accumulated and propagated after each iteration of our completion scheme defined in Equation 3. This can create large problems in rotationally symmetric objects, where the first

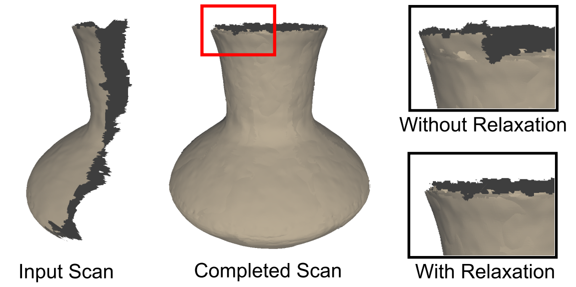

Figure 4: *Completion of a pottery vessel from a single partial scan using our approach. After multiple iterations, alignment errors are accumulated and propagated in the replicated geometry. A global error relaxation approach (multi-view ICP) is used to improve the alignment.*

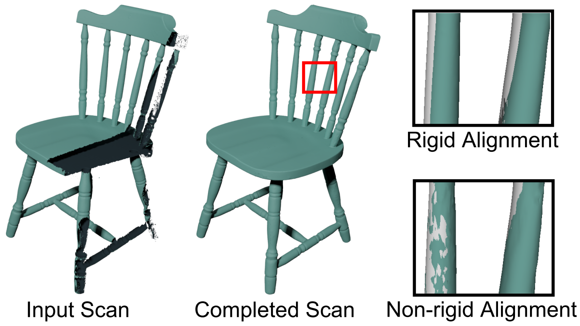

Figure 5: *Completion of a partial chair object with our method. Many objects exhibit approximate symmetries, which create alignment errors in the replicated geometry. Our approach properly handles these cases using non-rigid registration.*

and the last repeated segments might not align properly, as shown in Figure 4. For this reason, we perform a *global error relaxation* by applying a multi-view rigid *Iterative Closest Point* algorithm [CM92, BM92]. This step highly improves the alignment of the replicated geometry.

4.3. Non-Rigid Alignment

Many natural or man-made objects exhibit an approximate symmetry, due to deviations in the biological growth, imprecisions in the manufacturing process or stochastic fluctuations in physical processes. In such cases, the rigid symmetry transformation T cannot perfectly align the overlapping parts of the input object M_0 and $T(M_0)$. A typical example is shown in Figure 5, where our algorithm is used for the completion of a chair, which is not perfectly symmetric, thus the replicated geometry is not aligned correctly with the input one. To better handle these cases, we introduce a non-rigid registration method, after the global error relaxation that we have previously described. This step, deforms the replicated geometry in order to better match the input data set, improving the reliability of our approach in cases of approximate symmetry.

Various non-rigid registration methods in the bibliography can be used. In our implementation, we use an adaptation of the deformation model that was described by Pauly et al. [PMG*05], with improvements in the handling of large datasets and in the correspondence determination. To better handle large data sets (dense point clouds), the desired deformation is first computed in a low resolution version of the input surface, which is generated using uniform Poisson subsampling, and then is transferred to the full resolution data set. In order to improve the quality of the closest-point correspondences, we use an iterative rigidity relaxation scheme [LSP08], where the desired deformation is computed over multiple iterations, starting with a high rigidity value, which is then relaxed in the subsequent iterations.

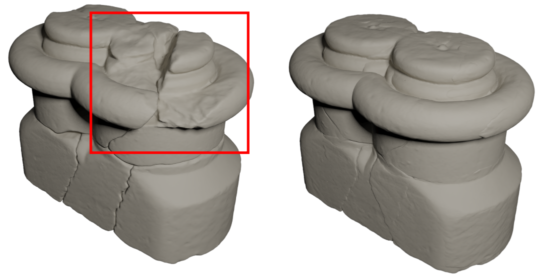

Figure 6: *Successful completion of a cultural heritage object. Left: The original object is damaged and hevily eroded, especially in the region highlighted by the red square. Right: Our method properly detects the underlying reflectional symmetry of the object and restores the damaged parts by replicating the geometry of the intact regions.*

5. Results

In this section we demonstrate our algorithm in the completion of cultural heritage objects and partial 3D scans with varying levels of partiality. We also evaluate the performance of our method for these tasks.

Since many cultural heritage objects exhibit some form of symmetry, one of the main motivations of our work was to exploit this fact in order to complete any potentially missing parts or restore heavily eroded surfaces. In this context, Figure 3 demonstrates how our algorithm can properly detect the *discrete* rotational symmetry in the partial *embrasure* shape and progressively complete the object after three iterations. Similarly, in Figure 6 our algorithm exploits the reflectional symmetry of the *column base* shape, in order to restore a heavily damaged or eroded region. These two objects were scanned from the archeological site of the Nidaros Cathedral in Trondheim. In Figure 7 we demonstrate our al-

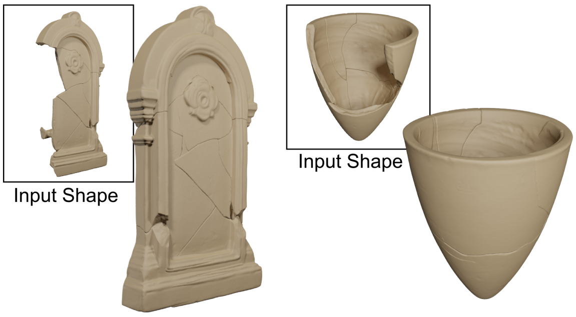

Figure 7: *Completion of partially reassembled objects. Our method handles both smooth feature-less and feature-rich shapes and successfully completes the missing parts.*

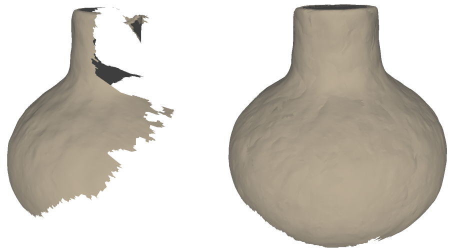

Figure 8: *Completion of a pottery vessel from a single partial scan. Even though the input scan covers only roughly 30% of the complete object, our approach detects the underlying rotational symmetry and successfully completes the missing geometry.*

gorithm in the completion of additional partially reassembled objects. These experiments indicate that our method can properly handle both feature-rich and also smooth featureless surfaces.

The reassembled objects that we examine in this section were broken into multiple parts and the datasets that we use are the direct output from the automatic fragment reassembly algorithm that was presented by Mavridis et al. [MAP15], without any preprocessing. This means that the internal fracture regions that connect the broken parts remain in the input dataset and are treated as outliers by our algorithm. This is a good demonstration of the robustness of our method. In contrast, the method proposed by Sipiran et al. [SGS14], requires a preprocessing step, in order to eliminate the internal contact surfaces and produce a manifold mesh.

It is worth noting that, when the input dataset includes damaged or eroded surfaces, we disable the non-rigid deformation stage in our pipeline. This is required, because the desired behaviour for the replicated geometry is to cover, and thus repair, any missing or eroded parts. In contrast, if the non-rigid deformation is enabled, the replicated geometry will be deformed in order to better match the underlying damaged parts, which defeats the purpose of the restoration. For this reason, the non-rigid registration is only enabled when the input dataset consists of partial 3D scans.

A similar application of our method is the completion of partial 3D scans. In Figures 4 and 8 we demonstrate the results of our algorithm in the completion of pottery vessels of various shapes from a single partial scan. The input partial scans cover only 20 to 30% of the complete object. Our algorithm properly detects the continuous rotational symmetry in the input objects and provides plausible completion results. It is worth noting that the completed pottery objects retain the details and irregularities of the input scans, which are not perfect surfaces of revolution. This completion approach for pottery objects is more general and provides more plausible results than previous methods [CM02] that extract a profile curve and a symmetry axis from the input data and predict the missing shape using surfaces of revolution. Figure 5 presents the successful completion of a chair object using reflectional symmetry. Figure 9 demonstrates our method on the completion of various partial shapes. It should be noted that, while only the fence example uses a purely translational symmetry, even the rotational and reflectional examples include a translational component in the calculated partial symmetries, because the centroid of the incomplete input is not the same as the centroid of the complete shape.

In Figure 10 we test our method in the completion of a wheel-cap object with varying degrees of partiality. When operating with fewer data, more iterations are required by our algorithm to compute the final object and as a result, the accumulated registration errors are higher. This is reflected in the RMS error measurements against the ground truth object. In practice, even when only 25% of the desirable object is provided as input, our method provides plausible completion results.

In Figure 10 we also provide a comparison with another symmetry-based completion method, proposed by Sipiran et al. [SGS14]. This approach automatically detects and uses the highest scoring symmetry plane to complete the missing part of an object. For the 25% and 50% partiality cases, this method augments the input data, but the object is still not complete. In this case, a natural choice is to apply the algorithm iteratively, until the object is completed. However, by definition the resulting object after the first iteration is perfectly symmetric, thus the highest scoring symmetry plane cannot be used to further augment the available data and the method naturally fails. The same problem also prevents the algorithm from working on the 75% partiality case. It is clear that a symmetry plane with a lower score should be used. Our approach can automatically select such plane using the *"k-sparse"* constraint, without relying on user input. Furthermore, our approach is not limited to planar symme-

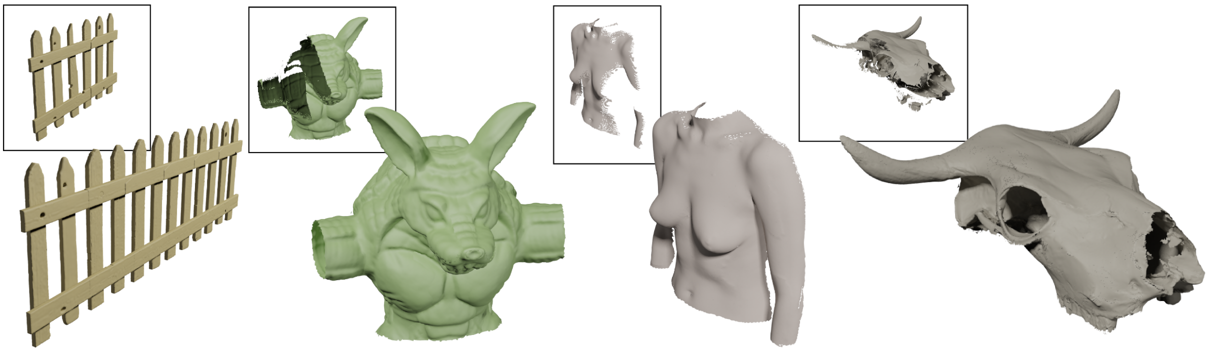

Figure 9: *Completion of various partial objects using our method. The insets show the partial input shapes. Translational symmetry is used for the fence shape (three iterations), while the other examples use reflectional symmetry.*

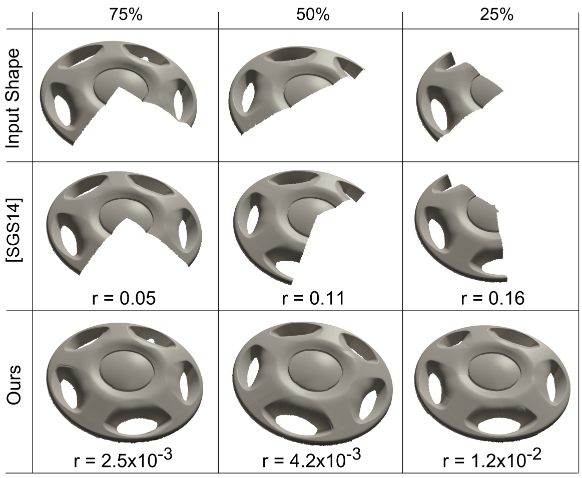

Figure 10: *Completion of a partial wheel cap object with varying degrees of partiality and comparison with [SGS14]. r denotes the Hausdorff RMS distance from the ground truth.*

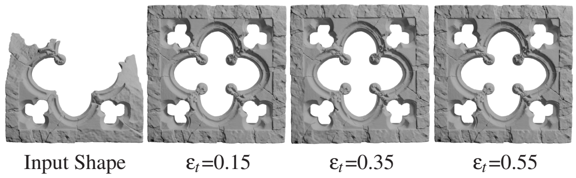

Figure 11: *When operating on objects with prominent geometric corners and features, our method is not sensitive to the parameter $\varepsilon_t \in [0,1]$. For smooth objects, the behavior is similar to the one illustrated in Figure 1.*

Dataset	# points	Congruent Sets	Run-time
Embrasure (Fig. 3)	496K	225K	47
Chair (Fig. 5)	209K	112K	64
Col-Base (Fig. 6)	260K	346K	77
Vessel (Fig. 8)	40K	56k	12

Table 1: *Runtime performance in seconds of our symmetry detection approach for datasets in this paper.*

try, and can properly perform the completion based on the rotational symmetry of the object.

Figure 11 explores the effect of parameter ε_t in the results. When operating on objects with prominent geometric corners and features, our method is not sensitive to the parameter ε_t. On the other hand, for smooth surfaces, the effect of ε_t is similar to the effect of the parameter k, as illustrated in Figure 1. As an example, when operating on smooth pottery scans, this parameter affects the number of iterations required to get the complete object.

5.1. Performance

In most application of our algorithm, the completion of partial objects can be performed offline, thus the runtime performance was not the main priority of our research. Table 1 shows the processing time for the data sets presented in this paper. While our method does not offer instantaneous completion, the overall processing times are practical. The runtime performance heavily depends on the number of points in the input dataset and the number of congruent 4-point sets that are detected during the symmetry detection. All measurements were performed on an *Intel Core i7-3820* CPU at 3.6GHz using a single processing thread.

5.2. Discussion and Limitations

Completion of the remaining holes. A common limitation of all completion methods that are based on self-similarity is that they cannot reproduce details or geometric structures that are completely missing from the input. In the example shown in Figure 12, the lower part of the vessel is completely missing from the input partial scan, thus it is also missing from the prediction that was computed with our symmetry-based completion algorithm.

To better handle such cases, our approach can be combined with traditional template-based completion, in order to fill any remaining missing parts using information from a similar template object. This requires non-rigidly deforming the template in order to match the output of the symmetry-based completion. Our system already incorporates a non-rigid registration component (Section 4.3), and this step can be performed by the same algorithm. The template object can be either provided manually or retrieved from a database of reference objects [SPS14].

When a template object is not available, a smoothness prior based method, such as Poisson reconstruction, can be used instead. Our symmetry-based completion is orthogonal to these approaches, and their combination produces a more reliable pipeline, that is able to properly handle a wider variety of input cases.

Semi-automatic completion. Equation 2 defines a global optimization problem that searches for the best (in terms of the residual distance - Equation 1) symmetry transformation that meets the specified constraints. However, it might be desirable to find more than one symmetry transformations and have the user manually select the one that is more appropriate for the specific object completion task. For example, this approach can provide more completion options to a cultural heritage expert. To this end, Super4PCS can be trivially modified to return the N best candidate transformations that meet the constraints, instead of returning only the first one, with the additional constraint that the distance between the returned transformations should be at least ε_t. All completion results in this paper were computed with the automatic approach.

Selective Replication. A limitation of our approach is that the iterative replication scheme that we have described in Equation 3 replicates the whole input, restricting our method to the completion of globally symmetric objects. However, in some cases, only parts of the objects should be replicated, e.g. to avoid repeating the handles or the spout of a vessel. To properly handle these cases of partially symmetric objects, our approach can be combined with a segmentation algorithm, in order to replicate only specific segments of the input. The selection of the segments that should be replicated can be either based on user input [GTB14], or can be determined automatically for some types of input objects [HKG11].

Feature-driven alignment. Our partial and approximate symmetry detection algorithm is based on a modified featureless registration algorithm (Section 3.2). The fact that we are not using features makes our algorithm suitable for use in smooth featureless surfaces, such as the ones often found in pottery, and also greatly simplifies the implementation. However, the use of features can potentially reduce the search space and improve the overall performance, when operating on feature-rich surfaces. Furthermore, features can further improve the alignment of the replicated geometry, when the desired local minima depend on small geometric details that can be potentially missed by our featureless voting scheme. To better handle these cases, an interesting direction of research for the future is to take advantage of features in our completion pipeline.

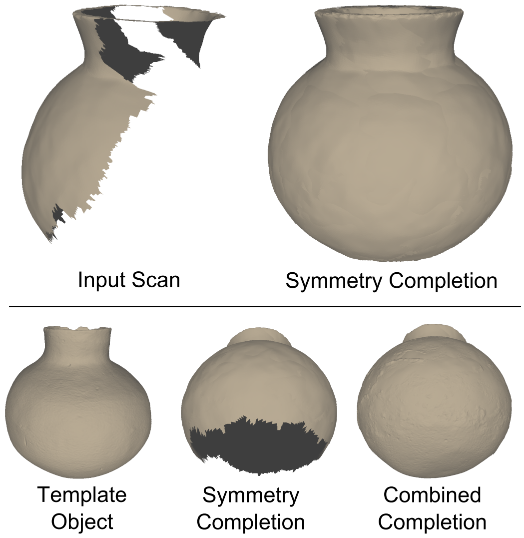

Figure 12: *Top: Completion of a pottery vessel from a single scan. Bottom: Symmetry-based completion cannot reproduce the lower part of this vessel, since it is completely missing from the input. Our method can be combined with template-based completion to fill the remaining parts.*

6. Conclusions

We have presented a new method for the completion of partial 3D objects, which is based on the detection of partial and approximate symmetries. The proposed algorithm uniformly exploits translational, rotational and reflectional symmetries and provides reliable completion results, even when more than half of the target object is missing. We have demonstrated that our approach is suitable for the completion of incomplete cultural heritage objects and partial 3D scans with varying levels of partiality.

Acknowledgements

We would like to thank the anonymous reviewers for their insightful comments and suggestions. The armadillo dataset in Figure 9 is from Stanford. The chair, body and scull scans in Figures 5 and 9 are from the Artec Group. The cultural heritage objects in Figures 3 and 6 were downloaded from the *Presious* project website. The pottery datasets in Figures 8 and 12 were provided by Athena Research Center. The template object in Figure 12 is from the Virtual Hampson Museum collection. This work was funded by EC FP7 STREP project "PRESIOUS", grant No. 600533.

References

[BAK10] BERMANIS A., AVERBUCH A., KELLER Y.: 3-D Symmetry Detection and Analysis Using the Pseudo-polar Fourier Transform. *Int. Journal of Computer Vision 90*, 2 (2010), 166–182. 2

[BM92] BESL P. J., MCKAY N. D.: Method for registration of 3-D shapes. In *Robotics-DL tentative* (1992), International Society for Optics and Photonics, pp. 586–606. 5

[CM92] CHEN Y., MEDIONI G.: Object modelling by registration of multiple range images. *Image and vision computing 10*, 3 (1992), 145–155. 5

[CM02] CAO Y., MUMFORD D.: Geometric structure estimation of axially symmetric pots from small fragments. In *In Proc. Signal Processing, Pattern Recognition, and Applications* (2002). 6

[ELF97] EGGERT D. W., LORUSSO A., FISHER R. B.: Estimating 3-D rigid body transformations: a comparison of four major algorithms. *Machine Vision and Applications 9*, 5-6 (1997), 272–290. 3

[GTB14] GUY E., THIERY J.-M., BOUBEKEUR T.: Simselect: similarity-based selection for 3d surfaces. *Computer Graphics Forum (Proc. EUROGRAPHICS 2014) 33*, 2 (2014), 165. 8

[HFG*06] HUANG Q.-X., FLÖRY S., GELFAND N., HOFER M., POTTMANN H.: Reassembling Fractured Objects by Geometric Matching. *ACM Trans. Graph. 25*, 3 (July 2006), 569–578. 1

[HGCO*12] HUANG H., GONG M., COHEN-OR D., OUYANG Y., TAN F., ZHANG H.: Field-guided registration for feature-conforming shape composition. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2012) 31* (2012), 171:1–171:11. 2

[HKG11] HUANG Q., KOLTUN V., GUIBAS L.: Joint shape segmentation with linear programming. *ACM Trans. Graph. 30*, 6 (Dec. 2011), 125:1–125:12. 8

[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. *Computer Graphics Forum 33*, 5 (2014), 45–54. 2

[JXCZ13] JIANG W., XU K., CHENG Z.-Q., ZHANG H.: Skeleton-based intrinsic symmetry detection on point clouds. *Graphical Models 75*, 4 (2013), 177–188. 2

[KFR04] KAZHDAN M., FUNKHOUSER T., RUSINKIEWICZ S.: Symmetry Descriptors and 3D Shape Matching. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing* (2004), SGP '04, pp. 115–123. 2

[KKP13] KAKARALA R., KALIAMOORTHI P., PREMACHANDRAN V.: Three-Dimensional Bilateral Symmetry Plane Estimation in the Phase Domain. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition* (2013), CVPR '13, IEEE Computer Society, pp. 249–256. 2

[KLAB15] KORMAN S., LITMAN R., AVIDAN S., BRONSTEIN A.: Probably approximately symmetric: Fast rigid symmetry detection with global guarantees. *Computer Graphics Forum 34*, 1 (2015), 2–13. 2

[LAGP09] LI H., ADAMS B., GUIBAS L. J., PAULY M.: Robust single-view geometry and motion reconstruction. *ACM Trans. Graph. 28*, 5 (Dec. 2009), 175:1–175:10. 2

[LCDF10] LIPMAN Y., CHEN X., DAUBECHIES I., FUNKHOUSER T.: Symmetry Factored Embedding and Distance. *ACM Trans. Graph. 29*, 4 (July 2010), 103:1–103:12. 2

[LSP08] LI H., SUMNER R. W., PAULY M.: Global correspondence optimization for non-rigid registration of depth scans. *Computer Graphics Forum (Proc. SGP'08) 27*, 5 (July 2008). 5

[MAM14] MELLADO N., AIGER D., MITRA N. J.: Super 4pcs fast global pointcloud registration via smart indexing. *Computer Graphics Forum 33*, 5 (2014), 205–215. 3

[MAP15] MAVRIDIS P., ANDREADIS A., PAPAIOANNOU G.: Fractured object reassembly via robust surface registration. In *Eurographics 2015 Short Papers* (2015), Eurographics 2015. 1, 6

[MGP06] MITRA N. J., GUIBAS L. J., PAULY M.: Partial and Approximate Symmetry Detection for 3D Geometry. *ACM Trans. Graph. 25*, 3 (July 2006), 560–568. 2

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

[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, 5

[SGS14] SIPIRAN I., GREGOR R., SCHRECK T.: Approximate symmetry detection in partial 3d meshes. *Computer Graphics Forum (proc. Pacific Graphics) 33* (2014), 131–140. 2, 6, 7

[SOG09] SUN J., OVSJANIKOV M., GUIBAS L. J.: A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. *Comput. Graph. Forum 28*, 5 (2009). 2

[SPS14] SAVELONAS M. A., PRATIKAKIS I., SFIKAS K.: An overview of partial 3d object retrieval methodologies. *Multimedia Tools and Applications, Springer* (2014), 1–26. 8

[THW*14] TEVS A., HUANG Q., WAND M., SEIDEL H.-P., GUIBAS L.: Relating shapes via Geometric Symmetries and Regularities. *ACM Trans. Graph. 33*, 4 (2014). 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

[Ume91] UMEYAMA S.: Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on pattern analysis and machine intelligence 13*, 4 (1991), 376–380. 3

[XZJ*12] XU K., ZHANG H., JIANG W., DYER R., CHENG Z., LIU L., CHEN B.: Multi-scale Partial Intrinsic Symmetry Detection. *ACM Trans. Graph. 31*, 6 (Nov. 2012), 181:1–181:11. 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 (2009), 138:1–138:10. 2

[ZPA95] ZABRODSKY H., PELEG S., AVNIR D.: Symmetry as a continuous feature. *Pattern Analysis and Machine Intelligence, IEEE Transactions on 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