

Approximate Symmetry Detection in Partial 3D Meshes

Ivan Sipiran¹ and Robert Gregor¹ and Tobias Schreck¹

¹Department of Computer and Information Sciences, University of Konstanz, Konstanz, Germany

Abstract

Symmetry is a common characteristic in natural and man-made objects. Its ubiquitous nature can be exploited to facilitate the analysis and processing of computational representations of real objects. In particular, in computer graphics, the detection of symmetries in 3D geometry has enabled a number of applications in modeling and reconstruction. However, the problem of symmetry detection in incomplete geometry remains a challenging task. In this paper, we propose a vote-based approach to detect symmetry in 3D shapes, with special interest in models with large missing parts. Our algorithm generates a set of candidate symmetries by matching local maxima of a surface function based on the heat diffusion in local domains, which guarantee robustness to missing data. In order to deal with local perturbations, we propose a multi-scale surface function that is useful to select a set of distinctive points over which the approximate symmetries are defined. In addition, we introduce a vote-based scheme that is aware of the partiality, and therefore reduces the number of false positive votes for the candidate symmetries. We show the effectiveness of our method in a varied set of 3D shapes and different levels of partiality. Furthermore, we show the applicability of our algorithm in the repair and completion of challenging reassembled objects in the context of cultural heritage.

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

In recent years, the study of symmetries has attracted the attention of the computer graphics and computer vision communities. It is mainly due to the fact that the symmetry may facilitate the computational understanding of visual representations, and hence it could give support to high level tasks such as recognition. In particular, the detection of symmetries in 3D geometry has a special interest since it is the first step in the analysis process, and therefore it needs to be conducted effectively and efficiently.

There are many proposals to tackle the problem of symmetry detection under the assumption that the complete object is available for analysis. This assumption is strong and may not hold for real applications in which a certain degree of missing geometry may be evident. For example in archaeology, scanned pieces could have been under deterioration for long time, or simply some portions are missing. Therefore, it is imperative to devise algorithms for the detection of symmetries taking into account the possible partiality of the input object.

Clearly, the problem of detection of symmetries in partial

data is not trivial. The main difficulty is that in presence of partiality, the center of mass of the input shape does not correspond to the center of mass of the real complete object. Additionally, we do not know a priori the amount of geometry that is missing, making the problem very hard. In this paper, we propose a method to detect symmetries in shapes with missing geometry. Our algorithm is based on a careful design of two procedures that jointly work to detect approximate global symmetries. The first component consists of the definition of a new surface function based on a heat diffusion process. This function characterizes local protrusions and allows us to select reliable points to find symmetric correspondences. The second component is a partiality-aware vote-based approach to validate the initial symmetric correspondences over which a final set of symmetries is detected. Moreover, we evaluate the robustness of our method in a challenging application: the symmetry-based repair and completion of cultural heritage objects.

The paper is organized as follows. Section 1.1 lists the contributions of our research. Section 2 discusses the literature related to our work. Section 3 presents an overview of our proposal. Section 4 defines our diffusion-based function

and discusses its behavior. Section 5 presents our partiality-aware vote scheme for validation of symmetries. Section 6 details our experiments and results, and also presents an application of our algorithm to repair and completion of cultural heritage objects. Finally, Section 7 draws our conclusions.

1.1. Contributions

The contributions of our paper can be summarized as follows:

- We propose a new diffusion-based feature that conveys a balanced characterization of globality and locality. Globality is necessary to find symmetric correspondences and locality is important to be robust to partial data. Potential symmetry planes are defined from local maxima of the new feature function.
- We propose a partiality-aware voting algorithm to validate the potential symmetry planes. Our algorithm computes a support for a plane in which the votes are expected to be consistent. Also, we propose a set of rules for discarding inconsistent votes.
- We present a novel application in the context of cultural heritage. We develop an algorithm to synthesize missing parts of reassembled objects using the knowledge obtained for our symmetry detection method.

2. Related Works

The study of symmetries has proliferated considerably in recent years in the computer graphics community, giving as consequence many proposals to tackle the problem of detection. A comprehensive enumeration of these techniques is out of the scope of this paper. Instead, we focus our discussion on methods related to global reflectional symmetries. For readers interested in an overall panorama on 3D symmetry, we recommend the survey by Mitra et al. [MPWC13].

A common approach is the evaluation of symmetries as a measure function. In their seminal work, Zabrodsky et al. [ZPA95] proposed a symmetry measure based on the exhaustive search in the rotational transformation space. In this direction, Kazhdan et al. [KFR04] formulated an efficient procedure to evaluate the symmetry measure using spherical harmonic coefficients. Similarly, Podolak et al. [PSG*06] devised a Monte-Carlo algorithm to sample points in which the symmetry measure evaluation is performed. Subsequently, a refinement step can be applied to get higher precision in transformations that are local maxima. Alike, Martinet et al. [MSHS06] used spherical harmonics to evaluate a symmetry measure based on shape moments. Also, Bermanis et al. [BAK10] proposed to evaluate the symmetry by comparing volumes through an angular difference function which can be efficiently computed using the 3D Pseudo-polar Fourier Transform. Analogously, Kakarala

et al. [KKP13] described an optimization method that detects symmetry planes which can be interpreted as a linear phase in the spherical harmonic domain. More recently, Korman et al. [KLAB14] proved an upper bound on the sampling density of the transformation group which guarantees to find symmetries with high probability. An obvious limitation of all these techniques is the assumption that the center of mass of the object is a fixed point of the transformation space. Nevertheless, this is not the case when the object is incomplete, where the center of mass is to be shifted.

Another interesting approach to detect symmetries is the vote-based scheme as proposed by Mitra et al. [MGP06]. The idea behind this approach is the use of self-correspondences that may convey evidence about the existence of symmetries. For instance, Mitra et al. [MGP06] proposed to match points with high curvature. Each pair of points induces a transformation which is stored as a high-dimensional point. A final clustering step in this space can provide the more frequent transformations which derive into symmetries. Similarly, Lipman et al. [LCDF10] developed a voting algorithm that selects the best transformation as possible for random pairs of points. The best alignment is stored in a symmetric correspondence matrix and its spectral decomposition is exploited to detect symmetries. Analogously, Xu et al. [XZJ*12] defined an intrinsic symmetry correspondence matrix. A multi-scale analysis over the scale of voting pairs of points is formulated in order to be aware of the scale of the symmetries. More recently, Tevs et al. [THW*14] proposed to quantify the relation between shapes based on the regularities of symmetric parts in 3D objects. The first step is a decomposition for a shape into a set of regions. Then, a graph is used to represent the relation between the regions in terms of symmetric transformations. The authors proposed an algorithm to match two graphs which provides correspondences between the regularities of given shapes.

Dealing with partial data. A few techniques have focused the attention in the problem of symmetry detection with missing data. Xu et al. [XZT*09] developed a voting algorithm to detect the intrinsic reflectional symmetry axis. Using the axis as a hint, a completion algorithm for missing geometry was shown. On the other hand, Jiang et al. [JXCZ13] proposed an algorithm to find intrinsic symmetries in point clouds by using the curve skeleton. A set of filters produces a good set of symmetric correspondences which are finally verified with spectral analysis. Although these two proposals showed results for partial data, the amount of missing geometry is small. In contrast, our approach is designed keeping in mind the existence of large missing parts, turning more challenging the problem to solve.

3. Overview

The core of our paper is the detection of symmetries in shapes with missing geometry. Before going into the details, our intention is to provide a clear scope for our method.

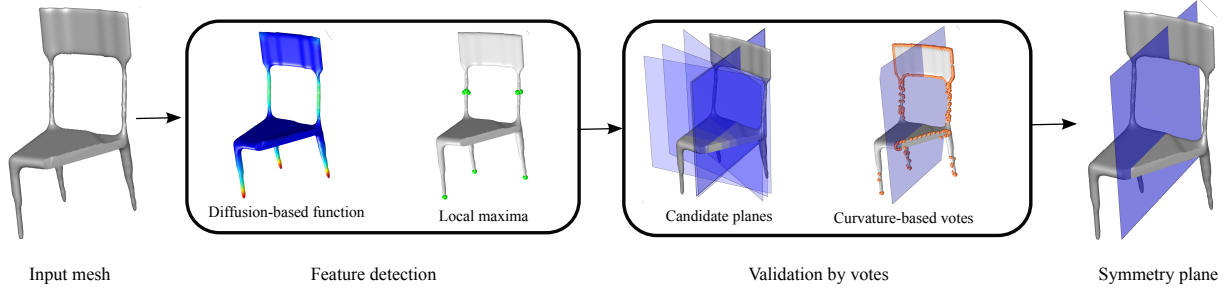


Figure 1: Our algorithm consists of two steps. First, a function is computed on the mesh such that the local maxima of this function are potential symmetric correspondences. Second, a set of candidate symmetry planes are generated and a further voting system determines the planes with the best chances to be symmetry planes. Both steps are designed to deal with the problem of missing geometry, giving a robust method as result.

According to the classification provided in [MPWC13], our method is intended to detect *global*, *approximate* and *extrinsic* symmetries. It is global because the entire object is involved in the symmetry evaluation. It is approximated in the sense that the detected symmetries are not guaranteed to be the exact symmetries. And it is extrinsic because we only focus in rigid objects. Moreover, in light of these observations, we characterize the symmetries through planes in the 3D space. That is, in this paper we only consider reflectional symmetries. Therefore, our problem can be stated as the detection of reflectional symmetry planes.

The main problem is that we cannot assume that the center of mass of a given object lies in a symmetry plane due to the missing geometry. Obviously, this is the reason about why to use a vote-based approach in our context. However, the missing geometry complicates the voting process because a point sampling based on simple local features (such as curvature, for instance) could not cover the surface properly.

Our method operates in two stages that try to attenuate the limitations imposed by the partiality. First, we aim at detecting a few reliable features that are effective to find symmetric correspondences and robust to missing data at the same time. We propose a new feature based on the heat diffusion on manifolds which will be proven to be effective to define local diffusion behaviors. More interestingly, pairs of matched points could give rise to potential symmetry planes.

Second, the potential symmetries require a validation in order to compute the final set of planes. We propose a vote-based algorithm that is aware of the partiality. For a potential symmetry plane, we analyze the influence of the partiality in the validation process. We develop the concept of *symmetry support* for a plane in order to look for evidence for this plane to be a symmetry. A graphic overview of our method can be seen in Fig. 1.

4. Multi-scale Local Features

The first step of our algorithm is the detection of robust local features. In order to use these features for symmetry detection of partial data, we must bear in mind two important aspects. First, the features need to preserve a certain degree of global awareness in order to enhance the matching of symmetric points. It would reduce the number of false positive matches, and therefore also the number of candidate symmetry planes. Second, the features need to characterize the local geometry in a local enough way so that it is possible to deal with the missing data. Although both requirements seem to be conflicting, we believe it is possible to formulate a balanced characterization to take advantage of both aspects.

Our formulation is built on top of the theory of heat diffusion on manifolds. The heat diffusion process over a compact manifold \mathcal{M} , possibly with boundary, is governed by the heat equation

$$\Delta_{\mathcal{M}} u(x, t) = -\frac{\partial u(x, t)}{\partial t} \quad (1)$$

where $\Delta_{\mathcal{M}}$ is the Laplace-Beltrami operator of \mathcal{M} and $u(\cdot, t)$ is the heat distribution over \mathcal{M} in time t .

The fundamental solution of the heat equation is $K_t(x, y)$ called the heat kernel. This represents a solution with a point heat source in x and can be considered as the amount of heat transferred from x to y at time t supposing that the heat source is x . For compact manifolds, the heat kernel can be expressed using the eigenvalues and eigenvectors of the Laplace-Beltrami operator as follows:

$$K_t(x, y) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \vec{v}_i(x) \vec{v}_i(y) \quad (2)$$

where λ_i is the i -th eigenvalue and $\vec{v}_i(\cdot)$ is the i -th eigenvector's entry corresponding to a given point.

It is well known that the heat kernel is isometric invariant, multi-scale, and stable against perturbations on the surface. In addition, restricting the heat kernel to the temporal domain and fixing the spatial variables, we can obtain a representation for each point on the manifold [SOG09]:

$$h(x, t) = K_t(x, x) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \vec{v}_i(x)^2 \quad (3)$$

The value $h(x, t)$ can be interpreted as the amount of heat that remains on the point x after time t . This interpretation suggests that points with high values of h for large values of t can be considered as keypoints. However, the use of high values of t is associated with a global behavior, which is not desired in our context. On the other hand, we could conversely evaluate $h(x, t)$ in low values of t , but this would make it sensitive to local perturbations. In order to give a balanced characterization and robustness, we propose a new function $\mathcal{H}(x, t) : \mathcal{M} \times \mathbb{R}^+ \rightarrow \mathbb{R}$ on the surface as follows

$$\mathcal{H}(x, t) = \int_0^t h(x, t) dt \quad (4)$$

The function $\mathcal{H}(x, t)$ associates the accumulation of heat up to time t to each point in the surface x . The intuition behind our formulation is that for low values of t , the accumulation in different times provides robustness to local perturbations, while still providing a global behavior. More formally, the evaluation of the integral in 4 yields

$$\begin{aligned} \mathcal{H}(x, t) &= \int_0^t h(x, t) dt \\ &= \int_0^t \left(\sum_{i=0}^{\infty} e^{-\lambda_i t} \vec{v}_i(x)^2 \right) dt \\ &= \sum_{i=0}^{\infty} \int_0^t e^{-\lambda_i t} \vec{v}_i(x)^2 dt \\ &= \sum_{i=0}^{\infty} \left(\int_0^t e^{-\lambda_i t} dt \right) \vec{v}_i(x)^2 \\ &= \sum_{i=0}^{\infty} \left(\frac{1 - e^{-\lambda_i t}}{\lambda_i} \right) \vec{v}_i(x)^2 \end{aligned} \quad (5)$$

The final result can be expressed in a simplified form as

$$\mathcal{H}_t(x) = \sum_{i=0}^{\infty} f(\lambda_i, t) \vec{v}_i(x)^2. \quad (6)$$

Note that Eq. 6 (in its general form as a representation of filters on eigenvectors) has been already reported by Litman and Bronstein [LB14] as a general formulation for spectral descriptors. In this notation, it is easy to identify the function

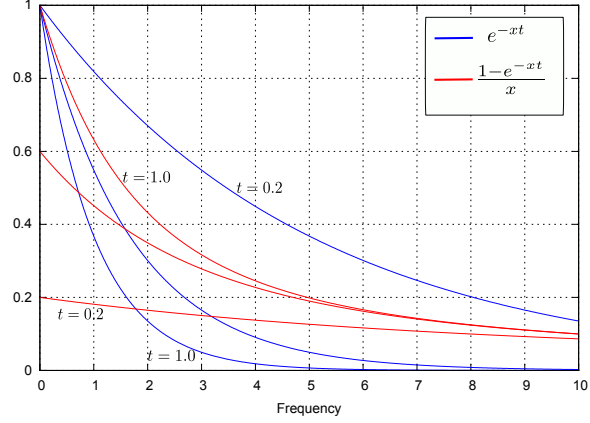


Figure 2: Kernel examples used in the computation of HKS and our feature \mathcal{H} . As t increases, the exponential decay of a low-pass filter is evident. As a result, the domain with relevant output values is considerably reduced, which implies a more global awareness. In contrast, our kernel shows a better distribution of values in a more broad domain. We believe this represents a good trade-off between global and local information.

$f(\dots)$ as a filter that weights the contributions of the shape frequencies (eigenvectors of the Laplace-Beltrami operator). For example the Heat Kernel Signatures [SOG09] use a low-pass filter and the Wave Kernel Signatures [ASC11] use a band-pass filter. Our proposal can be considered as a new filter that guarantees a balanced behavior of globality and locality.

In our case, $\mathcal{H}(x, t)$ has a different behavior. The amplitude of the function f changes as t changes. In contrast to the low-pass and band-pass filters in HKS and WKS, our filter penalizes low frequencies for low values of t , defining an upper bound over all possible values for f . The next proposition summarizes this point (see Appendix for a proof).

Proposition 1 Let \mathcal{M} be a compact Riemannian manifold and $\Delta_{\mathcal{M}}$ be its Laplace-Beltrami operator. Let $0 = \lambda_0 \leq \lambda_1 \leq \lambda_2 \leq \dots$ be the eigenvalues of $\Delta_{\mathcal{M}}$. Then

$$\frac{1 - e^{-\lambda_i t}}{\lambda_i} < t \quad (7)$$

for any $t > 0$ and $\lambda_i > 0$.

The most important point to remark is that t controls the influence of the eigenfunctions in the computation of $\mathcal{H}(x, t)$ according to the magnitude of their corresponding eigenvalues. Figure 2 illustrates the behavior of our proposed filter compared to the low-pass filter used in the computation of HKS. As comparison, we can refer to the curves at $t = 0.2$, where the attenuation of our filter is notorious in small frequencies compared to the values for a low-pass filter. In addition, our filter presents a better distribution of values in

all frequencies, compensating the exponential decay in the case of a low-pass filter. This can be traduced in a more balanced behavior with a good trade-off between globality and locality. This is also the explication of why our feature is multi-scale.

Figure 3 shows the result of computing our feature. The function $\mathcal{H}(x, t)$ is useful to identify local maxima which correspond to points with a high accumulation of local heat. In our example, it is clear that these points can be used to evaluate the symmetry, even when a considerable part of the geometry could be missing. Also as an example, note how our method is able to detect repeatable local structures, which is a desired property to take into account in matching in general.

For our symmetry detection algorithm, once we have computed the function $\mathcal{H}(x, t)$ for every point in the mesh, we select the local maxima in the 2-ring neighborhood. For all our experiments, t is set to 0.01 of the surface area.

5. Vote-based Symmetry Detection

The second step of our algorithm is the detection of symmetry planes. Let $S = \{s_1, s_2, \dots, s_n\}$ be the set of points detected in the previous stage. Points in S could be evidence of the existence of symmetries, so the matching between them is useful to generate good symmetry hypotheses. For this reason, we consider pairs of points in S as generators of our candidate planes. Let s_i and s_j be two points in S with $i \neq j$. These points define a unique plane as the set of equidistant points to both points. We define a plane with the coefficients of its general form, that is $P_{ij} = [a \ b \ c \ d]$, where the normal vector $\vec{N}_{P_{ij}} = (a, b, c) = (s_i - s_j) / \|s_i - s_j\|_2$. Furthermore, the value d can be obtained as $d = -\langle n, (s_i + s_j)/2 \rangle$ since the average point between s_i and s_j belongs to P_{ij} . It is worth to mention that the plane is considered as a candidate if and only if the generating points have a similar feature $\mathcal{H}(x, t)$. That is, P_{ij} is a candidate plane if $|\mathcal{H}(s_i, t) - \mathcal{H}(s_j, t)| \leq \beta$.

An important detail about the computation of the symmetry plane is the orientation of the plane. Given a pair of points, there are two planes (with opposite directions) that represent the bisector plane. In order to deal with this, we agree to always take a pair s_i and s_j such that $i < j$.

The next step consist of accumulating more evidence for each candidate plane in order to validate whether it is a real symmetry plane or not. For this purpose, we use the curvature in each point on the surface to produce a voting system. In this point, it is necessary to use the information of the candidate plane as well as taking into account that the input shape can have missing data. More specifically, if we assume P_{ij} as a symmetry plane of a partial mesh, it is expected that the amount of geometry on either side of the plane differs considerably. For this reason, we propose to compute a support for the analyzed candidate plane. The support consist of two planes parallel and equidistant to P_{ij}

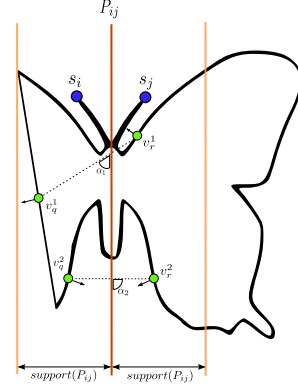


Figure 4: A pair of generating points s_i and s_j defines a bisector P_{ij} . The plane support guarantees the search of valid points for the voting scheme. In addition, voting points should preserve certain properties (as v_q^2 and v_r^2): consistent distance to P_{ij} , consistent normals and angle to bisector. Clearly, the pair formed by v_q^1 and v_r^1 do not hold these properties.

such that the part with less geometry is completely contained in this support. To compute the support, it is enough to compute the distance to the parallel planes and use it as reference to the following steps. Having the plane coefficients $P_{ij} = [a \ b \ c \ d]$, the signed distance of a point $v \in \mathcal{M}$ is defined as $dist_{P_{ij}}(v) = av_x + bv_y + cv_z + d$. The plane support is finally computed as

$$support(P_{ij}) = \min(|\max_{v \in \mathcal{M}} dist_{P_{ij}}(v)|, |\min_{v \in \mathcal{M}} dist_{P_{ij}}(v)|) \quad (8)$$

The support is important to filter the points that will be used to vote for a plane. Our selection of samples to vote is based on a simple rule that uses the curvature and the support. Given a plane P_{ij} , the samples to vote for it are the vertices v on the mesh such that $(K_1(v) + K_2(v))/2 > \delta$ and $|dist_{P_{ij}}(v)| < support(P_{ij})$, where K_1 and K_2 correspond to the maximum and minimum curvature, respectively. In all our experiments, we set δ to 0.01.

Subsequently, every pair of sample points is tested for evidence of the plane P_{ij} . Let us illustrate the ideas behind these tests before formally defining them. Figure 4 shows a 2D version of our problem. After generating the potential plane and its support, we are interested in evaluating certain characteristics of the voting pairs. For example, some interesting characteristics would be the consistent distance to the plane, the normal consistency, the angle formed by the plane normal and the segment between the voting pair, just to name a few. For example, in Figure 4, the lower pair fulfills all the aforementioned characteristics, and therefore this pair is a good evidence that the plane could coincide with a sym-

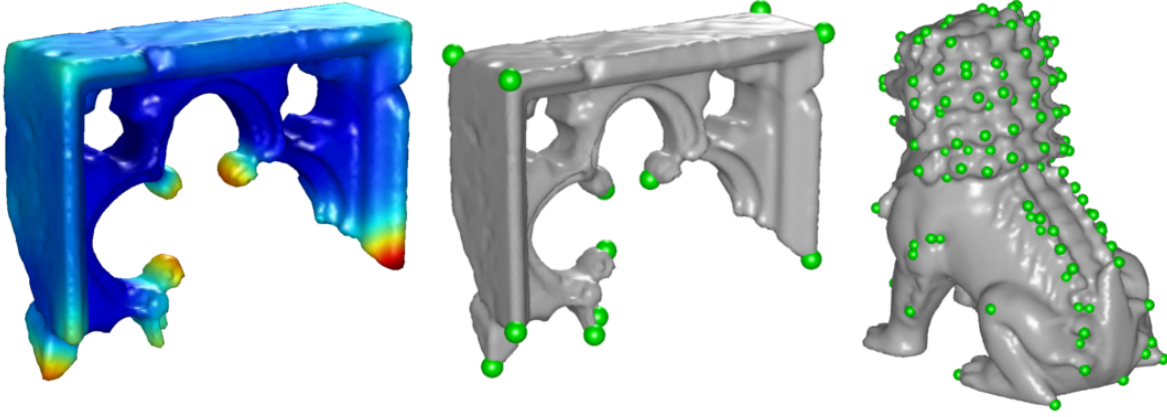


Figure 3: $\mathcal{H}(x, t)$ and its local maxima. Left: The feature values increase as the color goes from blue to yellow and to red. Center: the local maxima are potential references of symmetries, even with missing geometry. Right: local repeatable patterns could be useful in matching tasks.

metry plane. In contrast, the upper pair is not consistent in distance, normals and angles, hence it can be rejected from the analysis for the evaluated plane.

In our algorithm, given two sample points v_q and v_r , we test the following criteria in order to ensure that the pair is a real evidence:

- **Both sides of the plane.** Points should be in different sides of the plane. That is, the pair is rejected if

$$\text{sign}(\text{dist}_{P_{ij}}(v_q)) \neq \text{sign}(\text{dist}_{P_{ij}}(v_r)).$$

- **Coherence on plane distance.** Points should be approximately at the same distance from the plane. That is, the pair is rejected if

$$||\text{dist}_{P_{ij}}(v_q)| - |\text{dist}_{P_{ij}}(v_r)|| > \theta.$$

- **Orientation to plane.** Points should be disposed in such way that they form a line orthogonal to the plane as much as possible. That is, the pair is rejected if

$$(1 - |\langle \vec{N}_{P_{ij}}, (v_q - v_r) / \|v_q - v_r\|_2 \rangle|) > \alpha.$$

- **Consistency of normals.** Normals of points should also be symmetric. Here we only consider the consistency of directions. For instance, if the normal of point v_q is outward the plane P_{ij} , the normal of v_r should also be outward P_{ij} . Therefore, the pair is rejected if

$$\text{sign}(\langle \vec{N}_{P_{ij}}, \vec{N}_{v_q} \rangle) = \text{sign}(\langle \vec{N}_{P_{ij}}, \vec{N}_{v_r} \rangle).$$

- **Geometric similarity.** Points should have a similar characterization. That is, points are rejected if

$$|\mathcal{H}(v_q, t) - \mathcal{H}(v_r, t)| > \rho.$$

If a pair of samples v_q and v_r passes all test for a plane P_{ij} , the consistency of the tests can be used to compute a weight

which can be accumulated as evidence of symmetry for P_{ij} ($w(P_{ij}) = 0$ before running the algorithm). Formally,

$$w(P_{ij}) = w(P_{ij}) + w_d \times w_n \times w_s \quad (9)$$

where w_d , w_n and w_s are respectively related to distance, normal, and similarity consistency, and are defined as

$$w_d = \exp(-(|\text{dist}_{P_{ij}}(v_q)| - |\text{dist}_{P_{ij}}(v_r)|)^2)$$

$$w_n = \exp((|\langle \vec{N}_{P_{ij}}, (v_q - v_r) / \|v_q - v_r\|_2 \rangle| - 1)^2)$$

$$w_s = \exp(-(\mathcal{H}(v_q, t) - \mathcal{H}(v_r, t))^2).$$

Once the votes have been computed, we can use them to determine the plane or a set of planes depending on the application. In our results, we specify the selection rule for each shown example. Furthermore, in all our experiments, the thresholds for the test are set to $\theta = 0.1 \times$ the diagonal of the bounding box, $\alpha = \pi/18$ and $\rho = 0.4$.

6. Results and Applications

In this section, we present the results of our method and its application in the context of cultural heritage.

All our experiments were performed on watertight triangular meshes. We compute the Laplace-Beltrami operator using the approach of Belkin et al. [BSW08] and use only 300 eigenvalues and eigenvectors. For the computation of $\mathcal{H}(x, t)$, we discard eigenvalues equal to zero and their corresponding eigenvectors. For the computation of principal

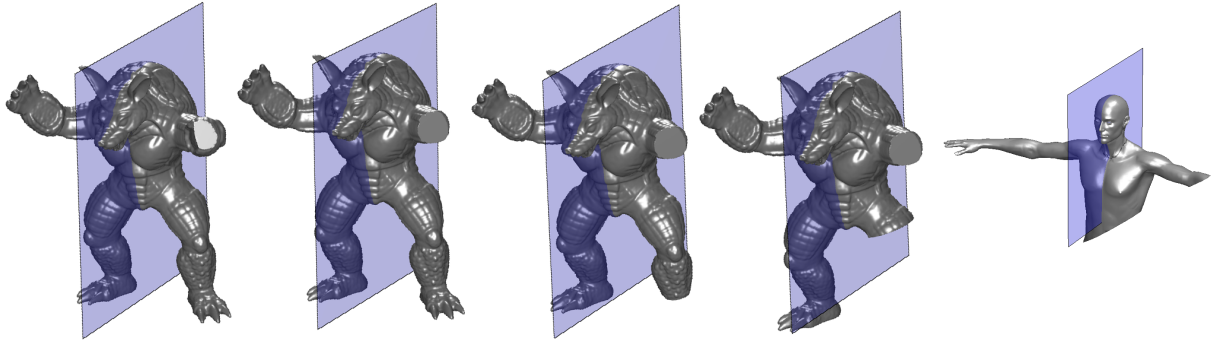


Figure 5: Symmetry planes detected in varying degrees of missing geometry.

curvatures, we use the method of normal cycles proposed in [CSM03].

Our first result shows the robustness of our method to missing geometry. Figure 5 depicts the detection of the symmetry plane in shapes with varying degree of missing data. In the case of the armadillo, we performed a boolean operation to gradually remove some parts. Interestingly, the detected symmetry plane remains approximately unchanged, despite the degree of partiality. On the other hand, for the human shape we removed approximately half of the shape with a boolean operation as well. For this example, our feature is still able to detect reliable local points in the head and ears, which leads to a good detection after the voting algorithm. For the results in Fig. 5, we selected the plane with the highest accumulated weight (see Eq. 9). Also, in order to test the reliability of our method to very large missing parts, we removed a large part of the Armadillo model and detect the symmetry plane (the result can be seen in Fig. 6). Note that we have removed geometry in both sides of the symmetry plane. This result demonstrates that our method can reliably combine the local feature detection and the support computation to detect symmetry planes in severe damaged objects.

An interesting and particular case is shown in Figure 7. The cake shape corresponds to a reassembled object from pieces, so there is missing data specially in regions with fractures. For this experiment, we applied a Poisson reconstruction [KBH06] to the aligned fragments in order to obtain a watertight mesh. Moreover, the object presents a rotational symmetry, which means that there are many symmetry planes. In this case, we noted that every potential symmetry plane obtained a high weight, and hence we decided to select the planes with a weight above 0.8 of the maximum weight. As can be seen, for similar cases, our method could be used to detect symmetry axes (the intersection of planes) as well.

More results are presented in Figure 8. The objects in this experiment were also obtained by a Poisson reconstruction of reassembled pieces. For the first example (gargoyle), we aimed at evaluating whether our algorithm is robust to the

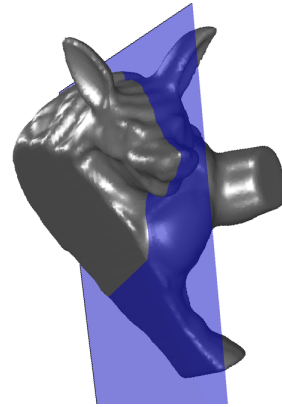


Figure 6: Our algorithm is robust to missing geometry in both sides of the symmetry plane. It takes advantage of the local features still present in the model.

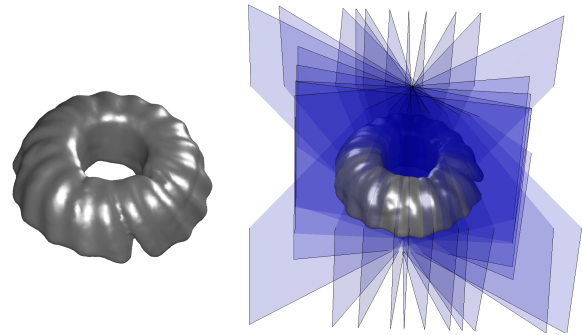


Figure 7: Left: a cake shape with many reflectional symmetry planes. Right: planes detected with our method which nearly share the symmetry axis of the shape.




Figure 8: Left and center: planes detected with our algorithm in shapes with varying level of detail. Right: the approximate main symmetry plane is detected in the head which has missing geometry and small local patterns.

level of detail. The low resolution mesh was obtained with a Poisson reconstruction in 6 octree levels and the high resolution was obtained with 10 octree levels. We can observe that our method was able to detect the approximate symmetry plane regardless of the resolution of the model. This fact can be attributed to our feature detection algorithm, which is still useful to detect local protrusions in low resolution models due to the use of accumulation of local heat distribution. On the other hand, the head model is a challenging shape because it contains smooth regions as well as a very local repeatable pattern. In addition, it also presents missing geometry which can be observed at the top of the head. Still, our method was able to handle the surface variations, mainly due to the robust detection of the pattern which induces the symmetry.

6.1. Completion of Cultural Heritage Objects

Our previous experiments have shown the effectiveness of our algorithm to detect symmetry planes. In this section, we present a challenging application that can take advantage of our results. The idea is to use the detected approximate symmetry planes to repair or complete (in some degree) the missing geometry of scanned cultural heritage objects. This scenario is useful to eventually repair the real object using the computational model as guidance or use recent advances in 3D printing to produce a repaired replica for exhibition.

The data obtained for this application belongs to the Nidaros Cathedral in Trondheim, Norway. Each piece was individually scanned and subsequently the fragments were aligned using a reassembly algorithm (similar to [HFG*06]). Next, we performed a Poisson reconstruction to obtain a watertight mesh that represents the reassembled object.

After detecting the approximate symmetry planes, it is possible to do the reflection transformation to synthesize

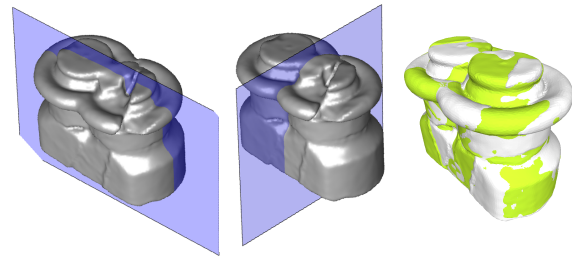


Figure 9: Column base with missing geometry due to fractures. The overlap between the original shape and the symmetric one shows the benefit of using our method to repair CH objects.

missing geometry. However, since the symmetries are approximate, the reflection does not necessarily match the object with itself. For this reason, a final local registration (an iterative re-weighting variant [RL01]) step is performed to completely align the object with its symmetric counterpart.

Figure 9 depicts a column base that contains two well-defined symmetry planes. In this case, we use the rule about preserving all the planes with weight above 0.8 of the maximum weight. As can be observed, our algorithm approximately detects the two symmetry planes. The result of applying the reflection and the local registration is also shown. The original shape and its symmetric are shown overlapped, with the colored being the result of the transformation. Note that this method allows us to recover some features that are lost in the original object due to fractures.

A more challenging shape is presented in Fig. 10. For this case, our algorithm detected three symmetries which are related to the three sphere-like local protrusions in the shape. Here we show three possible completions performed for each symmetry plane detected. In particular for this shape, our method takes advantage of the careful design of the two stages of our algorithm. For example, our feature detection stage selects symmetric points as can be seen in Figure 3 thanks to the balanced characterization of globality and locality. Also, the third symmetry plane may only be detected if we focus the voting scheme to the region determined by the plane support. Otherwise, there would be a lot of spurious votes making difficult the correct detection. It is worth to mention that the three symmetry planes correspond to the only three planes that remain after the selection of the generating points and the voting scheme.

Limitations

Our method is based on the ability of finding good candidate features. For this reason, our method is not suitable for feature-less shapes. In addition, if it is not possible to recognize at least a pair of symmetric features, the method could not compute candidate planes. Moreover, in feature-




Figure 10: A challenging embrasure shape. The resulting synthesized geometry is consistent with the original object, and therefore the completion generates plausible objects.

less shapes the voting scheme could not work due to the fact that it is based on high-curvature points. Nevertheless, we believe that our method may perform well in most of shapes in real applications as shown in the completion of cultural heritage objects.

7. Conclusions

In this paper, we presented an algorithm to detect symmetry planes in objects with partial data. Our method is composed of two stages: the feature detection and the vote-based validation. We proposed a new multi-scale diffusion-based surface function which exhibit a good trade-off between locality and globality. This balance has proved to be effective to select points in local protrusions and their symmetric correspondences. This fact facilitates the generation of the initial set of planes which can be considered as hypotheses that need to be validated. Precisely, the reliability of the potential planes is well complemented with the second stage where we aimed at building a voting system for validation. The definition of the plane support and the verification rules guarantee the reduction of false positive votes. As consequence, the combination of reliable multi-scale features and a partiality-aware voting scheme have proven to be effective to deal with the problem of missing geometry. In the future, we plan to extend our approach to detect intrinsic symmetries. Furthermore, we also plan to focus the attention in symmetries with respect to axes (1D) unlike the current proposal that only utilizes symmetries with respect to planes (2D).

Acknowledgments

We would like to thank to anonymous reviewers for their constructive comments. The Armadillo model is courtesy of the Stanford 3D Scanning Repository. The human shape was taken from the TOSCA dataset [BBK08]. The cake, gargoyle and head models are courtesy by Vienna University of Technology. This work was supported by EC FP7 STREP Project PRESIOUS, grant no. 600533.

References

- [ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *IEEE Int. Conf. in Computer Vision Workshops* (2011), pp. 1626–1633. 4
- [BAK10] BERMANIS A., AVERBUCH A., KELLER Y.: 3-d symmetry detection and analysis using the pseudo-polar fourier transform. *International Journal of Computer Vision* 90, 2 (2010), 166–182. 2
- [BBK08] BRONSTEIN A., BRONSTEIN M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*, 1 ed. Springer Publishing Company, Incorporated, 2008. 9
- [BSW08] BELKIN M., SUN J., WANG Y.: Discrete laplace operator on meshed surfaces. In *Proceedings of the Twenty-fourth Annual Symposium on Computational Geometry* (New York, NY, USA, 2008), SCG '08, ACM, pp. 278–287. 6
- [CSM03] COHEN-STEINER D., MORVAN J.-M.: Restricted delaunay triangulations and normal cycle. In *Proceedings of the Nineteenth Annual Symposium on Computational Geometry* (New York, NY, USA, 2003), SCG '03, ACM, pp. 312–321. 7
- [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. 8

- [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
- [KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson surface reconstruction. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing* (Aire-la-Ville, Switzerland, Switzerland, 2006), SGP '06, Eurographics Association, pp. 61–70. 7
- [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* (New York, NY, USA, 2004), SGP '04, ACM, pp. 115–123. 2
- [KKP13] KAKARALA R., KALIAMOORTHY P., PREMACHANDRAN V.: Three-dimensional bilateral symmetry plane estimation in the phase domain. In *Proceedings of the 2013 IEEE Conference on Computer Vision and Pattern Recognition* (Washington, DC, USA, 2013), CVPR '13, IEEE Computer Society, pp. 249–256. 2
- [KLAB14] KORMAN S., LITMAN R., AVIDAN S., BRONSTEIN A. M.: Probably approximately symmetric: Fast 3d symmetry detection with global guarantees. *CoRR abs/1403.6637* (2014). 2
- [LB14] LITMAN R., BRONSTEIN A.: Learning spectral descriptors for deformable shape correspondence. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 36, 1 (Jan 2014), 171–180. 4
- [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
- [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
- [MSHS06] MARTINET A., SOLER C., HOLZSCHUCH N., SIL-LION F. X.: Accurate detection of symmetries in 3d shapes. *ACM Trans. Graph.* 25, 2 (Apr. 2006), 439–464. 2
- [PSG*06] PODOLAK J., SHILANE P., GOLOVINSKIY A., RUSINKIEWICZ S., FUNKHOUSER T.: A planar-reflective symmetry transform for 3d shapes. *ACM Trans. Graph.* 25, 3 (July 2006), 549–559. 2
- [RL01] RUSINKIEWICZ S., LEVOY M.: Efficient variants of the icp algorithm. In *3-D Digital Imaging and Modeling, 2001. Proceedings. Third International Conference on* (2001), pp. 145–152. 8
- [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). 4
- [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
- [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 (Dec. 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

Appendix A

Proof of Proposition 1 . The proof consists in the demonstration that the function $f(x) = (1 - e^{-xt})/x$ is strictly decreasing and therefore the maximum value occurs when x approaches zero on the right. To prove that $f(x)$ is strictly decreasing, we need to prove that $f'(x) < 0$ for $x > 0$. The derivative $f'(x)$ yields

$$f'(x) = \frac{e^{-xt}(xt + 1 - e^{xt})}{x^2}.$$

Given that the terms e^{-xt} and x^2 are always positive for $x > 0$, then $f'(x) < 0$ if and only if

$$(xt + 1 - e^{xt}) < 0.$$

This inequality is true for any $x > 0$ and $t > 0$. This can be proved if we replace e^{xt} with its power series

$$e^{xt} = 1 + xt + \frac{(xt)^2}{2} + \frac{(xt)^3}{6} + O(x^4).$$

Now that we know that $f(x)$ is decreasing for any $x > 0$, the maximum value occurs in the neighborhood of $x = 0$. To calculate its maximum value, we need to evaluate the limit of the function when x tends to zero. Formally,

$$\lim_{x \rightarrow 0^+} \frac{1 - e^{-xt}}{x}$$

As the limit is indeterminate in $x = 0$, we apply the l'Hôpital's rule to evaluate it, which yields

$$\lim_{x \rightarrow 0^+} te^{-xt} = t.$$

This result proves that t is an upper bound of $f(x)$ when $x > 0$ and $t > 0$.