Characterization of Partial Intrinsic Symmetries

Aurela Shehu^{1,2}, Alan Brunton³, Stefanie Wuhrer¹, and Michael Wand⁴

¹ Cluster of Excellence MMCI, Saarland University, Germany

² Saarbrücken Graduate School of Computer Science, Germany

³ Fraunhofer Institute for Computer Graphics Research IGD, Germany

⁴ Utrecht University, Netherlands

Abstract. We present a mathematical framework and algorithm for characterizing and extracting partial intrinsic symmetries of surfaces, which is a fundamental building block for many modern geometry processing algorithms. Our goal is to compute all “significant” symmetry information of the shape, which we define as r -symmetries, i.e., we report all isometric self-maps within subsets of the shape that contain at least an intrinsic circle or radius r . By specifying r , the user has direct control over the scale at which symmetry should be detected. Unlike previous techniques, we do not rely on feature points, voting or probabilistic schemes. Rather than that, we bound computational efforts by splitting our algorithm into two phases. The first detects infinitesimal r -symmetries directly using a local differential analysis, and the second performs direct matching for the remaining discrete symmetries. We show that our algorithm can successfully characterize and extract intrinsic symmetries from a number of example shapes.

Keywords: symmetry, shape analysis, shape matching, intrinsic geometry, slippability analysis

1 Introduction

Shape symmetry is beneficial for many applications such as shape segmentation, geometry completion, beautification, shape synthesis, scan denoising and shape matching. A symmetry, or symmetric mapping of a shape, refers to a self-mapping that does not alter the geometry of the shape. *Extrinsic geometry* refers to properties of the shape, which depend on the coordinate space in which the shape is embedded. *Intrinsic geometry* refers to properties that can be measured on the surface, which are invariant to the embedding.

To detect extrinsic symmetries, i.e. self-mappings that do not alter the extrinsic geometry of the shape, one looks for rigid transformations, which map the shape to itself. However, in many applications we are interested in shapes that deform non-rigidly such that the surface does not tear or stretch significantly, such as locomotion of humans and animals, or folding of cloth. The symmetries present in such shapes in general cannot be detected extrinsically due to asymmetry of the embedding. We can circumvent this problem by considering

intrinsic geometry, which is invariant to isometric deformations. A self-mapping that does not alter the intrinsic geometry describes an *intrinsic symmetry*.

We can further distinguish between global and partial symmetries. A *global intrinsic symmetry* maps the entire shape to itself, whereas a *partial intrinsic symmetry* maps a part of the shape to another part of the shape, in a way that does not alter the intrinsic geometry of the part. In this work, we address the problem of detecting and characterizing partial intrinsic symmetries on a given shape. Shapes can be symmetric at various scales, and the desired symmetry scale may be application dependent. To model this, we use a symmetry scale r . For every geodesic disk of radius r on the shape, we aim to find whether there is a mapping of the disk on the shape, which does not induce any (significant) stretching. We call such a mapping an intrinsic r -symmetry.

The goal of our work is to report all r -symmetries, i.e., all mappings of large enough subsets of the shape to itself that are perfect isometries. The key challenge is that the space of partial intrinsic symmetries might be too large to systematically explore and report [24], [28]. Following Brunton et al. [6], we define partial correspondences by sets that are mapped with preservation of the local metric. Unlike Xu et al. [28], we do not require that the shortest geodesic paths to be maintained outside the mapped region, which leads to a much higher combinatorial complexity due to the global dependency on topology and boundaries of the subsets. In contrast, metric-preserving maps can be parametrized by only five degrees of freedom (two out of which do not even need fine sampling) [6], which makes an approximate enumeration of all isometries feasible in polynomial time. Our paper follows this idea but factors out infinitesimal symmetries, which dominate the run-time otherwise, and uses local descriptors to further narrow down the search.

Most methods that detect partial discrete intrinsic symmetries try to reduce the search space by computing features on the shape and further detect partial intrinsic symmetries in the feature space. While this approach reduces the complexity of the problem, it might miss some intrinsic r -symmetries. Further, existing methods have focused on either detecting partial infinitesimal symmetries [8], [26] or partial discrete symmetries [24], [28], [22], [13], [4], [12], [18], [23], [9], [15] but not both.

In contrast, we propose a method that finds discrete intrinsic r -symmetries and characterizes infinitesimal intrinsic r -symmetries of a shape in an efficient way. The search space is reduced by a mathematical analysis on infinitesimal r -symmetries, assuring that no information is thrown away, but simultaneously making the computations feasible. To our knowledge, this is the first method which provides a full characterization of intrinsic r -symmetries of a shape.

We provide a characterization of both infinitesimal and discrete partial point-wise intrinsic symmetries on 3D shapes, and present an algorithm for detecting such symmetries. Given as input a 3D shape and a radius r , our method tests first each point for infinitesimal r -symmetry and then for the ones with no infinitesimal r -symmetry it tests them for discrete r -symmetry since now we can enumerate the finite discrete r -symmetries. As a last step, for all the points that

belong to discrete r -symmetries we compute approximate equivalence classes. We evaluate our method on a standard symmetry dataset.

1.1 Definitions and Problem Statement

We consider a shape to be a smooth, orientable 2-manifold $\mathcal{M} \subset \mathbb{R}^3$ embedded in three-dimensional space, and denote the tangent space of $\mathcal{M}$ at point $\mathbf{s} \in \mathcal{M}$ by $T_{\mathbf{s}}\mathcal{M}$. In practice, $\mathcal{M}$ is discretized by a set of points that are connected by a neighborhood graph. We denote the geodesic distance between two points $\mathbf{s}, \mathbf{t} \in \mathcal{M}$ by $d_{\mathcal{M}}(\mathbf{s}, \mathbf{t})$.

Definition 1. *The points $\mathbf{s}$ and $\mathbf{t}$ on $\mathcal{M}$ are called r -symmetric if there is a mapping function $f : \mathcal{U} \rightarrow \mathcal{M}$, where $\mathcal{U}$ is the geodesic disk of radius r centered at $\mathbf{s}$, with $f(\mathbf{s}) = \mathbf{t}$ such that for all point pairs $(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{U}$, it is*

$$|d_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_j) - d_{\mathcal{M}}(f(\mathbf{x}_i), f(\mathbf{x}_j))| = 0. \quad (1)$$

Note that, for small enough r , $d_{\mathcal{U}}(\mathbf{x}_i, \mathbf{x}_j) = d_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_j)$ as $\mathcal{U}$ is convex.

Consider mapping functions of the type $f : \mathcal{M} \rightarrow \mathcal{M}$. If Equation (1) is satisfied under such a mapping, then $\mathbf{s}$ and $\mathbf{t}$ define a *global intrinsic symmetry*. We consider *partial intrinsic symmetries* with mapping functions of the type $f : \mathcal{U} \rightarrow \mathcal{M}$, where parts are defined as geodesic neighborhoods of points with fixed radius. The set of r -symmetric points form an equivalence relation.

We have guarantees only for exact symmetries and as these occur rarely in practice, we relax the symmetry condition by allowing for a stretching up to a fixed threshold ϵ .

Definition 2. *The points $\mathbf{s}$ and $\mathbf{t}$ are called (ϵ, r) -symmetric if there is a mapping function $f : \mathcal{U} \rightarrow \mathcal{M}$ with $f(\mathbf{s}) = \mathbf{t}$ such that for all point pairs $(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{U}$, it is*

$$|d_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_j) - d_{\mathcal{M}}(f(\mathbf{x}_i), f(\mathbf{x}_j))| \leq \epsilon. \quad (2)$$

The set of (ϵ, r) -symmetric points no longer form an equivalence relation.

When the mapping f encodes infinitesimal movements of points $\mathbf{x}_i$ in $\mathcal{U}$ and Equations (1) or (2) are satisfied, then $\mathbf{s}$ and $\mathbf{t}$ define an *infinitesimal (ϵ, r) -symmetry*. Note that in the literature, infinitesimal symmetries are often also called continuous symmetries because the isometries involved form a continuous set.

When the mapping f encodes movements by fixed geodesic distances of points $\mathbf{x}_i$ in $\mathcal{U}$ and Equations (1) or (2) are satisfied, then $\mathbf{s}$ and $\mathbf{t}$ define a *discrete (ϵ, r) -symmetry*. The reason for the term discrete is that the set of discrete global symmetries form a discrete set.

1.2 Overview

We now give a coarse outline of our algorithm for detecting intrinsic (ϵ, r) -symmetries. We start by detecting and factoring out all (possibly infinitely many)

Fig. 1: Illustrative example. Left to right: size of r , point slippability, symmetry-factored embedding distance from a corner of the cube.

infinitesimal (ϵ, r) -symmetries. We adapt methods that find infinitesimal symmetries in the extrinsic case [7], [10] to our scenario. Once the infinitesimal symmetries are factored out (e.g. leaving the white regions in Fig. 1, middle), there remain only a finite, and often manageable, number of discrete (ϵ, r) -symmetries. To extract these we use the fact that point-wise (ϵ, r) -symmetries arise from an ϵ -approximate isometric mapping (correspondence) between two regions $\mathcal{U}$ and $f(\mathcal{U})$, as defined in (2). Isometric mappings have a low number of degrees of freedom, which has been recently used to develop a direct region growing method for partial isometric correspondence [6]. We exploit this method, initialized with a sampling strategy based on spectral descriptors, to find the discrete symmetries. These steps are shown as illustration for a cube in Fig. 1, and explained in detail in the subsequent sections.

2 Related Work

There is a considerable amount of research on detecting extrinsic symmetry on 3D shapes and for an overview, we refer the reader to the survey from Mitra et al. [16]. Here we focus on research on detecting intrinsic symmetry.

To compute infinitesimal intrinsic symmetries, Ben Chen et al. [3] propose a method to compute approximate Killing vector fields on triangle meshes, which are tangent vector fields on the surface that generate isometric transformations. There are two methods that allow for the computation of partial infinitesimal symmetries. Solomon et al. [26] partition the shape into intrinsically symmetric building blocks and can be thought of as an extension of the method by Ben Chen et al. [3] to the partial case. While this method allows to “untangle” the global information into parts, it does not allow the detection of symmetries on a specific scale. Grushko et al. [8] find infinitesimal symmetries as the ones that lead to small stress of the generalized multi-dimensional scaling (MDS) energy [5]. It requires as input the target location of one point as boundary condition. Our method to detect infinitesimal partial symmetries can be seen as similar in spirit, but more general. We can find all directions by which a sample can move by an infinitesimal amount while its geodesic neighborhood stays symmetric, without the need for an initial target point location.

To compute global discrete intrinsic symmetries, several methods based on isometric invariants have been proposed. Kim et al. [11] use the average geodesic

function to find feature points and compute global discrete symmetries as Möbius transformations generated by these feature points. Possible symmetries generated by non-feature points will not be reported. Methods that use the Laplace-Beltrami operator embed the surface into a space where the global discrete symmetry detection has a reduced search space. Ovsjanikov et al. [21] reduce the space further by only considering non-repeating eigenvalues and by finding global intrinsic symmetries by transforming the problem to extrinsic symmetry detection in the embedding space. Ovsjanikov et al. [20] identify and factor out symmetries before finding correspondences between a pair of near isometric shapes and require a symmetry map on one shape from the user.

After detecting approximate global symmetries, it remains to recover the symmetry group from the approximate information. Lipman et al. [13] propose an embedding where symmetry is factored out to recover approximate cliques in the embedding space, which correspond to equivalence classes of symmetric points. Wang et al. [27] extend this approach by introducing an embedding where repeating eigenvalues of the Laplace-Beltrami operator are used to capture symmetry transformations. In this work, we modify the method of Lipman et al. to find approximately equivalent parts for visualization purposes.

As the complexity of the space of all partial symmetries is believed to be large [24], [28], existing methods proceed by reducing the search space after sampling the surface.

Some methods find partial discrete symmetries by minimizing a MDS energy. Raviv et al. [24] define as partial discrete symmetry a self-isometric part that minimizes the generalized MDS energy, and find symmetries by comparing pairs of samples that pass a pruning method. Raviv et al. [23] extend the previous approach by using a diffusion distance metric. Mitra et al. [15] detect intrinsically symmetric blocks on a given surface by flattening the surface using MDS, by solving the problem on the embedded image, and by projecting the computed patterns back to the original surface. This method is restricted to surfaces that are topologically equivalent to a plane.

Berner et al. [4] reduce the search space of possible partial discrete symmetries by using features. This method might miss possible symmetry candidates because of the feature selection. To overcome this problem, Lasowski et al. [12] find partial intrinsic symmetries by modeling a probability distribution over all possible correspondences. This method is sensitive to topological noise.

Partial discrete symmetries can alternatively be detected by a voting procedure. Xu et al. [28] decouple the problem into the two problems of symmetric point pair detection and symmetry scale detection. The specific difficulty of this method is the preservation of all shortest geodesic paths after a partial mapping to a larger context. This makes the solution space too large for a low-level, systematic exploration. Our model avoids this by removing these global guarantees for indirect effects due to geometry outside the region identified as symmetric. Mukhopadhyay et al. [18] detect overlapping intrinsically symmetric regions by a voting procedure of symmetric point pairs and characterize the computed symmetries. Jiang et al. [9] extend this type of voting procedure by coupling

it with a robust skeleton extraction method. All of these methods require complex optimizations before the voting step, and are sensitive to topological noise. Furthermore, in all of these methods, the scale(s) of the detected symmetries cannot be controlled by the user. Our method gives direct control of the scale of symmetry to be detected to the user.

To summarize, existing methods reduce the resulting search space, commonly using features, pruning, or voting techniques. In contrast, we observe that the space of all partial symmetries can be modeled explicitly as follows. There may be infinitely many infinitesimal partial symmetries, and points that have infinitesimal partial symmetries need to be detected and factored out to make partial symmetry detection feasible. Additionally, there may be finitely many discrete partial symmetries, each of which can be parameterized by a pair of point and tangent direction matches in a low-dimensional search space. This observation allows us to systematically detect infinitesimal and discrete partial symmetries in a way that is similar in spirit to work on detecting extrinsic partial symmetries for a micro-tile decomposition of a shape [10].

3 Infinitesimal r -Symmetry Detection

This section presents our approach to detect and characterize infinitesimal (ϵ, r) -symmetries.

3.1 Detecting r -Symmetries

We aim to find all directions $\mathbf{d}$, such that all points $\mathbf{x}_i \in \mathcal{U}$ can move by an infinitesimal amount in direction $\mathbf{d}$ without inducing any stretching. In contrast to the method of Ben-Chen et al. [3], we constraint the measured stretching to the shape part $\mathcal{U}$ instead of the entire shape. Here, $\mathbf{d} = [\delta_1^T \dots \delta_n^T]^T$ is a vector of length $2n$ that defines how each point $\mathbf{x}_i$ moves in its local tangent plane $T_{\mathbf{x}_i}\mathcal{M}$. We extend the notation from works on extrinsic symmetry detection [7] and call the directions $\mathbf{d}$ *r -slippage directions*.

To find the r -slippage directions $\mathbf{d}$, we construct a matrix $\mathbf{A}$ that has all r -slippage directions as eigenvectors with corresponding eigenvalue zero. This can be seen as a direct extension of the extrinsic slippage analysis of Gelfand and Guibas [7] to the intrinsic case. The main challenge in extending the slippage analysis to the intrinsic case is that the degrees of freedom that control the movement a part $\mathcal{U}$ increases from 6 in the extrinsic case (3 for rotation, 3 for translation) to $2n$ in the intrinsic case, where n denotes the number of samples in $\mathcal{U}$ (2 to indicate the movement of each sample in its tangent plane).

We start by considering the stretching of a mapping $f : \mathcal{U} \rightarrow \mathcal{M}$. Given a uniform sampling of the manifold $\mathcal{M}$, the stretching is

$$E_{SD}(f) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n (d_{\mathcal{M}}(f(\mathbf{x}_i), f(\mathbf{x}_j)) - d_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_j))^2, \quad (3)$$

where n is the number of samples in $\mathcal{U}$ and serves as an approximation of the area $A_{\mathcal{U}}$ of $\mathcal{U}$.

The energy expressed in Equation (3) is similar to the energy used in Generalized Multi-Dimensional Scaling (GMDS) [5]. GMDS tries to fit a surface $\mathcal{X}$ to a surface $\mathcal{Y}$ without distorting the intrinsic metric of $\mathcal{X}$. In our case, we try to fit the geodesic disk $\mathcal{U}$ to $f(\mathcal{U})$ without distorting the intrinsic metric of $\mathcal{U}$.

To find the r -slippage directions, we aim to minimize the energy given in Equation (3). Consider the identity mapping $f_0(\mathbf{x}) = \mathbf{x}$, $\mathbf{x} \in \mathcal{U}$. This mapping obviously minimizes $E_S(f_0) = 0$ and $E_{SD}(f_0) = 0$. Since f_0 minimizes E_{SD} , the gradient $\nabla_{f_0} E_{SD}(f)$ is zero as well. We write the gradient with respect to the mapping $f(\mathbf{x}_k)$ of a particular sample $\mathbf{x}_k$ as

$$\nabla_{f(\mathbf{x}_k)} E_{SD}(f) = \frac{\partial E_{SD}}{\partial f(\mathbf{x}_k)} = \frac{4}{n^2} \sum_{i=1}^n (d_{\mathcal{M}}(f(\mathbf{x}_k), f(\mathbf{x}_i)) - d_{\mathcal{M}}(\mathbf{x}_k, \mathbf{x}_i)) \nabla_k \mathbf{p}_{ik}, \quad (4)$$

with $\nabla_k \mathbf{p}_{ik} = \frac{\partial}{\partial f(\mathbf{x}_k)} (d_{\mathcal{M}}(f(\mathbf{x}_i), f(\mathbf{x}_k)))$, which, from the Eikonal equation, is a unit-vector in $T_{f(\mathbf{x}_k)} \mathcal{M}$.

Consider updating f_0 by an infinitesimal amount in direction $\mathbf{d}$. We are interested in finding directions $\mathbf{d}$, such that the mapping $f_0 + \mathbf{d}$ has a gradient of zero, which means $\nabla_{f(\mathbf{x}_k)} E_{SD}(f_0 + \mathbf{d}) = 0$.

The second derivative of E_{SD} with respect to the mapping f gives the rate of change of the first derivative under the mapping. Let $\mathbf{A}$ denote the second-order derivative matrix of f_0 that contains the block matrices

$$\begin{aligned} \mathbf{A}_{ki} &= \nabla_{f(\mathbf{x}_k)} \nabla_{f(\mathbf{x}_i)} E_{SD}(f_0) \\ \mathbf{A}_{kk} &= \nabla_{f(\mathbf{x}_k)} \nabla_{f(\mathbf{x}_k)} E_{SD}(f_0). \end{aligned} \quad (5)$$

For small $\mathbf{d}$, we can find the r -slippage directions $\mathbf{d}$ of $\mathcal{U}$ as the eigenvectors of $\mathbf{A}$ with corresponding eigenvalues zero, which means vectors $\mathbf{d}$ with

$$\nabla_{f(\mathbf{x}_k)} E_{SD}(f_0 + \mathbf{d}) \approx \mathbf{A} \mathbf{d} = \mathbf{0}. \quad (6)$$

The matrix $\mathbf{A}$ is symmetric, which implies that all its eigenvalues are real.

It remains to compute the second derivatives of E_{SD} with respect to the mapping. Taking the derivative of $\nabla_{f(\mathbf{x}_k)} E_{SD}(f)$ with respect to $f(\mathbf{x}_k)$ we get

$$\nabla_{f(\mathbf{x}_k)} \nabla_{f(\mathbf{x}_k)} E_{SD}(f) = \frac{4}{n^2} \sum_{i=1}^n \left(\nabla_k \mathbf{p}_{ik} \nabla_k \mathbf{p}_{ik}^T + (d_{\mathcal{M}}(f(\mathbf{x}_k), f(\mathbf{x}_i)) - d_{\mathcal{M}}(\mathbf{x}_k, \mathbf{x}_i)) \nabla_k \nabla_k \mathbf{p}_{ik}^T \right). \quad (7)$$

Taking the derivative of $\nabla_{f(\mathbf{x}_k)} E_{SD}(f)$ with respect to $f(\mathbf{x}_i)$, $i \neq k$, we get

$$\nabla_{f(\mathbf{x}_k)} \nabla_{f(\mathbf{x}_i)} E_{SD}(f) = \frac{4}{n^2} \left(\nabla_i \mathbf{p}_{ik} \nabla_k \mathbf{p}_{ik}^T + (d_{\mathcal{M}}(f(\mathbf{x}_k), f(\mathbf{x}_i)) - d_{\mathcal{M}}(\mathbf{x}_k, \mathbf{x}_i)) \nabla_i \nabla_k \mathbf{p}_{ik}^T \right). \quad (8)$$

In the following, we make the assumption that $\mathcal{M}$ is approximated locally by its tangent planes, which causes the second order derivative terms $\nabla_k \nabla_k \mathbf{p}_{ik}$ and $\nabla_i \nabla_k \mathbf{p}_{ik}$ in Equations (7) and (8) to vanish and lead to the following simplified derivatives

$$\begin{aligned} \nabla_{f(\mathbf{x}_k)} \nabla_{f(\mathbf{x}_k)} E_{SD}(f) &= \frac{4}{n^2} \sum_{i=1}^n \nabla_k \mathbf{p}_{ik} \nabla_k \mathbf{p}_{ik}^T \\ \nabla_{f(\mathbf{x}_k)} \nabla_{f(\mathbf{x}_i)} E_{SD}(f) &= \frac{4}{n^2} \nabla_i \mathbf{p}_{ik} \nabla_k \mathbf{p}_{ik}^T. \end{aligned} \quad (9)$$

3.2 Detecting (ϵ, r) -Symmetries

We now aim to compute all directions, by which a sample $\mathbf{s}$ can move by an infinitesimal amount while staying (ϵ, r) -symmetric. We call these directions (ϵ, r) -*slippage directions* in the following.

From the previous section (Equation (6)) we know that for $\epsilon = 0$, the r -slippage directions can be found by finding all eigenvectors $\mathbf{d}$ of the second derivative matrix $\mathbf{A}$ of f_0 with associated eigenvalues zero. In this section, we show that the (ϵ, r) -slippage directions of $\mathbf{s}$ can be found by finding all eigenvectors $\mathbf{d}$ of the second derivative matrix $\mathbf{A}$ of f_0 with associated eigenvalues λ with $|\lambda| \leq (\frac{2}{n^2}\epsilon^2)$. Note that there is no derivation of an upper bound in previous work related to slippage analysis [7].

In practice, we are not looking for the null space of A , but we consider all eigenvectors $\mathbf{d}$ with associated eigenvalues λ with $|\lambda|$ at most some threshold to be (ϵ, r) -slippage directions.

By integrating both sides of Equation 6, and using that the eigenvector $\mathbf{d}$ of $\mathbf{A}$ has at most unit length, we obtain

$$\int_{\mathbf{x}=0}^{\mathbf{d}} \nabla_{f(\mathbf{x}_k)} E_{SD}(f_0 + \mathbf{x}) d\mathbf{x} = \int_{\mathbf{x}=0}^{\mathbf{d}} \mathbf{A}\mathbf{x} d\mathbf{x}$$

$$E_{SD}(f_0 + \mathbf{d}) \leq \frac{1}{2} |\lambda| \quad (10)$$

Hence, using Equation 3, we can obtain the following loose upper bound on any individual stretch $|d_{\mathcal{M}}(\mathbf{x}_i + \delta_i, \mathbf{x}_j + \delta_j) - d_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_j)| \leq \frac{n}{\sqrt{2}} \sqrt{|\lambda|}$, $\forall i, j$, where δ_i is the part of the offset vector $\mathbf{d}$ in the tangent space of $\mathbf{x}_i$. This implies that by finding all eigenvectors $\mathbf{d}$ of $\mathbf{A}$ with associated eigenvalues λ with $|\lambda| \leq (\frac{2}{n^2}\epsilon^2)$, we can find the (ϵ, r) -slippage directions.

3.3 Characterization of Samples

With the results from the previous section, we can find all linearly independent (ϵ, r) -slippage directions for a sample $\mathbf{s}$ on $\mathcal{M}$. It is known that $\mathcal{M}$ can have at most three linearly independent (global) slippage directions [25]. Hence, we can characterize the sample $\mathbf{s}$ according to how many linearly independent (ϵ, r) -slippage directions $\mathbf{s}$ has. In the following, we call $\mathbf{s}$ *t-slippable* for $t = 0, 1, 2, 3$ if $\mathbf{s}$ has t linearly independent (ϵ, r) -slippage directions. Furthermore, we call $\mathbf{s}$ *slippable* if it has at least one (ϵ, r) -slippage direction, and *non-slippable* otherwise. Consider the illustrative example of a cube shown in the middle of Fig. 1, where white is 0-slippable, red 1-slippable, green 2-slippable and blue 3-slippable. A disk of radius r is shown in blue in the left of Fig. 1. As shown, a cube is 3-slippable almost everywhere except at the corners, where it is 0-slippable.

4 Discrete r -Symmetry Detection

In the previous section we presented our approach to detect infinitesimal (ϵ, r) -symmetries. In the following, we can factor out slippable points from further

analysis, thereby reducing the search space for discrete discrete (ϵ, r) -symmetries to all oriented matches of only the non-slippable points.

4.1 Discrete r -Symmetries by Isometric Region Growing

To test if two points have a discrete (ϵ, r) -symmetry, we use the method proposed by Brunton et al. [6]. Based on the observation that a map of a single point pair and its tangent space is sufficient to recover an isometry, they propose a low-dimensional representation to efficiently compute a mapping $f : \mathcal{V} \rightarrow \mathcal{M}$ that minimizes the stretch given in Equation (3). Here, $\mathcal{V}$ is grown to the largest possible size that allows a mapping with metric stretching at most ϵ .

From a conceptual point of view, detecting all discrete (ϵ, r) -symmetries would involve trying all possible point pairs with all possible direction alignments, and to check if the tangent spaces are related by an isometry up to a threshold ϵ . As suggested by Brunton et al. [6], one may compute features in order to reduce the search space. However, feature detection methods may miss certain symmetries. To reduce this risk, we sample a subset of the surface points that are left after removing all slippable points, and we test all pairs of samples.

To reduce the search space, we do not test point pairs that are close-by in terms of geodesic distance for symmetry. Furthermore, we prune the remaining point pairs and directions using an isometry-invariant shape descriptor. We choose the state of the art Wave Kernel Signature (WKS) [2] as descriptor, and disregard point pairs whose WKS do not agree up to a threshold τ . Computing the WKS globally over the entire surface is not consistent with our partial symmetry model. To remedy this, we compute the WKS over the geodesic r -neighborhood $\mathcal{U}$ of a point. This step is to speed up on exact data, i.e. shapes that contain exact isometric parts. In this case, all r -symmetries will be discovered but the method will try not to investigate the least promising ones.

For every point pair that passed the point pruning step, we proceed by pruning the direction pairs. Direction pairs pruning is done with the help of a descriptor, which we call *Wave Kernel Map* (WKM) in the following and is defined based on the wave kernel signature [2] in a similar way that the heat kernel map is defined based on the heat kernel signature [19]. Given that the wave kernel signature is defined as $\text{WKS}_\epsilon(\mathbf{x}, \mathbf{x}) = C_e \sum_k \phi_k^2(\mathbf{x}) e^{-\frac{(\epsilon - \log E_k)^2}{2\sigma}}$, we define the wave kernel map as

$$\text{WKM}_\epsilon(\mathbf{x}, \mathbf{y}) = C_e \sum_k \phi_k(\mathbf{x}) \phi_k(\mathbf{y}) e^{-\frac{(\epsilon - \log E_k)^2}{2\sigma}}. \quad (11)$$

Let $\mathbf{s}, \mathbf{t} \in \mathcal{M}$ be a point pair which passed the point pruning step, where $\mathcal{U}$ is a geodesic disk centered at $\mathbf{s}$, $f(\mathcal{U})$ is the corresponding patch of $\mathcal{U}$ under the mapping function f and $f(\mathbf{s}) = \mathbf{t}$. For all $\mathbf{x} \in \mathcal{U}$ we compute $\text{WKM}_\epsilon(\mathbf{s}, \mathbf{x})$ and for all $\mathbf{y} \in f(\mathcal{U})$ we compute $\text{WKM}_\epsilon(\mathbf{t}, \mathbf{y})$. Let $\mathbf{s}_{max}$ and $\mathbf{t}_{max}$ be

$$\begin{aligned} \mathbf{s}_{max} &= \max_x \text{WKM}_\epsilon(\mathbf{s}, \mathbf{x}) \\ \mathbf{t}_{max} &= \max_y \text{WKM}_\epsilon(\mathbf{t}, \mathbf{y}). \end{aligned} \quad (12)$$

We consider as direction pairs the union of the direction matches $\text{dir}_{\mathbf{s}}$ based on $\mathbf{s}_{max}$ and the direction matches $\text{dir}_{\mathbf{t}}$ based on $\mathbf{t}_{max}$ where

$$\begin{aligned}\text{dir}_{\mathbf{s}} &= \{(\mathbf{s}_{max}, t_{\text{dir}}) | t_{\text{dir}} : |\text{WKM}_e(\mathbf{t}, t_{\text{dir}}) - \text{WKM}_e(\mathbf{s}, \mathbf{s}_{max})| \leq \tau_d\} \\ \text{dir}_{\mathbf{t}} &= \{(\mathbf{t}_{max}, s_{\text{dir}}) | s_{\text{dir}} : |\text{WKM}_e(\mathbf{s}, s_{\text{dir}}) - \text{WKM}_e(\mathbf{t}, \mathbf{t}_{max})| \leq \tau_d\}\end{aligned}\quad (13)$$

and where τ_d is a direction pair threshold. These point pairs are used to encode the direction pairs from the point pair $\mathbf{s}, \mathbf{t}$. By disregarding all other possible direction matches, we can eliminate alignments that are not near-isometric.

After pruning the point and direction pairs, we use the method proposed by Brunton et al. [6] for each pair of points plus directions to find a mapping $f: \mathcal{V} \rightarrow \mathcal{M}$ that is isometric up to threshold ϵ . From this output, we can detect all points that are (ϵ, r) -symmetric by reporting every point $\mathbf{s} \in \mathcal{V}$ whose geodesic r -neighborhood $\mathcal{U}$ is entirely contained in $\mathcal{V}$ along with its corresponding point $\mathbf{t}$ and vice versa.

4.2 Approximate Equivalence Classes

Finally, we would like to detect classes of equivalent points under (ϵ, r) -symmetries. In the case of exact data, we look for r -symmetries ($\epsilon = 0$) and this step is not necessary since we can directly find exact equivalence classes. Considering $\epsilon > 0$, this problem is ill-posed as the set of all (ϵ, r) -symmetries do not form equivalence classes in general. However, since we assume that a threshold of ϵ is only required to counteract noise, there is an underlying equivalence class. To find this, we detect fuzzy equivalence classes of (ϵ, r) -symmetries using an approach similar to the one proposed by Lipman et al. [13]. In this way, transition consistencies will be amplified and everything else diminished.

We use our detected discrete (ϵ, r) -symmetries to construct a dissimilarity matrix $\mathbf{S} \in \mathbb{R}^{n \times n}$ as

$$S_{ij} = \min\{\min_{i'} d_{\mathcal{M}}(i, i'), \min_{j'} d_{\mathcal{M}}(j, j')\}, \quad (14)$$

where (i, j') and (j, i') define a discrete (ϵ, r) -symmetry and n is the number of samples on $\mathcal{M}$. We now can compute the symmetry correspondence matrix $\mathbf{C}$ as proposed by Lipman et al. [13] as $\tilde{\mathbf{C}}_{ij} = e^{-\left(\frac{S_{ij}}{\sigma_{diam}}\right)^2}$, where $diam$ is twice the intrinsic radius of the manifold $\mathcal{M}$. The symmetry correspondence matrix is derived from $\tilde{\mathbf{C}}$ by making each row sum up to one. After applying eigendecomposition to the matrix $\mathbf{C}$, we compute the symmetry factored embedding by using the eigenvalues and the eigenvectors of $\mathbf{C}$.

To visualize the approximate equivalence classes, we compute the symmetry-factored embedding (SFE) distance as Euclidean distance in embedding space and color the points based on the distance, where blue is small and red is large. Consider again the illustrative example of a cube in Fig. 1. The right of Fig. 1 shows the SFE distance from one of the corners of the cube. All of the corners have a small SFE distance to each other, and are therefore colored blue. The farther away a point is from the corner, the more red is its color.

5 Experiments

We evaluate our method on a common intrinsic symmetry dataset that is available from Xu et al. [28]. Our method requires as input a uniformly sampled manifold mesh. Therefore, as a preprocessing step, we resample the meshes uniformly using ReMesh [1]. Since most models do not consist of a single mesh, we assemble them into a single manifold mesh. We simplify scale invariance for our algorithm by normalizing the surface area to $A_{\mathcal{M}} = 100$.

5.1 Implementation details

In our implementation, we encode the tangent directions in an arbitrary but fixed coordinate system in $\mathbb{R}^3$. Hence, the matrix $\mathbf{A}$ used in Equation (6) has dimensionality $3n \times 3n$, and we encode $\nabla_i \mathbf{p}_{ik}$ and $\nabla_k \mathbf{p}_{ik}$ in $\mathbb{R}^3$. To prevent non-tangent directions to show up as (ϵ, r) -slippage directions, we add quadratic-constraints in the normal direction to the blocks of $\mathbf{A}$.

Both infinitesimal and discrete r -symmetry detection algorithms require geodesic distance information and we compute them as in Brunton et al. [6]. For efficiency, we detect symmetry on a subset of the vertices computed using Poisson disk sampling. The sample spacing s is fixed as $0.1r$ in our experiments.

In our experiments, the threshold used to characterize slippability is fixed to $0.03A_{\mathcal{M}}$ for $r = 0.1R$, where $R = \sqrt{A_{\mathcal{M}}/\pi}$ is the intrinsic radius of the surface. As r increases, we decrease the threshold quadratically to account for the change in area $A_{\mathcal{U}}$ (as in Fig. 2). Note that the threshold can be updated interactively.

Discrete symmetry characterization includes the computation of the wave kernel descriptor in order to prune out points which might not lead to symmetric parts. The parameters used for the wave kernel descriptor computations are $M = 100$ energy scales and number of eigenvalues equal to the rank of matrix being decomposed. To compute the discrete Laplace-Beltrami operator, the approach by Meyer et al. [14] is adopted resulting in a generalized eigenvalue problem which is solved with the Arnoldi method of ARPACK. To project original to subsampled mesh points, an approximate nearest neighbor search [17] is used.

5.2 Infinitesimal Symmetry

We color-code the slippability as follows: white is 0-slippable, red 1-slippable, green 2-slippable and blue 3-slippable. Fig. 2 shows the influence of r on point slippability. As expected, for larger r , points become less slippable.

Fig. 3 shows the influence of the threshold used to characterize slippability. As expected, for larger thresholds, points become more slippable. Note that for all thresholds, some points near the tips of all tentacles are found to be non-slippable, while the eyes are only classified as non-slippable for smaller thresholds. Consider for example a geodesic disk centered at a point between the octopus eyes. It can be thought of as topologically equivalent to a flat plane with two bumps. When big threshold is used then the stretching induced from the bumps is tolerated and therefore is evaluated by the method as a flat plane and

Fig. 2: Each pair shows symmetry size and resulting slippability. From left to right: $r = 0.1R$, $0.2R$ and $0.3R$ with thresholds $0.03A_M$, $0.0075A_M$ and $0.0033A_M$, respectively. Color scheme: white (0), red (1), green (2), blue (3).

colored hence blue. As we decrease the threshold, the stretching is tolerated less and the eyes start getting colored white and characterized as non-slippable.

Fig. 3: Left to right: symmetry size r , point slippability for thresholds $0.03A_M$, $0.003A_M$ and $0.0003A_M$. Color scheme: white (0), red (1), green (2), blue (3).

Fig. 4 shows infinitesimal symmetry detection results for different meshes for $r = 0.1R$. Note that symmetric parts receive the same color and that at this scale, very smooth surfaces only contain slippable points (see top left), while fine scale details of size below r result in 0-slippable points (see palm leaves).

5.3 Discrete Symmetry

We color-code the discrete symmetry results using the symmetry factored embedding distance from a marked point from blue (small distance) to red (large distance). For all of the following experiments, the symmetry radius is $r = 0.1R$, and the threshold for the infinitesimal symmetries is set to $0.03A_M$. The thresholds used to prune point and direction pairs using the wave kernel descriptor were set manually per model to ensure that not too many point and direction matches are processed. We set the thresholds to $\tau = 1e - 3$ and $\tau_d = 1e - 3$ for the octopus model and to $\tau = 1e - 18$ and $\tau_d = 1e - 10$ for the ballet model.

Fig. 4: For each model, left: symmetry size r , and right: point slippability. Color scheme: white (0), red (1), green (2), blue (3).

Fig. 5 shows the symmetry-factored embedding distance from the marked point on the octopus and the ballet models. For the octopus model, the non-slippable points are all near the tips of the tentacles, as shown in Fig. 3 (second from left). Note that for the discrete symmetry result, pairs of points close to the tips of all pairs of tentacles are correctly found to be (ϵ, r) -symmetric and hence, parts of all tentacle tips are shown in blue. For the ballet model, all dancer limbs that are not attached to other parts of the ballet model are found to be non-slippable, as shown in Fig. 4 (last row). Note that for the discrete symmetry result, pairs of points close to the tips of limbs that are not attached to other parts are correctly found to be (ϵ, r) -symmetric and shown in blue.

5.4 Discussion

Our method aims to detect partial infinitesimal and discrete symmetries and hence has different characteristics and applications than prior work that detects

Fig. 5: Symmetry factored embedding distances from marked points.

such symmetries globally [3], [20]. Considering partial symmetries significantly increases the space of possible solutions. We have found that intrinsic slippability is an ill-posed problem, and even more so than extrinsic slippability: Small perturbations in the intrinsic geometry do not cause a large change in the eigenvalues of the corresponding matrix of second derivatives. This in particular makes the slippability results strongly dependent on the threshold.

The presented algorithm requires a number of input parameters. We were able to fix most parameters for our experiments by normalizing the mesh areas. For infinitesimal symmetry detection, the slippage threshold needed to be adjusted manually for few models, which is supported at interactive rates. For discrete symmetry detection, the thresholds for pruning point and direction pairs needed to be adjusted manually per model to ensure computational efficiency.

6 Summary

For an orientable surface in $\mathbb{R}^3$, intrinsic r -symmetries are partial intrinsic symmetries: mappings of geodesic disks of radius r to other parts of the surface, which do not distort the intrinsic geometry of the disk. We have presented a method to extract discrete intrinsic (ϵ, r) -symmetries and characterize the infinitesimal intrinsic (ϵ, r) -symmetries of a shape. The algorithm is motivated by the observation that in the ideal case (no approximation), the set of symmetries can be exhaustively computed by direct sampling. Such symmetries are invariant to isometric deformations of the shape and topological changes such as contacts or holes, which are recognized as boundaries of partiality (see Fig. 5, right).

We have shown that our method can extract discrete intrinsic (ϵ, r) -symmetries and characterize infinitesimal (ϵ, r) -symmetries at each point on a standard symmetry dataset. We leave a large scale evaluation for future work.

Acknowledgments

We thank Tino Weinkauff for helpful discussions. This work has partially been funded by the Cluster of Excellence on *Multimodal Computing and Interaction* within the Excellence Initiative of the German Federal Government.

References

1. Attene, M., Falcidieno, B.: Remesh: An interactive environment to edit and repair triangle meshes. In: *Shape Modeling and Applications* (2006)
2. Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In: *International Conference on Computer Vision Workshops* (2011)
3. Ben-Chen, M., Butscher, A., Solomon, J., Guibas, L.: On discrete killing vector fields and patterns on surfaces. *Computer Graphics Forum* 25, 1701–1711 (2010)
4. Berner, A., Bokeloh, M., Wand, M., Schilling, A., Seidel, H.P.: Generalized intrinsic symmetry detection. Tech. rep., Max-Planck Institute for Informatics (2009)
5. Bronstein, A., Bronstein, M., Kimmel, R.: *Numerical Geometry of Non-Rigid Shapes*. Springer (2008)
6. Brunton, A., Wand, M., Wuhler, S., Seidel, H.P., Weinkauff, T.: A low-dimensional representation for robust partial isometric correspondences computation. *Graphical Models* 76, 70–85 (2014)
7. Gelfand, N., Guibas, L.: Shape segmentation using local slippage analysis. In: *Symposium on Geometry Processing* (2004)
8. Grushko, C., Raviv, D., Kimmel, R.: Intrinsic local symmetries: A computational framework. In: *Eurographics Workshop on 3D Object Retrieval* (2012)
9. Jiang, W., Xu, K., Chang, Z.Q., Zhang, H.: Skeleton-based intrinsic symmetry detection on point clouds. *Graphical Models* 75, 177–188 (2013)
10. Kalojanov, J., Bokeloh, M., Wand, M., Guibas, L., Seidel, H.P., Slusallek, P.: Microtiles: Extracting building blocks from correspondences. In: *Symposium on Geometry Processing* (2012)
11. Kim, V.G., Lipman, Y., Chen, X., Funkhouser, T.: Möbius transformations for global intrinsic symmetry analysis. *Computer Graphics Forum* 29, 1689–1700 (2010)
12. Lasowski, R., Tevs, A., Seidel, H.P., Wand, M.: A probabilistic framework for partial intrinsic symmetries in geometric data. In: *International Conference on Computer Vision* (2009)
13. Lipman, Y., Chen, X., Daubechies, I., Funkhouser, T.A.: Symmetry factored embedding and distance. *ACM Transactions on Graphics* 29, 103:1–12 (2010)
14. Meyer, M., Desbrun, M., Schröder, P., Barr, A.H.: *Mathematics and Visualization 3*, chap. Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. Springer (2002)
15. Mitra, N., Bronstein, A., Bronstein, M.: Intrinsic regularity detection in 3d geometry. In: *European Conference on Computer Vision* (2010)
16. Mitra, N., Pauly, M., Wand, M., Ceylan, D.: Symmetry in 3d geometry: Extraction and applications. In: *Eurographics State of the Art Report* (2012)
17. Mount, D., Arya, S.: ANN: A library for approximate nearest neighbor searching (2010), <http://www.cs.umd.edu/~mount/ANN/>
18. Mukhopadhyay, A., Bhandarkar, S., Porikli, F.: Detection and characterization of intrinsic symmetry. Tech. rep., arXiv 1309.7472 (2013)
19. Ovsjanikov, M., Mériçot, Q., Mmoli, F., Guibas, L.: One point isometric matching with the heat kernel. *Computer Graphics Forum* 29, 1555–1564 (2010)
20. Ovsjanikov, M., Mriçot, Q., Patraucean, V., Guibas, L.: Shape matching via quotient spaces. *Computer Graphics Forum* 32, 1–11 (2013)
21. Ovsjanikov, M., Sun, J., Guibas, L.: Global intrinsic symmetries of shapes. In: *Symposium on Geometry Processing* (2008)

- 22. Raviv, D., Bronstein, A., Bronstein, M., Kimmel, R.: Symmetries of non-rigid shapes. In: International Conference on Computer Vision (2007)
- 23. Raviv, D., Bronstein, A., Bronstein, M., Kimmel, R.: Diffusion symmetries of non-rigid shapes. In: International Symposium on 3D Data Processing, Visualization and Transmission (2010)
- 24. Raviv, D., Bronstein, A., Bronstein, M., Kimmel, R.: Full and partial symmetries of non-rigid shapes. *International Journal of Computer Vision* 89, 18–39 (2010)
- 25. Rinow, W.: Über Zusammenhänge zwischen der Differentialgeometrie im Großen und im Kleinen. *Mathematische Zeitschrift* 35, 512–528 (1932)
- 26. Solomon, J., Ben-Chen, M., Butscher, A., Guibas, L.: Discovery of intrinsic primitives on triangle meshes. In: Eurographics (2011)
- 27. Wang, H., Simari, P., Su, Z., Zhang, H.: Spectral global intrinsic symmetry invariant functions. *Graphics Interface* (2014)
- 28. Xu, K., Zhang, H., Jiang, W., Dyer, R., Cheng, Z., Liu, L., Chen, B.: Multi-scale partial intrinsic symmetry detection. *ACM Transactions on Graphics* 31, 181:1–11 (2012)