

University of West Bohemia
Faculty of Applied Sciences
Department of Computer Science and Engineering

Doctoral thesis

Symmetry Detection in Geometric Models

Declaration

I hereby declare that this doctoral thesis is completely my own work and that I used only the cited sources.

Plzeň, 5th March 2023

Ing. Lukáš Hruša

Acknowledgement

I would like to thank Prof. Dr. Ing. Ivana Kolingerová for patiently supervising my research and providing valuable advice, and without whom this work would have never been possible. I further thank all my coworkers who cooperated with me on my research and also my parents and my dear wife for supporting me throughout my studies.

This work was supported by the Czech Science Foundation, project GACR 21-08009K, Generalized Symmetries and Equivalences of Geometric Data, and by Ministry of Education, Youth and Sports of the Czech Republic, projects SGS-2019-016, Synthesis and Analysis of Geometric and Computing Models and SGS-2022-015, New Methods for Medical, Spatial and Communication Data.

Abstract

Symmetry is a common property of real world objects as well as of artificially created geometric models. The knowledge about symmetry of a given object can be useful in many applications in computer graphics and geometry processing, such as compression, object alignment, symmetric editing or completion of partial objects. In order to use the symmetry of any object in any given application, it first needs to be found. In this work we focus on the task of automatic symmetry detection in 3D objects and we also describe the link between symmetry detection and the problem of registration. Most importantly, we present our own contribution in these fields, namely in rigid surface registration, reflectional and rotational symmetry detection and plane space representation.

Abstrakt

Symetrie je častá vlastnost objektů reálného světa i uměle vytvořených geometrických modelů. Znalost symetrie daného objektu může být užitečná v mnoha aplikacích v počítačové grafice a zpracování geometrie, jako je komprese, zarovnávání objektů, symetrická editace nebo rekonstrukce neúplných objektů. Aby bylo možné symetrii daného objektu jakkoliv využít, musí být nejprve nalezena. Tato práce se zaměřuje na úlohu automatické detekce symetrie ve 3D objektech a také popisuje vazbu mezi detekcí symetrie a problémem registrace. Především však jsou v textu prezentovány naše vlastní příspěvky v těchto oblastech, zejména rigidní registrace povrchů, detekce zrcadlové a rotační symetrie a reprezentace prostoru rovin.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 2 | Background | 6 |
| 2.1 | Object and Symmetry Definition | 6 |
| 2.2 | Symmetry Types | 7 |
| 2.2.1 | Perfect and Approximate Symmetry | 7 |
| 2.2.2 | Global and Local Symmetry | 8 |
| 2.2.3 | Symmetry Type According to Transformation Group | 8 |
| 2.3 | Registration | 11 |
| 2.4 | Object And Other Data Representations | 12 |
| 2.4.1 | Discrete Point Set or Point Cloud | 13 |
| 2.4.2 | Triangle Mesh | 13 |
| 2.4.3 | Plane Representation | 14 |
| 2.4.4 | Others | 14 |
| 3 | Related Work | 15 |
| 3.1 | Reflectional Symmetry | 15 |
| 3.2 | Rotational Symmetry | 25 |
| 3.3 | More General Symmetry | 27 |
| 3.4 | Symmetry Detection Overview | 30 |
| 3.5 | Rigid Surface Registration | 32 |
| 3.6 | Summary | 34 |
| 4 | Consensus Evaluation in RANSAC Surface Registration | 36 |
| 4.1 | Model Registration Algorithm Description | 39 |
| 4.1.1 | Curvature Sampling | 39 |
| 4.1.2 | Candidate Rigid Transformations | 41 |
| 4.1.3 | Analyzing the Space of Rigid Transformations | 42 |
| 4.2 | Transformation Distance Metrics | 43 |
| 4.2.1 | Composed Metrics | 44 |
| 4.2.2 | Compound Metrics | 51 |
| 4.3 | Results | 55 |
| 4.3.1 | Comparing to LCP | 62 |
| 4.3.2 | Non-Centered Object and Density Visualization | 63 |
| 4.3.3 | Noisy Data | 63 |
| 4.3.4 | Comparing to Super4PCS | 63 |

| | | |
|----------|--|-----------|
| 4.3.5 | Limitations | 67 |
| 4.4 | Summary | 68 |
| 5 | Plane Space Representation in Mode-based Symmetry Plane Detection | 70 |
| 5.1 | Background | 71 |
| 5.1.1 | Candidate Creation Algorithm | 72 |
| 5.1.2 | Dependence on Scale and Position | 73 |
| 5.1.3 | Ground Truth | 73 |
| 5.2 | Plane Space Representations | 74 |
| 5.2.1 | Dual Representation in E^3 | 75 |
| 5.2.2 | 4D Vector Representation | 77 |
| 5.2.3 | Transformation Representation | 79 |
| 5.3 | Results | 83 |
| 5.3.1 | Theoretical Comparison | 87 |
| 5.4 | Summary | 88 |
| 6 | Symmetry Plane Detection Using Differentiable Symmetry Measure | 89 |
| 6.1 | Symmetry Measure | 90 |
| 6.1.1 | Similarity Function | 92 |
| 6.1.2 | Efficient Computation | 93 |
| 6.1.3 | Simplification | 93 |
| 6.1.4 | Locating Maxima | 94 |
| 6.2 | Proposed Symmetry Detection Method | 96 |
| 6.2.1 | Creating the Candidate Planes | 97 |
| 6.2.2 | Selecting the Best Candidates | 98 |
| 6.2.3 | Detecting Multiple Planes | 99 |
| 6.2.4 | Using the Weights | 99 |
| 6.2.5 | Using Gaussian Curvature and Normal Symmetry | 100 |
| 6.2.6 | Using Point-to-Point Distance | 102 |
| 6.3 | Results | 103 |
| 6.3.1 | Noisy Objects | 110 |
| 6.3.2 | Tests on a Larger Dataset | 112 |
| 6.3.3 | Detecting Multiple Planes | 113 |
| 6.3.4 | Results of the Modified Versions | 114 |
| 6.4 | Parameters | 116 |
| 6.5 | Another Application of the Symmetry Measure | 119 |
| 6.6 | Limitations | 120 |
| 6.7 | Summary | 120 |

| | | |
|----------|--|------------|
| 7 | Rotational Symmetry Detection Using Reflectional Symmetries and Quaternions | 122 |
| 7.1 | Rotational Symmetry Detection Method | 123 |
| 7.1.1 | Recapitulation of The Symmetry Plane Detection Method | 125 |
| 7.1.2 | Symmetry Measure | 127 |
| 7.1.3 | Candidate Creation and Pruning | 132 |
| 7.1.4 | Optimization | 133 |
| 7.1.5 | Identifying Circular Symmetry | 136 |
| 7.2 | Results | 136 |
| 7.2.1 | Comparisons | 142 |
| 7.3 | Limitations | 145 |
| 7.4 | Summary | 145 |
| 8 | Additional Contributions | 147 |
| 8.1 | Local Reflectional Symmetry Detection | 147 |
| 8.2 | Uses of Registration and Transformation Distance Metrics . | 150 |
| 8.3 | Uses of Symmetry Plane Detection | 152 |
| 8.4 | Other | 156 |
| 9 | Conclusion | 158 |
| A | Activities | 161 |
| A.1 | Publications on International Conferences | 161 |
| A.2 | Publications in Impacted Journals | 162 |
| A.3 | Other Topic-Related Publications | 162 |
| A.4 | Non-Related Publications | 163 |
| A.5 | Participation in Scientific Projects | 163 |
| A.6 | Oral Presentations | 163 |
| A.7 | Other | 164 |
| | Bibliography | 165 |

1 Introduction

Symmetry is a potentially very useful feature which many real world objects exhibit. An object has symmetry if there is an operation or transformation (rotation, reflection, etc.) that maps the object onto itself, i.e. the object is invariant under the given transformation. The information about symmetry is instrumental in a variety of applications. It can be used for object alignment [96] where an object is properly aligned with the x, y, z axes using the symmetry information, or in compression where some parts of the object can be left out to lower the data size and then filled in during decompression using their symmetric counterparts [109]. Another application is symmetrical editing [77] which allows editing a single part of the object while the other symmetric parts are being edited in the same way automatically. Symmetry can also be used to reconstruct incomplete objects [78, 105, 110–112, 117] by filling the missing parts using the information from the symmetric counterparts. However, for any given geometric object, the symmetry information is not known in advance and needs to be first extracted which means that the symmetry first needs to be detected, ideally automatically by a computer. This is the task on which we focus in this work - automatic symmetry detection, i.e. creating algorithms that can be used to automatically find symmetry in geometric objects or models.

Symmetry detection in general is not an easy task. There is never perfect symmetry in real world objects. They only exhibit approximate symmetry and there is no direct way for a computer to decide whether some transformation captures an approximate symmetry or not. It was already pointed out in [127] that symmetry is not a binary feature (is/is not) but rather a continuous one (how much symmetry there is).

When detecting approximate symmetries the goal is mostly to find symmetries that appear natural to a human observer. See, for example, the two human faces in Figure 1.1. The one on the right was created by simply mirroring the one on the left over the red line that passes through its nose and mouth and we can see that these two objects are quite similar so there is obviously some form of symmetry w.r.t. (with respect to) the mirroring operation. But we can also see that the symmetry is certainly not perfect, as can be expected since it is a 3D scan of a real human face and human faces never have perfect symmetry. Anyway, we would probably expect an automatic symmetry detection algorithm to be able to detect such symmetry. This also means that automatic symmetry detection is, in a certain way,




Figure 1.1: A 3D scanned human face [18] and its mirrored version (mirrored over the red line).

tied to how humans perceive symmetry in real world objects. Some experiments regarding the correlation between human and machine perception of symmetry in 3D objects were previously performed, e.g., in [95].

Although there are quite many methods for finding symmetries in geometric data, there are still some challenges and a lot of room for improvement in the field of symmetry detection. Several different types of symmetry exist and geometric objects can have varying shapes and other properties. They can also be damaged in different ways, e.g. by noise or by having some parts missing. No method exists that could reliably and efficiently find symmetries of all the different types in any object. Even the methods that specialize in detecting symmetries of certain types, or a single type, can fail on certain types of objects or objects damaged in certain ways. Some methods only detect strong or perfect symmetries, some only work with specific object representations and some are simply not very fast.

The main goal of this work is to enrich the field of symmetry detection by designing new more reliable and efficient methods or improving the existing ones to tackle the challenges of detecting weak or partial symmetries efficiently and reliably and to make symmetry detection an easier task to solve in general. In some parts of this text we will also talk about the problem of registration which is strongly related to symmetry detection as it can be simply described as the problem of finding symmetry between two objects. Symmetry detection can also be perceived as the process of registering an object onto itself, i.e. symmetry detection can be seen as special case of registration in a way. Therefore, the registration problem also fits well into the symmetry detection topic.

Our review of previous work suggests that the existing approaches are not always generic and flexible which can reduce their reusability. When designing the new methods, we put our effort into this aspect, so that in the future they can be modified or extended for solving more difficult problems or used for different purposes as easily as possible. Our contributions are following.

- **A new consensus evaluation strategy for RANSAC surface registration and analysis of rigid transformation metrics.**

We focused on the popular RANSAC (RANdom SAmple Consensus) scheme for rigid surface registration. We identified some issues with the commonly used consensus evaluation strategy based on overlap measurement including its sensitivity to parameter setting and proposed a different strategy based on **density peak estimation** in the candidate space which turned out to be more reliable. This strategy requires a metric in the space of rigid transformations so we **thoroughly analyzed different metrics** and discovered that the commonly used composed metrics have some theoretical drawbacks, such as depending on the input object’s pose, that we address by employing a metric that measures the distance of two transformations as the **difference of their effect on the input object**. The original version of this metric is not independent of object sampling density which we solved by proposing an improved variant that stems from computing the **transformation effect difference continuously**.

- **Comparison of plane space representations in mode-based symmetry plane detection.**

So-called mode-based approach is quite common in methods for symmetry plane detection in 3D. The results of such methods always, to some extent, depend on how the planes in the solution space are represented. However, it seems that the available literature does not give much guidance regarding the appropriateness of different plane space representations. Therefore, we provide a **theoretical analysis and practical comparison of various plane space representations** in context of mode-based symmetry plane detection including a newly proposed **representation based on dual quaternions**. The reported results can be useful for anyone designing new methods in the future.

- **A new method for symmetry plane detection.**

In general, there are quite a few existing methods for symmetry plane detection but they all have their limitations. Robust methods that could reliably find weak symmetries in general point sets (clouds) while being also robust to noise, computationally efficient and flexible still seem to be missing or their robustness, flexibility or computational efficiency is at least questionable. We address this issue by proposing a **new symmetry plane detection method** for point sets based on **gradient-based numerical optimization** of a simple **differentiable symmetry measure** in combination with **fast point cloud simplification** and **sparse sampling of the solution space**. The proposed method appears to outperform previous approaches in various aspects and seems to fulfill all the above mentioned conditions.

- **A new method for rotational symmetry detection.**

Unlike symmetry plane detection, the options for methods detecting rotational symmetries in 3D objects are quite sparse and those that are both general and robust to weak symmetries seem to be missing completely. To address this problem we propose using **pairs of potential symmetry planes to find rotational symmetries** in 3D point sets (clouds). We use this idea to design a **new method for rotational symmetry detection** by extending the previously proposed symmetry plane detection method. In this process we needed to find a differentiable parameterization of the rotation transformation in order to retain the possibility of using gradient-based optimization. We solved this by proposing a simple **quaternion-based parameterization** which stems from expressing the quaternion rotation formula in the vector space using elementary algebraic operations. To our knowledge, the resulting method is the most general existing method for rotational symmetry detection in 3D objects while also being robust to weak symmetry and noise and considerably fast.

- **Other small contributions.**

We also proposed a **new method for local reflectional symmetry detection** in 3D point sets. Some products of our research, mostly the symmetry plane detection method, were used/modified/extended by other researchers in their work on which the author of this thesis was often cooperation/consulting. This includes reflectional symmetry detection in 2D curves, detection of the spherical surface of reflectional

symmetry in 3D objects, reflectional symmetry detection in earth observation data, generalization of our symmetry plane detection method for non-uniform data and using the symmetry plane detection method for constructing the cutting plane of a human face or predicting the water level of a lake. Our previous work also helped the graphics group at the University of West Bohemia, Department of Computer Science and Engineering receive a research grant GACR 21-08009K, Generalized Symmetries and Equivalences of Geometric Data, Czech Science Foundation.

The rest of this text is organized as follows. Chapter 2 provides necessary background including definition of symmetry in geometric data and description of different symmetry types. Chapter 3 describes previous work in the field of symmetry detection and briefly in registration. Chapters 4 to 7 present our main contributions, first focusing on registration as a more general problem (also because parts of this research are referenced in the later chapters) and then we concentrate on the more specific areas of symmetry detection itself. Several other smaller contributions that are worth mentioning but are not significant enough to be contained in their own chapters are in Chapter 8 including a new method for local reflectional symmetry detection in point sets. Finally, Chapter 9 concludes the thesis.

2 Background

In this chapter we provide some background by defining symmetry in Euclidean data and describing different symmetry types in the 3D space. Then we also describe the task of registration which is strongly related to symmetry detection and we mention different representations of 3D objects.

2.1 Object and Symmetry Definition

We define a general object in an arbitrary Euclidean space E^d , where d is a positive whole number, in the following way.

Definition 1 *A set of points X is an object in E^d if $X \subset E^d$ and $X \neq \emptyset$.*

This means that we consider any non-empty set of points in E^d an object in E^d , regardless of whether the set is discrete or continuous. For example, an object in E^2 can be a 2D line or curve or any discrete set of 2D points. An object in E^3 can be a 3D line, 3D curve, a plane, a surface or a discrete set of 3D points.

It seems unclear whether a unique mathematical definition of symmetry is possible and whether all types of symmetry can be covered by a single definition [93]. However, symmetry of a geometric entity can be described as an invariance under a given geometric transformation. Therefore, for symmetry in d -dimensional Euclidean space we state a definition based on the one proposed in [93]:

Definition 2 *Having a function $F(\mathbf{x})$, $\mathbf{x} \in E^d$ and a transformation $T(\mathbf{x}) = \mathbf{y}$, $\mathbf{x}, \mathbf{y} \in E^d$ then F has symmetry w.r.t. the transformation T if $F(T(\mathbf{x})) = F(\mathbf{x})$ for all $\mathbf{x} \in E^d$.*

To make this definition applicable to objects, considering an object $X \subset E^d$, we can simply set F as

$$F(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \in X \\ 0 & \text{otherwise} \end{cases} . \quad (2.1)$$

Then we can alternatively define symmetry of an object in an arbitrary Euclidean space as follows.

Definition 3 Given an object $X \subset E^d$, $X \neq \emptyset$ and a transformation $T(\mathbf{x}) = \mathbf{y}$, $\mathbf{x}, \mathbf{y} \in E^d$ the object is symmetrical w.r.t. the transformation T if $T(\mathbf{x}) \in X$ for any $\mathbf{x} \in X$.

In other words, T represents symmetry of X if, after applying T on any point of X , the point ends up in any other point of X or in itself, meaning the transformation maps the object onto itself. When F is set as in Equation 2.1 this definition is equivalent to the first and more general one.

This work elaborates on the task of automatic computer symmetry detection. Although the definition of symmetry in the previous section applies for objects in a Euclidean space of arbitrary dimension, we mostly focus on detecting symmetries in 3D objects. Therefore, unless stated otherwise, in the rest of this work we only consider symmetries and objects in E^3 .

2.2 Symmetry Types

Symmetries in 3D objects can be classified based on different criteria some of which are described in this section.

2.2.1 Perfect and Approximate Symmetry

The symmetry definition in Section 2.1 is a definition of perfect symmetry where each point of the object is mapped exactly onto another point. However, perfect symmetry never occurs in real world objects and is quite rare even in artificially created digital models. In practice, objects only exhibit approximate symmetry where no transformation exists such that it maps the object onto itself perfectly but there might be transformations that map some portion of the object points "considerably close" to other points of the object. Therefore, when talking about symmetry detection in 3D objects, we usually mean detection of approximate symmetries, not perfect ones. Obviously, for approximate symmetries the conditions in the symmetry definition need to be relaxed in some way but no strict definition of approximate symmetry can be formed which in turn means that there is no exact way of deciding whether a transformation captures some approximate symmetry or not. This is what generally makes approximate symmetry detection quite a difficult task because we do not have an objective way for a computer to decide what is still perceived as approximate symmetry and what is not symmetry anymore.

We will use the term *strong symmetry* for approximate symmetry that is close to perfect symmetry, and the term *weak symmetry* for symmetry

farther from perfect. Note, however, that the terms *weak* and *strong* do not serve as exact definitions but rather to differentiate between the more and less approximate symmetries in the text in an intuitive way.

2.2.2 Global and Local Symmetry

If an object has symmetry as a whole, it is called global symmetry. However, very often there is only symmetry in some parts of the object, i.e. a transformation exists that maps some part(s) of the object onto another part(s), which is called local symmetry. There is again no strict definition of local symmetry and both global and local symmetries are usually approximate in practice.

2.2.3 Symmetry Type According to Transformation Group

When detecting symmetries in 3D objects we are looking for transformations that represent these symmetries and these transformations always belong to a certain transformation group. Depending on the transformation group the symmetries in E^3 can be classified into several types.

Reflectional Symmetry

Reflectional symmetries (sometimes called mirror symmetries or bilateral symmetries) are represented by transformations from a group of reflection transformations, i.e. transformations that perform reflection over an arbitrary plane. A reflectional symmetry is usually and most easily described by a single plane of reflection. A plane that captures reflectional symmetry is called a symmetry plane. This type of symmetry is probably the most often occurring symmetry type in both real world and artificial objects which makes reliable methods for detecting symmetry planes in 3D objects quite desirable. For examples of reflectional symmetries see Figure 2.1. The objects are all aligned in such a way that the symmetry plane, which is marked by the red line, is perpendicular to the image plane. Reflecting these objects over their symmetry planes would approximately map them onto themselves, therefore, these planes represent their approximate reflectional symmetries. The objects in Figure 1.1 are also examples of reflectional symmetry.




Figure 2.1: Examples of reflectional symmetry (figures taken from [53]).

Rotational Symmetry

Rotational symmetry is represented by a transformation that performs rotation by a given angle around a given axis. Such symmetry can be described by a rotation axis and a scalar value which defines the rotation angle. A special type of rotational symmetry is circular symmetry which can be represented only by the rotation axis where the symmetry holds for any rotation angle. A single circular symmetry basically represents infinitely many rotational symmetries. Some examples of approximate rotational and circular symmetries are in Figure 2.2. The top row shows the rotation axis for each object (marked by the black line) and the bottom row depicts the rotation angle (the red arc).




Figure 2.2: Examples of approximate rotational symmetry (figures taken from [49]).

Rigid Symmetry

Rigid symmetries are one of the most general symmetries that commonly occur in geometric models and they are represented by rigid transformations, i.e combinations of rotations and translations. Rigid symmetry can be described by a rotation matrix \mathbf{R} and a translation vector \mathbf{t} . Often symmetry is considered rigid even if the transformation contains reflection and in such a case \mathbf{R} can be either a rotation matrix or a product of a rotation matrix and a reflection matrix. When talking about rigid symmetries in 3D objects we often, but not always, consider local symmetries. Example of an object with local rigid symmetries can be seen in Figure 2.3, the symmetric parts are in the frames of the same color.



Figure 2.3: Example of an object with local approximate rigid symmetries (figure taken from [70]).

Other Symmetry Types

There are also other symmetry types such as point symmetry, which is represented by a reflection over a single point, or translational symmetry represented only by a translation. We can also consider more general symmetries represented by rigid transformations with additional uniform scaling or by general affine transformations. We could even have symmetries described by elastic transformations that can bend and stretch the object differently in different places which includes transformations that perform reflection over an arbitrary surface or curve. Such symmetries are hard to generate artificially or to even imagine, at least in 3-dimensional space. This is why we only show a 2D example in Figure 2.4 where reflection of the black object over the red curve can be considered its approximate symmetry.

If we do not put any constraints on the transformation type then we can basically for any object construct a transformation that captures its symmetry. These very general transformations are, however, very difficult to represent and have many degrees of freedom. Therefore, detecting such general symmetries is very hard.




Figure 2.4: Example of a general elastic symmetry in 2D - the black object has approximate reflectional symmetry w.r.t. the red curve.

Symmetries can be further classified as either extrinsic or intrinsic. Extrinsic symmetries are represented by geometric transformations in the Euclidean space. Until now we were only talking about extrinsic symmetries and all the symmetry types we have mentioned are types of extrinsic symmetry. Intrinsic symmetry is rather specific and does not really satisfy our definition of symmetry, but, just for completeness, we mention it anyway. Intrinsic symmetries are specific to surfaces and are described by intrinsic transformations on the surface rather than extrinsic transformations in the Euclidean space in which the surface is embedded. Usually intrinsic symmetries are defined as transformations that maintain geodesic distances between all points on the surface. Such symmetries are invariant under deformation of the surface, such as bending or folding, as long as the deformation does not change the geodesic distances of any points. There are several methods that are somewhat capable of detecting intrinsic symmetries, e.g. [87, 121, 122], however, this area is distinct from the focus of our research and by finding a good intrinsic symmetry, we do not necessarily get a meaningful extrinsic symmetry of the desired type. Also, intrinsic symmetries are represented by point-to-point or part-to-part correspondences which are rather impractical representations for extrinsic symmetries, which can usually be represented by simple geometric transformations. Therefore, in the rest of this work we will not be considering intrinsic symmetries in any way.

2.3 Registration

Registration is a field strongly related to symmetry detection. Having two input objects, the task of registration is to find a transformation that maps

one of the two objects onto the other one. So, only instead of mapping an object onto itself we are trying to map an object onto another object, therefore, registration can basically be understood as detecting symmetry between two objects. Another way of seeing this is that symmetry detection is a special case of registration where we are registering an object onto itself (the two input objects are the same), i.e. that registration is a more general extension of the symmetry detection task. Either way, symmetry detection and registration have very much in common and many principles and approaches used in registration methods can be also used in symmetry detection and vice versa.

Therefore, given the relation between these two areas, we are also going to consider registration in some parts of this work. Specifically, we will talk about rigid surface registration which is a task of finding a rigid transformation that maps a 3D surface onto a different 3D surface in such a way that some portion of the two surfaces overlaps. This is usually used for completing objects from partial 3D scans. An example of a correct result of a registration process can be seen in Figure 2.5.




Figure 2.5: Example of a correct registration result - (c) of the two surfaces - (a), (b) (figure taken from [46]).

2.4 Object And Other Data Representations

When trying to detect symmetries in 3D objects, we first need to define what representation of an object will be on the input. Throughout this text, we will work with two different representations of general 3D objects - a discrete point set (point cloud) and a triangle mesh. Here we describe these representations and also briefly mention other representations of objects and of some non-object geometric data.

2.4.1 Discrete Point Set or Point Cloud

A discrete point set, often called a point cloud, is very commonly used in geometry processing to represent a sampled 3D surface but it can also represent a sampled volume. Using the discrete point set representation, any 3D object X can be simply defined as

$$X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}, \mathbf{x}_i \in E^3, i = 1, \dots, n$$

where n is the number of the points in the point set. Discrete point sets are usually outputted e.g. by 3D scanners.

2.4.2 Triangle Mesh

The triangle mesh representation is very commonly used in computer graphics to describe a surface. A triangle mesh is often understood as an approximation of a continuous surface by a piece-wise planar surface where the planar pieces are triangles. Having an object X , to use the triangle mesh for its representation we need a set of vertices $X_v = \{\mathbf{x}_i \in E^3\}, i = 1.. \nu_X$ and a set of triangles $X_t = \{t_i^X\}, i = 1.. \tau_X$. For each triangle $t^X \in X_t$ three indices $k1, k2, k3$ are stored and the triangle is defined by the three vertices $\mathbf{x}_{k1}, \mathbf{x}_{k2}, \mathbf{x}_{k3} \in X_v$. The object is then defined using the triangle mesh representation as

$$X = \bigcup_{i=1}^{\tau_X} t_i^X$$

i.e. it is a union of all the triangles. We could, of course, only have the set of triangles without the set of vertices while describing the same exact object. This, however, would not be a triangle mesh but only a triangle set. Unlike the triangle set, the triangle mesh allows easily extracting adjacency information which is required for many operations, such as traversal or estimating differential quantities (e.g. curvature), which would be much more difficult to perform without it.

The triangle mesh is a more complex representation of an object than the discrete point set because it contains more information. A discrete point set can always be easily derived from a triangle mesh simply by taking its vertices or using some surface sampling technique (e.g. [3, 13]). But transforming a discrete set of points into a triangle mesh is not so straightforward and is considerably more difficult (various mesh reconstruction methods exist however, e.g. [73, 125, 130]). This also means that methods that are capable of working with the discrete point set representation are generally more widely applicable than those that require a triangle mesh on the input.

2.4.3 Plane Representation

In some parts of this text we will also work with planes in E^3 , therefore, here we describe how to represent them. Unless stated otherwise, we define an arbitrary plane P by its implicit equation $P : ax + by + cz + d = 0$ where $\mathbf{x} = [x, y, z]^T \in E^3$ is a general 3D point, a, b, c, d are the coefficients that define the plane and the points $\mathbf{x} \in E^3$ that satisfy this equation are the points of the plane P . The vector $[a, b, c]^T$ describes the normal of the plane and the d coefficient encodes the distance of the plane from the origin of the coordinate system. This representation is probably the simplest one but it is also ambiguous since for any real number $k \neq 0$, the coefficients a, b, c, d define the same plane as coefficients ka, kb, kc, kd .

2.4.4 Others

General objects in E^3 can also be represented in other ways, e.g. as parametric or implicit surfaces, parametric curves or binary functions which can represent volumes by having a value of 0 for points outside the object and 1 for points inside the object. Objects can also have values associated with each of their points. These values can be scalars or vectors and they can represent various quantities, such as colors or results of some physical measurements (temperature, pressure, etc.). When considering symmetries of such objects, sometimes, but not always, it is desired that apart from the positions of the points these associated values are also taken into account. In such a case Definition 3 of (perfect) symmetry of objects is no longer applicable and the more general Definition 2 needs to be used with the function F defined accordingly.

There are also other types of geometric data that do not satisfy our definition of an object but appear commonly in certain areas of computer science. We can have a general volumetric function which associates a value with each point in E^3 or in some subspace of E^3 . Or we can have a discrete volumetric grid where each cell has a value assigned to it - this is mainly used to represent medical images outputted by CT or MRI scanners. For such data we also need to use Definition 2 when talking about symmetries.

However, the representations mentioned in the previous two paragraphs will be used in this work only rarely.

3 Related Work

In this chapter we provide information about various existing symmetry detection methods. We mention some only briefly while some, which seem more important or interesting, will be described in more detail. Since the methods are quite diverse and based on various approaches it is quite difficult to find some way to classify them in order to make the chapter easier to read. Therefore, we divide the methods only according to the type of symmetry they are designed to detect. At the end of the chapter we also briefly describe several existing methods for rigid surface registration. Some parts of this chapter were taken from the author’s master’s thesis [43].

3.1 Reflectional Symmetry

As already mentioned, the most common type of symmetry being solved for 3D objects is reflectional symmetry. Such symmetry is captured by a single plane of reflection and there are many strategies of computing such a plane.

Some methods detect only planes that pass through some reference point such as the origin, centroid of the input object or its center of mass. Such methods are generally not very good at detecting weaker symmetries because weakly symmetrical objects can have the symmetry plane in an arbitrary position, not passing through the expected reference point.

Sun and Sherrah [115] proposed symmetry plane detection using orientation histogram and their method detects only planes that pass through the origin. The method uses the discrete version of the Extended Gaussian Image [42] called orientation histogram. It is obtained by dividing a unit sphere into hexagonal bins with values assigned according to the number of normal vectors of the input object facing in the given bin’s direction. An example of the orientation histogram of a given 3D model is shown in Fig. 3.1.

Several candidate planes are selected such that they pass through the origin and their normal vectors are facing approximately in the directions of the principal axes of the object. For each candidate the histogram is reflected over the given plane and a correlation of the reflected histogram with the original one is computed. The plane with the highest correlation is then selected as the strongest symmetry plane. The method can also detect rotational symmetries as will be mentioned in the next section.

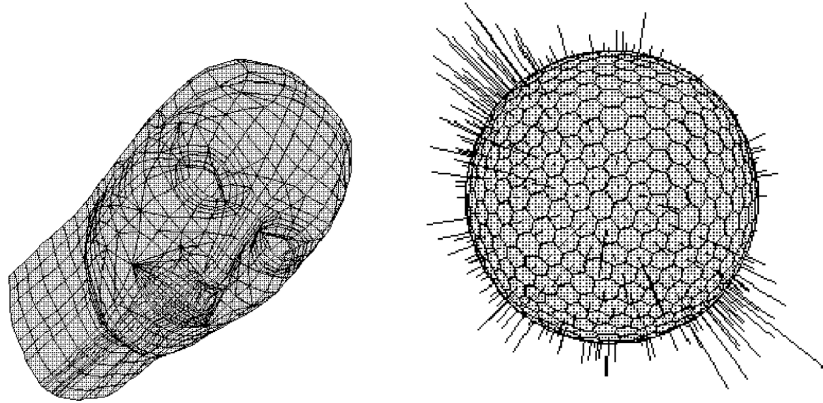


Figure 3.1: A 3D model of a human head (left) and its orientation histogram (right) (figure taken from [115]).




Figure 3.2: Symmetric ant model (left) and the star-shaped surface created by its approximation using the spherical harmonics (right) (figure taken from [58]).

Kakarala et al. [58] proposed a method that uses approximation of the input object with spherical harmonics, creating a star-shaped surface, and the symmetry plane detection is then performed on this new shape. Fig. 3.2 shows an example of the spherical harmonics approximation for a symmetric ant model.

The symmetry detection is based on the observation that if a real-valued function has symmetry across the origin then its Fourier transform is real-valued. The authors apply this observation on the spherical harmonics to derive an error function whose minimization leads to the symmetry plane. This method only detects planes that pass through the origin.

Korman et al. [60] designed a method that uses a distortion meas-




Figure 3.3: The viewpoint entropy distribution sphere for four different objects generated using 2562 viewpoints where the color represents the entropy value - blue: large entropy, green: mid-size entropy, red: small entropy (figure taken from [69]).

ure representing the amount of mismatched volume between the original shape and the transformed shape. The method detects only planes that pass through the origin but it can also detect rotational symmetries (with axes also passing through the origin). Since the evaluation of the distortion measure is too complex, the authors propose a randomized sampling procedure that gives approximately the same results with overwhelmingly high probability. The symmetry of an input object is then found by clever sampling of the transformation group based on the observation that the upper bound on the sampling density can be controlled by the maximum allowed distortion of the symmetry. The distortion is evaluated in each sample and the one with the lowest distortion can be selected as the strongest symmetry.

Li et al. [69] proposed a view-based method for symmetry plane detection on triangle meshes. A set of viewpoints is sampled on a sphere around the input model and in each viewpoint a camera is set to look at its center. Using these cameras the model is rendered from each viewpoint using orthogonal projection. The authors employ a viewpoint entropy which can be computed for each rendered view and depends on the areas of the faces that were rendered in the particular view. This entropy is computed for each viewpoint which creates a viewpoint entropy distribution sphere. Fig. 3.3 shows the viewpoint entropy distribution sphere for several objects generated using 2562 viewpoints.

The symmetry plane detection itself is then based on the observation that

the symmetry planes of the model and of the viewpoint entropy distribution sphere are the same. Therefore, the symmetry plane detection is performed on the viewpoint entropy distribution sphere which has the advantage of being independent of the vertex count of the input mesh. The authors only use 42 viewpoints for generating the distribution sphere and claim that it is enough. Candidate planes are created using pairs of viewpoints with matching values of the viewpoint entropy. For each candidate plane the rest of all pairs of viewpoints with matching entropy are verified to see whether they are symmetric with respect to the given plane. If the number of symmetric pairs is great enough the given candidate plane is declared a symmetry plane. This method only detects planes passing through or very near the center of the input object’s bounding sphere. However, it seems to outperform previously designed methods in terms of accuracy, speed and robustness to noise.

The following methods do not put any constraints on the detected planes and most of them can be used for detecting weaker symmetries.

Schiebener et al. [105] designed a method which detects the plane of symmetry of a point cloud and they use it for completion of partial 3D scans. Apart from the point cloud representing the input object, the method also needs a point cloud representing the object’s surrounding and the position from where the object was scanned. It relies on the fact that a 3D object usually stands on some kind of supporting structure, such as the ground, and the first step is detecting the supporting plane which represents such a structure. Candidates for these supporting planes are created using a RANSAC-based plane fitting algorithm and for each of these candidates, several symmetry plane candidates are created by sampling the space of planes orthogonal to the given supporting plane. The authors define rating of the candidate symmetry planes which uses the supporting plane and the scanner position to evaluate whether points of the input object, when reflected over a given candidate symmetry plane, end up in plausible locations. For example, the reflected points should appear above the supporting plane and mainly in occluded areas. The plane with the highest rating is declared the object’s plane of symmetry.

Fig. 3.4 shows an example scene containing a symmetric object in the middle with the estimated supporting plane and the estimated symmetry plane. This method only detects symmetry planes which are orthogonal to the supporting surface and, because it requires the surrounding and scanner position, it is only usable in specific applications.




Figure 3.4: A scene with a symmetric object (red) and its estimated symmetry plane (yellow vertical) and the estimated supporting plane of the scene (yellow horizontal)(figure taken from [105]).

Simari et al. [109] presented an algorithm for detecting local reflectional symmetries of 3D triangle meshes. First, a weighted covariance matrix and a weighted centroid are computed, with the weights set for each vertex according to the areas of its adjacent triangles. Eigenvectors of the matrix are used to define three orthogonal planes that pass through the centroid. For each of the planes all vertices are reflected over it and its cost is computed based on the distances of the reflected vertices from the original mesh (using the minimum point-to-triangle distance). Only the plane with the lowest cost is kept and a support region is defined as the region of the mesh which has a strong enough symmetry w.r.t. to this plane. Next, the weights of all vertices are updated to contain information of how far from the mesh the given vertex is when reflected over the plane - the farther, the smaller its weight. Also, weights of all vertices outside of the support region are set to 0. A new covariance matrix and a new centroid are computed using the new weights, a new plane is estimated and a new support region is found. These plane estimation and region finding steps are iterated until convergence is achieved. The vertices in the support region of the final plane are removed and the whole process is repeated for the remaining components of the mesh. This way a tree structure can be created called a folding tree which the authors propose to use for mesh compression.

Fig. 3.5 depicts three objects with the symmetry planes detected by this method in the process of creating the folding tree.




Figure 3.5: Symmetry planes detected on three different objects during the folding tree creation (figures taken from [109]).

Podolak et al. [96] proposed a Planar Reflective Symmetry Transform (PRST) for volumetric functions which associates each plane in the space of planes with a value determining how symmetrical the function is w.r.t. that plane. This value is defined as the distance of the function to the closest function which is perfectly symmetrical w.r.t. to the plane. The transform can be used to find the plane of symmetry of a discrete volumetric function represented by a uniform 3D grid. Fig. 3.6 shows several examples of PRST for 2D binary functions where the transform works as well only for 2D reflection lines instead of 3D reflection planes. The object in each image is represented by the black contour and the strength of each reflection line is determined by its shade of grey - the darker the line, the stronger symmetry of the object it represents in terms of the PRST.

The brute-force computation of the PRST has time complexity $\mathcal{O}(n^6)$ for a grid of $n \times n \times n$ voxels but the authors propose a way of finding the plane with the highest PRST in $\mathcal{O}(n^4 \log(n))$. To use this approach for symmetry detection on a sampled surface (triangle mesh, point cloud) the authors use the Gaussian Euclidean Distance Transform to convert the surface to a volumetric grid. To lower the computation cost for the surface data, the authors propose using a randomized Monte-Carlo algorithm to approximate the PRST. The space of planes is divided into discrete bins based on their normal orientation and distance from the origin. Pairs of points are randomly being selected, for each pair the symmetry plane of the two points is constructed and it votes for the plane represented by the corresponding bin. The planes with the largest vote count represent the most significant symmetry planes of the input object. Since these results are not necessarily very precise, the authors propose a final iterative refinement.




Figure 3.6: Visualization of the PRST in 2D examples - darker lines represent stronger reflectional symmetries in terms of the PRST, each image represents a binary function where the object is defined by the black contour (figures taken from [96]).

Although this approach works with sampled surface data, it still requires rasterizing the input object, which can be computationally expensive and seems unnecessary.

Caillière et al. [16] used Hough transform to detect symmetry planes of triangle meshes. This method is actually a modification of the more general clustering approach [83] that will be described in Section 3.3. It is overall also quite similar to the Monte-Carlo approach used in the previous method by Podolak et al. [96]. Points are paired based on similar curvature values and for each pair its symmetry plane is constructed. These planes vote for the planes represented by the corresponding bins in the space of planes and, in the end, the plane with the largest vote count is selected as the symmetry plane of the input mesh.

Speciale et al. [114] employ two symmetry plane detection methods for 3D shape reconstruction. Both the methods take depth information in voxel grid as input and one of them is essentially the implementation of the above described method of Podolak et al. [96]. The second method is a little different. High-gradient and high-curvature voxels are extracted and referred to as surface voxels. Pairs of surface voxels are randomly sampled and for each pair a symmetry plane candidate is constructed. Each candidate is



Figure 3.7: Example of the correspondence creation. Left image shows an input surface (the black curve) and a plane in a given iteration of the ICP-like process (the green line S). Right image shows also the reflection of the original surface over the plane (the grey curve) and the created correspondences (the small blue lines) (figure taken from [29]).

evaluated according to the number of inliers. A surface voxel is considered an inlier if, when reflected over the given candidate plane, it ends up in another surface voxel or in space with unknown voxel information where a potential occluded part of the symmetric object could be. Planes with the largest inlier count then represent the strongest symmetries. The authors claim that this approach, although similar to [96], can potentially be faster and has lower memory requirements than [96].

The methods of **Combès et al.** [26] and **Ecins et al.** [29] both use iterative ICP-like (see Section 3.5 for ICP) approaches for symmetry plane detection in point clouds. Such approaches can be roughly summarized as follows. Initial plane or planes are estimated in some way and for a given initial plane an iterative process is executed. In each iteration the points of the object are reflected over the plane and those that reflect too far from other points of the object are usually considered outliers and ignored in that iteration. For each of the remaining reflected points, the closest point of the object is found and a correspondence between the two points is created. Next, the plane is adjusted such that the corresponding pairs of points have the best possible reflectional symmetry w.r.t. the new plane - least square error is usually used for this purpose. These steps are iterated until convergence and the final plane usually represents some significant symmetry of the input object.

Fig. 3.7 demonstrates the correspondence creation process - the black curve represents an input surface, the green line denoted S is the reflection plane in a given iteration, the grey curve in the right image is the reflection of the surface over S and the small blue lines show the found correspondences.

Sipiran et al. [112] proposed a method specifically designed for global symmetry plane detection on incomplete objects represented by triangle meshes. First step is to detect local features of the given 3D shape which is done using the theory of heat diffusion on manifolds. A function is defined which associates the accumulation of heat up to a given time to each point on the surface. This function is computed using the eigenvalues and eigenvectors of the Laplace-Beltrami operator and is called a Heat Kernel. As the feature points, local maxima of the heat accumulation function are taken. Pairs of the feature points then generate the candidate symmetry planes but only pairs of points which have a similar value of the heat accumulation function are considered. The last stage of this algorithm is a voting process where other pairs of points are tested against the candidate planes. The more pairs of points are considered to be symmetrical with respect to a given plane the more votes the plane gets. There are several criteria designed by the authors, which are used to decide whether or not a plane is considered a plane of symmetry of a given pair of points. Also only points in which the mean curvature is larger than some threshold are used in the voting process. In the end the plane with the highest vote count can be declared the resulting plane of symmetry.

Fig. 3.8 shows the pipeline of the overall symmetry detection process. This method seems to provide very good results even when used on objects with very high level of missing parts. However, it only works on manifold triangle meshes and it does not work on featureless objects, which is quite constraining.

Cicconet et al. [22] solve the symmetry plane detection problem as a registration problem. They propose reflecting the input object over a fixed plane and then using some rigid registration algorithm to map the reflected object onto the original one. This approach results in a transformation that represents a reflection over some plane followed by some rigid transformation. Such a transformation does not generally represent reflection over a plane, so it needs to be extracted approximately from it, which seems unnecessarily more complex than finding the reflection plane directly. Furthermore, as elaborated in [99], the registration algorithms sometimes do not work well on symmetrical objects.

Most recently, **Nagar and Raman** [88] proposed using linear pairing and optimization on manifold to find the symmetry plane of a point cloud. This method works in a space of arbitrary number of dimensions, but due to time complexity $\mathcal{O}(n^{3.5})$ for n points, it is considerably slow. The same

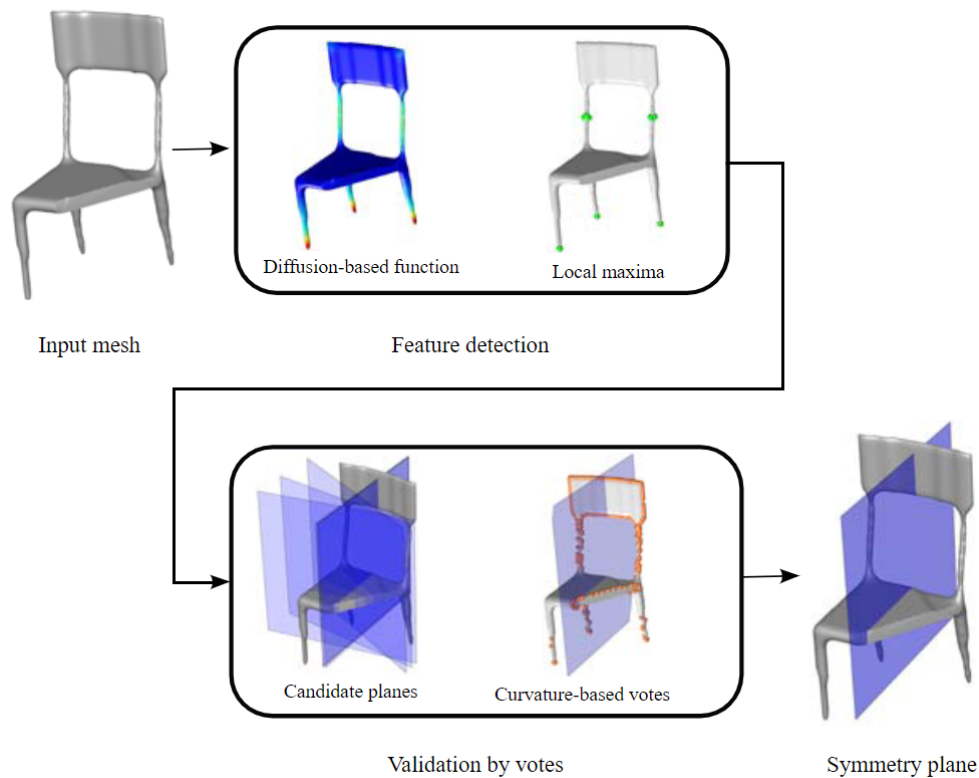


Figure 3.8: Pipeline of the symmetry detection process using the heat diffusion-based feature function and the voting approach (figure taken from [112] and modified).

authors then proposed a closed form solution to the symmetry plane detection problem with randomized initialization [89]. Apparently, at some stages this method expects the symmetry plane to pass through the center of mass of the input object which could possibly be problematic for some weakly symmetrical data.

There are also methods that approach symmetry plane detection by employing machine learning techniques, usually using neural networks. Some of the representatives of this approach include the methods of **Ji and Liu** [57], where the output of the network is followed by an ICP-like refinement, or **Gao et al.** [37] and **Shi et al.** [106] which can also detect rotational symmetries. The biggest advantage of such methods in general is that applying the trained neural network is usually very fast, especially when using a GPU implementation. On the other hand, they naturally require a set of training data and a training process, usually very slow, to be executed before any symmetry detection can be performed, which considerably limits their

applicability. Also, when trained on a given dataset, it is questionable how the network would perform on objects that significantly differ (in terms of shape, symmetry strength, noise level, sampling density, etc.) from objects in the training set.

3.2 Rotational Symmetry

Some of the methods mentioned above can also be used to detect rotational symmetries, but they exhibit the same problems and limitations for this purpose as for the symmetry plane detection. Specifically, the method of **Korman et al.** [60] finds rotational symmetries when the sampling of the transformation space includes rotations but it can only find rotations with axes passing through the origin. The method of **Sun and Sherrah** [115] can find rotational symmetries by using the approximate principal axes as rotation axis candidates instead of normal directions for symmetry plane candidates. For each of the candidate axes a set of rotation angles is sampled providing a set of candidate rotation transformations which are then evaluated in the same way as the reflection ones. This method can also find only rotations with axes passing through the origin. Some additional methods for rotational symmetry detection are mentioned in the following text.

Martinet et al. [77] designed a method for detection of rotational symmetries with possible reflections in triangle meshes using generalized moment functions. It is based on the observation that the generalized moments of a shape have at least the same symmetries as the shape itself and the symmetries of the moments can be computed efficiently. This method only detects symmetries with rotation axes passing through the input object's center of mass but it can also find pure reflectional symmetries (symmetry planes).

Gothandaraman et al. [39] proposed a method that rotates the input object by several fixed angles around the x , y , z axes and creates a 2D image of the object for each rotation by projecting it onto a fixed plane. SIFT [75] features are extracted for each of these images and are used to find a match among them. The information about matching images of different rotations of the input object are used to find its rotational symmetry with possible reflection including purely reflectional symmetry. This method does not seem to be very robust to approximate symmetries and only finds rotations with the axis passing through the origin.




Figure 3.9: Circles detected on a partial 3D rotationally symmetrical object (left) and the symmetry axis detected based on the circle set information (right)(figures taken from [110]).

Sipiran proposed two methods [110, 111] for rotational symmetry detection in partial 3D objects based on finding circular structures in the data (see Fig. 3.9). However, they only seem to detect symmetries of rotational shapes and do not detect discrete symmetries where the symmetry only holds for certain angles of rotation, i.e. they can only detect circular symmetries.

Ecins et al. [31] designed a symmetry-based segmentation method which combines symmetry detection steps with segmentation steps. The symmetry detection steps include detection of reflectional and rotational symmetries where the approach used for finding the reflectional symmetries seems similar (or identical) to their method [29] mentioned in the previous section.

The rotational symmetry detection is performed by minimizing an error function which for a given axis is computed as follows. For each point of the input object a plane is created such that it contains the axis and the given point. Then the angle between the normal in this point and the plane is computed (if its value is larger than 45° then 45° is used instead). These angle values computed for all the points are summed together providing the error function value. The authors claim that an axis that minimizes the error function can be found, for example, using the Levenberg-Marquardt [84] solver. The angle computation for a single point is detailed in Fig. 3.10a.

Three axes passing through the object's center of mass in the directions of its principal axes (see Fig. 3.10b) are used to initialize the solver and after the optimizations are finished the best of the three results is selected




Figure 3.10: The left image details the process of angle computation in the error function for a single point p with normal n when given an axis S - the angle is computed between the normal n and the plane spanned by S and p . The middle image depicts the three initial axes of a partial rotationally symmetrical object before optimization and the right image shows the best of the three axes after optimization (figures taken from [31]).

as the resulting symmetry axis (see Fig. 3.10c).

This method seems to work well for objects with approximate symmetries and even missing parts but it can only find the axis of rotational symmetry and only works for rotational shapes (it only finds circular symmetry). Also, it can be used for point clouds but a normal needs to be available in each point.

As already mentioned in the previous section, even rotational symmetries can be detected using machine learning techniques. The methods of **Gao et al.** [37] and **Shi et al.** [106] can also find rotational symmetries using neural networks. The advantages and problems of methods based on neural networks are the same regardless of whether detecting rotational or reflectional symmetries.

3.3 More General Symmetry

There are also methods designed for detecting more general symmetries but we only describe one such method in more detail because of its non-negligible significance in the overall field of symmetry detection. A few other methods will be only briefly mentioned, however, mentions and brief descriptions of

some additional, more general, symmetry detection methods can be found in the author’s technical report [44].

Probably the most commonly known and arguably most popular method for symmetry detection was designed by **Mitra et al.** [83] who used clustering in the transformation space for detecting quite general symmetries. The symmetry transformations can contain rotation, with or without reflection, translation and even uniform scaling. The method can also be restricted to only detect subgroups of these general symmetries. The input object is sampled and in each sampled point, the principal curvatures and principal directions are computed creating a local coordinate frame together with the normal vector. Candidate transformations are created by pairing the sampled points and aligning their local frames, but, since points where both principal curvatures are the same do not define unique transformations, such points are excluded beforehand. The scale component of each transformation is estimated from the ratio of the principal curvatures in the two points.

If the two generating points of the candidate transformation have too distinct signatures (represented by the curvature values), the transformation is discarded to ensure that only candidates that approximately match the local patches of the two points are kept. Since the candidate transformations map local patches onto each other, each of the transformations provides evidence of symmetry. The core idea of the method is that now the transformations that represent significant symmetries will appear most frequently forming modes in the candidate set, so the symmetry detection problem is approached by mode-seeking in the transformation space.

Of course, the transformations in the modes are generally not exactly the same, only similar, therefore, the mode-seeking is solved by clustering. Specifically the mean shift clustering algorithm [25]. The detected clusters then represent the symmetries of the input shape.

Fig. 3.11 shows a 3D scene at the top and the generated candidate transformations at the bottom - 280 transformations with 100 samples selected (left), 1262 transforms with 500 samples selected (right). The points in the transformation space were projected into 2D using multidimensional scaling. The two most significant clusters represent the two reflectional symmetries shown in Fig. 3.12.

Much later, **Shi et al.** [107] improved this method by designing and employing a more appropriate metric in the transformation space during the clustering stage. They also use a slightly different criterion for deciding what candidate transformations to discard. Apart from this, the method





Figure 3.11: A 3D scene in the top and the candidate transformations generated for it at the bottom projected into 2D with multidimensional scaling - the 100 black points generated 280 transformations (bottom left), the 500 yellow points generated 1262 transformations (bottom right) (figures taken from [83]).




Figure 3.12: The two symmetries corresponding to the most significant clusters from Fig.3.11 (figures taken from [83]).

[107] is almost identical to [83].

Although these clustering-based methods are very general, they are not very robust to noise and our experiments also suggest that they are rather sensitive to parameter setting.

Bokeloh et al. [11] proposed a method for finding rigid symmetries in point sets with normal vectors available for each point. The method uses feature line segments (found using normals and curvature estimates) connected to form a feature graph and it requires significant features to be present in the input data in order to find any symmetries.

Mavridis et al. [78] find global rigid symmetry of an object with missing parts by registering the object onto itself using the Super4PCS [80] rigid registration method (see also Section 3.5) with the additional modification that at least k percent of the points of the two objects should not be overlapping. This method obviously stands on the quality and properties of the selected registration method.

Li et al. [70] use co-occurrence analysis and clustering in the feature space to find rigid symmetries of a point cloud. The authors state that this method does not work well for objects with noise, irregular sampling or objects without significant features.

3.4 Symmetry Detection Overview

Here we provide an overview of the methods for symmetry detection described in the previous text. Tab. 3.1 contains most of the symmetry detection methods mentioned in the previous three sections, excluding the methods based on machine learning. For each method we state its authors together with the year it was published and citation of the publication. The first column after the method contains the symmetry type in terms of the detected transformation for each method. The second column states whether the given method finds global or local symmetry (or both).

The third column indicates the representation of input data which the method requires. Point set (point cloud) is considered the most generic representation and methods with this representation on the input are expected to work with any other representation from which a point set can be easily extracted using some sampling technique, such as a triangle mesh. By mesh representation we generally mean a triangle mesh. Some of the methods could probably also work on different types of polygonal meshes, which, however, are not very commonly used so we do not differentiate between different mesh types. A binary volume is a representation where for any

Table 3.1: Overview of the symmetry detection methods.

| Method | Transform type | Local / Global | Input representation | Weak symmetries |
|--------------------------------|------------------------------------|------------------|---|-----------------|
| Sun and Sherrah 1997 [115] | Reflection Rotation | Global | Anything with normals | No |
| Kakarala et al. 2013 [58] | Reflection | Global | Point set | No |
| Korman et al. 2015 [60] | Reflection Rotation | Global | Binary volume | No |
| Li et al. 2016 [69] | Reflection | Global | Mesh | No |
| Schiebener et al. 2016 [105] | Reflection | Global | Point set with surrounding + scanner position | Yes |
| Simari et al. 2006 [109] | Reflection | Global Local | Mesh | Yes |
| Podolak et al. 2006 [96] | Reflection | Global Local | Voxel grid (Point set, mesh) | Yes |
| Caillière et al. 2008 [16] | Reflection | Global Local? | Mesh | Yes? |
| Speciale et al. 2016 [114] | Reflection | Global Local | Depth information in voxel grid (Mesh) | Yes |
| Combès et al. 2008 [26] | Reflection | Global | Point set | Yes |
| Ecins et al. 2017 [29] | Reflection | Global | Mesh | Yes? |
| Sipiran et al. 2014 [112] | Reflection | Global | Manifold mesh | Yes |
| Cicconet et al. 2017 [22] | Reflection | Global | Point set | Yes? |
| Nagar and Raman 2019 [88] | Reflection | Global | Point set | Yes |
| Nagar and Raman 2020 [89] | Reflection | Global | Point set | Yes? |
| Martinet et al. 2006 [77] | Reflection Rotation | Global | Surface | No |
| Gothandaraman et al. 2020 [39] | Reflection Rotation | Global | Point set | No |
| Sipiran 2017 [110] | Rotation | Global | Manifold mesh | Yes |
| Sipiran 2018 [111] | Rotation | Global | Mesh | Yes |
| Ecins et al. 2018 [31] | Reflection Rotation | Global Local | Point set with normals | Yes |
| Mitra et al. 2006 [83] | Rigid + reflection + uniform scale | Global Local | Mesh | Yes |
| Shi et al. 2016 [107] | Rigid + reflection + uniform scale | Global Local | Mesh | Yes |
| Bokeloh et al. 2009 [11] | Rigid | Local | Point set with normals | Yes |
| Mavridis et al. 2015 [78] | Rigid | Global | Surface | Yes |
| Li et al. 2015 [70] | Rigid | Local | Point set | Yes |

point in the space it can be determined whether or not it lies inside the input object. By surface representation we mean a general representation of a surface which can also include a mesh. The text in brackets indicates additional representations the given method can be used for but only after some preprocessing that transforms it into the method’s native representation.

The fourth column states whether the given method can be used, to some extent, to find weak symmetries or possibly symmetries on objects with missing parts. The methods for reflectional or rotational symmetry detection which restrict the detected plane/axis so that it passes through some reference point are automatically assigned *No* in this column.

We note that all the information in the table was only extracted from the corresponding publications and we did not experimentally verify it in any way. If there is a question mark next to a line it means that we were not sure/convinced about this information based on what we could extract from the given paper.

3.5 Rigid Surface Registration

In this section we very briefly mention some methods that attempt to solve the problem of rigid surface registration where the goal is to map one surface onto another using a rigid transformation. The text in this section was mostly taken from our paper [46].

The registration process can be classified as either local registration or global registration, but in this case the terms local and global have a slightly different meaning than in context of symmetry. While a local registration only improves some initial alignment and makes it more precise, a true global registration finds an alignment without any assumptions on the initial position of the inputs. Typical registration pipeline consists of first using a global registration method to find a rough alignment and then employing some local registration method to improve its accuracy.

Perhaps the best known approach to local rigid surface registration is the Iterative Closest Point (ICP) algorithm [10, 20], and its derivatives [12, 97, 103] and generalizations [82]. The algorithm improves a given alignment by searching for point correspondences, usually based on local proximity, color similarity or other point properties. Based on the set of correspondences, an optimal rigid transformation is found that maps the points from the first object to their corresponding counterparts in the second object. These steps are iterated until convergence. Given a good initialization, the algorithm converges quite quickly, making it a good choice for local registration.

Algorithms for global registration are usually designed to be fully independent of the initial position of the inputs. Some algorithms decouple the rotation from translation, searching for the optimal rotation first. One possibility is using the Hough transform [7, 17]. Spherical maps of the input surfaces are created, where each point represents a direction, in which a certain property of the input data is computed. A rotation is then sought which maps the two spherical maps one onto another. Having the rotation, translation is found by projecting the rotated inputs onto coordinate axes and aligning the projections in each axis separately.

Another possibility of decoupling the rotation is to use the Extended Gaussian Image [42], approximated by a spherical histogram of normal orientations. Two such histograms can again be aligned first in the space of rotations [76], while the translation is found afterwards.

Another group of methods attempts to extend the notion of Phase Correlation [27], which has been successfully used for image alignment, to 3D point clouds. Usually, the input objects are resampled onto a regular grid, which is then transformed into frequency domain, where the aligning rotation can be found [15]. Unfortunately, the extension only works in a certain range of angles, and thus it cannot be applied for a general global registration.

Concepts of evolutionary algorithms were also applied to searching for the best aligning transformation [14, 21]. An initial population of possible alignments is created, and a complex evaluation function is used to determine the fitness of each phenotype. The most successful phenotypes are recombined and mutated to create the next generation of alignments. The method usually delivers a good match, however, its computational complexity is high.

Another global registration method has been proposed based on searching for pairs of congruent planar quadrilaterals in both input datasets [2], known as 4PCS (4-point congruent set). The method builds on a smart tree data structure which allows finding the congruent quadrilaterals based on their midpoint, called invariant, efficiently. An improved version of the method, the Super4PCS [80], eliminates some of the quadratic steps, such as finding all pairs of vertices at a given distance, by using another spatial structure built in a preprocessing step.

The Fast Global Registration [129] stands on the border between local and global approaches, despite its name. It uses Fast Point Feature Histogram (FPFH, [104]) features that are matched and a subset of valid matching pairs is iteratively refined, while each iteration also improves the aligning transformation.

Naturally, some of the concepts and approaches used in the methods

for symmetry detection that were described above could also be utilized in surface registration. On the other hand, some of the registration methods mentioned in this section could quite likely be modified to perform some form of symmetry detection as done e.g. in [78] or [22].

3.6 Summary

In this chapter we provided an overview of previous work done in the field of symmetry detection and briefly in rigid surface registration. The number of existing methods is certainly far from small and considering the age of the least recent ones the area has been developing for quite a time. Nevertheless, despite many of the methods being sophisticated and capable of providing good results, there is obviously still room for improvement.

For example, the possibilities for detecting rotational symmetries in 3D objects seem to be very limited - not just that the number of existing methods for this symmetry type is rather small but they also mostly have some major constraints such as only finding circular symmetries or only symmetries with the axis passing through some reference point. The number of methods for detecting reflectional symmetries is seemingly much larger but even here most of the methods are not robust enough to be used on weakly symmetrical objects or put some constraints on the input data such as requiring a mesh or other specific representation or at least estimated normals in each point. For the very few methods that appear to be capable of finding even weak reflectional symmetries in general point sets there are further questions about computational efficiency, robustness to noise, implementation simplicity or flexibility (possibility to easily modify the method for out-of-order situations, different types of data, different symmetry types etc.) which are often not even addressed in the corresponding publications.

One thing that can be noticed, however, is that most of the methods use their own unique approaches and are based on unique ideas. This means that, apart from several exceptions, each method usually advances the field in different direction. While this might be good from the scientific point of view because many novel methods are produced, it also suggest that the approaches these methods are based on are not very flexible because they are rarely reused. For example, when a method for some new specific type of data is needed (e.g. objects with missing parts) a new method is usually developed rather than extending/modifying an existing one. Therefore, in this work we will attempt to not just design new methods for symmetry detection that are better than the previous ones in terms of robustness, speed

and requirements on input data but also to base the methods on approaches that are flexible and generic so that they could be further extended, modified or generalized and reused in the future for other related problems.

4 Consensus Evaluation in RANSAC Surface Registration

In this chapter we describe our contribution to the field of rigid surface registration. The task of rigid surface registration is, having two surfaces on the input, find a rigid transformation (i.e., a combination of rotation and translation) that best aligns the first surface with the second one. The two surfaces might have generally different sizes and shapes but some portion of overlap between them (parts of the surfaces that are similar) is expected. The resulting rigid transformation should map the approximately matching parts of the two surfaces onto each other. For a specific example see Figure 2.5 in Section 2.3.

As already mentioned in Section 2.3, rigid surface registration is a field strongly related to symmetry detection since it can be understood as detecting rigid symmetry between two objects. The contribution and results presented in this chapter can therefore very easily find its application in the area of symmetry detection as well. The content of this chapter was previously published in [46].

Several registration algorithms mentioned in Section 3.5 (among others [2, 17, 76, 80]) can be interpreted as a particular implementation of the general RANSAC (Random Sample Consensus) approach, which is very commonly used in surface registration and other areas including symmetry detection. The overall scheme of a general RANSAC registration algorithm is following. By estimating the correspondences of points of the first surface with points of the second surface, a set of candidate transformations is constructed, each of which succeeds in aligning at least part of the input surfaces. Then, the consensus of the candidate transformations is evaluated to pick the best one. Some algorithms, including the current state of the art Super4PCS [80], seek consensus of points of the two input surfaces w.r.t. the candidate transformations. The candidate transformations are being applied on the first input surface and the consensus is determined for each candidate as the portion of the transformed surface that overlaps with the second surface. Generally, the candidate with the largest overlap is selected as the best one. However, we argue that this might not be the best possible approach.

The key problem of finding the consensus this way stems from the difficulty of evaluating the result. Usually, the user is expected to define a certain distance δ , and the quality of the alignment is defined as the size of the overlap, i.e. the portion of the input meshes that get mapped closer to each other than δ . Choosing a large δ leads to false matches, which align incorrect parts of the two objects. On the other hand, choosing a small δ leads to a narrow global optimum, which can be missed even with a rather dense sampling of the solution space, and a large number of local optima which complicate iterative refinement. For practical registration, δ cannot be determined universally, not even as a proportion of the input scale, as we demonstrate in our experiments. Moreover, since the solution involves an estimate of a rotation, it is possible that remote parts of the surface become severely misaligned, even when the rotation estimate is only off by a few degrees.

As a remedy to this problem, we propose *using information from all the candidates* to better analyze the space of rigid transformations. The key observation is that the optimal alignment *forms a density peak in the space of rigid transformations*. The problem can therefore be solved using some clustering algorithm (k-means, mean shift), interpreted as a facility location problem, or solved by efficiently estimating the *density* at each candidate location. In our experiments, the last choice is both the fastest and the most reliable. Similar idea was previously used by Mitra et al. for symmetry detection [83] and here we show how a related approach can be employed in surface registration.

In order to perform a clustering or a density estimation, a crucial ingredient is a *metric* that relates the elements of the given space, i.e. determines the amount of (dis)similarity of rigid transformations. This turns out to be a non-trivial problem, because the metric has to relate the different degrees of freedom of a rigid transformation (rotation, translation) to obtain a single distance value. This relation depends on the units used for translation and rotation, as well as on the scale of the object: for large objects, a unit rotation (for example one degree) has a bigger impact than a unit translation, and vice versa. For example, imagine a sphere centered at the origin with radius 10000. By rotating this sphere by 1° around the y axis all points on its surface that also lie in the xz plane travel along the arc of length $\frac{2\pi \cdot 10000}{360} \approx 174.53$. However, translating the sphere by a unit vector only causes every point on its surface to travel the distance of 1 which is hardly comparable.

Furthermore, unlike translation, rotation impacts differently different points of the object, i.e. points that are farthest from the rotation axis

travel the largest distances whereas points lying on the intersections of the axis with the object do not move at all.

There are several well-defined metrics for rotations [54], and the distance of two translations can be simply defined as the Euclidean distance of the two translation vectors. However, the problem of computing distances of combined transformations seems insufficiently addressed. In the related context of symmetry detection Mitra et al. [83] used the mean shift clustering algorithm with a metric which uses a weighting scheme that imperatively relates a rotation by π radians to a translation by one half of the body diagonal. In [113] the authors attempt to address the problem of measuring distances of 2D transformations by searching for geodesics on the manifold of rigid transformations and similar idea is used by [107] for transformations in 3D, however, the problem of proper scale is not addressed - in [107] the authors introduce two parameters that need to be fixed by the user to set the importance ratio between unit translation and unit rotation. A metric in $SE(3)$ is used by [8] to generate smooth transformations. In this case the problem of relating rotations and translations is bypassed by inheriting a metric from a higher-dimensional manifold $GL(3)$ and projecting the result back to $SE(3)$ which is not an optimal solution.

A common approach to constructing a metric for rigid transformations is to compute it as a weighed sum of some rotation metric and the Euclidean translation metric [65]. In this chapter, we analyze this approach and show that, although it is well applicable, it needs to be used with caution, because of several problems inherent to such metrics that may not be obvious. We also show how these problems can be mitigated or even eliminated using a different, data-induced transformation metric which embraces a different interpretation of the problem, and we demonstrate that in order to properly relate the rotation and translation component, the difference between two rigid transformations *necessarily must depend on the input data, and should be evaluated as such*. A metric in $SE(3)$ that fulfills this criterion has been defined in [98] to evaluate registration results, however, it has not been used in the registration (nor symmetry detection) process itself.

As a model case, we study a generic global RANSAC-based rigid surface registration scheme. We compare several metrics in the context of this algorithm, as well as different consensus evaluation strategies, and we report the results on multiple non-trivial registration datasets.

Our contribution is threefold:

- we analyze a global registration algorithm that is based on creating a set of candidate alignments, followed by finding a density peak in the

space of rigid transformations using an efficient data structure - the Vantage Point Tree,

- in the context of the algorithm, we test a variety of distance metrics, examining their properties theoretically and experimentally,
- we propose certain modifications to the compound metrics that improve their performance in the task at hand and potentially in other applications as well.

4.1 Model Registration Algorithm Description

As input, we have two triangle meshes P and Q with vertices $\mathbf{p}_i, i = 1..\nu_P$ and $\mathbf{q}_i, i = 1..\nu_Q$, and triangles $t_i^P, i = 1..\tau_P$ and $t_i^Q, i = 1..\tau_Q$. The result is a rigid transformation (an orientation-preserving isometry) that aligns Q to P . Since neither P nor Q represent the entire object, the overlap is only partial.

The model algorithm consists of the following steps:

1. A subset of the vertices of P is selected, so that their distribution is roughly uniform. For each sample, the principal curvatures κ_1, κ_2 , as well as the first principal direction \mathbf{e}_1 , are estimated and stored.
2. The mesh Q is also sampled in a roughly uniform manner and for each sample the best matching sample in P is identified based on the similarity of principal curvature estimations. A rigid transformation is constructed, which maps the corresponding points, and its quality is evaluated. High scoring transformations are stored for further processing.
3. The set of candidate transformations is interpreted as a point set in the $SE(3)$ space, and a density peak that can be interpreted as a candidate consensus is sought.

In the following, we describe each step in more detail.

4.1.1 Curvature Sampling

First, we need a feature vector that describes the local shape of the objects. It is possible to use various descriptors proposed in the literature [41]. For the purposes of our model scenario, we need a descriptor that is fast to




Figure 4.1: Neighborhood for curvature computation. Black lines represent original edges, while blue lines show their projection to the tangential plane, which is parameterized by $(\mathbf{l}_x, \mathbf{l}_y)$, represented by red arrows. Black arrows are 3D normals, while blue arrows are the normals projected to the tangential plane, i.e. $(\mathbf{n}_j \cdot \mathbf{l}_x, \mathbf{n}_j \cdot \mathbf{l}_y)$.

evaluate, yet considers larger local neighborhood than just a 1-ring (vertices in distance of 1 edge from the point in which the descriptor is being evaluated) in order to be resilient to noise. We choose to use a very simple custom local feature vector, only consisting of the two principal curvatures (see [119] for details about curvature) estimated from points in a local neighborhood, which is found using a breadth-first search (BFS) of the mesh connectivity. We stop the BFS at a certain Euclidean distance from the sample point.

Having a set of N neighbors $\mathbf{v}_j, j = 1..N$ of a vertex \mathbf{p}_i , we estimate the shape operator in a way similar to the method of Rusinkiewicz [102]. In contrast to his method, we do not attempt to obtain exact curvature, and thus we can omit some steps in order to simplify the computation. We parameterize the tangent plane, defined by the normal \mathbf{n}_i of the vertex \mathbf{p}_i , with a local orthonormal basis $(\mathbf{l}_x, \mathbf{l}_y)$ (see Fig. 4.1). We then use the property of the shape operator that for a tangential position difference vector it yields the corresponding change in normal, i.e. for each neighbor \mathbf{v}_j the following relationship represents two linear equations:

$$\begin{bmatrix} a & b \\ b & c \end{bmatrix} \begin{bmatrix} (\mathbf{v}_j - \mathbf{p}_i) \cdot \mathbf{l}_x \\ (\mathbf{v}_j - \mathbf{p}_i) \cdot \mathbf{l}_y \end{bmatrix} = \begin{bmatrix} (\mathbf{n}_j - \mathbf{n}_i) \cdot \mathbf{l}_x \\ (\mathbf{n}_j - \mathbf{n}_i) \cdot \mathbf{l}_y \end{bmatrix}.$$

Since $\mathbf{n}_i \cdot \mathbf{l}_x = 0$ and $\mathbf{n}_i \cdot \mathbf{l}_y = 0$ we get

$$\begin{bmatrix} a & b \\ b & c \end{bmatrix} \begin{bmatrix} (\mathbf{v}_j - \mathbf{p}_i) \cdot \mathbf{l}_x \\ (\mathbf{v}_j - \mathbf{p}_i) \cdot \mathbf{l}_y \end{bmatrix} = \begin{bmatrix} \mathbf{n}_j \cdot \mathbf{l}_x \\ \mathbf{n}_j \cdot \mathbf{l}_y \end{bmatrix}. \quad (4.1)$$

We obtain an overdetermined system of linear equations. Using the least squares we find the values a , b and c , which form an estimate of the shape

operator. Its eigenvalues are used as estimates of the principal curvatures and stored. We also compute the eigenvector corresponding to the larger eigenvalue and project it back to the tangent plane, yielding one principal direction \mathbf{e}_1 , which is also stored with the sample.

4.1.2 Candidate Rigid Transformations

We sample the vertices from P randomly, using a regular 3D grid over the distribution of points and ensuring that from each grid cell at most one vertex is selected. This way the sampling is usually close to uniform, even when the distribution of points in P is not.

Next, we build a set of candidate transformations. We construct a 2D k D-tree on the samples from P , using the curvature estimates as coordinates. Next, we sample the vertices of Q using the same strategy as with P , i.e. ensuring a roughly uniform sampling. For each sample point in Q , the principal curvatures are estimated using Eq. 4.1. Having the curvature estimates of a sample from Q , we use the k D-tree to locate the most similar sample from P .

Finally, having a potentially corresponding pair of vertices, we build a pair of rigid transformations T^+ and T^- that map the sample from Q to the sample from P , together with the normal, T^+ maps the principal direction \mathbf{e}_Q of the sample in Q to the principal direction \mathbf{e}_P of the sample in P , while T^- maps \mathbf{e}_Q to $-\mathbf{e}_P$.

For each generated transformation, we quickly verify that it aligns a substantial part of Q . In preprocessing, we build a uniform 3D binary grid over the bounding box of P , in which each cell value determines whether there is a point in P closer to the cell center than δ . We use point sample of Q reduced to about 400 vertices. Each sample vertex \mathbf{v}_q is transformed using the candidate transformation T , and the binary grid is used to determine whether or not there is a point \mathbf{v}_p in P such that $\|T(\mathbf{v}_q) - \mathbf{v}_p\| < \delta$. The percentage of points that meets this criterion is also known as Largest Common Pointset (LCP). If LCP is at least 3%, then the candidate is retained for further processing. We keep sampling Q until a certain number of candidate transformations is found, which pass the verification.

One possibility to choose the aligning transformation is to pick the candidate with the highest resulting overlap, i.e. the highest portion of samples fulfilling the verification condition. As discussed at the beginning of this chapter, this approach is sensitive to the choice of the δ parameter, and, therefore, we also use a different approach, exploiting the *distribution of the candidate transformations*.

4.1.3 Analyzing the Space of Rigid Transformations

In the set of candidate transformations, there are probably many transformations quite close to the optimum, but possibly not close enough to obtain a high score. Even in such a situation, however, it is possible to analyze the space of rigid transformations and find a good approximation of the optimum by looking for a density peak in the distribution of the candidate transformations.

The key ingredient needed for evaluating the density is a proper way of evaluating proximity. This problem is challenging, because of the non-trivial topology of the space, which makes its parameterization difficult. Common parameterizations, such as using 4×4 matrices in homogeneous coordinates or dual quaternions, are redundant in the number of parameters used to describe the transformation. The redundancy can be avoided for example by using Euler angles to describe the rotation, however, even then, similarity of two rigid transformations does not imply the similarity of the parameters that describe them. Choosing a proper metric in $SE(3)$ is a crucial task, and, therefore, we analyze different possibilities in the next section. At this point, we assume that the distances can be consistently measured and finish the overview of our model registration algorithm.

There are different ways of finding a density peak of vertices (rigid transformations in our case) in a vector space. A clustering algorithm can be used that groups similar vectors, and cluster centroids or modes often correspond to density peaks. There are various approaches to clustering, ranging from simple parametric algorithms, such as k-means, up to more complex procedures, such as mean shift clustering [25] or facility location [19]. In each case, the questions remain how to select the cluster which corresponds to the best alignment, how to select the best alignment within the cluster, and whether the clusters correspond to good alignments in the first place. Although some intuitive answers can be found for some of the clustering algorithms, we use a more direct way of finding the density peak, which turns out to be both more efficient computationally, as well as more robust/precise.

A sampling density estimation function can be defined using some kernel function K as

$$\rho(x) = \sum_i K(d(T_i, x)), \quad (4.2)$$

where $d(x, y)$ measures the distance of two samples. Various kernel functions can be used for K , for efficiency reasons we use a simple Gaussian $K(r) = e^{-(Dr)^2}$, where D is a spread parameter.

Instead of looking for the general location of the true global maximum of ρ , we only search for a candidate T^* , for which $\rho(T^*) \geq \rho(T_i)$ for all

candidate transformations T_i . Given the spread parameter D and some small threshold t , only samples within a radius $r = \sqrt{-\ln(t)}/D$ contribute significantly ($> t$) to the density.

Therefore, the task is to find for each candidate a set of candidates up to a distance r . Such goal is usually achieved using a k D-tree, however, such structure cannot be used in this setting, since the space we are working in is not Euclidean. Having a metric, however, allows using a more general acceleration structure - *the Vantage Point Tree* [126]. It is constructed as a binary tree, where each node represents a certain candidate c . Inner nodes have two child nodes: the *near* node contains all other candidates that are closer to c than a certain threshold, while the *far* node contains candidates that are farther from c than the threshold. The threshold is selected as the median of distances, so that the tree is balanced, and at a certain level the branching is terminated by a leaf node, which contains no more than a certain number of candidates.

Building such a structure is of $\mathcal{O}(n \log(n))$ complexity. With the structure built, proximity queries can be resolved quickly with $\mathcal{O}(\log(n))$ complexity in the average case by only investigating the branches which can contribute to the answer. Therefore, evaluating the density estimation function at each candidate location has an overall complexity of $\mathcal{O}(n \log(n))$ and it is easily feasible even with the 10^4 candidates used as default setting.

In principle, our approach is similar to what [83] proposed for symmetry detection, however, we use a *substantially different execution* with different metrics in a slightly different application. Our density peak location approach could be also applied in an arbitrary Mode-based method (see Chapter 5) for global symmetry detection. As an alternative, we have tested the k-means algorithm and the local search algorithm which solves the problem in the facility location interpretation. Both algorithms require a higher minimum number of candidates in order to robustly identify the correct alignment than the proposed exhaustive density evaluation. Moreover, k-means is more than 5 times slower and the local search for facility location is more than 48 times slower. Therefore, we argue that using an advanced clustering algorithm is not an efficient solution for finding the density peak representing the optimal alignment.

4.2 Transformation Distance Metrics

In order to find the density peak in the space of rigid transformations we need to define a metric $d(T_1, T_2)$ which measures the distance of transform-

ation T_1 from transformation T_2 . The metric must somehow relate the two parts of the transformation: the rotation and the translation. In this section we therefore discuss other transformation distance metrics that can be used for this purpose. Apart from surface registration, the presented knowledge about rigid transformations can be useful in different areas including symmetry detection, especially in the case of using the Mode-based methods (see Section 5) such as [83] to detect rigid symmetries.

4.2.1 Composed Metrics

Since a rigid transformation can be decomposed into a rotation and a translation, a metric can be defined as $d(T_1, T_2) = c_R d_R(\mathbf{R}_1, \mathbf{R}_2) + c_t \|\mathbf{t}_1 - \mathbf{t}_2\|$ where $d_R(\mathbf{R}_1, \mathbf{R}_2)$ is a metric for rotations, \mathbf{R}_1 and \mathbf{R}_2 are the rotations of T_1 and T_2 respectively, \mathbf{t}_1 and \mathbf{t}_2 are translation vectors of T_1 and T_2 respectively and c_R and c_t are customizable coefficients (a constant spread parameter D of the kernel will be used). This way of defining the metric was used e.g. in [65] or [4], a very similar approach was also used by Mitra et al. in [83] with the only difference that the terms are squared and there is an additional term for scaling.

An extensive analysis of several distance metrics for rotations where basically any of these metrics could be used as $d_R(\mathbf{R}_1, \mathbf{R}_2)$ can be found in [54]. Composed transformation distance metrics, as described above, exhibit certain impractical properties, some of which are independent of the choice of $d_R(\mathbf{R}_1, \mathbf{R}_2)$, as discussed next.

Order of Operations

A rigid transformation can be equivalently defined as either a rotation followed by a translation, i.e. $T(\mathbf{x}) = \mathbf{R}\mathbf{x} + \mathbf{t}$, or as a translation followed by a rotation, i.e. $T'(\mathbf{x}) = \mathbf{R}'(\mathbf{x} + \mathbf{t}')$. If T and T' describe the same transformation, following must hold:

$$\mathbf{R}\mathbf{x} + \mathbf{t} = \mathbf{R}'(\mathbf{x} + \mathbf{t}'). \quad (4.3)$$

This equation holds when $\mathbf{R} = \mathbf{R}'$ and $\mathbf{t}' = \mathbf{R}^T \mathbf{t}$, because then $\mathbf{R}'(\mathbf{x} + \mathbf{t}') = \mathbf{R}(\mathbf{x} + \mathbf{R}^T \mathbf{t}) = \mathbf{R}\mathbf{x} + \mathbf{t}$. Applying transformation T to an object corresponds to first changing its orientation to the target orientation by rotation and then translating the object to the target location. Applying T' corresponds to first translating the object to such a position that the subsequent rotation, which enforces the target orientation, also moves the object to the target location. The change of orientation is the same in both cases. Therefore, Equation

4.3 can be rewritten as follows

$$\begin{aligned}
\mathbf{R}\mathbf{x} + \mathbf{t} &= \mathbf{R}(\mathbf{x} + \mathbf{t}'), \\
\mathbf{R}\mathbf{x} + \mathbf{t} &= \mathbf{R}\mathbf{x} + \mathbf{R}\mathbf{t}', \\
\mathbf{t} = \mathbf{R}\mathbf{t}' &\Rightarrow \mathbf{R}^T\mathbf{t} = \mathbf{t}'.
\end{aligned} \tag{4.4}$$

Suppose there are two arbitrary rigid transformations T_1 and T_2 expressed in the rotation-first form, i.e. $\mathbf{R}_1\mathbf{x} + \mathbf{t}_1$, $\mathbf{R}_2\mathbf{x} + \mathbf{t}_2$, and translation-first form, i.e. $\mathbf{R}'_1(\mathbf{x} + \mathbf{t}'_1)$, $\mathbf{R}'_2(\mathbf{x} + \mathbf{t}'_2)$. Since $\mathbf{R}_1 = \mathbf{R}'_1$ and $\mathbf{R}_2 = \mathbf{R}'_2$, $d_R(\mathbf{R}_1, \mathbf{R}_2) = d_R(\mathbf{R}'_1, \mathbf{R}'_2)$ for any rotation metric d_R . The translation metrics $\|\mathbf{t}_1 - \mathbf{t}_2\|$ and $\|\mathbf{t}'_1 - \mathbf{t}'_2\|$ are, however, different. The first metric can be expanded as

$$\|\mathbf{t}_1 - \mathbf{t}_2\| = \sqrt{(\mathbf{t}_1 - \mathbf{t}_2)^T(\mathbf{t}_1 - \mathbf{t}_2)} = \sqrt{\mathbf{t}_1^T\mathbf{t}_1 - 2\mathbf{t}_1^T\mathbf{t}_2 + \mathbf{t}_2^T\mathbf{t}_2}. \tag{4.5}$$

From Equation 4.4 it follows that $\|\mathbf{t}'_1 - \mathbf{t}'_2\| = \|\mathbf{R}_1^T\mathbf{t}_1 - \mathbf{R}_2^T\mathbf{t}_2\|$ and the second metric can, therefore, be expanded as follows

$$\begin{aligned}
\|\mathbf{t}'_1 - \mathbf{t}'_2\| &= \|\mathbf{R}_1^T\mathbf{t}_1 - \mathbf{R}_2^T\mathbf{t}_2\| = \sqrt{(\mathbf{R}_1^T\mathbf{t}_1 - \mathbf{R}_2^T\mathbf{t}_2)^T(\mathbf{R}_1^T\mathbf{t}_1 - \mathbf{R}_2^T\mathbf{t}_2)} = \\
&= \sqrt{\mathbf{t}_1^T\mathbf{R}_1\mathbf{R}_1^T\mathbf{t}_1 - 2\mathbf{t}_1^T\mathbf{R}_1\mathbf{R}_2^T\mathbf{t}_2 + \mathbf{t}_2^T\mathbf{R}_2\mathbf{R}_2^T\mathbf{t}_2} = \\
&= \sqrt{\mathbf{t}_1^T\mathbf{t}_1 - 2\mathbf{t}_1^T\mathbf{R}_1\mathbf{R}_2^T\mathbf{t}_2 + \mathbf{t}_2^T\mathbf{t}_2}.
\end{aligned}$$

The only case when $\|\mathbf{t}_1 - \mathbf{t}_2\| = \|\mathbf{t}'_1 - \mathbf{t}'_2\|$ is when $\mathbf{R}_1 = \mathbf{R}_2$. The result of the translation metric and, therefore, of the composed transformation distance metric, depends on the order of rotation and translation, which we choose for describing the transformations. Furthermore, the difference between $\|\mathbf{t}_1 - \mathbf{t}_2\|$ and $\|\mathbf{t}'_1 - \mathbf{t}'_2\|$ grows with the increasing difference between \mathbf{R}_1 and \mathbf{R}_2 .

In principle there is no reason to prefer either representation and its choice seems rather arbitrary. However, for objects located near the origin, the rotation-first form may seem more appropriate, since regardless of the rotation, the direction of the translation vector *approximately* corresponds to the difference between the output and the starting position of the object. Also, the rotation-first form corresponds to the way rigid transformations are evaluated when using matrix multiplication in homogeneous coordinates for transformation evaluation. For these reasons, in the rest of this text, we only consider rigid transformations in the rotation-first form. Nevertheless, the statements made in the rest of this section can be similarly made about the translation-first form using analogous analysis and derivations.

Dependence on Position

Consider an input object Q_{in} and an arbitrary rigid transformation T that transforms Q_{in} into an output object $Q_{out} = T(Q_{in})$. Translating both Q_{in} and Q_{out} by the same arbitrary vector \mathbf{t}_0 results in two objects Q'_{in} and Q'_{out} respectively, with different absolute position but the same mutual position as Q_{in} and Q_{out} . To transform Q'_{in} into Q'_{out} a new transform T' must be defined such that $T'(Q'_{in}) = Q'_{out}$. Since translating Q'_{out} by $-\mathbf{t}_0$ results in Q_{out} , $T(\mathbf{x}) = T'(\mathbf{x} + \mathbf{t}_0) - \mathbf{t}_0$ must hold, which can be expanded as

$$\mathbf{R}\mathbf{x} + \mathbf{t} = \mathbf{R}'(\mathbf{x} + \mathbf{t}_0) + \mathbf{t}' - \mathbf{t}_0.$$

The change of orientation is equal for T and T' , and therefore $\mathbf{R} = \mathbf{R}'$, which implies

$$\begin{aligned} \mathbf{R}\mathbf{x} + \mathbf{t} &= \mathbf{R}(\mathbf{x} + \mathbf{t}_0) + \mathbf{t}' - \mathbf{t}_0, \\ \mathbf{t}' &= \mathbf{t} - \mathbf{R}\mathbf{t}_0 + \mathbf{t}_0. \end{aligned} \quad (4.6)$$

Consider two different rigid transformations T_1 and T_2 and some arbitrary vector \mathbf{t}_0 and two corresponding transformations T'_1 and T'_2 where $\mathbf{t}'_1 = \mathbf{t}_1 - \mathbf{R}_1\mathbf{t}_0 + \mathbf{t}_0$ and $\mathbf{t}'_2 = \mathbf{t}_2 - \mathbf{R}_2\mathbf{t}_0 + \mathbf{t}_0$ according to Equation 4.6. Since $\mathbf{R}_1 = \mathbf{R}'_1$ and $\mathbf{R}_2 = \mathbf{R}'_2$, it follows that $d_R(\mathbf{R}_1, \mathbf{R}_2) = d_R(\mathbf{R}'_1, \mathbf{R}'_2)$, but the translation metrics $\|\mathbf{t}_1 - \mathbf{t}_2\|$ and $\|\mathbf{t}'_1 - \mathbf{t}'_2\|$ are generally different. The first one expands as shown in Equation 4.5, while the second metric expands as

$$\begin{aligned} \|\mathbf{t}'_1 - \mathbf{t}'_2\| &= \|\mathbf{t}_1 - \mathbf{R}_1\mathbf{t}_0 + \mathbf{t}_0 - \mathbf{t}_2 + \mathbf{R}_2\mathbf{t}_0 - \mathbf{t}_0\| \\ &= \|\mathbf{t}_1 - \mathbf{R}_1\mathbf{t}_0 - \mathbf{t}_2 + \mathbf{R}_2\mathbf{t}_0\|, \end{aligned} \quad (4.7)$$

and thus $\|\mathbf{t}_1 - \mathbf{t}_2\| = \|\mathbf{t}'_1 - \mathbf{t}'_2\|$ only when $\mathbf{R}_1 = \mathbf{R}_2$ or when $\mathbf{t}_0 = \mathbf{0}$.

Suppose a different scenario, where only the input object Q_{in} is translated by \mathbf{t}_0 , creating Q'_{in} , and Q_{out} stays the same, i.e. $T(Q_{in}) = T'(Q'_{in}) = Q_{out}$. Now $T(\mathbf{x}) = T'(\mathbf{x} + \mathbf{t}_0)$ must hold, and it holds when $\mathbf{R} = \mathbf{R}'$ and $\mathbf{t}' = \mathbf{t} - \mathbf{R}\mathbf{t}_0$, the derivation is analogical to Eq. 4.6. If we now again consider two transformations T_1, T_2 and corresponding T'_1, T'_2 it is easy to see that $\|\mathbf{t}'_1 - \mathbf{t}'_2\| = \|\mathbf{t}_1 - \mathbf{R}_1\mathbf{t}_0 - \mathbf{t}_2 + \mathbf{R}_2\mathbf{t}_0\|$, which is the same expression as Eq. 4.7.

In a scenario where only Q_{out} is translated by \mathbf{t}_0 , yielding Q'_{out} , and Q_{in} stays the same ($T(Q_{in}) = Q_{out}, T'(Q_{in}) = Q'_{out}$). Now $T(\mathbf{x}) = T'(\mathbf{x}) - \mathbf{t}_0$ must hold and it is easily proven that it holds when $\mathbf{R} = \mathbf{R}'$ and $\mathbf{t}' = \mathbf{t} + \mathbf{t}_0$ because then $\mathbf{R}'\mathbf{x} + \mathbf{t}' - \mathbf{t}_0 = \mathbf{R}\mathbf{x} + \mathbf{t} + \mathbf{t}_0 - \mathbf{t}_0 = \mathbf{R}\mathbf{x} + \mathbf{t}$. For two transformations T_1, T_2 and corresponding T'_1, T'_2 we get

$$\|\mathbf{t}'_1 - \mathbf{t}'_2\| = \|\mathbf{t}_1 + \mathbf{t}_0 - \mathbf{t}_2 - \mathbf{t}_0\| = \|\mathbf{t}_1 - \mathbf{t}_2\|. \quad (4.8)$$

This has profound consequences in the context of rigid surface registration and possibly in other applications as well including symmetry detection. When the space of rigid transformations is being sampled and transformations are created by fitting points of Q onto points of P , the value of the Euclidean metric applied on the translation components of the transformations depends on the position of Q (Eq. 4.7), but does not depend on the position of P (Eq. 4.8). Having a general rigid transformation $T_1(\mathbf{x}) = \mathbf{R}_1\mathbf{x} + \mathbf{t}_1$ and $\mathbf{x} = \mathbf{0}$, then \mathbf{t}_1 exactly represents the change of position of the point caused by the transformation and if another transformation $T_2(\mathbf{x}) = \mathbf{R}_2\mathbf{x} + \mathbf{t}_2$ is defined then $\|\mathbf{t}_1 - \mathbf{t}_2\|$ exactly represents the difference between the position change of T_1 and the position change of T_2 . However, if \mathbf{x} changes by a non-zero vector \mathbf{t}_0 , then the value $\|\mathbf{t}_1 - \mathbf{t}_2\|$ starts to deviate from the difference of the position changes of T_1 and T_2 and, according to Eq. 4.7, the deviation grows infinitely with the distance of \mathbf{x} from the origin (the length of \mathbf{t}_0) and with the difference between the two rotations \mathbf{R}_1 and \mathbf{R}_2 (the angle between $\mathbf{R}_1\mathbf{t}_0$ and $\mathbf{R}_2\mathbf{t}_0$). This implies that generally the farther a point is from the origin, the worse the translation component \mathbf{t} of a rigid transformation describes the change of the point's position after applying the transformation, and the worse the Euclidean metric of the translation components describes the difference of these changes between the translation components of two arbitrary rigid transformations.

Using the composed transformation metric, which uses the Euclidean metric on the translation components, is therefore only meaningful when the transformed object is approximately centered at the origin. Otherwise the composed metric can behave unpredictably, having a great negative impact on the registration results.

Dependence on Orientation

A rotation metric d_R is right-invariant if for any two rotations $\mathbf{R}_1, \mathbf{R}_2$ and an arbitrary rotation \mathbf{R}_0 it is that $d_R(\mathbf{R}_1, \mathbf{R}_2) = d_R(\mathbf{R}_1\mathbf{R}_0, \mathbf{R}_2\mathbf{R}_0)$. It is left-invariant if $d_R(\mathbf{R}_1, \mathbf{R}_2) = d_R(\mathbf{R}_0\mathbf{R}_1, \mathbf{R}_0\mathbf{R}_2)$ and bi-invariant if $d_R(\mathbf{R}_1, \mathbf{R}_2) = d_R(\mathbf{R}_1\mathbf{R}_0, \mathbf{R}_2\mathbf{R}_0) = d_R(\mathbf{R}_0\mathbf{R}_1, \mathbf{R}_0\mathbf{R}_2)$.

Consider an input object Q_{in} and an arbitrary rigid transformation T that transforms Q_{in} into an output object $Q_{out} = T(Q_{in})$. Now imagine rotating both Q_{in} and Q_{out} using the same arbitrary rotation matrix \mathbf{R}_0 resulting in two objects Q'_{in} and Q'_{out} respectively, with different absolute orientation but the same mutual orientation as Q_{in} and Q_{out} . A transform T' transforms Q'_{in} into Q'_{out} , i.e. $T'(Q'_{in}) = Q'_{out}$. Since rotating Q'_{out} by \mathbf{R}_0^T results in Q_{out} , $T(\mathbf{x}) = \mathbf{R}_0^T T'(\mathbf{R}_0\mathbf{x})$ must hold, which can be expanded as

follows

$$\begin{aligned}\mathbf{R}\mathbf{x} + \mathbf{t} &= \mathbf{R}_0^T(\mathbf{R}'\mathbf{R}_0\mathbf{x} + \mathbf{t}'), \\ \mathbf{R}\mathbf{x} + \mathbf{t} &= \mathbf{R}_0^T\mathbf{R}'\mathbf{R}_0\mathbf{x} + \mathbf{R}_0^T\mathbf{t}'.\end{aligned}$$

This equation holds when $\mathbf{R}' = \mathbf{R}_0\mathbf{R}\mathbf{R}_0^T$ and $\mathbf{t}' = \mathbf{R}_0\mathbf{t}$.

Suppose we have two different rigid transformations T_1 and T_2 and two corresponding transformations T'_1 and T'_2 , where $\mathbf{t}'_1 = \mathbf{R}_0\mathbf{t}_1$, $\mathbf{t}'_2 = \mathbf{R}_0\mathbf{t}_2$, $\mathbf{R}'_1 = \mathbf{R}_0\mathbf{R}_1\mathbf{R}_0^T$ and $\mathbf{R}'_2 = \mathbf{R}_0\mathbf{R}_2\mathbf{R}_0^T$, and \mathbf{R}_0 is some arbitrary rotation matrix. Obviously, $\|\mathbf{t}_1 - \mathbf{t}_2\| = \|\mathbf{t}'_1 - \mathbf{t}'_2\|$ because $\|\mathbf{t}_1 - \mathbf{t}_2\| = \|\mathbf{R}_0\mathbf{t}_1 - \mathbf{R}_0\mathbf{t}_2\|$. In order to satisfy $d_R(\mathbf{R}_1, \mathbf{R}_2) = d_R(\mathbf{R}'_1, \mathbf{R}'_2)$ it must be that $d_R(\mathbf{R}_1, \mathbf{R}_2) = d_R(\mathbf{R}_0\mathbf{R}_1\mathbf{R}_0^T, \mathbf{R}_0\mathbf{R}_2\mathbf{R}_0^T)$ which holds if the rotation metric d_R is bi-invariant.

Suppose a different scenario, where only the input object Q_{in} is rotated by \mathbf{R}_0 , creating Q'_{in} , and Q_{out} stays the same, i.e. $T(Q_{in}) = T'(Q'_{in}) = Q_{out}$. Now $T(\mathbf{x}) = T'(\mathbf{R}_0\mathbf{x})$ must hold, which expands as $\mathbf{R}\mathbf{x} + \mathbf{t} = \mathbf{R}'\mathbf{R}_0\mathbf{x} + \mathbf{t}'$. This holds for $\mathbf{R}' = \mathbf{R}\mathbf{R}_0^T$ and $\mathbf{t}' = \mathbf{t}$. If we now again consider two transformations T_1 , T_2 and corresponding T'_1 , T'_2 , obviously $\|\mathbf{t}_1 - \mathbf{t}_2\| = \|\mathbf{t}'_1 - \mathbf{t}'_2\|$. For $d_R(\mathbf{R}_1, \mathbf{R}_2) = d_R(\mathbf{R}'_1, \mathbf{R}'_2)$ it must now be that $d_R(\mathbf{R}_1, \mathbf{R}_2) = d_R(\mathbf{R}_1\mathbf{R}_0^T, \mathbf{R}_2\mathbf{R}_0^T)$ which holds when the rotation metric d_R is right-invariant.

Finally, when only Q_{out} is rotated by \mathbf{R}_0 , creating Q'_{out} , and Q_{in} stays the same ($T(Q_{in}) = Q_{out}$, $T'(Q_{in}) = Q'_{out}$), $T(\mathbf{x}) = \mathbf{R}_0^T T'(\mathbf{x})$ must hold, which expands as $\mathbf{R}\mathbf{x} + \mathbf{t} = \mathbf{R}_0^T(\mathbf{R}'\mathbf{x} + \mathbf{t}')$. This implies $\mathbf{R}' = \mathbf{R}_0\mathbf{R}$ and $\mathbf{t}' = \mathbf{R}_0\mathbf{t}$ because then $\mathbf{R}_0^T(\mathbf{R}'\mathbf{x} + \mathbf{t}') = \mathbf{R}_0^T(\mathbf{R}_0\mathbf{R}\mathbf{x} + \mathbf{R}_0\mathbf{t}) = \mathbf{R}\mathbf{x} + \mathbf{t}$. For two transformations T_1 , T_2 and corresponding T'_1 , T'_2 , for the translation metric we now again get $\|\mathbf{t}_1 - \mathbf{t}_2\| = \|\mathbf{R}_0\mathbf{t}_1 - \mathbf{R}_0\mathbf{t}_2\|$, which holds. For $d_R(\mathbf{R}_1, \mathbf{R}_2) = d_R(\mathbf{R}'_1, \mathbf{R}'_2)$ we need $d_R(\mathbf{R}_1, \mathbf{R}_2) = d_R(\mathbf{R}_0\mathbf{R}_1, \mathbf{R}_0\mathbf{R}_2)$, which holds when the rotation metric d_R is left-invariant.

In context of rigid surface registration, when fitting points of object Q onto points of object P and a composed metric is used to measure distances between the created transformations, the value of the composed metric is independent of the initial orientation of Q only if the rotation metric is right-invariant and independent of the initial orientation of P only if the rotation metric is left-invariant. For a composed metric to be independent of the initial orientation of both Q and P , the rotation metric must be bi-invariant, even when the mutual orientation of Q and P remains unchanged.

Therefore we recommend only using rotation metrics that are bi-invariant, otherwise the composed metric might behave unpredictably, negatively impacting the registration results. This is because the distances of the candidate transformations could change just based on the initial orientation of the input objects, i.e. for one initial orientation two candidates might be

evaluated as distant while being evaluated as near for a different orientation. This can be compared to the following analogous situation: imagine a metric for vectors in E^2 that gives a different distance value for vectors $[1, 0]^T, [0, 1]^T$ than for vectors $[-1, 0]^T, [0, -1]^T$ only because their absolute orientation is different although their mutual pose is identical. Evaluating density with such a metric would be unpredictable and could easily lead to incorrect identification of the density maximum.

Rotation Metrics

We consider 6 rotation metrics that are described and analyzed in [54], denoted $\Phi_i, i = 1, \dots, 6$, but since $\Phi_6 = 2\Phi_3$, we exclude Φ_6 and only discuss $\Phi_i, i = 1, \dots, 5$. For two rotations, described by rotation matrices $\mathbf{R}_1, \mathbf{R}_2$, we denote $\mathbf{q}_1, \mathbf{q}_2$ the 4-dimensional vectors representing unit quaternions that correspond to the same rotations, and $(\alpha_1, \beta_1, \gamma_1), (\alpha_2, \beta_2, \gamma_2)$ the triplets of corresponding Euler angles.

Distance between Euler angles

$$\begin{aligned} d_R^{DEA}(\mathbf{R}_1, \mathbf{R}_2) &= \Phi_1((\alpha_1, \beta_1, \gamma_1), (\alpha_2, \beta_2, \gamma_2)) = \\ &= \sqrt{d(\alpha_1, \alpha_2)^2 + d(\beta_1, \beta_2)^2 + d(\gamma_1, \gamma_2)^2} \end{aligned}$$

where $d(a, b) = \min\{|a-b|, 2\pi-|a-b|\}$ and $\alpha, \gamma \in [-\pi, \pi], \beta \in [-\pi/2, \pi/2]$.

Distance between unit quaternions

$$d_R^{DQ}(\mathbf{R}_1, \mathbf{R}_2) = \Phi_2(\mathbf{q}_1, \mathbf{q}_2) = \min\{\|\mathbf{q}_1 - \mathbf{q}_2\|, \|\mathbf{q}_1 + \mathbf{q}_2\|\}$$

where \mathbf{q}_1 and \mathbf{q}_2 are treated as 4-dimensional vectors and $\|\cdot\|$ denotes the Euclidean norm.

Dot product of unit quaternions

$$d_R^{ADPQ}(\mathbf{R}_1, \mathbf{R}_2) = \Phi_3(\mathbf{q}_1, \mathbf{q}_2) = \arccos(|\mathbf{q}_1^T \mathbf{q}_2|)$$

and

$$d_R^{DPQ}(\mathbf{R}_1, \mathbf{R}_2) = \Phi_4(\mathbf{q}_1, \mathbf{q}_2) = 1 - |\mathbf{q}_1^T \mathbf{q}_2|$$

where \mathbf{q}_1 and \mathbf{q}_2 are treated as 4-dimensional vectors in both cases.

Deviation from identity matrix

$$d_R^{DIM}(\mathbf{R}_1, \mathbf{R}_2) = \Phi_5(\mathbf{R}_1, \mathbf{R}_2) = \|\mathbf{I} - \mathbf{R}_1 \mathbf{R}_2^T\|_F = \|\mathbf{R}_1 - \mathbf{R}_2\|_F$$

where $\|\cdot\|_F$ denotes the Frobenius norm and \mathbf{I} is the identity matrix.

All these metrics, except for d_R^{DEA} , are bi-invariant [54] and therefore well applicable in the context of rigid registration. For more detailed description and further analysis of these rotation metrics see [54]. The proof that d_R^{DEA} is not bi-invariant follows.

Consider general rotation matrices \mathbf{R}_1 , \mathbf{R}_2 and \mathbf{R}_0 and the corresponding Euler angles $(\alpha_1, \beta_1, \gamma_1)$, $(\alpha_2, \beta_2, \gamma_2)$ and $(\alpha_0, \beta_0, \gamma_0)$, respectively. If d_R^{DEA} is a bi-invariant rotation metric, then $d_R^{DEA}(\mathbf{R}_1, \mathbf{R}_2) = d_R^{DEA}(\mathbf{R}_1 \mathbf{R}_0, \mathbf{R}_2 \mathbf{R}_0) = d_R^{DEA}(\mathbf{R}_0 \mathbf{R}_1, \mathbf{R}_0 \mathbf{R}_2)$ must hold for any \mathbf{R}_1 , \mathbf{R}_2 , \mathbf{R}_0 . Now consider that the rotations are such that $\alpha_1 = 0$, $\beta_1 = 0$, $\gamma_1 = 90^\circ$, $\alpha_2 = 0$, $\beta_2 = -90^\circ$, $\gamma_2 = 0$, $\alpha_0 = 0$, $\beta_0 = 30^\circ$, $\gamma_0 = 60^\circ$. It can be easily shown that the Euler angles of the rotation defined by a matrix $\mathbf{R}_1 \mathbf{R}_0$ are $\alpha_{10} = -30^\circ$, $\beta_{10} = 0$, $\gamma_{10} = 150^\circ$ and those of $\mathbf{R}_2 \mathbf{R}_0$ are $\alpha_{20} = 0$, $\beta_{20} = -60^\circ$, $\gamma_{20} = 60^\circ$. Similarly, for $\mathbf{R}_0 \mathbf{R}_1$ we get $\alpha_{01} = 0$, $\beta_{01} = 30^\circ$, $\gamma_{01} = 150^\circ$ and for $\mathbf{R}_0 \mathbf{R}_2$ we get $\alpha_{02} = 73.896^\circ$, $\beta_{02} = -25.658^\circ$, $\gamma_{02} = 56.309^\circ$. It is not hard to show now that $d_R^{DEA}(\mathbf{R}_1, \mathbf{R}_2) = 127.279$, $d_R^{DEA}(\mathbf{R}_1 \mathbf{R}_0, \mathbf{R}_2 \mathbf{R}_0) = 112.249$ and $d_R^{DEA}(\mathbf{R}_0 \mathbf{R}_1, \mathbf{R}_0 \mathbf{R}_2) = 119.404$. This implies that in general $d_R^{DEA}(\mathbf{R}_1, \mathbf{R}_2) \neq d_R^{DEA}(\mathbf{R}_1 \mathbf{R}_0, \mathbf{R}_2 \mathbf{R}_0)$ and $d_R^{DEA}(\mathbf{R}_1, \mathbf{R}_2) \neq d_R^{DEA}(\mathbf{R}_0 \mathbf{R}_1, \mathbf{R}_0 \mathbf{R}_2)$ thus proving that the rotation metric d_R^{DEA} is neither left nor right-invariant, and therefore not bi-invariant.

Normalization

Since the values of all the rotation metrics described above are bounded on finite intervals, we normalize them so that their values are mapped to $[0, 1]$. We also normalize the translation metric by dividing it by the estimated size of Q which is computed as the average distance of its vertices from its centroid. The composed metric we use can therefore be expressed as $d(T_1, T_2) = c_R \frac{d_R(\mathbf{R}_1, \mathbf{R}_2)}{k_R} + c_t \frac{\|t_1 - t_2\|}{k_t}$ where k_R , k_t are the corresponding normalization constants.

Properties of Composed Metrics

In order to use a composed metric in the context of rigid registration, it is advisable to choose a bi-invariant rotation metric and only investigate transformations of a centered object. The centering can also be interpreted as integral part of the metric by evaluating each transformation composed

with a centering translation $d(T_1 \circ T_c, T_2 \circ T_c)$, where T_c moves the centroid of the input object to the origin. In this sense, *all composed metrics are data-dependent*, since the position of the centroid of the input object is necessary for their proper evaluation. All composed metrics have two free parameters, c_R and c_t . They influence the scale of the metric and, more importantly, they set the balance of the rotation and translation part.

4.2.2 Compound Metrics

Some of the problems of composed metrics can be eliminated by accepting and exploiting the fact that the metric *inevitably must* depend on the input data to which the transformations are applied, i.e. evaluating the similarity of rigid transformations by evaluating the *difference of their effect* on the input data.

Vertex Sum of Squares

A compound data-induced metric provides a single value which quantifies the difference between the effects of the two transformations to the input object. Such a metric was previously discussed in [98], however, it has only been used in order to evaluate the result of a registration, yet not within the registration process itself. For two rigid transformations T_1, T_2 it is defined as follows

$$d(T_1, T_2)^2 = \sum_{i=1}^{\nu_Q} \|T_1(\mathbf{q}_i) - T_2(\mathbf{q}_i)\|^2 \quad (4.9)$$

where \mathbf{q}_i are the vertices of the input object Q . This is in fact a special case of the Procrustes distance [61] used to measure distances between point clouds with known point-to-point correspondences. In the following, we show that for measuring distances between rigid transformations, it can be evaluated with $\mathcal{O}(1)$ time complexity rather than $\mathcal{O}(n)$.

Expressing T_1 as a rotation \mathbf{R}_1 and a translation \mathbf{t}_1 , and T_2 as a rotation \mathbf{R}_2 and a translation \mathbf{t}_2 , we can derive Eq. 4.10. For any vector \mathbf{v} and any matrix \mathbf{A} , it applies that

$$\mathbf{v}^T \mathbf{A} \mathbf{v} = \mathbf{A} : \mathbf{v} \mathbf{v}^T,$$

where $\mathbf{A} : \mathbf{B}$ denotes the Frobenius matrix product. Also

$$\mathbf{v}^T \mathbf{A} \mathbf{v} = \mathbf{A} : \mathbf{v} \mathbf{v}^T = \mathbf{A}^T : \mathbf{v} \mathbf{v}^T = \mathbf{v}^T \mathbf{A}^T \mathbf{v},$$

and thus the expression in Eq. 4.11 can be derived. Without loss of generality, we can shift Q so that $\sum_{i=1}^{\nu_Q} \mathbf{q}_i = \mathbf{0}$, producing the simplified expression in Eq. 4.12.

$$\begin{aligned}
d(T_1, T_2)^2 &= \sum_{i=1}^{\nu_Q} \|\mathbf{R}_1 \mathbf{q}_i + \mathbf{t}_1 - \mathbf{R}_2 \mathbf{q}_i - \mathbf{t}_2\|^2 = \\
&= \sum_{i=1}^{\nu_Q} (\mathbf{R}_1 \mathbf{q}_i + \mathbf{t}_1 - \mathbf{R}_2 \mathbf{q}_i - \mathbf{t}_2)^T (\mathbf{R}_1 \mathbf{q}_i + \mathbf{t}_1 - \mathbf{R}_2 \mathbf{q}_i - \mathbf{t}_2) = \\
&= \sum_{i=1}^{\nu_Q} (\mathbf{q}_i^T \mathbf{R}_1^T \mathbf{R}_1 \mathbf{q}_i + \mathbf{q}_i^T \mathbf{R}_1^T \mathbf{t}_1 - \mathbf{q}_i^T \mathbf{R}_1^T \mathbf{R}_2 \mathbf{q}_i - \mathbf{q}_i^T \mathbf{R}_1^T \mathbf{t}_2 + \mathbf{t}_1^T \mathbf{R}_1 \mathbf{q}_i + \mathbf{t}_1^T \mathbf{t}_1 \\
&\quad - \mathbf{t}_1^T \mathbf{R}_2 \mathbf{q}_i - \mathbf{t}_1^T \mathbf{t}_2 - \mathbf{q}_i^T \mathbf{R}_2^T \mathbf{R}_1 \mathbf{q}_i - \mathbf{q}_i^T \mathbf{R}_2^T \mathbf{t}_1 + \mathbf{q}_i^T \mathbf{R}_2^T \mathbf{R}_2 \mathbf{q}_i + \mathbf{q}_i^T \mathbf{R}_2^T \mathbf{t}_2 \\
&\quad - \mathbf{t}_2^T \mathbf{R}_1 \mathbf{q}_i - \mathbf{t}_2^T \mathbf{t}_1 + \mathbf{t}_2^T \mathbf{R}_2 \mathbf{q}_i + \mathbf{t}_2^T \mathbf{t}_2) \tag{4.10}
\end{aligned}$$

$$\begin{aligned}
d(T_1, T_2)^2 &= 2 \sum_{i=1}^{\nu_Q} \mathbf{q}_i^T \mathbf{q}_i + 2(\mathbf{t}_1 - \mathbf{t}_2)^T \mathbf{R}_1 \sum_{i=1}^{\nu_Q} \mathbf{q}_i + 2(\mathbf{t}_2 - \mathbf{t}_1)^T \mathbf{R}_2 \sum_{i=1}^{\nu_Q} \mathbf{q}_i \\
&\quad + \nu_Q \mathbf{t}_1^T \mathbf{t}_1 - 2\nu_Q \mathbf{t}_1^T \mathbf{t}_2 + \nu_Q \mathbf{t}_2^T \mathbf{t}_2 - 2\mathbf{R}_1^T \mathbf{R}_2 : \sum_{i=1}^{\nu_Q} \mathbf{q}_i \mathbf{q}_i^T \tag{4.11}
\end{aligned}$$

$$d(T_1, T_2)^2 = 2 \sum_{i=1}^{\nu_Q} \mathbf{q}_i^T \mathbf{q}_i + \nu_Q \mathbf{t}_1^T \mathbf{t}_1 - 2\nu_Q \mathbf{t}_1^T \mathbf{t}_2 + \nu_Q \mathbf{t}_2^T \mathbf{t}_2 - 2\mathbf{R}_1^T \mathbf{R}_2 : \sum_{i=1}^{\nu_Q} \mathbf{q}_i \mathbf{q}_i^T \tag{4.12}$$

$$\begin{aligned}
d(T_1, T_2)^2 &= 2 \sum_{i=1}^{\nu_Q} \mathbf{q}_i^T \mathbf{q}_i + \nu_Q \mathbf{t}_1^T \mathbf{t}_1 - 2\nu_Q \mathbf{t}_1^T \mathbf{t}_2 + \nu_Q \mathbf{t}_2^T \mathbf{t}_2 \\
&\quad - 2 \mathit{diag}(\mathbf{R}_1^T \mathbf{R}_2) \cdot \mathit{diag}\left(\sum_{i=1}^{\nu_Q} \mathbf{q}_i \mathbf{q}_i^T\right) \tag{4.13}
\end{aligned}$$

This approach can be taken one step further: rotate Q so that the matrix $\sum_{i=1}^{\nu_Q} \mathbf{q}_i \mathbf{q}_i^T$ becomes diagonal using EVD. Assuming Q has been rotated this way, we reach the final expression in Eq. 4.13, where $\mathit{diag}(\mathbf{M})$ stands for a vector of diagonal elements of a matrix \mathbf{M} . Neither the full matrix $\mathbf{R}_1^T \mathbf{R}_2$ nor $\sum_{i=1}^{\nu_Q} \mathbf{q}_i \mathbf{q}_i^T$ must be computed, only their diagonal elements that contribute to the result.

The final formulation only depends on a scalar value $\sum_{i=1}^{\nu_Q} \mathbf{q}_i^T \mathbf{q}_i$, and a vector $\mathit{diag}(\sum_{i=1}^{\nu_Q} \mathbf{q}_i \mathbf{q}_i^T)$, as has been noted but not exploited in [98]. These quantities can be precomputed for Q , which makes the distance evaluation independent of the size of Q , while still obtaining exactly the same result as when evaluating Eq. 4.9. This dissimilarity measure is fully independent of both rotation and translation and fulfills all properties of a metric.

Triangle Sum of Squares

One disadvantage of the vertex sum of squares is that its value depends on the sampling density of Q , which may be quite irregular. We propose addressing this issue by integrating the squared distance over all triangles rather than summing over vertices. Following the steps of Eq. 4.10-4.13, we

derive the value of the integral over a triangle t :

$$\begin{aligned} d_t(T_1, T_2)^2 &= \int_t \|\mathbf{R}_1 \mathbf{x} + \mathbf{t}_1 - \mathbf{R}_2 \mathbf{x} - \mathbf{t}_2\|^2 da \\ &= 2 \int_t \mathbf{x}^T \mathbf{x} da + a(t)(\mathbf{t}_1^T \mathbf{t}_1 - 2\mathbf{t}_1^T \mathbf{t}_2 + \mathbf{t}_2^T \mathbf{t}_2) \\ &\quad - 2 \text{diag}(\mathbf{R}_1^T \mathbf{R}_2) \cdot \text{diag}\left(\int_t \mathbf{x} \mathbf{x}^T da\right), \end{aligned}$$

where $a(t)$ is the area of triangle t . A full rigid transformation metric is obtained by summing over all mesh triangles:

$$d(T_1, T_2)^2 = \sum_{i=1}^{\nu_Q} d_{t_i^Q}(T_1, T_2)^2.$$

Note that computationally, the vertex sum and the triangle sum error metrics are equivalent, the only difference is the means of precomputing the constants needed for the evaluation. For a triangle $t = (A, B, C)$, it is not difficult to show that

$$\begin{aligned} \int_t \mathbf{x}^T \mathbf{x} da &= \int_t x^2 da + \int_t y^2 da + \int_t z^2 da, \\ \text{diag}\left(\int_t \mathbf{x} \mathbf{x}^T da\right) &= \left(\int_t x^2 da, \int_t y^2 da, \int_t z^2 da\right), \\ \int_t x^2 da &= J_t \frac{(A_x + B_x + C_x)C_x + A_x^2 + B_x^2 + A_x B_x}{12}, \\ J_t &= \|(B - A) \times (C - A)\|, \end{aligned}$$

while the remaining integrals of y^2 and z^2 are expressed analogously. Note that these simplified expressions can only be used for properly centered and pre-rotated mesh, the rotation is equivalent to the vertex sum, only again using integrals over triangles rather than sums. This does not, however, compromise the generality of the approach, since it can be equivalently expressed for general position, only resulting in a more expensive computation by a constant. Both vertex and triangle sum of squared distances depend on the input data, but only a few values suffice to evaluate the metrics in $\mathcal{O}(1)$ time.

L_1 Data Dependent Metrics

A typical problem of a squared distance error metric is its sensitivity to outliers: much more emphasis is given to larger distances than to small ones. This issue can be alleviated by using a different distance metric [12], such as L_1 , leading to:

$$d(T_1, T_2) = \sum_{i=1}^{\nu_Q} \|T_1(\mathbf{q}_i) - T_2(\mathbf{q}_i)\|, \quad (4.14)$$

or if a metric insensitive to irregular sampling is required,

$$d(T_1, T_2) = \sum_{i=1}^{\tau_Q} \int_{t_i^Q} \|T_1(\mathbf{x}) - T_2(\mathbf{x})\| da.$$

Notably, a very similar expression has been recently used in order to evaluate the smoothness (similarity) of local transformations needed for non-rigid alignment [71]. With the L_1 norm, it is substantially more difficult to evaluate such metric than when the L_2 norm was used. The vertex sum leads to linear computational complexity, while the triangle sum even leads to integrals that are difficult to evaluate in general. Therefore, instead of the full triangle sum L_1 metric, we propose using an approximation

$$d(T_1, T_2) = \sum_{i=1}^{\tau_Q} a(t_i^Q) \|T_1(\mathbf{c}(t_i^Q)) - T_2(\mathbf{c}(t_i^Q))\|, \quad (4.15)$$

where $\mathbf{c}(t)$ is the centroid of the triangle t . Even with this approximation, it is impractical to use either of the L_1 metrics in registration, because of their computational complexity. Nevertheless, L_1 metrics can be applied for evaluating the quality of the final alignment obtained by registration algorithms when a correct aligning transformation is known.

Alternatively, one could use the well known Hausdorff distance [40] computed between two versions of a mesh transformed by the two input transformations. Such approach, however, ignores the known point-to-point correspondences, and it also cannot be used in registration, because of the computational complexity of the Hausdorff distance evaluation [24].

Normalization

Normalized metrics based on the sum of squares can be written as

$$\frac{1}{r(Q)} \sqrt{\frac{1}{A(Q)} d(T_1, T_2)^2},$$

i.e. the normalization constant is $r(Q)\sqrt{A(Q)}$. The L_1 based metrics are normalized as

$$\frac{1}{r(Q)A(Q)} d(T_1, T_2),$$

where for the vertex-based metrics $A(Q) = \nu_q$ and for the triangle-based metrics $A(Q)$ is the surface area of Q , and $r(Q)$ is the estimated radius of Q . For the vertex-based metrics, $r(Q)$ is computed as the average distance of the vertices of Q from its centroid. For the triangle-based metrics, the average distance is computed for the centroids of all triangles weighted by

their areas and the same approach is used for computation of the centroid itself.

If the spread parameter D in the Gaussian kernel for the density peak location is constant, then the compound metrics have a single parameter c , which sets the global scale of the metric. The compound metrics are therefore generally evaluated as $\frac{c}{k}d(T_1, T_2)$ where d is a compound metric and k is the corresponding normalization constant.

4.3 Results

There are several parameters that influence the model algorithm, in our experiments, however, we found that the results are not very sensitive to the precise value of any of them. The parameters are:

- Geometric ring of radius 8 times the average edge length has been used as local neighborhood for the curvature estimation.
- 10^4 sample vertices of P were analyzed and their curvatures and principal directions were estimated.
- 10^4 candidate transformations were obtained by sampling Q .
- The distance δ used to determine the LCP of each candidate has been set to 2% of the mesh radius $r(Q)$.
- The minimum value of LCP for a candidate to be accepted has been set to 3%.

The registration algorithm was tested on 14 different realistic registration datasets. Datasets of varying character were chosen, with a range of properties that may cause problems with registration, such as small overlap, noise or potential to registration ambiguity. The reference implementation (current at the time of this research) of the state of the art algorithm Super4PCS [80] was only able to register 7 out of the 14 datasets, and even that was only achieved when the algorithm parameters were adjusted for different inputs. A detailed report on the performance of Super4PCS on the test data will be presented in Section 4.3.4.

We have tested our model algorithm with the composed metric with all the rotation metrics described in Section 4.2.1 and the two metrics based on sum of squares described in Section 4.2.2. We denote the composed metrics $d^{NAME}(T_1, T_2) = c_R \frac{d_R^{NAME}(\mathbf{R}_1, \mathbf{R}_2)}{k_R} + c_t \frac{\|t_1 - t_2\|}{k_t}$ where $NAME$ is the abbreviation for the corresponding rotation metric. Normalized metrics based on

vertex and triangle sum of squares are denoted d^{VSS} and d^{TSS} respectively. In all cases the object Q is centered at the origin.

For each of the 14 registration datasets the correct transformation is known and the error of the registration can therefore be measured as the difference between the resulting transformation and the correct one. To compute this error, we use the metric defined by Eq. (4.15), normalized as described in Section 4.2.2, with $c = 1$. We set a threshold ψ , below which we consider a registration successful. In our experiments, correct alignments have had error < 0.1 , in order to provide a certain tolerance we set $\psi = 0.15$. We also observed that in case of wrong alignment, the error is usually considerably larger, generally greater than 0.3. Therefore, choosing ψ between 0.1 and 0.3 does not make a big difference.

Before using the metrics, their parameters need to be set. In order to do that we designed and performed the following experiment. The spread parameter D in the Gaussian kernel was set to $D = 1$ and 900 different configurations of the c_R, c_t coefficients were created in the following way:

$$c_R = 1.5 \cdot 2^{i/3}, c_t = 1.5 \cdot 2^{j/3}$$

where $i, j = 0, 1, 2, \dots, 29$. The exponential progression was used because a good metric should not be very sensitive to the coefficient setting and, therefore, halving or doubling one of the coefficients should not impact the results very much. Since the metrics d^{VSS}, d^{TSS} only depend on a single parameter, we set it as $c = c_R$ while the value of c_t remains unused. With these metrics, the performance does not depend on j .

For each metric, we ran the registration on all the 14 datasets, 5 times at every combination of i and j , which gives 70 tests for every configuration in total. To visualize the results, we created a bitmap of 30×30 pixels, where the index of each pixel corresponds to a single configuration of i and j (i horizontal, j vertical from top to bottom). A pixel is white when the corresponding configuration was successful in all 70 registration cases, i.e. error $< \psi$, otherwise it is black.

For each metric, we computed the centroid of all the white pixels in the bitmap and we used the i, j indices of the centroid pixel to acquire the optimal values of c_R, c_t , which are shown in Table 4.1. The bitmaps are shown in Figure 4.2, the centroid pixels are colored red. Since the algorithm is non-deterministic, even for a good configuration of the coefficients, there is always a certain chance of failure. Therefore, a black pixel does not necessarily mean that the corresponding configuration is bad, especially when it is surrounded mostly by white pixels. On the other hand, if a black pixel




Figure 4.2: Bitmaps of the i, j configurations of the coefficient settings for all the metrics. White pixel indicates a successful registration in all 70 test cases, black pixel indicates at least one failure, red pixel shows the centroid of the white pixels and indicates the optimal configuration.

Table 4.1: Optimal coefficients of the transformation distance metrics.

| | DEA | DQ | ADPQ | DPQ | DIM | VSS | TSS |
|-------|------|------|------|------|------|------|------|
| c_R | 4.76 | 9.52 | 9.52 | 96 | 7.55 | 7.55 | 9.52 |
| c_t | 12 | 7.55 | 7.55 | 9.52 | 6 | - | - |

lies in a neighborhood of black pixels, it suggests a configuration far from optimal.

Figure 4.3a shows the registration result of the model algorithm using the d^{TSS} metric with $c = 9.52$ for the *Kac* dataset. The error of this registration was 0.007. For comparison, Figures 4.3b, 4.3c show two artificially created alignments with errors of 0.077 and 0.153 respectively, which are close to $\frac{\psi}{2}$ and ψ .

To perform a comparison of all the metrics, we ran the registration on each of the 14 datasets 1000 times, using the optimal parameters found previously. Table 4.2 shows the average error for each metric on every dataset and the fail count (FC), i.e. how many times of the 1000 tests the registration resulted in error $> \psi$. The error cells are colored according to the order of the error values for the given dataset, i.e. the cell with the largest error is red and the color transits through orange and yellow to green which indicates




Figure 4.3: Registration result of the model algorithm using the d^{TSS} metric with $c = 9.52$ (a) and its qualitative comparison to two artificial alignments with different errors (b), (c).

Table 4.2: Comparison of the transformation distance metrics. The Error column shows the average errors for corresponding datasets, the FC (fail count) column shows how many times the registration produced an alignment with error $> \psi$.

| Dataset | DEA | | DQ | | ADPQ | | DPQ | | DIM | | VSS | | TSS | |
|-----------|--------|-----|--------|-----|--------|-----|--------|-----|--------|-----|--------|-----|--------|-----|
| | Error | FC | Error | FC | Error | FC | Error | FC | Error | FC | Error | FC | Error | FC |
| Armadillo | 0.0683 | 23 | 0.0479 | 0 | 0.0485 | 0 | 0.0486 | 0 | 0.0463 | 0 | 0.0455 | 0 | 0.0480 | 0 |
| Bird | 0.0673 | 0 | 0.0672 | 0 | 0.0672 | 0 | 0.0669 | 0 | 0.0672 | 0 | 0.0673 | 0 | 0.0672 | 0 |
| Bubba | 0.0032 | 0 | 0.0035 | 0 | 0.0039 | 0 | 0.0059 | 0 | 0.0039 | 0 | 0.0034 | 0 | 0.0023 | 0 |
| Buddha | 0.0363 | 0 | 0.0255 | 0 | 0.0252 | 0 | 0.0246 | 0 | 0.0241 | 0 | 0.0242 | 0 | 0.0242 | 0 |
| Coati | 0.0283 | 0 | 0.0211 | 0 | 0.0208 | 0 | 0.0212 | 0 | 0.0202 | 0 | 0.0199 | 0 | 0.0200 | 0 |
| Dragon | 0.0258 | 0 | 0.0222 | 0 | 0.0221 | 0 | 0.0233 | 0 | 0.0221 | 0 | 0.0211 | 0 | 0.0209 | 0 |
| Eggs | 0.0758 | 98 | 0.0852 | 120 | 0.0868 | 123 | 0.0967 | 142 | 0.1003 | 143 | 0.1202 | 179 | 0.0915 | 140 |
| Head | 0.0079 | 0 | 0.0081 | 0 | 0.0082 | 0 | 0.0089 | 0 | 0.0083 | 0 | 0.0081 | 0 | 0.0075 | 0 |
| Hippo | 0.0760 | 32 | 0.0560 | 0 | 0.0558 | 0 | 0.0566 | 0 | 0.0556 | 0 | 0.0528 | 0 | 0.0538 | 0 |
| Kachel | 0.0187 | 0 | 0.0152 | 0 | 0.0156 | 0 | 0.0168 | 0 | 0.0153 | 0 | 0.0156 | 0 | 0.0155 | 0 |
| Oscar | 0.0044 | 0 | 0.0039 | 0 | 0.0040 | 0 | 0.0048 | 0 | 0.0039 | 0 | 0.0036 | 0 | 0.0034 | 0 |
| Suzanne | 0.0139 | 0 | 0.0133 | 0 | 0.0139 | 0 | 0.0144 | 0 | 0.0130 | 0 | 0.0132 | 0 | 0.0130 | 0 |
| Teeth | 0.0160 | 0 | 0.0137 | 0 | 0.0135 | 0 | 0.0148 | 0 | 0.0131 | 0 | 0.0129 | 0 | 0.0127 | 0 |
| Testbody | 0.0189 | 0 | 0.0175 | 0 | 0.0172 | 0 | 0.0197 | 0 | 0.0181 | 0 | 0.0186 | 0 | 0.0171 | 0 |
| Total | 0.0329 | 153 | 0.0286 | 120 | 0.0288 | 123 | 0.0302 | 142 | 0.0294 | 143 | 0.0305 | 179 | 0.0284 | 140 |

the cell with the lowest error. The colors do not correspond to the absolute error values, they only indicate their order for the given dataset. In the rest of this section, we refer to the datasets using a three letter abbreviation. Table 4.3 shows the vertex counts and triangle counts of both P and Q objects of the testing datasets.

The *Egg* dataset stands out from the others, because it is the only one with non-zero fail count for every metric, suggesting that this model is problematic in some way. The P and Q objects of *Egg* together with their correct

Table 4.3: Vertex counts and triangle counts of the P and Q objects of all the testing datasets.

| Dataset | Vertex count | | | Triangle count | | |
|---------|--------------|--------|---------|----------------|---------|---------|
| | P | Q | Total | P | Q | Total |
| Arm | 33902 | 36107 | 70009 | 65379 | 69985 | 135364 |
| Bir | 2656 | 3049 | 5705 | 5001 | 5758 | 10759 |
| Bub | 10912 | 14038 | 24950 | 21285 | 27643 | 48928 |
| Bud | 59544 | 51263 | 110807 | 116804 | 100640 | 217444 |
| Coa | 28107 | 28180 | 56287 | 53715 | 53829 | 107544 |
| Dra | 43181 | 39524 | 82705 | 83609 | 76980 | 160589 |
| Egg | 750729 | 761887 | 1512616 | 1480240 | 1501582 | 2981822 |
| Hea | 25827 | 24273 | 50100 | 50336 | 47316 | 97652 |
| Hip | 30519 | 21925 | 52444 | 59166 | 42254 | 101420 |
| Kac | 28393 | 28142 | 56535 | 56134 | 55583 | 111717 |
| Osc | 169862 | 166624 | 336486 | 336136 | 329268 | 665404 |
| Suz | 235735 | 236646 | 472381 | 469745 | 471528 | 941273 |
| Tee | 57024 | 61478 | 118502 | 106930 | 116652 | 223582 |
| Tes | 423715 | 422308 | 846023 | 832938 | 829762 | 1662700 |




Figure 4.4: The *Egg* dataset.

alignment are depicted in Figure 4.4. The most likely cause of the failure is the presence of large planar/spherical areas that lead to many incorrect matches, because of the simple curvature based descriptor. Surprisingly, the d^{DEA} metric has the lowest fail count as well as the lowest average error for *Egg*, however, it seems that under the circumstances of a failing descriptor, this result is caused by random influences. Otherwise, d^{DEA} has the largest average error and it is also the only metric with non-zero fail count for two additional datasets: *Arm* and *Hip*. The lowest total average error was achieved with the d^{TSS} metric with d^{DQ} and d^{ADPQ} having slightly larger, but very similar errors. Also, without the problematic *Egg* dataset, the total average errors of the metrics are d^{DEA} : 0.0296, d^{DQ} : 0.0242, d^{ADPQ} :




Figure 4.5: The Q object of the *Bub* dataset. The non-uniform distribution of the vertices is quite clear in the detailed view.

0.0243, d^{DPQ} : 0.0251, d^{DIM} : 0.0239, d^{VSS} : 0.0236, d^{TSS} : 0.0235, with both the compound metrics d^{VSS} and d^{TSS} having lower error than all the other metrics.

Figure 4.6 shows the ratio of the average error of d^{TSS} (chosen as reference because of its lowest average error) and the average errors of the remaining metrics for all 14 datasets. The figure, together with Table 4.2, shows that there are quite many cases where the d^{DEA} metric shows noticeably larger error than all the remaining metrics (*Arm*, *Bud*, *Coa*, *Dra*, *Hip*, *Kac*, *Tee*). This demonstrates how unreliable a composed metric can be when it uses the non-bi-invariant rotation metric based on distance of Euler angles and, therefore, application of such a metric is not advisable. On several datasets, the d^{DPQ} metric also shows slightly, but noticeably larger error than the other metrics (*Tee*, *Hea*, *Kac*, *Osc*), except for the *Bub* dataset, where the error of d^{DPQ} is much larger in terms of relative comparison.

For the *Bub* dataset, the d^{TSS} metric has the lowest average error and, in relative comparison, even substantially lower than d^{VSS} , even though d^{VSS} and d^{TSS} are based on the same concept. The explanation might be that the *Bub* dataset has quite non-uniformly distributed vertices as shown in Figure 4.5. In such a case, the d^{TSS} metric is expected to work more precisely than d^{VSS} , because it accounts for the triangle size distribution, having more exact information about the object’s shape. Similar reason might be behind the lower error of d^{TSS} for other datasets, although there the non-uniform distribution is not so pronounced as with the *Bub* dataset.

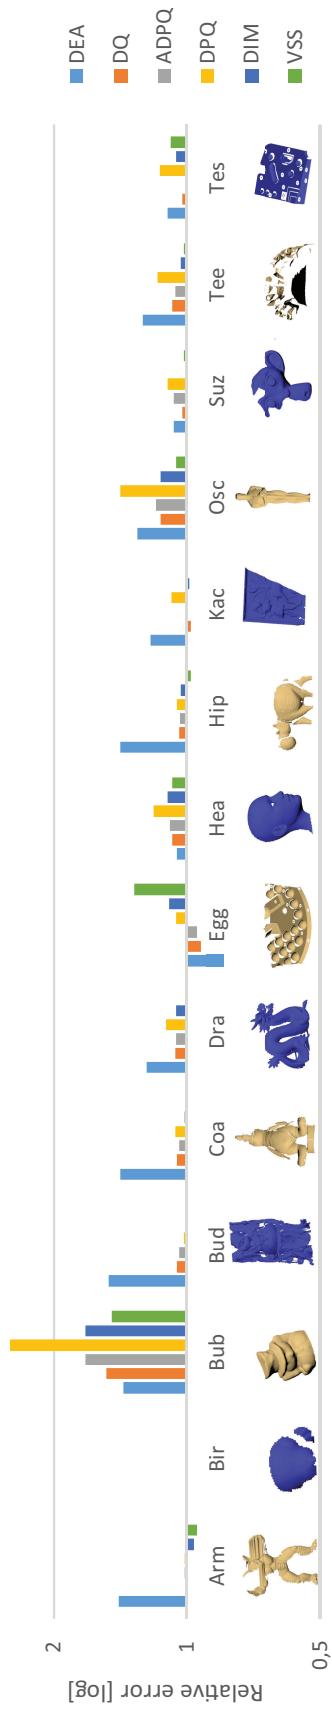


Figure 4.6: Average registration error of each metric normalized by the average error of d^{TSS} for each dataset, shown in log scale.

Table 4.4: Fail counts of the registration with LCP used for evaluation of the consensus with $\delta = \delta' r(Q)$ for different values of δ' .

| Dataset | δ' | | | | |
|---------|-----------|-------|------|------|------|
| | 0.0025 | 0.005 | 0.01 | 0.02 | 0.04 |
| Arm | 619 | 8 | 0 | 0 | 0 |
| Bir | 887 | 757 | 633 | 983 | 1000 |
| Bub | 0 | 0 | 0 | 16 | 1000 |
| Bud | 112 | 0 | 0 | 0 | 0 |
| Coa | 59 | 0 | 0 | 0 | 0 |
| Dra | 330 | 1 | 0 | 0 | 0 |
| Egg | 215 | 4 | 0 | 0 | 0 |
| Hea | 0 | 0 | 0 | 0 | 645 |
| Hip | 708 | 324 | 219 | 76 | 3 |
| Kac | 368 | 0 | 0 | 0 | 327 |
| Osc | 0 | 0 | 0 | 0 | 0 |
| Suz | 90 | 8 | 6 | 438 | 1000 |
| Tee | 0 | 0 | 0 | 0 | 0 |
| Tes | 375 | 55 | 1 | 0 | 0 |
| Total | 3763 | 1157 | 859 | 1513 | 3975 |

4.3.1 Comparing to LCP

We also show results of the registration algorithm where LCP was used instead of the density in the transformation space for evaluating the consensus. The candidate transformation for which the LCP value was the largest was selected as the result of the registration. Such an approach is quite common and was used e.g. by the state of the art registration algorithm Super4PCS [80].

We experimented with the δ parameter of LCP chosen as $\delta = \delta' r(Q)$ where $r(Q)$ is the radius of Q , and we tested several values of δ' . Table 4.4 shows the fail counts of the registration with LCP used for the consensus evaluation with different values of δ' (run 1000 times on each dataset). There is no universal value of δ' and the total fail count is non-negligible for all of the tested values. Although increasing or decreasing δ' leads to decrease of fail count for some datasets, at the same time it causes increase of fail counts for other datasets. This demonstrates that the results of registration algorithms using LCP to select the best candidate transformation depend quite strongly on the choice of δ . On the other hand, when using the density peak location in the space of transformations, although there are some para-

meters to set (either one or two, depending on the metric), the registration is not very sensitive to their setting, as shown in Figure 4.2 (note that the coefficient growth is *exponential* in both directions).

4.3.2 Non-Centered Object and Density Visualization

To verify the dependence of the metrics on position as discussed in Section 4.2.1, the registration was executed 1000 times on the *Arm* dataset with the composed metric d^{DIM} , but after centering Q , it was further translated by a vector $\xi \cdot r(Q) \cdot [1, 1, 1]^T$ where ξ determines the final distance of Q from the origin. For $\xi = 0$, the fail count is 0 as in Table 4.2. For $\xi = 2, 5, 15$ and 100, the fail counts were 24, 205, 648 and 865 respectively. Centering Q is therefore truly crucial for reliable results when using a composed metric and the farther the object is from the origin, the worse results can be expected. In order to visualize the sampling in $SE(3)$ we use a standard multidimensional scaling. First, we compute 3D coordinates that match the candidate distances provided by a metric as closely as possible, obtaining a 3D point cloud. The point cloud is converted to volume data representing the density as defined in Eq. 4.2 and visualized by standard direct volume rendering, choosing an informative viewpoint manually. The visualizations of the decentered dataset registrations are shown in Figure 4.7.

Fig. 4.8 shows another such visualization. It demonstrates that a wrong alignment may produce the highest LCP. On the other hand, a cluster of candidates of rather mediocre LCP is obtained in the vicinity of the correct alignment.

4.3.3 Noisy Data

The plot in Figure 7.8 shows how the model registration algorithm (using d^{TSS} , $c = 9.52$) behaves on the *Arm* dataset in presence of Gaussian noise with varying standard deviation. The standard deviation is relative to $r(Q)$ and the figure also shows the resulting alignment for deviation of 0.016. The error grows quite slowly up to deviation of 0.018, while for higher values the noise prevents correct registration. Similar results were obtained for other datasets.

4.3.4 Comparing to Super4PCS

Table 4.5 shows the errors of the Super4PCS registration method, i.e. of its implementation that was the latest at the time of the research (April 8,

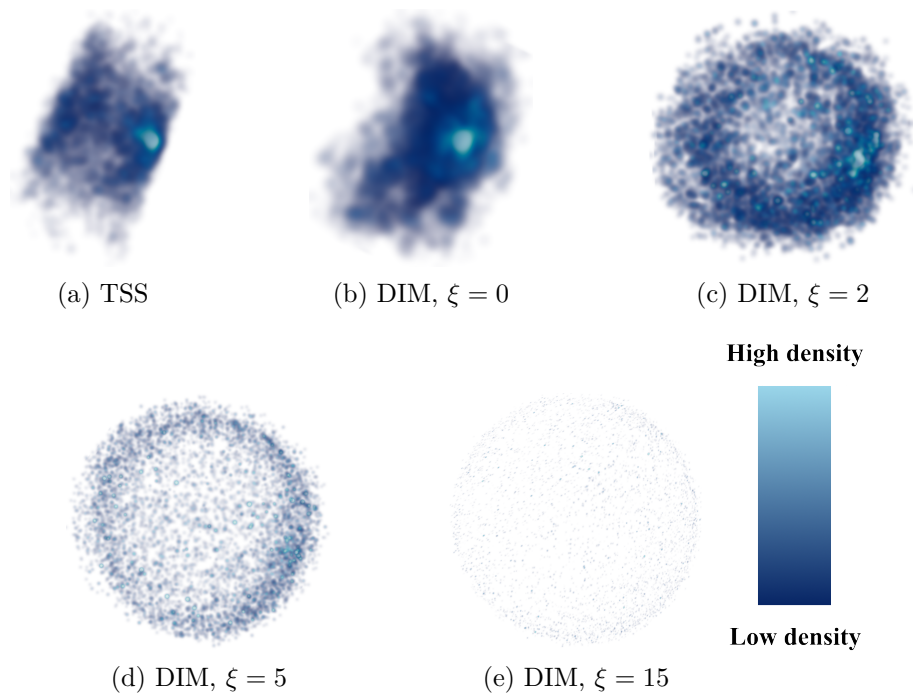


Figure 4.7: Projection of the candidates into E^3 using MDS for the Arm dataset with TSS and DIM with different values of ξ .




Figure 4.8: 2D visualization of the space of rigid transformation candidates. The hue reflects the score of each candidate, while the brightness reflects the estimated density. The alignments on the left, even though they have the highest score, represent a wrong alignment. A much better alignment on the right can be identified by looking for a density peak.

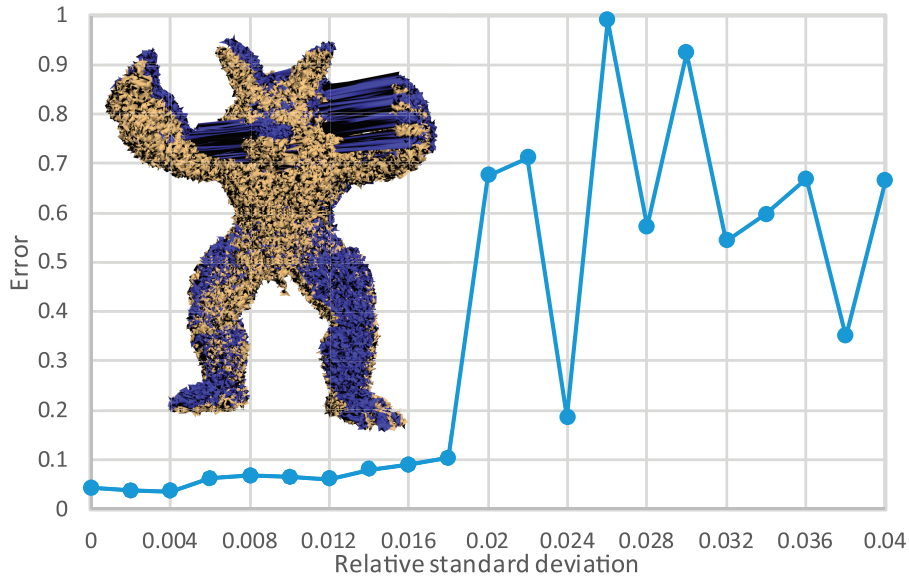


Figure 4.9: Dependence of registration error on Gaussian noise with varying standard deviation relative to $r(Q)$ for the Arm dataset. The depicted alignment was achieved with deviation of 0.016.

2019). We ran Super4PCS with different values of the δ parameter of its LCP evaluation for all the datasets. The value of delta is set relative to the radius of Q as $\delta = \delta' r(Q)$. The cells with error $\leq \psi$ are marked bold, in all other cases the registration was considered a failure. Names of the datasets for which the registration was succesful for at least one value of δ' are also marked bold, for all the other datasets the registration failed for all the values.

The largest number of successful registrations (7) was achieved using $\delta' = 0.106, 0.121$ and 0.184 . Since there is quite a large gap between 0.121 and 0.184 , we select $\delta' = 0.11$ as approximately optimal because it is close to both 0.121 and 0.106 . Table 4.6 shows comparison of Super4PCS with $\delta' = 0.11$ to the model RANSAC algorithm where d^{TSS} was used as the metric with $c = 9.52$. For each dataset, ICP [10] alignment was used after the global registration with the same δ setting as Super4PCS ($\delta' = 0.11$). We only show binary evaluation of whether the final registration was successful instead of showing the error, since in the case of successful alignment the final error is dictated by the ICP rather than by the global registration, and in the case of incorrect result the error value is irrelevant. The running times do not include ICP and the measurements were done on a computer with CPU *Intel Core i7-4770* (clock rate 3.4 GHz, 4 cores, L1 cache 256 kB, L2

Table 4.5: Errors of Super4PCS with different values of $\delta = \delta'r(Q)$ for all the datasets.

| δ' | Arm | Bir | Bub | Bud | Coa | Dra | Egg | Tee | Hea | Hip | Kac | Osc | Tes | Suz |
|-----------|--------------|-------|-------|--------------|--------------|--------------|--------------|-------|-------|--------------|-------|--------------|--------------|-------|
| 0.010 | 1.257 | 1.941 | 1.306 | 1.594 | 1.387 | 1.581 | 1.361 | 1.416 | 1.704 | 1.366 | 1.801 | 1.075 | 1.816 | 1.366 |
| 0.011 | 0.220 | 1.941 | 1.303 | 1.588 | 1.383 | 1.456 | 1.348 | 1.396 | 1.712 | 1.382 | 1.834 | 0.543 | 1.822 | 1.376 |
| 0.013 | 1.214 | 1.941 | 1.318 | 1.598 | 1.387 | 1.457 | 1.363 | 1.380 | 1.705 | 1.353 | 1.826 | 1.073 | 1.402 | 1.341 |
| 0.015 | 1.242 | 1.941 | 1.622 | 1.776 | 1.245 | 1.974 | 1.346 | 1.383 | 1.703 | 1.417 | 1.841 | 1.073 | 1.814 | 1.596 |
| 0.017 | 0.099 | 1.855 | 1.296 | 1.590 | 1.385 | 0.418 | 1.380 | 1.411 | 1.710 | 1.357 | 1.828 | 2.016 | 1.824 | 1.550 |
| 0.020 | 0.334 | 0.890 | 1.464 | 1.592 | 1.389 | 0.066 | 0.166 | 1.405 | 1.696 | 1.048 | 1.761 | 1.073 | 0.391 | 2.103 |
| 0.023 | 0.134 | 1.319 | 2.262 | 0.373 | 0.225 | 0.587 | 1.363 | 1.398 | 1.718 | 1.340 | 1.852 | 0.151 | 1.327 | 2.099 |
| 0.026 | 0.232 | 1.577 | 1.464 | 0.800 | 1.380 | 0.197 | 1.151 | 1.397 | 1.926 | 1.357 | 1.844 | 0.081 | 1.413 | 1.562 |
| 0.030 | 0.131 | 1.026 | 1.432 | 0.273 | 1.380 | 2.068 | 1.295 | 1.399 | 1.915 | 1.096 | 1.783 | 0.067 | 0.207 | 2.039 |
| 0.035 | 0.080 | 1.742 | 1.454 | 1.589 | 1.380 | 0.090 | 1.568 | 1.422 | 1.780 | 0.443 | 1.779 | 0.146 | 0.162 | 1.752 |
| 0.040 | 0.285 | 1.277 | 1.495 | 0.280 | 1.378 | 0.039 | 1.909 | 1.442 | 1.671 | 0.106 | 1.825 | 0.135 | 0.122 | 1.825 |
| 0.046 | 0.121 | 1.318 | 1.397 | 0.209 | 0.258 | 0.060 | 0.769 | 1.443 | 1.915 | 0.094 | 1.930 | 0.161 | 1.974 | 1.822 |
| 0.053 | 0.047 | 1.729 | 1.485 | 0.091 | 0.484 | 0.136 | 0.557 | 1.431 | 1.964 | 0.350 | 1.848 | 0.037 | 1.275 | 1.876 |
| 0.061 | 0.070 | 1.411 | 1.307 | 0.064 | 0.120 | 0.065 | 1.606 | 1.414 | 1.997 | 0.153 | 1.721 | 0.035 | 0.143 | 1.888 |
| 0.070 | 0.047 | 1.421 | 1.482 | 0.178 | 0.065 | 0.035 | 0.207 | 1.410 | 1.755 | 0.083 | 1.593 | 0.044 | 0.237 | 1.877 |
| 0.080 | 0.079 | 1.570 | 1.360 | 0.071 | 0.134 | 0.102 | 0.383 | 1.420 | 1.554 | 1.685 | 1.858 | 0.079 | 0.105 | 1.771 |
| 0.092 | 0.045 | 1.028 | 1.525 | 0.087 | 0.051 | 0.059 | 0.519 | 1.471 | 1.638 | 0.168 | 1.452 | 0.105 | 1.163 | 2.003 |
| 0.106 | 0.071 | 1.481 | 1.548 | 0.052 | 0.122 | 0.086 | 0.490 | 1.472 | 1.946 | 0.055 | 1.474 | 0.074 | 0.040 | 2.030 |
| 0.121 | 0.067 | 0.531 | 1.345 | 0.134 | 0.104 | 0.084 | 0.121 | 0.467 | 1.623 | 0.117 | 1.758 | 0.165 | 0.076 | 1.979 |
| 0.139 | 0.082 | 1.558 | 1.467 | 0.075 | 0.164 | 0.126 | 0.157 | 1.558 | 1.994 | 0.227 | 1.237 | 0.094 | 0.086 | 1.837 |
| 0.160 | 0.103 | 1.473 | 1.472 | 0.062 | 0.048 | 0.069 | 0.528 | 1.503 | 1.666 | 0.332 | 1.639 | 0.109 | 0.072 | 1.781 |
| 0.184 | 0.102 | 1.410 | 1.521 | 0.073 | 0.071 | 0.128 | 0.548 | 1.585 | 1.957 | 0.140 | 1.857 | 0.149 | 0.117 | 1.723 |
| 0.211 | 0.119 | 1.475 | 1.505 | 0.161 | 0.141 | 0.136 | 0.314 | 1.500 | 1.612 | 0.251 | 1.639 | 0.171 | 0.148 | 1.779 |
| 0.243 | 0.065 | 1.352 | 1.521 | 0.118 | 0.038 | 0.178 | 0.383 | 0.468 | 1.947 | 0.174 | 1.621 | 0.260 | 0.148 | 1.769 |
| 0.279 | 0.162 | 1.358 | 1.536 | 0.152 | 0.178 | 0.160 | 0.401 | 1.797 | 1.750 | 0.201 | 1.764 | 0.147 | 0.076 | 1.909 |
| 0.320 | 0.090 | 1.430 | 1.518 | 0.117 | 0.120 | 1.228 | 0.386 | 1.731 | 1.907 | 0.309 | 1.843 | 0.127 | 2.002 | 1.835 |
| 0.368 | 0.256 | 1.250 | 1.533 | 0.183 | 0.117 | 0.202 | 0.442 | 1.936 | 1.920 | 0.223 | 1.898 | 0.382 | 0.342 | 1.767 |
| 0.422 | 1.149 | 1.431 | 1.620 | 1.938 | 0.099 | 1.412 | 0.502 | 2.003 | 1.895 | 0.130 | 1.509 | 0.603 | 1.349 | 1.836 |
| 0.485 | 0.354 | 1.222 | 1.622 | 1.964 | 0.371 | 1.610 | 0.462 | 2.028 | 1.785 | 0.427 | 1.382 | 0.244 | 1.387 | 1.826 |
| 0.557 | 1.676 | 1.464 | 1.628 | 0.372 | 0.185 | 0.115 | 0.676 | 2.358 | 2.043 | 0.391 | 1.800 | 0.611 | 1.385 | 1.789 |
| 0.640 | 2.128 | 1.365 | 1.533 | 0.446 | 0.115 | 0.313 | 0.746 | 2.422 | 1.837 | 0.328 | 1.405 | 0.652 | 1.360 | 1.621 |
| 0.735 | 2.017 | 1.497 | 1.334 | 0.557 | 0.259 | 0.333 | 1.901 | 0.329 | 1.690 | 0.283 | 1.463 | 0.491 | 1.529 | 1.748 |
| 0.844 | 1.462 | 1.278 | 1.552 | 1.187 | 0.426 | 2.011 | 1.941 | 2.626 | 2.069 | 1.747 | 1.428 | 0.643 | 1.365 | 2.560 |
| 0.970 | 1.161 | 1.284 | 0.834 | 0.869 | 0.303 | 0.819 | 1.929 | 3.507 | 1.779 | 1.680 | 2.208 | 0.720 | 1.800 | 1.514 |
| 1.114 | 2.713 | 2.683 | 1.575 | 1.917 | 0.554 | 2.168 | 0.927 | 4.015 | 1.934 | 1.435 | 1.944 | 0.602 | 1.675 | 1.688 |

cache 1 MB, L3 cache 8 MB) and 16 GB of memory with clock rate of 1.6 GHz with *Windows 10 64-bit* operating system.

If the global registration resulted in a good alignment, then the subsequent ICP further strongly increased its precision for most datasets. The only exception was the *Bir* dataset, where the model registration algorithm actually found a very good alignment, but the ICP made it much worse, resulting in an unsuccessful registration. Super4PCS for this dataset resulted in a bad alignment even without ICP. In total, including ICP did not lead to an improvement of the registration using Super4PCS, which was again successful with 7 datasets. The running times of Super4PCS are comparable to those of our model registration algorithm.

Table 4.6: Comparison of the model RANSAC registration algorithm using d^{TSS} with $c = 9.52$ to Super4PCS with $\delta' = 0.11$, after the global registration ICP was performed also with $\delta' = 0.11$, the running times do not include ICP.

| Dataset | Success | | Time [ms] | |
|---------|---------|-------|-----------|-------|
| | Model | S4PCS | Model | S4PCS |
| Arm | yes | yes | 2406 | 9807 |
| Bir | no | no | 820 | 298 |
| Bub | yes | no | 1337 | 596 |
| Bud | yes | yes | 2987 | 7203 |
| Coa | yes | yes | 1958 | 701 |
| Dra | yes | yes | 3071 | 9084 |
| Egg | yes | no | 20654 | 19592 |
| Hea | yes | no | 1369 | 9693 |
| Hip | yes | yes | 1495 | 1372 |
| Kac | yes | no | 1496 | 635 |
| Osc | yes | yes | 5018 | 4182 |
| Suz | yes | no | 6641 | 5525 |
| Tee | yes | no | 2361 | 1432 |
| Tes | yes | yes | 10926 | 10382 |

4.3.5 Limitations

Since the model RANSAC algorithm is feature-based, it naturally tends to fail if very strong noise is present in P and Q because in such a case the curvature based features lose reliability. For similar reason, the algorithm does not behave well in case of objects with repetitive patterns or planar areas (e.g. the *Egg* dataset) because there are multiple different spots with similar curvatures, which creates a large potential for registration ambiguity.

The results of the algorithm depend on the size of the overlap of the inputs, both when using the composite and the compound metrics. Either the centering step for composite metrics or the pre-computations for compound metrics may lead to skewed results when the overlap is small, which in turn may negatively influence the registration reliability. This may occur in practice, when for example Q is a result of merging several partial scans, and has therefore only a small overlap with P . In general, it is therefore advisable to choose as Q the smaller of the two input models, since it has a larger relative portion of overlap.

4.4 Summary

A method of identifying consensus in RANSAC surface registration has been discussed. Instead of searching for a consensus of data points, we look for a consensus of candidate solutions. The consensus is identified with a density peak in the $SE(3)$ space, which is found using the vantage point tree data structure.

In order to locate the density peak, a proper metric is needed in the solution space. We have discussed multiple choices known in the literature and proposed their modified versions. While composed metrics can be applied for this purpose, they suffer from fundamental drawbacks, such as the necessity of setting two free parameters and the lack of translation independence that is merely mitigated by centering the input object. Using the distance of Euler angles as a rotation metric is particularly discouraged, because of its unreliable behaviour.

On the other hand, the equally fast *compound metrics* d^{VSS} and d^{TSS} behave well, regardless of the position of Q , and they exhibit the lowest total average error except for the *Eggs* dataset, where the failure is likely caused by the used feature vector rather than by the metric. Additionally, the proposed d^{TSS} metric is also independent of the sampling density, which makes it the method of choice not only in this, but also in other applications.

Our model registration algorithm turns out to work on par or better than the state of the art algorithm Super4PCS both in terms of speed and reliability. The algorithm is able to identify the best alignment even in situations when the LCP score of a wrong transformation is better than that of the near-optimal alignments. The registration is truly global, fully independent of the initial location of the inputs.

The reliability could be further improved by incorporating more advanced building blocks, such as better local feature vectors or more accurate candidate filtering, since we have chosen rather simple solutions in order to test the consensus evaluation step. As with all global registration methods, a refinement using a variant of ICP is expected to improve the result. The reference implementation of our algorithm is available for download at [45].

The proposed consensus identification strategy could be applied in other registration or symmetry detection algorithms as well, particularly easily in those that can be interpreted in the RANSAC framework, such as Super4PCS. Moreover, it could also be used in other applications, such as dynamic surface segmentation based on motion similarity and others.

We have shown that even the composed metrics are data dependent in certain sense, relying on the knowledge of the input data centroid. Com-

pound metrics based on the sum of squares retain more information about the input objects in the form of precomputed constants, however, the information is still quite limited. An interesting research direction could be investigating the possibility of formulating metrics that take even more information about the inputs into account, however, without compromising the computational complexity as in the case of L_1 metrics.

5 Plane Space Representation in Mode-based Symmetry Plane Detection

Some of the symmetry detection methods described in Chapter 3 can be interpreted as an implementation of one specific and very popular approach. It is to create a number of candidate transformations by matching different points or parts of the input object and then find those transformations that occur most often in the transformation space. This can also be described as seeking modes (places of the highest density) in the transformation space so we call this approach Mode-based symmetry detection and it can be classified as a kind of the more general RANSAC approach. For reference, see e.g. the method of Mitra et al. [83], which is most likely the first method that used this approach and one of the most commonly known methods for symmetry detection in general. This approach can be applied in algorithms detecting symmetries of various types, however, in this chapter we only focus on its application for detecting the planes of symmetry (reflectional symmetries) of 3D objects. The content of this chapter was previously published in [50].

The method [83] together with its newer and improved variant [107] are the two more general representatives of the Mode-based approach that can be used to find symmetries of quite a general type. In their case the symmetry transformations can contain rotation, with or without reflection, translation and uniform scaling. The transformation type, however, can be restricted to any subgroup of these general transformations, so these methods can be used to detect reflectional symmetries (symmetry planes) as well. To find the modes both these methods use clustering. Some form of the clustering method [83] was used to detect symmetry planes e.g. in [72], on models of damaged skulls, or in [79] and [62], suggesting its popularity for symmetry plane detection.

In Chapter 4 we presented a Mode-based approach used in rigid surface registration. In this case the candidate space contained rigid transformations and to find a single mode, a density peak estimation algorithm was used. This approach could also possibly be used to find the global plane of symmetry of an object, only with planes as candidates instead of rigid transformations. The Hough transform-based method [16] and the Monte Carlo algorithm used in [96] for 3D surfaces can also be interpreted as Mode-based

approaches, only designed to detect specifically symmetry planes. However, instead of finding the mode in a continuous environment, they divide the space of planes into discrete bins to count plane occurrences.

Regardless of the specific algorithm, any Mode-based method for symmetry plane detection requires defining some representation of the space of planes and the result of the method will always somewhat depend on the representation selected. The important aspect of the plane space representation is how well distances between points in the space correspond to the actual similarity/dissimilarity of the planes in E^3 . In context of symmetry detection, planes can also be understood as transformations reflecting points over the given plane.

A useful observation is that the mode(s) can be found in an arbitrary non-Euclidean space only using distances between the points in the space, e.g. as described in [120] or in Section 4.1.3 where we also described how proximity queries in non-Euclidean spaces can be accelerated using the Vantage Point Tree data structure [126]. In Section 4.2 we thoroughly analyzed the problem of computing distances between rigid transformations, however, for reflection transformations or planes the same problem does not seem to be sufficiently addressed in the literature despite the fact that the Mode-based approach is quite popular in the field of symmetry plane detection.

In this chapter we describe and analyze several different representations of the space of planes and for each representation we discuss possible ways of defining a *reasonable* distance function that could be used in Mode-based symmetry plane detection. We compare these distance functions to a single distance function that we consider the ground truth but which cannot be used in practice because of its large computation cost.

Some of the representations we describe could also be more appropriate for different applications, such as visualization. In general, the information about the plane space representations described below can be useful in any other application, outside the scope of symmetry detection, where some form of plane representation is needed, although the presented distance functions might require some adjustments according to the specific application. We therefore believe that researchers from various fields, not restricted to symmetry plane detection, could benefit from the results of our research.

5.1 Background

Here we provide some necessary background for the rest of this chapter. A general plane P can be defined by the four a, b, c, d coefficients as described

in Section 2.4.3. Here we denote $\mathbf{n} = [a, b, c]^T$ the normal vector of the plane and for the purpose of this chapter we always consider the coefficients to be normalized such that $\|\mathbf{n}\| = 1$ in which case d is the signed distance of the plane from the origin. A function $\mathbf{r}_P(\mathbf{x}) \in E^3$ that reflects an arbitrary point $\mathbf{x} \in E^3$ over the plane P can be then defined as shown in Eq. 5.1.

$$\mathbf{r}_P(\mathbf{x}) = \mathbf{x} - 2(\mathbf{n}^T \mathbf{x} + d)\mathbf{n} \quad (5.1)$$

5.1.1 Candidate Creation Algorithm

In order to demonstrate and compare the plane space representations on realistic data we use the following model algorithm for creating planes as candidates for symmetry planes of an arbitrary set of points (point cloud). We use the point set representation because it is more general than a triangle mesh and, therefore, makes the information in this section useful in wider range of applications.

Suppose a set of points $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$, $\mathbf{x}_i \in E^3, i = 1, 2, \dots, N$. We first create a 3D uniform grid with the cell size $\frac{l_{avg}}{\delta} \times \frac{l_{avg}}{\delta} \times \frac{l_{avg}}{\delta}$ where l_{avg} is the estimated size of the object represented by X computed as the average distance of the points of X from their centroid. We mark each cell as either occupied if any point from X falls into it or unoccupied if no point from X falls into it. Then we start randomly selecting pairs of points from X and for each of these randomly selected pairs $\mathbf{x}_i, \mathbf{x}_j$ we create a plane P such that $\mathbf{r}_P(\mathbf{x}_i) = \mathbf{x}_j$.

Now, we need candidates for which there is at least some chance that they could capture some symmetry of the object, otherwise there would be too much clutter in the candidate space and no significant mode would be detectable. Therefore, we perform a quick check to determine whether P is a plausible candidate. This is done by randomly selecting another five points from X , reflecting them over P and checking whether all of them end up in an occupied cell of the previously created grid. If they do P is accepted as a candidate. If at least one of the five points reflects into an unoccupied cell then P is rejected. We keep iterating this process until we have k accepted candidates. The key idea behind the Mode-based approach is that now there should be significant modes in the candidate space of planes corresponding to the strongest symmetries of the input point set X .

In order to have a well detectable mode which also well represents the symmetry plane of the input point set, the parameters k and δ need to be set sufficiently large. However, setting them too large would create a very pronounced mode consisting of very good candidates but the time consumption of the candidate creation process would be overwhelming. If not stated

otherwise we set $\delta = 5$ and $k = 2000$ which, in our experiments, proved to be sufficiently large while leaving the speed of the candidate creation step acceptable.

5.1.2 Dependence on Scale and Position

Unlike the a, b, c coefficients, which are bounded on finite interval $\langle -1; 1 \rangle$, the value of the d coefficient of any candidate plane depends on the overall scale and position of the input object represented by the point set X . The dependence on scale is obvious because d represents the distance of the given candidate plane from the origin and if the size of the input object changes, the span of the d coefficient will change as well. However, the a, b, c coefficients will stay the same.

The dependence on position is less obvious. Imagine we translate the input object (all points in X) by some arbitrary vector \mathbf{t} , then for an arbitrary candidate plane P its d coefficient will change by $\mathbf{t}^T \mathbf{n}$ against the value it would have if it was created at the original position. Since the change of d does not depend only on \mathbf{t} but also on the orientation of the given plane, the change of d is inconsistent throughout the candidate planes and this inconsistency is the more significant the farther the input object gets from the origin. For example, suppose a trivial case where an object consists of a single point \mathbf{x} which lies at the origin and suppose a set of all planes passing through \mathbf{x} . The d coefficient of all these planes is 0. If we translate \mathbf{x} by \mathbf{t} , the d coefficients of the planes that pass through \mathbf{x} will now span from $-\|\mathbf{t}\|$ to $\|\mathbf{t}\|$. Therefore, the position of the input object, i.e. its distance from the origin, again influences the span of d but does not influence the span of a, b, c .

Many of the distance functions for planes, presented later in this text, are negatively influenced by the significantly different span of d and a, b, c . To mitigate this problem, before creating the candidate planes, we always translate the input object so that its centroid is at the origin and, where necessary, we also normalize d by l_{avg} to make the span of d similar to the span of a, b, c . For those distance functions where the translation to origin is not necessary, this fact will be pointed out explicitly.

5.1.3 Ground Truth

It was already mentioned that in symmetry detection, any candidate plane can be described as reflection transformation as defined by Eq. 5.1. As was pointed out in Chapter 4, distance between transformations cannot be

well defined without the context (the object on which the transformations are applied), which is consistent with what was described in Section 5.1.2. Therefore, the distance function for planes that we consider the most meaningful is the point-based version of the metric used for error evaluation of registration results in Chapter 4 (see Eq. 4.14), only with reflection transformations instead of rigid ones. Given two arbitrary planes P_1 and P_2 the distance function measures the exact difference between the effects of the reflections defined by P_1 and P_2 on the input object. Since such a function gives us the exact evaluation of how differently the input object is affected by the two reflections, we consider it the ground truth distance function, we denote it $D_{GT}(P_1, P_2)$ and it is defined as shown in Eq. 5.2 where $\mathbf{x}_i \in X, i = 1, \dots, N$.

$$D_{GT}(P_1, P_2) = \sum_{i=1}^N \|\mathbf{r}_{P_1}(\mathbf{x}_i) - \mathbf{r}_{P_2}(\mathbf{x}_i)\| \quad (5.2)$$

The D_{GT} distance function is not effected by the position of the input object, so it does not require the translation to origin, and the object size only effects its overall scale. Unfortunately, the time complexity of computing D_{GT} is $\mathcal{O}(N)$ where N is the point count of the input object, which makes it too computationally expensive and, therefore, virtually unusable in any Mode-based symmetry detection algorithm. However, we can compare it to the other distance functions described below and measure how close they get to it.

5.2 Plane Space Representations

In this section we describe various different ways of representing the space of planes in E^3 and for each we also describe possible distance functions that can be used in an arbitrary Mode-based symmetry plane detection algorithm. Furthermore, we use the algorithm described in Section 7.1.3 to create a set of candidate symmetry planes of the armadillo object shown in Figure 5.1 and we visualize them in various representations of the plane space. The black line in the figure represents the symmetry plane that we consider the correct one and the object is rotated in such a way that this plane is perpendicular to the plane in which the figure is rendered. We purposely selected an object that is not perfectly symmetrical but still exhibits noticeable reflectional symmetry. Although the object is represented by a triangle mesh, which makes it easier to visualize, only its vertices are used as the points for the candidate creation.




Figure 5.1: Model object with its correct symmetry plane.

5.2.1 Dual Representation in E^3

Although the implicit equation of a plane has four coefficients, there are actually only three degrees of freedom when defining a plane because the space of planes is a 3-dimensional manifold embedded in 4-dimensional space. Therefore, we can use a dual representation of an arbitrary plane as a point in E^3 . We denote $\rho(P) \in E^3$ a dual representation of a plane P . Euclidean metric could then be used to compute the distance between two planes P_1, P_2 as $D_\rho(P_1, P_2) = \|\rho(P_1) - \rho(P_2)\|$.

One possibility to represent a plane in E^3 is to encode its orientation into a vector in E^3 with the same direction as the plane normal vector, and the plane distance from origin into the length of this vector. Such dual representation can be defined as $\rho_1(P) = d\mathbf{n}$. Obviously, for $d \rightarrow 0$ such representation gets ambiguous because all planes with $d = 0$ are shrunk into a single point. To solve this problem the value of d can be shifted by a constant μ so that these planes get spread on the surface of a sphere with radius μ instead of being all at the origin. We set $\mu = \frac{1}{2}l_{avg}$ so that rotating the normal by π and changing d by l_{avg} make approximately similar change in position of the point in the dual space. The dual representation is therefore finally defined as

$$\rho_1(P) = \begin{cases} (d + \frac{1}{2}l_{avg})\mathbf{n} & d \geq 0 \\ (d - \frac{1}{2}l_{avg})\mathbf{n} & d < 0 \end{cases}.$$

Distances in such dual space still do not very well correspond to similarities/disimilarities of the actual planes. Mainly, two planes with d close to 0 and similar normal vectors can be on the other sides of the sphere, and therefore more than 2μ apart, although they are actually very similar. However, such representation can be very good for visualization because each point in the dual space represents the plane quite intuitively.




Figure 5.2: Dual representations ρ_1 , ρ_2 and ρ_3 of the candidate symmetry planes. The colors represent density (the darker, the larger density), the red spot corresponds to the correct symmetry plane.

Figure 5.2a shows the generated candidates on the armadillo model in the dual E^3 space transformed with ρ_1 . The colored spots represent the candidates, their absolute coordinates are not important, only their distances. The darker spots correspond to larger density of the points in the space, the red spot corresponds to the correct plane from Figure 5.1. The point of view for projecting the 3D visualization into the 2D image was selected manually to maximize the information in the figure. It can be seen that the correct plane is in a noticeable mode (dense spot) but this mode is split on the surface of the sphere that corresponds to $d = 0$ and its non-negligible part is on the other side. This is quite undesirable because it makes the mode much less significant than it would be if the two parts were together in the space.

Another duality, also called polar duality (described e.g. in [35]), uses normalization of the plane coefficients such that $d = 1$ and then only uses the a, b, c coefficients as coordinates in E^3 . So instead of multiplying the normal vector by d , this time we divide it by d . This again poses a problem for $d \rightarrow 0$ which makes the dual points approach infinity and planes with $d = 0$ cannot be represented at all by this duality. We solve this issue in the same way as with ρ_1 by shifting the d coefficient and we define this dual representation as

$$\rho_2(P) = \begin{cases} \frac{1}{(d + \frac{1}{2}l_{avg})} \mathbf{n} & d \geq 0 \\ \frac{1}{(d - \frac{1}{2}l_{avg})} \mathbf{n} & d < 0 \end{cases}.$$

Figure 5.2b shows the candidates transformed by ρ_2 into the dual E^3 space. The correct plane is located in a noticeable mode which is again split into

two separate parts that are very far from each other. In this case there are also other significant modes that correspond to very different planes.

Another duality commonly used in computational geometry expresses a plane using its coefficients in explicit representation [9]. There are three possible explicit representations of a plane in E^3 :

$$\begin{aligned} x &= -\frac{b}{a}y - \frac{c}{a}z - \frac{d}{a}, \\ y &= -\frac{a}{b}x - \frac{c}{b}z - \frac{d}{b}, \\ z &= -\frac{a}{c}x - \frac{b}{c}y - \frac{d}{c}. \end{aligned}$$

For demonstration, we select the first one, the dual representation is then defined as $\rho_3(P) = [\frac{b}{a}, \frac{c}{a}, \frac{d}{l_{avg} \cdot a}]$. The division of d by l_{avg} is necessary for normalizing the span of d . Such duality obviously cannot represent planes parallel to the x -axis and planes with $a \rightarrow 0$ approach infinity in the dual space. We could possibly solve this by shifting a but this time, we do not include l_{avg} into the shift because the span of a does not depend on the size of the input object, so we get

$$\rho_3(P) = \begin{cases} [\frac{b}{a+\frac{1}{2}}, \frac{c}{a+\frac{1}{2}}, \frac{d}{l_{avg}(a+\frac{1}{2})}] & a \geq 0 \\ [\frac{b}{a-\frac{1}{2}}, \frac{c}{a-\frac{1}{2}}, \frac{d}{l_{avg}(a-\frac{1}{2})}] & a < 0 \end{cases}.$$

Figure 5.2c shows the candidates in the dual space transformed by ρ_3 and in this case there do not seem to be any significant modes.

In general, the dual representations appear not to be very appropriate for representing planes in any Mode-based symmetry detection algorithm because they all contain singularities. Although this problem can always be solved by shifting the value of some coefficient by a constant, the choice of this constant is rather arbitrary and even then the distances between points in the dual space might not well correspond to similarities of the planes. However, the dual representations can quite easily be used for visualizing the candidates because the dual points are 3-dimensional.

5.2.2 4D Vector Representation

Probably the most intuitive way of representing a plane is by a 4D vector of the plane coefficients. Given a plane P we represent it by a vector $\mathbf{p} = [a, b, c, \frac{d}{l_{avg}}]^T$. In such a space we can easily define a distance function as the Euclidean distance of the two 4D vectors. However, \mathbf{p} and $-\mathbf{p}$ represent the

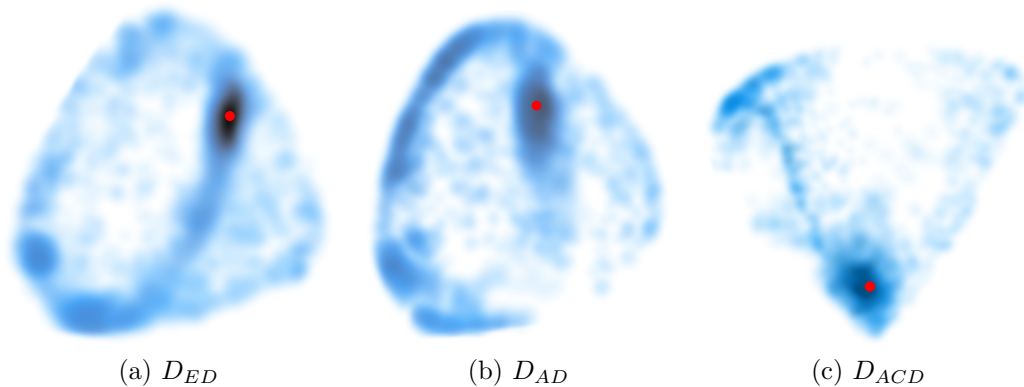


Figure 5.3: The candidates represented by 4D vectors projected into E^3 with MDS using different distance functions.

same plane so we need to take this into account. The Euclidean distance function is therefore defined as

$$D_{ED}(P_1, P_2) = \begin{cases} \|\mathbf{p}_1 - \mathbf{p}_2\| & \mathbf{p}_1^T \mathbf{p}_2 \geq 0 \\ \|\mathbf{p}_1 + \mathbf{p}_2\| & \mathbf{p}_1^T \mathbf{p}_2 < 0 \end{cases}.$$

In this case the points cannot be visualized directly, so we use the multidimensional scaling (MDS) technique to transform the points into E^3 so that they maintain their distances, w.r.t. the given distance function, as well as possible. However, the projection into E^3 might cause some imprecision in the visualization. Figure 5.3a shows the candidate planes projected into E^3 with MDS using the D_{ED} distance function and there is a very significant mode visible around the correct symmetry plane.

The distances in 4D vector space of planes can also be measured as angles between the vectors because the length of the vector \mathbf{p} does not influence the plane P it represents. The angle distance function can be defined as

$$D_{AD}(P_1, P_2) = \arccos \left(\frac{|\mathbf{p}_1^T \mathbf{p}_2|}{\|\mathbf{p}_1\| \|\mathbf{p}_2\|} \right).$$

Figure 5.3b shows the candidates after using MDS with the D_{AD} distance function and the correct plane is again placed inside a noticeable mode.

We can also use only the cosine of the angle and measure its deviation from 1. The angle cosine distance function can be defined as

$$D_{ACD}(P_1, P_2) = 1 - \frac{|\mathbf{p}_1^T \mathbf{p}_2|}{\|\mathbf{p}_1\| \|\mathbf{p}_2\|}$$

and its visualization using MDS is shown in Figure 5.3c. There is again a noticeable mode around the correct plane.

Obviously, the 4D representation of the plane space is much more appropriate for any Mode-based symmetry detection algorithm than the dual representations in E^3 . However, they are not as convenient for visualization because in order to show the points they first need to be projected into a lower dimensional Euclidean space which causes a loss of information.

5.2.3 Transformation Representation

As already mentioned, the space of planes can be understood as the space of reflection transformations and, therefore, the distance between arbitrary two planes P_1 and P_2 can be defined as the distance between the two reflection transformations \mathbf{r}_{P_1} and \mathbf{r}_{P_2} defined according to Eq. 5.1. One way of doing this is using the compound metric that was evaluated as the most suitable for rigid transformations in Chapter 4. Since in this context we expect a point cloud on the input, not a triangle mesh, we use its point-based version (see eq. 4.9). It is based on sum of squared distances between the transformed points and for reflection transformations is defined as

$$D_{SSD}(P_1, P_2) = \sqrt{\sum_{i=1}^N \|\mathbf{r}_{P_1}(\mathbf{x}_i) - \mathbf{r}_{P_2}(\mathbf{x}_i)\|^2}$$

where $\mathbf{x}_i \in X, i = 1, \dots, N$. There is a notable similarity between D_{SSD} and the ground truth distance function D_{GT} (see Eq. 5.2). However, as already described in Section 4.2.2 for rigid transformations, there are two major differences. First, D_{SSD} uses squared distances instead of absolute ones, favouring smaller displacements over larger ones, which leads to different distances. Second, unlike D_{GT} , D_{SSD} can be computed in $\mathcal{O}(1)$, given an $\mathcal{O}(N)$ preprocessing is performed beforehand. To achieve this, the transformations must be expressed as $\mathbf{M}\mathbf{x} + \mathbf{t}$ where \mathbf{M} is an orthogonal transformation matrix, \mathbf{t} is an arbitrary translation vector and \mathbf{x} is the transformed point. For the details of deriving the $\mathcal{O}(1)$ computation we refer to Section 4.2.2. Directly from Eq. 5.1 we get

$$\mathbf{r}_P(\mathbf{x}) = \mathbf{x} - 2\mathbf{nn}^T\mathbf{x} - 2d\mathbf{n} = (\mathbf{I} - 2\mathbf{nn}^T)\mathbf{x} - 2d\mathbf{n}$$

where \mathbf{I} is the identity matrix. If we now denote $\mathbf{M} = (\mathbf{I} - 2\mathbf{nn}^T)$ and $\mathbf{t} = -2d\mathbf{n}$, then the reflection transformation can be expressed as

$$\mathbf{r}_P(\mathbf{x}) = \mathbf{M}\mathbf{x} + \mathbf{t}.$$

Since the matrix \mathbf{M} is orthogonal (and also symmetric), we can use the same derivation as in Section 4.2.2 to compute D_{SSD} in $\mathcal{O}(1)$ with $\mathcal{O}(N)$

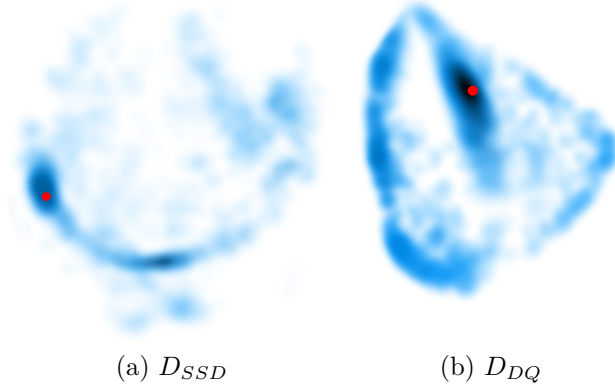


Figure 5.4: The candidates represented as transformations projected into E^3 with MDS using different distance functions.

preprocessing. The D_{SSD} distance function, as well as D_{GT} , is not effected by the position of the input object, so the initial translation to the origin is not required, and the object size only effects the overall scale of the distance function.

Figure 5.4a shows the candidates projected into E^3 using MDS with the D_{SSD} distance function and the correct plane is again in a significant mode. There is another smaller significant mode visible in the figure, however, this can very likely be caused by the distortion of the MDS projection.

Dual Quaternions

Dual quaternions combine the concepts of quaternions and dual numbers and are commonly used, mainly in robotics, to represent rigid transformations. However, in the following text we show how they can also be used to represent a transformation of reflection over an arbitrary plane. A general quaternion is defined as $Q = q_0 + q_1i + q_2j + q_3k$ where the i, j, k units multiply according to the following rules

$$i^2 = j^2 = k^2 = ijk = -1, ij = k = -ji, jk = i = -kj, ki = j = -ik.$$

Quaternion multiplication is not commutative in general. A conjugate Q^* of a quaternion Q is defined as $Q^* = q_0 - q_1i - q_2j - q_3k$. If $q_0 = 0$ then Q is so called pure quaternion in which case it is that $Q^* = -Q$. For two arbitrary quaternions Q_1 and Q_2 it is that $(Q_1 + Q_2)^* = Q_1^* + Q_2^*$ and $(Q_1Q_2)^* = Q_2^*Q_1^*$. A size or norm of a general quaternion Q is defined as $\|Q\| = \sqrt{QQ^*}$ which is always a non-negative real number. We denote $v(\mathbf{x}) = xi + yj + zk$ a quaternion that represents an arbitrary point $\mathbf{x} = [x, y, z]^T \in E^3$. If \mathbf{u} is an arbitrary unit vector and we set $Q = \cos\alpha + v(\mathbf{u})\sin\alpha$, then $Qv(\mathbf{x})Q^*$

represents the point \mathbf{x} rotated by angle 2α around the axis that passes through the origin and has the direction of \mathbf{u} . Also, for any unit vector \mathbf{u} and any α , Q is a unit quaternion, i.e. $QQ^* = Q^*Q = 1$. Similarly, if we set $Q = v(\mathbf{u})$ (as a special case of $\alpha = \frac{\pi}{2}$) then $Qv(\mathbf{x})Q$ represents the point \mathbf{x} reflected over the plane with normal \mathbf{u} that passes through the origin. Notice, that this time Q at the right side is not conjugated. For details about quaternions we refer to [38].

A dual quaternion is defined as

$$Q_d = Q + \epsilon Q_\epsilon = q_0 + q_1i + q_2j + q_3k + \epsilon(q_{\epsilon 0} + q_{\epsilon 1}i + q_{\epsilon 2}j + q_{\epsilon 3}k)$$

where Q and Q_ϵ are quaternions and ϵ is the dual unit which commutes with the quaternion units i, j, k and it is that $\epsilon^2 = 0$. A quaternion conjugate of Q_d is defined as $Q_d^* = Q^* + \epsilon Q_\epsilon^*$, a dual conjugate of Q_d is defined as $\overline{Q_d} = Q - \epsilon Q_\epsilon$. These conjugations can be combined into $\overline{Q_d^*} = Q^* - \epsilon Q_\epsilon^*$.

We denote $v_d(\mathbf{x}) = 1 + \epsilon v(\mathbf{x}) = 1 + \epsilon(xi + yj + zk)$ a dual quaternion that represents an arbitrary point $\mathbf{x} = [x, y, z]^T \in E^3$. If Q is a quaternion that represents rotation and $Q_\epsilon = \frac{v(\mathbf{t})Q}{2}$ where $\mathbf{t} = [t_x, t_y, t_z]^T$ is an arbitrary translation vector then for $Q_d = Q + \epsilon Q_\epsilon$, it can be shown that the expression $Q_d v_d(\mathbf{x}) \overline{Q_d^*}$ represents a rigid transformation in the following way. Since $\overline{Q_d^*} = Q^* - \epsilon \frac{(v(\mathbf{t})Q)^*}{2}$ and $(v(\mathbf{t})Q)^* = Q^* v(\mathbf{t})^*$ the expression can be expanded as

$$Q_d v_d(\mathbf{x}) \overline{Q_d^*} = (Q + \epsilon \frac{v(\mathbf{t})Q}{2})(1 + \epsilon v(\mathbf{x}))(Q^* - \epsilon \frac{Q^* v(\mathbf{t})^*}{2}).$$

By multiplying the brackets while respecting that $\epsilon^2 = 0$ we get

$$Q_d v_d(\mathbf{x}) \overline{Q_d^*} = QQ^* - \epsilon \frac{QQ^* v(\mathbf{t})^*}{2} + \epsilon Qv(\mathbf{x})Q^* + \epsilon \frac{v(\mathbf{t})QQ^*}{2}$$

and since $QQ^* = 1$ we further get

$$Q_d v_d(\mathbf{x}) \overline{Q_d^*} = 1 - \epsilon \frac{v(\mathbf{t})^*}{2} + \epsilon Qv(\mathbf{x})Q^* + \epsilon \frac{v(\mathbf{t})}{2}.$$

Because $v(\mathbf{t})$ is a pure quaternion and therefore $v(\mathbf{t})^* = -v(\mathbf{t})$, the expression finally yields the form

$$Q_d v_d(\mathbf{x}) \overline{Q_d^*} = 1 + \epsilon Qv(\mathbf{x})Q^* + \epsilon v(\mathbf{t}) = 1 + \epsilon(Qv(\mathbf{x})Q^* + v(\mathbf{t}))$$

which represents the point \mathbf{x} rotated using the quaternion Q and then translated by \mathbf{t} . This shows how dual quaternions can be used for representing and computing rigid transformations. Note that Q_d represents the same transformation as $-Q_d$ with the identity being represented by either 1 or

-1. Also, the transformations can be concatenated by multiplying the corresponding dual quaternions and if Q_d represents a rigid transformation then Q_d^* represents its inverse. This in turn means that given two dual quaternions Q_{d1}, Q_{d2} representing rigid transformations, these transformations are the same only if $Q_{d1}Q_{d2}^* = 1$ or $Q_{d1}Q_{d2}^* = -1$. For details about dual quaternions we refer to [100] or [94].

Consider now a plane P and a dual quaternion $Q_d = Q + \epsilon Q_\epsilon$ defined such that $Q = v(\mathbf{n})$ and $Q_\epsilon = \frac{v(\mathbf{t})Q}{2}$ where $\mathbf{t} = -2d\mathbf{n}$. Now Q_d represents a transformation that first rotates by π around the axis that passes through the origin and has the direction of \mathbf{n} , and then translates by $-2d\mathbf{n}$. However, if we apply the transformation on $-\mathbf{x}$ instead of \mathbf{x} , it can be easily shown that the transformation expression expands as

$$Q_d v_d(-\mathbf{x}) \overline{Q_d^*} = 1 - \epsilon Q v(\mathbf{x}) Q^* + \epsilon v(\mathbf{t}).$$

Because $Q = v(\mathbf{n})$ is a pure quaternion, hence $Q^* = -Q$, we can adjust the expression as

$$\begin{aligned} Q_d v_d(-\mathbf{x}) \overline{Q_d^*} &= 1 + \epsilon Q v(\mathbf{x}) Q + \epsilon v(\mathbf{t}) = 1 + \epsilon (Q v(\mathbf{x}) Q + v(\mathbf{t})) = \\ &= 1 + \epsilon (v(\mathbf{n}) v(\mathbf{x}) v(\mathbf{n}) - v(2d\mathbf{n})) = v_d(\mathbf{r}_P(\mathbf{x})) \end{aligned}$$

which exactly represents $\mathbf{r}_P(\mathbf{x})$. This shows that a dual quaternion can also represent a reflection transformation by representing a rigid transformation that transforms $-\mathbf{x}$ to $\mathbf{r}_P(\mathbf{x})$. Therefore, to measure distances between reflection transformations we can use a distance function for dual quaternions.

We denote $vec(Q_d) = [q_0, q_1, q_2, q_3, q_{\epsilon 0}, q_{\epsilon 1}, q_{\epsilon 2}, q_{\epsilon 3}]^T \in E^8$ an 8-dimensional vector that is equivalent to Q_d . Given a plane P , we create the corresponding dual quaternion Q_d such that $Q = v(\mathbf{n})$ and $Q_\epsilon = \frac{v(-2d\mathbf{n})Q}{2l_{avg}}$, i.e. $Q_d = v(\mathbf{n}) + \epsilon \frac{v(-2d\mathbf{n})v(\mathbf{n})}{2l_{avg}}$. The division by l_{avg} is again to normalize the d coefficient (which of course cannot be done if the dual quaternion is intended to be used for computing the reflection transformation). Since $v(-2d\mathbf{n}) = -2d v(\mathbf{n})$ and because $v(\mathbf{n})$ is a pure unit quaternion we can get that

$$Q_\epsilon = \frac{-2d v(\mathbf{n})v(\mathbf{n})}{2l_{avg}} = \frac{d v(\mathbf{n})^* v(\mathbf{n})}{l_{avg}} = \frac{d}{l_{avg}}$$

so Q_d can be finally expressed as shown in Eq. 5.3.

$$Q_d = v(\mathbf{n}) + \epsilon \frac{d}{l_{avg}} = ai + bj + ck + \epsilon \frac{d}{l_{avg}} \quad (5.3)$$

Notice that the transition from the implicit equation of a plane to the corresponding dual quaternion is very straightforward because the a, b, c, d

coefficients of the plane exactly correspond to four of the eight values of the dual quaternion and the remaining four values are 0.

There are two common distance functions for dual quaternions. The first one uses differences between the equivalent 8-dimensional vectors [94]. Suppose two arbitrary planes P_1 and P_2 represented by dual quaternions Q_{d1} and Q_{d2} respectively. Such distance function can be defined as

$$\min\{\|vec(Q_{d1}) - vec(Q_{d2})\|, \|vec(Q_{d1}) + vec(Q_{d2})\|\}$$

but given Eq. 5.3 this is exactly the same as D_{ED} . The second distance function [33] uses a difference transformation $Q_{d1}Q_{d2}^*$ and computes its distance from the identity, i.e. from 1 or -1 . It is defined as

$$D_{DQ}(P_1, P_2) = \min\{\|vec(1 - Q_{d1}Q_{d2}^*)\|, \|vec(1 + Q_{d1}Q_{d2}^*)\|\}.$$

Figure 5.4b shows the candidates projected into E^3 using MDS with D_{DQ} and the correct plane is in an obvious mode.

5.3 Results

We compared the distance functions by generating the candidate symmetry planes of a given object (using the model algorithm described in Section 7.1.3), computing distances between them using the given distance function and comparing them to distances computed using the ground truth distance function. We did this for the six different test objects shown in Figure 5.5, taken from datasets [32, 67]. The objects are represented by triangle meshes for easier visualization, but we again only used their vertices as the input points for the candidate creation process. The Armadillo and Bunny objects are simplified because the computation of D_{GT} on their original version would be too timely for the experiments we performed.

Let $C = \{P_1, P_2, \dots, P_k\}$, $k = 2000$ be the set of candidate planes created for a given input object. The error of a given distance function D against the ground truth is defined as

$$Err(D) = \frac{1}{Count(k)} \sum_{i=1}^k \sum_{j=i+1}^k \left| \frac{D_{GT}(P_i, P_j)}{Avrg(D_{GT})} - \frac{D(P_i, P_j)}{Avrg(D)} \right|$$

where

$$Avrg(D) = \frac{1}{Count(k)} \sum_{i=1}^k \sum_{j=i+1}^k D(P_i, P_j)$$

is the average distance between candidates in C and $Count(k) = \frac{1}{2}(k^2 - k)$ is the total number of candidate pairs used for the computation. The

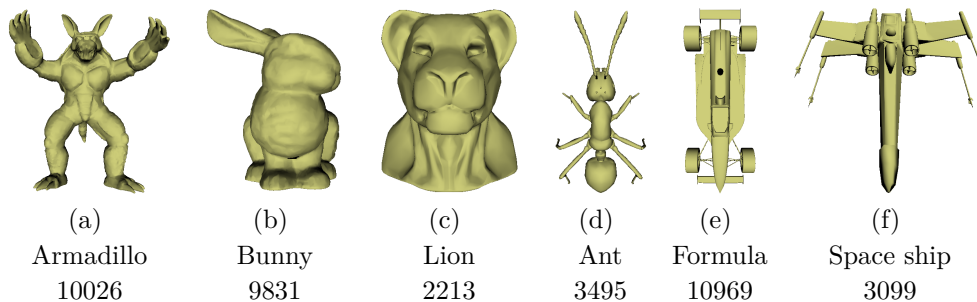


Figure 5.5: The test objects used to generate the candidate sets for comparing the distance functions. The number under the name of each object expresses its point count.

Table 5.1: Errors of the distance functions for the candidate sets for different objects.

| | Arm | Bun | Ant | For | Lio | Shi | Average |
|--------------|-------|-------|-------|-------|-------|-------|---------|
| D_{ED} | 0.120 | 0.277 | 0.163 | 0.093 | 0.130 | 0.234 | 0.169 |
| D_{AD} | 0.133 | 0.281 | 0.157 | 0.098 | 0.144 | 0.236 | 0.174 |
| D_{ACD} | 0.299 | 0.388 | 0.264 | 0.250 | 0.306 | 0.352 | 0.309 |
| D_{SSD} | 0.012 | 0.023 | 0.009 | 0.014 | 0.011 | 0.012 | 0.013 |
| D_{DQ} | 0.118 | 0.277 | 0.162 | 0.093 | 0.129 | 0.232 | 0.168 |
| D_{ρ_1} | 0.382 | 0.399 | 0.503 | 0.596 | 0.326 | 0.425 | 0.438 |
| D_{ρ_2} | 0.401 | 0.408 | 0.488 | 0.563 | 0.360 | 0.489 | 0.451 |
| D_{ρ_3} | 0.280 | 0.446 | 0.269 | 0.730 | 0.362 | 0.447 | 0.422 |

normalization by Avg is used because the overall scales of the distance functions do not matter so the differences are computed after both D_{GT} and D are divided by their mean values.

Table 5.1 shows the errors of all the distance functions described above for all the test objects. For completeness, we include the dual representations in the comparison.

The smallest error is obviously achieved using D_{SSD} which is probably due to D_{SSD} and D_{GT} being based on the same principle. However, it is still rather surprising that the D_{SSD} function which uses squared distances is so similar to D_{GT} that uses absolute distances. The D_{ED} , D_{AD} and D_{QD} all exhibit very similar errors (with D_{QD} usually having the lowest of these three) which are overall lower than those of D_{ACD} and the distances in the dual spaces, but in case of D_{ACD} this can be explained by its resemblance to the cosine function ($D_{ACD} = 1 - \cos(D_{AD})$). The function D_{ρ_3} exhibits

similar or lower error than D_{ACD} on some objects (Arm, Ant) but also considerably larger error on different ones (For, Shi) which suggests that D_{ρ_3} is quite unpredictable.

The graphs in Figures 5.6, 5.7, 5.8 show the relation between D_{GT} and the other distance functions. We generated 50 candidates on the Armadillo object and for each pair of the candidates we put its distance computed using D_{GT} on the horizontal axis and the distance computed using a given different distance function on the vertical axis. We divide each value by the mean of the given distance function computed by *Avg*. If some distance function D was exactly the same as D_{GT} (apart from overall scale) then there would be a perfect linear dependency and the points for D would lie on a perfect line in the graph.

Figure 5.6 shows the relations of D_{SSD} , D_{ED} and D_{ACD} to D_{GT} . The functions D_{AD} and D_{DQ} are not included in this figure because they are too similar to D_{ED} , this similarity is shown separately in Figure 5.7. The dual representations are also shown separately in Figure 5.8. For different objects the graphs are slightly different but overall very similar.

There is an obvious almost linear dependency between D_{GT} and D_{SSD} (see Fig. 5.6), however, D_{ED} , D_{AD} and D_{DQ} exhibit relation to D_{GT} that is also quite near linear dependency (see Fig. 5.7). For D_{ACD} the resemblance to cosine is visible in the graph. On the other hand, the dual representations show rather unstable behavior (see Fig. 5.8). This is mostly caused by the shift in some of the coordinates in the dual space but without this shift all the dual representations would suffer from singularities which is even worse and makes them virtually unusable. It is possible that shifting by a different constant could lead to better results, at least for some objects, however, the choice of the shifting constant is arbitrary and there does not seem to be any reasonable way to set it appropriately.

Table 5.2 shows the Pearson correlations [86] between all pairs of the distance functions (including D_{GT}) for the data shown in Figures 5.6, 5.7, 5.8. Value of 1 indicates perfect linear dependency and the closer the value is to 0 the weaker the linear dependency is. Expectedly, D_{SSD} shows the best linear correlation with D_{GT} . However, the correlations of D_{ED} , D_{AD} , D_{DQ} and even D_{ACD} with D_{GT} are all rather high, all above 0.9. On the other hand, the distances in the dual spaces exhibit mostly low correlation with D_{GT} . Notably, the correlations among D_{ED} , D_{AD} and D_{DQ} are all very high, confirming that these three distance functions are indeed all very similar.




Figure 5.6: Relations between $D_{SSD}/D_{ED}/D_{ACD}$ and D_{GT} for the Armadillo object.




Figure 5.7: Relations between $D_{DQ}/D_{ED}/D_{AD}$ and D_{GT} for the Armadillo object.




Figure 5.8: Relations between $D_{\rho_1}/D_{\rho_2}/D_{\rho_3}$ and D_{GT} for the Armadillo object.

Table 5.2: Pearson correlations of the distance functions for the Armadillo object.

| | GT | ED | AD | ACD | SSD | DQ | ρ_1 | ρ_2 | ρ_3 |
|----------|--------|--------|--------|--------|--------|--------|----------|----------|----------|
| GT | 1.0000 | 0.9723 | 0.9644 | 0.9120 | 0.9998 | 0.9738 | 0.5361 | 0.3265 | 0.7238 |
| ED | 0.9723 | 1.0000 | 0.9989 | 0.9679 | 0.9713 | 0.9998 | 0.5537 | 0.3460 | 0.7621 |
| AD | 0.9644 | 0.9989 | 1.0000 | 0.9767 | 0.9635 | 0.9983 | 0.5499 | 0.3474 | 0.7548 |
| ACD | 0.9120 | 0.9679 | 0.9767 | 1.0000 | 0.9105 | 0.9664 | 0.5196 | 0.3111 | 0.7262 |
| SSD | 0.9998 | 0.9713 | 0.9635 | 0.9105 | 1.0000 | 0.9728 | 0.5351 | 0.3286 | 0.7214 |
| DQ | 0.9738 | 0.9998 | 0.9983 | 0.9664 | 0.9728 | 1.0000 | 0.5516 | 0.3423 | 0.7637 |
| ρ_1 | 0.5361 | 0.5537 | 0.5499 | 0.5196 | 0.5351 | 0.5516 | 1.0000 | 0.9248 | 0.4217 |
| ρ_2 | 0.3265 | 0.3460 | 0.3474 | 0.3111 | 0.3286 | 0.3423 | 0.9248 | 1.0000 | 0.1934 |
| ρ_3 | 0.7238 | 0.7621 | 0.7548 | 0.7262 | 0.7214 | 0.7637 | 0.4217 | 0.1934 | 1.0000 |

5.3.1 Theoretical Comparison

Based on the results, the most appropriate representation of the space of planes in any Mode-based symmetry detection method is the transformation representation with the D_{SSD} distance function. But the results also suggest that, except for the dual representations, all the distance functions are rather similar and none of them deviates significantly from D_{GT} which makes all of them well applicable. One only has to keep in mind that all the distance functions except D_{SSD} require translating the input object to the origin, otherwise the normalization of the d coefficient would have to be done differently. But if the practical results are put aside, there are yet some theoretical differences between the various representations.

Since the dual and the 4D vector representations are basically Euclidean, using these representations the candidates can easily be stored in some Euclidean data structure, such as a k D-tree or a grid, that can be used for fast proximity queries if needed. In case of the D_{AD} and D_{ACD} distance functions some structure can possibly be built using the polar coordinates in 4D. Also, there are quite many possible algorithms for mode-seeking in Euclidean data. The transformation representations and the D_{SSD} and D_{DQ} distance functions are non-Euclidean and therefore slightly more restrictive in terms of the accelerating data structures and possible mode seeking algorithms that can be used. However, as already pointed out, such possibilities exist [120, 126] (see also Section 4.1.3). Also, the implementation of the D_{SSD} and D_{DQ} is more complex since D_{SSD} requires some matrix representation and basic operations of matrix algebra and D_{DQ} requires implementing the operations of dual quaternion algebra. On the other hand, the dual and 4D vector representations only require the ordinary algebraic operations with real numbers and vectors.

We also note that, although the D_{DQ} distance function does not seem to bring any considerable improvement over the other simpler distance func-

tions, the idea of representing reflections by dual quaternions seems novel and can possibly find its use in different applications or if some better distance functions for dual quaternions occur in the future. It could also be used when creating new symmetry detection algorithms to represent the planes, possibly taking advantage of the dual quaternion algebra.

5.4 Summary

We have described several representations of the space of planes that can be used in any Mode-based algorithm for symmetry plane detection and we have also described how distances can be computed in the various space representations. We have shown that the 3-dimensional dual space representations are not very appropriate for this purpose but they can easily be used for visualization purposes. In order to represent the space of planes appropriately in the Mode-based symmetry detection, spaces of higher dimensionality need to be used and the transformation representation, which appears to be the most appropriate one, is even non-Euclidean. However, the results suggest that, apart from the 3D dual spaces, all the plane space representations are well applicable in the context of Mode-based symmetry detection, although, as described above, there are some theoretical differences between them that need to be taken into account.

6 Symmetry Plane Detection Using Differentiable Symmetry Measure

In the previous chapter we analyzed different ways of representing planes in methods for symmetry plane detection that are based on the popular Mode-based approach. In this chapter, we describe a new method for detecting the global symmetry plane of a 3D object which is based on a completely different approach. The content of this chapter was previously published in [53].

Since reflectional symmetry is probably the most often occurring symmetry in real world objects, having a reliable and robust method for symmetry plane detection in 3D data can be very useful for many applications. In general, it is desirable for the method to be capable of detecting symmetries in objects that are only approximately symmetrical (see e.g. Fig. 6.1a) or symmetries in objects damaged by noise (see Fig. 6.1b). Some applications, such as object reconstruction [105, 112, 117], even require the symmetry detection to work on objects where some parts are missing, which strongly disrupts the symmetry and makes it much more difficult to detect. See, for example, the objects in Figures 6.1c and 6.1d. Although the global reflectional symmetry in them is quite weak, a human observer is still able to see it and a good symmetry detection method should be able to find it.

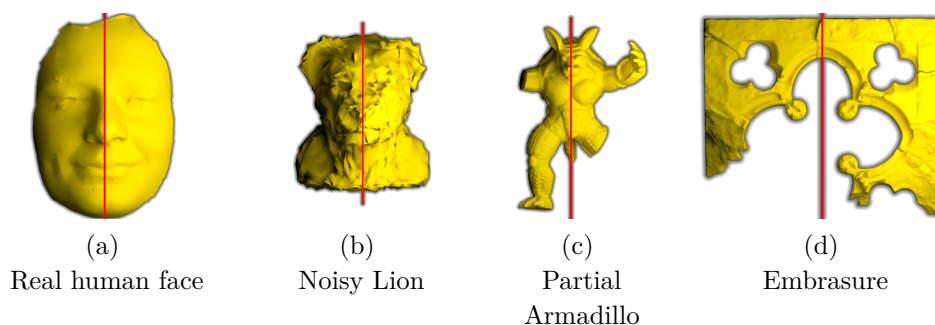


Figure 6.1: An approximately symmetrical object [18] - (a) a noisy object - (b) and two objects with missing parts - (c), (d).

In the following, we propose a surprisingly simple, differentiable symmetry measure, usable for evaluating symmetry of a 3D object, and we use

it to design a robust and flexible global symmetry plane detection method. The proposed method can be used on perfectly as well as approximately symmetrical objects and we show that it is also capable of detecting the plane of symmetry for objects with extensive missing parts, as well as for noisy objects. Furthermore, the proposed method takes a general discrete set of points on the input and therefore puts virtually no constraints on the input data, which can be very useful because a more advanced object representation, such as a manifold mesh, is not always available.

Many previous approaches use some form of descriptors (e.g. [16, 83, 107, 112]), mostly curvature-based, i.e. the quality of their results depends on the quality of the descriptor, which in turn depends on the properties of the input object. The proposed method works very well without descriptors, yet it is also very flexible, so that if more information about the input object is available, the method can be easily extended to use this information in its favor. Therefore, any descriptor can be used as the additional information for the method, as will be shown, but in general, this is not necessary. This altogether makes the proposed method very robust, flexible and widely applicable. We will also show that the method is mostly superior to other existing methods in terms of robustness, accuracy and speed. Although we only consider a 3D case, the proposed symmetry measure and the method built on it can easily be extended into more dimensions (or fewer, if needed) and the symmetry measure can also be generalized for different types of symmetry.

6.1 Symmetry Measure

We again represent a general plane P as described in Section 2.4.3 and we denote $\mathbf{p} = [a, b, c, d]^T$ a 4D vector of the plane coefficients. In this chapter we do not implicitly consider the plane coefficients to be normalized in any way. A vector function $\mathbf{r}(\mathbf{p}, \mathbf{x}) \in E^3$ that reflects a point $\mathbf{x} = [x, y, z]^T \in E^3$ over a plane P represented by \mathbf{p} is defined as

$$\mathbf{r}(\mathbf{p}, \mathbf{x}) = \mathbf{x} - 2 \frac{\mathbf{n}_{\mathbf{p}}^T \mathbf{x} + d}{\mathbf{n}_{\mathbf{p}}^T \mathbf{n}_{\mathbf{p}}} \mathbf{n}_{\mathbf{p}}$$

where $\mathbf{n}_{\mathbf{p}} = [a, b, c]^T$ is the normal vector of the plane P . The components of the function $\mathbf{r}(\mathbf{p}, \mathbf{x})$ are continuous and differentiable w.r.t. \mathbf{p} except for $\mathbf{p} = [0, 0, 0, d]^T$, which does not represent a valid plane. Consider a set of points $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$, $\mathbf{x}_i \in E^3$, $i = 1, \dots, n$ which represents a sampled 3D object. In theory, the points in X can be samples acquired from arbitrary

3D object representations, including volumetric ones, and all that follows is applicable regardless of the original representation from which the point set was extracted. However, we only experimented with surface data so in this work we consider X to represent a sampled 3D surface, e.g. vertices of a polygonal mesh, samples of a parametric surface, points of a raw point cloud, etc. We propose a symmetry measure that gives an evaluation of how much the point set X is symmetrical with respect to a given plane P , represented by \mathbf{p} . The measure is defined as follows

$$s_X(\mathbf{p}) = \sum_{i=1}^n \sum_{j=1}^n w_{ij} \varphi(\|\mathbf{r}(\mathbf{p}, \mathbf{x}_i) - \mathbf{x}_j\|). \quad (6.1)$$

The function $\varphi(l)$ is some radial function such that $\varphi(0) = 1$ and its value decreases and approaches 0 as l increases, w_{ij} are weights of point pairs and will be discussed below. By default, the weights are not used, i.e. $w_{ij} = 1$ for all i, j . The function φ is called a similarity function because it transforms distance of two points into their similarity. The symmetry measure s_X considers all possible pairs of points in X . For each pair $\mathbf{x}_i, \mathbf{x}_j \in X$ the point \mathbf{x}_i is reflected over the plane P , represented by \mathbf{p} , and its distance from \mathbf{x}_j is computed and transformed into similarity using φ . These similarities are summed together for all pairs, giving the symmetry measure value. The idea behind the symmetry detection method proposed below is that maximizing the symmetry measure $s_X(\mathbf{p})$ for \mathbf{p} should force as many points as possible to reflect over P as close as possible to other points. This idea well fits in the nature of the task of global approximate symmetry detection, which is mainly to balance the size of the matching parts of the object (how many points reflect close to other points) and the precision of the match (how close to other points they reflect).

The idea of using some form of symmetry measure to find symmetries in shapes is not new (see e.g. [60, 96, 127]). However, the symmetry measure we propose provides several significant advantages that allow for designing a very fast, robust and flexible symmetry plane detection method. First of all, if a proper φ function is used, the measure is differentiable. This is a very useful property since it allows analytically computing its gradient and provides a natural way of quickly finding its maximum using some fast gradient-based optimization method. It can also be computed efficiently, while remaining differentiable, as described below. Furthermore, the measure maintains the relations between every two points, which allows using the weights w_{ij} to adjust the importance of given point pairs based on additional information, if some is available, as will be demonstrated later. Although using the weights is certainly not necessary in majority of cases, in some spe-

cific situations it can be useful and this option makes the symmetry measure very flexible. The measure is well suitable for approximate or weak symmetries and it can be computed for any set of points, not requiring a closed or manifold surface or any other specific object representation. It can also be easily extended into higher dimensions by just adding more coordinates to the vector \mathbf{p} and easily generalized for different symmetry types by replacing the \mathbf{r} function with a different transformation.

6.1.1 Similarity Function

The symmetry measure $s_X(\mathbf{p})$ is differentiable w.r.t. \mathbf{p} (except for $\mathbf{p} = [0, 0, 0, d]^T$) when $\varphi(l)$ is differentiable for $l \in \langle 0; \infty \rangle$ and $\frac{d}{dl}\varphi(0) = 0$. This holds, for example, for the Gaussian function and also for most of the Wendland's functions [123]. Although the Gaussian function is simple and easy to implement, for reasons that will be obvious soon, we rather used the following modified Wendland's function instead

$$\varphi(l) = \begin{cases} (1 - \frac{1}{2.6}\alpha l)^5(8(\frac{1}{2.6}\alpha l)^2 + 5\frac{1}{2.6}\alpha l + 1) & \alpha l \leq 2.6 \\ 0 & \alpha l > 2.6 \end{cases}.$$

The value α is the shape parameter of the function. The multiplier $\frac{1}{2.6}$ is our modification which ensures that the function has a similar shape and spread to the Gaussian ($e^{-(\alpha l)^2}$) for the same value of α . The main difference between the Gaussian and our Wendland's function is that the Wendland's function equals 0 for $\alpha l > 2.6$, i.e. it is locally supported. This means that the contribution of any point $\mathbf{x}_i \in X$ to the value of $s_X(\mathbf{p})$ is fully determined by the points of X that are not farther than $\frac{2.6}{\alpha}$ from $\mathbf{r}(\mathbf{p}, \mathbf{x}_i)$ which will prove very useful in Section 6.1.2.

If no two points of X are closer than $2\frac{2.6}{\alpha}$ then the contribution is determined by at most one point of X which is the closest to $\mathbf{r}(\mathbf{p}, \mathbf{x}_i)$ and the maximum value of this contribution is 1. Therefore, the maximum possible value of $s_X(\mathbf{p})$ is n and it can only occur in the case when each point of X reflects over P precisely to another point of X , which can only happen in the case of perfect symmetry. This implies that if we set $\alpha \geq 2\frac{2.6}{l_{min}}$, where l_{min} is the smallest distance between two points of X , and X is perfectly symmetrical, then the global maximum of $s_X(\mathbf{p})$ will be always in the plane of perfect symmetry.

However, with such a value of α , the function $s_X(\mathbf{p})$ would have many local maxima, especially if l_{min} was small, making it difficult to optimize. Also, in practice, objects are never perfectly symmetrical and we therefore

aim to detect approximate symmetries rather than perfect ones, so we instead set α according to the size of the input object as $\alpha = \frac{15}{l_{avg}}$, which in most cases makes the span of φ considerably larger and $s_X(\mathbf{p})$ smoother and easier to optimize. The value l_{avg} is the average distance of the points in X from their centroid and the value 15 was chosen as approximately optimal based on vast experiments. We also note that the proposed symmetry detection method is not very sensitive to this value, as will be shown in Section 6.4. Nevertheless, in case of finding perfect symmetry, setting $\alpha = 2\frac{2.6}{l_{min}}$ could still be used for some final refinement but we do not include this into our symmetry detection method.

6.1.2 Efficient Computation

A brute force computation of $s_X(\mathbf{p})$ has time complexity of $\mathcal{O}(n^2)$ but for many pairs $\mathbf{x}_i, \mathbf{x}_j \in X$ the similarity $\varphi(\|\mathbf{r}(\mathbf{p}, \mathbf{x}_i) - \mathbf{x}_j\|)$ is 0, so it only needs to be computed for pairs where $\|\mathbf{r}(\mathbf{p}, \mathbf{x}_i) - \mathbf{x}_j\| \leq \frac{2.6}{\alpha}$. We use a uniform grid with the cell size $\frac{2.6}{\alpha} \times \frac{2.6}{\alpha} \times \frac{2.6}{\alpha}$. During the computation of $s_X(\mathbf{p})$, after a point \mathbf{x}_i is reflected over the given plane and ends up in a cell C , only points in C and cells adjacent to C are used for the symmetry measure computation. This way, due to the locality of the Wendland’s function, $s_X(\mathbf{p})$ can be computed efficiently and remain first-order differentiable. Higher-order differentiability can easily be achieved using some higher-order differentiable Wendland’s function, but in our application we only need the first-order differentiability and the function we use is less computationally expensive than the higher-order differentiable ones.

In general, however, any other locally supported function that fulfills the above mentioned differentiability conditions can be used without affecting the efficient computation and differentiability of the symmetry measure provided its spread parameter and the grid cell size are set correspondingly. This also adds to the flexibility of the symmetry measure.

6.1.3 Simplification

Using the above described computation, the symmetry measure can well be used to quantify symmetry for a given plane even when the input object has quite a large number of points. But for the purpose of symmetry detection, where the symmetry measure must be evaluated repeatedly, it can still be too computationally expensive. However, in extensive experiments we observed that the symmetry measure well represents the symmetry of a point set even after the set gets simplified to a rather low number of points, given a proper

simplification method is used. One way to simplify a set of points is to downsample it randomly but this approach is insufficient, since it results in a point set that does not represent the shape of the original object very well. Instead, we use the following simplification algorithm, which is quite simple and very fast. A 3D grid is created for the input point set with the cell size $\frac{l_{avg}}{k} \times \frac{l_{avg}}{k} \times \frac{l_{avg}}{k}$ and each occupied cell gives one point of the simplified point set by averaging all points contained in the cell. It is desired to simplify the point set to approximately m points, so the simplification of the original point set is repeated several times with increasing value of k until the resulting point count reaches at least m , usually it gets slightly above m .

The left column of Figure 6.2 visualizes the symmetry measure for three objects simplified to approximately 1000 points using the above described algorithm. Each point $\mathbf{q} \in \langle -1; 1 \rangle \times \langle -1; 1 \rangle \times \langle -1; 1 \rangle$ in the 3D space corresponds to a single plane with normal vector $\mathbf{n}_p = \frac{1}{\|\mathbf{q}\|} \mathbf{q}$ and $d = l_{avg} \cdot (\|\mathbf{q}\| - \frac{1}{2})$. The symmetry measure increases as the color goes from purple to light blue. The $\frac{1}{2}$ shift in the d coefficient causes the empty sphere in the middle, we only consider $d \geq 0$. For each object, the viewpoint was chosen to maximize the perceived information. For visualization simplicity, the objects are represented by triangle meshes but only their vertices were used as the original point sets and input for the simplification. In Figures 6.2a and 6.2b it appears as if there are two significant maxima but this is caused by planes with normal \mathbf{n}_p being the same as planes with normal $-\mathbf{n}_p$ for $d = 0$ and, in these two cases, the plane corresponding to the global maximum passes very near the origin, so d is close to 0. This causes the light spot of the global maximum to be split into two spots on the opposite sides of the empty sphere whose surface corresponds to $d = 0$. In Figure 6.2c, there is only one spot of the global maximum because the plane it represents does not pass so near the origin. Next to each visualization, there is the corresponding object together with the plane that represents the global maximum of the symmetry measure. In all three cases the global maximum corresponds to a very good symmetry plane of the given object, even though it is simplified to only approximately 1000 points. The original point counts are in the captions.

6.1.4 Locating Maxima

To find a local maximum of the symmetry measure, we employ a quasi-Newton optimization method L-BFGS [74], which uses the gradient of the symmetry measure, exploiting its differentiability, and it usually converges,




Figure 6.2: The symmetry measure visualized for Lion - (a), Armadillo [67] - (b) and its partial version - (c), simplified to approximately 1000 points (original point counts are in the captions), the symmetry measure increases as the color goes from purple to light blue, the objects are shown with the plane in the global maximum of the measure, the right column shows the convergence regions for the given objects.

to a sufficient precision, in only several iterations. Computation of the gradient can be accelerated in the same way as the computation of the symmetry measure itself (see Section 6.1.2). This is because the first-order derivative of $\varphi(l)$ also equals 0 for $\alpha l > 2.6$. We also tried using the Nelder-Mead optimization method [90], which does not require the gradient. It also worked but needed a much larger number of iterations making the optimization noticeably slower, which suggests the usefulness of the analytically computed gradient. Furthermore, differentiability can be useful even in case of non-gradient methods, because it makes the symmetry measure free of sudden changes, providing more stability to the optimization.

Before starting the L-BFGS optimization, the point set should be translated together with the initial plane somewhere near the origin because in a large distance from the origin (large d) even a slight change of the plane normal vector direction can cause a significant change of the plane position, which can negatively influence the convergence of the optimization method. In order to find the global maximum of some function, generally the initial point needs to be somewhat close to it. The right column of Figure 6.2 shows a visualization of the plane space, for the same three objects using the same parameterization as in the top row, where the red spots correspond to a region of planes from which the L-BFGS optimization converged to the global maximum - the convergence region. The objects were again simplified to approximately 1000 points. For Lion and Armadillo, the convergence regions are very large. For partial Armadillo, due to its much weaker symmetry, the region is smaller but still considerably big, with the angle between normal vectors of any two planes on the opposite sides of the region exceeding 45° . This indicates that in order to find the global maximum, the starting plane of the optimization does not need to be particularly close to it, even in the case of quite weakly symmetrical objects. This property of the symmetry measure will be exploited in the proposed symmetry detection method.

6.2 Proposed Symmetry Detection Method

The proposed global symmetry plane detection method roughly follows the common RANSAC scheme, which is very often used in symmetry detection and related areas. The key idea of RANSAC is generating a large number of candidate solutions based on the input data, which creates a good enough probability of a sufficiently good solution being among them. Each of the candidates is evaluated for its fitness, so creating unnecessarily large number of candidates can easily lead to a time consuming computation. But in our

case, since the numerical optimization converges to the global maximum of the symmetry measure from quite a large distance, we do not need to find a solution that close to the best one. We only need to find one in the considerably big convergence region, suggesting that a rather sparse sampling of the candidate space will be sufficient. Determining with certainty if a candidate lies in the convergence region is impossible, but a good indication is the symmetry measure itself. It can be expected that planes near the global maximum have a higher symmetry measure than those far from it.

The overview of the proposed method follows. Several candidate symmetry planes are created and only a small number of them, with the largest symmetry measure, are selected as having the largest potential of being in the convergence region. Then the optimization is started from these few planes and the resulting plane with the largest final symmetry measure is selected. The other planes to which the optimizations converge can be used as secondary planes in case of objects with multiple significant symmetries, as will be described later. For reasons described in Section 6.1.4, before the symmetry detection itself, the input set of points is translated so that its centroid is at the origin and in the end the inverse translation is applied to the resulting symmetry plane(s).

6.2.1 Creating the Candidate Planes

Meaningful candidate planes could be created by taking each pair of points $\mathbf{x}_i, \mathbf{x}_j \in X$, $i \neq j$ and creating the plane of symmetry of these two points. But using this approach directly on X results in an overwhelming number of planes, at least when X consists of more than a few tens of points. Therefore, we first simplify X using the algorithm described in Section 6.1.3 with $m = 100$, creating a new set of points X_{cand} with approximately 100 to 110 points. The candidate plane creation is then performed on X_{cand} , creating approximately 5000 to 6000 candidate planes, which is still unnecessarily many. However, we observe that many of the candidates are very similar, which means there is no need to evaluate all of them.

Candidate Pruning

We use the following distance function to measure the distance between two planes represented by $\mathbf{p}_u, \mathbf{p}_v$

$$D(\mathbf{p}_u, \mathbf{p}_v) = \begin{cases} \|\hat{\mathbf{p}}_u - \hat{\mathbf{p}}_v\| & \mathbf{n}_{\mathbf{p}_u}^T \mathbf{n}_{\mathbf{p}_v} \geq 0 \\ \|\hat{\mathbf{p}}_u + \hat{\mathbf{p}}_v\| & \mathbf{n}_{\mathbf{p}_u}^T \mathbf{n}_{\mathbf{p}_v} < 0 \end{cases}$$

where $\hat{\mathbf{p}} = \frac{1}{\|\mathbf{n}_{\mathbf{p}}\|} [a, b, c, \frac{d}{l_{avg}}]^T$ and $\mathbf{n}_{\mathbf{p}_u}$, $\mathbf{n}_{\mathbf{p}_v}$ are the normal vectors of the planes. This is basically the D_{ED} distance function from Chapter 5 only modified for non-normalized planes. A theoretically more appropriate distance function can be derived when representing the candidates as reflection transformations as described in Chapter 5 (see also Chapter 4 and [98]). However, as also shown in Chapter 5, the distance function we use here shows very similar behavior in practice (given the object is centered, which we did at the beginning), is simpler, easier to implement and does not require conversion to matrix represented transformations. When creating the candidates, if a newly created candidate \mathbf{p}_v is closer than $\delta = 0.1$ to the closest previously created candidate \mathbf{p}_u , the new one is not added to the candidate set. Instead, \mathbf{p}_u is replaced with $avg(\frac{1}{\|\mathbf{n}_{\mathbf{p}_u}\|} \mathbf{p}_u, \frac{1}{\|\mathbf{n}_{\mathbf{p}_v}\|} \mathbf{p}_v)$, computed as

$$avg(\mathbf{p}_u, \mathbf{p}_v) = \begin{cases} \mathbf{p}_u + \mathbf{p}_v & \mathbf{n}_{\mathbf{p}_u}^T \mathbf{n}_{\mathbf{p}_v} \geq 0 \\ \mathbf{p}_u - \mathbf{p}_v & \mathbf{n}_{\mathbf{p}_u}^T \mathbf{n}_{\mathbf{p}_v} < 0 \end{cases}.$$

If \mathbf{p}_u is already an average of ψ candidates, $\psi > 1$, then \mathbf{p}_v is only added to the average which is computed as $avg(\frac{1}{\|\mathbf{n}_{\mathbf{p}_u}\|} \mathbf{p}_u, \frac{1}{\psi \|\mathbf{n}_{\mathbf{p}_v}\|} \mathbf{p}_v)$. To accelerate locating the closest candidate, we use a 4D grid with cell size $\delta \times \delta \times \delta \times \delta$ where for each candidate \mathbf{p} , the vector $\hat{\mathbf{p}}$ is stored. This is plausible because D is a Euclidean distance. The grid query is similar to the one described in Section 6.1.2 for the 3D grid, only differing in that when locating the closest candidate to \mathbf{p} , the query is executed for both $\hat{\mathbf{p}}$ and $-\hat{\mathbf{p}}$ because they represent the same plane. We observe that candidates created by averaging a very low number of planes can quite safely be considered outliers. Therefore, we further remove all candidates created by averaging less than $\psi_{min} = 4$ planes, which basically means that there are at most 3 pairs of points in X_{cand} that are roughly symmetric w.r.t. the given candidate plane, which is too few to generate a meaningful candidate. The remaining candidates represent the final candidate set, which usually consists of only a few hundreds of planes, so the candidate space is sampled rather sparsely but still sufficiently.

6.2.2 Selecting the Best Candidates

Once the candidates are created, we use the algorithm from Section 6.1.3 with $m = 1000$ to create another simplified version of X with approximately 1000 points, denoted X_{simp} . Then we compute the symmetry measure $s_{X_{simp}}$ for all the candidate planes. In case of objects with weaker or multiple less significant symmetries, it is not safe to only select the best candidate to initialize the optimization because its relatively large symmetry measure

can be caused by its proximity to some significant local maximum rather than the global one. So we select S candidates with the largest symmetry measure and start the optimization from all of them. As the candidate pruning ensures that the candidates are not too similar, it is unlikely that the optimization converges to the same plane from all of them. In the end, we get S local maxima of the symmetry measure, among which we select the one with the largest symmetry measure and declare it the resulting plane of symmetry. We use $S = 5$.

6.2.3 Detecting Multiple Planes

Since the above described approach provides S local maxima of the symmetry measure the method can be easily adjusted to find multiple symmetry planes of the input object. It is very likely that two or more planes among the S local maxima will be very similar, therefore, to avoid detecting the same plane multiple times, we do the following. We sort the S planes representing the local maxima according to their symmetry measure in a descending order and we iterate through the sorted list from the first to the last plane. The first plane is always accepted as one of the resulting symmetry planes and any other plane is only accepted if its distance D from the closest already accepted plane is larger than 0.25. The method can find at most S significant symmetries this way. Therefore, when used to detect multiple planes, it is advisable to increase the value of S for the price of slightly larger time consumption. We use $S = 20$ instead of the original 5, which is sufficient in most applications. Because the local maxima can represent symmetry planes of varying significance, we further recommend accepting only planes with symmetry measure above some predefined threshold. The value of the threshold depends on the intended application but for a general purpose we recommend a value about 70% of the symmetry measure of the best detected plane.

6.2.4 Using the Weights

Until now we have been ignoring the weights w_{ij} in the symmetry measure s_X (see Equation 6.1) and considered all of them equal to 1. But when using the weights the method can be much more flexible. The weight w_{ij} can be expressed as $w_{ij} = w_{ij}^s w_{ij}^d(\mathbf{p})$ where w_{ij}^s is a static weight and $w_{ij}^d(\mathbf{p})$ is a dynamic weight which depends on the plane P represented by \mathbf{p} . Now the

symmetry measure s_X can be expressed as

$$s_X(\mathbf{p}) = \sum_{i=1}^n \sum_{j=1}^n w_{ij}^s w_{ij}^d(\mathbf{p}) \varphi(\|\mathbf{r}(\mathbf{p}, \mathbf{x}_i) - \mathbf{x}_j\|).$$

The dynamic weights can, for example, represent the symmetry of normal vectors or directions of principal curvatures in corresponding pairs of points with respect to a given plane. The static weights can be set to represent the importance of given pairs of points. In other words, the more it is desired for the point \mathbf{x}_i to end up in or near the point \mathbf{x}_j after reflecting it over the plane of symmetry the higher the static weight w_{ij}^s should be. The importance can be set manually by the user or as a similarity of some kind of feature function values in the given two points. As this feature function, for example, some type of curvature can be used, since it can be expected that in two symmetrical points there are similar curvature values.

In the rest of this section we show two modifications of the proposed method that use the weights in a beneficial manner. The first modification uses similarity of Gaussian curvature values to set the static weights and the symmetry of normal vectors for the dynamic weights. The Gaussian curvature as the feature function was chosen because it has previously shown itself useful for quantifying point similarity on 3D shapes [47] in context of symmetry detection. This modification will be described in Section 6.2.5. How this weighting can be used will also be shown in Section 6.3.4. The second modification only uses the static weights, which are set according to point-to-point distance, to force the method to detect a plane perpendicular to the largest dimension of the object. It will be described in Section 6.2.6. The results and possible purpose of both these modifications will be shown in Section 6.3.4.

6.2.5 Using Gaussian Curvature and Normal Symmetry

Here we describe the first modification which uses the Gaussian curvature values and normal vectors. The curvature values can help to identify similarity of features in the input object, which can be useful in case of objects with severe partiality where the whole object is completely asymmetrical but some features still exhibit symmetry. The normal vectors can further help to make the resulting symmetry more accurate. Let us now suppose that the input object is not represented by a general set of points but by a manifold triangle mesh on which the Gaussian curvature and a unit normal vector can be computed for each vertex. The normal vectors of all vertices

are computed by summing the normal vectors of triangles adjacent to the given vertex and normalizing the resulting vector. The values of Gaussian curvature are computed as described in [81].

Apart from the set of points $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ that represent vertices of the mesh we now also have a set of unit normal vectors $N = \{\mathbf{n}_1, \mathbf{n}_2, \dots, \mathbf{n}_n\}$ and a set of Gaussian curvatures $G = \{g_1, g_2, \dots, g_n\}$ where \mathbf{n}_i is the unit normal vector in the point \mathbf{x}_i and g_i is the Gaussian curvature in the point \mathbf{x}_i . It also needs to be defined how the normal vectors and the values of Gaussian curvature will be determined when the simplification algorithm is applied. When a new point of the simplified point set is created by averaging the points in a given cell of the simplification grid, its normal vector is determined by averaging the normal vectors in all points in the cell and normalizing the resulting vector. Its Gaussian curvature is taken from the point in the cell for which the absolute value of its Gaussian curvature is the largest.

Symmetry of Normal Vectors

The dynamic weights are set using the symmetry of normal vectors in vertices. In order to measure the reflectional symmetry of two unit normal vectors in two points, we first have to define a function $\mathbf{r}_n(\mathbf{p}, \mathbf{n})$ which reflects a unit normal vector \mathbf{n} over the plane P represented by \mathbf{p} . This function is defined as

$$\mathbf{r}_n(\mathbf{p}, \mathbf{n}) = \mathbf{n} - 2 \frac{\mathbf{n}_p^T \mathbf{n}}{\mathbf{n}_p^T \mathbf{n}_p} \mathbf{n}_p.$$

The symmetry of two normals \mathbf{n}_i and \mathbf{n}_j is defined as the similarity of $\mathbf{r}_n(\mathbf{p}, \mathbf{n}_i)$ and \mathbf{n}_j . To quantify such similarity we apply the similarity function $\varphi(l)$ on the angle between $\mathbf{r}_n(\mathbf{p}, \mathbf{n}_i)$ and \mathbf{n}_j . As φ our Wendland's function is used with the shape parameter α set as $\alpha = 4$. This value of α was chosen so that for angle $\frac{\pi}{16}$ (that is 11.25°) the similarity is approximately 0.5 because it is reasonable that only for low angles the similarity is significant (close to 1), otherwise the value could be quite close to 1 even for rather asymmetrical pairs of points. The dynamic weights $w_{ij}^d(\mathbf{p})$ are therefore defined as

$$w_{ij}^d(\mathbf{p}) = \varphi(\arccos(\mathbf{r}_n(\mathbf{p}, \mathbf{n}_i)^T \mathbf{n}_j)), \quad \text{with } \alpha = 4.$$

Similarity of Gaussian Curvature Values

The static weights are set using the Gaussian curvature values as follows

$$w_{ij}^s = \begin{cases} \frac{\min(|g_i|, |g_j|)}{\max(|g_i|, |g_j|)} & |g_i| \geq \frac{g_{avg}}{h} \wedge |g_j| \geq \frac{g_{avg}}{h} \wedge g_i g_j > 0 \\ 0 & \text{otherwise} \end{cases}.$$

The weight is non-zero only when both curvatures g_i and g_j have the same sign and their absolute values are both greater than the threshold $\frac{g_{avrg}}{h}$ where g_{avrg} is the average of absolute values of Gaussian curvatures in all points and h is a constant which we set as $h = 100$. This ensures that Gaussian curvatures with very small absolute values are not considered because they are usually present in points without significant features.

Additional Changes

When such weighting is used, we do not use the default candidate pruning, because it does not reflect the importance of given point pairs defined by the weights. Instead, we use the weights to prune the set of candidate planes in a smarter way. When creating a candidate plane P , represented by \mathbf{p} , as a symmetry plane of $\mathbf{x}_i \in X_{cand}$ and $\mathbf{x}_j \in X_{cand}$, we test whether the weights w_{ij}^s and $w_{ij}^d(\mathbf{p})$, computed using N_{cand} and G_{cand} , are above certain thresholds and if not, the given plane is not considered a candidate anymore. The sets N_{cand} and G_{cand} are the sets N and G (respectively) corresponding to the simplified point set X_{cand} . The normal vector information and Gaussian curvature information are quite damaged by the simplification process and therefore the thresholds for the weights should not be set very large. Specifically we consider the plane P , created as a symmetry plane of $\mathbf{x}_i \in X_{cand}$ and $\mathbf{x}_j \in X_{cand}$, a candidate plane if $w_{ij}^s > 0$ and $w_{ij}^d(\mathbf{p}) > 0.25$. Also, since with this weighting the symmetry measure gets more complex (less smooth) and more difficult to optimize, we use $m = 200$ instead of $m = 100$, when simplifying X to create X_{cand} , to sample the space of candidate planes more densely.

6.2.6 Using Point-to-Point Distance

Here we describe the second modification that uses the point-to-point distances to set the weights. There are many objects which are symmetrical with respect to more than one plane and in some cases only some of these planes are desired to be detected. For example, imagine an object with considerably different sizes in different dimensions (width, height, depth). In such cases, the user might want to always detect the plane along the smaller dimensions and perpendicular to the largest dimension to maximize the object's span across the detected symmetry plane. In order to achieve such behavior a very simple weighting can be applied by setting the static weights to reflect the distance between the corresponding points. Specifically, we set them as $w_{ij}^s = \frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{l_{avrg}}$. The dynamic weights remain unused, i.e. $w_{ij}^d(\mathbf{p}) = 1$.

When such weights are plugged into the symmetry measure, pairs of points with larger mutual distance will be treated as more important. Therefore, the method will try to find such a symmetry plane for which there is the best symmetry between the pairs with the largest distance. The rest of the method remains unmodified in this case. Similar weighting could also be used to achieve the opposite effect and force the method to detect planes along the largest dimension and perpendicular to the smaller ones, but we did not perform any experiments in this way.

6.3 Results

The proposed method was implemented in *C#* and its results were acquired on a computer with CPU *Intel® Core™ i7-4770* and 16 GB of memory running a *Windows 10* operating system. Until stated otherwise, we present results of detecting the single most significant symmetry plane, because this is what our method is primarily designed for, using the basic version of the proposed method where the weights in the symmetry measure are not used ($w_{ij} = 1$). Later we will also show some results of multiple plane detection and of the two modifications mentioned above. We always use the proposed method with its default parameter values introduced throughout the previous text.

We compare our method to the View-based method [69]. Although it is not very appropriate for weakly symmetrical objects, it shows good accuracy for objects with stronger symmetry, is fast and robust to noise and works on an arbitrary triangle mesh, not requiring any further property such as manifoldness. Therefore, we consider it one of the state-of-the-art methods for symmetry plane detection. Moreover, its implementation is publicly available [68]. In our experiments we used the implementation of the View-based method from [68] with its default parameter setting proposed in [69].

We further compare our method to the Clustering-based method by Shi et al. [107], which is an improved version of one of the most commonly known symmetry detection methods by Mitra et al. [83]. Apart from using a more appropriate metric in the transformation space, [107] is very similar to [83], so comparison to [107] can, in some sense, be considered a comparison to [83] as well. The Clustering-based method is designed to detect symmetries of much more general types, not just reflectional ones. We used the implementation of [107] provided to us by its authors and adjusted by themselves to specifically detect reflectional symmetries (symmetry planes). If the method found multiple symmetries, the most significant one was selected as the res-

ult according to the authors’ instructions. The default parameter values were also provided to us by the authors. The Clustering-based method was run on *Linux Mint 18.3*.

There are other relevant methods for symmetry plane detection but mainly because their implementation is not available (as also mentioned by [69]), we compare to [69] and [107]. For fair comparison to the View-based and Clustering-based methods (which require a triangle mesh) and because triangle meshes are the most common way of representing 3D objects in computer graphics, the test objects we used are all triangle meshes. The basic version (with $w_{ij} = 1$) of our method only uses the vertices of each mesh as the input set of points. However, unlike [69] and [107], the basic version of our method could as well be used for different representations including raw point clouds, since it only works with points and does not require connectivity information.

Figure 6.3 shows 26 objects of varying shapes and properties, together with their symmetry planes detected by the proposed method. The objects are always rendered so that the detected plane (marked by the line) is perpendicular to the plane of the figure. They were taken from various datasets [1, 32, 67, 108, 116] and the four faces were provided to us by the authors of the *Fidentis* project [18]. The point counts of all the objects are in Table 6.1.

The Figures 6.3a - 6.3i contain artificial and mostly strongly symmetrical objects, 6.3j - 6.3u show realistic 3D-scanned objects and 6.3v - 6.3z objects with missing parts, usually damaged artificially by clipping, except for the Embrasure, which was damaged naturally before it was scanned. The proposed method detected correct symmetry planes for all the objects, including the incomplete ones, where detecting a naturally appearing symmetry plane is generally quite challenging.

Under each object its point count is shown and there are also the V and C letters, which stand for the View-based and Clustering-based method, respectively, followed by either a check mark - ✓, in case the corresponding method provided a correct symmetry plane, mostly visually similar to the one our method detected, or a cross mark - ✗, in case the result of the method was wrong or considerably imprecise. With the Clustering-based method we further use the check mark in brackets - (✓), which means that it did not work with its default parameter configuration but we were able to find such a configuration which made it provide a plausible symmetry plane for the given object. The View-based method was only tested with its default parameter values since changing them did not lead to noticeable improvement in its results, only to larger computation time. The implementation of the View-

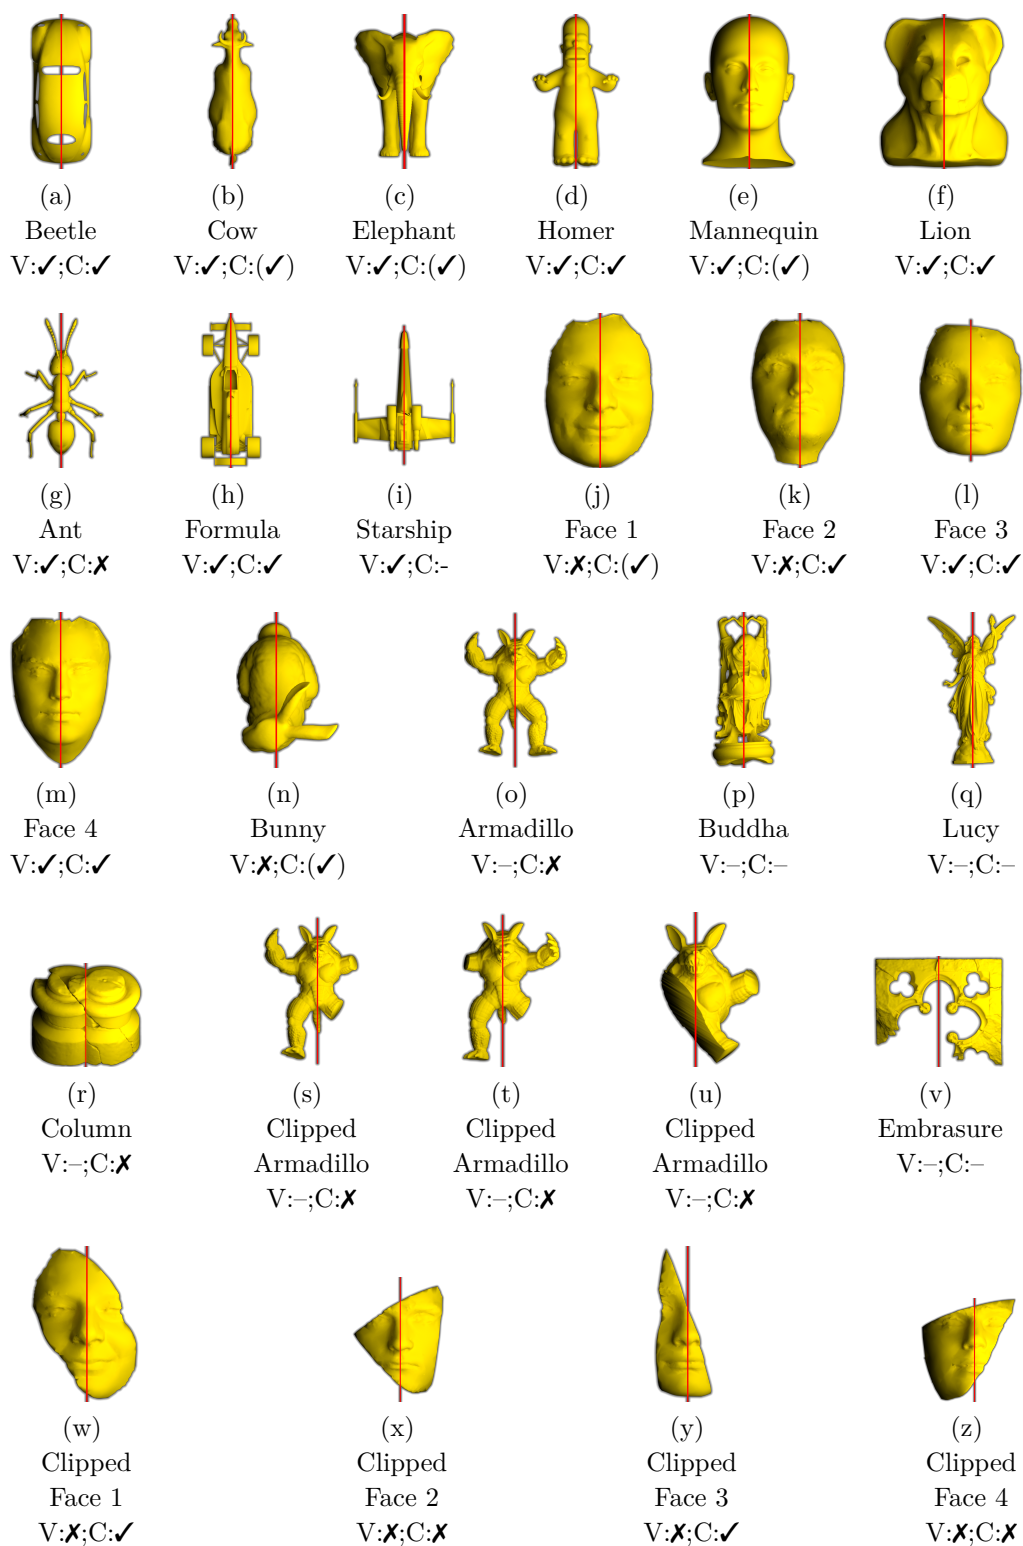


Figure 6.3: Several objects with their symmetry planes detected by the proposed method, the symbols after the V/C letter state wheter the View-based/Clustering-based method detected a plausible plane for the given object.

Table 6.1: The number of points of each of the testing objects.

| Object | Point count | Object | Point count | Object | Point count |
|-----------|-------------|-----------|-------------|----------------------------|-------------|
| Beetle | 988 | Face 1 | 2500 | Clipped Armadillo (first) | 135601 |
| Cow | 2903 | Face 2 | 2500 | Clipped Armadillo (second) | 134989 |
| Elephant | 19753 | Face 3 | 2500 | Clipped Armadillo (third) | 64647 |
| Homer | 5103 | Face 4 | 2500 | Embrasure | 117535 |
| Mannequin | 6737 | Bunny | 34834 | Clipped face 1 | 2224 |
| Lion | 2213 | Armadillo | 172974 | Clipped face 2 | 1782 |
| Ant | 3495 | Buddha | 543103 | Clipped face 3 | 1401 |
| Formula | 10969 | Lucy | 750001 | Clipped face 4 | 1602 |
| Starship | 3099 | Column | 69512 | | |

based method crashed, for unknown reason, when used on any object with more than approximately 35000 vertices and the one of the Clustering-based method also crashed for several objects, probably due to non-manifoldness. These cases are indicated by a dash.

The View-based method detected correct symmetry planes on 11 out of 18 of the objects which its implementation was able to process. It seems to work flawlessly on the strongly symmetrical objects (9 out of 9 correct) but expectedly the objects with weaker symmetries (2 out of 5 correct) and missing parts (0 out of 4 correct) are problematic for this method. The Clustering-based method was able to correctly detect planes on 14 out of 22 of the objects its implementation processed but only on 9 of them it worked well with its default parameter configuration, on the other 5 it only worked after tuning the parameters for the specific objects. It suggests that this method is rather sensitive to its parameter setting but it shows a larger potential on weakly symmetrical and incomplete objects than the View-based method.

Figure 6.4 shows two selected fail cases of the View-based and two of the Clustering-based method - one of a noticeably imprecise detection and one of a completely wrong detection. These results were obtained using the default parameter settings. On the Ant object the Clustering-based method found no symmetry.

The results imply that the proposed method is considerably more robust than both the View-based and Clustering-based methods, mainly when used on objects with weaker symmetry and missing parts.

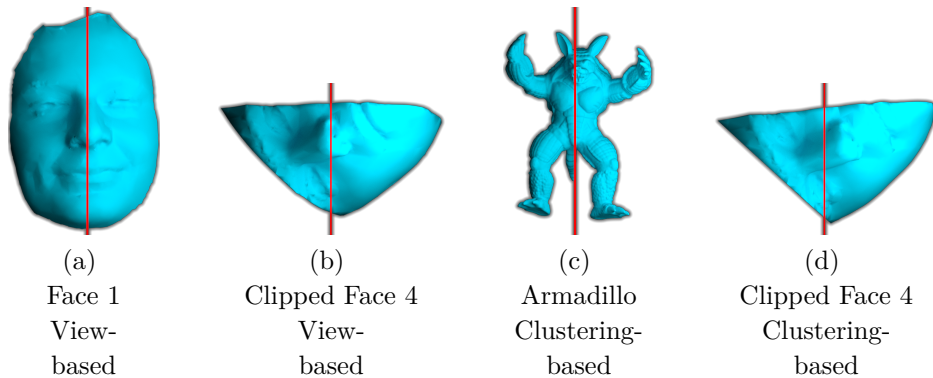


Figure 6.4: Two examples of imprecisely - (a), (c) and two of completely incorrectly - (b), (d) detected planes of the View-based - (a), (b) and the Clustering-based - (c), (d) methods.

For objects with missing parts, comparison to [112] would be suitable because it is a state-of-the-art method for symmetry plane detection for objects with missing parts. Unfortunately, despite our effort, we were unable to acquire its implementation and there are no numerical results in [112], so no fair comparison to this method is possible. But the proposed method was tested with good results on several objects which are the same or very similar to those used in [112] - the objects in Figures 6.3r, 6.3s, 6.3t, 6.3u and 6.3v. However, [112] requires a manifold triangle mesh and, due to the use of the Heat Kernel descriptor, does not work on featureless objects. These are quite constraining properties that make the method [112] much less general and less applicable than our method.

It will now be shown that even when used on objects where the View-based and Clustering-based methods detect a plausible symmetry plane, the proposed method mostly shows higher accuracy, which could not be seen in the visual results but can be demonstrated by a numerical measurement.

To measure the error of a detected symmetry plane, we apply the Metro [24] distance measure to evaluate the distance between the original object and the object created by reflecting the original one over the detected plane. The same error measurement was used by the authors of the View-based method in [69]. The Metro distance provides two values, the mean and the max error. The max error is the Hausdorff distance, which is very sensitive to outliers and often shows larger error for correct symmetry planes than for obviously incorrect ones. For this reason, we consider the max error inappropriate for measuring symmetry detection accuracy and we only use the mean error.

Table 6.2: Comparison of the proposed, View-based [69] and Clustering-based [107] methods in terms of Metro [24] mean error, the lowest error for each object is marked bold, the second lowest is in italic, ∞ indicates no symmetry was found, the average error excludes Ant and Starship.

| Object | Metro mean error | | | |
|-----------|-------------------|-------------------------|-----------------------------------|-----------------------------------|
| | View-based method | Clustering-based method | Proposed method (aligned objects) | Proposed method (rotated objects) |
| Beetle | 0.00619 | <i>0.00586</i> | 0.00055 | 0.00055 |
| Cow | 0.00730 | (0.00116) | 0.00047 | <i>0.00072</i> |
| Elephant | 0.00030 | (0.00357) | <i>0.00035</i> | 0.00046 |
| Homer | 0.00125 | 0.00627 | 0.00090 | <i>0.00097</i> |
| Mannequin | 0.00960 | (0.00838) | <i>0.00905</i> | 0.00838 |
| Lion | 0.00717 | 0.00368 | <i>0.00107</i> | 0.00101 |
| Ant | 0.00138 | ∞ | 0.00042 | <i>0.00055</i> |
| Formula | 0.00124 | 0.00523 | 0.00035 | <i>0.00044</i> |
| Starship | 0.00424 | N/A | 0.00154 | <i>0.00157</i> |
| Face 1 | 0.03086 | (0.01778) | <i>0.00815</i> | 0.00812 |
| Face 2 | 0.05604 | 0.02569 | <i>0.01156</i> | 0.01191 |
| Face 3 | 0.01704 | 0.01922 | 0.00955 | <i>0.00943</i> |
| Face 4 | 0.01716 | 0.02245 | <i>0.00571</i> | 0.00567 |
| Average | 0.01401 | 0.01084 | <i>0.00434</i> | 0.00433 |

Table 6.2 shows the error values of the symmetry planes detected by the View-based and Clustering-based methods and by the proposed method on several objects from Figure 6.3. Since the proposed method is not fully rotation invariant, we present two error values for each object - one for an axis-aligned object and another for a randomly rotated version of the same object. These two values are always very similar, implying that the accuracy of the proposed method does not depend much on the orientation of the object. The errors in brackets were obtained after parameter tuning consistently with what is stated in Figure 6.3. The proposed method shows mostly the lowest error (usually several times lower than the other two methods), except for the Elephant, where the View-based method is slightly more precise, and the Mannequin, where the Clustering-based method shows slightly smaller error than the proposed method on the aligned version, but it was achieved after parameter tuning, without which it did not find a plausible

Table 6.3: Running times (in seconds) of the proposed method and the implementations of the View-based [69] and Clustering-based [107] methods, the shortest time in each row is marked bold, the second shortest is in italic.

| Object | Point count | Time [s] | | | |
|-----------|-------------|-------------------|-------------------------|-------------------------------|------------------------------|
| | | View-based method | Clustering-based method | Proposed method single-thread | Proposed method multi-thread |
| Beetle | 988 | <i>0.41</i> | 16.9 | 0.54 | 0.25 |
| Lion | 2213 | <i>0.45</i> | 18.0 | 0.51 | 0.23 |
| Face 1 | 2500 | <i>0.45</i> | 18.9 | <i>0.45</i> | 0.21 |
| Face 2 | 2500 | <i>0.45</i> | 20.6 | 0.46 | 0.22 |
| Face 3 | 2500 | <i>0.45</i> | 19.3 | 0.48 | 0.23 |
| Face 4 | 2500 | <i>0.45</i> | 18.1 | 0.51 | 0.27 |
| Cow | 2903 | <i>0.46</i> | 21.7 | 0.63 | 0.31 |
| Starship | 3099 | <i>0.48</i> | N/A | 0.50 | 0.23 |
| Ant | 3495 | <i>0.49</i> | 17.6 | 0.65 | 0.29 |
| Homer | 5103 | 0.55 | 19.1 | <i>0.47</i> | 0.23 |
| Manneq. | 6737 | 0.60 | 30.1 | <i>0.42</i> | 0.22 |
| Formula | 10969 | 0.69 | 23.7 | <i>0.66</i> | 0.30 |
| Elephant | 19753 | 1.00 | 32.7 | <i>0.46</i> | 0.25 |
| Bunny | 34834 | 1.44 | 145.6 | <i>0.58</i> | 0.25 |
| Armadillo | 172974 | N/A | 1190.0 | <i>0.64</i> | 0.38 |
| Buddha | 543103 | N/A | N/A | <i>0.78</i> | 0.55 |
| Lucy | 750001 | N/A | N/A | <i>0.90</i> | 0.68 |

plane. The average error in the bottom row does not include Ant and Starship and it is considerably smaller for the proposed method than for the other two methods.

Table 6.3 shows the running times of the proposed method and the implementations of View-based and Clustering-based (with default parameters) methods for several objects. Since some parts of our method can be easily parallelized, we show times of both single-thread and multi-thread implementations. The Clustering-based method appears incomparably slower than the other two methods, so we further exclude it from the timing comparison.

The graph in Figure 6.5 shows timing comparison of the proposed method to the implementation of the View-based method with respect to the point

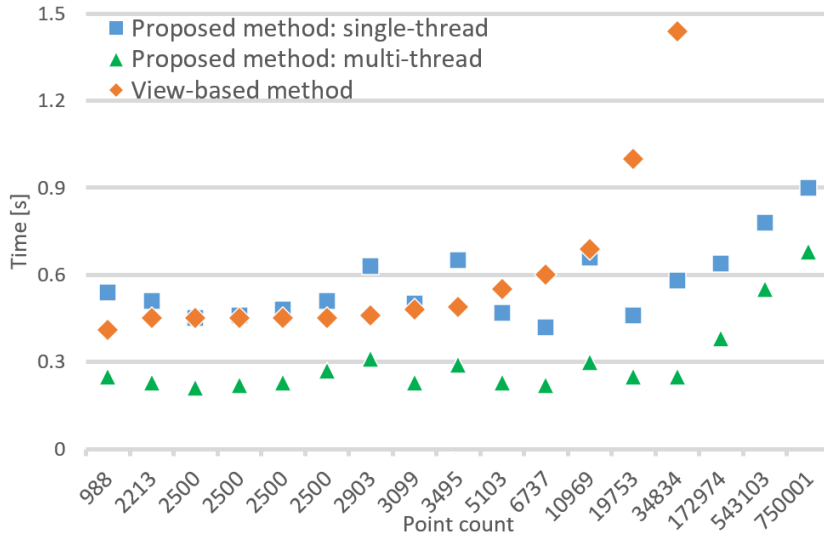


Figure 6.5: Comparison of the running times of the proposed method and the implementation of the View-based method [69].

count. The multi-threaded implementation of our method is faster than the View-based method. The single-threaded version is comparable to the View-based method for objects with point count up to approximately 10000 but after this point the running time of our method grows much less, which is due to the fast simplification (see Section 6.1.3).

Unfortunately, the running times of the View-based method for the last three objects are not available, however, [69] states that the method processed an object with 467252 points in 48.2s on a computer with CPU *Intel® Xeon® X5675* (clock rate 3.07 GHz) which is roughly comparable to our machine. Since our method processed an object with 750001 points in 0.9s, this suggests that for objects with larger point count our method is significantly faster than the View-based method.

6.3.1 Noisy Objects

We created noisy versions of four different objects by adding a random vector $[rand_x, rand_y, rand_z]^T \cdot l_{avg} \cdot mag$ to each point of the object, where $rand_x$, $rand_y$ and $rand_z$ are uniform random values from $\langle -1; 1 \rangle$ and mag is a constant which determines the noise magnitude. Figure 6.6 shows the symmetry planes detected by the proposed method on these objects, where the noise was created with $mag = 0.05$ for the objects in the top row, and with $mag = 0.1$ for those in the bottom row. Although the noise is very strong, the detected symmetry planes are still visually correct. The Clustering-based method found a plausible plane only for the Beetle with




Figure 6.6: The symmetry planes detected by the proposed method on four objects with added noise with $mag = 0.05$ - top row, and with $mag = 0.1$ - bottom row.

$mag = 0.05$ but only after parameter tuning, suggesting that this method is not suitable for objects with significant noise, which is also admitted in [107]. The View-based method only failed on the Lion with $mag = 0.05$, which can be considered an exception, implying that this method is comparable to our method in robustness to noise. The failure of the View-based and two selected failures of the Clustering-based method are depicted in Figure 6.7.

In the above mentioned cases the noise was uniformly distributed across the whole surface, so the random deviations of all points approximately




Figure 6.7: Two imprecisely - (a), (b) and one incorrectly - (c) detected planes on the noisy objects by the View-based - (a) and the Clustering-based - (b), (c) methods.




Figure 6.8: The symmetry planes detected by the proposed method - top row, and by the View-based method - bottom row, on four objects with non-uniformly distributed noise.

cancel out throughout the object and its symmetry actually does not suffer that much. Figure 6.8 depicts the same four objects with added noise with $mag = 0.05$ but the noise is only added to one half of each object. The top row shows the planes detected by the proposed method, which are all visually correct, and the bottom row shows the planes detected by the View-based method, which are noticeably inaccurate. These results imply that the proposed method is more robust to non-uniformly distributed noise than the View-based method and, therefore, more robust to noise in general.

6.3.2 Tests on a Larger Dataset

We tested our method on the *Thing10K* [128] dataset, which consists of 10000 3D objects represented by triangle meshes. Results of our method for all objects in the dataset can be found here [52]. From the dataset we extracted all objects that could be processed by the implementations of both the View-based and Clustering-based methods (manifold with up to 30000 vertices). We randomly selected 100 objects from this set and replaced objects without any significant reflectional symmetry with new random ones, which we repeated until there were no objects to replace. We ran the proposed, View-based and Clustering-based (with default parameters) methods for all the resulting 100 objects and for each method we manually observed the results and counted correct detections.

The results are in Table 6.4. Cases where the detected plane captured a

Table 6.4: Number of Correct/Imprecise/Incorrect detections for the 100 random objects from Thingi10k [128].

| | View-based | Clustering-based | Proposed | Proposed (sampled) |
|-----------|------------|------------------|----------|--------------------|
| Incorrect | 10 | 30 | 3 | 0 |
| Imprecise | 6 | 13 | 2 | 1 |
| Correct | 84 | 57 | 95 | 99 |

significant symmetry are marked as Correct. Those that were quite close to some correct symmetry plane, but with a noticeable imprecision, are marked as Imprecise, and those that were completely wrong as Incorrect. The incorrect detections also include cases where the Clustering-based method found no symmetry (12 cases). The proposed method shows the highest number of correct detections with 3 incorrect and 2 imprecise detections. This is caused by some objects in the dataset being heavily undersampled in some areas and oversampled in others, e.g. typical CAD models that often contain densely sampled details but their planar or otherwise monolithic areas are only covered by large triangles and therefore very few or no vertices. Since for our method we only used the vertices and not the triangles, this property sometimes forced the method to detect a symmetry of some small oversampled detail, which was globally incorrect. However, this can easily and quickly be resolved by uniformly sampling the surface instead of only taking the vertices as the input point set. We used the *MeshLab's* [23] *stratified triangle sampling* to extract a better quality point set for each of the 100 objects, after which our method achieved 99 correct detections (see the last column of Table 6.4).

6.3.3 Detecting Multiple Planes

The method can also be used to detect multiple symmetries, as described in Section 6.2.3. We used it on three different objects and only accepted planes with symmetry measure larger than 70% of the best plane's symmetry measure. The objects with the detected planes are shown in Figure 6.9. The number under each figure denotes the symmetry measure of the plane relative to the measure of the best plane for the given object. The computation times in the caption are of the single-/multi-thread implementation.

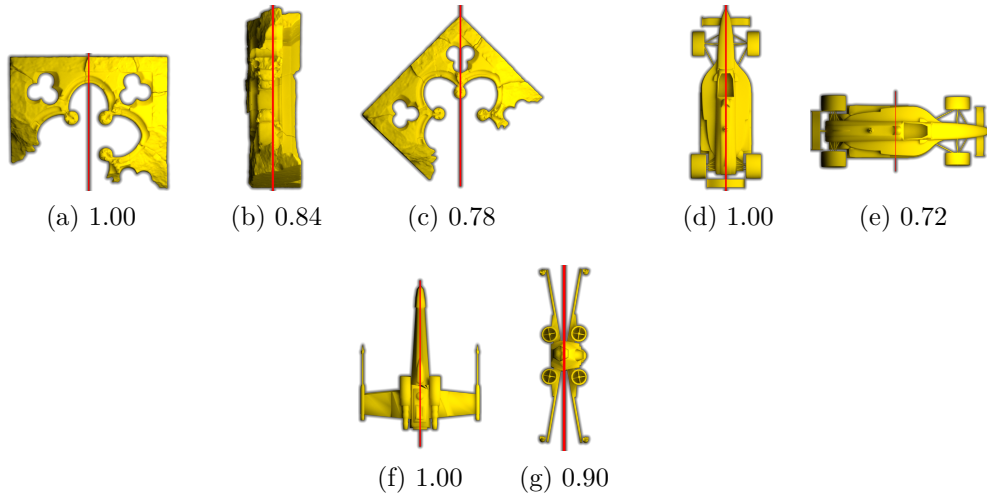


Figure 6.9: All the planes detected by the proposed method on the Embrasure - (a), (b), (c) (time: 0.89/0.40 s), Formula - (d), (e) (time: 1.66/0.58 s) and Starship - (f), (g) (time: 0.84/0.32 s).

6.3.4 Results of the Modified Versions

We mentioned two modifications of our method in Section 6.2.4 which both employ the weights in the symmetry measure in some way. Their results follow.

Gaussian Curvature and Normal Symmetry Weighting

Curvatures in general are very good for detecting local features of 3D models, especially those representing human faces or human heads [64]. In the following text, we will show that the first modified version of our method, which uses the Gaussian curvature similarity and the normal vector symmetry to set the weights w_{ij} (see Section 6.2.5), can be used to detect the symmetry plane of very small parts of human faces, as long as at least some local features are preserved.

Figure 6.10 shows the symmetry planes detected using this modification on heavily damaged versions of two of the scanned human faces. The figure also contains the original non-damaged faces for comparison. In all three cases the symmetry planes are detected correctly, despite the fact that the objects contain very little symmetry information, because there are still some symmetrical local features preserved. Also, the normal vector symmetry weighting makes the symmetry detection more accurate. When the non-weighted version was used on these objects, it found a completely incorrect plane in all cases.




Figure 6.10: Faces 2 and 3 with symmetry planes - (a), (c) and small parts of Face 2 - (b) (980 points, time: 1.84/0.62 s) and Face 3 - (d) (1026 points, time: 1.72/0.72 s), (e) (803 points, time: 1.52/0.52 s) with symmetry planes detected by the proposed method with the Gaussian curvature and normal symmetry weighting.




Figure 6.11: A cylindrical object and its symmetry plane detected by the non-weighted version of the proposed method - (a) and symmetry plane detected by the proposed method with the distance weighting for this and two other objects- (b), (c), (d).

Distance Weighting

The second modification, which uses point-to-point distances as weights (see Section 6.2.6), can be used to force the method to detect a plane along the smaller and perpendicular to the larger dimension of the input object. For example, see the cylindrical object depicted in Figure 6.11a together with the plane detected on it by the non-weighted version of the proposed method. In some cases, the user might prefer to find the plane shown in Figure 6.11b instead, although the object is slightly less symmetrical with respect to this plane. This plane is exactly the one detected by the proposed method with the distance weighting. Figures 6.11c and 6.11d depict additional two objects with the symmetry planes detected using the distance weighted modification of our method. In both cases the plane captures some non-negligible symmetry in the object (the wheels in the Formula and the fenders in the Beetle) and is also perpendicular to its largest dimension. This is despite the fact

that in both cases there is one considerably stronger symmetry along the largest dimension, which the distance weighting suppresses.

6.4 Parameters

The basic version of the proposed method (with $w_{ij} = 1$, detecting a single plane) has several parameters whose default values were presented throughout this chapter at the first mention of each one. The default values were set based on extensive testing and the method does not require any parameter tuning in vast majority of cases. Nevertheless, there could be situations where the knowledge of the parameters' meaning could be useful. Most of the parameters represent the typical trade-off between reliability or accuracy and computational efficiency.

- The target point count of X_{simp} influences the precision of the detection. The larger the point count, the more accurately the symmetry measure represents the symmetry of the object and the more precise the detection is. At the same time increasing this value makes the method slower. The default is 1000.
- The target point count of X_{cand} determines the number of created symmetry candidates which grows quadratically with it. With more candidates, there is bigger chance of finding one sufficiently close to the global maximum of the symmetry measure but it also takes more time to prune and evaluate them. The default is 100.
- The δ parameter determines the roughness of the candidate pruning step. With lower δ more candidates are kept, increasing chance of good candidates being among them but also making the evaluation step more timely. Lower δ also requires selecting larger value of S (see below). The default is 0.1.
- Candidates that were created by averaging less than ψ_{min} planes during the pruning step are further removed. With larger ψ_{min} less candidates are kept and the evaluation step is less timely but there is also bigger chance of removing good candidates, decreasing reliability. The default is 4.
- In the end, S best candidates are selected to start the optimization from. Increasing S makes the last step more timely but also increases the chance of locating the global maximum of the symmetry measure. The default is 5.

The only parameter which is not part of the efficiency-to-quality trade-off is the α spread parameter of the similarity function φ in the symmetry measure. Although its value also influences the time consumption, it is not true that changing it to make the method slower also makes it provide better results and vice versa. We set it as $\alpha = \frac{15}{l_{avg}}$ by default and below we will show that the result of our method is not very sensitive to its value.

Above we showed that for several objects the Clustering-based method [107] requires parameter tuning. We observed that the results of the method are the most sensitive to the value of the parameter which sets the threshold for candidate pruning, let us denote it ϵ . The method creates a number of candidates by matching pairs of points and for each candidate it performs a quick relevance test and only keeps candidates that pass the test. In [107] this is done by using the candidate to transform the local neighborhood of one of the two generating points and computing its average distance to the local neighborhood of the second point. If the distance is lower than ϵ the candidate is kept, otherwise removed. In the original Clustering-based method by Mitra et al. [83] the test is done using distances in signature space represented by curvature values where the distance between the signatures of the two generating points must be lower than some threshold in order to accept the candidate. This approach, according to [107], is less robust to noise than the one they use. However, the curvature values are also computed using local neighborhoods of the points so the two approaches are essentially very similar. Nevertheless, the sensitivity to the setting of the threshold parameter implies that the result of the clustering step is very sensitive to the quality of the candidates on its input.

Figures 6.12 and 6.13 show the sensitivity of the proposed method (top row) to the value of α and of the Clustering-based method [107] (bottom row) to the value of ϵ . For both methods we used different multiples of the default parameter values. We selected two objects where the Clustering-based method works with its default parameter setting but for different objects the optimal value of ϵ can be different from the default one. For the two lowest values of ϵ the Clustering-based method does not find any symmetry for either of the two objects. Increasing ϵ makes it provide incorrect results. The proposed method detects good symmetries for all the values of α except for the lowest one where it start being imprecise because the spread of φ starts getting too large and φ ceases to well represent point similarity.

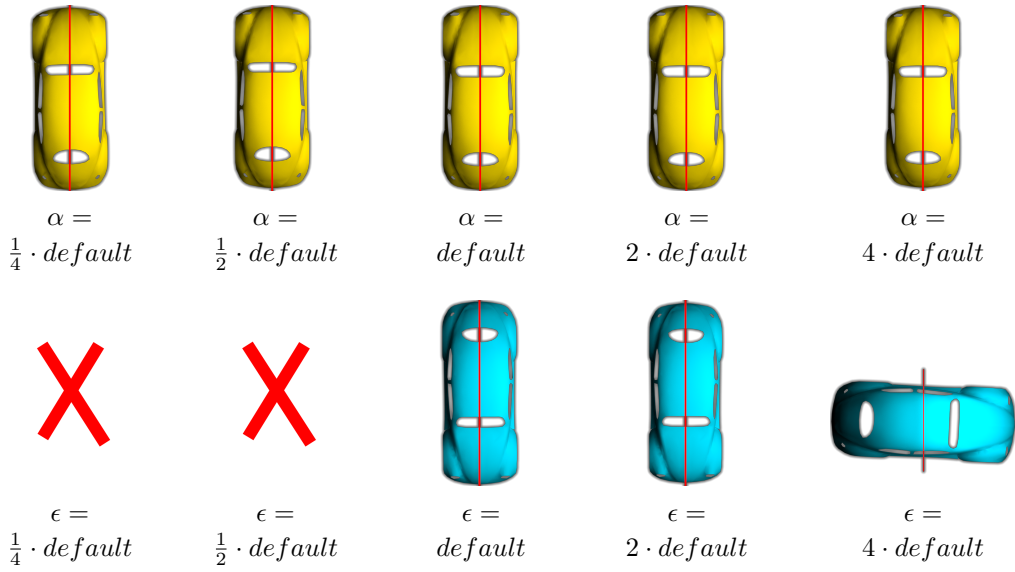


Figure 6.12: Results of the proposed method after changing the value of α - top row, and of the Clustering-based method [107] after changing the value of ϵ - bottom row, for the Beetle object.

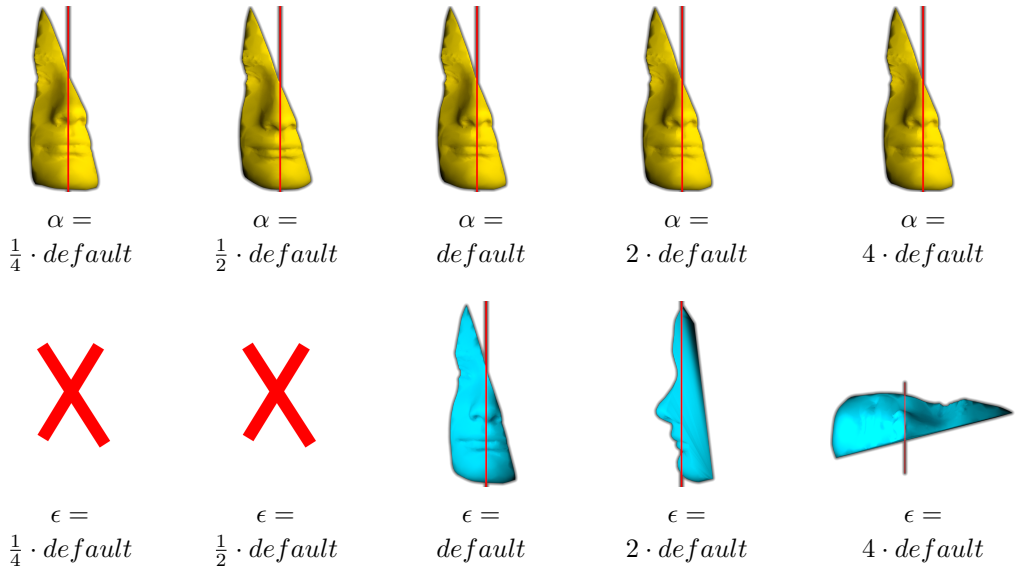


Figure 6.13: Results of the proposed method after changing the value of α - top row, and of the Clustering-based method [107] after changing the value of ϵ - bottom row, for the Clipped Face 3.

6.5 Another Application of the Symmetry Measure

In many applications (e.g. designing fields on symmetric surfaces [92]) it is desirable not only to find the symmetry plane but also to know which points are actually symmetrical with respect to the detected plane and, eventually, how much symmetry there is between them. Many of the existing methods, including the View-based and Clustering-based methods, do not implicitly provide such information. On the other hand, the symmetry measure used in the proposed method is essentially a sum of many terms where each represents the amount of symmetry between given two points, which can be used if needed. In some cases, for a given point $\mathbf{x}_i \in X$, we might want to find such a point $\mathbf{x}_j \in X$ for which there is the strongest symmetry between \mathbf{x}_i and \mathbf{x}_j w.r.t. a given plane P represented by \mathbf{p} . This only means finding such a point \mathbf{x}_j , for which the value $w_{ij}\varphi(\|\mathbf{r}(\mathbf{p}, \mathbf{x}_i) - \mathbf{x}_j\|)$ is the largest. Furthermore, if each point \mathbf{x}_i is colored according to this value, the symmetry of the object can be visualized as shown in Figure 6.14. The lighter the color, the larger the symmetry value. Figures 6.14a, 6.14b, 6.14c, 6.14d show the coloring for $w_{ij} = 1$. Figure 6.14e shows it for the case when the weights w_{ij} are set according to the distance weighting, resulting in almost no symmetry near the plane and the most symmetry in the wheels. Unlike some other methods, the concept of the proposed method also gives us a straightforward way to derive the symmetry information, with respect to any plane, for arbitrary points of the input object.




Figure 6.14: Four objects with their points colored according to the amount of symmetry with respect to their symmetry plane with no weights used - (a), (b), (c), (d) and using the distance weighting - (e) (darker - less symmetry, lighter - more symmetry).

6.6 Limitations

Although the proposed method was shown to be fast and robust in several ways, it still has some limitations and drawbacks. First, as mentioned in Section 6.3.2, the method does not handle well objects with very non-uniform point sampling. For triangle meshes this can be easily resolved using some uniform sampling technique but for raw point clouds it might be problematic and could be addressed in the future, possibly using the weights in the symmetry measure.

Second, when using the method for detecting multiple symmetries, a threshold needs to be set on the symmetry measure to only accept symmetries that are strong enough to be meaningful. There is, however, not yet a general way of determining this threshold for arbitrary input data. Finding it is also left for future research efforts. Last, we showed that our method handles well objects with missing parts and objects with noise. However, our preliminary experiments show that if a single object has both significant missing parts and strong noise, the method often tends to fail. This might be possible to solve by some parameter tuning which would probably also lead to increased computational cost.

6.7 Summary

We proposed a new method for symmetry plane detection, based on a simple and flexible differentiable symmetry measure, and we have shown that it works very accurately for both strongly and weakly symmetrical objects and also for objects with significant noise or missing parts. We further demonstrated that in terms of robustness, accuracy and speed the proposed method outperforms the View-based [69] method and the Clustering-based [107] method which represents an improved version of the well known method of Mitra et al. [83].

Tests on a larger dataset showed that the proposed method exhibits good reliability on random objects of varying properties but they also revealed that it is not good at handling objects with significantly irregular sampling. We also showed how our method can be modified to detect multiple symmetries of a single object. Implementation of the proposed method is available for download at [52].

There are many possible ways of how the method can be modified, extended or generalized. The weights in the symmetry measure make the method very flexible and there are situations when using them can be very beneficial. We showed two specific ways of setting the weights and we also showed their

results to prove their usefulness. Further examining the possibilities of these weights and finding more beneficial uses for them is a viable direction for future research.

7 Rotational Symmetry Detection Using Reflectional Symmetries and Quaternions

While there are quite many existing methods for finding reflectional symmetries in 3D shapes, many of which are relatively robust in different ways (see Section 3.1), the options for finding rotational symmetries appear to be significantly sparser and most of them have some major limitations (see Section 3.2). Some can only find circular symmetries which means they only work on rotational shapes and only find the axis, not the angle. Other methods constrain the detected symmetries by having the axis pass through some reference point, such as the object's centroid which limits their applicability in cases of weaker symmetries or missing parts.

There are, of course, also methods that are capable of detecting more general symmetries (see Section 3.3). However, modifying such methods to find symmetries defined by a rotation around an axis would probably be difficult or even impossible because such a transformation is rather specific and different from, e.g., a general rigid transform with which such methods usually can work. On top of this, these more general methods are usually also quite computationally expensive.

To our knowledge, no existing method for detecting rotational symmetries in 3D objects can find both the axis and angle of rotation and be robust, at least to some extent, to weak symmetries and missing parts at the same time. However, as mentioned above, there are quite a few approaches for finding reflectional symmetries that are robust in several ways including our method presented in Chapter 6. This implies that finding a way of using potential reflectional symmetries to find plausible rotational symmetries would be highly beneficial because many of the existing approaches for reflectional symmetries in 3D data could be extended to find the rotational ones. In this chapter we present a new method for finding rotational symmetry in 3D objects by using this very idea to extend our previously described method for reflectional symmetry detection. The content of this chapter was previously published in [49] and its main contributions are following:

1. We propose a simple way of creating plausible candidates for rotational symmetries by combining suitable pairs of approximate reflectional symmetries.
2. We generalize the symmetry measure and the step of its maximization from Chapter 6 to work for rotational symmetries using a very simple, yet efficient, quaternion-based parameterization of the rotation transformation that seems novel in this area and stems from algebraically converting the quaternion rotation formula to a vector expression.
3. We combine these two contributions to extend the overall approach from Chapter 6 into a robust and efficient method for detecting rotational symmetry in a 3D point set.

Note that contributions 1 and 2 are both very well usable on their own and can be implemented in a different context as well. However, we believe that their combination provides a great basis for an overall very good solution to rotational symmetry detection.

7.1 Rotational Symmetry Detection Method

Let $\mathbf{r}(\mathbf{p}, \mathbf{x})$ be a transformation that reflects an arbitrary point \mathbf{x} over a given plane \mathbf{p} . If we have two reflections $\mathbf{r}_1 = \mathbf{r}(\mathbf{p}_1, \mathbf{x})$, $\mathbf{r}_2 = \mathbf{r}(\mathbf{p}_2, \mathbf{x})$ with non-parallel planes \mathbf{p}_1 , \mathbf{p}_2 , composing \mathbf{r}_1 and \mathbf{r}_2 results in rotation around the axis, which is an intersection of the planes \mathbf{p}_1 , \mathbf{p}_2 , by the angle which is a double of the oriented angle between the planes [36]. In many cases a rotational symmetry in a 3D object follows this rule and is actually a composition of two reflectional symmetries. If an object has perfect reflectional symmetry w.r.t. both planes \mathbf{p}_1 and \mathbf{p}_2 , the rotations $\mathbf{r}(\mathbf{p}_1, \mathbf{r}(\mathbf{p}_2, \mathbf{x}))$ and $\mathbf{r}(\mathbf{p}_2, \mathbf{r}(\mathbf{p}_1, \mathbf{x}))$ must also be its perfect symmetries - the object reflected over both \mathbf{p}_1 or \mathbf{p}_2 remains unchanged, i.e. it also remains unchanged when reflected over both these planes in succession (note, however, that the converse implication does not hold, i.e. not all rotationally symmetric objects have necessarily reflectional symmetries).

An object with a rotational symmetry often has multiple reflectional symmetries where the reflection planes intersect in the axis of the rotational symmetry and the rotation angle of the symmetry can be derived from the planes using the rule above. Let us demonstrate the idea on simple 2D cases (we replace reflection planes with lines and an axis of rotation with a center of rotation). Figure 7.1 shows three objects with several marked lines of reflectional symmetries. In Fig. 7.1a, when the triangle is reflected over the




Figure 7.1: 2D examples of reflectional symmetries where composing reflections over the lines results in rotational symmetries.

symmetry line **a** (turning the point \mathbf{x} into \mathbf{x}') and then over **b** (turning \mathbf{x}' into \mathbf{x}'') the resulting transformation is a rotation by 120° around the center O which lies at the intersection of the two lines that make an angle of 60° . Such a rotation is a symmetry of this object. Similarly, in Fig. 7.1b we get symmetry with 90° rotation around the center O for lines **a** and **b** or **b** and **c** and 180° for **a** and **c**. In Fig. 7.1c we get 72° rotational symmetry around O with **a** and **b** or **b** and **c** and 144° with **a** and **c**.

This idea is applicable in 3D as well. Therefore, if we have a reliable algorithm for finding planes of reflectional symmetries or plausible symmetry plane candidates, we can use them to generate meaningful candidates for rotational symmetries. It also holds when the symmetries are only approximate, i.e., no perfect reflectional symmetry exists. We use the method from Chapter 6 as the underlying symmetry plane detection algorithm and we propose its extension for finding rotational symmetries in 3D objects using the idea described above. We generate potential rotational symmetries from the reflectional ones and evaluate their fitness to select the best ones (see Section 7.1.3), in the end we perform refinement in the rotation space (see Section 7.1.4) using a proposed quaternion-based parameterization (see Section 7.1.2). Note, however, that this core idea could possibly be implemented with many other different symmetry plane detection approaches, as well as the proposed rotation refinement. A brief overview of the method from Chapter 6 follows.

7.1.1 Recapitulation of The Symmetry Plane Detection Method

In this section we briefly recapitulate the underlying symmetry plane detection method that is used as basis for our rotational symmetry detection and was described thoroughly in Chapter 6. The method takes a set of points in E^3 as input and tries to find its best plane of reflectional symmetry. This is done by maximizing an objective function called symmetry measure. The symmetry measure $s_X(\mathbf{T})$ is for a point set $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$, $\mathbf{x}_i \in E^3$, $i = 1, \dots, n$ defined as

$$s_X(\mathbf{T}) = \sum_{i=1}^n \sum_{j=1}^n \varphi(\|\mathbf{T}(\mathbf{x}_i) - \mathbf{x}_j\|) \quad (7.1)$$

where $\varphi(l)$ is so-called similarity function which transforms distance into similarity (the larger the distance l the lower the value of $\varphi(l)$) and \mathbf{T} is a geometric transformation which in this case is $\mathbf{T}(\mathbf{x}) = \mathbf{r}(\mathbf{p}, \mathbf{x})$. To simplify some parts of the brief description we set $\varphi(l) = \hat{\varphi}(\alpha l)$ where α is the spread parameter of φ and $\hat{\varphi}$ is just φ with its spread parameter set to 1. The final function φ is then just $\hat{\varphi}$ with its argument scaled by the spread parameter α which is by default set as $\alpha = \frac{15}{l_{avg}}$ where l_{avg} is the average distance of the points in the input point set from their centroid (see Section 6.1.1). The goal is to maximize s_X w.r.t. the parameters of the plane \mathbf{p} . The computation of s_X basically comes down to reflecting each point in X over \mathbf{p} , computing its similarity to all points of X and summing all these similarities. The more similar the reflected point set X is to the original X the larger the value of s_X should be. As $\hat{\varphi}(t)$ a specific Wendland's [123] function is used such that it closely resembles the Gaussian for $t \geq 0$ but is locally supported. It is a piece-wise polynomial differentiable function and it is that $\hat{\varphi}(0) = 1$, $\frac{d}{dt}\hat{\varphi}(0) = 0$ and $\hat{\varphi}(t) = 0$ for $t > 2.6$. See Section 6.1.1 for details and the exact definition of φ . With such a choice of φ the measure s_X is differentiable w.r.t. the parameters of the transformation \mathbf{T} if \mathbf{T} is parameterized differentially. It can also be made higher-order differentiable if needed by choosing a higher-order differentiable Wendland's function.

The overall process of finding the best symmetry plane has several steps:

1. Two different point sets are created by simplifying the input set of points X using a fast grid-based simplification algorithm. One denoted X_{simp} with target point count 1000, and one denoted X_{cand} with target point count 100. The actual point counts are usually slightly larger than these values.

2. The point set X_{cand} is used to create candidates for symmetry planes. This is done by taking every pair of points in X_{cand} and creating the symmetry plane of that pair. If a new candidate is similar (distance smaller than δ) to some already created candidate it is not added into the candidate set but merged with the similar one. The number of resulting candidates is usually in the order of hundreds.
3. Each candidate is evaluated using the symmetry measure but instead of computing it for the original point set X it is computed for the simplified X_{simp} which has approximately 1000 points. Such point count is still sufficient to well represent the symmetry of the input object and the computation of the symmetry measure is usually significantly faster than for the original point set X . Since the similarity function φ is locally supported the points outside its support region can be excluded to further decrease the computational cost. This is done using a 3D grid as an auxiliary data structure. Now, S candidates with the largest symmetry measure $s_{X_{simp}}$ are selected for further processing. By default $S = 5$.
4. Local numerical optimization using the gradient-based L-BFGS method [74] is started from all the S candidates from the previous step to maximize their symmetry measure $s_{X_{simp}}$. This provides S local maxima of the measure among which the one with the largest measure value is selected as the resulting symmetry plane of the input point set X . The other local maxima can potentially be used as secondary symmetries if the input object has more than one significant reflectional symmetry.

After the optimization step, we could theoretically use the two best symmetries and if they were perfect, their combination would provide a perfect rotational symmetry. However, we aim for a robust method that also works for objects with weak symmetries, significant noise, or even for incomplete objects. In such cases testing only one rotation created by two reflectional symmetries does not have to be reliable enough. Furthermore, better results can be expected when optimizing the final rotation instead of optimizing the reflectional symmetries used to create it. This is why we omit step 4 completely.

Instead, we propose a method for detecting rotational symmetry of a set of points in E^3 built upon the first 3 steps of the method described in Chapter 6. We use the best reflectional symmetry candidates outputted by step 3 and pair them to create candidates for rotational symmetries. Then we use the symmetry measure from Eq. 7.1, modified to work with rotations,

to evaluate the candidates and select a few best ones. As the final step, we perform the optimization directly in the space of rotations which is much more natural and, as will be shown in Section 7.1.4, can well compensate for the imprecision of the non-optimized reflectional symmetry candidates. Intuitively, this approach seems not to work for the detection of rotational symmetries in objects that are not reflectionally symmetric. However, we will show in Section 7.2 that it mostly works even in such cases. The details of the overall process of finding rotational symmetries will be provided in the rest of this section, as well as the description of any changes we made in the first three steps of the underlying symmetry plane detection method.

7.1.2 Symmetry Measure

In the following text we describe our modification of the symmetry measure to take rotation as the input transformation instead of reflection. We discovered that for finding rotational symmetries the symmetry measure needs to represent the symmetry of the input point set slightly more precisely than in the case of reflectional ones. This is probably because the rotation transformation has more degrees of freedom than a reflection. Therefore, we set the target point count of X_{simp} to 3000 instead of 1000 for the simplification algorithm and we use φ with slightly smaller radius by setting $\alpha = \frac{20}{l_{avg}}$ instead of $\alpha = \frac{15}{l_{avg}}$. Both these parameter values were set based on vast experiments to maximize the balance between the quality of the results and computation cost.

The symmetry measure allows finding its significant maxima by applying a rather sparse sampling of the candidate space and being able to converge to them from quite large distances using numerical optimization (see Section 6.1.4). This saves a lot of time in the candidate evaluation step because the candidate count does not need to be very large. Since the symmetry measure is differentiable, and mainly because its gradient can be computed analytically, any gradient-based optimization technique can be employed to locate the maxima quickly. Such techniques are usually faster than those that do not require the gradient because they often need a much smaller number of iterations. Furthermore, the differentiability makes the measure smooth, free of discontinuities and sudden changes in any direction which can make the optimization easier and faster in general, even with non-gradient methods. Therefore, when modifying the symmetry measure to work with rotations instead of reflections, it is useful to keep it differentiable with the possibility to compute its gradient analytically, retaining the option of using gradient-based methods and generally keeping the optimization efficient.

Rotation Parameterization

As mentioned, the symmetry measure is differentiable if the transformation \mathbf{T} is parameterized differentiably. In our case, \mathbf{T} is a rotation around an axis. Rotations are usually represented by matrices but in our case this would be cumbersome due to the need to keep the rotation matrix orthogonal throughout the parameterization. Angle-based parameterizations, on the other hand, commonly suffer from singularities, such as poles, which might negatively influence the stability of the optimization. We would also like to keep the parameterization simple, easy to implement and possible to differentiate analytically. To fulfill these conditions, we propose using the following and rather simple quaternion-based representation instead.

First, we briefly remind the basics of quaternion algebra. Let $Q = a + bi + cj + dk$ be a quaternion where i, j, k are elements of the quaternion basis and $i^2 = j^2 = k^2 = -1, ij = k = -ji, jk = i = -kj, ki = j = -ik$. A quaternion conjugate of Q is defined as $Q^* = a - bi - cj - dk$ and the length of a quaternion is given as $\|Q\| = \|Q^*\| = \sqrt{QQ^*} = \sqrt{a^2 + b^2 + c^2 + d^2}$ (for more details about quaternions we again refer to [38]). We denote $v(\mathbf{x}) = xi + yj + zk$ a quaternion that represents a vector or point $\mathbf{x} = [x, y, z]^T \in E^3$. Inversely, we denote $\mathbf{q}(u) = [x, y, z]^T \in E^3$ a vector representation of a quaternion $u = xi + yj + zk$.

Any quaternion can be expressed as $Q = a + v([b, c, d]^T)$. Subsequently, any unit quaternion can be expressed as $Q = \cos(\beta) + \sin(\beta)v(\mathbf{u})$ where \mathbf{u} is a unit vector. Then the expression $Qv(\mathbf{x})Q^*$ represents the point \mathbf{x} rotated by the angle 2β around the axis that passes through the origin in the direction of the vector \mathbf{u} . By denoting $a = \cos(\beta)$, $\mathbf{v} = [b, c, d]^T = \sin(\beta)\mathbf{u}$, $Q = a + v(\mathbf{v})$ we can expand this as

$$\begin{aligned} Qv(\mathbf{x})Q^* &= (a + v(\mathbf{v}))v(\mathbf{x})(a - v(\mathbf{v})) = \\ &a^2v(\mathbf{x}) - av(\mathbf{x})v(\mathbf{v}) + av(\mathbf{v})v(\mathbf{x}) - v(\mathbf{v})v(\mathbf{x})v(\mathbf{v}). \end{aligned}$$

Now using the rule that the product of two vectors \mathbf{a}, \mathbf{b} in quaternion algebra is $v(\mathbf{a})v(\mathbf{b}) = -\mathbf{a} \cdot \mathbf{b} + v(\mathbf{a} \times \mathbf{b})$ (see [38]) we can further expand the expression as

$$a^2v(\mathbf{x}) - a(-\mathbf{x} \cdot \mathbf{v} + v(\mathbf{x} \times \mathbf{v})) + a(-\mathbf{v} \cdot \mathbf{x} + v(\mathbf{v} \times \mathbf{x})) - (-\mathbf{v} \cdot \mathbf{x} + v(\mathbf{v} \times \mathbf{x}))v(\mathbf{v}).$$

As the cross product is anticommutative we can replace $v(\mathbf{v} \times \mathbf{x})$ with $-v(\mathbf{x} \times \mathbf{v})$ and since the terms $a(\mathbf{x} \cdot \mathbf{v})$ and $-a(\mathbf{v} \cdot \mathbf{x})$ cancel out we can continue with the expansion as follows

$$a^2v(\mathbf{x}) - 2av(\mathbf{x} \times \mathbf{v}) + (\mathbf{v} \cdot \mathbf{x})v(\mathbf{v}) + v(\mathbf{x} \times \mathbf{v})v(\mathbf{v}).$$

The last term can be further expanded into

$$v(\mathbf{x} \times \mathbf{v})v(\mathbf{v}) = -(\mathbf{x} \times \mathbf{v}) \cdot \mathbf{v} + v((\mathbf{x} \times \mathbf{v}) \times \mathbf{v})$$

and because the vectors $\mathbf{x} \times \mathbf{v}$ and \mathbf{v} are orthogonal the first term $-(\mathbf{x} \times \mathbf{v}) \cdot \mathbf{v}$ is 0. Using the well known vector triple product rule, i.e. having three vectors \mathbf{a} , \mathbf{b} , \mathbf{c} it is $(\mathbf{a} \times \mathbf{b}) \times \mathbf{c} = (\mathbf{a} \cdot \mathbf{c})\mathbf{b} - (\mathbf{b} \cdot \mathbf{c})\mathbf{a}$, we get

$$v(\mathbf{x} \times \mathbf{v})v(\mathbf{v}) = v((\mathbf{x} \cdot \mathbf{v})\mathbf{v} - (\mathbf{v} \cdot \mathbf{v})\mathbf{x}) = (\mathbf{x} \cdot \mathbf{v})v(\mathbf{v}) - (\mathbf{v} \cdot \mathbf{v})v(\mathbf{x}).$$

By plugging this back into the rotation formula we get

$$Qv(\mathbf{x})Q^* = a^2v(\mathbf{x}) - 2av(\mathbf{x} \times \mathbf{v}) + 2(\mathbf{x} \cdot \mathbf{v})v(\mathbf{v}) - (\mathbf{v} \cdot \mathbf{v})v(\mathbf{x})$$

and to use the transformation in the symmetry measure it needs to be expressed in the vector space which leads to

$$\mathbf{q}(Qv(\mathbf{x})Q^*) = a^2\mathbf{x} - 2a(\mathbf{x} \times \mathbf{v}) + 2(\mathbf{x} \cdot \mathbf{v})\mathbf{v} - (\mathbf{v} \cdot \mathbf{v})\mathbf{x}. \quad (7.2)$$

The vector expression can be easily differentiated and is useful also for implementation purposes because vector operations are easier to implement then operations in quaternion algebra.

Since in general the rotation axis does not pass through the origin we can select an arbitrary point $\mathbf{s} = [s_x, s_y, s_z]^T$ that lies on the axis and express the rotation as $\mathbf{q}(Qv(\mathbf{x} - \mathbf{s})Q^*) + \mathbf{s}$. The rotation transformation can be parameterized using the coefficients of the quaternion Q , i.e. a, b, c, d , and the coordinates of the point on the axis \mathbf{s} , i.e. s_x, s_y, s_z . A general quaternion $Q = a + bi + cj + dk$ is not unit and since only unit quaternions represent rotations we need to incorporate normalization into the parameterization. We therefore define the rotation transformation as

$$\begin{aligned} \mathbf{T}(\mathbf{x}) &= \mathbf{rot}(Q, \mathbf{s}, \mathbf{x}) = \mathbf{q} \left(\frac{1}{\|Q\|} Qv(\mathbf{x} - \mathbf{s}) \frac{1}{\|Q\|} Q^* \right) + \mathbf{s} \\ &= \mathbf{q} \left(\frac{1}{\|Q\|^2} Qv(\mathbf{x} - \mathbf{s}) Q^* \right) + \mathbf{s} = \frac{1}{a^2 + b^2 + c^2 + d^2} \mathbf{q}(Qv(\mathbf{x} - \mathbf{s})Q^*) + \mathbf{s}. \end{aligned}$$

By again denoting $\mathbf{v} = [b, c, d]^T$, $Q = a + v(\mathbf{v})$ and using the vector expression from Eq. 7.2 we can finally rewrite the transformation as

$$\begin{aligned} \mathbf{rot}(Q, \mathbf{s}, \mathbf{x}) &= \frac{1}{a^2 + \mathbf{v} \cdot \mathbf{v}} \mathbf{q}(Qv(\mathbf{x} - \mathbf{s})Q^*) + \mathbf{s} = \\ &= \frac{1}{a^2 + \mathbf{v} \cdot \mathbf{v}} \left(a^2(\mathbf{x} - \mathbf{s}) - 2a((\mathbf{x} - \mathbf{s}) \times \mathbf{v}) + 2((\mathbf{x} - \mathbf{s}) \cdot \mathbf{v})\mathbf{v} - (\mathbf{v} \cdot \mathbf{v})(\mathbf{x} - \mathbf{s}) \right) + \mathbf{s}. \end{aligned}$$

This easy-to-implement formula is obviously differentiable w.r.t. all seven parameters $a, b, c, d, s_x, s_y, s_z$ (except for $a = b = c = d = 0$ which is not a rotation) and it only uses basic algebraic operations.

As a result, by substituting $\mathbf{T}(\mathbf{x}) = \mathbf{rot}(Q, \mathbf{s}, \mathbf{x})$ we get a symmetry measure in the form

$$s_X(\mathbf{rot}) = s_X(Q, \mathbf{s}) = \sum_{i=1}^n \sum_{j=1}^n \varphi(\|\mathbf{rot}(Q, \mathbf{s}, \mathbf{x}_i) - \mathbf{x}_j\|) \quad (7.3)$$

with the following useful properties all of which are only possible due to the simple, yet efficient, quaternion-based parameterization of the rotation transformation:

1. It is still differentiable w.r.t. all parameters of the transformation and its gradient can be analytically computed.
2. Apart from the square root function needed to compute Euclidean distances it only uses basic algebraic operations (φ is piece-wise polynomial) - no need for trigonometric functions and conversions between angles and their sines/cosines.
3. It lacks major singularities, such as poles that appear in the case of angle-based representations of the rotation, allowing for more stable computation and optimization.

Small Angle Penalization

The use of the symmetry measure from Eq. 7.3 often tends to rotations with a random axis and the rotation angle as small as possible, often near 0. The method obviously finds the rotation as close as possible to the identity which is not a symmetry but would provide the best possible match. We discovered that this phenomenon is wholly suppressed if rotations with angles with an absolute value below 30° are forbidden.

While this is easy to ensure when creating the candidates (see later in Section 7.1.3) modifying the symmetry measure to enforce this in the optimization step (see later in Section 7.1.4) is slightly more difficult since we still want to keep it differentiable. Therefore, we use the symmetry measure in the following form:

$$s_X(Q, \mathbf{s}) = \left(\sum_{i=1}^n \sum_{j=1}^n \varphi(\|\mathbf{rot}(Q, \mathbf{s}, \mathbf{x}_i) - \mathbf{x}_j\|) \right) Pen(Q) \quad (7.4)$$

where $Pen(Q)$ is a function that penalizes angles with small absolute values. Let γ be the rotation angle, we need $Pen(Q) = 0$ for $|\gamma| < 30^\circ$ but we




Figure 7.2: The penalization function for $\cos(\beta)$ transiting through t_1 and t_2 , the function is even and continues to be 1 to the left (up until $\cos(\beta) = -t_1$) and 0 to the right.

need it to start decreasing at some larger angle, we choose 43° for this threshold, i.e. $Pen(Q) = 1$ for $|\gamma| > 43^\circ$ and for $|\gamma| \in \langle 30^\circ, 43^\circ \rangle$ it decreases with $|\gamma|$. The 43° value was chosen so that the method is still capable of detecting symmetries with rotation angle around 45° and the 2° difference is to provide some safe margin. We could define Pen for the angle γ but to avoid using trigonometric functions we instead define Pen around the cosine of $\beta = \frac{1}{2}\gamma$ which can be easily obtained from the quaternion Q as $\cos(\beta) = \frac{a}{\sqrt{a^2 + \mathbf{v} \cdot \mathbf{v}}}$. We set $t_1 = \cos(21.5^\circ)$, $t_2 = \cos(15^\circ)$. The rotation is by 2β for $\cos(\beta)$ so for $-\cos(\beta)$ it is by $360^\circ - 2\beta$, i.e. by 2β in the opposite direction. Therefore, Pen needs to satisfy $Pen(Q) = 1$ for $|\cos(\beta)| < t_1$, $Pen(Q) = 0$ for $|\cos(\beta)| > t_2$ and be decreasing with increasing $|\cos(\beta)|$ for $|\cos(\beta)| \in \langle t_1, t_2 \rangle$. Our similarity function $\hat{\varphi}(t)$ is differentiable and behaves in the needed way for $t \in \langle 0, 2.6 \rangle$ so we just need to transform $\langle t_1, t_2 \rangle$ to this interval. So we define

$$\hat{Pen}(t) = \begin{cases} 1 & t \in \langle -t_1, t_1 \rangle \\ \hat{\varphi}\left(\left(|t| - t_1\right)\frac{2.6}{t_2 - t_1}\right) & otherwise \end{cases}$$

which is differentiable for any t because $\hat{\varphi}(t)$ is differentiable, $\frac{d}{dt}\hat{\varphi}(0) = 0$ so the transition through t_1 and $-t_1$ to the constant 1 is differentiable as well and the constant between $-t_1$ and t_1 is differentiable trivially (see Fig. 7.2). Now, we simply define $Pen(Q) = \hat{Pen}\left(\frac{a}{\sqrt{a^2 + \mathbf{v} \cdot \mathbf{v}}}\right)$ which is differentiable w.r.t all four parameters of Q (again except for $a = b = c = d = 0$) and therefore the symmetry measure in Eq. 7.4 is differentiable as well. Fig. 7.3 shows how the Pen function changes with the rotation angle γ .

Of course, this basically prevents the method from detecting symmetries with rotation angles below 43° . However, for any rotational symmetry with the rotation angle $|\gamma| < 180^\circ$ there should be also significant symmetries with the same rotation axis and rotation angles with different k -multiples of γ , some larger then 43° . Therefore, in the case of symmetry with $|\gamma| < 43^\circ$




Figure 7.3: The penalization function w.r.t. the rotation angle $\gamma = 2\beta$, the function continues to be 1 to both sides up until $\gamma = 180^\circ / -180^\circ$.

the method should find a corresponding and similarly good symmetry with some rotation angle $k\gamma$, $k \in \mathbb{Z}$, $|k\gamma| \geq 43^\circ$, so we do not actually lose the symmetry. The corresponding lower angle symmetry could possibly be detected additionally by dividing the symmetry angle by different integers k and testing the resulting rotations but we leave this for future work.

7.1.3 Candidate Creation and Pruning

The first stage of our method consists of creating candidates for rotational symmetries. At the beginning we simply run steps 1, 2 and 3 of the original symmetry plane detection method (see Section 7.1.1) only with X_{simp} and α changed as described at the beginning of Section 7.1.2. Since we will use pairs of planes outputted by step 3 to create the rotation candidates and we need a decent amount of them, the default value of $S = 5$ is too small, so we increase it to $S = 30$ which proved to be more than enough in most cases. We therefore get a set of candidate planes $\{\mathbf{p}_i\}$, $i = 1, \dots, 30$. For each pair of planes \mathbf{p}_i , \mathbf{p}_j , $i \neq j$ we compute the line of their intersection and an angle $\gamma = 2 \arccos(\mathbf{n}_i \cdot \mathbf{n}_j)$ where \mathbf{n}_i , \mathbf{n}_j are unit normal vectors of planes \mathbf{p}_i , \mathbf{p}_j respectively. Now if $|\gamma| > 30^\circ$ we create two candidate rotations around the intersection line by angles γ and $-\gamma$. The point \mathbf{s} on the rotation axis is chosen so that it is the closest point on the axis to the centroid of the input object. From now on, we consider the rotation angle to be always positive where the rotations in the opposite direction are defined using the opposite axis direction.

Similar to the original symmetry plane detection method, we incorporate pruning into the candidate creation process to reduce the final candidate count, i.e., to save time in the candidate evaluation step, and to ensure that there are no too similar candidates which would result in the optimization (see the next section) being started multiple times from almost the same place. Checking the similarity of two rotation transformations could be done

using the theoretically appropriate distance function based on vertex sum of squares (see Section 4.2.2). However, it proved to be unnecessarily complex and computationally expensive mainly for the need to transform the candidates from quaternion-based to matrix-based representation. We discovered that the following much simpler approach provides almost identical results. When a new rotation candidate $\mathbf{c} = (Q, \mathbf{s})$ is created we set its weight $w(\mathbf{c}) = 1$ and find whether any previously created candidate $\mathbf{c}_i = (Q_i, \mathbf{s}_i)$ exists such that

$$\begin{aligned} \|Q - Q_i \operatorname{sgn}(Q \cdot Q_i)\| &< \epsilon_Q, \\ \|\mathbf{s} - \mathbf{s}_i\| &< \epsilon_s \end{aligned}$$

where $Q \cdot Q_i$ represents a quaternion dot product which corresponds to dotting the four-dimensional vectors with the same coordinates as the quaternion coefficients. If so, we merge the candidates into a new one $\mathbf{c}_{new} = (Q_{new}, \mathbf{s}_{new})$, $w(\mathbf{c}_{new}) = w(\mathbf{c}) + w(\mathbf{c}_i)$ where \mathbf{s}_{new} is a weighted average of \mathbf{s} and \mathbf{s}_i with $\frac{w(\mathbf{c})}{w(\mathbf{c})+w(\mathbf{c}_i)}$ and $\frac{w(\mathbf{c}_i)}{w(\mathbf{c})+w(\mathbf{c}_i)}$ as the corresponding weights. Similarly, Q_{new} is computed using the standard SLERP quaternion interpolation [38] between Q and Q_i with the same weights. We set the thresholds experimentally as $\epsilon_Q = 0.05$, $\epsilon_s = 0.05l_{avg}$, but we note that the method is not very sensitive to the exact value of either of them. We recenter \mathbf{s}_{new} such that it is again the closest point on the axis of \mathbf{c}_{new} to the centroid and the same process is then performed for \mathbf{c}_{new} to see if it should be merged with some other candidate.

The number of resulting candidates is usually in the lower or mid-hundreds. We compute the symmetry measure from Eq. 7.4 with X_{simp} for each one of the candidates. The calculation of $s_{X_{simp}}$ is accelerated in the same way as in Chapter 6 by exploiting the locality of φ and using a 3D grid to find the points that could be in its support region. Finally, we select M candidates with the largest measure where $M = 5$ by default.

7.1.4 Optimization

As the final step, local optimization using the L-BFGS [74] method is started from the M candidates outputted by the previous step to find M local maxima of the symmetry measure $s_{X_{simp}}$ from Eq. 7.4. Computation of the gradient of $s_{X_{simp}}$ is accelerated using the 3D grid in the same way as the computation of $s_{X_{simp}}$ itself. The one local maximum with the largest measure is then selected and can be declared the resulting rotational symmetry of the input point set X and considered the output of our method. The other local maxima can potentially be used as secondary symmetries.




Figure 7.4: The process of obtaining the rotational symmetry from the reflectional symmetry candidates, (a) shows the original object, (b) and (c) depict the candidate symmetry planes with the object colored according to the strength of the symmetry (the lighter the color, the stronger the symmetry in the given point), (d) shows the initial rotational symmetry candidate created by the two planes, the object is now colored according to the rotational symmetry, (e) shows the final rotational symmetry acquired by the optimization started from the initial one, the rotation angles for (d) and (e) are in the captions.

Fig. 7.4 demonstrates the overall process of acquiring the final rotational symmetry from the initial symmetry plane candidates. Fig. 7.4a shows the original object; note that it is strongly but not perfectly symmetrical. For visualization simplicity, the object is represented by a triangle mesh but we used only its vertices as the input point set. Figures 7.4b and 7.4c show the two initial symmetry plane candidates respectively - the object is shown from two different viewpoints (side view, bottom-top view) in such a way that the candidate plane (marked by the red line) is perpendicular to the plane of the figure. The object is colored so that the points that are symmetric w.r.t. to the plane are light (the lighter the color, the stronger the symmetry) while the others are dark. Notice that although there is non-negligible reflectional symmetry w.r.t. both these planes, neither of them is very precise. Fig. 7.4d shows the initial candidate rotation created by the two planes. The coloring now shows symmetry w.r.t. to the rotation transformation, the rotation axis is shown in the figure and the angle is in the caption. Obviously, the symmetry is not very good - the axis is tilted

and the angle is more than 3° off (60° is expected for this object). Finally, Fig. 7.4e depicts the resulting rotational symmetry outputted by the numerical optimization started from the initial one. We note that the candidate symmetry planes are often more precise than those in Figures 7.4b, 7.4c and that this scenario is not entirely common. However, it demonstrates how robustly the symmetry measure, combined with the optimization approach, extends the core idea of combining candidates for reflectional symmetries into a solution for detecting rotational symmetries. From two rather weak reflectional symmetries we acquired an imprecise rotational symmetry candidate, which was then turned into a very strong symmetry by optimizing (maximizing) the symmetry measure in the space of rotations, which can be done quite easily using the proposed quaternion-based rotation parameterization. This robustness is very useful if the input object has rotational but not reflectional symmetries, as detailed later in Section 7.2.

The method attempts to find the best possible rotational symmetry. Suppose the input object has multiple similarly significant symmetries with the same axis and different k -multiples of (almost) the same angle. In that case, the method chooses mostly randomly which one is the best. For example, if there is a strong symmetry with rotation angle 60° , the resulting symmetry would have the corresponding axis and the rotation angle will basically be a random choice between 60° , 120° and 180° . The method simply chooses the k -multiple for which there is the best match according to the symmetry measure, even if the difference is marginal compared to the other k -multiples. However, from the user perspective, the smallest possible k -multiple might be preferred in these scenarios. We observe that such rotation is very often present among the other four detected local maxima of the symmetry measure. Therefore, our method has one more optional step which we execute by default. We simply check all the four other detected secondary rotations and if any of them has almost the same symmetry measure as the best detected rotation (we allow at most 1% difference) but has a smaller rotation angle, we chose it as the output instead. If multiple rotations satisfy this condition, we choose the one with the smallest rotation angle. Even with this step the method cannot find rotations with angles below 43° (see Section 7.1.2). So if the smallest k -multiple of the angle of the best symmetry is, e.g., 30° the method will most likely output its 60° equivalent.

In the case of perfect symmetry, the rotation angle is always $\frac{360^\circ}{k}$, $k \in \mathbb{Z}$ and, therefore, the resulting angle of the detected rotation could always be rounded to the nearest such value. However, this is not the case with approximate/weak symmetries. Since we do not seek the rotation axis and angle separately but rather search for a compound rotation, changing the

angle alone would also require finding a new axis that provides optimal symmetry with the new angle. Even then, the new rotation with the rounded angle generally does not have to provide a better symmetric match than the original rotation for non-perfectly symmetrical objects. This is why we leave the resulting angle as outputted by the optimization as part of the resulting rotation but the rounding step can be very easily implemented if preferred by the user.

7.1.5 Identifying Circular Symmetry

Our method cannot straightforwardly detect circular symmetries because they cannot be represented by specific geometric transformations and the symmetry measure cannot be applied for them. Therefore, if the input object possesses circular symmetry, our method will identify its axis and the rotation angle will be mostly arbitrary. However, we perform a simple check to identify whether the outputted symmetry could be contained in a potential circular symmetry. We take the axis of the outputted rotation and generate rotations with different angles γ around this very axis. We use all prime numbers between 43° and 180° as the values of γ so that there are no reoccurring multiples of some smaller angles. We compute the symmetry measure for all these rotations, and if their average is at least $\frac{2}{3}$ of the symmetry measure of the outputted rotational symmetry, we report potentially significant circular symmetry around the axis of the outputted rotation. The $\frac{2}{3}$ threshold might seem very low but note that we aim our method also for objects with missing parts and in such cases, the average measure would be noticeably lowered because large areas can be misaligned for some angles.

7.2 Results

The proposed method was implemented in C#, all the experiments were performed on a computer with CPU Intel® Core i7-10700F (frequency: 2.9GHz, turbo boost: 4.8GHz, 8 cores, L1 cache: 512kB, L2 cache: 2MB, L3 cache: 16MB) and 32GB memory (frequency: 3.2GHz) running a Windows 10 operating system. Figure 7.5 shows the results of our method on several objects from various different datasets [30, 34, 59, 128]. For each object, there is a side view with the detected axis and a top-down view with the axis in the center and the red arc depicting the rotation angle of the detected symmetry. The angle is also in the caption, sometimes followed by a circle which means that the method indicated potential circular symmetry for the given

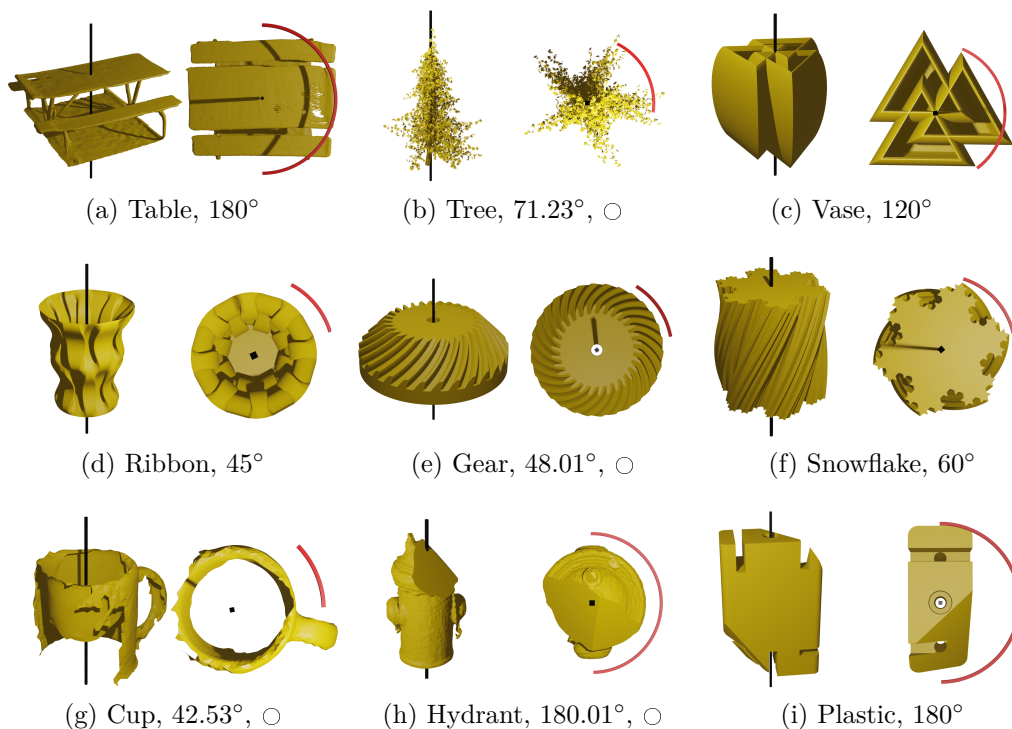


Figure 7.5: Results of the proposed method on several objects with both strong and weak symmetries and with missing parts, each figure shows a side view with the detected axis, a top-down view with the axis in the center and a red arc depicting the detected rotation angle which is also in the caption, the circles indicate detected possible circular symmetry.

object. All the objects are represented by triangle meshes because this is the most common object representation in computer graphics and it eases visualization. However, we only take the vertices as the input point set for our method. Furthermore, some of these objects are 3D scans of real objects and were probably created by triangulating raw point clouds (for the Cup we are certain of this) so using their vertices is the same as running the method on the original point cloud. Since the proposed method is not entirely rotationally invariant (mainly due to the grid-based simplification) all the objects were randomly rotated before starting the symmetry detection.

There are several objects with strong symmetries, some with weaker symmetries and even some incomplete objects (the bottom row). The Cup was incomplete originally, for Plastic and Hydrant the missing parts were removed artificially by clipping. For the Cup and Gear the detected angle is rather arbitrary because there is a quite strong circular symmetry in both cases so other angles would be plausible as long as the axis is the same. The

similar applies for the Hydrant where the approximate 180° angle corresponds to the strongest symmetry but different angles with the given axis would still provide pleasing approximate symmetries. In fact, for different random rotations of the Hydrant our method sometimes detects the symmetry with a different arbitrary angle, still providing a plausible symmetry just not the best one possible. This can, however, be easily mitigated by increasing the M parameter, determining the number of best candidates selected for the optimization, from the default 5 to approximately 10 or more. For the Tree object the angle of the detected symmetry is 71.23° , which seems correct since an angle around 72° can be expected, and circular symmetry is indicated. This is due to its weak symmetry and overall conic shape which might not be obvious from the top-down view but is quite clear from the side view.

Notice that some of the objects do not have reflectional symmetries (Vase, Snowflake and somewhat even Gear) but the proposed method still finds a very good rotational symmetry for all of them. This is because the symmetry plane candidates that approximately intersect in the axis of rotational symmetry still have a noticeably larger symmetry measure than other arbitrary planes and they are more than likely to be among the best planes passed to the rotation creation step. Let us demonstrate this on a very simple shape depicted in the top row of Fig. 7.6. Fig. 7.6a shows the detected rotational symmetry axis for this object, the detected angle is 180° . Figures 7.6b and 7.6c show the two planes that were used to create the initial candidate whose optimization lead to the final symmetry, whereas Fig. 7.6d shows a random plane passing through the object, the symmetry coloring is the same as in Fig. 7.4. Unlike the random plane, the first two planes capture symmetry in a rather large portion of the object's points, although its whole shape actually does not have a reflectional symmetry, which explains why they end up among the best symmetry plane candidates. Of course, the rotational symmetry created by such planes in similar cases cannot be expected to be very precise but this is solved by the robustness of the optimization step (see Section 7.1.4).

The bottom row of Fig. 7.6 provides the same demonstration for an S-shaped object with even less reflectional symmetry showing that the proposed method can truly find even rotational symmetries which are not created by significant reflectional symmetries. However, it should be noted that the probability of failure is naturally larger in such cases because the chance of *hitting* the right rotation candidate is more driven by random influence than by the reflectional symmetry of the object. Increasing the values of some parameters such as S (number of selected reflections) and M (rota-

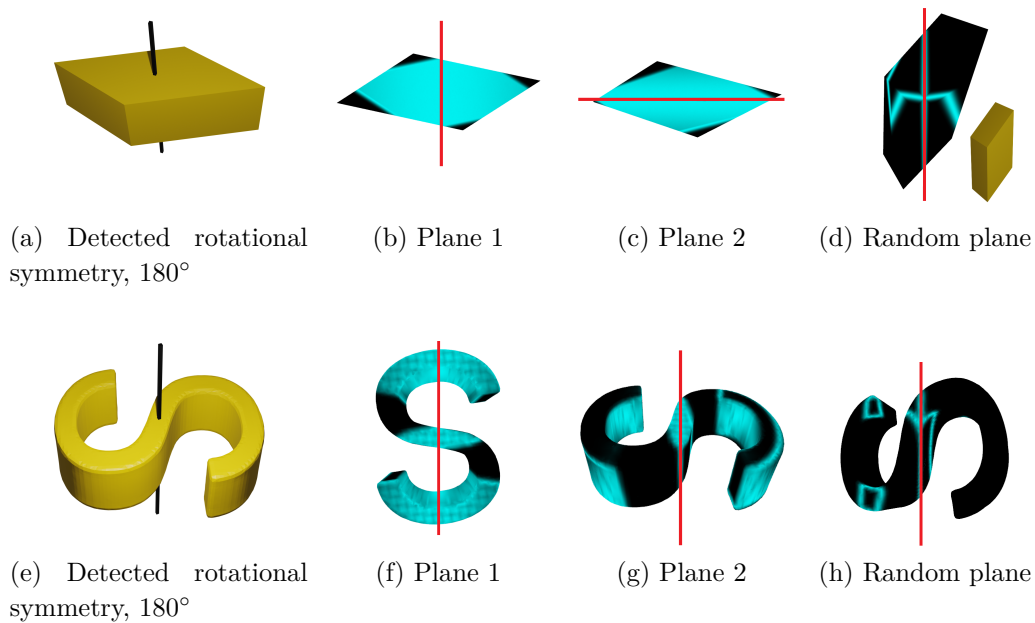


Figure 7.6: Two different shapes without reflectional symmetries, the first figure in each row shows the detected rotational symmetry, the second and third show the two symmetry plane candidates, which created the initial rotation, with symmetry coloring (the objects are oriented so that the planes are orthogonal to the figure plane) and the last figure depicts a random plane with symmetry coloring (the small yellow image in the top row provides perspective).

tions) can generally improve the chance of success but for the cost of longer computation time.

Table 7.1 demonstrates the computational efficiency of the proposed method. As some parts of the method can be easily parallelized, we show the computation time of both single-thread (s.t.) and multi-thread (m.t.) implementation. The running time obviously does not depend much on the object's point count, which is ensured by the fast simplification that is also included in the measurements. The single-thread implementation runs mostly under 3 seconds but since nowadays multi-core CPUs are standard, we are confident that our method can mostly run under or around 1 second, as suggested by the last column.

For some objects, other significant symmetries can be found by looking at the remaining local maxima of the symmetry measure outputted by the optimization step. Fig. 7.7 shows a dodecahedron-like object with the first three detected rotational symmetries. The fourth symmetry has a similar

Table 7.1: Computation times of the proposed method, first column shows the point counts of the objects, second column shows the computation times of the single-thread and the last column of the multi-thread implementation, times include simplification.

| Object | Point count | Time s.t. [ms] | Time m.t. [ms] |
|-----------|-------------|----------------|----------------|
| Plastic | 617 | 492 | 209 |
| Vase | 3115 | 2459 | 1117 |
| Snowflake | 3264 | 2109 | 490 |
| Cup | 4104 | 2074 | 542 |
| Gear | 14566 | 3075 | 858 |
| Ribbon | 16160 | 1814 | 494 |
| Hydrant | 72012 | 2467 | 644 |
| Tree | 134194 | 1380 | 467 |
| Table | 151026 | 2227 | 683 |




Figure 7.7: The best three detected rotational symmetries of a dodecahedron-like objects.

angle to the second but different axis (however, the images would be almost identical) and the fifth is similar to the third. By increasing M and optimizing more candidates we could possibly obtain even more different symmetries.

To demonstrate the robustness of our method to noise we added the vector $[rand_x, rand_y, rand_z]^T \cdot l_{avg} \cdot mag$ to every point of the input object

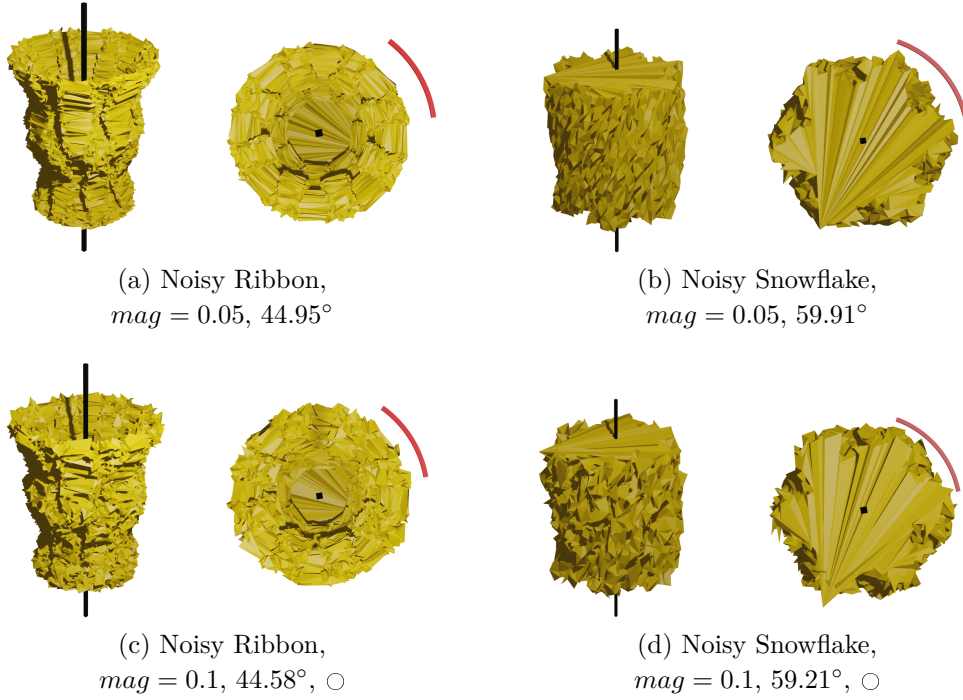


Figure 7.8: Ribbon and Snowflake with added uniform noise, in (a,b) with $mag = 0.05$ and in (c,d) with $mag = 0.1$.

where $rand_x, rand_y, rand_z$ are uniform random values from $\langle -1, 1 \rangle$, l_{avg} is the average distance of points of the object from their centroid and mag determines the noise magnitude. Fig. 7.8 depicts the detected symmetries for the Ribbon and Snowflake objects with $mag = 0.05$ (top) and $mag = 0.1$ (bottom). For $mag = 0.1$ the method indicates circular symmetries because the noise noticeably decreases the precision of the present symmetry so the differences between symmetries with different angles are much smaller, presuming the axis is correct. Still, the detected symmetries are very good for all these noisy objects.

The chart in Fig. 7.9 shows how the noise magnitude influences the precision of the symmetry detection. We used four different objects (Snowflake, Ribbon, Vase and Gear) and we normalized them in size such that $l_{avg} = 1$. For each object we added noise of varying magnitude and applied our method to find rotational symmetry. Then we used the resulting rotation to transform the original noise-free object and measured its distance from its non-transformed version using the Metro mean error [24] similarly as we did in Section 6.3 or as done in [69]. Due to the normalization, the units of both the error and noise magnitude are percents of l_{avg} . The error starts growing noticeably only after the noise magnitude reaches roughly 5%,

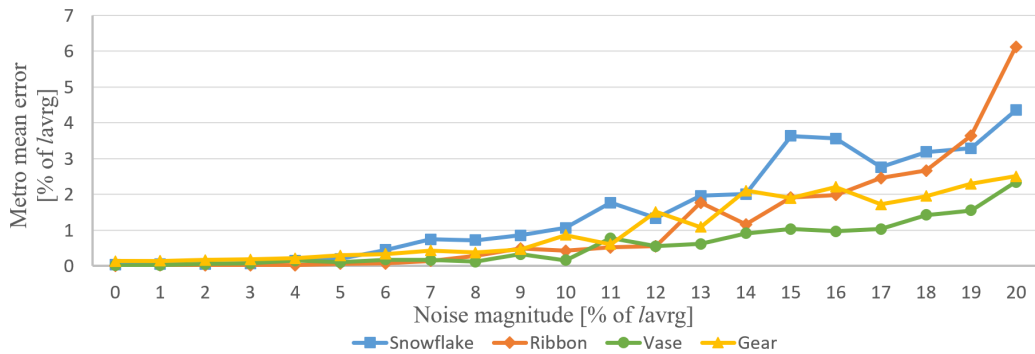


Figure 7.9: The influence of noise magnitude on the precision of our method measured by Metro mean error [24].

which corresponds to the top of Fig. 7.8, and a significant increase in error starts only after roughly 10%, which corresponds to the bottom of Fig. 7.8. Overall, the error seems to grow rather slowly with the noise magnitude, confirming the robustness of the proposed method to noise.

7.2.1 Comparisons

As already mentioned, not many methods for rotational symmetry detection exist. Moreover, their implementation is usually not publically available, at least not to our knowledge, and no benchmark datasets with ground truth rotational symmetries seem to exist. More importantly, most of the methods aim at some specific type of data either in terms of required shape, object representation or some other properties. Therefore, providing any qualitative or quantitative comparison to the existing methods would be very difficult on its own, let alone providing an objective one. Nevertheless, we at least show an experimental comparison to a simple yet surprisingly well-performing PCA (Principal Component Analysis [124]) method and we provide a side-by-side comparison to the existing methods based on information available in the corresponding papers.

The PCA algorithm outputs three principal directions. Provided the input point set has strong rotational symmetry, one of the three directions should be the direction of its symmetry axis [115] that passes through its centroid. We manually selected the one of the three possible axes that seemed to be the closest to a plausible symmetry axis. For the strongly symmetrical shapes (Ribbon, Gear, Snowflake, Vase or even Table) this approach provides a surprisingly accurate axis of rotational symmetry. However, when used on weakly symmetrical objects, this method failed to detect a plausible axis or it was at least noticeably inaccurate, as can be seen in Fig.




Figure 7.10: Results of the PCA method, the best of the three possible axes was chosen manually, the small yellow images show the corresponding results of our method from Fig. 7.5 for easier comparison.

7.10. This renders any PCA-based method practically unusable for detecting weaker symmetries, at least without aggressive refinement. Also, PCA obviously cannot find the rotation angle, and in practice, some algorithm would be needed to select the best of the three possible axes automatically.

We selected 100 random objects from the Thingi10k dataset [128] (10000 objects represented by triangle meshes of varying shapes and properties), we removed all objects that obviously did not exhibit any rotational symmetry and replaced them with new random objects and we kept repeating this until we acquired 100 random objects with at least some form of rotational symmetry. Then we used both our method and PCA to find rotational symmetry for each of these 100 objects. As mentioned in Chapter 6 the symmetry measure is not very robust to strongly non-uniform point distribution (e.g., typical CAD models) and PCA has the same problem. However, since the objects are meshes, this can be easily solved by some surface sampling technique. As the Thingi10k dataset contains quite a few such problematic objects we first used MeshLab’s [23] stratified triangle sampling to resample the objects and we applied both symmetry detection methods on the resampled point sets. We manually observed the detected symmetries and determined their correctness as objectively as possible, the results are reported in Table 7.2. If the resulting symmetry axis was plausible (one with a significant symmetry around it, not necessarily the best possible one for the object), it was marked as correct. If it was close to a plausible axis but with noticeable imprecision, it was marked as imprecise and in other cases it was marked as incorrect. The angle was marked as correct if the axis was correct or imprecise and it was within 1° of a correct angle w.r.t. the given axis or in case of circular symmetry. If the deviation was greater than 1° but up to 5° , it was marked as imprecise, otherwise incorrect (in case of incorrect axis

Table 7.2: Comparison of the proposed method to PCA on 100 randomly selected resampled objects with some rotational symmetry.

| | Axis correct | Axis imprecise | Axis incorrect | Angle correct | Angle imprecise | Angle incorrect |
|----------|--------------|----------------|----------------|---------------|-----------------|-----------------|
| Proposed | 92 | 5 | 3 | 93 | 2 | 5 |
| PCA | 78 | 15 | 7 | - | - | - |

Table 7.3: Side-by-side comparison of existing methods for rotational symmetry detection in 3D including the proposed method.

| Method | General axis | Can detect angle | Can handle weak symmetry or missing parts | Input object representation | Supported shape |
|--------------------------------|--------------|------------------|---|-----------------------------|-----------------|
| Proposed | Yes | Yes | Yes | Point set | Any |
| Sun and Sherrah 1997 [115] | No | Yes | No | Any with normals | Any |
| Martinet et al. 2006 [77] | No | Yes | No | Surface | Any |
| Korman et al. 2015 [60] | No | Yes | No | Binary volume | Any |
| Sipiran 2017 [110] | Yes | No | Yes | Mesh | Rotational |
| Sipiran 2018 [111] | Yes | No | Yes | Mesh | Rotational |
| Ecins et al. 2018 [31] | Yes | No | Yes | Point set with normals | Rotational |
| Gothandaraman et al. 2020 [39] | No | Yes | No | Point set | Any |
| PCA | No | No | No | Point set | Any |

the angle was marked as incorrect automatically). Note that, since PCA always finds three axes and we selected the best one of them manually for each object, the chance of at least one of them being imprecise at worst is quite large, which explains its rather low number of incorrect axis detections (still larger than our method nevertheless). However, the number of correct ones is noticeably larger for our method, which further has the advantage of also finding the angle of the rotational symmetry.

Table 7.3 shows side-by-side theoretical comparison of the methods for detecting rotational symmetries in 3D mentioned in Section 3.2 including ours and PCA. The first column indicates whether the given method can detect symmetry with a general axis or only with an axis passing through some reference points which severely limits applicability for detecting weaker symmetries. The second column indicates whether the method also finds the rotation angle of the symmetry while the third column states if the method

can handle weakly symmetrical objects or missing parts. The fourth column says what object representation is required on the input where point set is the most general one while a triangle mesh or other surface representations are the least general ones and binary volume stands for any representation where it can be determined whether a point lies inside or outside the input shape. At last, the fifth column states what type of shape the method requires to work properly. Note, however, that we did not experimentally verify any of this information (except for PCA and our method) and we state what we derived from the corresponding publications. The proposed method appears to be the most general one of all, detecting symmetry with general axis plus the rotation angle, handles well objects with weak symmetries and missing parts, only requires a set of points on the input and does not put any constraints on the shape of the input object.

7.3 Limitations

The proposed method also has some limitations. First of all, it obviously depends on the quality and properties of the selected underlying symmetry plane detection approach, in our case the method from Chapter 6. As already mentioned, the symmetry measure is not suitable for objects with very non-uniformly distributed points and the method can be expected to fail on such data. However, in the case of triangle mesh or other surface representation, it can easily be resolved using some surface sampling technique. Also, our method might be rather fast but still significantly slower than the original symmetry plane detection algorithm from Chapter 6. This is primarily due to the increased target point count of X_{simp} from 1000 to 3000. Therefore, in the future it would be more than appropriate to find a way of keeping the method sufficiently robust without this change. At last, it can be expected that the method will fail with non-negligible probability in cases when the rotational symmetry is not composed of reflectional ones. But, while this might be true theoretically, it turns out that in practice the method works even in these cases as we showed and explained throughout the previous sections.

7.4 Summary

We proposed a new method for detecting rotational symmetries in 3D objects using a simple idea of combining potential reflectional symmetries and modifying the optimization step from Chapter 6 for rotations with a quaternion-

based parameterization which seems to be novel in the symmetry detection field. We showed that the method is well capable of detecting both strong and weak symmetries including partial symmetries for objects with missing parts and is robust to noise as well as computationally efficient. The robustness and precision of the proposed method could potentially be improved even more by adjusting some of its parameters, such as the target point count of the simplified point set or the number of best-selected rotation candidates. This would, however, come with the cost of larger computation time and as the results suggest, it is unnecessary in most cases. These adjustments can be used in the opposite way as well, i.e. if precision and robustness are not a priority, the method can be made faster.

Furthermore, to our knowledge, the proposed method is the most general one of the existing methods for rotational symmetry detection. Also, the main theoretical contributions of this chapter could be used in other context as well and possibly combined with different approaches for symmetry plane detection. The implementation of our method can be downloaded from [52].

The parameters of the proposed method are overall the same as those of the underlying symmetry plane detection method and their default values have been introduced throughout the paper. The meaning and effect of each parameter have already been described rather thoroughly in Section 6.4. The only additional parameters in the proposed method are ϵ_Q , ϵ_s (see Section 7.1.3), which have the same effect on the candidate rotations as the δ parameter (see Section 7.1.1) has on candidate planes in the original method, and M (see Sections 7.1.3 and 7.1.4) which has the same effect as S (see Section 7.1.1) in the original method.

As described in Chapter 6, the symmetry measure allows incorporating weights that can be used to adjust importance of given point pairs based on some additional information, such as similarity of some local descriptors. In the future, it would be meaningful to examine this possibility also in the context of our rotational symmetry detection method. Also, it might be possible to extend our core idea to find symmetries w.r.t. isometries which can be composed of three reflections. We leave this for future work as well.

8 Additional Contributions

In this chapter we provide an overview of some of our additional, mostly smaller, contributions that we do not believe to be significant enough to be contained in their own chapters but are still worth mentioning. These contributions are often in the form of cases where other researchers used our previously described methods or discoveries in their own work or proceeded to extend/modify them in some way, sometimes in partial cooperation with the author of this thesis.

8.1 Local Reflectional Symmetry Detection

Building on our previous work from Chapters 4 and 6 and the work of Mitra et al. [83] we designed a method for finding local reflectional symmetries in 3D point clouds. The content of this section was previously published in [51]. We only describe the method briefly and show some of its results.

We proposed a very simple local shape descriptor for point cloud data based on the Principal Component Analysis [124] (PCA) which can be used for symmetric point matching. It stems from taking points in a neighborhood of fixed size around a given point and computing the PCA of the points contained in this neighborhood. The PCA provides three orthogonal eigenvectors that correspond to the principal directions of the points in the neighborhood, and three eigenvalues that correspond to the variances of the points in the directions of the given eigenvectors.

The eigenvalues provide a compact information about the shape of the local neighborhood that is invariant under rotation, translation and, if considered in absolute values, reflection. Since it is expected that two points with similar shapes of their neighborhoods will have similar absolute values of the eigenvalues provided by the PCA descriptor, the eigenvalues can be used for finding potentially symmetric pairs of points. When representing the set of absolute values of the three eigenvalues by vectors in E^3 , finding the potentially symmetric pairs of points (points with similar values of these 3D vectors) can be done efficiently using any Euclidean data structure - we use a uniform 3D grid.

Any pair of points matched in this way provides evidence of some symmetry but this symmetry does not have to be reflectional. However, the existence of potential reflectional symmetry can be verified by checking existence of reflectional symmetry between one or more of the corresponding




Figure 8.1: A scene with a reflectionally symmetric couch in the middle colored using the eigenvalues of the PCA descriptor as RGB values (data source: [34]).

eigenvectors of the two neighborhoods w.r.t. the plane which is constructed as a symmetry plane of the given pair of points - if the two neighborhoods are symmetric w.r.t. the plane the eigenvectors of the neighborhoods will be symmetric as well. The pairs that do not pass this check are discarded and the others are kept as potentially reflectionally symmetric pairs of points and the symmetry plane of each of these pairs represents a symmetry candidate. The set of all candidate planes is kept for further processing. The symmetry pairing is combined with point cloud simplification (the same approach as described in Section 6.1.3) to ensure smaller computation cost. Also, before the pairing process is performed, we first find and remove points whose neighborhood is roughly planar because planar neighborhoods add too much potential for false positive matches.

Fig. 8.1 shows a 3D scene where each point is colored using the eigenvalues of the PCA shape descriptor - the RGB values of the colors are set using the absolute values of the eigenvalues and it can be noticed that the symmetric counterparts of the couch have noticeably similar colors. Fig. 8.2 depicts the same scene and the pairs of points that were matched as potentially symmetric and passed the reflectional symmetry check - they are connected by the black lines. The points are still colored using the PCA eigenvalues but the grey color now represents the points that were removed as planar.

We follow the theory, described at the beginning of Chapter 5 and used, e.g., by [83] that the significant symmetries should now form modes in the




Figure 8.2: A scene with the PCA-based descriptor symmetric matching - the symmetrically matched pairs of points are connected by the black lines, the coloring corresponds to the descriptor eigenvalues, the grey color indicates planar areas.

space of candidate transformations - planes in our case. To find the modes we use the density peak location algorithm we introduced in Section 4.1.3 modified to find multiple modes instead of just the strongest one. After finding the most significant density peak, we remove all candidates that contributed to its density value and run the location algorithm again. This is repeated multiple times until the density at the detected peak is not significant enough. As the distance function for planes in the density peak location we use the D_{ED} function from Section 5.2.2. Because the distance function is Euclidean, we use a 4D uniform grid as accelerating data structure in the same way as used in Section 6.2.1 instead of the *Vantage Point Tree* that was originally used in Section 4.1.3.

Each of the detected modes now represents a possible local symmetry plane. For every one of these planes we extract the set of points that are approximately symmetric w.r.t. that plane and use them to represent the symmetric region of the plane, i.e. we have a set of possibly reflectionally symmetric regions (can be overlapping) and for each region we have one plane which is an estimate of its symmetry plane. The symmetry plane estimate for a given region might not be accurate after only the density peak estimation. Therefore, we now treat each of the regions as globally symmetric and use the symmetry measure from Section 6.1 and the optimization step from Section 6.1.4 initialized with the symmetry plane estimate of the given region and with the region points as the input point set. After acquiring the results of the optimization for all the regions we further discard symmetries whose symmetry measure is too small which suggests that they are probably results of false positive matching. The remaining optimized planes now represent

the final detected reflectional symmetries.

Fig. 8.3 shows several results of the method for scenes from the dataset [34]. The scenes are represented by triangle meshes but we only used the vertices as the input point set, the black rectangles in each scene represent all the detected symmetries the method found. The caption of each result states the point count of the scene and the computation time of the method. In the last scene (Fig. 8.3f), there are some smaller objects with reflectional symmetries which the method did not detect suggesting that it has its limitations. On the other scenes in the figure the method provided satisfactory results. The method also appears to be fairly fast and, because there is still potential for further optimizations (including multi-threaded implementation), it could possibly be even much faster in the future. For detailed description of the method, its parameter setting, discussion of its limitations and for more results we refer to [51].

8.2 Uses of Registration and Transformation Distance Metrics

In this section we mention some cases where the product of our research described in Chapter 4 appeared in the work of other researchers.

The registration method described in this chapter, using the Triangle Sum of Squares metric (see Section 4.2.2) for transformations in the density peak estimation step, was extended by Arvanitis et al. [6] to work on noisy point clouds. The triangles for the metric are defined using closest neighbors of points. They then used this extended version of the method in their own approach for identifying objects in partially scanned and cluttered point clouds. The registration method is used to align each model from a set of target objects to a given input scene.

Dvořák et al. [28] used the Vertex Sum of Squares transformation metric, analyzed thoroughly in Section 4.2.2, in their tracking method for time-varying surfaces. At last, Miroslav Levora, when working on his bachelor's thesis [66], reimplemented the entire registration algorithm in the C programming language (the original implementation was in C#) using various optimizations. The new implementation was approximately three times faster than the original one.

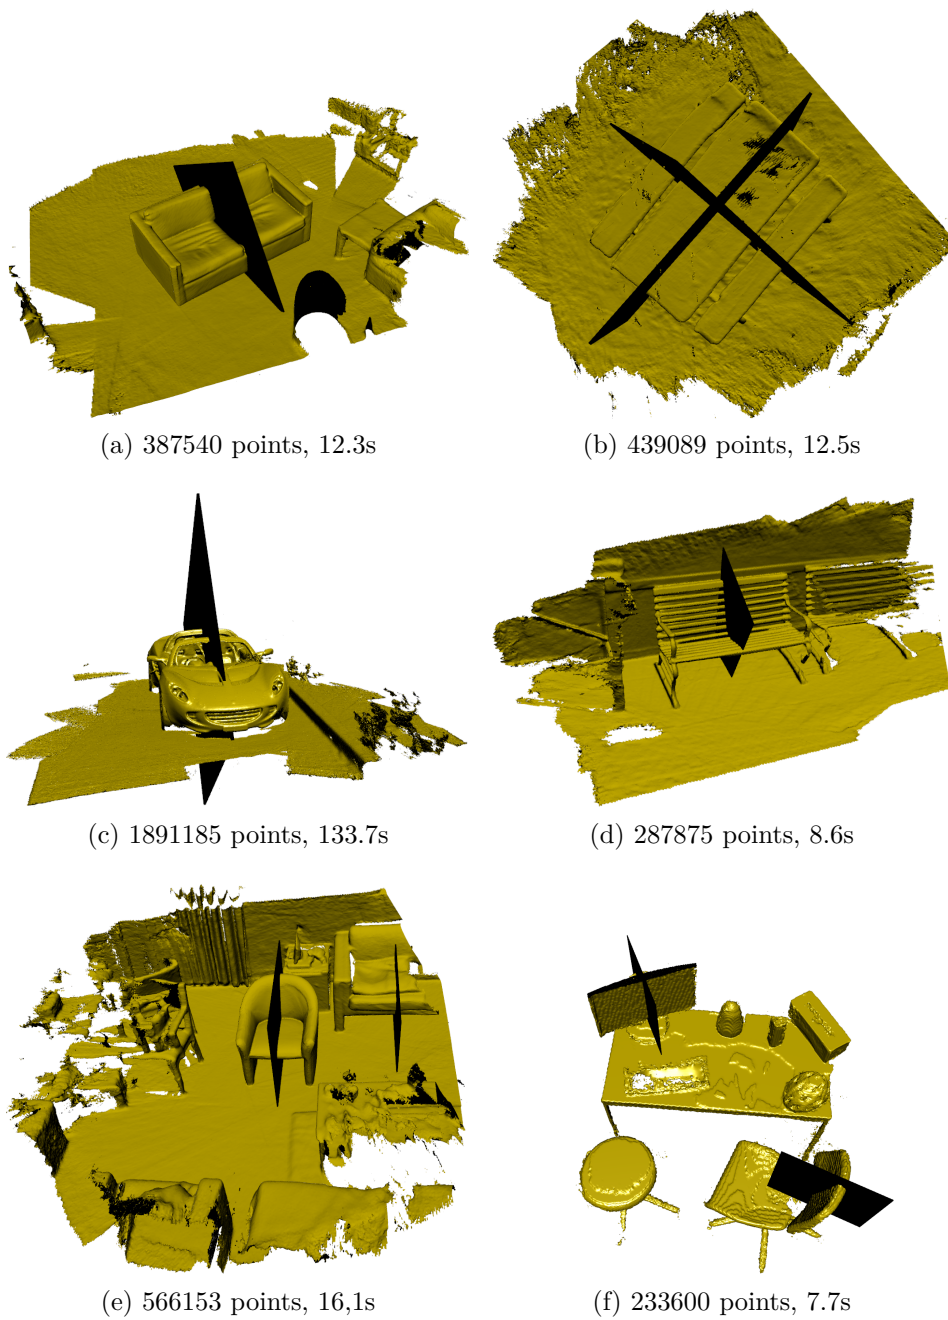


Figure 8.3: The results of our local symmetry plane detection method on several 3D scanned scenes from the [34] dataset, the detected symmetry planes are marked by the black rectangles, the captions state the point counts of the scenes and computation times.

8.3 Uses of Symmetry Plane Detection

Here we mention several cases where the material from Chapter 6 was used or extended/modified by others.

A researcher named Renan Sales Barros from the University of Amsterdam contacted us in 2019 regarding the implementation of our symmetry plane detection method. He was developing a method for detecting subtle brain hemorrhages in CT-scanned human heads. He needed to do alignment of the heads and wanted to use the symmetry information for this purpose. Since the scan often does not contain the entire head he needed a robust symmetry plane detection method that would work on objects with some missing parts and asked us for the implementation of our method. As the content of Chapter 6 was not published at the time, we provided him the implementation of an older version of the method which was consistent with the one described in the author's Master's thesis [43]. He sent us four images of the results he acquired using our method on his data. These results are shown in Fig. 8.4, notice that the head in the rightmost image at the top is missing its whole face.

Roughly a year after he informed us that he did not use our method after all but was still interested in our symmetry detection research and recent improvements.

Eliška Mouryová [85] in her Bachelor's thesis managed to modify our symmetry plane detection method, with good results, for detecting symmetries in planar (2D) curves. The author of this doctoral thesis was consulting/cooperating on this research. Fig. 8.5 shows the symmetry axes detected by her method on three planar curves (top row) and on the same data damaged by noise (bottom row). The red points represent the input data, the black lines are the detected symmetry axes.

Eliška Mouryová is also currently working, in cooperation with the author of this thesis, on another modification of our method which aims to find the spherical surface of reflectional symmetry instead of a plane of reflectional symmetry. Apart from some other modifications, the transformation of over-plane reflection was changed in the method to a transformation that reflects points over a given spherical surface. The goal is to find a spherical surface over which the input 3D point set would reflect to approximately fit onto itself. So far, there are only preliminary results for spheres with fixed radius but they look promising. These results are shown in Fig. 8.6, the yellow points represent the input point set and red points demonstrate the detected spherical surface of symmetry.

Štěpánka Krutinová, in her Bachelor's thesis [63], attempted to use our

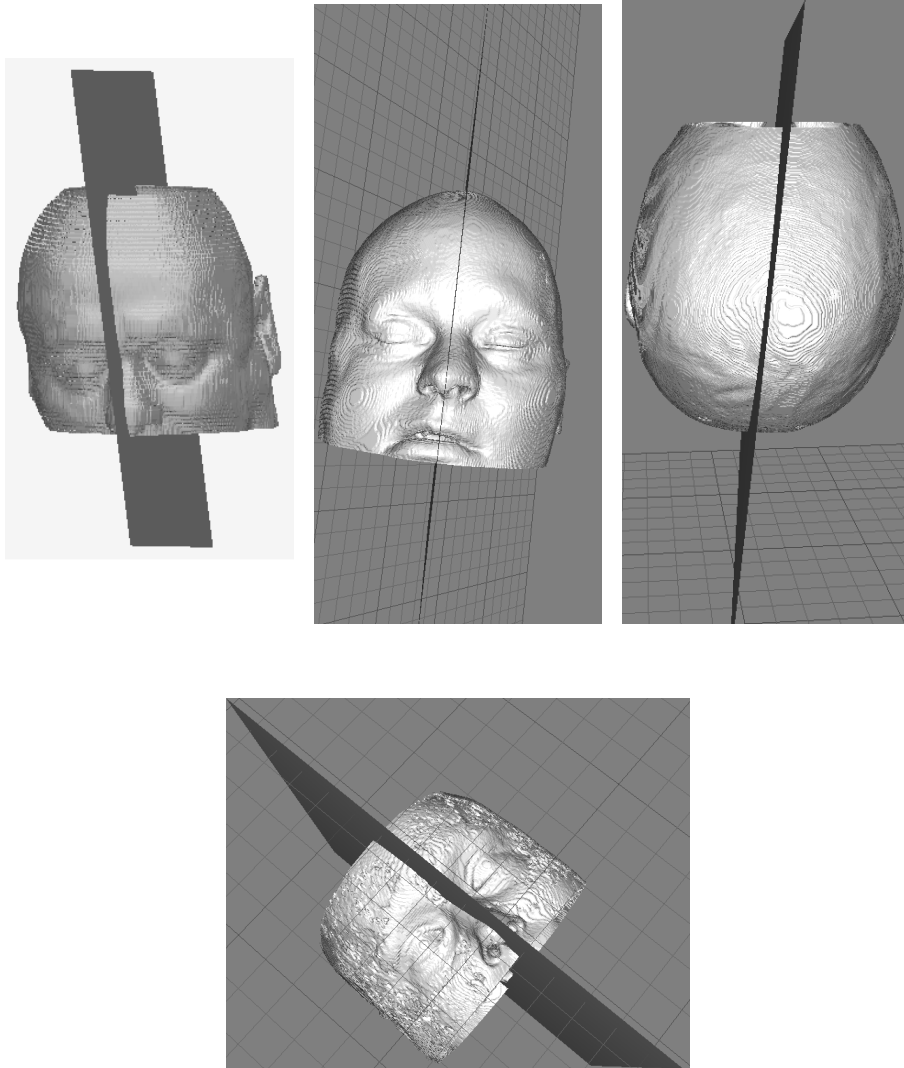


Figure 8.4: Results of an older version of our symmetry plane detection method on four CT-scanned human heads (figures provided by Renan Sales Barros).




Figure 8.5: Symmetry axes detected on 2D objects using the method described in [85] which is a modification of our symmetry plane detection algorithm (Figures taken from [85]).




Figure 8.6: Preliminary results of symmetry sphere detection (figures provided by Eliška Mourycová).

symmetry plane detection method to find symmetries of interesting parts of large scenes acquired using the Earth observation techniques. The author of this doctoral thesis was consulting on this research. It appeared that the method only provides satisfactory results when the interesting points of the scenes were selected manually. However, some promising results were acquired when the interesting parts were selected using Laplacian filtering, see Fig. 8.7. The image on the left shows a scene with the manually selected points colored red and the symmetry plane (the black line) detected on the selected points. These points probably represent a road or a street. The image on the right shows only the points of this scene that were selected automatically by the Laplacian filtering and their detected symmetry plane.




Figure 8.7: Left - a scene with manually selected points (red) and their symmetry plane (black line), right - points of the scene selected by Laplacian filtering and their symmetry plane (figures taken from [63]).




Figure 8.8: Symmetry planes found when assigning larger weights to the selected (red) areas (figures taken from [5]).

Ondřej Anděl, in his Master’s thesis [5], also used our method for finding symmetries in the Earth observation data, this time by utilizing the weights in the symmetry measure (see Section 6.1). A given area of the scene is selected and pairs of points in this area are assigned greater weights than the rest of the points. Selection of these areas is done in different ways including manual selection or k-means clustering. Fig. 8.8 shows a scene with three different areas selected as clusters by the k-means clustering algorithm (the red parts). The black lines show the detected symmetries for the case when the selected areas have five times greater weights than the rest of the scene.

Kateřina Uhrová in her Master’s thesis [118] experimented with different parameter settings of our symmetry plane detection method and also with different uses of the weights in the symmetry measure (see Section 6.1).




Figure 8.9: Non-uniformly sampled objects with planes detected by our symmetry plane detection method without using the weights (top row) and with the weights set using point density and point-to-point distances (bottom row) (figures taken from [118]).

The author of this doctoral thesis was consulting on this research. The results were interesting and especially the use of the weights for mitigating the effect of non-uniform point sampling, which poses a problem for the method, appeared promising, see Fig. 8.9. The top row shows objects with very non-uniform vertex sampling (only vertices are used as the input for the method) and the planes detected by the method without using the weights. The bottom row shows the same objects with planes detected with the weights set using point density in combination with point-to-point distances. The results in the bottom row are obviously better.

Dominik Ráček used our symmetry plane detection method in his bachelor’s thesis [101] for constructing cutting planes of 3D models of human faces. He used an implementation of the method which is currently present in the production version of FIDENTIS Analyst II project [91] which is being developed by researches at the Masaryk University, Faculty of Informatics and uses our method for symmetry plane detection on human faces.

At last, David Jesenko et al. [55, 56] used the method for prediction of water level of lake Cerknica in Slovenia. Using the fact that the symmetry of the lake changes based on the current water level, our method was used to find the symmetry plane of the lake which was then used to predict the water level.

8.4 Other

Apart from the above content of this chapter there are some other contributions that can be mentioned.

The implementation of our reflectional and rotational symmetry detection methods (described in Chapters 6 and 7) are currently being used by a few other students in their bachelor's theses.

Previous work of the author of this thesis, mainly on the topic of reflectional symmetry detection (Chapter 6) helped the graphics group at the University of West Bohemia, Department of Computer Science and Engineering receive a research grant GACR 21-08009K, Generalized Symmetries and Equivalences of Geometric Data, Czech Science Foundation. At the moment of submitting the thesis the graphics group is working on the project with the Department of Mathematics and a research group at the University in Maribor in Slovenia and the author of this thesis is participating in it.

At last, based on the work of Ondřej Anděl mentioned in the previous text, the research group at our department managed to establish cooperation with researchers at the Czech Technical University in Prague regarding user testing of human perception of symmetry.

9 Conclusion

In this work we focused on automatic detection of symmetries in geometric objects, mainly in 3D space. We first provided some necessary theoretical background to understand the meaning of symmetries and described the link between symmetry detection and the problem of registration. Next, we described various existing methods and approaches for detecting symmetries of different types and briefly also for rigid surface registration. The largest part of the text then presented our own contributions to the field of rigid surface registration and symmetry plane detection.

First, we proposed a new approach to consensus evaluation for RANSAC-based registration methods using density peak location in the transformation space which was inspired by previous work in symmetry detection. We tested this approach in a model RANSAC registration algorithm and we compared it to the conventional evaluation strategy based on overlap measurement. Since the density peak location algorithm requires defining a distance metric in the space of rigid transformations, we tested and analyzed several different metrics, both practically and theoretically, and presented the results. We also proposed an improvement of the most suitable of the existing metrics. This new improved metric now seems to represent the best choice for measuring distances between rigid transformations of 3D triangle meshes.

In the context of symmetry plane detection we first focused on what we call the Mode-based approach which can actually also include the density peak evaluation algorithm we proposed previously for surface registration. We analyzed and tested various ways of representing planes in the methods that implement the Mode-based approach for symmetry plane detection because the choice of the plane representation can impact the results considerably. We mainly focused on possible distance functions that can be defined for the various representations to measure similarities/dissimilarities between the planes and we reported the results. As part of this we also proposed a new way of representing planes using dual quaternions and using the dual quaternion algebra to compute reflections over planes.

We further proposed a new method for symmetry plane detection in discrete point sets (point clouds) based on a differentiable symmetry measure and a gradient-based numerical optimization. The method showed itself to be very fast, robust and also quite general. It handles well noisy objects and even objects where significant parts are missing and is mostly superior to other existing methods.

Next, we proposed using pairs of potential reflectional symmetries and a quaternion-based parameterization of the rotation transformation to find rotational symmetries in 3D point sets. The method also utilizes the previously described symmetry measure and appears to be the most robust and general method for finding rotational symmetries in 3D objects among the methods that are known to us.

At last, we mentioned some of our smaller contributions including a new method for finding local reflectional symmetries in 3D scenes represented by point sets using a simple PCA-based local shape descriptor.

Most notably, the ideas and approaches we used in our methods appear to be flexible enough to be reused in the future, in the form of extensions or modifications, to solve even greater challenges. We proved this in Chapter 7 where we used the overall approach and core ideas from the symmetry plane detection method presented in Chapter 6 and by performing several novel modifications we turned it into a method for detecting rotational symmetries. However, the overall approach remained the same - use sparse sampling of the transformation space so that the number of evaluated candidates is not unnecessarily large, select only a few candidates that have a chance of being close to good symmetries, apply numerical optimization to maximize a specific differentiable symmetry measure computed on a properly simplified version of the object and select the best local maximum found. Extensions for other symmetry types would very likely be possible as well since the transformation in the symmetry measure can be changed to any other type. Therefore, it could be possible to modify the method for detecting rigid symmetries, where the transformations could be differentially parameterized by dual quaternions, or even general affine symmetries where the parameterization could be done on the standard matrix representation directly. This would, however, require redesigning the process of creating candidate symmetries. For this purpose, some form of a local shape descriptor could be used but a more convenient approach would be using the 4-point congruent sets [2, 80] and their affine invariant properties.

In Section 8.3 we also showed a modification of the method for symmetry detection in 2D curves and preliminary results of another modification that attempts finding the spherical surface of reflectional symmetry in 3D point sets instead of the symmetry plane. Although these modifications were not done by the author of this work directly, they confirm the flexibility of the overall approach. A possible way of future research could include extending the approach even further to allow detecting symmetries that are represented by reflections over more general surfaces such as general quadratic surfaces.

The weights in the symmetry measure (see Section 6.1) also showed to be a useful tool for adjustments and modifications. We showed two possible uses for them already in Chapter 6 but some different uses of the weights were found by others as described Section 8.3. All of these uses, especially the one for mitigating the effect of non-uniform sampling, suggest that the weights provide a great tool for additional adjustments and extensions of the approach and increase its flexibility even more. The weights could also be possibly used for detecting symmetries in volumetric data, rather than point clouds, by assigning static weights to pairs of voxels based on the similarity of their values. Finding more beneficial ways of employing the weights in the symmetry measure is another potential way of future work.

By our opinion, showing that a single approach to symmetry detection can exist that is robust, reliable, accurate, computationally efficient and at the same time is flexible and generic enough to allow modifications and extensions in so many different ways as suggested above, is the largest contribution of this thesis.

We believe that the material presented in Chapter 4 regarding rigid surface registration, RANSAC consensus evaluation and transformation distance metrics, is also well reusable in other applications as was partly shown in Section 8.2. Furthermore, the method for finding local reflectional symmetries in 3D point clouds, briefly described in Section 8.1, utilizes content of three different chapters of this work. It uses a modified version of the density peak location algorithm from Chapter 4 to find modes in the the candidate space of planes. From Chapter 5 it uses one of the analyzed distance functions for planes in the density peak location process and from Chapter 6 it uses mainly the point cloud simplification and optimization of the symmetry measure on the locally symmetric parts. The distance function for planes from Chapter 5 is also used in the symmetry plane detection method in Chapter 6 in the candidate pruning step to find similar planes. This further suggests the flexibility a reusability of the research outcomes presented throughout this thesis.

A Activities

A.1 Publications on International Conferences

- Hruda, L., Kolingerová, I., and Lávička, M. Plane space representation in context of mode-based symmetry plane detection. In *Computational Science – ICCS 2020* (Cham, 2020), V. V. Krzhizhanovskaya, G. Závodszy, M. H. Lees, J. J. Dongarra, P. M. A. Sloot, S. Brissos, and J. Teixeira, Eds., Springer International Publishing, pp. 509–523. [50] **(80%)**

My contribution: theoretical analysis of the plane space representations and distance functions, derivation of the dual quaternion representation, evaluation methodology, implementation, experiments, text writing.

- Jesenko, D., Hruda, L., Kolingerová, I., Žalik, B., and Podgorelec, D. Detection of water levels in lake cerknica using sentinel-2 data and symmetry. In 3rd International Conference on Advances in Signal Processing and Artificial Intelligence (ASPAI' 2021) (2021), pp. 66–69. [55] **(10%)**

My contribution: consulting, proofreading, writing a small part of the text.

- Hruda, L., Kolingerová, I., and Podgorelec, D. Local reflectional symmetry detection in point clouds using a simple pca-based shape descriptor. In *Proceedings of the 18th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 1: GRAPP* (2023), pp. 52–63. [51] **(90%)**

My contribution: core idea, design of the overall method, implementation, experiments, text writing.

A.2 Publications in Impacted Journals

- Hruda, L., Dvořák, J., and Váša, L. On evaluating consensus in ransac surface registration. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 175–186. [46] **(50%)**, **citations:** 1 (WoS / Scopus / Scholar, excludes publications with any co-author intersection)

My contribution: theoretical analysis of the transformation distance metrics, derivation of the metric based on triangle sum of squares (this was also done in parallel by doc. Váša), implementation of the metrics and their integration into the registration algorithm, experiments and comparisons, significant participation on text writing.

- Hruda, L., Kolingerová, I., and Váša, L. Robust, fast and flexible symmetry plane detection based on differentiable symmetry measure. *The Visual Computer* (Jan 2021). [53] **(80%)**, **citations:** 3 (WoS: 2, Scopus: 1, Scholar: 2, excludes publications with any co-author intersection)

My contribution: core idea, design and implementation of the overall method, experiments, comparisons, text writing.

- Hruda, L., Kolingerová, I., Lávicka, M., and Maňák, M. Rotational symmetry detection in 3d using reflectional symmetry candidates and quaternion-based rotation parameterization. *Computer Aided Geometric Design* 98 (2022), 102138. [49] **(80%)**

My contribution: theoretical and practical verification of the core idea, derivation of the quaternion-based rotation parameterization, design and implementation of the overall method, experiments, comparisons, text writing.

A.3 Other Topic-Related Publications

- Hruda, L., and Dvořák, J. Estimating approximate plane of symmetry of 3d triangle meshes. In Proc. *Central European Seminar on Computer Graphics* (Smolenice, Slovakia, 2017). [47] **(50%)**

My contribution: participation on designing the method, implementing the method, performing experiments and text writing.

- Jesenko, D., Hruđa, L., Kolingerová, I., Źalik, B., and Podgorelec, D. Symmetry-based method for water level prediction using sentinel 2 data. *Sensors & Transducers* 256, 2 (2022), 12–18. [56] **(10%)**

My contribution: consulting, proofreading, writing a small part of the text.

A.4 Non-Related Publications

- Hruđa, L., and Kohout, J. Generic caching library and its use for vtk-based real-time simulation and visualization systems. In *VISIGRAPP (1: GRAPP)* (2018), pp. 154–164. [48] **(90%)**

My contribution: design and implementation of the caching library, part of the experiments, writing major part of the text.

- Dvořák, J., Káčereková, Z., Vaněček, P., Hruđa, L., and Váša, L. As-rigid-as-possible volume tracking for time-varying surfaces. *Computers & Graphics* 102 (2022), 329–338. [28] **(5%)**, **citations:** 1 (WoS / Scopus / Scholar, excludes publications with any co-author intersection)

My contribution: consulting, proofreading.

A.5 Participation in Scientific Projects

- GACR 21-08009K, Generalized Symmetries and Equivalences of Geometric Data, Czech Science Foundation
- SGS-2019-016, Synthesis and Analysis of Geometric and Computing Models, Ministry of Education, Youth and Sports of the Czech Republic
- SGS-2022-015, New Methods for Medical, Spatial and Communication Data, Ministry of Education, Youth and Sports of the Czech Republic

A.6 Oral Presentations

- Hledání symetrie v geometrických modelech (Language: Czech), 13. 11. 2018, seminar of the graphics group at the University of West

Bohemia, Faculty of Applied Sciences, Pilsen

- Hledání symetrie v geometrických modelech (Language: Czech), 6. 12. 2018, finale of the Czechoslovakian IT SPY 2018 master's thesis competition, Prague
- On evaluating consensus in RANSAC surface registration (Language: English), 2. 7. 2019, training presentation for the SGP 2019 conference and a seminar of the graphics group at the University of West Bohemia, Faculty of Applied Sciences, Pilsen
- On evaluating consensus in RANSAC surface registration (Language: English), 10. 7. 2019, Symposium on Geometry Processing 2019, Milan, Italy
- Local reflectional symmetry detection in point clouds using a simple PCA-based shape descriptor (Language: English), 21. 2. 2023, GRAPP 2023, online presentation in the form of a pre-recorded video

A.7 Other

- 2nd place in the Czechoslovakian IT SPY 2018 master's thesis competition
- GRAPP 2023 Best Student Paper Award
- *h*-index of 2 according to WoS and Scopus (excludes self-cites but includes cites by publications with the same co-authors)

Bibliography

- [1] Aim@shape.
<http://visionair.ge.imati.cnr.it/ontologies/shapes/>, 2014.
Accessed: 2018-11-17.
- [2] AIGER, D., MITRA, N. J., AND COHEN-OR, D. 4-points congruent sets for robust surface registration. *ACM Transactions on Graphics* 27, 3 (2008), #85, 1–10.
- [3] ALEXA, M., RUSINKIEWICZ, S., NEHAB, D., AND SHILANE, P. Stratified point sampling of 3d models. In *Proc. Eurographics Symp. on Point-Based Graphics* (2004), pp. 49–56.
- [4] AMATO, N. M., BAYAZIT, O. B., DALE, L. K., JONES, C., AND VALLEJO, D. Choosing good distance metrics and local planners for probabilistic roadmap methods. In *Proceedings. 1998 IEEE International Conference on Robotics and Automation (Cat. No. 98CH36146)* (1998), vol. 1, IEEE, pp. 630–637.
- [5] ANDĚL, O. Finding the symmetry of geodata, master’s thesis at the university of west bohemia, faculty of applied sciences.
- [6] ARVANITIS, G., ZACHARAKI, E. I., VÁŠA, L., AND MOUSTAKAS, K. Broad-to-narrow registration and identification of 3d objects in partially scanned and cluttered point clouds. *IEEE Transactions on Multimedia* 24 (2021), 2230–2245.
- [7] BALLARD, D. H. Readings in computer vision: Issues, problems, principles, and paradigms. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1987, ch. Generalizing the Hough Transform to Detect Arbitrary Shapes, pp. 714–725.
- [8] BELTA, C., AND KUMAR, V. Euclidean metrics for motion generation on $se(3)$. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science* 216, 1 (2002), 47–60.
- [9] BERG, DE, M., CHEONG, O., KREVELD, VAN, M., AND OVERMARS, M. *Computational geometry : algorithms and applications*, 3rd ed ed. Springer, Germany, 2008.
- [10] BESL, P. J., AND MCKAY, N. D. A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.* 14, 2 (Feb. 1992), 239–256.

- [11] BOKELOH, M., BERNER, A., WAND, M., SEIDEL, H.-P., AND SCHILLING, A. Symmetry detection using feature lines. *Computer Graphics Forum* 28, 2 (2009), 697–706.
- [12] BOUAZIZ, S., TAGLIASACCHI, A., AND PAULY, M. Sparse iterative closest point. In *Proceedings of the Eleventh Eurographics/ACMSIGGRAPH Symposium on Geometry Processing* (Aire-la-Ville, Switzerland, Switzerland, 2013), SGP '13, Eurographics Association, pp. 113–123.
- [13] BOUBEKEUR, T., AND ALEXA, M. Mesh simplification by stochastic sampling and topological clustering. *Computers & Graphics* 33, 3 (2009), 241–249.
- [14] BRUNNSTRÖM, K., AND STODDART, A. J. Genetic algorithms for free-form surface matching. In *Proc. 13th Int Pattern Recognition Conf* (1996), vol. 4, pp. 689–693.
- [15] BÜLOW, H., AND BIRK, A. Spectral 6-dof registration of noisy 3d range data with partial overlap. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* 35 (2013), 954–969.
- [16] CAILLIÈRE, D., DENIS, F., PELE, D., AND BASKURT, A. 3d mirror symmetry detection using hough transform. In *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on* (2008), IEEE, pp. 1772–1775.
- [17] CENSI, A., AND CARPIN, S. Hsm3d: Feature-less global 6dof scan-matching in the hough/radon domain. In *Proceedings of the 2009 IEEE International Conference on Robotics and Automation* (Piscataway, NJ, USA, 2009), ICRA'09, IEEE Press, pp. 1585–1592.
- [18] CHALÁS, I., URBANOVÁ, P., KOTULANOVÁ, Z., JANDOVÁ, M., KRÁLÍK, M., KOZLÍKOVÁ, B., AND SOCHOR, J. Forensic 3d facial identification software (fidentis). In *Proceedings of the 20th World Meeting of the International Association of Forensic Sciences* (2014).
- [19] CHARIKAR, M., AND GUHA, S. Improved combinatorial algorithms for the facility location and k-median problems. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science* (1999), FOCS '99, IEEE Computer Society, pp. 378–.
- [20] CHEN, Y., AND MEDIONI, G. Object modelling by registration of multiple range images. *Image Vision Comput.* 10, 3 (Apr. 1992), 145–155.
- [21] CHOW, C. K., TSUI, H.-T., AND LEE, T. Surface registration using a dynamic genetic algorithm. *Pattern Recognition* 37, 1 (2004), 105–117.

- [22] CICCONE, M., HILDEBRAND, D. G., AND ELLIOTT, H. Finding mirror symmetry via registration and optimal symmetric pairwise assignment of curves: Algorithm and results. In *Proceedings of the IEEE International Conference on Computer Vision* (2017), pp. 1759–1763.
- [23] CIGNONI, P., CALLIERI, M., CORSINI, M., DELLEPIANE, M., GANOVELLI, F., AND RANZUGLIA, G. Meshlab: an open-source mesh processing tool. In *Eurographics Italian chapter conference* (2008), vol. 2008, pp. 129–136.
- [24] CIGNONI, P., ROCCHINI, C., AND SCOPIGNO, R. Metro: measuring error on simplified surfaces. In *Computer Graphics Forum* (1998), vol. 17, Wiley Online Library, pp. 167–174.
- [25] COMANICIU, D., AND MEER, P. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on pattern analysis and machine intelligence* 24, 5 (2002), 603–619.
- [26] COMBÈS, B., HENNESSY, R., WADDINGTON, J., ROBERTS, N., AND PRIMA, S. Automatic symmetry plane estimation of bilateral objects in point clouds. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on* (2008), IEEE, pp. 1–8.
- [27] DE CASTRO, E., AND MORANDI, C. Registration of translated and rotated images using finite fourier transforms. *IEEE Trans. Pattern Anal. Mach. Intell.* 9, 5 (May 1987), 700–703.
- [28] DVOŘÁK, J., KÁČERKOVÁ, Z., VANĚČEK, P., HRUDA, L., AND VÁŠA, L. As-rigid-as-possible volume tracking for time-varying surfaces. *Computers & Graphics* 102 (2022), 329–338.
- [29] ECINS, A., FERMULLER, C., AND ALOIMONOS, Y. Detecting reflectional symmetries in 3d data through symmetrical fitting. In *Proceedings of the IEEE International Conference on Computer Vision* (2017), pp. 1779–1783.
- [30] ECINS, A., FERMÜLLER, C., AND ALOIMONOS, Y. Cluttered tabletop dataset. <http://www.aecins.umiacs.io/projects/symseg/>, 2018. Accessed: 2021-12-02.
- [31] ECINS, A., FERMÜLLER, C., AND ALOIMONOS, Y. Seeing behind the scene: Using symmetry to reason about objects in cluttered environments. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2018), IEEE, pp. 7193–7200.

- [32] FANG, R., GODIL, A., LI, X., AND WAGAN, A. A new shape benchmark for 3d object retrieval. In *International Symposium on Visual Computing* (2008), Springer, pp. 381–392.
- [33] FIGUEREDO, L., ADORNO, B. V., ISHIHARA, J. Y., AND BORGES, G. A. Robust kinematic control of manipulator robots using dual quaternion representation. In *2013 IEEE International Conference on Robotics and Automation* (2013), IEEE, pp. 1949–1955.
- [34] FUNK, C., LEE, S., OSWALD, M. R., TSOVKAS, S., SHEN, W., COHEN, A., DICKINSON, S., AND LIU, Y. 2017 iccv challenge: Detecting symmetry in the wild. In *Proceedings of the IEEE International Conference on Computer Vision* (2017), pp. 1692–1701.
- [35] GALLIER, J. Notes on convex sets, polytopes, polyhedra, combinatorial topology, voronoi diagrams and delaunay triangulations. *arXiv preprint arXiv:0805.0292* (2008).
- [36] GALLIER, J. The cartan–dieudonné theorem. In *Geometric Methods and Applications*. Springer, 2011, pp. 231–280.
- [37] GAO, L., ZHANG, L.-X., MENG, H.-Y., REN, Y.-H., LAI, Y.-K., AND KOBBELT, L. Prs-net: Planar reflective symmetry detection net for 3d models. *IEEE Transactions on Visualization and Computer Graphics* 27, 6 (2020), 3007–3018.
- [38] GOLDMAN, R. *Rethinking Quaternions: Theory and Computation*. Morgan and Claypool Publishers, 2010.
- [39] GOTHANDARAMAN, R., JHA, R., AND MUTHUSWAMY, S. Reflectional and rotational symmetry detection of cad models based on point cloud processing. In *2020 IEEE 4th Conference on Information & Communication Technology (CICT)* (2020), IEEE, pp. 1–5.
- [40] HAUSDORFF, F. *Grundzüge der Mengenlehre*. Chelsea Pub. Co, New York, 1949.
- [41] HEIDER, P., PIERRE-PIERRE, A., LI, R., AND GRIMM, C. Local shape descriptors, a survey and evaluation. In *Proceedings of the 4th Eurographics Conference on 3D Object Retrieval* (2011), 3DOR '11, Eurographics Association, pp. 49–56.
- [42] HORN, B. K. P. Extended gaussian images. *Proceedings of the IEEE* 72, 2 (1984), 1671–1686.
- [43] HRUDA, L. Symmetry detection in geometric models, master’s thesis at the university of west bohemia, faculty of applied sciences.

- [44] HRUDA, L. Symmetry detection in geometric models. <https://www.kiv.zcu.cz/site/documents/verejne/vyzkum/publikace/technicke-zpravy/2021/Rigo-Hruda-2021.pdf>, 03 2021.
- [45] HRUDA, L., DVOŘÁK, J., AND VÁŠA, L. On evaluating consensus in ransac surface registration: Supplementary material. <http://meshcompression.org/sgp2019>.
- [46] HRUDA, L., DVOŘÁK, J., AND VÁŠA, L. On evaluating consensus in ransac surface registration. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 175–186.
- [47] HRUDA, L., AND DVOŘÁK, J. Estimating approximate plane of symmetry of 3d triangle meshes. In *Proc. Central European Seminar on Computer Graphics* (Smolenice, Slovakia, 2017).
- [48] HRUDA, L., AND KOHOUT, J. Generic caching library and its use for vtk-based real-time simulation and visualization systems. In *VISIGRAPP (1: GRAPP)* (2018), pp. 154–164.
- [49] HRUDA, L., KOLINGEROVÁ, I., LÁVIČKA, M., AND MAŇÁK, M. Rotational symmetry detection in 3d using reflectional symmetry candidates and quaternion-based rotation parameterization. *Computer Aided Geometric Design* 98 (2022), 102138.
- [50] HRUDA, L., KOLINGEROVÁ, I., AND LÁVIČKA, M. Plane space representation in context of mode-based symmetry plane detection. In *Computational Science – ICCS 2020* (Cham, 2020), V. V. Krzhizhanovskaya, G. Závodszy, M. H. Lees, J. J. Dongarra, P. M. A. Slood, S. Brissos, and J. Teixeira, Eds., Springer International Publishing, pp. 509–523.
- [51] HRUDA, L., KOLINGEROVÁ, I., AND PODGORELEC, D. Local reflectional symmetry detection in point clouds using a simple pca-based shape descriptor. In *Proceedings of the 18th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 1: GRAPP* (2023), pp. 52–63.
- [52] HRUDA, L., KOLINGEROVÁ, I., AND VÁŠA, L. Robust, fast and flexible symmetry plane detection based on differentiable symmetry measure: Supplementary material. <http://meshcompression.org/tvcj-2020>.
- [53] HRUDA, L., KOLINGEROVÁ, I., AND VÁŠA, L. Robust, fast and flexible symmetry plane detection based on differentiable symmetry measure. *The Visual Computer* 38, 2 (2022), 555–571.

- [54] HUYNH, D. Q. Metrics for 3d rotations: Comparison and analysis. *Journal of Mathematical Imaging and Vision* 35, 2 (2009), 155–164.
- [55] JESENKO, D., HRUDA, L., KOLINGEROVÁ, I., ŽALIK, B., AND PODGORELEC, D. Detection of water levels in lake cerknica using sentinel-2 data and symmetry. In *3rd International Conference on Advances in Signal Processing and Artificial Intelligence (ASPAI' 2021)* (2021), pp. 66–69.
- [56] JESENKO, D., HRUDA, L., KOLINGEROVÁ, I., ŽALIK, B., AND PODGORELEC, D. Symmetry-based method for water level prediction using sentinel 2 data. *Sensors & Transducers* 256, 2 (2022), 12–18.
- [57] JI, P., AND LIU, X. A fast and efficient 3d reflection symmetry detector based on neural networks. *Multimedia Tools and Applications* 78, 24 (2019), 35471–35492.
- [58] KAKARALA, R., KALIAMOORTHY, P., AND PREMACHANDRAN, V. Three-dimensional bilateral symmetry plane estimation in the phase domain. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2013), pp. 249–256.
- [59] KOHEK, Š., STRNAD, D., ŽALIK, B., AND KOLMANIČ, S. Interactive synthesis and visualization of self-organizing trees for large-scale forest succession simulation. *Multimedia Systems* 25, 3 (2019), 213–227.
- [60] KORMAN, S., LITMAN, R., AVIDAN, S., AND BRONSTEIN, A. Probably approximately symmetric: Fast rigid symmetry detection with global guarantees. In *Computer Graphics Forum* (2015), vol. 34, Wiley Online Library, pp. 2–13.
- [61] KOSCHAT, M. A., AND SWAYNE, D. F. A weighted procrustes criterion. *Psychometrika* 56, 2 (1991), 229–239.
- [62] KROEMER, O., AMOR, H. B., EWERTON, M., AND PETERS, J. Point cloud completion using extrusions. In *2012 12th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2012)* (2012), IEEE, pp. 680–685.
- [63] KRUTINOVÁ, Š. Využití symetrie pro geodata, bachelor’s thesis at the university of west bohemia, faculty of applied sciences.
- [64] KUBÁSKOVÁ, K. Recognition of important features of triangulated human head models. In *Proc. Central European Seminar on Computer Graphics* (Smolenice, Slovakia, 2016).

- [65] KUFFNER, J. J. Effective sampling and distance metrics for 3d rigid body path planning. In *ICRA* (2004), pp. 3993–3998.
- [66] LEVORA, M. Efektivní implementace registrace 3d povrchů, bachelor’s thesis at the university of west bohemia, faculty of applied sciences.
- [67] LEVOY, M., GERTH, J., CURLESS, B., AND PULL, K. The stanford 3d scanning repository.
<http://www.graphics.stanford.edu/data/3Dscanrep/> (2005).
- [68] LI, B. Bo li homepage.
<https://sites.google.com/site/libohomepage/homepage>. Accessed: 2020-03-26.
- [69] LI, B., JOHAN, H., YE, Y., AND LU, Y. Efficient 3d reflection symmetry detection: A view-based approach. *Graphical Models* 83 (2016), 2–14.
- [70] LI, C., WAND, M., WU, X., AND SEIDEL, H.-P. Approximate 3D partial symmetry detection using co-occurrence analysis. In *2015 International Conference on 3D Vision* (2015), IEEE, pp. 425–433.
- [71] LI, K., YANG, J., LAI, Y., AND GUO, D. Robust non-rigid registration with reweighted position and transformation sparsity. *IEEE Trans. Vis. Comput. Graph.* 25, 6 (2019), 2255–2269.
- [72] LI, X., YIN, Z., WEI, L., WAN, S., YU, W., AND LI, M. Symmetry and template guided completion of damaged skulls. *Computers & Graphics* 35, 4 (2011), 885–893.
- [73] LIN, H.-W., TAI, C.-L., AND WANG, G.-J. A mesh reconstruction algorithm driven by an intrinsic property of a point cloud. *Computer-Aided Design* 36, 1 (2004), 1–9.
- [74] LIU, D. C., AND NOCEDAL, J. On the limited memory bfgs method for large scale optimization. *Mathematical programming* 45, 1-3 (1989), 503–528.
- [75] LOWE, D. G. Object recognition from local scale-invariant features. In *Proceedings of the seventh IEEE international conference on computer vision* (1999), vol. 2, Ieee, pp. 1150–1157.
- [76] MAKADIA, A., PATTERSON, A. I., AND DANILIDIS, K. Fully automatic registration of 3d point clouds. In *Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Volume 1* (Washington, DC, USA, 2006), CVPR ’06, IEEE Computer Society, pp. 1297–1304.

- [77] MARTINET, A., SOLER, C., HOLZSCHUCH, N., AND SILLION, F. X. Accurate detection of symmetries in 3d shapes. *ACM Transactions on Graphics (TOG)* 25, 2 (2006), 439–464.
- [78] MAVRIDIS, P., SIPIRAN, I., ANDREADIS, A., AND PAPAIOANNOU, G. Object completion using k-sparse optimization. In *Computer Graphics Forum* (2015), vol. 34, Wiley Online Library, pp. 13–21.
- [79] MCCRAE, J., SINGH, K., AND MITRA, N. J. Slices: a shape-proxy based on planar sections. *ACM Trans. Graph.* 30, 6 (2011), 168.
- [80] MELLADO, N., AIGER, D., AND MITRA, N. J. Super 4pcs fast global pointcloud registration via smart indexing. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 205–215.
- [81] MEYER, M., DESBRUN, M., SCHRÖDER, P., AND BARR, A. H. Discrete differential-geometry operators for triangulated 2-manifolds. In *Visualization and mathematics III*. Springer, 2003, pp. 35–57.
- [82] MITRA, N. J., GELFAND, N., POTTMANN, H., AND GUIBAS, L. Registration of point cloud data from a geometric optimization perspective. In *Symposium on Geometry Processing* (2004), pp. 23–31.
- [83] MITRA, N. J., GUIBAS, L. J., AND PAULY, M. Partial and approximate symmetry detection for 3d geometry. In *ACM Transactions on Graphics (TOG)* (2006), vol. 25, ACM, pp. 560–568.
- [84] MORÉ, J. J. The levenberg-marquardt algorithm: implementation and theory. In *Numerical analysis*. Springer, 1978, pp. 105–116.
- [85] MOURYCOVÁ, E. Hledání symetrie funkcí s využitím symetrie množiny bodů, bachelor’s thesis at the university of west bohemia, faculty of applied sciences.
- [86] MUDELSEE, M. Estimating pearson’s correlation coefficient with bootstrap confidence interval from serially dependent time series. *Mathematical Geology* 35, 6 (2003), 651–665.
- [87] NAGAR, R., AND RAMAN, S. Fast and accurate intrinsic symmetry detection. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2018), pp. 417–434.
- [88] NAGAR, R., AND RAMAN, S. Detecting approximate reflection symmetry in a point set using optimization on manifold. *IEEE Transactions on Signal Processing* 67, 6 (2019), 1582–1595.

- [89] NAGAR, R., AND RAMAN, S. 3dsymm: robust and accurate 3d reflection symmetry detection. *Pattern Recognition* (2020), 107483.
- [90] NELDER, J. A., AND MEAD, R. A simplex method for function minimization. *The computer journal* 7, 4 (1965), 308–313.
- [91] OŠLEJŠEK, R., SOCHOR, J., AND URBANOVÁ, P. Fidentis analyst ii. <https://www.radek-oslejsek.cz/it/fidentis-analyst-2/>, 2022.
- [92] PANOZZO, D., LIPMAN, Y., PUPPO, E., AND ZORIN, D. Fields on symmetric surfaces. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 111.
- [93] PETITJEAN, M. A DEFINITION OF SYMMETRY. *Symmetry: Culture and Science* 18, 2-3 (2007), 99 – 119.
- [94] PHAM, H.-L., PERDEREAU, V., ADORNO, B. V., AND FRAISSE, P. Position and orientation control of robot manipulators using dual quaternion feedback. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems* (2010), IEEE, pp. 658–663.
- [95] PLUMED, R., COMPANY, P., AND VARLEY, P. Study on human perception of bilateral symmetry planes. <http://www.regeo.uji.es/publicaciones/Regeo10.pdf>, 2016.
- [96] PODOLAK, J., SHILANE, P., GOLOVINSKIY, A., RUSINKIEWICZ, S., AND FUNKHOUSER, T. A planar-reflective symmetry transform for 3d shapes. *ACM Transactions on Graphics (TOG)* 25, 3 (2006), 549–559.
- [97] POMERLEAU, F., COLAS, F., SIEGWART, R., AND MAGNENAT, S. Comparing icp variants on real-world data sets. *Auton. Robots* 34, 3 (Apr. 2013), 133–148.
- [98] POTTMANN, H., HUANG, Q.-X., YANG, Y.-L., AND HU, S.-M. Geometry and convergence analysis of algorithms for registration of 3d shapes. *International Journal of Computer Vision* 67, 3 (2006), 277–296.
- [99] PRANTL., M., VÁŠA., L., AND KOLINGEROVÁ., I. Symmetry-aware registration of human faces. In *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 1: GRAPP*, (2019), INSTICC, SciTePress, pp. 185–192.
- [100] PROSKOVA, J. Discovery of dual quaternions for geodesy. *J Geom Graph* 16, 2 (2012), 195–209.

- [101] RÁČEK, D. 3d prierezová rovina pre analýzu ľudských tvárí, bachelor's thesis at the masaryk university, faculty of informatics.
- [102] RUSINKIEWICZ, S. Estimating curvatures and their derivatives on triangle meshes. In *Symposium on 3D Data Processing, Visualization, and Transmission* (Sept. 2004).
- [103] RUSINKIEWICZ, S., AND LEVOY, M. Efficient variants of the ICP algorithm. In *Third International Conference on 3D Digital Imaging and Modeling (3DIM)* (June 2001).
- [104] RUSU, R. B., BLODOW, N., AND BEETZ, M. Fast point feature histograms (fpfh) for 3d registration. In *Proceedings of the 2009 IEEE International Conference on Robotics and Automation* (Piscataway, NJ, USA, 2009), ICRA '09, IEEE Press, pp. 1848–1853.
- [105] SCHIEBENER, D., SCHMIDT, A., VAHRENKAMP, N., AND ASFOUR, T. Heuristic 3d object shape completion based on symmetry and scene context. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on* (2016), IEEE, pp. 74–81.
- [106] SHI, Y., HUANG, J., ZHANG, H., XU, X., RUSINKIEWICZ, S., AND XU, K. Symmetrynet: learning to predict reflectional and rotational symmetries of 3d shapes from single-view rgb-d images. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–14.
- [107] SHI, Z., ALLIEZ, P., DESBRUN, M., BAO, H., AND HUANG, J. Symmetry and orbit detection via lie-algebra voting. In *Computer Graphics Forum* (2016), vol. 35, Wiley Online Library, pp. 217–227.
- [108] SHILANE, P., MIN, P., KAZHDAN, M., AND FUNKHOUSER, T. The princeton shape benchmark. In *Shape modeling applications, 2004. Proceedings* (2004), IEEE, pp. 167–178.
- [109] SIMARI, P., KALOGERAKIS, E., AND SINGH, K. Folding meshes: Hierarchical mesh segmentation based on planar symmetry. In *Symposium on geometry processing* (2006), vol. 256, pp. 111–119.
- [110] SIPIRAN, I. Analysis of partial axial symmetry on 3D surfaces and its application in the restoration of cultural heritage objects. In *Proceedings of the IEEE International Conference on Computer Vision Workshops* (2017), pp. 2925–2933.
- [111] SIPIRAN, I. Completion of cultural heritage objects with rotational symmetry. In *Proceedings of the 11th Eurographics Workshop on 3D Object Retrieval* (2018), pp. 87–93.

- [112] SÍPIRAN, I., GREGOR, R., AND SCHRECK, T. Approximate symmetry detection in partial 3d meshes. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 131–140.
- [113] SMEETS, D., KEUSTERMANS, J., HERMANS, J., VANDERMEULEN, D., AND SUETENS, P. Feature-based piecewise rigid registration in 2-d medical images. In *ISBI* (2012), IEEE, pp. 696–699.
- [114] SPECIALE, P., OSWALD, M. R., COHEN, A., AND POLLEFEYS, M. A symmetry prior for convex variational 3d reconstruction. In *European Conference on Computer Vision* (2016), Springer, pp. 313–328.
- [115] SUN, C., AND SHERRAH, J. 3d symmetry detection using the extended gaussian image. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19, 2 (1997), 164–168.
- [116] T., T., AND RIEKE-ZAPP, D. PRESIOUS 3d cultural heritage differential scans, 2013.
- [117] THRUN, S., AND WEGBREIT, B. Shape from symmetry. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on* (2005), vol. 2, IEEE, pp. 1824–1831.
- [118] UHROVÁ, K. Zobecněná symetrie geometrických dat, master’s thesis at the university of west bohemia, faculty of applied sciences.
- [119] VÁŠA, L., VANĚČEK, P., PRANTL, M., SKORKOVSKÁ, V., MARTÍNEK, P., AND KOLINGEROVÁ, I. Mesh statistics for robust curvature estimation. In *Computer Graphics Forum* (2016), vol. 35, Wiley Online Library, pp. 271–280.
- [120] VEDALDI, A., AND SOATTO, S. Quick shift and kernel methods for mode seeking. In *European conference on computer vision* (2008), Springer, pp. 705–718.
- [121] WANG, H., AND HUANG, H. Group representation of global intrinsic symmetries. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 51–61.
- [122] WANG, W., MA, J., XU, P., AND CHU, Y. Intrinsic symmetry detection on 3d models with skeleton-guided combination of extrinsic symmetries. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 617–628.
- [123] WENDLAND, H. Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Advances in computational Mathematics* 4, 1 (1995), 389–396.

- [124] WOLD, S., ESBENSEN, K., AND GELADI, P. Principal component analysis. *Chemometrics and intelligent laboratory systems 2*, 1-3 (1987), 37–52.
- [125] WONGWAEN, N., TIENDEE, S., AND SINTHANAYOTHIN, C. Method of 3d mesh reconstruction from point cloud using elementary vector and geometry analysis. In *2012 8th International Conference on Information Science and Digital Content Technology (ICIDT2012)* (2012), vol. 1, IEEE, pp. 156–159.
- [126] YIANILOS, P. N. Data structures and algorithms for nearest neighbor search in general metric spaces. In *Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms* (Philadelphia, PA, USA, 1993), SODA '93, Society for Industrial and Applied Mathematics, pp. 311–321.
- [127] ZABRODSKY, H., PELEG, S., AND AVNIR, D. Symmetry as a continuous feature. *IEEE Transactions on pattern analysis and machine intelligence* 17, 12 (1995), 1154–1166.
- [128] ZHOU, Q., AND JACOBSON, A. Thingi10k: A dataset of 10,000 3d-printing models. *arXiv preprint arXiv:1605.04797* (2016).
- [129] ZHOU, Q., PARK, J., AND KOLTUN, V. Fast global registration. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II* (2016), pp. 766–782.
- [130] ZHOU, Y., SHEN, S., AND HU, Z. Detail preserved surface reconstruction from point cloud. *Sensors* 19, 6 (2019), 1278.