

A Low-Dimensional Representation for Robust Partial Isometric Correspondences Computation

Alan Brunton* Michael Wand* Stefanie Wuhrer[†] Hans-Peter Seidel* Tino Weinkauf*

Abstract

Intrinsic shape matching has become the standard approach for pose invariant correspondence estimation among deformable shapes. Most existing approaches assume global consistency. While global isometric matching is well understood, only a few heuristic solutions are known for partial matching. Partial matching is particularly important for robustness to topological noise, which is a common problem in real-world scanner data. We introduce a new approach to partial isometric matching based on the observation that isometries are fully determined by local information: a map of a single point and its tangent space fixes an isometry. We develop a new representation for partial isometric maps based on equivalence classes of correspondences between pairs of points and their tangent-spaces. We apply our approach to register partial point clouds and compare it to the state-of-the-art methods, where we obtain significant improvements over global methods for real-world data and stronger guarantees than previous partial matching algorithms.

1 Introduction

Modern computer graphics has experienced a paradigm shift: Traditional manual modeling is increasingly complemented by data-driven techniques where measured data, such as 3D scans, are used as a basis for building 3D models. An important data source are 3D scans of deformable models, such as humans or animals in varying poses. Today, there exist sophisticated scanning setups for acquiring moving geometry in real-time [1–4] and there are even consumer devices on the market such as the Microsoft KinectTM. This leads to new applications such as virtual cinematography [5], or the creation of data-driven generative shape models of deformable objects [6–9]. Finding correspondences among such data is a fundamental problem for all of these applications: Almost any further processing, such as the registration of partial scans into a complete shape, editing of sequences, or statistical analysis, requires dense correspondences between surface points.

Matching deformable shapes is in many cases a difficult problem: If we permit rather general deformations this

might require many parameters that have to be explored during matching. The size of this search space is exponential with respect to the available degrees of freedom. However, for the important special case of a single object in different poses, we can often assume that the deformation is approximately isometric, i.e., preserving the intrinsic metric structure. Concretely, the distances along surfaces of objects such as humans, animals, plants, or cloths do not change a lot without serious injury or damage. This restriction leads to a strongly constrained search space. Lipman et al. [10] argue that isometries between topological disks are a special case of conformal mappings, thereby limiting the degrees of freedom to six (three point-to-point correspondences are sufficient). Ovsjanikov et al. [11] show that a single point correspondence is sufficient for a special class of shapes where the spectrum of the Laplace-Beltrami operator is not degenerate, thus showing that there are only two degrees of freedom in this special case. Additional cues, such as carefully selected constellations of local features [12], can even reduce the complexity for many shapes to the point of leaving a trivial search space, eliminating all degrees of freedom. As approximate isometry makes the correspondence problem feasible while still permitting significant pose changes, many of the recent shape matching algorithms are based on this assumption [10–19].

However, the isometric matching problem is not yet solved: Because of the intrinsic view of the geometry, it is naturally sensitive to topological noise. In case of holes, missing data, or contacts, intrinsic distances become distorted and thus no longer constitute invariants that can be exploited for matching. One solution is to replace the notion of distance. For example, by using diffusion distances, or variants thereof, one can reduce the sensitivity to topological artifacts [17, 20, 21] when the pieces of geometry that cause the problem are small in relation to the overall shape. Nonetheless, these invariants still break down in case of large artifacts (wide contacts, large holes, as shown for example in Section 5 of this paper). Unfortunately, real-world 3D scanner data, one of the main practical application areas, is almost universally troubled by substantial artifacts of this kind.

Formulated more generally, we have to address the problem of *partial matching*, where not the whole manifold can be brought into correspondence by a (near-) isometric mapping but only an excerpt of the surface can be matched. In this case, typical invariants (geodesic paths, Laplacian eigenfunctions) become unreliable because they utilize global information. Importantly, the excerpts of the

*Max Planck Institute for Informatics, Germany, {abrunton,mwand,weinkauf,hpseidel}@mpi-inf.mpg.de

[†]Cluster of Excellence, Multi-Modal Computing and Interaction, Saarland University, Germany, swuhrer@mmci.uni-saarland.de

surfaces that can be matched are not known a priori (otherwise, we could just restrict a traditional method accordingly) but need to be determined along with the matching. This seems to re-introduce prohibitive complexity as we now have to choose from an exponential number of such subsets [22]. The most important contribution of this paper is to show that with an appropriate matching model this is not the case. The search space is not much larger than in the global problem and we give an efficient algorithm for computing such matches.

The core of our method is based on the observation from differential geometry that an isometric map can be fully specified by using local information up to first order only: An isometry between Riemannian manifolds is fixed by a single point correspondence and an orthogonal map of the tangent spaces (see for example [23, p. 201]). In the case of surfaces, this means that the map of a point and a local direction (plus orientation, in case of unoriented surfaces) is sufficient to determine an isometry. We will sketch a constructive proof in Section 3 that directly yields a propagation algorithm for computing matches: Starting with a single oriented point match, correspondence information is incrementally propagated to the neighborhood, thereby flood-filling a partially consistent region of isometric geometry. Being local in nature, the method handles partial matches naturally and is robust to topological noise, which is reported naturally as boundary of partiality.

From a structural point of view, we can understand partial isometries among smooth, connected manifolds as equivalence classes of mappings of a pair of points and their tangent spaces on the two surfaces involved. In the case of oriented 2-manifolds, each such object has six degrees of freedom (a point and an angle around the surface normal for each for source and target surface, respectively). The mapping is captured by a equivalence class of such 6-tuples. This set is redundant in that one degree of freedom (the sum of the angles) can be chosen arbitrarily and two further parameters (either the starting or the end point) are only needed to select the partial region to be mapped. This leaves three degrees of freedom that need to be actually explored densely, with false starting points being rejected in $\mathcal{O}(1)$ time.

In this view, we can perform a relaxation: In order to also find *approximate* isometries, we can cast this problem as the task of finding *approximate equivalence classes*. Kim et al. [18] have used this idea in the context of globally consistent isometric mappings in order to efficiently find approximate isometries. In our paper, we demonstrate how our partial matching framework can be adapted to perform the same task in a partial matching scenario. We perform agglomerative clustering [24] in the space of near-isometric mappings, which are concisely represented using the tuples introduced above.

We validate our algorithm on standard benchmark data sets and raw scanner data, and compare the results to previous work. We show a significant improvement in quality over global methods in shapes with topological noise. Our algorithm yields similar or better results as previous heuristics for partial matching, but with stronger guaran-

tees of discovering existing isometries as outlined above.

In summary, our paper proposes a systematic framework and new algorithms for extending isometric matching to the case of partial consistency, thereby making the following specific contributions:

- We characterize partial isometries of shapes by single-point maps up to first order, which yields a tight bound on the inherent degrees of freedom.
- This leads to a novel matching algorithm that provides a systematic approach to the general setting of partial intrinsic matching, where both surfaces may be incomplete, including robustness to strong topological noise.
- By interpreting partial matching as a problem of finding approximate equivalent classes in our novel representation, we obtain an algorithm for *approximate* partial isometric matching.

Our algorithmic pipeline for approximate partial isometric matching is summarized as follows. We identify distinctive feature points on both surfaces. We compute oriented point correspondences by matching feature descriptors, with orientation determined by nearby features. From oriented point matches, we perform local metric propagation, stopping the propagation when the stretching becomes too large. This gives us a set of partial isometries, covering different, but possibly overlapping regions of both surfaces. We define a dissimilarity measure between partial isometries, and use this to cluster the partial matches into equivalence classes. The cluster (equivalence class) with the smallest total intra-cluster distance is merged by taking the geodesic centroid of all candidate correspondences for each point on the source manifold.

The steps in this pipeline exhibit sensitivity to certain challenges. Most prominent are the sensitivity of feature matching to surface noise, missing data and topological noise, and the sensitivity of the clustering step to under-sampling of the space of isometric mappings, which can be exacerbated by problems in the feature matching stage, in addition to the lack of a proper distance metric between partial matches. The difficulty of reliable feature matching on real data can be alleviated by the metric propagation algorithm, which will typically produce smaller partial isometries for incorrect feature matches than for correct ones since we expect the local metric to vary significantly for different parts of the surface. However, in the presence of strong continuous intrinsic symmetries this assumption breaks down. We thoroughly discuss and explore in which situations and to what extent these challenges create problems in the final result in Section 5.

2 Complexity of Isometric Shape Matching

In this section, we discuss different isometric matching models and their implications on the difficulty for finding

a globally optimal solution to the matching problem. To the best of our knowledge, this has not yet been analyzed explicitly in literature. We do not aim to give an extensive review of surface registration methods; for this, we refer the reader to recent surveys such as van Kaick et al. [25] or Tam et al. [26].

We start by introducing some formal notations.

Manifolds: We consider smooth, orientable 2-manifolds $\mathcal{M} \subset \mathbb{R}^3$ embedded in three-dimensional space. In order to represent partial data (such as scans with acquisition holes), we permit boundaries, denoted by $\partial\mathcal{M}$. The orientation of \mathcal{M} might (optionally) be prescribed by oriented surface normals $\mathbf{n}(\mathbf{x})$, $\mathbf{x} \in \mathcal{M}$, pointing outwards.

Tangent space: We further use $T_{\mathbf{x}}\mathcal{M}$ to denote the tangent space of \mathcal{M} at point $\mathbf{x} \in \mathcal{M}$. For its representation, we choose two arbitrary but fixed orthogonal tangent vectors $\mathbf{u}(\mathbf{x})$, $\mathbf{v}(\mathbf{x})$, i.e.: $T_{\mathbf{x}}\mathcal{M} = \text{span}\{\mathbf{u}(\mathbf{x}), \mathbf{v}(\mathbf{x})\}$.

Distances: We use $\text{dist}_{\mathcal{M}}(\mathbf{x}, \mathbf{y})$ for $\mathbf{x}, \mathbf{y} \in \mathcal{M}$ to denote the *intrinsic* or *geodesic distance* between the two points \mathbf{x} and \mathbf{y} . A *geodesic* connecting \mathbf{x} and \mathbf{y} is a curve that has no geodesic curvature, which means that the derivative of the curve at a point \mathbf{s} projected to $T_{\mathbf{s}}\mathcal{M}$ is zero. We call the shortest geodesic connecting \mathbf{x} and \mathbf{y} a *shortest geodesic path* in \mathcal{M} .

Mappings and isometries: Consider two manifolds \mathcal{S} and \mathcal{T} , and a mapping $f : \mathcal{U} \rightarrow \mathcal{T}$ from an open subset $\mathcal{U} \subseteq \mathcal{S}$ to \mathcal{T} . Let $f_{\mathbf{u}}(\mathbf{x})$ and $f_{\mathbf{v}}(\mathbf{x})$ denote the partial derivatives of f with respect to the tangent space directions \mathbf{u} and \mathbf{v} in \mathbb{R}^3 of \mathcal{S} . The first fundamental form $\mathbb{I}_f(\mathbf{x})$ of f at point $\mathbf{x} \in \mathcal{U}$ is then given by:

$$\mathbb{I}_f(\mathbf{x}) = \begin{pmatrix} f_{\mathbf{u}}(\mathbf{x}) \cdot f_{\mathbf{u}}(\mathbf{x}) & f_{\mathbf{u}}(\mathbf{x}) \cdot f_{\mathbf{v}}(\mathbf{x}) \\ f_{\mathbf{v}}(\mathbf{x}) \cdot f_{\mathbf{u}}(\mathbf{x}) & f_{\mathbf{v}}(\mathbf{x}) \cdot f_{\mathbf{v}}(\mathbf{x}) \end{pmatrix}.$$

The function f is an *isometry* if and only if $\mathbb{I}_f(\mathbf{x}) = \mathbf{I}$ for all $\mathbf{x} \in \mathcal{U}$.

Equipped with this notation, we will now identify and analyze three classes of approaches for finding surface correspondences. The difference is in how the notion of *approximate* isometry is handled, leading to different complexity characteristics and algorithms.

2.1 Global Approximate Isometry

The global model assumes that *geodesic distances* are *global invariants* of the shape, being consistent at least up to an error margin $\nu > 0$ that accounts only for a small fraction of the object size. This means, the energy

$$E_{\text{global}} = \sup_{\mathbf{x}, \mathbf{y} \in \mathcal{U}} |\text{dist}_{\mathcal{S}}(\mathbf{x}, \mathbf{y}) - \text{dist}_{\mathcal{T}}(f(\mathbf{x}), f(\mathbf{y}))| \quad (1)$$

must be smaller than ν . For two points $\mathbf{x}, \mathbf{y} \in \mathcal{U}$ Eq. 1 corresponds to an additive error of at most ν , i.e.

$$|\text{dist}_{\mathcal{S}}(\mathbf{x}, \mathbf{y}) - \text{dist}_{\mathcal{T}}(f(\mathbf{x}), f(\mathbf{y}))| \leq \nu. \quad (2)$$

The global consistency criterion is sometimes modified to allow for a multiplicative error of ν instead, as

$$\max \left(\frac{\text{dist}_{\mathcal{S}}(\mathbf{x}, \mathbf{y})}{\text{dist}_{\mathcal{T}}(f(\mathbf{x}), f(\mathbf{y}))}, \frac{\text{dist}_{\mathcal{T}}(f(\mathbf{x}), f(\mathbf{y}))}{\text{dist}_{\mathcal{S}}(\mathbf{x}, \mathbf{y})} \right) \leq (1 + \nu). \quad (3)$$

The former error model considers absolute errors, while the latter one considers relative errors.

In case of exact isometry, i.e., $\nu = 0$, the set of matching candidates becomes strongly constrained. The isometry assumption has been used to embed the intrinsic geometry of a shape in a Euclidean space using multi-dimensional scaling, such that embeddings of isometric shapes become identical, which facilitates shape recognition and matching [27–29]. Alternatively, the geometry of shape \mathcal{S} can be embedded into shape \mathcal{T} using generalized multi-dimensional scaling, thereby computing a cross-parameterization directly [14].

Exact isometry results in a set of matching candidates with few degrees of freedom. Lipman and Funkhouser [10] have noticed that isometries are special cases of conformal maps, thereby having only the degrees of freedom given by the Möbius group, which are fixed by three pointwise matches on spherical topologies. Ovsjanikov et al. [11] have shown that even a single point match is sufficient to fix an isometry if the Laplace-Beltrami spectrum of \mathcal{S} is non-degenerate. In this paper, we exploit a different way to uniquely describe an isometry: fixing one point, a tangential direction, and the surface orientation is necessary and sufficient to specify an isometric mapping [23]. This provides the fewest possible degrees of freedom while still covering all cases including shapes with global intrinsic symmetries.

In case of small, global error margins, statistical triangulation algorithms can be applied that compute all correspondences from a few landmark matches [15, 16] or, similarly, by voting for several approximate solutions [10, 11]. Depending on the geometry of the shape, errors can become amplified so that a bit more than only the minimal set of initial correspondences are required [12]. Nonetheless, matching according to this model is efficient and can be considered a more or less solved problem.

The global approximate isometry criterion has also been employed to study matching partial overlaps of complete surfaces. Van Kaick et al. [30] use a pair of features to define a map that captures a local geodesic region between the two features, and show how this descriptor can be used for partial matching. Here, using two features instead of one is crucial because two features encode orientation information on the surface, while one feature does not.

The drawback of the global consistency model is its sensitivity to topological noise. To make globally consistent approaches more robust w.r.t. topological changes, several approaches have been proposed to perform a band-limited analysis in the eigenspace of the manifold’s Laplace-Beltrami operator [17, 20, 21, 31]. Unlike prior approaches based on embedding the intrinsic geometry of the shape directly [14, 27–29], these approaches successfully handle small topological errors. However, they break down

in the presence of large artifacts, such as wide contacts or missing pieces. A notable exception is the heuristic region growing approach by Sharma et al. [32]. It tries to match points with similar spectral signatures using an expectation-maximization framework, which has been shown to perform well in the presence of large contacts. Despite good practical performance, the algorithm is heuristic in nature and it remains unclear under which conditions it will find a correct solution. In particular, if local descriptors are not unique, the greedy region growing might fail and the EM-algorithm does not guarantee to recover a correct global match. In contrast, our region growing algorithm, which is based on propagation of metric information rather than potentially ambiguous descriptor matching, comes with the theoretical guarantee to find correct results for exact isometries while requiring an initialization from a small search space with very few degrees of freedom.

2.2 Global Approximate Isometry in Partial Regions

The global consistency model is incompatible with the notion of partial matching, since distances have to be measured on the complete surfaces \mathcal{S} and \mathcal{T} , which might not be available. Xu et al. [22] modify the criterion in Eq. 1 by restricting paths to the partially matched region \mathcal{U} :

$$E_{\text{partial}} = \sup_{\mathbf{x}, \mathbf{y} \in \mathcal{U}} |\text{dist}_{\mathcal{U}}(\mathbf{x}, \mathbf{y}) - \text{dist}_{f(\mathcal{U})}(f(\mathbf{x}), f(\mathbf{y}))|. \quad (4)$$

Note that E_{partial} considers again absolute errors

$$|\text{dist}_{\mathcal{U}}(\mathbf{x}, \mathbf{y}) - \text{dist}_{f(\mathcal{U})}(f(\mathbf{x}), f(\mathbf{y}))| \leq \nu \quad (5)$$

and can be modified to consider relative errors as in Eq. 3.

The drawback is that the shortest geodesic paths and thus the energy depend on the shape of the domain, which makes it difficult to optimize E_{partial} ; changing the domain \mathcal{U} influences which pairs of points are mapped in a geodesically consistent way. For this reason, Xu et al. restrict their method to considering geodesically convex regions \mathcal{U} , which are defined as regions where for any two points \mathbf{x}, \mathbf{y} in \mathcal{U} , there exists a shortest geodesic path between \mathbf{x} and \mathbf{y} in \mathcal{U} . In this special case, both E_{partial} and E_{global} measure shortest paths on \mathcal{S} and \mathcal{T} . The solution proposed by Xu et al. optimizes the scale of consistent regions and the consistent points separately, which leads to a rather complex algorithm. Further, the notion of a scale parameter is not canonically related to the original matching problem.

Sahilloglu and Yemez [33] consider the case where one of the surfaces is complete and the other an incomplete, deformed part of that surface. Using a coarse sampling and matching strategy between shape extremities, they can directly estimate a scale parameter between the two surfaces, which allows them to define a scale-invariant isometric distortion measure. This results in one-sided partial dense intrinsic matching up to an arbitrary scale. Their approach also allows matching of semantically similar, but

non-isometric complete surfaces. Our approach allows both surfaces to be incomplete, at the cost that they must be scaled consistently beforehand. Given real-world scanners often provide data in known units, our method is compatible with the scenario of matching surfaces acquired with different modalities. While in this paper we do not explore the latter scenario, our approach would be compatible with this task given a reliable way to estimate scale change during metric propagation, possibly using shape extremities or other intrinsic features.

Bronstein et al. [34] introduced a general framework to evaluate partial similarity using Pareto optimality. In case of partial intrinsic shape matching, this method aims to find large parts of two surfaces that are similar to each other, where similarity is measured according to Eq. 2. In practice, the parts are found using generalized multi-dimensional scaling. Raviv et al. [35] use a similar technique to find partial intrinsic symmetries.

2.3 Local Approximate Isometry

Another common way to relax the requirement of exact isometry towards approximate matching is to maintain the metric tensor in a least-squares-sense. Again, let $f : \mathcal{U} \rightarrow \mathcal{T}, \mathcal{U} \subseteq \mathcal{S}$ be a mapping between two manifolds, where \mathcal{U} is the open subset of \mathcal{S} that should be mapped to \mathcal{T} in a distortion minimizing way. We can measure the distortions for example by minimizing a matrix norm of the Green deformation tensor (difference of the first fundamental form to identity):

$$E_{\text{local}}(f) = \int_{\mathcal{U}} \|\mathbb{I}_f(\mathbf{x}) - \mathbf{I}\|_F^2 d\mathbf{x}. \quad (6)$$

This criterion is purely local and thus well suited for partial matching. It is worth noting that extrinsic elastic deformation techniques, such as [36–38] are closely related: They either include the preservation of the curvature tensor in the objective function to maintain the extrinsic shape, or apply Eq. 6 to the 3-manifold of the embedding Euclidean volume [39, 40]. All of these methods are designed for partial matching.

The problem with both intrinsic and extrinsic elastic matching models is that the search space becomes very large, rendering any approach based on exhaustive search prohibitively expensive. The structure of the search space can be approximately understood by a linearized analysis. In order to understand the degrees of freedom of the local matching model, we can use the tool of modal analysis of such elastic models, first introduced by Pentland and Williams [41] to the field of computer graphics (the aim of their paper was actually to speed up the simulation of extrinsic elastic deformations of solid objects in three-space). Modal analysis represent the deformation of an object as a linear combination of the eigenmodes of the object’s vibration, which are found using a spectral analysis of a linearized deformation energy.

Typically, these energies are related to the Laplacian of the deformed domain, thus leading to eigenvalues that are

only decreasing rather slowly. Many modes need to be retained to represent the space of low-energy (i.e., plausible) deformations adequately. If the local matching is not very stiff, the size of the search space explodes. This is intuitive: Permitting local deformations creates a large variety of permissible shape variants, for example by adding different local dents and combinations of those everywhere. Due to this large search space, it is impractical to match shapes purely based on the deformation model of Eq. 6. In order to reduce the large search space for elastic matching, existing methods use additional constraints, such as, template models [40], temporal coherence [39], or fairly large sets of coherent feature correspondences [37].

Recently, Kim et al. [18] proposed a new paradigm for the elastic matching problem. First, the method computes multiple dense maps between two shapes by assuming a global near-isometric deformation model. Multiple maps are obtained by fixing different triples of corresponding points for the computation of global conformal maps [10]. Second, the method computes local weights for pairs of maps, depending on their local adherence to isometry, and performs a spectral analysis to combine multiple weighted maps into a single global map. The key idea is to find cliques of similar mappings by clustering approximately compatible maps. The method was shown to perform well in many interesting cases. However, it cannot handle partial mapping; in particular constellations with large topological noise cannot be handled: Each global conformal map can be highly distorted due to the lack of global consistency. This introduces distortions in many of the local matches, and it is not always possible to remove these distortions in the final blended result (as demonstrated in Section 5).

Our clustering method (Section 4.4) is based on the same idea, but it combines *partial maps* instead of *global maps*, therefore avoiding the mentioned problems of artifacts due to only partial consistency. On the technical side, the main challenge is that we cannot measure the distance between all pairs of candidate maps but only between actually overlapping ones. This is a problem for the original spectral clustering, which we substitute by an alternative technique geared towards sparse pairwise constraints [24].

2.4 Previous Work on Local Isometry

In previous work, there have been a number of attempts to find local isometric mappings, similar to our approach. However, these did not consider mappings between general surfaces but only local planar parametrizations.

Different techniques have been proposed to locally parameterize the intrinsic geometry of a surface to a plane using a local approximate isometry criterion. Schmidt et al. [42] used exponential maps to transport a local coordinate system along a surface for the purpose of texture mapping. More recently, Schmidt [43] used transported exponential maps to produce a parameterization of a local surface patch to the plane that has low metric distortion. The local surface patch is provided through user interaction as

an input stroke on the surface. Malvaer et al. [44] propose an alternative parameterization based on an extension of polar coordinates to surfaces. In our work, we cross-parameterize local surface patches from \mathcal{S} to \mathcal{T} . This cross-parameterization is more challenging than a parameterization to the plane due to the arbitrary geometry of \mathcal{T} . Our cross-parameterization task is further complicated as in our application scenario, both \mathcal{S} and \mathcal{T} are scanned point clouds with missing surface information and scanner noise. Hence, the reviewed methods cannot be applied in a straight forward manner in our application.

3 Local Metric Matching

We saw that for general surface matching using a global approximate isometry criterion for partial matching (Eq. 4) is difficult since the domain changes, and that using a local approximate isometry criterion (Eq. 6) is difficult since this results in a large search space. In this section, we outline our new method: starting from a point $\mathbf{s} \in \mathcal{S}$ and attached direction in $T_{\mathbf{s}}\mathcal{S}$ with known correspondence $f(\mathbf{s}) \in \mathcal{T}$ and corresponding attached direction in $T_{f(\mathbf{s})}\mathcal{T}$, we find the largest domain \mathcal{U} such that Eq. 5 is satisfied.

The key assumption of our approach is that \mathcal{S} and \mathcal{T} can be matched using a global approximate isometry in some partial region \mathcal{U} containing \mathbf{s} , implying that Eq. 5 holds. Our goal is to find the largest region \mathcal{U} for which this assumption is satisfied. This assumption holds in scenarios where we know that \mathcal{S} and \mathcal{T} are actually related by a near-isometric map but the data does not comprise all of the original input and/or contains additional unrelated geometry or contacts. The most important practical example where this assumption holds is the acquisition of a surface that deforms near-isometrically but the scanning equipment introduces areas of missing data (shadowed to the scanner by other object parts) and cannot correctly resolve the topology in contact areas (such as the hand of a person being in contact with the body).

Our approach can be viewed as a hybrid approach between global and local approximate isometric matching. We use the assumption that \mathcal{S} and \mathcal{T} can be matched using a global approximate isometry in some partial region \mathcal{U} , and we compute \mathcal{U} by growing a region using a local approximate isometry criterion. This allows us to combine the advantage of the global methods of having a low-dimensional search space with the advantage of the local methods of being well suited to describe partial isometric matches.

Let Θ denote the parameter domain of all partial isometric matchings between \mathcal{S} and \mathcal{T} . Our goal is to compute all partial isometries $\{f_{\theta}, \mathcal{U}_{\theta}\}_{\theta \in \Theta}$ that map maximal subsets \mathcal{U} of \mathcal{S} to \mathcal{T} . The vector $\theta \in \Theta$ parametrizes the set of all such mappings.

In the following, we will discuss that:

- Isometric deformations of \mathcal{S} have three degrees of freedom.

- A partial isometry can be parametrized by $\Theta = \mathcal{S} \times SO(1) \times \mathcal{T} \times SO(1)$, where $SO(1)$ denotes the unit circle.
- There exists a global representation redundancy that identifies all choices of parameters $\mathbf{s} \in \mathcal{U} \subset \mathcal{S}$ and $\mathbf{d}_\mathbf{s} \in T_\mathbf{s}\mathcal{S}$.

3.1 Three Degrees of Freedom

For a (sufficiently) smooth Riemannian manifold \mathcal{S} fixing one point \mathbf{s} on \mathcal{S} , a tangential direction $\mathbf{d}_\mathbf{s}$ in $T_\mathbf{s}\mathcal{S}$, and a surface normal $\mathbf{n}_\mathbf{s}$ at \mathbf{s} suffices to specify an isometric mapping [23]. For details on the smoothness criteria, refer to [45]. The following proof sketch aims to give the intuition behind this statement for 2-dimensional surfaces embedded in \mathbb{R}^3 by showing that we can transfer the local coordinate frame defined by $\mathbf{d}_\mathbf{s}$ and $\mathbf{n}_\mathbf{s}$ to any point on \mathcal{S} in a canonical way.

We start with two definitions. The *injectivity radius* $\rho_\mathbf{s}\mathcal{S}$ at \mathbf{s} is the largest radius, such that for any point \mathbf{x} on \mathcal{S} with geodesic distance less than $\rho_\mathbf{s}\mathcal{S}$ from \mathbf{s} , there is a unique shortest geodesic path between \mathbf{s} and \mathbf{x} . The *injectivity radius of \mathcal{S}* is defined as $\rho\mathcal{S} = \inf_{\mathbf{s} \in \mathcal{S}} (\rho_\mathbf{s}\mathcal{S})$.

For closed surfaces, the injectivity radius can be bounded from below by the minimum of $\pi/\sqrt{\sup K}$, where K is the Gaussian curvature, and half of the length of the smallest periodic geodesic [23, Thm. 10]. It follows that $\rho\mathcal{S} > 0$ holds for closed Riemannian surfaces with finite Gaussian curvature. For general surfaces with boundary, the injectivity radius may become zero. Note that in this paper, we consider Riemannian surfaces with finite Gaussian curvature with boundaries. The boundaries are present because the closed Riemannian surface of interest is only partially observed by the acquisition device. In this special case, the injectivity radius is still positive.

Imagine that \mathcal{S} is covered by overlapping regions of intrinsic radius less than $\rho\mathcal{S}$. These regions are all topologically equivalent to disks. Consider a disk D containing a point \mathbf{s} as shown in Figure 1. In a small neighborhood of \mathbf{s} , D is arbitrarily close to the tangent plane $T_\mathbf{s}\mathcal{S}$. Note that \mathbf{s} , $\mathbf{d}_\mathbf{s}$, and $\mathbf{n}_\mathbf{s}$ fix a local orthonormal coordinate frame in \mathbb{R}^3 . This coordinate frame can be transported to a point $\mathbf{x} : \mathbf{x} \in D, \mathbf{x} \neq \mathbf{s}$ along the (unique) shortest geodesic path $P_{\mathbf{s}\mathbf{x}}$ from \mathbf{s} to \mathbf{x} in D by parallel transport, which can be thought of as repeatedly projecting the direction $\mathbf{d}_\mathbf{s}$ to the tangent planes of consecutive points along $P_{\mathbf{s}\mathbf{x}}$ in infinitesimally small steps (for details on parallel transport see for example Berger [23, Chapter 3.1]). Let $\mathbf{d}_\mathbf{x}$ denote the transported direction. The transferred direction lies in the tangent plane $T_\mathbf{x}\mathcal{S}$, and can again be used to fix a local orthonormal coordinate frame at \mathbf{x} . This transfer can be repeated by chaining together disks until every point \mathbf{y} on \mathcal{S} was assigned a fixed tangential direction $\mathbf{d}_\mathbf{y} \in T_\mathbf{y}\mathcal{S}$ by parallel transport along a path connecting \mathbf{s} and \mathbf{y} that consists of an arbitrary but fixed sequence of (unique) shortest geodesic paths within the chained disks. Note that by construction, we only use intrinsic information to propagate the direction $\mathbf{d}_\mathbf{s}$ to the entire surface. Hence, the resulting




Figure 1: Propagating a local coordinate system along \mathcal{S} .

local coordinate frames are invariant with respect to isometric deformations of \mathcal{S} when encoded relative to fixed local coordinate systems $\mathbf{u}(\mathbf{y}), \mathbf{v}(\mathbf{y}), \mathbf{n}(\mathbf{y})$ on \mathcal{S} .

This implies that any isometric deformation of an *oriented* surface \mathcal{S} can be specified using three degrees of freedom: one point \mathbf{s} on \mathcal{S} (accounting for two degrees of freedom) and a direction in the tangent plane of \mathbf{s} .

3.2 Representation

We can use the fact that any isometric deformation of an oriented surface can be specified using three degrees of freedom to derive a representation θ for intrinsic mappings. Specifically, identifying one corresponding point and one corresponding tangential direction completely determines an isometric mapping between two oriented surfaces. More formally, to define an isometric mapping between (subsets of) \mathcal{S} and \mathcal{T} , it suffices to specify a point $\mathbf{s} \in \mathcal{S}$, its intrinsically corresponding point $\mathbf{t} \in \mathcal{T}$, a tangential direction $\mathbf{d}_\mathbf{s}$ in $T_\mathbf{s}\mathcal{S}$, and its intrinsically corresponding tangential direction $\mathbf{d}_\mathbf{t}$ in $T_\mathbf{t}\mathcal{T}$.

Starting from this information, and assuming that \mathcal{S} and \mathcal{T} are isometric, we can propagate the correspondence information by mapping the metric structure of \mathcal{S} onto \mathcal{T} as follows. Starting from the corresponding points \mathbf{s} and \mathbf{t} along with the corresponding directions $\mathbf{d}_\mathbf{s}$ and $\mathbf{d}_\mathbf{t}$, we can propagate the correspondence information to a sufficiently small geodesic neighborhood $N_\mathbf{s}$ of \mathbf{s} by simultaneously walking along corresponding geodesic paths starting at \mathbf{s} and \mathbf{t} , respectively, and by matching points that are reached at the same time. Here, $N_\mathbf{s}$ is sufficiently small if its geodesic radius is below $\rho\mathcal{S}$. Once the correspondence information is computed for $N_\mathbf{s}$, we continue propagating the correspondence information from a point close to the boundary of $N_\mathbf{s}$ to its geodesic neighborhood, and iterate until every point on \mathcal{S} has a correspondence.

The assumption that \mathcal{S} and \mathcal{T} are (near-) isometric can also be used to detect the boundary of the largest region $\mathcal{U} \subseteq \mathcal{S}$ containing \mathbf{s} for which the mapping is near-isometric. The reason is that the propagation algorithm allows us to measure the difference in intrinsic geometry in newly mapped parts of \mathcal{S} and \mathcal{T} directly. Hence, we can stop the region growing algorithm if a newly added correspondence would induce a stretching larger than ν .

Figure 2 illustrates the near-isometric region growing process. The plane on the left (\mathcal{S}) and the plane with a hill on

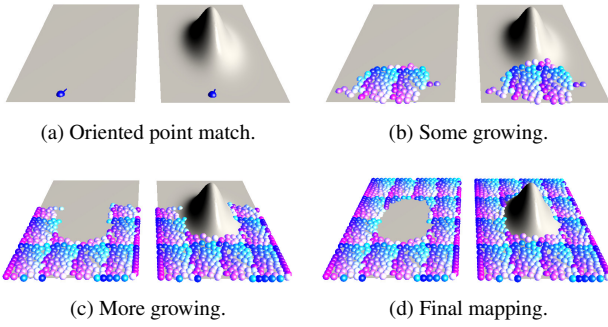


Figure 2: Illustration of the near-isometric region growing process. Corresponding points share the same color.

the right (\mathcal{T}) are isometric except for the hill. Beginning with a point and direction match in Figure 2a, the isometric region grows outward in all directions where the isometry condition is locally satisfied. This way, the largest near-isometric partial match is identified, as shown in Figure 2d. Note that this region can have a complex topology and geometry.

3.3 Representation Redundancy

The representation discussed above contains redundant information: Infinitely many θ may represent the same near-isometric mapping.

The first degree of redundancy is the choice of the direction \mathbf{d}_s in the tangent plane $T_s\mathcal{S}$. Changing this direction merely rotates the field of directions \mathbf{d}_x in $T_x\mathcal{S}$ (and similarly, the field of corresponding directions \mathbf{d}_y in $T_y\mathcal{T}$). Hence, the choice of this direction has no influence on the final result. To remove this redundancy, we start from an arbitrary but fixed direction \mathbf{d}_s and precompute and fix \mathbf{d}_x for all x on \mathcal{S} .

The second redundancy is the choice of the start point s . Let $\mathcal{U} \subseteq \mathcal{S}$ be the maximal set within which an isometry $f: \mathcal{U} \rightarrow \mathcal{T}$ can be constructed. We can replace s in θ by any other point $s' \in \mathcal{U}$. See Fig. 3. This requires an update of the remaining parameters. The direction $\mathbf{d}_{s'}$ is set to the precomputed value (see above), and the remaining parameters can be updated using the computed mapping f . More specifically, $\mathbf{t}' = f(s')$ and the direction of $\mathbf{d}_{t'}$ is set to the direction of $f(s' + \epsilon\mathbf{d}_{s'}) - f(s')$, where ϵ is set small enough that $s' + \epsilon\mathbf{d}_{s'}$ is in \mathcal{U} . Thus, in the case of a global isometry, or when we know beforehand the isometric region \mathcal{U} , the mapping f has three degrees of freedom. In the general partial isometry case, however, where \mathcal{U} is unknown, the starting point $s \in \mathcal{S}$ is not fully redundant; it still selects the equivalence class that represents a certain partial map; computationally, this is the patch to be matched by the propagation algorithm. This does not increase the complexity strongly as we just have to restart the matching in case multiple partial matches exist. Ideally, we would sample one starting point per partial isometric region, which in practice will be far fewer than there are samples on \mathcal{S} . While we do not determine this lower-bound beforehand, we maintain low complexity




Figure 3: For a given set \mathcal{U} and a corresponding isometry (shown in blue), the choice of the starting point s is arbitrary.

by using features to identify potential starting points, and marking a starting point s as redundant if another starting point s' produces an isometric region \mathcal{U} containing s . In case of a mismatch, it can be discovered quickly if the target area does not match, which will become evident within constant time.

In summary, all mappings represented by $s', \mathbf{d}_{s'}, \mathbf{t}', \mathbf{d}_{t'}$ form an equivalence class in the parameter space Θ . Since we can remove one degree of redundancy by fixing the tangential directions on \mathcal{S} , for each mapping f we have two degrees of freedom that vary among equivalent maps (the choice of the start point on \mathcal{S}), which along with the manifold structure of \mathcal{U} implies that each equivalence class forms a 2-manifold in Θ , which can be computed directly using the propagation algorithm introduced in Section 3.2.

In practice, we can take advantage of this representational redundancy as follows. Since \mathcal{S} and \mathcal{T} are discretized and corrupted by noise, the error of the correspondence information computed using the propagation algorithm increases with increasing distance from the start point of the propagation. Hence, it is possible that the propagation stops prematurely due to discretization artifacts and the influence of noise, thereby identifying a region \mathcal{U} that is smaller than the correct solution. Thanks to the redundancy in the representation, we can start the propagation algorithm from multiple oriented point pairs, identify a set of equivalent mapping functions f_i , and then into a single mapping function f that covers a larger area of \mathcal{S} and is less influenced by noise than the individual f_i . The following section discusses a direct implementation of this theoretically motivated method, which we will use to compute correspondences of noisy scanner data.

4 Pairwise Intrinsic Matching

From the preceding analysis, we derive an algorithm for computing partial near-isometries between two surfaces \mathcal{S} and \mathcal{T} . We start our discussion by outlining how surfaces are represented and how basic steps of the algorithm are implemented (Section 4.1). Our direct, non-optimized implementation is based on enumerating the non-equivalent choices by selecting different oriented point matches (Section 4.2), growing the isometric region locally from there until no more points locally satisfy ν (Section 4.3), and finally clustering the partial maps into equivalence classes (Section 4.4). Figure 4 gives a graphical overview of our matching pipeline.

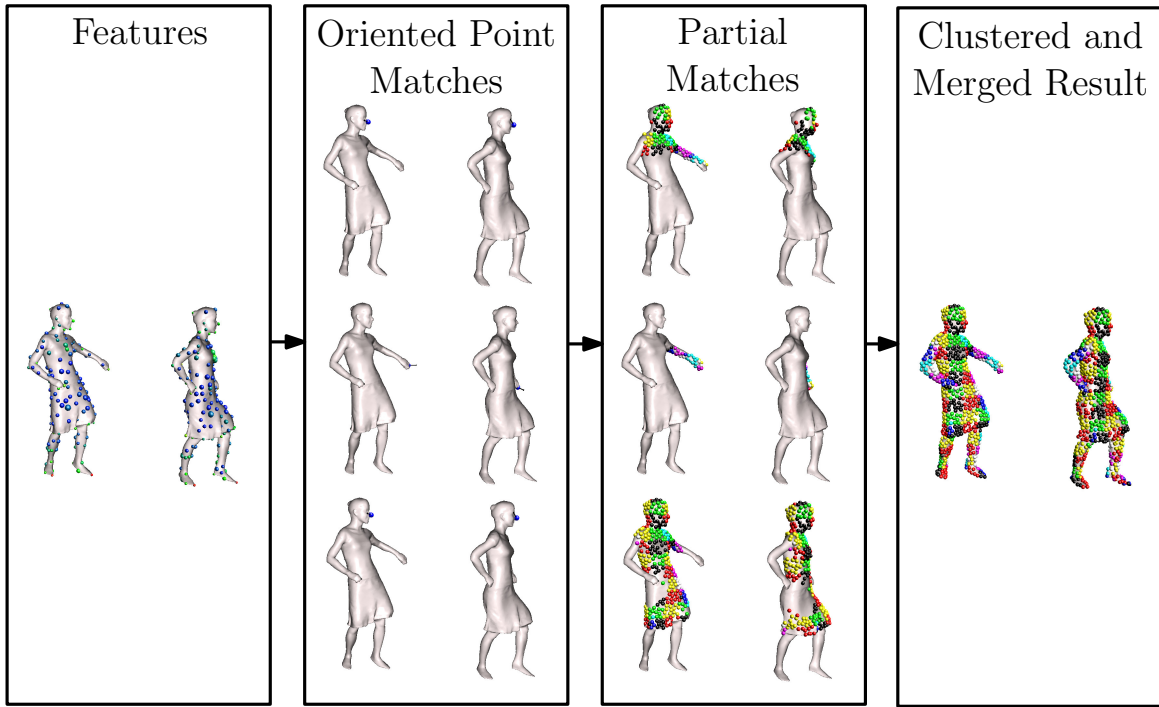


Figure 4: Overview of the pairwise matching pipeline. First, features are detected, second, oriented point matches are computed, third, starting from the oriented point matches, partial isometric matches are found, and finally, different partial matches are clustered and merged.

4.1 Surface Representation

In our implementation, the surfaces \mathcal{S} and \mathcal{T} are either represented as oriented point clouds or meshes. While most parts of the algorithm can be extended to discretized surfaces in a straight forward manner, for some parts of the algorithm we require a continuous surface representation. We obtain a continuous surface representation as implicit moving-least-squares (MLS) manifolds using the robust method of Öztireli et al. [46]. Using this method, a discretized surface \mathcal{S} is represented continuously as the zero level-set of an implicit function derived from the oriented vertices of \mathcal{S} .

Using a MLS representation allows a point to be projected to a continuous representation of \mathcal{S} in the case where \mathcal{S} is given as oriented point cloud or mesh. In the following, whenever we refer to projecting a point \mathbf{x} onto a discretized surface \mathcal{S} , this projection is implemented as projecting \mathbf{x} to the MLS representation derived from \mathcal{S} .

One basic operation needed by our algorithm is the computation of geodesic distances and paths on \mathcal{S} . In some parts of the algorithms, rough estimates of geodesic distances and paths suffice, and these are computed using Dijkstra’s algorithm. In other parts of the algorithm, it is crucial to have accurate estimates of geodesic distances and paths. In these cases, we initialize a geodesic path P to the Dijkstra path and refine P by minimizing the length of P using the constraint that P must not leave \mathcal{S} . This optimization is carried out iteratively using a conjugate gradient method. After each step, P is projected back to \mathcal{S} . In the following, we refer to these refined geodesic distances and paths as *smoothed geodesic distances and paths*.

A core part of our approach is to compare distances measured on the target surface to distances measured on the source for corresponding points. Comparing geodesics between all pairs of correspondences quickly becomes prohibitively expensive and would limit the practical applicability of our algorithm. To remedy this, we construct a topology hierarchy on \mathcal{S} similar to [47] as follows. We define level 0 of the hierarchy to be the original set of vertices and their connectivity—either the original triangle mesh or the k -nn graph for a point cloud. (In all our experiments involving point clouds, we set $k = 8$.) The sample spacing ϵ_0 is defined as the average edge length. Level 1 of the hierarchy is constructed by selecting an evenly spaced subset of the level 0 vertices and connecting vertices in a topology preserving way. Subsequent levels $j + 1$ are constructed in the same way as a subset of level j . At each level the subset for the next level is determined by doubling the desired sample spacing, $\epsilon_{j+1} = 2\epsilon_j$. At a coarser level of this hierarchy fewer vertices are connected to any others, but at a greater distance to each other. In our algorithm, we only consider geodesic distances between vertices that share an edge in at least one level of the hierarchy, which results in a sparse set of distances to be optimized, even for a dense set of correspondences.

This hierarchical structure ensures that locality is respected, as only geodesics between points that are neighbors in some level of the hierarchy are considered. Having a structure that respects locality is crucial to allow for partial matching. The number of levels in the hierarchy determines the trade-off between local and global isometry constraints. We use 5 levels in all experiments in this paper.

Note that during the execution of the matching algorithm, the vertices on \mathcal{S} are fixed, as we aim to find a correspondence for every vertex of \mathcal{S} on \mathcal{T} (if such a correspondence exists). Hence, we can precompute all geodesic distances on \mathcal{S} for vertices connected in some level of the hierarchy. This way, only distances on \mathcal{T} need to be updated during the metric optimization.

4.2 Finding Oriented Point Matches

In principle, by exhaustively trying all possible starting points and directions, our algorithm will recover all partial isometries. However, since this is infeasible, we reduce the search space using sparse feature matches. This section outlines an algorithm to find features matches. However, note that this part is not a novel contribution of this paper and only given for completeness; in principle, any feature matches can be used to reduce the search space, as demonstrated in Section 5, where we show a result using image-based feature matches from a multi-view capture setup. We denote the feature descriptor at a point $\mathbf{s} \in \mathcal{S}$ with a vector $\mathbf{D}_{\mathcal{S}}(\mathbf{s})$, and similarly for points on \mathcal{T} .

We identify features on \mathcal{S} and \mathcal{T} using the geodesic fingerprint descriptor $\mathbf{D}_{\mathcal{S}}(\mathbf{s})$ [48], which compares how an isocontour of the geodesic distance from a point $\mathbf{s} \in \mathcal{S}$ deviates in length from the isocontour of the same 2D Euclidean distance. We use 10 isocontours in all experiments, for geodesic radii between r_{\min} and r_{\max} . The values of r_{\min} and r_{\max} need to be varied slightly depending on the amount of noise in the input data, and are discussed in Section 5. We define the *distinctiveness* $F_{\mathcal{S}}(\mathbf{s})$ of a point $\mathbf{s} \in \mathcal{S}$ as the sum of L_1 distances to the rest of the vertices of \mathcal{S} in descriptor space. Features are points that locally maximize distinctiveness. The left-most box in Figure 4 shows features color-coded by distinctiveness (red for most distinct, blue for least distinct).

To find feature matches, we begin by computing the Cartesian product of L_2 descriptor distances between features on \mathcal{S} and \mathcal{T} . The vast majority of these are not correct mappings between the surfaces. We filter these potential feature matches using both the L_2 distances between descriptors and the distinctiveness of the features. More precisely, for each feature \mathbf{s} on \mathcal{S} , we only consider the K features of \mathcal{T} that have the closest descriptor matches to \mathbf{s} . (We use $K = 10$ in all our experiments.) The *initial dissimilarity* $\delta_{\mathbf{s},\mathbf{t}}^{(init)}$ between two features $\mathbf{s} \in \mathcal{S}$ and $\mathbf{t} \in \mathcal{T}$ is defined as

$$\delta_{\mathbf{s},\mathbf{t}}^{(init)} = -\log(F_{\mathcal{S}}(\mathbf{s})) + \omega_{\mathbf{D}} \|\mathbf{D}_{\mathcal{S}}(\mathbf{s}) - \mathbf{D}_{\mathcal{T}}(\mathbf{t})\|_2^2, \quad (7)$$

where $\omega_{\mathbf{D}}$ is a weight. (We use $\omega_{\mathbf{D}} = 400$ in all our experiments.) This dissimilarity is minimized for features with high distinctiveness that are similar in descriptor space.

This procedure often produces a good set of sorted feature matches on clean data. For noisy data from real scanners, however, the descriptors will be less discriminative, and considering spatial relations between features in addition to descriptors and distinctiveness leads to more reliable feature matches.

To do this, we iteratively build (possibly overlapping) clusters C_i of consistent feature matches. In each itera-

tion, the feature match (\mathbf{s}, \mathbf{t}) with the next lowest $\delta_{\mathbf{s},\mathbf{t}}^{(init)}$ is chosen as a starting point for C_i . The cluster C_i is built by repeatedly adding the close-by feature match that has the lowest dissimilarity to C_i . More precisely, let $C_i = \{(\mathbf{s}_j, \mathbf{t}_j)\}$. In the next step, all features $\mathbf{s}' \notin C_i$ that are neighbors of \mathbf{s}_j in the topology hierarchy are considered, along with their K potential matches $\mathbf{t}' \in \mathcal{T}$. The dissimilarity $\delta_{C_i,(\mathbf{s}',\mathbf{t}')}$ between a feature match $(\mathbf{s}', \mathbf{t}')$ and cluster C_i is defined as

$$\delta_{C_i,(\mathbf{s}',\mathbf{t}')} = \delta_{\mathbf{s},\mathbf{t}}^{(init)} + \omega_C \sum_{(\mathbf{s}_j,\mathbf{t}_j)} |\text{dist}_{\mathcal{S}}(\mathbf{s}_j, \mathbf{s}') - \text{dist}_{\mathcal{T}}(\mathbf{t}_j, \mathbf{t}')|^2, \quad (8)$$

where ω_C is a weight. (In all our experiments, we set ω_C to one over $8 \times$ the square of the average edge length on \mathcal{S} .) We repeatedly add the match $(\mathbf{s}', \mathbf{t}')$ with smallest dissimilarity $\delta_{C_i,(\mathbf{s}',\mathbf{t}')}$ to C_i as long as $\delta_{C_i,(\mathbf{s}',\mathbf{t}')}$ is below a threshold. (We use threshold $11.5 \approx -\log 10^{-5}$ in all our experiments.) We stop adding new clusters once the sum of the cardinalities of the clusters C_i exceeds the initial number of feature matches.

Note that the above clustering scheme is equivalent to modeling both the descriptor distances and the summed stretches of geodesic distances of a feature match to a cluster as normally distributed. Hence, the above stopping criteria correspond to stopping the clustering once the joint probability of a match belonging to a cluster becomes small.

After the clustering, we place the feature matches (\mathbf{s}, \mathbf{t}) in a min-priority queue, so that we start isometric region growing first from the matches we expect to be most reliable. A feature match (\mathbf{s}, \mathbf{t}) is assigned priority ∞ if (\mathbf{s}, \mathbf{t}) is not part of any cluster C_i and priority $\min_{C_i: (\mathbf{s},\mathbf{t}) \in C_i} \delta_{C_i,(\mathbf{s},\mathbf{t})}$ otherwise. We repeatedly take the minimum element from queue and use it to generate partial isometric mappings using the region growing of Section 4.3.

However, so far we have only established a positional correspondence between features, and we need a directional correspondence as well to fix the partial isometry we wish to grow. To establish direction, we build a min-priority queue as outlined above, but only with the subset of features in the neighborhood of the current positional match in the topology hierarchy. Matches of neighboring features allow us to find corresponding tangent plane directions from corresponding smoothed geodesic paths between feature points.

Once a partial isometry has been computed, we increase the priority of feature matches that are redundant given the already computed partial match. If the same source and target points are already matched, by our model, the result of running the growing again from those same points will be equivalent.

To keep the run-time of our method bounded, we stop after a fixed number of oriented point matches have been tried. (We use 200 in all our experiments.) A more general stopping criterion could be devised based on determining the maximal coverage of \mathcal{S} and \mathcal{T} subject to a consistent mapping. The second box from the left of Figure 4 shows three oriented feature matches found using this method.

4.3 Isometric Region Growing

Starting from an oriented point match $\mathbf{s}, \mathbf{d}_s, \mathbf{t}, \mathbf{d}_t$, we grow the region \mathcal{U} by adding matches incrementally in the local neighborhood of the boundary of \mathcal{U} . For a new point \mathbf{u} near the boundary of \mathcal{U} such that $\mathbf{u} \notin \mathcal{U}$, let $N_1(\mathbf{u})$ denote the neighborhood of \mathbf{u} in level 1 of our topology hierarchy. We use parallel transport along corresponding directions in \mathcal{S} and \mathcal{T} emanating from an oriented point match $\mathbf{s}', \mathbf{d}_{s'}, \mathbf{t}', \mathbf{d}_{t'}$, where $\mathbf{s}' \in \mathcal{U} \cap N_1(\mathbf{u})$ is in a local neighborhood of \mathbf{u} . We know the full path between \mathbf{s}' and \mathbf{u} on \mathcal{S} , and from $\mathbf{d}_{s'}$ and $\mathbf{d}_{t'}$ we know the corresponding start direction on \mathcal{T} . Hence, we can transport the start direction along corresponding paths on \mathcal{S} and \mathcal{T} until we have traveled $\text{dist}_{\mathcal{S}}(\mathbf{s}', \mathbf{u})$. In practice, we implement parallel transport by moving in small steps in the tangent plane, and by reprojecting the resulting point to the surface and updating the tangent to the path. We use smoothed geodesic paths to transport the position of the new match on \mathcal{T} to reduce discretization errors for differently sampled surfaces. For robustness, we use all available oriented point matches in $N_1(\mathbf{u})$ and take the Riemannian center of mass [49] of the transported points, where all transported points have equal mass.

The newly matched source point \mathbf{u} is added to \mathcal{U} only if it locally respects the stretch factor ν as given in Eq. 5. This is necessary for two reasons. First, subsequent matches will be estimated based on the assumption that the existing matches are near-isometric. Second, as discussed next, we use a nonlinear least-squares optimization to refine the positions of the matched points on \mathcal{T} , which effectively distributes the error evenly over the matched region. To avoid introducing errors, it is therefore important to verify that each newly added point match locally respects the stretch factor ν .

To reduce the effect of quantization errors and noise, we optimize the metric matching of \mathcal{U} using a non-linear optimization technique [14] every time the area of \mathcal{U} has doubled. This means more frequent optimizations at the start of the growing process. This optimization reduces the amount of drift as it re-aligns the matched regions while taking into account long geodesics, between neighbors in the top level of our topology hierarchy, in \mathcal{U} . For increased efficiency, we only consider edges in the topology hierarchy of \mathcal{S} (explained in Section 4.1) during the optimization.

An important difference to the global approach used by Bronstein et al. [14] is that we optimize the metric only using geodesic paths which are entirely *within* \mathcal{U} and $f(\mathcal{U})$. This models the isometry criterion in partial regions and is crucial to handling topology changes and missing data. Note that such defects may cause large differences between $\text{dist}_{\mathcal{S}}(\mathbf{x}, \mathbf{y})$ and $\text{dist}_{\mathcal{U}}(\mathbf{x}, \mathbf{y})$, as well as $\text{dist}_{\mathcal{T}}(f(\mathbf{x}), f(\mathbf{y}))$ and $\text{dist}_{f(\mathcal{U})}(f(\mathbf{x}), f(\mathbf{y}))$, respectively. The second box from the right in Figure 4 shows the first three partial isometries found by growing isometrically from oriented feature matches in our priority queue.

The proper value for the stretching threshold ν depends on a number of factors: material properties, the resolu-

tion at which the surface is sampled (as it affects the accuracy of the Dijkstra paths), and the noise of the acquisition process. In our experiments, we do not consider material properties, and we assume that the acquisition noise has an equal influence on both source and target. Hence, we set ν to $0.5\epsilon_0$ to account for quantization effects in the computation of geodesics.

4.4 Combining Equivalent Partial Maps

It remains to identify and merge a set of partial mappings that represent the same mapping function f . If the surfaces were related by exact isometries and noise was negligible, the following step would not be required. However, for real-world data, the identification of functions that are approximately equivalent improves the results substantially.

The problem at this point similar to the blending problem by Kim et al. [18]. Recall that their approach uses a spectral method to find blending weights for different maps. This is a good approach in their case as blending weights are given as the solution to a quadratic energy function. Note that since in our model, equivalent mappings form a 2-manifold in parameter space Θ , and the parameter space Θ is non-linear, it is not appropriate to use a spectral approach.

However, we can take advantage of the property that equivalent mappings form a 2-manifold in Θ . We employ the agglomerative clustering algorithm of Zhang et al. [24] for discovering manifold structures in high-dimensional data based on the in-degree and out-degree of the nearest-neighbor graph of points in high-dimensional space. In our case, we consider the nearest neighbor graph of partial isometric mappings, where the dissimilarity between these points in Θ is measured as follows.

We compare different maps f_i and f_j using a dissimilarity measure based on their domains \mathcal{U}_i and \mathcal{U}_j :

$$d_{\Theta}(f_i, f_j) = W_1 \int_{\mathcal{U}_{ij}} \text{dist}_{\mathcal{T}}(f_i(\mathbf{x}), f_j(\mathbf{x})) d\mathbf{x} + W_2 \int_{f_i(\mathcal{U}_{ij}) \cap f_j(\mathcal{U}_{ij})} \text{dist}_{\mathcal{S}}(f_i^{-1}(\mathbf{y}), f_j^{-1}(\mathbf{y})) d\mathbf{y}^{(9)}$$

where $\mathcal{U}_{ij} = \mathcal{U}_i \cap \mathcal{U}_j$, $W_1 = \frac{1}{A(\mathcal{U}_{ij})}$, $W_2 = \frac{1}{A(f_i(\mathcal{U}_{ij}) \cap f_j(\mathcal{U}_{ij}))}$, and $A(\cdot)$ denotes the surface area. In practice, we compute a discrete version of this dissimilarity by replacing integrals over regions by sums over vertices in the region. We cluster different mappings together until the maximum affinity between any two clusters is greater than a threshold ρ . Affinity is computed from the weighted graph degree between mappings, where the weights have a double-exponential fall-off as d_{Θ} increases. See Zhang et al. [24] for details. While the direct relation of ρ to the allowed stretch is not easily defined, it should be lower when we want to allow for greater stretching between partial maps. We only have to adjust this value in a few cases in our experiments, as discussed in Section 5. Following the clustering, we select for our final mapping the cluster with the highest intra-cluster affinity, or connectivity, as proposed by Zhang et

al. [24]. This most often correlates with the cluster that covers the largest portion of the surface.

We merge clustered maps by computing a weighted geodesic average on \mathcal{T} for each source point in the union of the mapped regions as

$$f(\mathbf{x}) = \arg \min_{y \in \mathcal{T}} \sum_i w_i(\mathbf{x}) \text{dist}_{\mathcal{T}}(y, f_i(\mathbf{x})), \quad (10)$$

where f_i are the clustered partial isometries, which we want to merge into f . The weights are computed as an exponential distribution in the geodesic distance from the starting point match s_i as $w_i(\mathbf{x}) = \exp(-\lambda_{ds} \text{dist}_{\mathcal{U}_i}(\mathbf{x}, s_i))$, reflecting that we expect errors to accumulate as the growing proceeds because of discretization artifacts, data noise, and deviations from isometry. We set $\lambda_{ds} = 1/(5\epsilon_1)$, where ϵ_1 is the sample spacing of the level 1 of the topology hierarchy described in Section 4.1. Note that this is equivalent to finding the Riemannian center of mass [49] of the estimates $f_i(\mathbf{x})$. The right-most box of Figure 4 shows the final clustered and merged result.

5 Experiments

To validate our theoretical analysis, we evaluate a direct implementation of our algorithm and compare our results to state of the art approaches. In the following, we demonstrate that our method achieves results that are either comparable or superior to the state of the art, which demonstrates that our algorithm not only has theoretical advantages, but is applicable in practice as well.

5.1 Implementation Details

We implemented the algorithm described in Section 4 in C++, and conducted our evaluations on a standard laptop PC. During the evaluation, all but two types of parameters are fixed. The first type of parameters that is varied is $\{r_{min}, r_{max}\}$, which controls the size of the neighborhood used to compute surface descriptors. If the radii are set higher, then the method is more robust with respect to noise at the cost of potentially missing features of small scale. In our experiments, we only use two settings for these parameters. The first setting, $r_{min} = 0.9R$ and $r_{max} = 1.7R$, is for relatively clean data, and the second setting, $r_{min} = 1.5R$ and $r_{max} = 3.4R$, is for noisy data. Here, R is set to 5% of the diameter of \mathcal{S} . The second parameter that is varied is the threshold ρ used to control the clustering. Lower values of ρ allow for less isometric patches to be clustered together. We vary $\rho \in \{0, 1, 1.9\}$ in our experiments.

For the examples discussed below, our algorithm takes between 30 minutes and 8 hours to compute the final result. To give an idea of the distribution of the time, we discuss the running time for one pair of models (the space-carved samba models shown in Fig. 9) in more detail. For this pair, finding oriented feature matches takes about 2

minutes, growing partial mappings takes about 1.5 hours, clustering the mappings takes about 9 minutes, and merging the patches takes about 44 minutes. Hence, the total time to compute the results is about 2.4 hours. Note however, that the running time of our method depends significantly on the distinctiveness of the intrinsic geometry of the surfaces, relative to the noise level. For example, the template-fitted samba models shown in Fig. 6 take about 1.5 hours in total, despite having more than twice as many vertices as the space-carved versions, because the approximate isometry criterion is more discriminative (geodesic distances are less perturbed by surface noise).

5.2 Comparison to State of the Art

We compare the performance of our algorithm against four existing methods, namely heat kernel maps (HKM) [11], blended intrinsic maps (BIM) [18], the method of Sharma et al. [32], and the method of Tevs et al. [50]. We select these methods for comparison because they represent the state of the art for matching between surfaces. More specifically, HKM is derived from the theoretical complexity of isometric mappings, BIM can match surfaces that exhibit local deviations from isometry, and the methods of Sharma et al. and Tevs et al. are heuristics that have been specifically designed for matching partial data with topological noise. For HKM we use our own implementation, which uses two feature correspondences to initialize the mapping, for BIM we use the authors' implementation, for comparisons with Sharma et al., we run our code on their data, and for the method of Tevs et al., the authors were kind enough to run their algorithm on our data.

We show comparative evaluations on a variety of types of data. The first type is a synthetic data set that helps in understanding the major difference between our approach and HKM, and illustrates the partial isometric matching model. The second type is data acquired using either a laser scanner or an image-based reconstruction system that was processed by fitting a template to the data. In this case, we treat the result of the template fitting as ground truth. The third, and most challenging, type is unprocessed real-world data acquired using different acquisition systems.

To compare our approach to previous methods, we use two evaluation methodologies. In cases where ground truth is available, we compare quantitatively by evaluating the accuracy of different results with respect to the ground truth. For data that has no ground truth, we rely on visual evaluation. Our visualization scheme is as follows. A texture on \mathcal{S} is mapped to corresponding points on \mathcal{T} , and regions of \mathcal{T} that have no correspondence in \mathcal{S} are colored red. As a texture we combine constant coloring of semantically distinct parts (where applicable) with a checkerboard pattern. This type of texture simultaneously shows both global semantic accuracy and fine-scale distortion.

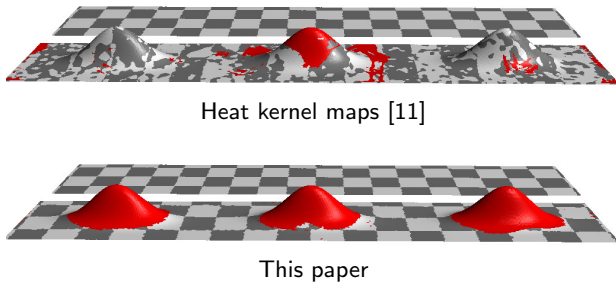


Figure 5: Results of our method and HKM on globally non-isometric data. Red indicates unmatched area.

5.3 Synthetic Data

We start by comparing our algorithm against HKM using the synthetic example shown in Fig. 5, where we map from a plane to a plane with three peaks. This experiment illustrates the difference between a global isometric matching model and partial one: the two surfaces are globally non-isometric, but the planar part is isometric. In this example, we use two ground truth correspondences to initialize both algorithms to remove differences due to different feature matching approaches. Since HKM aims to map the shapes using a global isometry, the result maps planar parts to the peaks, while our method successfully detects the largest part of the surfaces that can be isometrically mapped.

5.4 Template-Fitted Scan Data

Next, we consider acquisitions of real-world data that was processed by fitting a template shape to the raw data. For all experiments in this section, we use $r_{min} = 0.9R$ and $r_{max} = 1.7R$.

We first use two frames of the samba sequence by Vlasic et al. [51]. These frames are locally very close to isometric, but globally have high non-isometric distortion due to the dress being connected to the legs. We therefore set $\rho = 0$ in the clustering step. Vlasic et al. provide a processed version of the frames, where a template was fitted to the data. This processing ensures that the models have the same topology, and the processed data can be used as ground truth correspondence.

We compare our method to HKM and BIM using the processed frames of the samba sequence. Figure 6 shows the results. Note that while HKM leads to a result with visual artifacts, the results of BIM and our method are visually pleasing. Furthermore, since we have ground truth correspondences, we show the cumulative error distributions for all three methods in Figure 7. Here, geodesic error is measured as a fraction of the square root of the surface area of \mathcal{T} . Note that our method numerically outperforms the two other methods.

The experiments conducted so far have shown that HKM, while being based on a solid theoretical foundation, leads to results of low quality when the aim is to find dense correspondences in datasets that contain noise and non-isometric distortion. Hence, in the following, we exclude

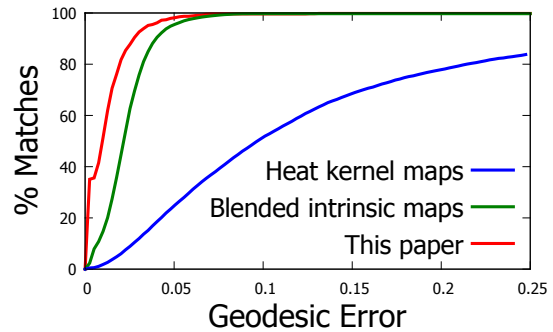


Figure 7: Cumulative error distributions for BIM, HKM, and our method based on the known ground truth. The models are shown in Figure 6.




Figure 8: Cumulative error distributions for BIM and our method based on the known ground truth for SCAPE dataset.

this method from our comparisons.

Second, we test our algorithm on the SCAPE dataset [52] consisting of 71 scans of a male scanned in different postures. In our experiment, we match the neutral posture to all 70 remaining postures using BIM and our approach. The cumulative error distributions for all 70 mappings are shown in Figure 8. This data differs from the samba frames in that it is globally near-isometric, but contains local areas of high-distortion (at joints for example). For this reason it provides a different kind of near-isometric test, and poses a greater challenge for our method, which does not exploit global assumptions. This is reflected in our error curve being below that of BIM. We set $\rho = 1.9$ in this experiment to reduce the influence of partial isometric patches that were thrown off by high local distortions.

5.5 Raw Scan Data

Finally, we consider raw scan data acquired using different acquisition systems. This type of data is noisy and incomplete, and is therefore significantly more challenging to match than the data used in the previous experiments. By using data from a variety of acquisition systems, we show our methods robustness to different types of acquisition noise. We also show our method's robustness to other acquisition artifacts that violate global isometry: large holes and contacts.




Figure 6: Comparison to HKM and BIM on models with ground truth. For error rates see Figure 7. Red indicates unmatched area; all further colors and the checker board have been painted on the source surface in order to visualize correspondences.

First, we use the same two frames of the samba sequence by Vlasic et al. [51] that were used above. However, this time, we consider the geometry reconstructed by a space-carving algorithm rather than template fitting. In this case, the frames we choose have different topology (as the hands merge with the body at the hips in one of the models) and are severely corrupted by noise. For this reason, we set $r_{min} = 1.5R$, $r_{max} = 3.4R$, and $\rho = 1$. We use these models to compare our method to BIM, Tevs et al.’s method, and Sharma et al.’s method. Figure 9 shows the results. Note that the results using BIM and the method by Tevs et al. match parts of the body of \mathcal{S} to the arms and legs of \mathcal{T} , thereby leading to visual artifacts. (See enlarged areas in Figure 9.) The method of Sharma et al., which is especially well suited to the scenario of handling contacts, leads to a result that covers the surfaces well. For this example, our method is run using two feature sets: first, using the standard features of our method and second, using the same image-based features used by Sharma et al. Our method produces a visually accurate mapping in both cases. However, when using standard features, the result of our method does not cover the right foot of the target surface, while the entire target surface is covered when using image-based features. Note that both the result by Sharma et al. and our results detect the contacts correctly and stop the growing in these regions. Hence, all areas of the surface are matched well. The improved performance with image-based features illustrates some of the technical challenges. The high level of surface noise makes matching geometric features difficult, which results in poorer sampling of the space of isometries, which in turn results in a poorer clustering result. We note however, that using the same features as Sharma et al., we obtain equal coverage and accuracy, and that when using purely geometric information, we outperform the other purely geometric methods tested.

Second, we compare our method to BIM and Tevs et al.’s method using two models of the BU-3DFE face database [53]. The two models contain numerous small holes and outliers. Furthermore, the models have different topology because the mouth is closed in one model and open in the other one. For these reasons, we set $r_{min} = 0.9R$, $r_{max} = 1.7R$, and $\rho = 1$. Figure 10 shows the results. Note that our method obtains a map-

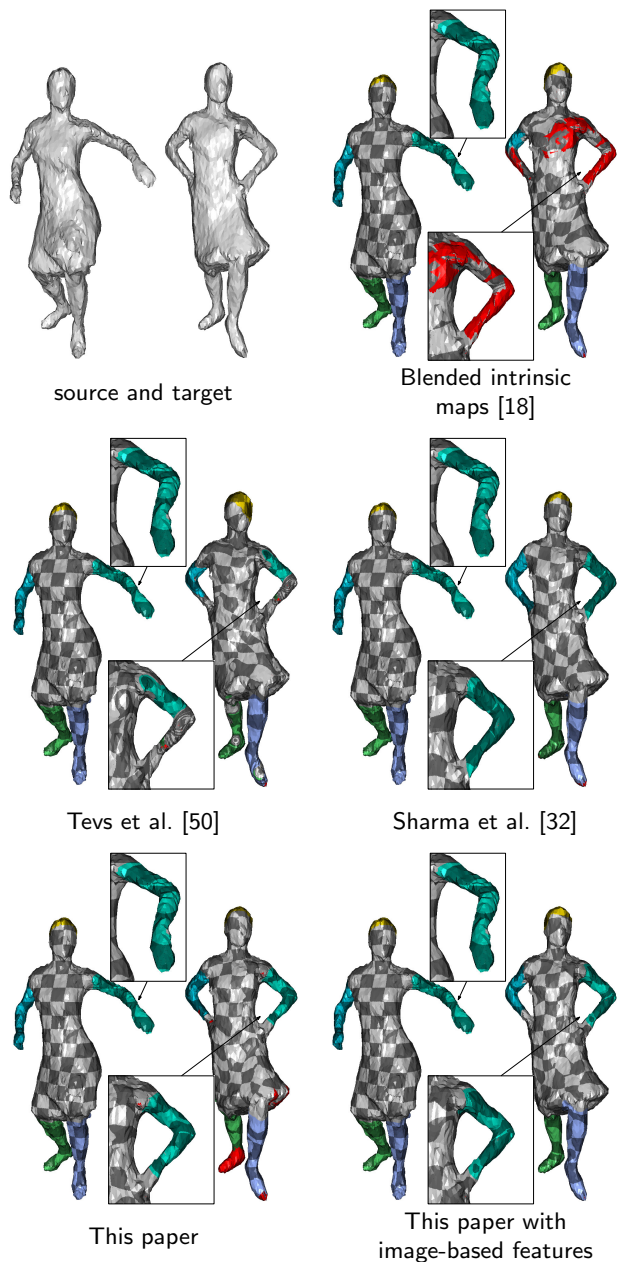


Figure 9: Comparison in presence of large contacts. Each pair shows \mathcal{S} on the left and \mathcal{T} on the right. Different data than in Figure 6. Red again indicates unmatched area

ping with higher visual accuracy than both BIM and the method of Tevs et al. In the case of BIM, this is to be expected, since the topological change breaks the global isometry. (Note the lips mapped to the side of the face.) The method of Tevs et al. produces a much more accurate result, but with still significant overall distortion and outliers (speckle-like effect). Our method leverages local information to fix partial isometries, and produces a mapping that largely preserves the semantic coloring with significantly lower distortion and without outliers.

Third, we compare our method to that of Tevs et al. on two point clouds acquired using a laser scanner. The two models contain numerous holes and outliers. As the models are point clouds and BIM requires input meshes, we do not compare our result to BIM for this experiment. For these models, we set $r_{min} = 1.5R$, $r_{max} = 3.4R$, and $\rho = 1$. The results are shown in Figure 11. As can be seen, the method of Tevs et al. obtains better coverage, however our method has fewer outliers—*islands of incorrectly mapped points within larger smoothly mapped regions* (see the enlarged parts of Figure 11). The suboptimal coverage of our method is due to the difficulty in matching features on surfaces plagued by missing data. A feature descriptor less sensitive to holes could improve this.

5.6 Limitations

In the previous sections, we have demonstrated that our method not only has theoretical advantages, but also computes results that improve upon the state of the art results in challenging cases where the input data is a pair of raw scans with topological noise.

We now discuss some limitations of our algorithm. Our algorithm is based on growing near-isometric mappings between partial regions of two surfaces and then clustering consistent mappings together. In this way, our algorithm enumerates a set of near-isometric mappings. For models that exhibit a large number of partial intrinsic symmetries, this technique enumerates a large set of near-isometric mappings where many of the mappings are inconsistent with each other. Since we stop growing new near-isometric regions after the first 200 oriented feature point matches have been considered, for shapes with a large number of partial near-isometric symmetries, it may happen that there is no cluster of consistent mappings covering a large area of \mathcal{S} .

An example where the clustering step fails to identify a cluster of consistent mappings covering a large area of \mathcal{S} is shown in Figure 12. The two frames of the flashkick dataset [54] contain a large topological change due to a merge of the subject’s pants in one of the models. Note that our algorithm correctly stops the growing of single partial mapping in this area, as shown in Figure 12. However, since the legs and core of the target body are intrinsically symmetric (similar to cylindrical), many inconsistent partial mappings are found by the algorithm, and they cannot be clustered in a consistent way.

We should also note that the computational costs are quite

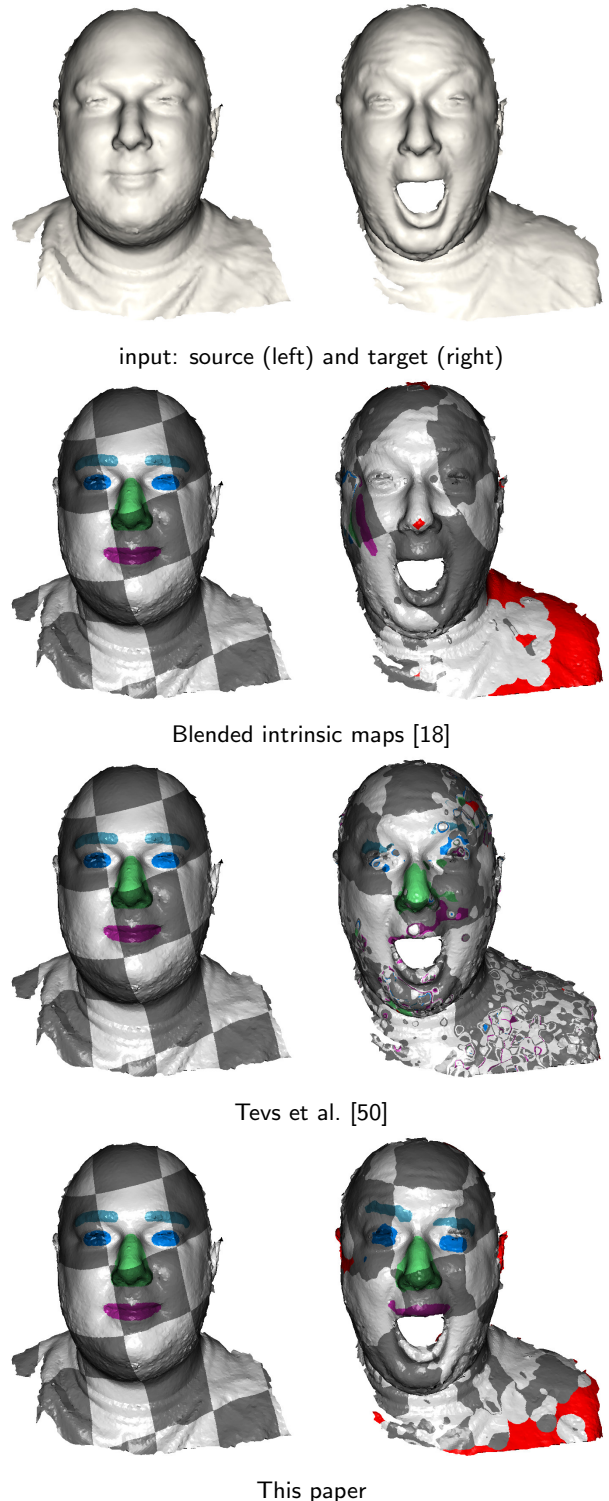


Figure 10: Comparison on data with significant topological changes and acquisition noise. Each pair shows \mathcal{S} on the left and \mathcal{T} on the right. Unmatched area marked in red.




Figure 11: Comparison on data with significant holes and acquisition noise. The first row shows the method of Tevs et al., while the second row shows our method. Within each row, the front and back are shown, in each case with S on the left and T on the right. Unmatched area shown in red.




Figure 12: A single partial mapping for an example with large topological changes. The red area is unmatched.

high. This is partially due to unoptimized code, but an algorithmic shortcoming is the rather simple feature-matching algorithm for finding start positions and directions. Optimizing this was not the focus of this work; our method should rather be understood as an alternative for the dense matching step where it can replace previous approaches based on conformal maps ([10, 18] and follow-ups), heat-kernel maps ([11] and follow-ups), or geodesic triangulation with landmarks ([15, 16] and follow-ups) in order to handle partial matching.

For future work, to address this limitation, we plan to combine our approach with an approach that detects continuous intrinsic symmetries [55], thereby reducing the search space for the initial feature matching and allowing us to efficiently enumerate all partial matches. Similarly, our framework could be extended to find partial intrinsic symmetries within a single object.

6 Conclusion

We have analyzed the complexity of the isometric matching problem under global and local isometry assumptions and based on this analysis we have introduced a new approach to solve the partial isometric matching problem using a representation for partial isometries that is both low-dimensional and redundant. Underpinning this is the fundamental observation that isometric mappings can be determined using purely local information and have only three degrees of freedom on 2-manifolds. The local metric propagation algorithm we derived from this observation is designed to handle topological noise that could affect large portions of the model, including both large holes and contacts. The redundancy in the representation can be exploited to increase robustness of and to combine partial matches. We have shown how a direct implementation of this theoretical framework can be used to match challenging surfaces with different types of topological noise.

The insights gained by studying the partial isometric matching problem have the potential to impact other shape processing tasks. For instance, the representation for partial isometric matches introduced in this paper can be used to derive new algorithms to detect partial symmetries of shapes. For future work, we plan to further investigate this option.

Acknowledgements

The authors thank Art Tevs, Aurela Shehu, Waqar Khan and Avinash Sharma for their help in conducting the comparative evaluation, Vladimir Kim for making his code available, and Art Tevs, Silke Jansen, Alexander Berner, Qi-Xing Huang, and Leonidas Guibas for discussions. We also thank the anonymous reviewers for their valuable comments that we used to improve the paper.

References

- [1] J. Davis, D. Nehab, R. Ramamoorthi, and S. Rusinkiewicz, "Spacetime stereo: A unifying framework for depth from triangulation," *IEEE PAMI*, vol. 27, no. 2, pp. 296–302, 2005.
- [2] S. König and S. Gumhold, "Image-based motion compensation for structured light scanning of dynamic scenes," *Int. J. of Int. Sys. Tech. App.*, vol. 5, no. 3/4, pp. 434–441, 2008.
- [3] T. Weise, B. Leibe, and L. V. Gool, "Fast 3d scanning with automatic motion compensation," in *Proc. CVPR*, 2007, pp. 1–8.
- [4] D. Vlastic, P. Peers, I. Baran, P. Debevec, J. Popović, S. Rusinkiewicz, and W. Matusik, "Dynamic shape capture using multi-view photometric stereo," *TOG*, vol. 28, no. 5, pp. 174:1–174:11, 2009.
- [5] P. Debevec, "Virtual cinematography: Relighting through computation," *IEEE Computer*, vol. 39, no. 8, pp. 57–65, August 2006.
- [6] V. Blanz and T. Vetter, "A morphable model for the synthesis of 3d faces," in *Proc. SIGGRAPH*, 1999, pp. 187–194.
- [7] B. Allen, B. Curless, and Z. Popović, "The space of human body shapes: Reconstruction and parameterization from range scans," *TOG*, vol. 22, no. 3, pp. 587–594, 2003.
- [8] N. Hasler, C. Stoll, M. Sunkel, B. Rosenhahn, and H.-P. Seidel, "A statistical model of human pose and body shape," *CGF (Proc. Eurographics)*, vol. 28, no. 2, 2009.
- [9] T. Weise, S. Bouaziz, H. Li, and M. Pauly, "Realtime performance-based facial animation," *TOG*, vol. 30, no. 4, pp. 77:1–77:10, 2011.
- [10] Y. Lipman and T. Funkhouser, "Möbius voting for surface correspondence," *TOG*, vol. 28, no. 3, pp. 72:1–72:12, 2009.
- [11] M. Ovsjanikov, Q. Mérigot, F. Méholi, and L. J. Guibas, "One point isometric matching with the heat kernel," *CGF*, vol. 29, no. 5, pp. 1555–1564, 2010.
- [12] A. Tevs, A. Berner, M. Wand, I. Ihrke, and H.-P. Seidel, "Intrinsic shape matching by planned landmark sampling," *CGF*, vol. 29, no. 2, pp. 543–552, 2011.
- [13] D. Anguelov, P. Srinivasan, H.-C. Pang, D. Koller, S. Thrun, and J. Davis, "The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces," in *Proc. NIPS*, 2005, pp. 33–40.
- [14] A. M. Bronstein, M. M. Bronstein, and R. Kimmel, "Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching," *PNAS*, vol. 103, no. 5, pp. 1168–1172, 2006.
- [15] Q.-X. Huang, B. Adams, M. Wicke, and L. J. Guibas, "Non-rigid registration under isometric deformations," *CGF*, vol. 27, no. 5, pp. 1449–1457, 2008.
- [16] A. Tevs, M. Bokeloh, M. Wand, A. Schilling, and H.-P. Seidel, "Isometric registration of ambiguous and partial data," in *Proc. CVPR*, 2009, pp. 1185–1192.
- [17] J. Sun, M. Ovsjanikov, and L. Guibas, "A concise and provably informative multi-scale signature based on heat diffusion," *CGF*, vol. 28, no. 5, pp. 1383–1392, 2009.
- [18] V. G. Kim, Y. Lipman, and T. Funkhouser, "Blended intrinsic maps," *TOG*, vol. 30, no. 4, pp. 79:1–79:12, 2011.
- [19] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas, "Functional maps: a flexible representation of maps between shapes," *TOG*, vol. 31, no. 4, pp. 30:1–30:11, 2012.
- [20] A. M. Bronstein, M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro, "A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching," *IJCV*, vol. 89, no. 2-3, pp. 266–286, 2010.
- [21] Y. Lipman, R. M. Rustamov, and T. A. Funkhouser, "Biharmonic distance," *TOG*, vol. 29, no. 3, pp. 27:1–27:11, 2010.
- [22] K. Xu, H. Zhang, W. Jiang, R. Dyer, Z. Cheng, L. Liu, and B. Chen, "Multi-scale partial intrinsic symmetry detection," *TOG*, vol. 31, no. 6, pp. 181:1–181:11, 2012.
- [23] M. Berger, *A Panoramic View of Riemannian Geometry*. Springer, 2002.
- [24] W. Zhang, X. Wang, D. Zhao, and X. Tang, "Graph degree linkage: Agglomerative clustering on a directed graph," in *Proc. ECCV*, 2012, pp. 428–441.
- [25] O. van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or, "A survey on shape correspondence," *CGF*, vol. 30, no. 6, pp. 1681–1707, 2011.
- [26] G. Tam, Z.-Q. Cheng, Y.-K. Lai, F. Langbein, Y. Liu, D. Marshall, R. Martin, X.-F. Sun, and P. Rosin, "Registration of 3d point clouds and meshes: A survey from rigid to non-rigid," *TVCG*, To appear.
- [27] A. Elad and R. Kimmel, "On bending invariant signatures for surfaces," *TPAMI*, vol. 25, no. 10, pp. 1285–1295, 2003.
- [28] V. Jain, H. Zhang, and O. van Kaick, "Non-rigid spectral correspondence of triangle meshes," *IJSM*, vol. 13, no. 1, pp. 101–124, 2007.
- [29] S. Wuhler, Z. B. Azouz, C. Shu, and P. Bose, "Posture invariant correspondence of incomplete triangular manifolds," *IJSM*, vol. 13, no. 2, pp. 139–157, 2007.
- [30] O. van Kaick, H. Zhang, and G. Hamarneh, "Bilateral maps for partial matching," *CGF*, vol. to appear, 2013.

- [31] M. Aubry, U. Schlickewei, and D. Cremers, “The wave kernel signature: A quantum mechanical approach to shape analysis,” in *Proc. ICCV Workshops*, 2011, pp. 1626–1633.
- [32] A. Sharma, R. P. Horaud, J. Cech, and E. Boyer, “Topologically-robust 3d shape matching based on diffusion geometry and seed growing,” in *Proc. CVPR*, 2011, pp. 2481–2488.
- [33] Y. Sahillioglu and Y. Yemez, “Scale normalization for isometric shape matching,” *Computer Graphics Forum (Proc. Pacific Graphics)*, vol. 31, no. 7, 2012.
- [34] A. Bronstein, M. Bronstein, A. Bruckstein, and R. Kimmel, “Partial similarity of objects, or how to compare a centaur to a horse,” *IJCV*, vol. 84, no. 2, pp. 163–183, 2009.
- [35] D. Raviv, A. Bronstein, M. Bronstein, and R. Kimmel, “Full and partial symmetries of non-rigid shapes,” *IJCV*, vol. 89, no. 1, pp. 18–39, 2010.
- [36] D. Hähnel, S. Thrun, and W. Burgard, “An extension of the ICP algorithm for modeling nonrigid objects with mobile robots,” in *Proc. IJCAI*, 2003, pp. 915–920.
- [37] H. Zhang, A. Sheffer, D. Cohen-Or, Q. Zhou, O. van Kaick, and A. Tagliasacchi, “Deformation-driven shape correspondence,” *CGF*, vol. 27, no. 5, pp. 1431–1439, 2008.
- [38] H. Li, R. W. Sumner, and M. Pauly, “Global correspondence optimization for non-rigid registration of depth scans,” *CGF*, vol. 27, no. 5, pp. 1421–1430, 2008.
- [39] M. Wand, B. Adams, M. Ovsjanikov, A. Berner, M. Bokeloh, P. Jenke, L. Guibas, H.-P. Seidel, and A. Schilling, “Efficient reconstruction of nonrigid shape and motion from real-time 3d scanner data,” *TOG*, vol. 28, no. 2, pp. 1–15, 2009.
- [40] H. Li, B. Adams, L. J. Guibas, and M. Pauly, “Robust single-view geometry and motion reconstruction,” *TOG*, vol. 28, no. 5, pp. 175:1–175:10, 2009.
- [41] A. Pentland and J. Williams, “Good vibrations: modal dynamics for graphics and animation,” in *Proc. SIGGRAPH*, 1989, pp. 215–222.
- [42] R. Schmidt, C. Grimm, and B. Wyvill, “Interactive decal compositing with discrete exponential maps,” *TOG*, vol. 25, no. 3, pp. 605–613, 2006.
- [43] R. Schmidt, “Stroke parametrization,” *CGF*, vol. 32, no. 3, 2013.
- [44] E. L. Malvaer and M. Reimers, “Geodesic polar coordinates on polygonal meshes,” *CGF*, vol. 31, no. 8, pp. 2423–2435, 2012.
- [45] R. Palais, “On the differentiability of isometries,” *Proc. of the AMS*, vol. 8, no. 4, pp. 805–807, 1957.
- [46] A. Öztireli, G. Guennebaud, and M. Gross, “Feature preserving point set surfaces based on non-linear kernel regression,” *CGF*, vol. 28, no. 2, 2009.
- [47] F. Mévoli and G. Sapiro, “A theoretical and computational framework for isometry invariant recognition of point cloud data,” *FoCM*, vol. 5, no. 3, pp. 313–347, 2005.
- [48] Y. Sun and M. Abidi, “Surface matching by 3d point’s fingerprint,” in *Proc. ICCV*, 2001, pp. 263–269.
- [49] R. Rustamov, “Barycentric coordinates on surfaces,” *CGF*, vol. 29, no. 5, 2010.
- [50] A. Tevs, A. Berner, M. Wand, I. Ihrke, M. Bokeloh, J. Kerber, and H.-P. Seidel, “Animation cartography - intrinsic reconstruction of shape and motion,” *TOG*, vol. 31, no. 2, p. 15, April 2012.
- [51] D. Vlastic, I. Baran, W. Matusik, and J. Popović, “Articulated mesh animation from multi-view silhouettes,” *TOG*, vol. 27, no. 3, pp. 97:1–97:9, 2008.
- [52] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis, “Scape: Shape completion and animation of people,” *TOG*, vol. 24, no. 3, pp. 408–416, 2005.
- [53] L. Yin, X. Wei, Y. Sun, J. Wang, and M. J. Rosato, “A 3d facial expression database for facial behavior research,” in *Proc. FG*, 2006, pp. 211–216.
- [54] J. Starck and A. Hilton, “Surface capture for performance based animation,” *CG&A*, vol. 27, no. 3, pp. 21–31, 2007.
- [55] M. Ben-Chen, A. Butscher, J. Solomon, and L. Guibas, “On discrete killing vector fields and patterns on surfaces,” *CGF*, vol. 29, no. 5, 2010.