

# Scale Normalization for Isometric Shape Matching

Y. Sahillioğlu and Y. Yemez

Department of Computer Engineering, Koç University, Istanbul, Turkey

---

## Abstract

*We address the scale problem inherent to isometric shape correspondence in a combinatorial matching framework. We consider a particular setting of the general correspondence problem where one of the two shapes to be matched is an isometric (or nearly isometric) part of the other up to an arbitrary scale. We resolve the scale ambiguity by finding a coarse matching between shape extremities based on a novel scale-invariant isometric distortion measure. The proposed algorithm also supports (partial) dense matching, that alleviates the symmetric flip problem due to initial coarse sampling. We test the performance of our matching algorithm on several shape datasets in comparison to state of the art. Our method proves useful, not only for partial matching, but also for complete matching of semantically similar hybrid shape pairs whose maximum geodesic distances may not be compatible, a case that would fail most of the conventional isometric shape matchers.*

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Computer Graphics—Computational Geometry and Object Modeling[Matching]

---

## 1. Introduction

Shape correspondence is a fundamental problem in both computer vision and graphics with various applications such as deformation transfer, statistical shape analysis, shape retrieval and spatio-temporal shape modeling [BBK08]. 3D shape correspondence methods seek on two given shapes for pairs of surface points that are similar or semantically equivalent. Isometry is an important clue in achieving this goal since similar shape parts usually have similar metric structures. Although shape matching can be achieved by enforcing geodesic consistencies or by searching for mappings with minimum isometric distortion, the arbitrary scale of shapes usually poses an important challenge to overcome for partial or complete shape correspondence, as we address in this paper.

There are different ways of dealing with the scale problem in the partial and/or complete shape correspondence literature. Some methods simply assume that the shapes come in compatible scales, which is rather a very strong assumption, but that might sometimes prove useful especially in the case of matching 3D scan data [BBK06, GMGP05, HAWG08, TBW\*09]. Other methods are feature-based and rely solely on local shape descriptors which are usually designed to be scale-invariant [FS06, ZBVH09, BK10]. Local shape similarity is an important clue for shape correspondence, especially

in the case of non-isometric deformations, but otherwise it is considered as less reliable than global shape information such as isometry. They may not perform well for example when the shapes to be matched exhibit large variations in their local geometry, or when there are many points that are locally similar.


If two shapes are isometric (perfectly, nearly, or partially) and come in different scales, in order to be able to incorporate metric similarities, the shapes have to be normalized into the same scale prior to the matching process. There exist two different approaches to achieve this. The first and simpler approach is to scale the original geometry with respect to some global intrinsic property such as maximum geodesic distance [SY10, ZSCO\*08], maximum centrality [ACOT\*10], or total surface area [OMMG10]. This strategy may work satisfactorily in the case of complete shape correspondence but otherwise for partial matching its success depends highly on the global similarity between the shapes. Even in the case of complete correspondence, if the given shapes are only nearly isometric, it is indeed possible that global intrinsic properties significantly deviate from one shape to the other, yielding inaccurate normalization. The second approach is to transform input shapes into a different domain where the scaling problem is implicitly handled. Euclidean embedding is one such transform which is commonly used for isometric shape correspondence, but it is rather a global technique and does

not well apply to partial correspondence [JZ06,MHK\*08]. A better alternative for partial matching is based on the Möbius transformation that can be used for conformal embedding of the given shapes into a canonical coordinate frame on the complex plane [LF09,ZWW\*10], where deviations from isometry can approximately be computed based on mutually closest points. The Möbius Voting method in [LF09] considers the most general setting of the partial correspondence problem where shapes to be matched can also have uncommon (non-similar) parts. This voting-based approach can generate reliable but sparse correspondences, suffering mainly from spurious votes of the extra uncommon parts in the shape models, and hence performs poorly in terms of dense and shape extremity matching. We also note that the Möbius Voting approach is restricted to sphere topology and prone to errors due to embedding approximation.

An alternative to geodesic metric is the diffusion metric which is less accurate for measuring isometric distortion but generally considered as more robust to topological noise. Local scale differences are however difficult to handle using diffusion-based metrics. The commute-time metric for example addresses the scale problem only globally [WBBP11], hence cannot be used for the partial matching problem, i.e., in a setting where one of the shapes to be matched includes a scaled part of the other. In [PBB11], a scale-invariant version of the heat kernel signature is used to address the part matching problem, which however requires setting a time scale parameter that itself depends on the shape scale.

In the most general setting of the correspondence problem, the ambiguity of shape scale is a local issue that cannot be resolved relying on some global shape intrinsics. This is demonstrated in Fig. 1. While two completely isometric shapes can be brought to the same scale trivially by normalizing their maximum geodesic distances (Fig. 1a), scale normalization for partially overlapping shapes (and even for hybrid shapes which are nearly isometric) is not that straightforward (Fig. 1b). In the latter case, the scale ambiguity can be resolved based on local geodesic distances computed on a sparse set of trusted correspondences (Fig. 1c), which is the approach that we follow in this paper. Alternatively, a small set of trusted correspondences can be used to define an Euclidean embedding, such as Möbius transformation in the extended complex plane [LF09], that implicitly handles the scale problem (Fig. 1d).

In this paper, we consider the shape correspondence problem in the particular setting where one of the shapes to be matched is an isometric part of the other up to an arbitrary scale (note that this setting also includes the problem of complete shape matching). We address the scale problem in a combinatorial framework that minimizes a scale-invariant isometric distortion function in the 3D Euclidean space. We first sample shape extremities from the mesh representations of the given shape pair and then find a coarse map in between via combinatorial search, that also extends to dense matching. The output of the proposed method is hence a sparse or



**Figure 1:** (a) Scale normalization of perfectly isometric shapes based on maximum geodesic distances (yellow paths), which does not apply to partially isometric shapes, e.g., normalized distances (red paths) are not the same on both shapes (b). Scale normalization for partial matching can be achieved based on trusted correspondences, either by normalizing (green) geodesic paths (c) or by defining an Euclidean embedding (d).

dense optimal set of correspondences between the surfaces of the given partially or completely isometric shapes.

We note that the source code and the executables for the method that we present in this paper are publicly available in <http://home.ku.edu.tr/~yyemez/scalennormalization>.

## 2. Isometric distortion

We describe two different isometric distortion measures in the sequel,  $D_{\text{iso}}^{(1)}$  and  $D_{\text{iso}}^{(2)}$ , that can be used interchangeably in our combinatorial shape matching framework. The first measure can be considered as the one that we propose as a novel measure whereas the latter rather serves as a baseline measure to compare with the former. These measures will be compared experimentally later in Section 7.

### 2.1. Scale-invariant isometric distortion

Let  $S$  and  $T$  represent two sets of points sampled from the given source and target shapes, respectively. Suppose also that a mapping  $\xi : S \rightarrow T$  (a relation in the most general setting) is given. Then the isometric distortion of this mapping can be measured by the scale-invariant function  $D_{\text{iso}}^{(1)}(\xi)$  as follows:

$$D_{\text{iso}}^{(1)}(\xi) = \frac{1}{|\xi|} \sum_{(s_i, t_j)} d_{\text{iso}}^{(1)}(s_i, t_j), \quad (1)$$

where  $d_{\text{iso}}^{(1)}(s_i, t_j)$  is the contribution of the individual correspondence  $(s_i, t_j)$  to the overall isometric distortion:


$$d_{\text{iso}}^{(1)}(s_i, t_j) = \frac{1}{\binom{|\xi'|}{2}} \sum_{((s_a, t_b), (s_c, t_d)) \in \mathcal{C}(\xi')} |\rho(s_i, t_j; s_a, t_b) - \rho(s_i, t_j; s_c, t_d)| \quad (2)$$

with  $\xi' = \xi - \{(s_i, t_j)\}$  and  $\mathcal{C}(\xi')$  denoting the set of all pairwise combinations from  $\xi'$ . The ratio function  $\rho(s_i, t_j; s_k, t_l)$

is then written in terms of geodesic distances, for a given  $(s_k, t_l) \in \xi$ :

$$\rho(s_i, t_j; s_k, t_l) = \max \left( \frac{g(s_i, s_k)}{g(t_j, t_l)}, \frac{g(t_j, t_l)}{g(s_i, s_k)} \right) \quad (3)$$

where  $g(\cdot, \cdot)$  is the geodesic distance between two surface points. This definition of isometric distortion is based on the observation that the ratios between geodesic distances on a surface remain unchanged under scaling and isometric deformations. Hence if  $T$  and  $S$  are sampled consistently from the given arbitrarily scaled (partially) isometric shapes, then one can find an optimal mapping  $\xi^*$  in between such that  $D_{\text{iso}}^{(1)}(\xi^*) = 0$ . This is illustrated in Fig. 2.



**Figure 2:** Demonstration of the scale-invariant distortion measure defined in (1). The ratios between geodesic distances remain invariant under scaling and isometric deformation:  $\rho(s_i, t_j; s_a, t_b) = \rho(s_i, t_j; s_c, t_d)$ .

### 2.2. Isometric distortion with normalized geodesics

The isometric distortion  $D_{\text{iso}}^{(2)}(\xi)$  for a given map  $\xi : S \rightarrow T$  between point sets  $S$  and  $T$  is computed as follows:

$$D_{\text{iso}}^{(2)}(\xi) = \frac{1}{|\xi|} \sum_{(s_i, t_j) \in \xi} d_{\text{iso}}^{(2)}(s_i, t_j) \quad (4)$$

where  $d_{\text{iso}}^{(2)}(s_i, t_j)$  is the contribution of the individual correspondence  $(s_i, t_j)$  to the overall isometric distortion:

$$d_{\text{iso}}^{(2)}(s_i, t_j) = \frac{1}{|\xi'|} \sum_{(s_l, t_m) \in \xi'} |g_n(s_i, s_l) - g_n(t_j, t_m)| \quad (5)$$

where  $\xi'$  is the list of correspondences to be traversed, which is set to be as  $\xi' = \xi - \{(s_i, t_j)\}$  unless stated otherwise. The function  $g_n(\cdot, \cdot)$  is the geodesic distance between two samples, normalized by the local maximum geodesic distance, that is, the distance between the two farthest points in  $S$  (or  $T$ ). Note that this definition of isometric distortion requires a consistent joint-sampling between  $S$  and  $T$  to enable reliable normalization of pairwise geodesics (see also Fig. 3).

Unlike  $D_{\text{iso}}^{(1)}(\xi)$  that averages over all  $\binom{|\xi|-1}{2}$  pairs of available correspondences in  $\xi$  for each individual isometric distortion computation of a constituent match,  $D_{\text{iso}}^{(2)}(\xi)$  relies only on two farthest points to obtain locally normalized geodesics. Although this suggests a saving in the order of  $|\xi|$  for the time complexity of isometric distortion computation, the use of the latter measure may easily induce a




**Figure 3:** Demonstration of the distortion measure with normalized geodesics defined in (4). Suppose that sample sets  $S$  and  $T$  each contains three surface points to be matched. The geodesics used in computation of  $d_{\text{iso}}^{(2)}(s_i, t_j)$  (red paths on the right) are normalized w.r.t. the local maximum geodesics (yellow paths on the left).

failure: If the two farthest point pair in one sample set does not align well with the pair in the other, the normalization of geodesics becomes inconsistent, and the distortion measure cannot be computed reliably. Although this is less of a problem in the case of complete shape matching (when compared to partial matching), complete matching of semantically similar shapes may still be problematic since farthest point pairs defining maximum geodesics are not necessarily consistent for nearly isometric shapes, such as for Cat vs. Wolf with long and short tails, and for Gorilla vs. Human with long and short arms, being two specific examples from our experiments.

### 3. Feature point selection

We use shape extremities as feature points since they are the most salient points of a surface, which are easy to identify while searching for a sparse set of correspondences. To find shape extremities on a given shape, we first compute the integral geodesic distance function [HSKK01] for every vertex of the mesh representation and mark the one yielding the maximum function value which is expected to be at the most prominent tip of the shape. The farthest point sampling procedure of [ELPZ97] started from this marked vertex then selects a number of samples on the surface. We set the number of samples large enough to provide sufficient coverage on the surface, and small enough for computational efficiency of the combinatorial framework that will be described in the next section. The feature vertices resulting from this sparse sampling process constitute the sets  $S$  (source) and  $T$  (target) to be matched (see Fig. 4).




**Figure 4:** Feature vertices on two models (10 samples).

#### 4. Combinatorial matching

Once we have the feature sets,  $S$  and  $T$ , extracted from the source and target shapes, we search for an optimal partial mapping  $\xi^*$  from  $S$  to  $T$  with minimum distortion. In our setting, one of the shapes, say  $S$ , is an isometric part of  $T$  up to a scale factor. However, since two shapes are never perfectly isometric, even partly, due to imperfections of the modeling process and geometry discretization errors, it is not usually possible to find a zero distortion mapping, hence the goal rather becomes minimization of the isometric distortion function, that is either  $D_{\text{iso}}^{(1)}(\xi)$  defined in (1) or  $D_{\text{iso}}^{(2)}(\xi)$  in (4). We also note that when matching a shape part with a complete model, the geodesics may slightly differ between similar surface points from one shape to the other due to the cut regions.

We minimize the isometric distortion function via combinatorial search over all possible mappings  $\xi$  between  $S$  and  $T$ . In order to make this combinatorial search problem tractable and to reduce the search space, we consider  $M = 5$  evenly-spaced vertices from  $S$ . Since, in our setting, the shape  $S$  is assumed to be a part of  $T$ , these extreme vertices are expected to match (at least roughly) with a subset of  $T$ . Hence we need to compute the isometric distortion for  $M! \binom{|T|}{M}$  different possible (one-to-one) mappings, and the mapping that yields the minimum distortion is selected as the optimal mapping  $\xi^*$  (see Fig. 5).



**Figure 5:** Overview of the combinatorial matching process.  $M$  extreme vertices of the source shape are matched with  $|T|$  extremities from the target. Hence  $M!$  possible permutations of  $\binom{|T|}{M}$  different combinations are tested with the source, and the one with minimum distortion gives the optimal mapping  $\xi^*$  (right).

#### 5. Extension to dense correspondence

We now extend the sparse correspondence  $\xi^* : S \rightarrow T$  between matched shape extremities to a denser map. This support for dense correspondence also allows us to alleviate the symmetric flip problem that isometric coarse correspondence methods usually suffer from (e.g., left arm/leg is matched to right arm/leg between two human shapes due to sparse sampling).

The first step towards dense matching is to bring the shapes to the same scale using the trusted  $\xi^*$  that provides

us with the factor  $\kappa$  to scale the target mesh:

$$\kappa = \frac{1}{\binom{|\xi^*|}{2}} \sum_{((s_a, t_b), (s_c, t_d)) \in \mathcal{C}(\xi^*)} \frac{g(s_a, s_c)}{g(t_b, t_d)} \quad (6)$$

With both shapes at the same scale, in the second step, we take evenly-spaced dense samples on each with consistent spacing. The dense sampling algorithm that we use is very similar to the one described in [SY10]. We set the sampling radius to  $r = 0.17 \sqrt{A/\pi}$ , where  $A$  is the surface area of the target shape, that ensures sampling of about 100 dense samples. An arbitrary vertex is selected as the initial sample, and all the vertices lying within its patch of radius  $r$  are marked. The next sample is then selected arbitrarily from the set of unmarked vertices. When this is repeated until no unmarked vertex is left, we obtain a partitioning of the shape surface into samples that are at least  $r$  apart from each other. We denote the dense sample sets on source and target shapes by  $\hat{S}$  and  $\hat{T}$ , respectively, to which we also append  $S$  and  $T$  themselves since they cover the salient extreme points. Note that the number of samples on the source shape can be significantly less than the number of samples on the target since the former is assumed to be an isometric part of the latter in our setting.

The third and last step begins with filling in a cost matrix  $C$  where each entry  $c_{ij}$  represents the cost of matching a sample  $\hat{s}_i$  in  $\hat{S}$  to a sample  $\hat{t}_j$  in  $\hat{T}$ . We build  $C$  by setting  $c_{i,j} = d_{\text{iso}}^{(2)}(\hat{s}_i, \hat{t}_j)$  via (5), based on the trusted correspondences in  $\xi^*$ , hence  $\xi' = \xi^*$ . We then perform minimum-weight perfect matching [Kol09] on  $C$ , that reveals the desired dense map  $\hat{\xi}^*$ . Since the cardinalities of the disjoint sets must match for a perfect matching, if  $|\hat{S}| \neq |\hat{T}|$ , we introduce virtual vertices with connector edges of  $\infty$  weights.

To address the symmetric flip problem, we iterate the third step above  $K$  times, each time setting  $\xi'$  to one of the best  $K$  sparse correspondences. Note that our combinatorial matching framework allows us to sort out the  $K$  least distorted sparse maps,  $\xi_1, \xi_2, \dots, \xi_K$ , where  $\xi_1 = \xi^*$ . Hence we compute  $K$  dense maps  $\hat{\xi}_1, \hat{\xi}_2, \dots, \hat{\xi}_K$ . We then promote the sparse map among the best  $K$ , which generates the dense correspondence with minimum distortion, as our final sparse correspondence output:

$$\hat{\xi}^* \leftarrow \arg \min_{\xi_1, \xi_2, \dots, \xi_K} \hat{D}_{\text{iso}}(\hat{\xi}) \quad (7)$$

where  $\hat{D}_{\text{iso}}$  is similar to  $D_{\text{iso}}^{(2)}$  except that it uses geodesics normalized by the same overall maximum geodesic over the shapes, which is the maximum geodesic of the target. In this way, the coarse correspondences with similar distortion values, which are possibly symmetrically flipped, are compared at a finer resolution where their isometric distortions can be more accurately computed. We use  $K = 5$  in all our experiments, which is sufficient to differentiate the true mapping most of the time.



## 6. Computational Complexity

We analyze the computational complexity of  $D_{\text{iso}}^{(1)}(\xi)$  and  $D_{\text{iso}}^{(2)}(\xi)$  computations, and the combinatorial part matching algorithm separately. Given a mapping  $\xi$  of size  $N$ , it takes  $N \cdot \binom{N-1}{2}$  iterations of constant work to compute each  $d_{\text{iso}}^{(1)}(s_i, t_j)$ , and hence  $O(N^3)$  time for  $D_{\text{iso}}^{(1)}(\xi)$  computation. Similarly,  $D_{\text{iso}}^{(2)}(\xi)$  is achieved in  $O(N^2)$  time as  $N \cdot (N-1)$  iterations are required.

The partial matching algorithm, on the other hand, when fed with two original meshes having  $V_S$  and  $V_T$  vertices, starts with feature point selection of  $O(V \log V)$  time, where  $V = \max(V_S, V_T)$ . The combinatorial search algorithm then demands  $\binom{|T|}{M} M!$  evaluations of the function  $D_{\text{iso}}^{(1)}(\xi)$  or  $D_{\text{iso}}^{(2)}(\xi)$ , where  $M$  is the size of the sought mapping, i.e., the size of the reduced source feature set. Hence the overall algorithmic complexity is  $O(\binom{|T|}{M} M! M^3)$  or  $O(\binom{|T|}{M} M! M^2)$ , which gives reasonable computation time in practice thanks to the feature sampling that leads to small  $|T|$  values as well as the choice  $M = 5$ .

Establishing a dense map of size  $U$  adds another  $O(V \log V)$  for dense sampling plus  $O(U^2 \log U)$  time for minimum-weight perfect matching, both of which are dominated by the preceding combinatorial part matching algorithm (we set  $U = 100$  in our experiments).

## 7. Experimental results

We have tested the performance of our shape correspondence algorithm on several shape datasets. The first three datasets are uniformly sampled fixed-connectivity mesh sequences, each representing the motion of an articulated object, which we refer to as *Horse*, *Dog*, *Wolf*, *Cat* (from TOSCA shape benchmark [BBK08]), *Jumping Man* [SMP03], and *Dancing Man* [dAST\*08]. The fourth dataset is a reconstructed pose sequence of a human actor from the SCAPE benchmark [ASK\*05], which contains 71 different non-uniformly sampled fixed-connectivity models. We have created partial models from some of these datasets by manually cropping and arbitrarily rescaling the original complete models. Representative partial models can be observed throughout the figures.

The last dataset is from the SHREC'11 shape benchmark [BBB\*11]. The part that we use from this dataset includes a null (reference) shape (*Human*) plus three classes named as isometry, partial, and scaling, each containing 5 uniform high-resolution human models with arbitrary connectivity. The isometry class contains isometric deformations of the null shape and the scaling class includes models from the isometry class in different scales, whereas the partial class is composed by parts of the complete models from the isometry class. We have also used the female class from TOSCA, which contains 10 models, to match with the partial class of SHREC'11 to expand our experiments on nearly isometric

shape pairs. Besides, we match *Gorilla* from TOSCA with *Human*, which verifies, along with *Cat* vs. *Wolf*, the accuracy of our algorithm on semantically similar hybrid shape pairs whose maximum geodesic distances do not coincide, a case that would fail most of the conventional isometric shape matchers.

We evaluate the quality of the obtained sparse mappings by using the distortion measure  $D_{\text{iso}}^{(1)}$  or  $D_{\text{iso}}^{(2)}$ , depending on which measure is employed to find the given mapping. We use the distortion measure  $\hat{D}_{\text{iso}}$  to evaluate the dense mappings. We also define the following ground-truth distortion  $D_{\text{grd}}$  to measure deviation from ground-truth correspondences:

$$D_{\text{grd}}(\xi) = \frac{1}{|\xi|} \sum_{(s_i, t_j) \in \xi} g(\vartheta(s_i, t_j)), \quad (8)$$

where  $\vartheta(s_i)$  stands for the ground-truth correspondence of  $s_i$  on target as either known a priori or computed automatically by aligning the cropped model with its complete version. The maximum geodesic on the target model is normalized to 1 in order to simplify the interpretation of this measure.


We have compared our algorithm with the Möbius Voting (MV) method of [LF09]. For comparison tests, we have run the publicly available code of MV with its default settings of 100 samples and 1M votes, and evaluated its performance based on the  $M = 5$  samples that are closest to the shape extremity samples used for our algorithm. We also compare the full dense correspondences obtained by MV to our dense correspondences. In the visualization of our results, we highlight the generating/trusted coarse correspondences with larger spheres and bold lines, and the generated dense maps with smaller spheres. Yellow lines represent the worst matches w.r.t. ground-truth distortion (or isometric distortion in case the former is not available).

Figures 6-9 display various examples from our partial correspondence results obtained by using the isometric distortion measure  $D_{\text{iso}}^{(1)}$ , whereas in Fig. 10, we visually compare the performance of our algorithm with the MV method. Since Möbius Voting algorithm is restricted to sphere topology, while our method is not, for the sake of comparison, we have patched up all the partial models at their cut regions except for standard SHREC'11 shapes. Whether cut regions are patched up or not, the geodesic information on a partial surface slightly changes with respect to its complete version, introducing some extra imperfection to the correspondence problem; yet we observe that the correspondences that we obtain using our method are very satisfactory, and generally more accurate especially at shape extremities when compared to MV (see Fig. 10). The MV algorithm can actually generate accurate correspondences only for a small number of correspondence pairs, those with high confidence values. However the vertices corresponding to these pairs are not generally well distributed on the surface, being arbitrarily located, as can also be observed from Fig. 10. We also ob-

serve that the matching results of MV at shape extremities may be very inaccurate.


We provide quantitative performance evaluation of our method for dense and extremity-based sparse matching in comparison to MV in Table 1, where we also compare the two isometric distortion functions that we have employed for scale normalization. We note that in addition to partial matching (first 4 rows), both methods naturally support complete matching (last 4 rows). Note also that the computation of dense distortion  $\hat{D}_{iso}$ , as used in Eq. 7, is based on trusted coarse correspondences obtained either by using the distortion measure  $D_{iso}^{(1)}$  (column named as "with  $D_{iso}^{(1)}$ ") or  $D_{iso}^{(2)}$  (column named as "with  $D_{iso}^{(2)}$ ").

For part matching experiments (Figures 6-8), we have respectively 6, 7, 7, and 5 partial models for *Jumping Man*, *Dancing Man*, *Horse*, and *SHREC11* sequences. In the test suite for each dataset, each of the  $C$  complete models is matched with  $P$  random partial models from the same class, where the pair  $(C, P)$  is (18, 2), (9, 3), (7, 3), and (11, 5), respectively. For *SHREC11* evaluation, in addition to matching the 11 models from the null shape, isometry, and scaling classes to all 5 models from the partial class, we also compute the mapping from each of the 10 female models from TOSCA to a random partial class model in *SHREC11*. In all cases, we exclude from evaluation, the shape pairs that result in symmetric flips when matched using our method and/or Möbius Voting. As a result, the performance measures given in Table 1 are computed over 19, 22, 15, and 39 pairs, respectively on four datasets. Note that, a denser map that captures more intrinsic geometry of the shapes tends to overcome the symmetric flip problems observed at a coarser resolution, which in turn corrects the symmetric flip between the initial extreme samples as well (Fig. 7-bottom).




**Figure 6:** Partial correspondences between two different *Jumping Man* pairs at different scales (left half), and similarly for *Dancing Man* pairs (right half).  $S$  and  $T$  denote the partial and complete models, respectively, which applies to all subsequent figures that involve partial models.

As for complete matching experiments (Fig. 9), each of the 9 and 10 complete models is matched with a random complete model from the same class for *Dog* and *SCAPE* sequences, respectively. We also perform complete matching across classes to address the incompatible maximum geodesic issue with 5 *Cat-Wolf* and 4 *Gorilla-Human* pairs.




**Figure 7:** (Left) Partial correspondences obtained on two different *Horse* pairs. (Framed) The sparse partial mapping obtained on a *Horse* pair initially has a symmetric flip problem (left), which is resolved at a denser resolution (right).



**Figure 8:** Partial correspondence on *SHREC11*: Mappings to a partial class model  $S$  from a null shape, a scaling class shape, and two female class shapes in *TOSCA* (left to right).


In Table 1, we observe that the distortion measure  $D_{iso}^{(1)}$  outperforms  $D_{iso}^{(2)}$  especially for the part matching problem (first 4 rows). The performance difference here is mostly due to the inconsistency of trusted farthest point pairs used for geodesic normalization with the latter measure, which is not likely to occur in complete matching where  $D_{iso}^{(2)}$  performs as good as  $D_{iso}^{(1)}$  for shapes with perfectly compatible maximum geodesic distances (rows 5 – 6) but falls behind for hybrid pairs (rows 7 – 8). We also observe that our method outperforms the MV method in sparse extremity matching as well as dense matching. Note that MV cannot be evaluated



**Figure 9:** Complete correspondences between a *Dog* pair (top left) and two *SCAPE* pairs (second column) as well as between a *Cat-Wolf* pair (bottom left) and two *Gorilla-Human* pairs (third column).

| Pair                            | Our extremity matching                 |  | Our dense matching          |                             | MV extremity matching                  | MV dense matching                      |
|---------------------------------|--|--|-----------------------------|-----------------------------|--|--|
|                                 | $D_{\text{iso}}^{(2)}, D_{\text{grd}}$ | $D_{\text{iso}}^{(1)}, D_{\text{grd}}$ | with $D_{\text{iso}}^{(2)}$ | with $D_{\text{iso}}^{(1)}$ | $D_{\text{iso}}^{(1)}, D_{\text{grd}}$ | $\hat{D}_{\text{iso}}, D_{\text{grd}}$ |
| <i>Jumping Man (partial)</i>    | .030, .083                             | <b>.219, .072</b>                      | .019, .139                  | <u>.017, .089</u>           | <b>2.465, .253</b>                     | <u><b>.051, .140</b></u>               |
| <i>Dancing Man (partial)</i>    | .025, .068                             | <b>.159, .037</b>                      | .011, .082                  | <u>.011, .055</u>           | <b>1.965, .428</b>                     | <u><b>.068, .241</b></u>               |
| <i>Horse (partial)</i>          | .027, .056                             | <b>.114, .028</b>                      | .012, .055                  | <u>.011, .048</u>           | <b>2.463, .244</b>                     | <u><b>.051, .121</b></u>               |
| <i>SHREC11 (partial)</i>        | .044, n/a                              | <b>.207, n/a</b>                       | .038, n/a                   | <u>.034, n/a</u>            | n/a, n/a                               | n/a, n/a                               |
| <i>Dog (complete)</i>           | .038, .032                             | <b>.191, .032</b>                      | .013, .037                  | <u>.013, .037</u>           | <b>2.530, .210</b>                     | <u><b>.044, .101</b></u>               |
| <i>SCAPE (complete)</i>         | .032, .026                             | <b>.143, .026</b>                      | .010, .058                  | <u>.010, .058</u>           | <b>2.744, .348</b>                     | <u><b>.069, .139</b></u>               |
| <i>Cat-Wolf (complete)</i>      | .077, n/a                              | <b>.372, n/a</b>                       | .039, n/a                   | <u>.027, n/a</u>            | <b>3.367, n/a</b>                      | <u><b>.101, n/a</b></u>                |
| <i>Gorilla-Human (complete)</i> | .079, n/a                              | <b>.376, n/a</b>                       | .041, n/a                   | <u>.021, n/a</u>            | <b>3.879, n/a</b>                      | <u><b>.091, n/a</b></u>                |

**Table 1:** Quantitative performance analysis of our method in comparison with Möbius Voting (MV) as bold vs. bold and underlined bold vs. underlined bold.




**Figure 10:** Partial correspondences obtained by our method and by MV respectively on the left and right hand sides of each box for four different datasets. MV matches corresponding to our sparse matches are shown with larger spheres, whereas smaller spheres represent the remaining matches.

on *SHREC11* since the models in the partial class contain holes on the surface. The ground-truth correspondence information is not available for hybrid pairs and *SHREC11*.

The execution times of our isometric part matching algorithm on a 2.53GHz PC is, for the highest-resolution dataset *SHREC11* of 50K vertices, 30 seconds for feature selection, followed by 0.6 seconds combinatorial matching of  $M = 5$  samples and then 0.4 seconds dense matching of  $U = 100$  samples. The respective seconds for the lowest-resolution dataset *SCAPE* of 12.5K vertices are 3.6, 0.3, and 0.3. Note that the cubic complexity of the promoted scale-invariant measure  $D_{\text{iso}}^{(1)}$  becomes negligible compared to the quadratic  $D_{\text{iso}}^{(2)}$  with  $M = 5$ .

When we investigate how the choice of  $M$ , i.e., the size of the sought coarse mapping, affects the output, we observe that small values, e.g.,  $M = 3$ , may cause unstable results as the individual distortions making up  $D_{\text{iso}}^{(1)}(\mathcal{S})$  become less reliable, whereas with relatively high values such as  $M = 7$ , other than computational load concerns, inaccuracies may arise due to joint sampling that starts to produce incompat-

ible point pairs to be matched. In Fig. 11, we see that although the initial samples which may contain some inconsistent points are matched as accurately as possible in all cases, the quality of the final dense correspondences may degrade due to these inconsistencies with  $D_{\text{grd}} = .201, .058$ , and  $.107$  values on average over *SCAPE* dataset for  $M = 3, 5$ , and  $7$ , respectively. We note that the proper choice of  $M$  can also be considered as model or database dependent. On the *Dog* dataset for example, which contains shapes with more extremities, the choice of  $M = 5$  or  $M = 7$  does not affect the performance with  $D_{\text{grd}} = .037$  and  $.038$  values on average, respectively (for the complete matching scenario).



**Figure 11:** Different choices of  $M$  in creation of the initial coarse mapping on a *SCAPE* shape pair.

## 8. Conclusion

We have introduced a novel scale normalization method, in comparison with a state of the art method, that proves useful for partial/complete shape correspondence under a particular setting where one of the shapes to be matched is a scaled and isometric part of the other. Our method uses an isometric distortion measure in a combinatorial framework to establish a trusted coarse correspondence based on which the shapes are brought to the same scale, hence handling the scale problem. Dense matching then enables comparison of our approach with state of the art as well as alleviates the symmetric flip problem due to initial coarse sampling. Our isometric matching solution naturally covers the complete matching problem, where the novel distortion measure has proved to be very effective while matching hybrid pairs of semantically similar shapes whose maximum geodesic distances do not necessarily coincide, a case that would fail most conventional isometric shape matchers.



As a partial shape matcher, there is still room for improvement in this framework. We assume that the sample set on the partial model is (approximately) a subset of the complete model sample set which currently has a cardinality of 10. While this setting is sufficient when the partial model is at least the half of the complete model, as in general and in all examples of this work, the assumption breaks when the part is too small, e.g., Hand vs. Human, since in this case the complete model will probably lack most of the samples featured in the partial model. A simple solution is just to populate the target samples to cover all extremities, instead of only 10 evenly-spaced ones, which would work up to 20 samples due to combinatorial complexity.

As a scale normalization tool, we see the work complete after thorough comparisons and experiments. The most straightforward example to an application for which our method can be used in its current form is part retrieval, i.e., searching a 3D database for shapes containing a surface part given as query. We also note that the proposed scale-invariant isometric distortion measure can be embedded into other frameworks that address the partial shape correspondence problem in a more general setting, such as matching shape pairs which are partially isometric but both having parts which are not in common. That would however require the envisagement of a mechanism to eliminate the shape outliers (the uncommon parts) from the global pool of shape vertices, such as the voting approach employed in [LF09], which will be the topic of our further research.

### Acknowledgement

This work has been supported by TUBITAK under the project EEEAG-109E274.

### References

- [ACOT\*10] AU O. K.-C., COHEN-OR D., TAI C.-L., FU H., ZHENG Y.: Electors voting for fast automatic shape correspondence. *Computer Graphics Forum (Proc. Eurographics)* 29, 2 (2010), 645–654. 1
- [ASK\*05] ANGUELOV D., SRINIVASAN P., KOLLER D., THRUN S., RODGERS J., DAVIS J.: Scape: shape completion and animation of people. *ACM Trans. Graph.* 24, 3 (2005), 408–416. 5
- [BBB\*11] BOYER E., BRONSTEIN A., BRONSTEIN M., BUSTOS B., ET AL.: Shrec 2011: robust feature detection and description benchmark. *Eurographics Workshop on 3D Object Retrieval* (2011). 5
- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Scientific Computing* 28, 5 (2006), 1812–1836. 1
- [BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008. 1, 5
- [BK10] BRONSTEIN M. M., KOKKINOS I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. *Proc. Computer Vision and Pattern Recognition (CVPR)* (2010). 1
- [dAST\*08] DE AGUIAR E., STOLL C., THEOBALT C., AHMED N., SEIDEL H. P., THRUN S.: Performance capture from sparse multi-view video. *Proc. SIGGRAPH* (2008). 5
- [ELPZ97] ELДАР Y., LINDENBAUM M., PORAT M., ZEEVI Y.: The farthest point strategy for progressive image sampling. *IEEE Trans. Image Processing* 6 (1997), 1305–1315. 3
- [FS06] FUNKHOUSER T., SHILANE P.: Partial matching of 3d shapes with priority-driven search. *Symposium on Geometry Processing (SGP)* (2006). 1
- [GMGP05] GELFAND N., MITRA N., GUIBAS L., POTTMAN H.: Robust global registration. *Symposium on Geometry Processing (SGP)* (2005). 1
- [HAWG08] HUANG Q., ADAMS B., WICKE M., GUIBAS L.: Non-rigid registration under isometric deformations. *Symposium on Geometry Processing (SGP)* (2008), 1149–1458. 1
- [HSKK01] HILAGA M., SHINAGAWA Y., KOHMURA T., KUNII T.: Topology matching for fully automatic similarity estimation of 3d shapes. *Proc. SIGGRAPH* (2001). 3
- [JZ06] JAIN V., ZHANG H.: Robust 3d shape correspondence in the spectral domain. *IEEE Int. Conf. on Shape Modeling and Applications (SMI)* (2006), 118–129. 2
- [Kol09] KOLMOGOROV V.: Blossom V: A new implementation of a minimum cost perfect matching algorithm. *Mathematical Programming Computation (MPC)* (2009), 43–67. 4
- [LF09] LIPMAN Y., FUNKHOUSER T.: Mobius voting for surface correspondence. *ACM Trans. Graph.* 28, 3 (2009). 2, 5, 8
- [MHK\*08] MATEUS D., HORAUD R., KNOSSOW D., CUZ-ZOLIN F., BOYER E.: Articulated shape matching using laplacian eigenfunctions and unsupervised point registration. *Proc. Computer Vision and Pattern Recognition (CVPR)* (2008), 1–8. 2
- [OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L.: One point isometric matching with the heat kernel. *Computer Graphics Forum* 29, 5 (2010), 1555–1564. 1
- [PBB11] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M.: A correspondence-less approach to matching of deformable shapes. *Proc. Scale Space and Variational Methods (SSVM)* (2011). 2
- [SMP03] SAND P., MCMILLAN L., POPOVIC J.: Continuous capture of skin deformation. *Int. Conf. on Computer Graphics and Interactive Techniques* (2003). 5
- [SY10] SAHILLIOĞLU Y., YEMEZ Y.: 3d shape correspondence by isometry-driven greedy optimization. *Proc. Computer Vision and Pattern Recognition (CVPR)* (2010), 453–458. 1, 4
- [TBW\*09] TEVS A., BOKELOH M., WAND M., SCHILLING A., SEIDEL H.-P.: Isometric registration of ambiguous and partial data. *Proc. Computer Vision and Pattern Recognition (CVPR)* (2009). 1
- [WBBP11] WANG C., BRONSTEIN M. M., BRONSTEIN A. M., PARAGIOS N.: Discrete minimum distortion correspondence problems for non-rigid shape matching. *Proc. Scale Space and Variational Methods (SSVM)* (2011). 2
- [ZBVH09] ZAHARESCU A., BOYER E., VARANASI K., HORAUD R.: Surface feature detection and description with applications to mesh matching. *Proc. Computer Vision and Pattern Recognition (CVPR)* (2009). 1
- [ZSCO\*08] ZHANG H., SHEFFER A., COHEN-OR D., ZHOU Q., VAN KAICK O., TAGLIASACCHIL A.: Deformation-driven shape correspondence. *Computer Graphics Forum* (2008), 1431–1439. 1
- [ZWW\*10] ZENG Y., WANG C., WANG Y., GU X., SAMARAS D., PARAGIOS N.: Dense non-rigid surface registration using high-order graph matching. *Proc. Computer Vision and Pattern Recognition (CVPR)* (2010). 2