Correspondences of Persistent Feature Points on Near-Isometric Surfaces

Ying Yang[1,2], David Günther[1,3], Stefanie Wuhrer[3,1], Alan Brunton[3,4]
Ioannis Ivrissimtzis[2], Hans-Peter Seidel[1], Tino Weinkauf[1] ⋆

[1]MPI Informatik [2]Durham University [3]Saarland University [4]University of Ottawa

Abstract We present a full pipeline for finding corresponding points between two surfaces based on conceptually simple and computationally efficient components. Our pipeline begins with robust and stable extraction of feature points from the surfaces. We then find a set of near isometric correspondences between the feature points by solving an optimization problem using established components. The performance is evaluated on a large number of 3D models from the following perspectives: robustness w.r.t. isometric deformation, robustness w.r.t. noise and incomplete surfaces, partial matching, and anisometric deformation.

1 Introduction

Intrinsic surface correspondence computation is a highly active research area [1,2]. It is a fundamental aspect of applications such as morphing, texture transfer, geometry synthesis and animation. In general, the search space of intrinsic correspondences between two surfaces is too large to be computationally tractable. A well-established way to reduce the search space is to extract distinctive features from both surfaces and compute correspondence between these. Such features typically also have the benefit of being more reliable to match because of their distinctiveness. Consequently, the search space is reduced in a strategic way. We present an algorithmic pipeline that is able to match robustly feature points between two nearly isometric surfaces.

To effectively match feature points, they should not only be distinctive, but also intuitive and visually meaningful. This is important in visually evaluating correspondence quality on real data, where no ground truth correspondence is available for numerical evaluation. We can get a set of points with these properties from Gaussian curvature. However, Gaussian curvature is greatly affected by noise, resulting in minima and maxima of which only a small subset describe meaningful features. With this in mind, we extract reliable feature sets by using topological persistence [3].

We extract features using topological persistence and compute correspondences using feature descriptors and near isometric matching. In so doing, we

⋆ We thank Art Tevs and Michael Wand for helpful discussions. This work has partially been funded by the MMCI within the Excellence Initiative of the German Federal Government.

present a novel pipeline for feature matching that has the following properties: (1) it is based on established components; (2) it produces accurate and stable correspondences; (3) it is conceptually direct and simple; and (4) it is computationally efficient. We provide an extensive evaluation of our pipeline.

2 Related Work

There are many recent works on feature extraction and surface correspondence, so for conciseness we will only review the most relevant here. For a more exhaustive comparison of correspondence methods, we refer the reader to a recent survey [1] and a recent competition [2].

Feature Extraction: Heat Kernel Signature (HKS) [4] organizes information about the intrinsic geometry of a shape in a multi-scale way that is stable under perturbations of the shape. Features are detected as local maxima of the HKS for large scales. A later variant used persistent homology to filter out unstable features [5]. While we also use persistence to filter features, we apply discrete Morse theory to Gaussian curvature, thus making our approach less computationally costly and conceptually simpler. The difference of Gaussians and histogram of oriented gradients feature operators have been adapted to meshes [6], and applied to matching. These methods require a multi-scale neighborhood structure, whereas our features are efficient to compute and direct, only needing a fixed neighborhood to compute Gaussian curvature.

Correspondence: Möbius voting [7] uses the observation that isometries are a subset of the Möbius group to devise a method for automatic sparse surface correspondence. A high-order Markov Random Field (MRF) formulation of graph matching based on Möbius transforms has also been proposed for both sparse and dense correspondence [8]. Blended intrinsic maps [9] find per-point blending weights for multiple low-dimensional intrinsic maps computed using Möbius voting; these maps are then blended by linear interpolation. As noted by the authors, this approach is limited in its ability to match partially corresponding surfaces due to its global nature. We use these methods for comparison because they are the current state-of-the-art and have demonstrated equal or superior performance to competing algorithms [9]. A feature-based dense correspondence method [10] starts by computing sparse feature correspondences, and then uses a MRF and front-propagation to compute dense correspondence. It is a well-established technique to extract feature points and explore permutations of matches to find a combination with minimal alignment and deformation error [10,11,12]. Two of these methods [10,12] use the geodesic integral or average to extract features, which are computationally more expensive than Gaussian curvature. The other [11] uses principal curvatures, which are not isometry invariant. Many methods make use of the isometry assumption, often using some kind of embedding [1,13], which are often indirect and expensive to compute. If the embedding is global then the method can be expected to have difficulty with partial matching. Other methods based on the isometry assumption con-

sider HKS as local surface descriptor. Given one pair of corresponding points, full correspondence can be computed for two isometric surfaces using HKS [14].

3 The Matching Pipeline

This section presents our pipeline to find feature point correspondences on two near isometric surfaces S and $\widetilde{S}$. We call two surfaces nearly isometric if the ratio of any corresponding geodesic distances is bounded by a constant threshold τ. The idea of this pipeline is to extract features as the most dominant extrema of the Gaussian curvature in terms of persistence, an established technique in the visualization community (Section 3.1), and to then find near isometric correspondences between these feature sets using modifications of established algorithms in the vision and geometry processing communities (Section 3.2).

3.1 Feature Points and Persistence

In this work, we interpret feature points as extremal points of a curvature field. Since we assume isometry, we use Gaussian curvature. In recent years, several techniques to compute this quantity were proposed. We use a simple quadratic least-square fitting to the underlying point cloud to compute the Gaussian curvature [15]. However, our pipeline does not depend on this choice.

We consider the scalar field formed by the Gaussian curvature on a surface. Points of minimal and maximal Gaussian curvature are critical points of this scalar field. A robust and consistent way to compute critical points is by means of discrete Morse theory [16]. We use the algorithm by Robins et al. [17] to compute the critical points in a combinatorial fashion. Note that the computed critical points are in a one-to-one correspondence to the topological changes of the lower level sets of the formed scalar field [17].

Numerical issues in the curvature computation and noise may create spurious critical points, which challenge the upcoming matching. To distinguish noise-induced and dominant critical points, we make use of an established importance measure for critical points: *persistence* [3]. It measures the "life time" of connected components and loops considering an evolution of the lower level sets. We denote the most dominant minima and maxima of the Gaussian curvature fields on the surfaces S and $\widetilde{S}$ as feature points X and $\widetilde{X}$, respectively.

3.2 Computing Correspondences of Feature Points

Correspondences between feature points are found as follows. For each feature point, we construct a vector based on the geodesic distances between it and a set of sample points on the surface. We measure the similarity of these vectors and find initial correspondences by solving a minimization problem. We then enforce isometric consistency of the set of correspondence pairs using graph matching, pruning inconsistent matches. Finally, a post-matching method finds additional matches that are consistent with the established correspondences.

Initial Correspondences: We first find an initial correspondence between the feature sets by matching the spatial distribution of feature points. Looking from one feature point $\mathbf{x}$ on S to a uniquely defined set of reference points Y, the distribution of those points depends on the point of view of $\mathbf{x}$ and is unique up to intrinsic symmetry. When measuring the distribution with an isometric quantity such as the geodesic distance, this point of view is invariant under isometric transformations. We represent the view point dependent distribution of Y in a quantitative manner by constructing two sets of reference points Y and $\widetilde{Y}$ of cardinality R from S and $\widetilde{S}$ using *geodesic farthest point sampling* [13] and by considering the quantity $f(\mathbf{x}, \mathbf{y}) = 1/(1 + g(\mathbf{x}, \mathbf{y}))$, where $\mathbf{x} \in X$, $\mathbf{y} \in Y$, and $g(\mathbf{x}, \mathbf{y})$ denotes the geodesic distance between $\mathbf{x}$ and $\mathbf{y}$ on S. The function f measures the influence of the reference points on each feature point, and is designed to allow for partial matching as nearby points are weighed more than distant points, and the local neighborhood therefore has a greater influence.

The *feature vector* $\mathbf{f}_Y$ for a given feature point $\mathbf{x} \in X$ is given by the collection of $f(\mathbf{x}, \mathbf{y}_j)$ for all reference points $\mathbf{y}_j \in Y$ in non-decreasing order. Consider two surfaces S and $\widetilde{S}$ and their respective feature points X and $\widetilde{X}$. Assuming $\mathbf{x} \in X$ is the correspondence of $\widetilde{\mathbf{x}} \in \widetilde{X}$, the corresponding feature vectors $\mathbf{f}_Y(\mathbf{x})$ and $\mathbf{f}_{\widetilde{Y}}(\widetilde{\mathbf{x}})$ are expected to be similar. Hence, we measure the dissimilarity Ψ of two feature vectors by their normalized L^1-distance. Computing the dissimilarity between all feature vectors of X and $\widetilde{X}$ yields a dissimilarity matrix. A good correspondence is found if the sum of all its dissimilarities is small. The aim is therefore to find a minimum assignment through column and/or row permutation to minimize the trace of the dissimilarity matrix. To solve this optimization problem, various optimization algorithms, such as [18], can be used. We in this paper use the Hungarian algorithm [19], which results in a set of correspondences Σ_1.

Isometric Correspondences: In the following, we remove the pairs in Σ_1 that are not consistent with the assumption that deformations should be approximately isometric. We aim to find the largest set Σ_2 of consistent correspondences. These correspondences can be found using a kernel extraction method as proposed by Leordeanu and Hebert [20] and used by Huang et al. [11]. Let $(\mathbf{c}_i, \widetilde{\mathbf{c}}_i)$ denote the i-th correspondence in Σ_1. Any two consistent correspondences $\{\mathbf{c}_i, \widetilde{\mathbf{c}}_i\}$ and $\{\mathbf{c}_j, \widetilde{\mathbf{c}}_j\}$ should satisfy the following near-isometry constraint: the minimum c_{ij} of the two ratios $g(\mathbf{c}_i, \mathbf{c}_j)/g(\widetilde{\mathbf{c}}_i, \widetilde{\mathbf{c}}_j)$ and $g(\widetilde{\mathbf{c}}_i, \widetilde{\mathbf{c}}_j)/g(\mathbf{c}_i, \mathbf{c}_j)$ should be larger than the stretching tolerance τ $(0 < \tau < 1)$. It is known that a set of correspondences satisfying this condition can be found using a spectral method on a matrix $\mathbf{M}$ that depends on c_{ij} and τ. For more details, refer to [11].

Final Correspondences: As the set Σ_2 might not contain all near-isometric feature point matches, we add additional pairs of feature points in the final step. The additional correspondences are found based on a modified geodesic triangulation technique. Let $X_R \subset X$ and $\widetilde{X}_R \subset \widetilde{X}$ denote the sets of the "rejected" feature points for which correspondences have not been found yet. For each point in X_R, we compute a feature vector w.r.t. the matched points similar to above. The only difference is that the feature vector is now ordered w.r.t. an arbitrary but fixed order of the correspondences in Σ_2. We add a new correspondence

Table 1. Parameter settings for all models used in the tests.

	Cat	*Cat (topo. noise)*	*Centaur*	*David*	*Dog*	*Horse*	*Wolf*	*Face 1*	*Face 2*	*Face 3*
κ	0.05	0.008	0.05	0.05	0.04	0.03	0.05	0.04	0.10	0.02
τ	0.72	0.72	0.83	0.83	0.72	0.72	0.72	0.72	0.72	0.72

pair if the feature vectors of two points are symmetric nearest neighbors under the dissimilarity measure Ψ and the new pair respects the isometric threshold τ w.r.t. all correspondences in Σ_2. All pairs that fulfill these conditions are added to the set Σ_3 of feature correspondences, which is initialized by Σ_2.

4 Evaluation

This section validates the proposed pipeline. We implement the pipeline using MATLAB and C++ and test it on a standard PC. We use code from Surazhsky et al. [21] to compute geodesics and code from Cao[1] for the Hungarian algorithm. Our non-optimized implementation takes about 3 minutes to find corresponding points for the *Cat* model and about 1.5 minutes for the *Centaur* model.

We evaluate the algorithm on a large number of 3D models of the TOSCA [13] database and some models of the BU-3DFE [22] database. Similar to Bronstein et al. [2], we define the correspondence error $\mathcal{C}$ as follows:

$$\mathcal{C} = \frac{1}{|\Sigma_3| \cdot d_g} \min \left\{ \sum_{i=1}^{|\Sigma_3|} g(\mathbf{c}_i, \mathbf{c}'_i), \sum_{i=1}^{|\Sigma_3|} g(\mathbf{c}_i, \mathbf{c}''_i) \right\}, \tag{1}$$

where $|\Sigma_3|$ is the cardinality of Σ_3, d_g is the geodesic diameter of neutral pose S, $(\mathbf{c}_i, \widetilde{\mathbf{c}}_i)$ is a correspondence pair in Σ_3, $\mathbf{c}'_i$ and $\mathbf{c}''_i$ are the ground truth correspondence and the symmetric ground truth correspondence of $\widetilde{\mathbf{c}}_i$ in S, respectively, and the geodesic distances $g(\mathbf{c}_i, \mathbf{c}'_i)$ and $g(\mathbf{c}_i, \mathbf{c}''_i)$ are measured on S. Here, the symmetric ground truth $\mathbf{c}''_i$ is defined as the ground truth mapping of $\mathbf{c}_i$ to its intrinsically symmetric part on the shape, given by flipping the left and right sides of the model.

Our algorithm involves three parameters: the persistence threshold κ, the cardinality R of the set Y of sample points and the deformation threshold τ. Regarding parameter settings, we fix $R = 800$ and each of the other two parameters at one consistent value per model, with the exception of the topological noise example, as shown in Table 1.

Fig. 1 (left) shows the influence of the κ and τ parameter values on the results of matching two *Cat* models. A correspondence $\mathbf{c}_i$ is considered close to its ground truth $\mathbf{c}'_i$ if $g(\mathbf{c}_i, \mathbf{c}'_i) < 0.05 d_g$, close to its symmetric ground truth $\mathbf{c}''_i$ if $g(\mathbf{c}_i, \mathbf{c}''_i) < 0.05 d_g$, and a mismatch otherwise. As expected, as the persistence threshold κ increases, the number of features decreases and as the stretching threshold τ increases, the number of matched pairs decreases.

[1] http://www.mathworks.com/matlabcentral/fileexchange/20328, 2008

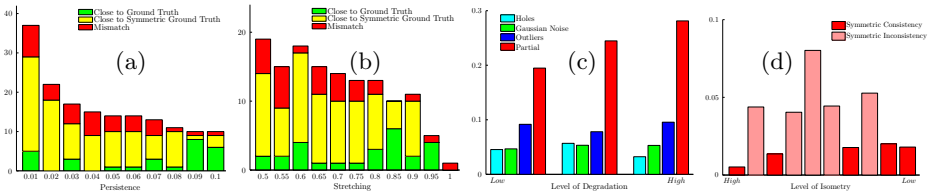

Figure 1. (a,b): Influence of parameter values on matching two clean *Cat* models. The x-axes show the thresholds and the y-axes show the number of matches. (c,d): Correspondence error $\mathcal{C}$ for the *Cat* model with different types and levels of degradation. Each bar in (c) is the mean over all model pairs, while each bar in (d) is for one pair.

4.1 Synthetic Evaluation

We evaluate the robustness of our algorithm on TOSCA models from four perspectives: isometric deformation, different categories of noise, different object matching and partial matching. Whenever we match two shapes from the same object class, we match the deformed/noisy model to the clean shape of the same object class in neutral pose. To evaluate the robustness against noise, we artificially introduced five different kinds of noise to some models. First, we added three levels of *Gaussian noise* to the deformed versions of the *Cat*, *Centaur*, *David*, *Dog*, *Horse*, and *Wolf* models (38 models total). The variances of Gaussian noise used in the experiments are 20%, 40% and 60% of the model's bounding ball radius. Second, we added three levels of *outliers* to the aforementioned 38 models by moving a vertex in the direction of its outer normal with probability 0.004 by varying the strength of the offset. The outliers are modeled as a type of shot noise that is typically present in scanner data from multi-view camera systems. The models are corrupted by moving a vertex in the direction of its outer normal with probability 0.004. We use three levels of outliers by varying the strength of the offset. Third, we added three levels of *holes* to the deformed versions of the *Cat* model (10 models total). The first level removes the one-ring neighborhood of a set of vertices distributed over the surface. The second and third levels enlarge the holes by removing all triangles that are on the boundary of the model. Fourth, we remove parts of the models in three levels to simulate *partial matching*. For each model in neutral pose, we remove a part by cutting the model with a plane parallel to the symmetry plane of the model. The three levels remove $17\%, 33\%$, and 50% of the model's bounding box, respectively. The partial models are then deformed into all other poses to generate all partial models. Finally, we added *topological noise* to one of the *Cat* models.

Fig. 1 (right) shows the correspondence errors $\mathcal{C}$ for the *Cat* models with near-isometric deformations and different types of noise. Note that the correspondence quality does not degrade significantly for increasing levels of Gaussian noise, outliers or holes. As expected, for increasing levels of partial matching, the quality of the correspondence degrades more than for the other types of noise. However, even in case where 50% of the surface was removed, the average cor-

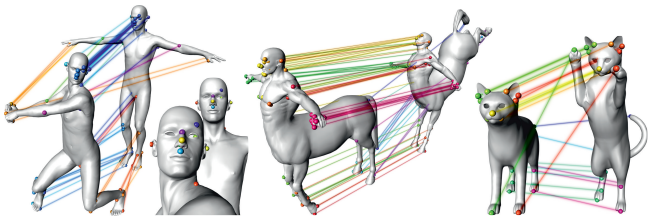

Figure 2. Matching results for different models, where correspondences are shown in the same color and connected by a line.

respondence error is below 30% of the geodesic diameter. Similar plots for the other model classes are shown in Appendix B of the supplementary material.

Non-Isometric Deformation: Fig. 1 (right) shows $\mathcal{C}$ for the correspondences computed between pairs of *Cat* models. For all of the models that have a mean correspondence error above 0.03, we encounter the following problem. Some points on S correspond to points close to their ground truth correspondences on $\widetilde{S}$, while other points on S correspond to their symmetric ground truth correspondences on $\widetilde{S}$. Hence, while all correspondences are locally acceptable, the correspondence map is globally inconsistent, which leads to a large value of $\mathcal{C}$. We call this problem *symmetric inconsistency* in the following, and the matching of the legs of the *Dog* and *Wolf* models in Fig. 4 shows an example. However, the symmetric inconsistency only affects very few feature points as can be seen in the bar plots of Fig. 1 (left). The majority of feature points are correctly matched w.r.t. the ground truth or the symmetric ground truth. Fig. 2 shows some qualitative results.

Gaussian Noise: As Gaussian noise will change the intrinsic geometry of the shape, we adjust the parameter τ depending on the specific level of noise. Appendix A of the supplementary material discusses how to relax τ. Basically, τ decreases with the increase of the noise level. This is because stronger noise will create a greater deformation than weaker noise does. Fig. 3 illustrates matching a *Centaur* model degraded by Gaussian noise and its corresponding clean model in neutral pose. As demonstrated by both the small correspondence errors in Fig. 1 (right) and the qualitative results in Fig. 3, our method is able to match feature points even in the presence of Gaussian noise.

Outliers: Fig. 3 shows an example of matching a *Horse* degraded by outliers and its corresponding clean model in neutral pose. Both the numerical evaluation in Fig. 1 (right) and the qualitative results in Fig. 3 demonstrate that our algorithm is able to find high-quality correspondences, when applied to data with outliers.

Holes: Fig. 1 (right) shows that our method still provides comparable performance as for clean models in terms of correspondence quality, although the existence of holes might potentially result in significant changes in geodesic paths. Fig. 3 illustrates qualitatively that feature points are correctly matched.

Partial Matching: Fig. 3 shows an example of matching a partial *Cat* model to a complete *Cat* model. The feature points appearing in both models are visually

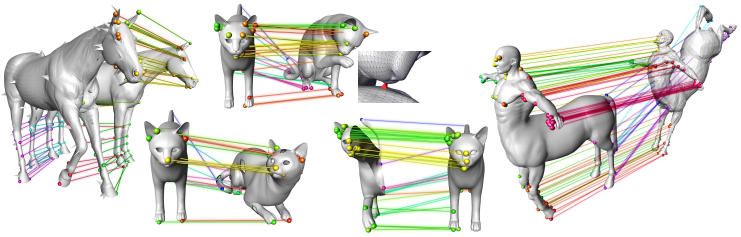

Figure 3. Matching results for different models corrupted by synthetic noise. From left to right: outliers, holes, topological noise, partial information, and Gaussian noise.

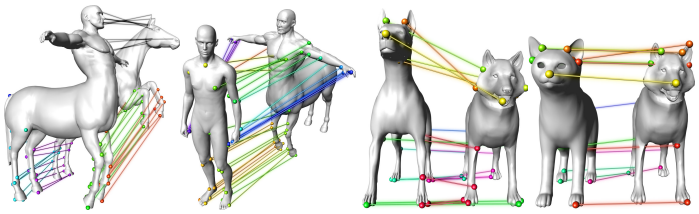

Figure 4. Matching between different object classes.

matched correctly. In a second experiment, we are interested in finding corresponding pairs of vertices for shapes from different object classes.The partial matching results are illustrated in Fig. 4 (left). We observe that feature points describing semantically the same region are correctly matched. For instance, observe the hands and upper body between *Centaur* and *David*. However, feature points could not be matched correctly in regions of the surfaces that are semantically different, as expected. This can be seen in Fig. 4 at the head of *Horse* and the head of *Centaur*.

Topological Noise: Topological noise significantly changes the intrinsic geometry of the surface, and is thus expected to cause problems for our algorithm. Fig. 3 shows the correspondences on a *Cat* model with topological noise.

Different Object Matching: Finding correspondences between two objects of different classes is challenging, since the surfaces are far from isometric. However, our pipeline is able to match most of the feature points correctly, as can be seen in Fig. 4 (right).

4.2 Practical and Comparative Evaluation

To assess the real-world applicability of our algorithm, we compute correspondences between different 3D face scans from the BU-3DFE database. These tests are challenging because the meshes have inconsistent topology and different local shape features. The results here are presented visually, as no ground truth is available to evaluate numerically. The matching results are shown in Fig. 5.

We compare the proposed method with two state-of-the-art 3D matching algorithms: Möbius voting [7] and blended intrinsic maps [9] using the code

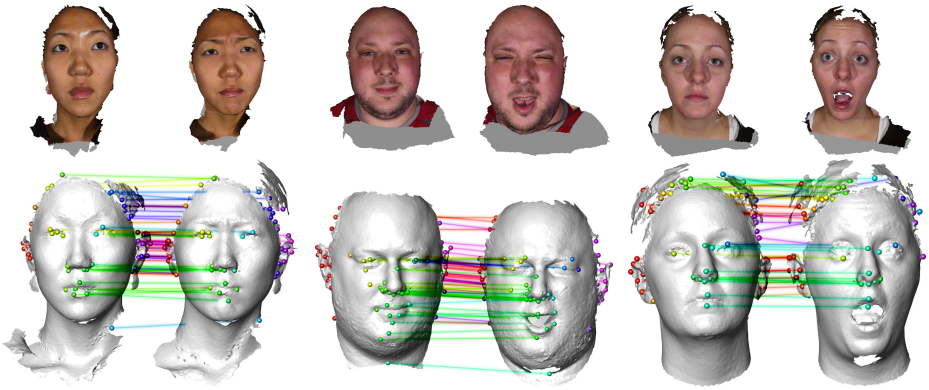

Figure 5. Correspondences of scans with different facial expressions. The top row shows the textured raw scans and the bottom row shows our results.

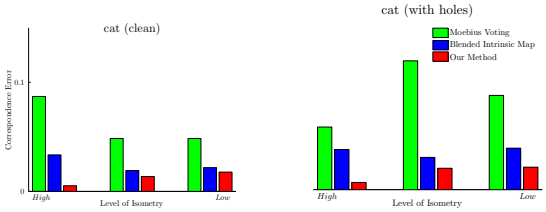

Figure 6. Comparison for *Cat* without noise and with second level of holes.

released by the authors[2]. Note that the implementation of Möbius voting may not reflect all the details of the original implementation. We compare to these two methods on pairs of models of the same object from TOSCA for which our method does not encounter the symmetric inconsistency problem. Fig. 6 shows the correspondence errors $\mathcal{C}$ for models with non-isometric deformations and with holes. Note that our method compares favorably to previous approaches. We observed this trend for different models and types of noise.

Limitations of the proposed method include the symmetric inconsistency problem and difficulty handling topological noise and non-isometric deformation. These limitations are due to the heavy dependence on geodesics.

5 Conclusion

We have presented a feature-based approach to find corresponding points for any two given surfaces. The performance is evaluated on a large database, which contains not only clean models, but also data acquired with 3D scanners. We show that our method is robust against isometric deformations and different types of noise, including Gaussian noise, outliers, holes, topological noise and

[2] http://www.cs.princeton.edu/ vk/CorrsCode/, 2011

scanner noise. Correct correspondences are found even in the case of partial matching. Future work includes overcoming the symmetric inconsistency issue encountered when matching surfaces with intrinsic symmetries.

References

1. van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. CGF **30** (2011) 1681–1707
2. Bronstein, A., Bronstein, M., Castellani, U., Dubrovina, A., Guibas, L., Horaud, R., Kimmel, R., Knossow, D., von Lavante, E., Mateus, D., Ovsjanikov, M., Sharma, A.: SHREC 2010: robust correspondence benchmark. In: 3DOR. (2010)
3. Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. DCG **28** (2002) 511 – 533
4. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. In: SGP. (2009) 1383–1392
5. Dey, T., Li, K., Luo, C., Ranjan, P., Safa, I., Wang, Y.: Persistent heat signature for pose-oblivious matching of incomplete models. CGF **29** (2010) 1545–1554
6. Zaharescu, A., Boyer, E., Varanasi, K., Horaud, R.: Surface feature detection and description with applications to mesh matching. In: CVPR. (2009) 373–380
7. Lipman, Y., Funkhouser, T.: Möbius voting for surface correspondence. TOG (Proc. SIGGRAPH) **28** (2009) 72:1–72:12
8. Zeng, Y., Wang, C., Wang, Y., Gu, X., Samaras, D., Paragios, N.: Dense non-rigid surface registration using high-order graph matching. In: CVPR. (2010) 382–389
9. Kim, V., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. TOG **30** (2011) 79:1–79:12
10. Tung, T., Matsuyama, T.: Dynamic surface matching by geodesic mapping for 3d animation transfer. In: CVPR. (2010) 1402 – 1409
11. Huang, Q., Adams, B., Wicke, M., Guibas, L.J.: Non-rigid registration under isometric deformations. In: SGP. (2008) 1149–1458
12. Zhang, H., Sheffer, A., Cohen-Or, D., Zhou, Q., van Kaick, O., Tagliasacchi, A.: Deformation-driven shape correspondence. CGF (Proc. SGP) **27** (2008) 1393–1402
13. Bronstein, A., Bronstein, M., Bronstein, M., Kimmel, R.: Numerical geometry of non-rigid shapes. Springer (2008)
14. Ovsjanikov, M., Merigot, Q., Memoli, F., Guibas, L.: One point isometric matching with the heat kernel. CGF (Proc. SGP) **29** (2010) 1555–1564
15. Cazals, F., Pouget, M.: Estimating differential quantities using polynomial fitting of osculating jets. In: SGP. (2003) 177 – 187
16. Forman, R.: Morse theory for cell-complexes. Adv. in Math. **134** (1998) 90–145
17. Robins, V., Wood, P., Sheppard, A.: Theory and algorithms for constructing discrete morse complexes from grayscale digital images. TPAMI **33** (2011) 1646 –1658
18. Stošić, M., Marques, M., Costeira, J.: Convex solution of a permutation problem. Linear Algebra and its Applications **434** (2011) 361–369
19. Kuhn, H.: The hungarian method for the assignment problem. Naval research logistics quarterly **2** (1955) 83–97
20. Leordeanu, M., Hebert, M.: A spectral technique for correspondence problems using pairwise constraints. In: ICCV. (2005) 1482 –1489
21. Surazhsky, V., Surazhsky, T., Kirsanov, D., Gortler, S., Hoppe, H.: Fast exact and approximate geodesics on meshes. TOG (Proc. SIGGRAPH) **24** (2005) 553–560
22. Yin, L., Wei, X., Sun, Y., Wang, J., Rosato, M.: A 3d facial expression database for facial behavior research. In: FG. (2006) 211–216