

Reliable Isometric Point Correspondence from Depth

Emel Küpcü, Yücel Yemez
 Computer Engineering Department, Koç University
 Istanbul, Turkey
 {ekupcu, yyemez}@ku.edu.tr

Abstract

We propose a new iterative isometric point correspondence method that relies on diffusion distance to handle challenges posed by commodity depth sensors, which usually provide incomplete and noisy surface data exhibiting holes and gaps. We formulate the correspondence problem as finding an optimal partial mapping between two given point sets, that minimizes deviation from isometry. Our algorithm starts with an initial rough correspondence between keypoints, obtained via a standard descriptor matching technique. This initial correspondence is then pruned and updated by iterating a perfect matching algorithm until convergence to find as many reliable correspondences as possible. For shapes with intrinsic symmetries such as human models, we additionally provide a symmetry aware extension to improve our formulation. The experiments show that our method provides state of the art performance over depth frames exhibiting occlusions, large deformations and topological noise.

1. Introduction

Depth sensors have become commodity in the last half decade, and this has opened up new opportunities in the field of computer vision and graphics as well as new challenges. Finding correspondences from depth is a key step for the success of various tasks in 3D computer vision, such as registration [10] and reconstruction [27].

Although the field of 3D shape correspondence has become quite mature in the last decade, finding reliable correspondences from depth, especially for non-rigid objects, is still an open problem. The first challenge is that the depth data is incomplete by acquisition since objects can be sensed only from one direction; hence correspondences exist only partially. Second, while estimating correspondences on shapes with intrinsic symmetries such as human body, symmetric flip issues commonly complicate the matching process. Third, noisy data provided by commodity depth sensors, exhibiting holes and large gaps, makes estimation of geodesic distances on the surface geometry very

difficult. Moreover, when objects undergo non-rigid deformation, their topology can change drastically, which makes computation of geodesic distances inconsistent between the poses; hence reliable isometric point-based matching techniques that are able to overcome the problematic structure of depth frames are currently less mature compared to mesh-based techniques.

The most common and generic approach for non-rigid point correspondence is to match individual surface points based on local shape descriptors [4, 9, 20, 48, 56]. However, since a local approach discards global shape cues such as isometry, it can easily yield incorrect correspondences especially when the shapes exhibit large variations in their local geometry, or when there are many points that are locally similar. A number of works in the literature address the problem of isometric point matching [2, 16, 18, 26, 32, 34, 55]. Most of these works however perform poorly in the case of noisy and incomplete data since they rely on fitting intermediate mesh-based representations to point clouds and/or computation of geodesic distances. In this paper, we present a new mesh-free point-based method which can estimate reliable sparse correspondences on non-rigid objects undergoing large isometric deformations from noisy and incomplete depth data. We show that starting from an initial correspondence obtained by any standard descriptor matching technique, it is possible to iteratively prune and update the initial matching, and obtain a set of sparse but reliable correspondences even with challenging noisy and incomplete data under occlusion. We additionally provide a symmetry-aware extension to our formulation in order to further boost our correspondence results on shapes with intrinsic symmetries such as human models. In our experimental results, we have cases where the initial correspondences are completely incorrect with 0% precision and we still obtain up to 100% precision. On average, we improve the ground truth error of the initial correspondences by 23x.

We formulate the correspondence problem as finding an optimal partial mapping between two given point sets, minimizing deviation from isometry. We measure deviation from isometry based on a diffusion-based distance metric

that we compute in a robust manner over noisy point clouds using an approximation of graph Laplacian. At the core of our method, we make use of an iterative pruning algorithm that outputs as many *reliable* correspondences as possible.

Related work. Isometric deformations are the most common forms of non-rigidity. There exist a mature literature on isometric 3D correspondence methods that can find sparse [6, 45, 47, 51, 63] and/or dense [7, 11, 57, 21, 38, 44, 49, 52] accurate correspondences between complete mesh representations of 3D shapes. Yet, finding correspondences between representations with partial similarity is still an active research topic with solutions devised to specifically handle partially correspondent mesh models [6, 5, 14, 29, 30, 39, 51, 58, 61, 62].

Many isometric correspondence methods for mesh structures rely on geodesic distance information [6, 11, 51]. However, conventional ways of computing geodesic distances such as shortest path algorithms become invalid on noisy surfaces with holes and gaps. A better alternative for noisy mesh data is employing diffusion-based distance [7, 13, 45, 62]. Diffusion distance takes into account all paths between two keypoints, thereby reducing the negative effects of topological noise and incomplete data on reliable estimation of distances.

There are isometric point correspondence estimation techniques that handle unorganized point cloud data [2, 8, 16, 18, 23, 26, 46, 55, 60]. Most of these methods, except [2, 16, 23, 46, 55], are not actually mesh-free techniques, relying on intermediate mesh representations fit to input point clouds, so that geodesic distances can be computed.

The mesh-free method presented by Guo *et al.* [16] addresses the correspondence problem through piecewise rigid point registration by discovering parts in an iterative process. Their method hence relies heavily on correct estimation of rigid parts as well as approximation of geodesic distances by k -nearest neighbor graph distances. This is especially problematic in the case of occlusions as with depth data provided by commodity sensors. The base correspondence estimation part of our algorithm looks in spirit like the RANSAC-like method [55] which however uses a probabilistic approach relying again on geodesic distances. Another related method is the point-based dense correspondence technique of [2], which uses medial diffusion to deal with incomplete data. However, their medial axis prior is mostly dependent on shape topology and intolerant to large missing data, making it inapplicable to partial depth data with severe occlusions. Kovnatsky *et al.* [23] extend the initial functional map framework proposed in [38] so as to handle data with missing parts and partially similar models. Rodola *et al.* [46] then take the functional map framework one step further to target partial correspondence problem, but they do not explicitly address topological change and symmetry problems, both of which are targeted by our ap-

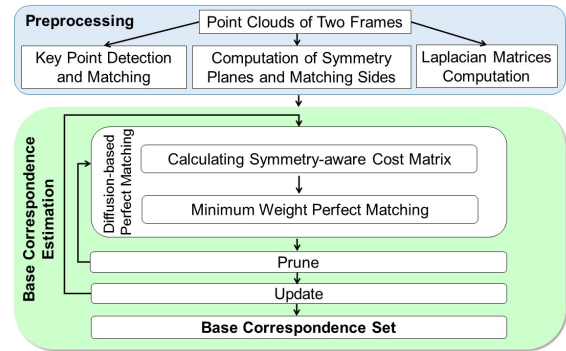


Figure 1. Block diagram of our point correspondence algorithm.

proach. They employ functional maps by calculating Laplacian over point clouds or meshes in order to provide dense correspondences, while we aim to provide sparse but reliable correspondences.

Probabilistic non-rigid registration techniques also as a by product provide us with point correspondences between point clouds [3, 12, 17, 19, 33, 34, 36]. They find correspondences by optimizing a global objective to align point sets. Myronenko and Song [36] and Ma *et al.* [33] introduce non-rigid point registration methods that estimate parameters of transformations using Gaussian mixture models (GMMs). In addition to global cues, Ma *et al.* [33] incorporate local features to take into account similarity of neighboring structure of the points. While point registration methods generate correspondences by matching all the points available, our focus is on finding partial mappings between point clouds based on isometric cues, with as many reliable correspondences as possible.

The symmetric flip problem is inherent to all isometric matching methods. There exist several shape correspondence methods in the literature, addressing explicitly this problem such as in [15, 38, 39, 50, 61, 62, 63]. These methods are however all mesh-based. The symmetric flip problem is also closely related to the problem of symmetry detection, which is addressed in various works [25, 28, 35, 40, 43, 53, 59]. In our correspondence estimation scheme, we employ the method of Lipman *et al.* [28] which detects global intrinsic symmetries on noisy, partial models with non-rigid deformation. Their methodology of generating the symmetry orbit for each point provides us with reflectional symmetry planes that we make use of to resolve symmetric flip ambiguities in frontal human shapes under a point-based setting.

2. Overview

Our reliable isometric point correspondence algorithm consists of two main parts as visualized in Figure 1: preprocessing and base correspondence estimation. The input is a pair of point cloud representations obtained from depth frames of the object of interest. In the preprocessing step,

we first apply a standard keypoint detection technique, and then match the detected points using a standard descriptor based matching algorithm. The resulting matching serves as the initial correspondence to be further improved. In this step, we also compute the graph Laplacian matrices of both point clouds, which we will later use to calculate diffusion based distances between keypoints. An optional task in this step is computation of the symmetry planes and matching the corresponding sides, to be carried out only if the given shapes are symmetric such as human shapes.

During base correspondence estimation, the initial correspondence set computed in the preprocessing step is iteratively pruned and updated employing a diffusion-based perfect matching technique until convergence. This step contains two nested loops, the outer one for update and the inner one for pruning. For the diffusion-based matching task of the inner loop, we construct an isometric cost matrix and apply perfect matching. Each entry of the cost matrix represents the deviation of a given correspondence pair from isometry, which can be computed using diffusion distances only if a set of (base) correspondences is known a priori; hence the need for the initial correspondence obtained via descriptor matching. The diffusion-based matching yields an updated correspondence set that usually includes incorrect matchings due to incomplete and noisy nature of the data, and thus needs to be pruned. The pruning iterations eliminate these outliers one by one based on an isometric error criterion, each time repeating the diffusion-based matching, hence resulting in a smaller but more reliable set of correspondences. This relatively more reliable set replaces the existing correspondences at the beginning of each outer loop iteration, and is gradually improved until convergence.

To address the symmetrical flip problem that inherently occurs while matching symmetric shapes, we introduce a symmetry-aware version of the cost matrix used during the perfect matching phase. This version penalizes matching of the points on the non-corresponding sides of the shapes in proportion to their distances from the symmetry planes.

3. Preprocessing

3.1. Point-based Laplacian

To calculate the diffusion distance between two given points on a point cloud, we compute the graph Laplacian directly on point representations in a similar way as described in Belkin and Niyogi [1]. To do this, we first find, for each point x_n in the point set $X = \{x_1, x_2, \dots, x_N\}$, the K nearest neighboring points $\{y_1, y_2, \dots, y_K\}$ within a predesignated fixed distance threshold and generate a graph structure. We then calculate a weight $w(i, k)$ between x_i and each y_k in the neighborhood (all initially set to 0) by $w(i, k) = e^{-\frac{\|x_n - y_k\|^2}{\epsilon}}$, where the parameter $\epsilon = \frac{1}{N} \sum_{n=1}^N \max_k \|x_n - y_k\|$, which is related to the av-

erage extent of the neighborhood [42]. We then construct a diagonal matrix D , where each diagonal entry (i, i) is the sum of the weights $w(i, k)$ over k for a given point x_i . Finally, the Laplacian matrix is $L = D - W$. We denote the source and target point clouds extracted from depth frames by P^S and P^T , and their corresponding Laplacian matrices by L_S and L_T , respectively. The smallest M eigenvalues $\lambda_S = \{\lambda_{S,1}, \dots, \lambda_{S,M}\}$, $\lambda_T = \{\lambda_{T,1}, \dots, \lambda_{T,M}\}$ and the corresponding eigenvectors $\phi_S = \{\phi_{S,1}, \dots, \phi_{S,M}\}$, $\phi_T = \{\phi_{T,1}, \dots, \phi_{T,M}\}$ are used for computing diffusion distances as explained in Section 4.1, where M is an experimentally chosen parameter.

3.2. Keypoint detection and matching

For initialization, we detect keypoints on the given pair of point clouds and match them using descriptor matching. Any point-based 3D keypoint detection algorithm, such as SIFT [31] or intrinsic shape signature (ISS) [64], can be used for this purpose. We represent the keypoint sets for the source and target with $S = \{s_1, s_2, \dots, s_{|S|}\}$ and $T = \{t_1, t_2, \dots, t_{|T|}\}$, respectively. We match the detected keypoints based on the Euclidean distance between the descriptors, in two directions: from target to source and from source to target. The intersection of the resulting correspondence sets, referred to as *reciprocal correspondences* [41], is used as the initial base correspondence set B_0 in the next step. For our implementation, we tried various descriptors such as FPFH[48], and eventually chose to employ the SHOT descriptor [56] due to its accuracy and speed.

4. Base correspondence estimation

4.1. Diffusion-based perfect matching

The detected keypoints can be reconsidered for a better matching based on global isometric clues. For this purpose, we construct an isometric cost matrix C , where each entry c_{ij} represents the deviation from isometry of a candidate correspondence pair. To compute deviations from isometry, we need to rely on a set of known correspondences. We refer to this set as base correspondence denoted with $B = \{(b_1^S, b_1^T), (b_2^S, b_2^T), \dots, (b_{|B|}^S, b_{|B|}^T)\}$. The set B is initially set to B_0 , that is, the correspondence obtained in the preprocessing step via local descriptor matching. We also employ diffusion distance [13, 24] as in Eq. (1) for the calculation of the isometric cost.

$$d_{X,t}(x, y) = \sum_{m=1}^M e^{-2\lambda_m t} (\phi_m(x) - \phi_m(y))^2, \quad (1)$$

where the distance is calculated between two points x and y at time t , $\{\lambda_m\}$ and $\{\phi_m\}$ are the (smallest) M eigenvalues and the corresponding eigenvectors of the Laplacian matrix.

The isometric cost of matching a keypoint s_i on the source keypoint set S with t_j on the target T is then calculated as follows:

$$c_{ij} = \frac{1}{|B|} \sum_{(b_i^S, b_l^T) \in B} |d_S(s_i, b_l^S) - d_T(t_j, b_l^T)|. \quad (2)$$

The resulting cost matrix C is bipartite, so we can apply the Blossom V algorithm, minimum-weight perfect matching algorithm of Kolmogorov [22] to match all the keypoints from scratch, similar to [51]. Each entry $c_{ij} \in C$ is normalized to be in the range $[0, 1)$ by $c_{ij} \leftarrow (1 - e^{-c_{ij}})$.

Note that the number of keypoints may vary between point sets, whereas the perfect matching algorithm we employ requires a square cost matrix. Thus, we add virtual nodes to the smaller keypoint set and set the corresponding cost for non-existent pairs to infinity. The perfect matching algorithm results in a one-to-one and onto mapping, from which we then remove the pairs including virtual nodes.

4.2. Extension for symmetric shapes

The isometric cost calculated using diffusion distance as given in Eq. 2 fails to handle the symmetric flip problem which is inherent to all isometric correspondence methods. Hence while dealing with objects with intrinsic symmetries, we propose to use a modified version of this cost function so as to penalize symmetrically flipped correspondences.

To this effect, we first find the symmetry plane on each point cloud by implementing the method of Lipman *et al.* [28]. We then set the orientations of these planes such that the angle between their normals is less than 90 degrees. This setting allows us to match the sides of the shapes. For example, in the case of human shapes, this strategy enables us to differentiate the left and right hand sides of the human shapes assuming that they are both in frontal pose. Both facing backward pose would also work.

We incorporate this orientation information into our cost calculation by penalizing matching of the points that are on non-matching sides of the two shapes, in proportion to their distances from the symmetry planes:

$$c'_{ij} = c_{ij} * (1 + \alpha * \frac{d_{sym}(s_i, t_j) - \min_{k,l} d_{sym}(s_k, t_l)}{\max_{k,l} d_{sym}(s_k, t_l)}) \quad (3)$$

where $d_{sym}(s_i, t_j)$ is the joint distance from points $\{s_i, t_j\}$ to the corresponding symmetry planes calculated by $d_{sym}(s_i, t_j) = \max(d_{sym}^S(s_i), d_{sym}^T(t_j))$, where $d_{sym}^S(s_i)$ and $d_{sym}^T(t_j)$ are the distances to the corresponding symmetry planes. The penalization constant $\alpha > 0$ is experimentally set as $\alpha = 10$. Note that for $\alpha = 0$, the symmetry-aware cost c'_{ij} becomes identical to c_{ij} .

4.3. Iterative pruning and update

The perfect matching process results in a one-to-one and onto mapping B . We iteratively modify this matching B to make it as reliable as possible. We achieve this goal with a nested loop of *pruning* and *update* iterations. While the

inner loop prunes unreliable correspondences based on an isometric error criterion, the outer iterations gradually update the initial correspondence B_0 with which we initialize the base correspondence set in the first place. At the end of these iterations, we expect to end up with a partial one-to-one mapping that establishes a reliable but possibly sparse correspondence between keypoints.

We compute the isometric error, $E_{iso}(b_i^S, b_l^T)$, of a given correspondence pair $(b_i^S, b_l^T) \in B$, in terms of its deviation from isometry with respect to other available pairs in B :

$$E_{iso}(b_i^S, b_l^T) = \frac{1}{|B| - 1} \sum_{\substack{(b_i^S, b_l^S) \in B \\ i \neq l}} |d_S(b_i^S, b_l^S) - d_T(b_i^T, b_l^T)| \quad (4)$$

If $(b_i^S, b_l^T) \in B$ is a correct matching pair, its isometric error is expected to be close to zero. Hence the correspondence set B can be pruned by eliminating the pairs having relatively larger errors compared to others. The reliability of the isometric error defined in Eq. 4 depends on the correctness of B itself. Thus we perform pruning in an iterative scheme, one pair (the worst one) at a time, and each time we reinvoke the diffusion-based perfect matching algorithm with the pruned base correspondence set. At each iteration, we also remove the keypoints of the eliminated pair from the keypoint sets S and T . The pruning and perfect matching tasks are iterated until the gap between maximum and minimum isometric errors over the pairs becomes small enough according to a predesignated threshold value τ .

The modified correspondence set B resulting from the iterative pruning algorithm is smaller but usually much more reliable than the input correspondence. Hence it can be used to update the initial base correspondence for the next run of the iterative pruning algorithm in the outer loop. At the beginning of each outer iteration, the initial base correspondence is renewed with the current B , whereas the keypoint sets, S and T , are set back to their original content. Hence, the keypoints that are discarded during iterative pruning due to mismatches are reconsidered for other possible matches based on a more reliable estimation of isometric errors. The outer iterations terminate when the mean isometric error converges, i.e., when there is no further improvement on the base correspondence set B .

Computational complexity. The complexity of our correspondence estimation algorithm is dominated by the minimum-weight perfect matching algorithm with $O(N_0^2 \log N_0)$ cost [22], where $N_0 = \max(|S|, |T|)$, the number of key points, which is better than $O(Q^2 N_0)$ complexity of the PR-GLS algorithm [33] (one of our baselines), where $Q \sim N_0$ in practice. On the other hand, the CPD algorithm [36], which can be seen as precedent to the PR-GLS method, has $O(N_0)$ time complexity.

5. Results

Datasets. We evaluate the performance of our point correspondence estimation algorithm on two depth datasets: The Berkeley motion human action dataset (MHAD) [37, 54], and the dataset (Human) that we collected, both containing noisy and incomplete depth data of freely moving subjects, captured using Kinect v1. We convert all depth frames to 3D point cloud representations and discard color.

We picked 11 depth frames (Figure 2(a)) exhibiting large non-rigid deformations with respect to each other with visual inspection, and generated their 55 pair combinations using MHAD dataset. Each frame has approximately 26K points and 43 ground-truth marker positions.

The other dataset (Human) that we collected contains RGB-D frames of a human subject exhibiting larger non-rigid motion than the MHAD contains. The Human dataset, in Figure 2(b), includes 6 frames with 33 manually selected ground-truth keypoints out of approximately 22K points on each, and thus 15 model pairs.




Figure 2. The frames of each dataset with ground-truth points displayed in green.

Evaluation. The baselines are the SHOT descriptor matching algorithm as in Section 3.2 using a publicly available implementation (pointclouds.org), and two other state of the art methods for non-rigid point registration with publicly available codes: the CPD method of Myronenko and Song [36] and the PR-GLS method of Ma *et al.* [33]. We compare our method with these three point-based state of the art methods based on two evaluation metrics: 1) deviation from isometry, and 2) deviation from ground-truth. The deviation from isometry, i.e., the isometric error, for a given correspondence set is computed by averaging E_{iso} measure given in Eq. 4 over all pairs in the set. The deviation from ground-truth, i.e., the ground-truth error is calculated by averaging the deviations over all pairs with respect to the ground-truth correspondence. The ground-truth error $E_{\text{grd}}^{\text{base}}$ for a given base correspondence pair (b_i^S, b_i^T) is simply given by $E_{\text{grd}}^{\text{base}}(b_i^S, b_i^T) = d_T(g_i^T, b_i^T)$, where g_i^T represents the ground-truth match for b_i^S . Each of these metrics

is then averaged over the whole dataset.

To make a fair comparison between our method and the baseline methods, we equalize the number of matchings for each pair, by selecting the best R correspondences resulting from each method, where R is the size of the correspondence set with the fewer number of pairs.

Implementation details. We normalize the coordinates of each point cloud in a given dataset so that all the points lie within the unit sphere centered at the origin. For computation of Laplacian matrices, we select up to $K = 120$ neighboring points within a distance threshold. To calculate the distance threshold, we take the mean distance between each point and its K closest neighbors for each point cloud. Then, we set the distance threshold as the maximum of those mean distances within the dataset. If a point has less than $K/4$ neighbors within that distance, we pick the closest $K/4$ neighbors for that vertex regardless of the distance threshold. This helps avoid singularity problems in eigen analysis of the Laplacian matrix. For diffusion distance, we use the smallest $M = 20$ eigenvalues and the corresponding eigenvectors of the Laplacian matrix. Also, the time step parameter t is incremented from 1 to 600 to compute the average diffusion distance. The error threshold coefficients for base correspondence algorithm is set experimentally as $\tau = 2.3$.

Correspondence estimation results. Our initial experiments have shown that incorporation of the symmetry-aware cost function defined in Eq. 3 increases the precision results between 14% and 20% with up to 2.9x ground-truth error improvement. Hence all the results provided in this section are obtained by using the symmetry-aware extension. Note that all the shapes in both datasets are in approximately frontal pose exhibiting reflectional symmetries.

In Table 1, we observe that our algorithm significantly outperforms the baseline algorithm in terms of isometric and ground truth errors. This is especially pronounced in the human dataset which contains very large deformations, mainly because local similarities in this case quickly drop due to occlusions resulting from large motion, whereas the MHAD dataset exhibits more rigidity between model pairs. Therefore, the baseline methods have higher success on MHAD compared to the Human dataset.

In Figure 3, we observe that the worst matchings of our method on the MHAD dataset are reasonably close to the true matchings, while the other methods contain correspondences that are symmetrically flipped or very inaccurate such as matches between hand/foot and hip/head. Our algorithm improves the initial correspondences obtained between ground-truth keypoints via descriptor matching, yielding mostly reliable matchings, pruning relatively unreliable ones, finding correct correspondences even on the left arm with large deformation.

We test the performance of our base correspondence al-

Method	Ground Truth Error ($\times 10^{-5}$)		Isometric Error ($\times 10^{-5}$)		Isometric Error with Key Points ($\times 10^{-5}$)	
	MHAD	Human	MHAD	Human	MHAD	Human
Descriptor matching	4.35	9.93	6.04	7.83	7.88	7.75
CPD	0.53	7.27	4.29	10.30	6.24	9.24
PR-GLS	1.37	5.48	3.84	6.57	6.01	8.51
Our method	0.19	0.44	0.52	0.32	0.97	0.86

Table 1. Quantitative evaluation of our correspondence estimation in comparison to baseline methods.

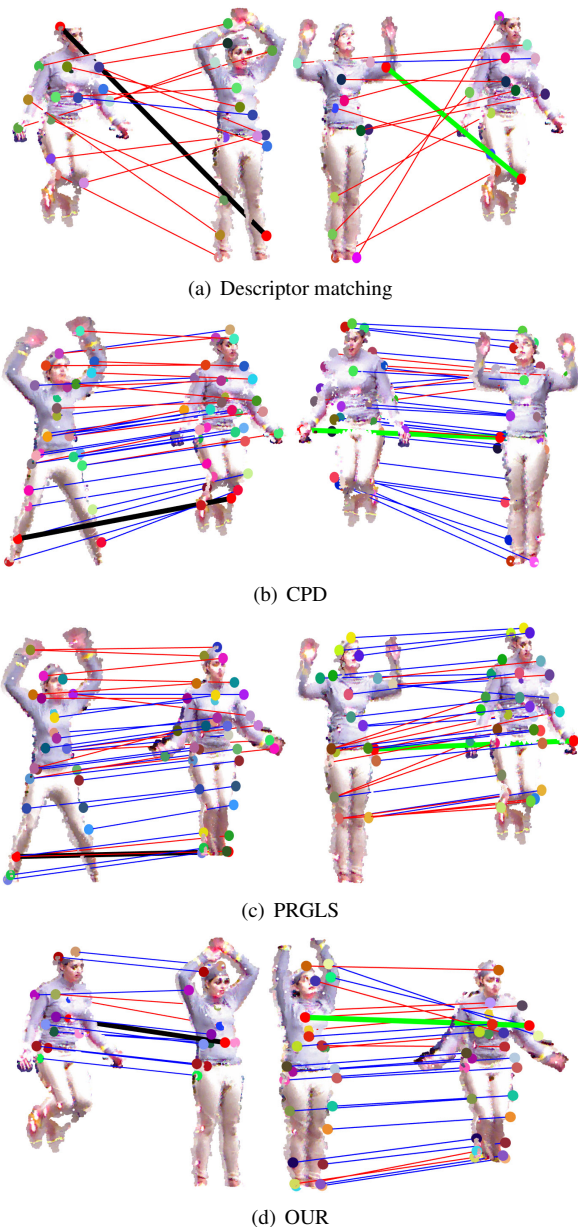


Figure 3. The worst matches on the MHAD dataset for each method according to ground-truth (left column, black lines) and isometric errors (right column, green lines).

gorithm also with automatically detected keypoints. For keypoint detection, we use Intrinsic Shape Signatures (ISS) method [64], though any keypoint detection algorithm could be employed for this purpose. The average number of base correspondences found using this automatic keypoint detection algorithm is 15 for the Human and 20 for MHAD dataset. In this case, we can evaluate the matching results of each method based only on mean isometric error provided in the last two columns of Table 1. Our algorithm still provides more reliable matching results compared to baseline methods in terms of isometric error. We visualize an example pair result of each method in Figure 4, which further supports our improvement over the other methods. We find visually correct correspondences even on the right arm and left leg, where the nonrigid motion is large.

6. Conclusion

We have proposed an isometric mesh-free diffusion-based method to find reliable sparse correspondences between point clouds generated from partial depth data exhibiting noise, large deformations, and occlusions. Our experiments have shown that our method provides state of the art performance on such challenging datasets, particularly on those exhibiting large deformations. We have also provided a symmetry-aware extension that significantly boosts our performance on symmetric human shapes which are in approximately frontal pose. We stress that our method focuses on finding as many reliable correspondences as possible, pruning whenever reliable matching is not possible. Extending our work so as to find denser correspondences and generalizing our symmetry-aware extension will be our future work.

7. Acknowledgements

This work was supported by the Scientific and Technological Research Council of Turkey (TUBITAK) Grants 114E628 and 215E201.

References

- [1] M. Belkin and P. Niyogi. Towards a theoretical foundation for laplacian-based manifold methods. *Journal of Computer and System Sciences*, 74(8):1289–1308, 2008. 3




Figure 4. Example correspondence estimation results using keypoints of ground-truth (top) and automatic key point detection (bottom).

- [2] M. Berger and C. T. Silva. Nonrigid matching of undersampled shapes via medial diffusion. *Computer Graphics Forum*, 31(5):1587–1596, 2012. 1, 2
- [3] S. D. Billings, E. M. Boctor, and R. H. Taylor. Iterative most-likely point registration (impl): a robust algorithm for computing optimal shape alignment. *PloS one*, 10(3):e0117688, 2015. 2
- [4] D. Boscaini, J. Masci, E. Rodolà, M. M. Bronstein, and D. Cremers. Anisotropic diffusion descriptors. *Computer Graphics Forum*, 35(2):431–441, 2016. 1
- [5] A. M. Bronstein and M. M. Bronstein. Not only size matters: regularized partial matching of nonrigid shapes. In *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW'08. IEEE Computer Society Conference on*, pages 1–6. IEEE, 2008. 2
- [6] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences*, 103(5):1168–1172, 2006. 2
- [7] 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. *International Journal of Computer Vision*, 89(2-3):266–286, 2010. 2
- [8] A. Brunton, M. Wand, S. Wuhler, H.-P. Seidel, and T. Weinkauff. A low-dimensional representation for robust partial isometric correspondences computation. *Graphical Models*, 76(2):70–85, 2014. 2
- [9] M. Carrière, S. Y. Oudot, and M. Ovsjanikov. Stable topological signatures for points on 3d shapes. *Computer Graphics Forum*, 34(5):1–12, 2015. 1
- [10] W. Chang and M. Zwicker. Global registration of dynamic range scans for articulated model reconstruction. *ACM Transactions on Graphics (TOG)*, 30(3):26, 2011. 1
- [11] Q. Chen and V. Koltun. Robust nonrigid registration by convex optimization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2039–2047, 2015. 2
- [12] H. Chui and A. Rangarajan. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*, 89(2):114–141, 2003. 2
- [13] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. 2, 3
- [14] L. Cosmo, E. Rodola, J. Masci, A. Torsello, and M. M. Bronstein. Matching deformable objects in clutter. In *3D Vision (3DV), 2016 Fourth International Conference on*, pages 1–10. IEEE, 2016. 2
- [15] A. Dubrovina and R. Kimmel. Approximately isometric shape correspondence by matching pointwise spectral features and global geodesic structures. *Advances in Adaptive Data Analysis*, 3(01n02):203–228, 2011. 2
- [16] H. Guo, D. Zhu, and P. Mordohai. Correspondence estimation for non-rigid point clouds with automatic part discovery. *The Visual Computer*, pages 1–14, 2015. 1, 2
- [17] R. Horaud, F. Forbes, M. Yguel, G. Dewaele, and J. Zhang. Rigid and articulated point registration with expectation conditional maximization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(3):587–602, 2011. 2
- [18] Q.-X. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. *Computer Graphics Forum*, 27(5):1449–1457, 2008. 1, 2
- [19] B. Jian and B. C. Vemuri. Robust point set registration using gaussian mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(8):1633–1645, 2011. 2
- [20] A. E. Johnson. *Spin-Images: A Representation for 3-D Surface Matching*. PhD thesis, CMU-RI-TR-97-47, 1997. 1
- [21] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. *ACM Transactions on Graphics*, 30(4):79, 2011. 2
- [22] V. Kolmogorov. Blossom V: A new implementation of a minimum cost perfect matching algorithm. *Mathematical Programming Computation*, 1(1):43–67, 2009. 4
- [23] A. Kovnatsky, M. M. Bronstein, X. Bresson, and P. Vandergheynst. Functional correspondence by matrix completion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 905–914, 2015. 2
- [24] S. S. Lafon. *Diffusion maps and geometric harmonics*. PhD thesis, Yale University, 2004. 3
- [25] B. Li, H. Johan, Y. Ye, and Y. Lu. Efficient 3d reflection symmetry detection: A view-based approach. *Graphical Models*, 83:2–14, 2016. 2
- [26] H. Li, R. W. Sumner, and M. Pauly. Global correspondence optimization for non-rigid registration of depth scans. *Computer Graphics Forum*, 27(5):1421–1430, 2008. 1, 2
- [27] M. Liao, Q. Zhang, H. Wang, R. Yang, and M. Gong. Modeling deformable objects from a single depth camera. In *ICCV*, 2009. 1
- [28] Y. Lipman, X. Chen, I. Daubechies, and T. Funkhouser. Symmetry factored embedding and distance. In *ACM SIG-*

- GRAPH 2010 Papers*, SIGGRAPH '10, pages 103:1–103:12, New York, NY, USA, 2010. ACM. 2, 4
- [29] Y. Lipman and T. Funkhouser. Möbius voting for surface correspondence. *ACM Transactions on Graphics (TOG)*, 28(3):72, 2009. 2
- [30] O. Litany, E. Rodolà, A. M. Bronstein, M. M. Bronstein, and D. Cremers. Non-rigid puzzles. In *Proceedings of the Symposium on Geometry Processing*, SGP '16, pages 135–143, Goslar Germany, Germany, 2016. Eurographics Association. 2
- [31] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. 3
- [32] J. Ma, J. Zhao, J. Tian, A. L. Yuille, and Z. Tu. Robust point matching via vector field consensus. *IEEE Transactions on Image Processing*, 23(4):1706–1721, 2014. 1
- [33] J. Ma, J. Zhao, and A. L. Yuille. Non-rigid point set registration by preserving global and local structures. *IEEE Transactions on image Processing*, 25(1):53–64, 2016. 2, 4, 5
- [34] D. Mateus, R. Horaud, D. Knossow, F. Cuzzolin, and E. Boyer. Articulated shape matching using Laplacian eigenfunctions and unsupervised point registration. In *CVPR*, 2008. 1, 2
- [35] N. J. Mitra, L. J. Guibas, and M. Pauly. Partial and approximate symmetry detection for 3d geometry. *ACM Transactions on Graphics (TOG)*, 25(3):560–568, 2006. 2
- [36] A. Myronenko and X. Song. Point set registration: Coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2262–2275, 2010. 2, 4, 5
- [37] F. Ofli, R. Chaudhry, G. Kurillo, R. Vidal, and R. Bajcsy. Berkeley MHAD: A comprehensive multimodal human action database. In *WACV*, 2013. 5
- [38] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):30, 2012. 2
- [39] M. Ovsjanikov, Q. Mérigot, F. Méholi, and L. Guibas. One point isometric matching with the heat kernel. *Computer Graphics Forum*, 29(5):1555–1564, 2010. 2
- [40] M. Ovsjanikov, J. Sun, and L. Guibas. Global intrinsic symmetries of shapes. In *Computer graphics forum*, volume 27, pages 1341–1348. Wiley Online Library, 2008. 2
- [41] T. Pajdla and L. Van Gool. Matching of 3-d curves using semi-differential invariants. In *ICCV*, 1995. 3
- [42] F. Petronetto, A. Paiva, E. S. Helou, D. Stewart, and L. G. Nonato. Mesh-free discrete Laplace-Beltrami operator. *Computer Graphics Forum*, 32(6):214–226, 2013. 3
- [43] D. Raviv, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Full and partial symmetries of non-rigid shapes. *International journal of computer vision*, 89(1):18–39, 2010. 2
- [44] D. Raviv, A. Dubrovina, and R. Kimmel. Hierarchical framework for shape correspondence. *Numerical Mathematics: Theory, Methods and Applications*, 6(01):245–261, 2013. 2
- [45] E. Rodola, A. M. Bronstein, A. Albarelli, F. Bergamasco, and A. Torsello. A game-theoretic approach to deformable shape matching. In *CVPR*, 2012. 2
- [46] E. Rodolà, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial functional correspondence. In *Computer Graphics Forum*, volume 36, pages 222–236. Wiley Online Library, 2017. 2
- [47] E. Rodola, A. Torsello, T. Harada, Y. Kuniyoshi, and D. Cremers. Elastic net constraints for shape matching. In *ICCV*, 2013. 2
- [48] R. B. Rusu, N. Blodow, and M. Beetz. Fast point feature histograms (fpfh) for 3d registration. In *ICRA*, 2009. 1, 3
- [49] Y. Sahillioğlu and Y. Yemez. Coarse-to-fine combinatorial matching for dense isometric shape correspondence. *Computer Graphics Forum*, 30(5):1461–1470, 2011. 2
- [50] Y. Sahillioğlu and Y. Yemez. Coarse-to-fine isometric shape correspondence by tracking symmetric flips. *Computer Graphics Forum*, 32(1):177–189, 2013. 2
- [51] Y. Sahillioğlu and Y. Yemez. Partial 3-d correspondence from shape extremities. *Computer Graphics Forum*, 33(6):63–76, 2014. 2, 4
- [52] A. Sharma, R. Horaud, J. Cech, and E. Boyer. Topologically-robust 3d shape matching based on diffusion geometry and seed growing. In *CVPR*, 2011. 2
- [53] I. Sipiran, R. Gregor, and T. Schreck. Approximate symmetry detection in partial 3d meshes. *Computer Graphics Forum*, 33(7):131–140, 2014. 2
- [54] B. Teleimmersion Lab, University of California. Berkeley multimodal human action database (MHAD). <http://teleimmersion.citris-uc.org>. 5
- [55] A. Tevs, M. Bokeloh, M. Wand, A. Schilling, and H.-P. Seidel. Isometric registration of ambiguous and partial data. In *CVPR*, 2009. 1, 2
- [56] F. Tombari, S. Salti, and L. Di Stefano. Unique signatures of histograms for local surface description. In *ECCV*, 2010. 1, 3
- [57] O. Van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 30(6):1681–1707, 2011. 2
- [58] C. Wang, M. M. Bronstein, A. M. Bronstein, and N. Paragios. Discrete minimum distortion correspondence problems for non-rigid shape matching. In *SSVM*, 2011. 2
- [59] H. Wang, P. Simari, Z. Su, and H. Zhang. Spectral global intrinsic symmetry invariant functions. In *Proceedings of Graphics Interface 2014*, pages 209–215. Canadian Information Processing Society, 2014. 2
- [60] L. Wei, Q. Huang, D. Ceylan, E. Vouga, and H. Li. Dense human body correspondences using convolutional networks. In *CVPR*, 2016. 2
- [61] Y. Yoshiyasu, E. Yoshida, and L. Guibas. Symmetry aware embedding for shape correspondence. *Computers & Graphics*, 60:9–22, 2016. 2
- [62] Y. Yoshiyasu, E. Yoshida, K. Yokoi, and R. Sagawa. Symmetry-aware nonrigid matching of incomplete 3d surfaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4193–4200, 2014. 2
- [63] Z. Zhang, K. Yin, and K. W. Foong. Symmetry robust descriptor for non-rigid surface matching, 2013. 2
- [64] Y. Zhong. Intrinsic shape signatures: A shape descriptor for 3d object recognition. In *ICCV Workshops*, 2009. 3, 6