

Supplementary Material: Non-rigid Point Cloud Registration with Neural Deformation Pyramid

I. Introduction

Supplemental material includes

- Document for outlier rejection approach, see [II.](#)
- Document for data filter details in 4DMatch benchmark, see [III.](#)
- **Code.zip** that includes the code of this paper. It contains a README.md file to show the following:
 - anonymous link for data downloading
 - anonymous link for pretrained weights of correspondence model
 - instructions for reproducing the main results presented in this paper
 - license of the data/code
- **Demo.mp4** that shows the optimization process of our non-rigid registration method.

II. Outlier rejection using Transformer

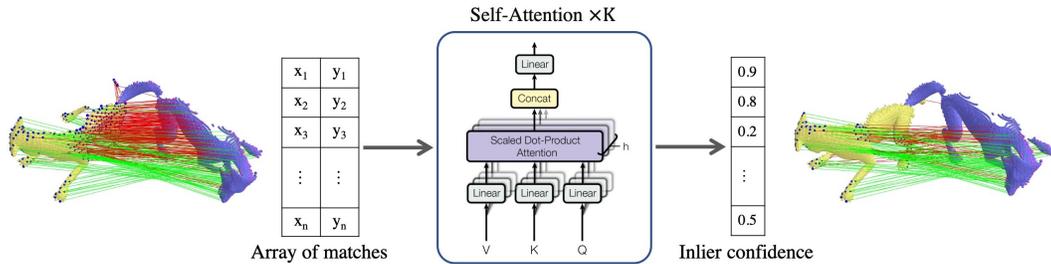


Figure I: Outlier rejection network using Transformer. The red/green lines indicate inliers/outliers.

It is necessary to reject poor correspondences for robust non-rigid registration. Existing learning-based approaches, such as DGR [7] and PointDSC [2], formulate outlier rejection as a binary classification problem, i.e. they use a network to predict the outlier probability of each corresponding point pair. Inspired by these works, we use Transformer to reject outliers. As discussed in DGR [7], inlier correspondences between two rigid scans should form a manifold geometry in 6D space. Similarly, inlier correspondences in non-rigid scenes should also form a submanifold in 6D, because natural deformations are usually locally rigid. Therefore, the motivation is to learn the geometry of inliers in 6D space using transformers.

Fig. I shows the overview of our outlier rejection method. It takes as input an array of 6D coordinates, denoting the corresponding point pairs, and predicts the inlier probability of each correspondence.

We obtain the input correspondence using Leopard [21]. The network consists of a sequence of multi-head self-attention (SA) layers [39], with the head number set to 8. To leverage the spatial consistency of the correspondence, we modify the self-attention layer using the spatial-consistency-guided attention as proposed in PointDSC [2]. The output of the last SA layer is activated through the *Sigmoid* function that maps the output to range $[0, 1]$ as the per-correspondence confidence score. We use weighted binary cross-entropy (BCE) as the loss function for training. The network is trained, validated, and tested on the 4DMatch [21] split. When testing, we treat correspondence with a confidence score below the threshold $\theta = 0.3$ as outliers. Fig. II shows the influence of the threshold θ for non-rigid registration.

Number of SA layer. Tab. I shows the ablation study of number of self-attention (SA) layer. Precision grows with the number of the SA layer. The recall rate becomes stable after using 3 or more layers. By default, our method uses 9 SA layers.

Table I: Outlier rejection results on 4DLoMatch. Following [21], inlier threshold is set to 4cm.

		Number of SA layer							
	Input from Leopard[21]	1	2	3	4	5	6	7	8
Precision	55.7 (inlier rate)	63.2	64.7	64.3	74.6	74.3	74.9	76.3	77.1
Recall		58.9	77.0	86.0	80.2	80.2	83.4	80.3	80.1

Outlier rejection confidence. The precision and recall rate in Tab. I can not reflect the end-to-end performance for non-rigid registration. We want to find a proper outlier rejection threshold for non-rigid registration. Fig. II shows the ablation study of the confidence threshold for non-rigid registration accuracy. The optimal confidence threshold lies around 20%-30%. A threshold higher than 60% leads to worse results than not rejecting outliers.

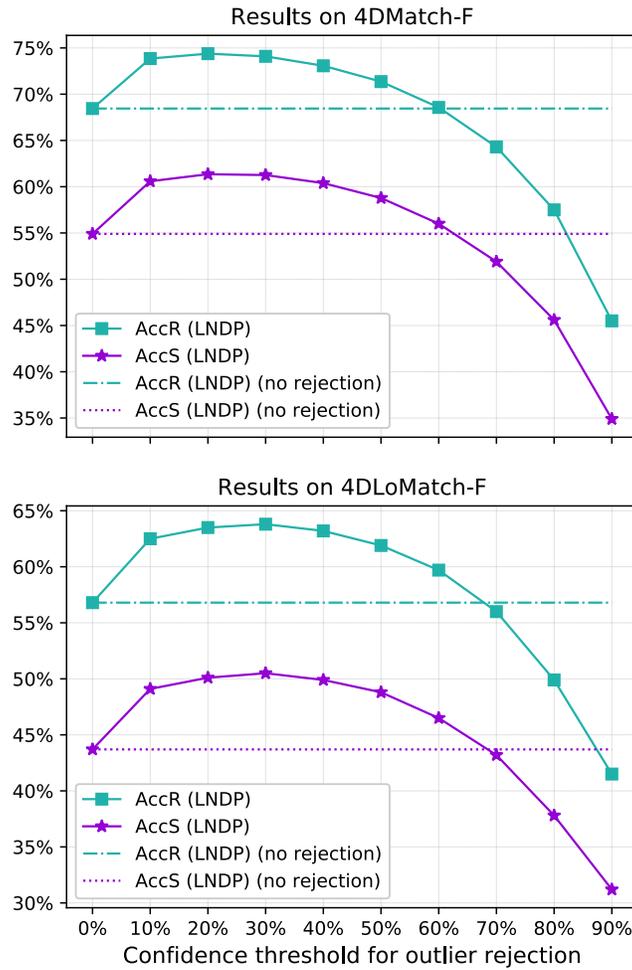


Figure II: Ablation study of outlier rejection threshold for non-rigid registration.

III. 4DMatch benchmark filtering.

We observe that the original 4DMatch/4DLoMatch benchmarks contain a certain amount of entries that are dominated by rigid motion, i.e. the best rigid fitting already registers most of the points. For fair benchmarking of non-rigid registration, we probabilistically remove point cloud pairs that have near-rigid movements. In the end, we removed around 50% of the entries. Fig. III shows the histogram of the benchmark before and after filtering.

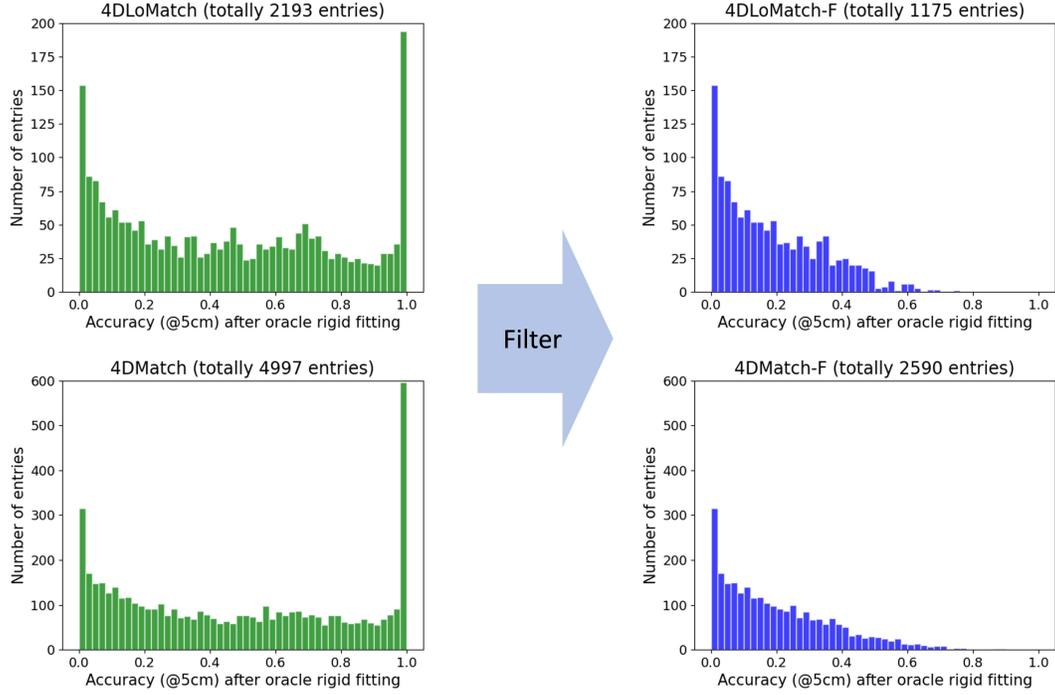


Figure III: Histogram of point cloud registration Accuracy (at 5cm) via the oracle rigid fitting. The left shows the pre-filtered 4DMatch 4DLoMatch benchmarks, the right shows the filtered benchmarks.