PointCLM: A Contrastive Learning-based Framework for Multi-instance Point Cloud Registration

Mingzhi Yuan^{1,3}, Zhihao Li^{1,3}, Qiuye Jin^{1,3}, Xinrong Chen^{1,2,3}, and Manning Wang^{1,3}, Magamma Barrow Ba

¹ Digital Medical Research Center, School of Basic Medical Sciences, Fudan University, Shanghai 200032, China

² Academy for Engineering and Technology, Fudan University, Shanghai 200032, China

³ Shanghai Key Laboratory of Medical Image Computing and Computer Assisted Intervention, Shanghai 200032, China

 $\{mzyuan20, lizhihao21, qyjin18, chenxinrong, mnwang\}$ @fudan.edu.cn

Abstract. Multi-instance point cloud registration is the problem of estimating multiple poses of source point cloud instances within a target point cloud. Solving this problem is challenging since inlier correspondences of one instance constitute outliers of all the other instances. Existing methods often rely on time-consuming hypothesis sampling or features leveraging spatial consistency, resulting in limited performance. In this paper, we propose PointCLM, a contrastive learning-based framework for multi-instance point cloud registration. We first utilize contrastive learning to learn well-distributed deep representations for the input putative correspondences. Then based on these representations, we propose a outlier pruning strategy and a clustering strategy to efficiently remove outliers and assign the remaining correspondences to correct instances. Our method outperforms the state-of-the-art methods on both synthetic and real datasets by a large margin. The code will be made publicly available at http://github.com/phdymz/PointCLM.

Keywords: Multi-instance point cloud registration, Multi-model fitting, Contrastive learning

1 Introduction

3D point cloud registration is a fundamental task in computer vision [21,9,50], and most studies mainly focus on pairwise registration. However, in real applications, target scene may contain multiple repeated instances, and we need to estimate multiple rigid transformations between a source point cloud and these repeated instances in the target point cloud. An example is illustrated in Figure 1. This problem is named as multi-instance point cloud registration and it is more challenging than pairwise point cloud registration.

There exist two solutions to the multi-instance point cloud registration problem. One solution is to use an instance detector [32,37] to detect instances in the



Fig. 1. Given a source point cloud of a 3d object, pairwise point cloud registration (left) focuses on estimating a single rigid transformation between the source point cloud and the target point cloud, while multi-instance point cloud registration (right) aims to estimate the 6d poses of the same objects within the target point cloud.

target point cloud, and turn this problem into multiple pairwise registrations. However, this approach can only detect known classes in the training set, and the registration performance is limited by the instance detector. Another solution is via multi-model fitting [42,29,23,31]. This approach starts from building putative correspondences based on local features, followed by estimating multiple transformations from noisy correspondences using multi-model fitting algorithms. Multi-model fitting problem has been studied for decades. The basic idea of most multi-model fitting methods is to sample a series of hypotheses and then perform preference analysis [15,16,54,48] or consensus analysis [43,29,30,35]. Most of these multi-model fitting algorithms are traditional methods, which rely on a large number of sampling for generating hypotheses and sophisticated strategies for selecting real models, resulting in large computational cost.

Recently, a series of deep learning-based methods have achieved excellent performance in pairwise point cloud registration, including feature matching [13,4,20,46], outlier correspondences rejection [52,12,3,24], etc. Multi-instance point cloud registration also faces the interference of outliers, which inspires researchers to extend deep learning-based methods to the multi-instance case. However, in the multi-instance case, the inliers of one instance constitute outliers of all the other instances, so we need not only to identify outliers, but also to predict which instance the inliers belong to. To the best of our knowledge, the only existing deep learning-based method [23] that can be used in multiinstance point cloud registration is by sequential binary classification. However, this method has limited performance because it does not fully explore the interaction between correspondences. There is also a traditional method [42] which clusters the correspondences based on spatial consistency [25]. However, due to the ambiguity of spatial consistency, this method cannot assign the correspondences very efficiently when outliers are located close to inliers. Therefore, an intuitive idea is to learn a more discriminative representation for correspondence so that not only the outliers can be easily pruned but also the inliers of different instances are separable in the feature space.

In this paper, we propose a contrastive learning-based method to learn deep representations for correspondences, based on which, we can not only remove outliers in the putative correspondences, but also correctly assign inliers to each instance. Specifically, we first select inlier correspondences belonging to the same instance to build positive pairs and select one correspondence from an instance and another correspondence not belonging to that instance to build negative pairs. Then we train a feature extractor with contrastive loss to make the correspondences in positive pairs near each other and the correspondences in negative pairs fall apart from each other in the feature space. This makes the inliers of each instance form a certain scale of clusters in the feature space, while the outliers are scattered. Based on this distribution, we prune the input putative correspondences according to the density over the feature space and obtain a correspondence set that has little outliers. After that, we perform a spectral clustering based on feature similarity and spatial consistency [25,3], in which the number of instances can be automatically determined and the remaining correspondences can be correctly assigned to each instance. Finally, we calculate the transformation for each instance using the correspondences assigned to it. Our main contributions are as follows:

- We propose a contrastive learning-based strategy that makes inlier correspondences and outlier correspondences well distributed in the feature space;
- We propose a pruning strategy based on feature similarity and spatial consistency to remove outliers and then utilize spectral clustering to assign the remaining correspondences to correct instances;
- Our method outperforms existing state-of-the-art methods in multi-instance point cloud registration on both synthetic and real datasets.

2 Related Work

Point cloud registration has long been a fundamental task in computer vision and robotic, which can be roughly divided into direct methods [9,39,36,1] and feature-based methods [4,20,46,17]. In recent years, thanks to the development of deep learning, many feature-based methods achieved state-of-the-art performance. These methods commonly produce correspondences by feature matching and then remove outliers to estimate transformations robustly. Despite the rapid development of deep features [13,4,20,46], the correspondences generated by feature matching still contain outliers. Therefore, removing outliers is of great significance in point cloud registration. In the past, many traditional methods have been proposed to remove outliers, including RANSAC-based methods [15,5,26], branch and bound-based methods [10], and many others [40,49]. A comprehensive review of these methods can be found in [21,50,51]. Recently, a series of learning-based methods [3,24,52,12] have been proposed and achieved remarkable results in outlier removal. For example, Yi et al. used a PointNet-style [38] network with instance normalization to predict outliers [52], which has been widely used as a backbone network for predicting outliers [53,41]. Choy et al. used a sparse convolution-based network to classify putative correspondences into inliers and outliers [12]. Based on the assumption of spatial compatibility between inliers, Bai et al. incorporated geometric prior into deep neural network and designed a non-local layer [3] to better aggregate features, which achieved outstanding results. The above methods are all designed for pairwise registration. However, unlike pairwise registration, inliers of one instance constitute outliers of all the other instances in multi-instance point cloud registration. Such pseudo outliers make it difficult to directly generalize the above binary classification models to the case of multi-instance point cloud registration. Inspired by success of previous works in pairwise registration, we introduce deep learning into multi-instance point cloud registration and propose a method that can not only remove outliers but also assign inlier correspondences to each instance.

Multi-model fitting aims to fit multiple models from noisy data, such as fitting multiple planes [6] in a point cloud, estimating fundamental matrices in motion segmentation [18], calculating rigid transformations in multi-instance point cloud registration [42], etc. Since inliers of one instance constitute outliers of all the other instances, multi-model fitting is more challenging than single-model fitting. Existing multi-model fitting methods can be roughly divided into two categories. The first category fits models sequentially [22,7,8,23], which relies on repeatedly sampling and selecting models. For example, sequential RANSAC [22] detects instances in a sequential manner by repeatedly running RANSAC to recover a single instance and then removing its inliers from the input. Progressive-X and Progressive-X+ [7,8] use a better performing Graph-cut RANSAC [5] as a sampler to generate hypotheses. CONSAC [23] introduced deep models into multi-model fitting for the first time, using a network similar to PointNet [38] to guide sampling. The second category fits multiple models simultaneously [42,43,29,30,31]. For example, many preference analysis-based methods [43,29,30] initially sample a series of hypotheses and then cluster input points according to the residuals of the hypotheses. RansaCov [31] formulates the multi-model fitting as a maximum coverage problem and provides two strategies to solve it approximately. ECC [42] utilizes the spatial consistency [25] of point cloud rigid transformation and clusters correspondences in a bottom-up manner based on a distance-invariant matrix. Although spatial consistency performs efficiently in [42], the lack of orientation constraints makes the distance-invariant matrix still ambiguous in some cases, especially when outliers are close to inliers. In this paper, a novel deep representation is integrated with the spatial consistency to achieve better results.

3 Problem Formulation

We use X and Y to denote the source and target point clouds, respectively. The source point cloud consists of an instance of a 3D model, and the target point cloud contains M instances of the same model, where these instances may be sampled from a part of the 3D model. By matching local features [13,4,20,46], we can generate putative correspondences between the two point clouds. A correspondence is denoted $c_i = (x_i, y_i) \in \mathbb{R}^6$, where $x_i \in X$, $y_i \in Y$ are the coordinates of a pair of 3D keypoints from the two point clouds. Our objective is to divide the putative correspondence set $C = \{c_i\}_{i=1}^N$ into M + 1 subsets C_o, C_1, \ldots, C_M satisfying $C = C_o \cup C_1 \cup \ldots \cup C_M$, where C_o denotes the predicted outlier set and C_m denotes the inlier set for the m-th predicted instance.

When we know the true instance number is M, recovering M rigid transformations $\{R_m, t_m\}_{m=1}^M$ from the two point clouds is to minimize the objective function:

$$\min_{\{R_m, t_m\}_{m=1}^M} \frac{1}{M} \sum_{m=1}^M \sum_{(x_{mi}, y_{mi}) \in C_m^{gt}} \frac{1}{|C_m^{gt}|} \|y_{mi} - R_m x_{mi} - t_m\|^2$$
(1)

where C_m^{gt} denotes the ground truth inlier set of the *m*-th instance, and $|C_m^{gt}|$ denotes the number of inliers in C_m^{gt} . The above problem is very challenging, because the inliers in C_i^{gt} constitute outliers in C_j^{gt} for $i \neq j$. In practice, the problem becomes even more difficult because the true instance number in the target point cloud is often unknown in prior, which is the case that we deal with in this paper.



Fig. 2. The pipeline of the proposed multi-instance point cloud registration framework PointCLM. It takes putative correspondence as input, and output M rigid transformations. The green lines and red lines represent inliers and outliers, respectively. After clustering, the correspondences in different clusters are visualized in different colors. The green bounding boxes in output transformations represent the ground truth poses of instances in the target point cloud and the red bounding boxes represent our predictions. The transformed point clouds in the target point cloud are visualized in blue.

4 Method

In this section, we present our framework for multi-instance point cloud registration, which is illustrated in Figure 2. Our framework takes the putative correspondences generated by feature matching as input, and it first uses the Feature

Extractor trained by contrastive learning to extract deep representations for the input correspondences (section 4.1). Then, we prune the correspondences according to both spatial consistency and the similarity of their deep representations (section 4.2). Finally, we cluster the remaining correspondences and estimate the transformations for multiple instances using the clustering results (section 4.3).

4.1 Feature Extractor

The first stage of our framework embeds the N input putative correspondences $C = \{c_i\}_{i=1}^N$ into a feature space to obtain well-distributed d-dimension representations $F = \{f_i \in \mathbb{R}^d\}_{i=1}^N$ for the following pruning and clustering. Here, we adopt the SCNonlocal module in [3] as our feature extractor, which consists of 12 repetitive blocks. As shown in Figure 2(b), each block consists of a shared Perceptron layer, a BatchNorm layer with ReLU and a nonlocal layer. The nonlocal layer integrates the spatial consistency [25] of rigid transformation. Before calculating the features of each layer, a spatial consistency matrix β is first calculated:

$$\beta_{ij} = \left[1 - \frac{d_{ij}^2}{\sigma_d^2}\right]_+, d_{ij} = |||x_i - x_j|| - ||y_i - y_j|||$$
(2)

where σ_d is a distance parameter to control the sensitivity to length difference, [.]₊ denotes a clamp function max(x, 0) to make $\beta_{ij} \geq 0$. The nonlocal layers aggregate the intermediate features using β :

$$f_i^{k+1} = f_i^k + MLP\left(\sum_{j=1}^{|C|} \operatorname{soft} \max_j(\alpha\beta)g\left(f_j^k\right)\right)$$
(3)

where α denotes the embed dot-product similarity between the intermediate feature representations f_i^k and f_j^k in the k-th blocks, and $g(\cdot)$ denotes a linear projection function. More details about the network can be found in [3].

We utilize contrastive learning to train our feature extractor to obtain welldistributed deep representations $F = \{f_i \in \mathbb{R}^d\}_{i=1}^N$. Concretely, for an anchor correspondence $c_{mi} \in C_m^{gt}$, we define the other correspondences in C_m^{gt} as its positive samples, and define the correspondences in $C \setminus C_m^{gt}$ as its negative samples. We define the negative sample with the smallest Euclidean distance to the anchor correspondence in the feature space as the hardest negative sample. During each iteration, we exhaust all inliers as anchors and select their hardest negative samples to build hardest negative pair set \mathcal{N} and randomly select positive samples to build positive pair set \mathcal{P} . Our contrastive loss is formulated as:

$$\mathcal{L} = \sum_{(i,j)\in\mathcal{P}} \left[D\left(f_i, f_j\right) - m_p \right]_+^2 / |\mathcal{P}| + \sum_{(i,j)\in\mathcal{N}} \left[m_n - D\left(f_i, f_j\right) \right]_+^2 / |\mathcal{N}|$$
(4)

where m_p and m_n are margins for positive and negative pairs, which prevent the network from overfitting [28]. Since the inlier rate of input correspondences is commonly not high and we can downsample the input correspondences to a suitable scale, our exhaustion over all negative samples to find the hardest one for each minibatch is feasible. Our ablation experiments show that utilizing hardest negative pairs plays an important role in the high performance of the proposed method.

By optimizing the loss in equation 4, the representations of correspondences belonging to different instances are easily separable in the feature space, while the representations of correspondences belonging to the same instance are clustered together. We then use these well-distributed representations for both pruning and clustering. Our experiments show that the deep representations provide information that spatial consistency does not have, which makes pruning and clustering more effective.

4.2 Pruning

The input correspondences tend to contain a large proportion of outliers, which severely devastates the following inlier correspondence clustering and transformation estimation. Therefore, an intuitive idea is to first prune the input correspondences to remove outliers. A series of pruning strategies have been proposed in the past. For example, the input can be pruned based on similarity threshold [19] or confidence [53]. In this paper, we propose a density-based pruning strategy which prunes the input correspondences according to the deep representation learned in Section 4.1 along with spatial consistency constraints.

As introduced in Section 4.1, through contrastive learning, correspondences satisfying the same rigid transformation have similar representations in the feature space and they are far away from other correspondences. This means that the feature density around inliers are higher than that around outliers [44]. Based on this idea, we design a density-based pruning strategy as shown in Figure 2(c). We treat the isolated points in red as outliers and remove it, while the clustered points in other colors are reserved. The detail of our pruning method is as follows.

We first calculate the feature similarity matrix S^F between correspondences using the extracted deep representations $F = \{f_i \in \mathbb{R}^d\}_{i=1}^N$:

$$S_{ij}^F = \langle f_i, f_j \rangle \tag{5}$$

where $\langle \cdot \rangle$ represents dot product. After that, we use the feature similarity matrix S^F and the spatial consistency matrix β to calculate the similarity matrix S:

$$S = S^F \otimes \beta \tag{6}$$

where \otimes represents element-wise product, and the spatial consistency matrix β has been calculated before.

Second, we set a threshold τ_S and use it to binarize the similarity matrix S to a binary similarity matrix \hat{S} . When the *i*-th correspondence and the *j*-th correspondence satisfy sufficient spatial consistency and similarity in the feature space at the same time, $\hat{S}_{ij} = 1$, otherwise $\hat{S}_{ij} = 0$. Now, we can treat the input

correspondences as the nodes of a graph, and \hat{S} as the graph's adjacency matrix. The inlier sets can be regarded as subgraphs with a certain scale while outliers appear as isolated points or small-scale subgraphs.

Finally, we sum the rows of the binary similarity matrix \hat{S} , and select the correspondences with row-sum values larger than a threshold τ_N . As illustrated in Figure 2(a), our pruning method can effectively remove most or even all outliers. Our experiments also demonstrate that the pruning step is crucial for our method. We also provide details about how to select the thresholds τ_S and τ_N and ablation experiments in Supplementary Material.

4.3 Clustering and Transformations Estimation

After pruning, we obtain a clean set of correspondences, which is almost free of outliers. The next step is to divide these correspondences into multiple subsets belonging to different instances and estimate the final rigid transformations for all the instances. The correspondence division can be considered as a clustering problem, and the number of instances should be equal to the number of clusters. Here, we use a spectral clustering algorithm [45,27], which can determine the number of clusters automatically. The algorithm consists of the following steps:

Step 1: Recompute the binary similarity matrix \hat{S}_n using the feature similarity and spatial consistency of the remaining correspondences, where *n* represents the number of correspondences after pruning;

Step 2: Calculate the normalized Laplacian matrix \hat{L}_n of the matrix \hat{S}_n , and calculate the eigenvalues $\lambda_1(\hat{L}_n) \leq \lambda_2(\hat{L}_n) \leq \ldots \leq \lambda_n(\hat{L}_n)$ of the normalized Laplacian matrix;

Step 3: Determine the number of clusters M by the following formula:

$$M = \arg\max_{k} \left\{ \lambda_{k+1} \left(\hat{L}_n \right) - \lambda_k \left(\hat{L}_n \right) \right\}$$
(7)

Step 4: Apply spectral clustering [45] with M clusters.

Since the deep representations have a good distribution in the feature space, the binary similarity matrix \hat{S}_n conforms to the ideal matrix defined in [27]. Therefore, the number of clusters determined by equation 7 is reliable, which is proved in [27].

After the above steps, the remaining correspondences are assigned to different clusters as shown in Figure 2(a), and then we can use a solver such as RANSAC [15] to estimate the rigid transformation of each instance. Because our pruning and clustering strategy are very efficient, the inlier rate of each instance is high, and only a few dozen RANSAC iterations are enough to achieve outstanding performance.

5 Experiment

We conduct experiments on both synthetic and real datasets and compare our PointCLM to state-of-the-art methods [42,29,23,31]. The following sections are

organized as follows. First, we illustrate our experimental settings including our implementation, datasets, competitors and evaluation metrics in section 5.1. Next, we conduct experiments on the synthetic and the real datasets in section 5.2 and 5.3, respectively. We further conduct comprehensive ablation studies in section 5.4 to illustrate the efficiency of our PointCLM and the importance of each component. More details are provided in the supplementary material, including the construction of the datasets, more qualitative evaluation and some experiments for hyperparameter choices.

5.1 Experimental Settings

Implementation: We implement our network in Pytorch [33] and implement spectral clustering with sklearn [34]. We randomly downsample the input correspondences to 1000. The distance parameter σ_d is set to 0.05 for the synthetic dataset and 0.1 for the real dataset. We set the dimension d of the deep representation to 128. Threshold τ_S is set to 0.85. Threshold τ_N is set to 10 for the synthetic dataset and 20 for the real dataset. The margins m_p and m_n are set to 0.1 and 1.4, respectively. Our batchsize is set to 16. We optimize the network using the ADAM optimizer with an initial learning rate of 0.01 and train the network for 15K iterations. All the experiments are conducted on a single RTX1080Ti graphic card with Intel Core i7-7800X CPU.

Datasets: We conduct experiments on both synthetic and real datasets. Our synthetic dataset is constructed from ModelNet40 [47], which consists of 12311 meshed CAD models from 40 categories. To construct our synthetic dataset, for each model, we uniformly downsample 1024 points from it to form the source point cloud, and then rotate and translate it 5-10 times repeatedly to generate multiple instances. The instances are mixed with noise points to form the target point cloud as shown in Figure 3. The rotation along each axis is uniformly sampled in [0,180°] and the translation is in [0,5]. Our input correspondences for synthetic dataset are randomly generated by mixing the ground truth correspondences with outliers. We control the inlier ratio per instance at approximate 2%. We generate 12311 such synthetic source-target point cloud pairs using 12311 models. We use 9843 pairs for training and 2468 pairs for testing. We randomly set aside 10% pairs in the training set for validation.

Our real dataset is Scan2CAD [2], which is constructed using ShapeNet [11] and ScanNet [14]. This dataset uses the CAD models in ShapeNet to replace the point clouds in the real scanned scene and provides accurate annotations, including models' categories, rotations and translations, etc. The dataset provides 1506 annotated scenes, and each scene contains at least one class of instances. Therefore, we make full use of these annotations and split the scenes containing multiple kinds of instances into multiple source-target point cloud pairs for multi-instance registration. In this way, we get 2184 pairs of point clouds, most of which contain 2-5 instances of the same class in the target point clouds. We divide the samples into training set, validation set and test set according to the ratio of 7:1:2. We use fine-tuned FCGF [13] to produce local features and

generate putative correspondences by feature matching. More details about our dataset construction are provided in the supplementary material.

Competitors: We compare our PointCLM with four state-of-the-art methods, including T-linkage [29], RansaCov [31], CONSAC [23] and ECC [42]. T-linkage is a typical algorithm based on preference analysis, which samples a series of hypotheses and clusters the inputs based on the residuals of the hypotheses. RansaCov regards multi-model fitting as a maximum coverage problem, which can be approximately solved in greedy strategy or using relaxed linear programming. Here we use the former strategy due to its effectiveness. CONSAC uses a deep network to guide the sampling process, and we train this network using the same training set as our network. ECC clusters the input correspondences based on spatial consistency and doesn't need hypothesis sampling as above methods. For a fair comparison, we not only use the same input, but also fine-tune the above methods both on GPU and CPU, and choose the ones with the best performance for comparison.

Evaluation metrics: We first define rotation error RE and translation error TE as follow:

$$RE(R) = \arccos\left(\frac{\operatorname{Tr}(R^{\mathrm{T}}R^{*}) - 1}{2}\right), TE(t) = \|t - t^{*}\|_{2}$$
(8)

where R^* and t^* are the ground truth rotation and translation. We consider the instances with both rotation error and translation error below the thresholds to be successfully registered and our thresholds for RE and TE are 15° and 0.1, respectively. Since both our method and the above comparing methods predict multiple rigid transformations, we use mean recall (MR), mean precision (MP) and mean F1 Score (MF) as evaluation metric. For a pair of source point cloud and target point cloud, we define instance recall as $\frac{n_{success}^{pred}}{M^{grt}}$ and instance precision as $\frac{n_{success}^{pred}}{M^{pred}}$, where $n_{success}^{pred}$ denotes the number of successfully registered instances in prediction, M^{gt} denotes the ground truth number of instances, and M^{pred} denotes the number of predicted transformations. The instance F1 Score is the harmonic mean of the instance recall, and instance F1 Score of each sample in test set and average them to obtain our final evaluation metrics, the mean recall (MR), the mean precision (MP), and the mean F1 Score (MF).

5.2 Experiment on Synthetic Dataset

We first compare our method with other competitors on the synthetic dataset and the results are show in Table 1. The sampling-based methods such as T-Linkage, RansaCov, and CONSAC do not achieve good performance due to extremely low inlier ratio. Benefiting from spatial consistency, ECC performs effectively. Nevertheless, our PointCLM surpasses the second best method ECC by a large margin in all evaluation metrics.

We provide a set of visualizations to qualitatively evaluate our PointCLM and compare it with other competitors in Figure 3. The first row of Figure 3

PointCLM 11

	ModelNet40 (synthetic dataset)				Scan2CAD (real dataset)			
	MR(%)	$\mathrm{MP}(\%)$	MF(%)	$\operatorname{Time}(s)$	MR(%)	MP(%)	MF(%)	$\operatorname{Time}(s)$
T-linkage [29]	0.61	1.48	0.87	3.89	34.99	46.86	40.07	6.64
RansaCov [31]	0.73	5.33	1.29	0.14	60.50	33.28	42.94	0.07
CONSAC [23]	1.00	7.45	1.77	0.61	55.48	53.34	54.39	0.39
ECC [42]	82.90	92.92	87.63	3.56	64.66	69.73	67.10	1.84
PointCLM	92.60	99.69	96.01	0.06	78.10	70.64	74.18	0.10

Table 1. Registration results on synthetic dataset (left) and real dataset (right).

shows the input correspondences, and our pruning and clustering results. Figure 3(b) shows that our PointCLM surprisingly removes all outliers, and the remaining correspondences are well clustered as shown in Fig. 3 (c). The second row of Figure 3 shows the registration results of our proposed method and the competitors. It can be seen that both T-Linkage and RansaCov fail to register any instances. For the six instances in the target point cloud, CONSAC only registers one instance successfully. It is worth noting that although ECC register four instances successfully, but it fails to register the two tables in the lower right corner. This is because the two tables are mixed together and spatial consistency is not enough to distinguish them. However, our method successfully registers all instances.



Fig. 3. Results on synthetic dataset. In (a) and (b), the green lines and red lines represent inlier correspondences and outlier correspondences, respectively. In (c), the correspondences in each clusters after clustering are visualized in different colors. In (d-f), the green bounding boxes represent the ground truth poses of instances in the target point cloud and the red bounding boxes represent predictions. The transformed point clouds in the target point cloud are visualized in blue.

5.3 Experiment on Real Dataset

We then compare our PointCLM with other competitors on Scan2CAD [2]. As shown in Table 1, our PointCLM outperforms all the competitors on all three evaluation metrics, MR, MP, and MF and it is also competitive in speed. The performance of ECC and PointCLM is lower than that on the synthetic dataset while the performance of the other methods is higher than that on the synthetic dataset. This is due to the change of the distribution of the instances' inlier ratio.

We also provide a set of visualizations to qualitatively evaluate our Point-CLM and compare it with the other competitors. The first row of Figure 4 shows the input correspondences, and our pruning and clustering results. Figure 4(b) shows that our method removes almost all outliers, ensuring that the following clustering performs efficiently as shown in Figure 4(c). For the five instances contained in the target point cloud, T-Linkage and CONSAC successfully register two instances, but one prediction of T-Linkage has large errors. RansaCov successfully registers three instances. Since three chairs in the target point cloud are close to each other, ECC does not successfully register all these instances. Our method successfully accomplishes the registration of all instances.



Fig. 4. Results on real dataset. In (a) and (b), the green lines and red lines represent inlier correspondences and outlier correspondences, respectively. In (c), the correspondences in each clusters after clustering are visualized in different colors. In (d-f), the green bounding boxes represent the ground truth poses of instances in the target point cloud and the red bounding boxes represent predictions. The transformed point clouds in the target point cloud are visualized in blue.

5.4 Ablation Studies

In this section we provide comprehensive ablation studies to illustrate the effectiveness of each component. All our ablation studies are performed on Scan2CAD. **Ablation on deep representation:** To study the effectiveness of our adopted deep representations, we compare the performance of our framework with and

Deep representation	Pruning	MR(%)	$\mathrm{MP}(\%)$	$\mathrm{MF}(\%)$	$\operatorname{Time}(s)$
	\checkmark	76.61	65.05	70.36	0.17
\checkmark		62.23	32.77	42.93	1.24
\checkmark	\checkmark	78.10	70.64	74.18	0.10

Table 2. Ablation study results on deep representation and pruning.

without deep representation. The version without deep representation deletes Feature Extractor and set $S = \beta$, which relies only on spatial consistency for pruning and clustering. The comparison results are shown in Table 2. It can be seen that both the accuracy and the speed metrics are improved by using the deep representations. Without deep representations, less correspondences are pruned and more correspondences need to be clustered, which increase the runtime of our method. Although the framework without deep representations has lower performance in MR, MP and MF, it is still a competitive baseline, which is suitable for the case without training data. This indicates the effectiveness of the pruning and clustering strategies proposed in this paper.

Ablation on pruning: To quantitatively study the effectiveness of our pruning strategy, we simply remove the pruning strategy in our method, and then compare the performance before and after removal. Table. 2 shows that the performance of our method drops sharply if the pruning step is removed, which is due to the fact that the noisy binary similarity matrix does not conform to the ideal matrix defined in [27] and spectral clustering cannot group the correspondences correctly. This result reveals that pruning is crucial for subsequent clustering.

Ablation on RANSAC iterations: In addition, we also test the performance of our model with different numbers of RANSAC iterations and the results are shown in Table 3. It can be seen that fairly good results can be achieved with only five iterations, and the performances with 50 and 500 iterations are very close. These results indicate that the pruning and clustering steps greatly improve the inlier ratio of each instance, so reliable results can be estimated using only a small number of iterations.

RANSAC iterations	MR(%)	MP(%)	MF(%)
5	77.39	67.76	72.26
50	78.10	70.64	74.18
500	79.33	70.64	74.73

Table 3. Influence of the number of RANSAC iterations.

Ablation on hardest negative pairs: Exhausting the negative pairs to find the hardest ones for each minibatch incurs larger computational overhead, but our ablation shows it worthwhile. We compare it with the same model trained

using randomly selected negative pairs. The results of using the hardest negative pairs as in the proposed method and using randomly selected negative pairs are shown in Table 4. All accuracy metrics are improved by using hardest negative pairs in contrastive learning. By comparing Table 4 and Table 2, we can find that using random negative pairs in contrastive representative learning achieves better results than not using the deep representation, but the margin is small. In addition, we collect the average cosine similarity within positive pairs and top-K hardest negative pairs in the feature space in Table 4. The result shows the representation trained using hardest negative pairs are better separated in the feature space, which is more discriminative.

Table 4. Ablation experiment results on training with hardest negative pairs.

	$\mathrm{MR}(\%)$	$\mathrm{MP}(\%)$	$\mathrm{MF}(\%)$	$\operatorname{Positive}(\%)$	$\operatorname{Top-1}(\%)$	$\operatorname{Top-10}(\%)$
Random	76.85	67.50	71.87	95.43	96.91	91.37
Hardest	78.10	70.64	74.18	83.96	61.32	49.65

6 Conclusion

In this paper, we propose a novel framework to address the multi-instance point cloud registration problem. We use contrastive learning to learn well-distributed deep representations for input correspondences, based on which we develop a pruning and a clustering strategy to remove outlier correspondences efficiently and assign the remaining inlier correspondences to correct instances. Then, the transformation from the source point cloud to each instance can be easily estimated. Extensive experiments on both synthetic and real datasets demonstrate the effectiveness of our framework and its superiority over existing solutions. We think the proposed representation learning and the outlier pruning strategy has the potential to be used in pairwise point cloud registration.

Acknowledgements This work was supported by the National Natural Science Foundation of China under Grant 62076070.

References

- Aoki, Y., Goforth, H., Srivatsan, R.A., Lucey, S.: Pointnetlk: Robust & efficient point cloud registration using pointnet. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7163–7172 (2019)
- Avetisyan, A., Dahnert, M., Dai, A., Savva, M., Chang, A.X., Nießner, M.: Scan2cad: Learning cad model alignment in rgb-d scans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2614– 2623 (2019)

- Bai, X., Luo, Z., Zhou, L., Chen, H., Li, L., Hu, Z., Fu, H., Tai, C.L.: Pointdsc: Robust point cloud registration using deep spatial consistency. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15859–15869 (2021)
- Bai, X., Luo, Z., Zhou, L., Fu, H., Quan, L., Tai, C.L.: D3feat: Joint learning of dense detection and description of 3d local features. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6359– 6367 (2020)
- Barath, D., Matas, J.: Graph-cut ransac. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6733–6741 (2018)
- Barath, D., Matas, J.: Multi-class model fitting by energy minimization and modeseeking. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 221–236 (2018)
- Barath, D., Matas, J.: Progressive-x: Efficient, anytime, multi-model fitting algorithm. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3780–3788 (2019)
- 8. Barath, D., Rozumny, D., Eichhardt, I., Hajder, L., Matas, J.: Progressive-x+: Clustering in the consensus space. arXiv preprint arXiv:2103.13875 (2021)
- Besl, P.J., McKay, N.D.: Method for registration of 3-d shapes. In: Sensor fusion IV: control paradigms and data structures. vol. 1611, pp. 586–606. Spie (1992)
- Bustos, A.P., Chin, T.J.: Guaranteed outlier removal for point cloud registration with correspondences. IEEE transactions on pattern analysis and machine intelligence 40(12), 2868–2882 (2017)
- Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
- Choy, C., Dong, W., Koltun, V.: Deep global registration. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2514–2523 (2020)
- Choy, C., Park, J., Koltun, V.: Fully convolutional geometric features. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8958– 8966 (2019)
- Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5828–5839 (2017)
- Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Communications of the ACM 24(6), 381–395 (1981)
- Fouhey, D.F., Scharstein, A.D.: Multi-model estimation in the presence of outliers. Bachelorsthesis, Middlebury College, Middlebury (2011)
- Fu, K., Liu, S., Luo, X., Wang, M.: Robust point cloud registration framework based on deep graph matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8893–8902 (2021)
- Hartley, R.I.: In defense of the eight-point algorithm. IEEE Transactions on pattern analysis and machine intelligence 19(6), 580–593 (1997)
- Heckel, R., Bölcskei, H.: Subspace clustering via thresholding and spectral clustering. In: 2013 IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 3263–3267. IEEE (2013)
- Huang, S., Gojcic, Z., Usvyatsov, M., Wieser, A., Schindler, K.: Predator: Registration of 3d point clouds with low overlap. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4267–4276 (2021)

- 16 M. Yuan al.
- Huang, X., Mei, G., Zhang, J., Abbas, R.: A comprehensive survey on point cloud registration. arXiv preprint arXiv:2103.02690 (2021)
- 22. Kanazawa, Y., Kawakami, H.: Detection of planar regions with uncalibrated stereo using distributions of feature points. In: BMVC. pp. 1–10. Citeseer (2004)
- Kluger, F., Brachmann, E., Ackermann, H., Rother, C., Yang, M.Y., Rosenhahn, B.: Consac: Robust multi-model fitting by conditional sample consensus. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4634–4643 (2020)
- Lee, J., Kim, S., Cho, M., Park, J.: Deep hough voting for robust global registration. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15994–16003 (2021)
- 25. Leordeanu, M., Hebert, M.: A spectral technique for correspondence problems using pairwise constraints (2005)
- Li, J., Hu, Q., Ai, M.: Gesac: Robust graph enhanced sample consensus for point cloud registration. ISPRS Journal of Photogrammetry and Remote Sensing 167, 363–374 (2020)
- Li, Z., Liu, J., Chen, S., Tang, X.: Noise robust spectral clustering. In: 2007 IEEE 11th International Conference on Computer Vision. pp. 1–8. IEEE (2007)
- Lin, J., Morere, O., Chandrasekhar, V., Veillard, A., Goh, H.: Deephash: Getting regularization, depth and fine-tuning right. arXiv preprint arXiv:1501.04711 (2015)
- Magri, L., Fusiello, A.: T-linkage: A continuous relaxation of j-linkage for multimodel fitting. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3954–3961 (2014)
- Magri, L., Fusiello, A.: Robust multiple model fitting with preference analysis and low-rank approximation. In: BMVC. vol. 20, p. 12 (2015)
- Magri, L., Fusiello, A.: Multiple model fitting as a set coverage problem. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3318–3326 (2016)
- Misra, I., Girdhar, R., Joulin, A.: An end-to-end transformer model for 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2906–2917 (2021)
- 33. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
- 34. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al.: Scikit-learn: Machine learning in python. the Journal of machine Learning research 12, 2825–2830 (2011)
- Pham, T.T., Chin, T.J., Yu, J., Suter, D.: The random cluster model for robust geometric fitting. IEEE transactions on pattern analysis and machine intelligence 36(8), 1658–1671 (2014)
- Pomerleau, F., Colas, F., Siegwart, R.: A review of point cloud registration algorithms for mobile robotics. Foundations and Trends in Robotics 4(1), 1–104 (2015)
- Qi, C.R., Litany, O., He, K., Guibas, L.J.: Deep hough voting for 3d object detection in point clouds. In: proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9277–9286 (2019)
- Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 652–660 (2017)
- Rusinkiewicz, S., Levoy, M.: Efficient variants of the icp algorithm. In: Proceedings third international conference on 3-D digital imaging and modeling. pp. 145–152. IEEE (2001)

- Stechschulte, J., Ahmed, N., Heckman, C.: Robust low-overlap 3-d point cloud registration for outlier rejection. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 7143–7149. IEEE (2019)
- Sun, W., Jiang, W., Trulls, E., Tagliasacchi, A., Yi, K.M.: Acne: Attentive context normalization for robust permutation-equivariant learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11286– 11295 (2020)
- 42. Tang, W., Zou, D.: Multi-instance point cloud registration by efficient correspondence clustering. arXiv preprint arXiv:2111.14582 (2021)
- 43. Toldo, R., Fusiello, A.: Robust multiple structures estimation with j-linkage. In: European conference on computer vision. pp. 537–547. Springer (2008)
- Torr, P.H., Nasuto, S.J., Bishop, J.M.: Napsac: High noise, high dimensional robust estimation-it's in the bag. In: British Machine Vision Conference (BMVC). vol. 2, p. 3 (2002)
- Von Luxburg, U.: A tutorial on spectral clustering. Statistics and computing 17(4), 395–416 (2007)
- Wang, H., Liu, Y., Dong, Z., Wang, W., Yang, B.: You only hypothesize once: Point cloud registration with rotation-equivariant descriptors. arXiv preprint arXiv:2109.00182 (2021)
- 47. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1912–1920 (2015)
- Xu, L., Oja, E., Kultanen, P.: A new curve detection method: randomized hough transform (rht). Pattern recognition letters 11(5), 331–338 (1990)
- Yang, H., Shi, J., Carlone, L.: Teaser: Fast and certifiable point cloud registration. IEEE Transactions on Robotics 37(2), 314–333 (2020)
- Yang, J., Xian, K., Wang, P., Zhang, Y.: A performance evaluation of correspondence grouping methods for 3d rigid data matching. IEEE transactions on pattern analysis and machine intelligence 43(6), 1859–1874 (2019)
- Yang, J., Xian, K., Xiao, Y., Cao, Z.: Performance evaluation of 3d correspondence grouping algorithms. In: 2017 International Conference on 3D Vision (3DV). pp. 467–476. IEEE (2017)
- Yi, K.M., Trulls, E., Ono, Y., Lepetit, V., Salzmann, M., Fua, P.: Learning to find good correspondences. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2666–2674 (2018)
- Zhao, C., Ge, Y., Zhu, F., Zhao, R., Li, H., Salzmann, M.: Progressive correspondence pruning by consensus learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6464–6473 (2021)
- Zuliani, M., Kenney, C.S., Manjunath, B.: The multiransac algorithm and its application to detect planar homographies. In: IEEE International Conference on Image Processing 2005. vol. 3, pp. III–153. IEEE (2005)