

Progressive Correspondence Regenerator for Robust 3D Registration

Guiyu Zhao^{1,2,4*}, Sheng Ao^{1,2,3*}, Ye Zhang^{3†}, Kai Xu⁵, Yulan Guo³

¹Fujian Key Laboratory of Sensing and Computing for Smart Cities, Xiamen University

²Key Laboratory of Multimedia Trusted Perception and Efficient Computing,
Ministry of Education of China, Xiamen University, China ³Sun Yat-sen University

⁴Beijing Institute of Technology ⁵National University of Defense Technology

Abstract

Obtaining enough high-quality correspondences is crucial for robust registration. Existing correspondence refinement methods mostly follow the paradigm of outlier removal, which either fails to correctly identify the accurate correspondences under extreme outlier ratios, or select too few correct correspondences to support robust registration. To address this challenge, we propose a novel approach named Regor, which is a progressive correspondence regenerator that generates higher-quality matches whilst sufficiently robust for numerous outliers. In each iteration, we first apply prior-guided local grouping and generalized mutual matching to generate the local region correspondences. A powerful center-aware three-point consistency is then presented to achieve local correspondence correction, instead of removal. Further, we employ global correspondence refinement to obtain accurate correspondences from a global perspective. Through progressive iterations, this process yields a large number of high-quality correspondences. Extensive experiments on both indoor and outdoor datasets demonstrate that the proposed Regor significantly outperforms existing outlier removal techniques. More critically, our approach obtain 10 times more correct correspondences than outlier removal methods. As a result, our method is able to achieve robust registration even with weak features. The code is available at [Regor].

1. Introduction

Point cloud registration [8, 21] is a fundamental and critical task in 3D computer vision, with widespread applications across various domains such as robotic localization [36, 42] and remote sensing [21, 35]. Point cloud registration aims to estimate a rigid transformation to align two point clouds scanned from different perspectives. With the advancement




Figure 1. Difference between our method and existing outlier removal techniques. Outlier removal methods perform top-down pruning, producing an optimized set of correspondences that is a subset of the initial correspondences. When the initial inliers are few, this approach can obtain only a limited number of inliers, which is insufficient for robust pose estimation. In contrast, our method takes a bottom-up approach, using a regeneration strategy to generate more inliers, thereby effectively addressing this issue.

of feature metric learning, feature-based point cloud registration methods [1, 20, 30, 43] have garnered significant attention and achieved promising results. However, in scenarios characterized by low overlap and high noise, direct feature matching often leads to a large number of outliers, which poses a great challenge for point cloud registration.

Currently, the mainstream approach to address the issue of numerous outliers in point cloud registration is to filter wrong correspondences. Existing outlier removal methods [3, 6, 22, 27, 47] can be categorized into two types: geometry-based methods and learning-based methods. Geometry-based methods [6, 14, 28, 47] typically remove outliers according to whether point correspondences satisfy geometric consistency. Learning-based methods [3,

*These authors contributed equally.

†Corresponding Author.

[22, 27] treat this task as a classification problem to distinguish between outliers and inliers. However, in challenging scenarios, these correspondence filtering methods all suffer from a common problem: when initial correspondences only contains few inliers, estimating the correct rigid transformation between two point clouds is extremely difficult even if all inliers are identified.

In this paper, we aim to design a new 3D registration architecture, which is robust to few inliers and is able to estimate precise 6DoF pose. As shown in Figure 1, different from existing outlier removal methods [3, 6, 7, 47] that focus on pruning correspondences, our approach clearly regenerate more high-quality correspondences to achieve robust registration. The key insight is to progressively generate correspondences and leverage the positional priors of correspondences from the previous iteration to guide more refined local regeneration, thereby obtaining more high-quality point correspondences.

Our method, named Regor, follows a progressive iterative framework composed of three main modules: Local Grouping and Rematching, Local Correspondence Refinement and Global Correspondence Refinement. In each iteration, **Local Grouping and Rematching** leverages *prior-guided local grouping* to gradually reduce the matching space and *generalized mutual matching* to perform reliable local matching. The **Local Correspondence Refinement** introduces an effective *center-aware three-point consistency* to mine inliers and update local correspondences. The **Global Correspondence Refinement** is employed to perform further optimization from a global perspective. Note that, each iteration gradually increases the number of inliers, instead of only removing outliers. Our method achieves state-of-the-art performance on both the 3DMatch [46] and KITTI [16] datasets. Notably, the proposed Regor demonstrates a strong ability to generate correspondences, yielding up to 10 times more correct matches than existing outlier removal methods [6, 7, 47]. The main contributions of this paper are as follows:

- We propose a novel registration approach named Regor to obtain massive high-quality point correspondences, overcoming the challenge of scarce inliers by progressively generating correspondences.
- Our method achieves state-of-the-art performance. In particular, the number of correct correspondences is 10 times higher than that of outlier removal methods.
- Benefiting from the proposed regenerator, our approach achieves robust registration even with weak features.

2. Related Work

2.1. 3D Feature Matching

Feature matching is a crucial step in correspondence-based point cloud registration [1, 20, 30, 32, 43, 45, 46]. Accord-

ing to the approach of feature extraction, feature matching methods can be categorized into two types: traditional descriptor-based methods [31–33] and learning-based methods [1, 20, 30, 43, 45]. Before the widespread adoption of deep learning, features were manually designed descriptors tailored to represent local information. They are classified into LRF-based methods [18, 33] and LRF-free methods [31, 32] based on whether a local reference frame (LRF) is needed. With the advancements in deep learning, numerous learning-based methods have emerged. Based on the representation strategy, they can be further divided into patch-based methods [1, 17, 49] and fragment-based methods [9, 15, 20, 43]. SpinNet [1] and BUFFER [2] have achieved excellent generalization ability through 3D cylindrical convolution. Additionally, CoFiNet [43] and Geo-Transformer [30] propose a coarse-to-fine matching strategy and a Transformer-based framework, which has garnered widespread attention. Subsequent methods [23, 37, 44, 45] have further improved performance by incorporating position encoding and prior information into this framework. ODIN [23] achieves the best performance through diffusion strategy and global optimization.

2.2. Geometry-based Outlier Removal

Geometry-based outlier removal involves fitting models from noisy correspondences using the geometric properties of 3D scenes to estimate poses. The most classic method is random sampling consensus (RANSAC) [14], which samples correspondences multiple times and validates to obtain the best pose, thus mitigating the impact of noisy correspondences. Subsequently, numerous variants of RANSAC [4, 5, 12, 34] have been proposed to address issues of time consumption and instability. TEASER [38] introduces the truncated least squares (TLS) cost to solve for pose and employ rotation invariant measurements to handle outliers. Then, SC²-PCR [6, 7] propose second-order consistency for robust outlier removal. In recent years, methods [39, 47] based on geometric consistency have been widely applied in point cloud registration, leveraging graph-theoretic approaches [13] and geometric consistency to remove outliers, thereby achieving robust registration.

2.3. Learning-based Outlier Removal

Recent researches [3, 22, 29] have integrated deep learning into outlier removal and pose estimation, showcasing promising performance through training. Inspired by image matching [41], 3DRegNet [29] adapts CN-Net [41] to remove outliers in point cloud registration. Subsequently, PointDSC [3] embeds the spatial consistency into feature maps and employs neural spectral matching to estimate confidence of correspondences, achieving robust registration. VBReg [22] utilizes variational Bayesian inference to achieve better outlier removal. Hunter [40] addresses severe




Figure 2. Overall framework of our method. We extract features from the original point cloud, obtaining features $\mathbf{F}^P$ and $\mathbf{F}^Q$ as input for our method. Subsequently, a progressive process is applied to iteratively regenerate more accurate and denser correspondences $\mathcal{G}^t$. At each iteration, the output correspondences $\mathcal{G}^{t-1}$ from the previous iteration serve as input. Firstly, prior-guided local grouping is employed to sample seed corresponding points and form local correspondence regions $\mathbf{P}_i^t$ and $\mathbf{Q}_i^t$. Then, for each pair of local correspondence regions, generalized mutual matching is performed to get new correspondences. Next, these correspondences are refined locally and globally using our center-aware three-point consistency, followed by a merging operation $\oplus$ of local correspondences $\mathcal{G}_i^t$ using a hash table. Finally, using these refined correspondences, we achieve robust and accurate transformation estimation $\mathbf{T}\{\mathbf{R}, \mathbf{t}\}$ only with SVD.

outlier issues by introducing higher-order consistency using hypergraphs. However, while both geometry-based and learning-based methods have shown effectiveness in outlier removal, they inherently rely on the initial correspondences and only trim down on them, which poses challenges when initial inliers are scarce.

3. Method

3.1. Problem Formulation

Point cloud registration is aligning two point clouds $\mathbf{P}$ and $\mathbf{Q}$, captured from different perspectives, by estimating a pose transformation $\mathbf{T}\{\mathbf{R}, \mathbf{t}\}$ where $\mathbf{R}$ is a rotation matrix and $\mathbf{t}$ is a translation vector. Feature-based point cloud registration estimates correspondences through feature extraction and matching, then minimizes the average Euclidean distance between these correspondences, as shown in Eq. 1:

$$\arg \min_{\mathbf{R} \in SO(3), \mathbf{t} \in \mathbb{R}^3} \sum_{(\mathbf{p}_{x_i}, \mathbf{q}_{y_i}) \in \mathcal{I}} \|\mathbf{R} \cdot \mathbf{p}_{x_i} + \mathbf{t} - \mathbf{q}_{y_i}\|_2^2. \quad (1)$$

In this study, unlike outlier removal methods that perform “subtraction” filtering on initial correspondences, resulting in a subset of initial correspondences highly dependent on them, our aim is to generate more higher-quality correspondences $\mathcal{G}^t$ where $t = 1, 2, \dots, m$ represents the current iterative stage. We define the regeneration process as $\mathcal{G}^t = \Theta(\mathcal{G}^{t-1}; \mathbf{P}, \mathbf{Q}, \varphi^t)$ where Θ is the operation of the correspondence regeneration, and φ^t represents the parameter used in the regeneration of the t th stage.

3.2. Progressive Regeneration: Key Insight

Existing methods focus on pruning the initial correspondences to remove outliers and retain inliers. Due to common problems such as weak feature representation and low overlap, the proportion of correct initial correspondences is typically very small. Even all correct correspondences are identified, achieving accurate registration remains challenging. To this end, we design a novel progressive correspondence regeneration strategy to address the scarcity of initial inliers and the frequent failures of correspondence filtering under extreme outlier ratios.

The key idea of our method is to progressively regenerate better correspondences in a radius-variable local sphere. We use the positional priors of correspondences generated in the previous stage to guide the creation of the local sphere. In each local sphere, we regenerate more correspondences with higher quality. Compared to direct global matching, re-matching and refining within these local spheres significantly reduce the problem’s scale and search space, resulting in more precise matches and reduced computational time. As the quality of the correspondence gets higher, we progressively reduce the size of the local spheres to achieve exact convergence of the correspondences. This approach not only generates a large number of new dense correspondences but also incrementally enhances their accuracy.

3.3. Local Grouping and Rematching

In this module, we implement the core of each iteration: local correspondence generation. First, we introduce prior-guided local grouping to leverage priors from the previous




Figure 3. Illustration of GMM. The orange, blue, green, and red lines are the MNN correspondences, NN correspondences, correct correspondences, and wrong correspondences, respectively.

iteration. Then, we propose a generalized mutual matching technique to ensure the generation of a sufficient number of inliers even under high outlier rates.

Prior-guided Local Grouping. The input to our method is the correspondence $\mathcal{G}^{t-1}$ of the previous stage. Efficiency is crucial for our method due to the need for iterative processing. Therefore, we sample $\mathcal{G}^{t-1}$ using efficient random sampling [19] to derive seed correspondences $\tilde{\mathcal{G}}^{t-1}$. We then utilize the positional priors of each seed correspondence to guide the construction of local regions. Specifically, we conduct a radius r^t nearest neighbor search centered on each seed correspondence $\tilde{g}_i^{t-1} \in \tilde{\mathcal{G}}^{t-1}$ to obtain local point clouds $\mathbf{P}_i^t$ and $\mathbf{Q}_i^t$, $i = 1, 2, \dots, n$:

$$\mathbf{P}_i^t = \text{RNN}(\mathbf{p}_i^{t-1}; r^t, \tilde{\mathcal{G}}^{t-1}), \mathbf{Q}_i^t = \text{RNN}(\mathbf{q}_i^{t-1}; r^t, \tilde{\mathcal{G}}^{t-1}), \quad (2)$$

where $\text{RNN}(\mathbf{p}; r, \mathcal{G})$ is a radius nearest neighbor search on $\mathcal{G}$ with $\mathbf{p}$ as the center and r as the radius. $\mathbf{p}_i^{t-1}$ and $\mathbf{q}_i^{t-1}$ satisfy $\tilde{g}_i^{t-1} = (\mathbf{p}_i^{t-1}, \mathbf{q}_i^{t-1})$. We perform the above grouping operation for each pair of seed corresponding points and obtain n pairs of local region $(\mathbf{P}_i^t, \mathbf{Q}_i^t)$.

Generalized Mutual Matching. The next step is to generate new point correspondences from each pair of local region by feature matching. However, the features in small local regions are usually similar, which results in traditional nearest-neighbor (NN) matching [11] producing a large number of false correspondences. Unfortunately, mutual NN matching [24, 26] may even fail to find any correspondence at all. To address this issue, we propose a generalized mutual matching to relax the strict mutual constraint, thereby improving the robustness of the local correspondences.

Given point cloud $\mathbf{Q}_i^t$ as the query, the NN matching matrix $\mathbf{M}_1^{P \rightarrow Q}$ and multi-nearest-neighbor (MNN) matching matrix $\mathbf{M}_2^{P \rightarrow Q}$ are obtained by performing a k-nearest-

neighbor (KNN) search in feature space, respectively:

$$\begin{aligned} \mathbf{M}_1^{P \rightarrow Q}(m, n) &= \text{NN}(\mathbf{F}(\mathbf{p}_m), \mathbf{F}(\mathbf{q}_n)), \\ \mathbf{M}_2^{P \rightarrow Q}(m, n) &= \text{MNN}(\mathbf{F}(\mathbf{p}_m), \mathbf{F}(\mathbf{q}_n)), \end{aligned} \quad (3)$$

where $\forall \mathbf{p}_m \in \mathbf{P}_i^t, \mathbf{q}_n \in \mathbf{Q}_i^t$ and $\mathbf{F}(\cdot)$ represents the mapping from points to features. $\text{NN}(\cdot)$ and $\text{MNN}(\cdot)$ represent the indicator functions for NN and MNN matching, respectively. If the matching condition is satisfied, the value is set to 1. Then, we conduct mutual matching from $\mathbf{P}_i^t$ to $\mathbf{Q}_i^t$ and from $\mathbf{Q}_i^t$ to $\mathbf{P}_i^t$, resulting in four matching matrices $\mathbf{M}_1^{P \rightarrow Q}$, $\mathbf{M}_2^{P \rightarrow Q}$, $\mathbf{M}_1^{Q \rightarrow P}$, and $\mathbf{M}_2^{Q \rightarrow P}$.

By computing the Hadamard product between the NN and MNN matching matrices, we obtain the mutual matching matrices $\mathbf{M}_*^{P \rightarrow Q}$, $\mathbf{M}_*^{Q \rightarrow P}$. Further, we perform an element-wise logical OR operation to relax the constraint and get the generalized mutual matching matrix:

$$\mathbf{M}^* = (\mathbf{M}_1^{P \rightarrow Q} \odot \mathbf{M}_2^{Q \rightarrow P}) \otimes (\mathbf{M}_1^{Q \rightarrow P} \odot \mathbf{M}_2^{P \rightarrow Q}), \quad (4)$$

where operation $\odot$ and $\otimes$ denote Hadamard product and element-wise logical OR, respectively. It is important to note that, to avoid a large number of outliers, we discard the strategy of directly applying mutual matching only based on MNN. Finally, using the matching matrix $\mathbf{M}^*$, we establish n sets of local correspondences $\mathcal{G}_i^t$.

3.4. Local Correspondence Refinement

Because some corresponding local regions may not overlap exactly, the rematching in local regions still introduces some erroneous correspondences. Therefore, we further enhance the quality of the regenerated correspondences through local correspondence refinement.




Figure 4. Illustration of the local correspondence correction. Yellow and blue areas represent a pair of local point clouds. The purple line indicates the correct correspondence after correction.

Center-aware Three-point Consistency. Note that, our seed correspondences $g_i^{t-1} = (\mathbf{p}_i^{t-1}, \mathbf{q}_i^{t-1})$ (central points) are derived from the previous stage and exhibit high accuracy. Therefore, they can serve as prior information to

guide the optimization of local correspondence refinement. Based on this, we propose an efficient center-aware three-point consistency (CTC) to search inliers.

We first utilize efficient translation and rotation invariants [38] to describe the consistency $s_\sigma(g_i, g_j)$ between two point pairs $g_i = (\mathbf{p}_i, \mathbf{q}_i)$ and $g_j = (\mathbf{p}_j, \mathbf{q}_j)$:

$$s_\sigma(g_i, g_j) = \mathbb{1} \left(\left| \|\mathbf{p}_i - \mathbf{p}_j\|_2 - \|\mathbf{q}_i - \mathbf{q}_j\|_2 \right| \leq \sigma \right). \quad (5)$$

Next, we leverage the prior information of accurate center points to compute the center-aware consistency between g_j and g_k , denoted as $s_\sigma(g_j, g_i^{t-1}) \cdot s_\sigma(g_i^{t-1}, g_k)$. Additionally, we account for the possibility that seed correspondences may not always be correct. Relying solely on the aforementioned criterion could allow errors in the seed correspondences to negatively impact the optimization of the entire region. To mitigate this, we introduce a strict point-pair constraint $s_{\frac{\sigma}{2}}(g_j, g_k)$. CTC is computed as:

$$s_\Delta(g_j, g_k) = (s_\sigma(g_j, g_i^{t-1}) \cdot s_\sigma(g_i^{t-1}, g_k)) \parallel s_{\frac{\sigma}{2}}(g_j, g_k), \quad (6)$$

where $\parallel$ represents operator OR and it means if either condition is satisfied, the correspondence is considered correct. Therefore, we get the score matrix $\mathbf{S}_i^t = [s_\Delta(g_j, g_k)]_{jk} \in \mathbb{R}^{N \times N}$ of the local correspondences $\mathcal{G}_i^t \in \mathbb{R}^{N \times 2}$. To evaluate the overall quality of the local region, we define a local overall score:

$$\text{Score}(\mathcal{G}_i^t) = \frac{\|\mathbf{S}_i^t\|_1}{a \cdot N}, \quad (7)$$

where $\|\cdot\|_1$ represents the L_1 -norm of a matrix and $a \in (0, 1)$ denotes the threshold for the proportion of correct correspondences. Essentially, we consider regions where the proportion of correct correspondences exceeds a as correct correspondence regions. It is proved in the Appendix.

Local Correspondence Correction. Within each correct corresponding region, we identify a set of precise correspondences $\hat{\mathcal{G}}_i^t$ to guide the refinement of correspondences across the entire local region. Specifically, based on the CTC score, we identify precise correspondences $\hat{\mathcal{G}}_i^t$ with high scores. Subsequently, using these correspondences, we estimate the local transformation $\mathbf{T}_i^t\{\mathbf{R}_i^t, \mathbf{t}_i^t\}$:

$$\mathbf{R}_i^t, \mathbf{t}_i^t = \min_{\mathbf{R}, \mathbf{t}} \sum_{(\mathbf{p}_j, \mathbf{q}_j) \in \hat{\mathcal{G}}_i^t} \|\mathbf{R} \cdot \mathbf{p}_j + \mathbf{t} - \mathbf{q}_j\|_2^2. \quad (8)$$

Leveraging this local pose $\mathbf{T}_i^t\{\mathbf{R}_i^t, \mathbf{t}_i^t\}$, we search for the nearest points $\mathbf{q}'$ in the target local point cloud $\mathbf{Q}_i^t$ for each point $\mathbf{p}$ in the source local point cloud $\mathbf{P}_i^t$:

$$\begin{aligned} \mathbf{q}' &= \arg \min_{\mathbf{q}_j \in \mathbf{Q}_i^t} \|\mathbf{R}_i^t \mathbf{p} + \mathbf{t}_i^t - \mathbf{q}_j\|_2, \\ s.t. & \|\mathbf{R}_i^t \mathbf{p} + \mathbf{t}_i^t - \mathbf{q}'\|_2 \leq \sigma_d, \forall (\mathbf{p}, \mathbf{q}) \in \mathcal{G}_i^t. \end{aligned} \quad (9)$$

Through the one-way refinement of $\mathbf{Q}$, we update our local correspondences $\mathcal{G}_i^t \leftarrow \Omega_l(\mathcal{G}_i^t) = \{(\mathbf{p}, \mathbf{q}')\}$ where Ω_l represents the operation of local correspondence correction.

3.5. Global Correspondence Refinement

Through local refinement, we obtain n sets of refined local correspondences. However, consistency across sets is not fully aligned. Therefore, we need a global merging and refinement step to complete the global correspondence optimization. Since there is no seed correspondence at the global scale, we use second-order consistency [6].

Correspondence Merging. After correcting the local correspondences, we merge them into global correspondences $\mathcal{G}^t$. It is worth noting that some local regions overlap with each other. Therefore, we utilize a hash table $\mathbb{H}$ to implement correspondence merging. By using point indices as hash table keys, we efficiently merge local correspondences and remove duplicates to obtain the global correspondences $\mathcal{G}^t = \bigcup_{i=1}^n \mathcal{G}_i^t$. With this strategy, we further enhance the efficiency of our method and save memory overhead.

Global Correspondence Correction. Although we achieve the correspondence correction in local regions, if the seed correspondences $\tilde{\mathcal{G}}^{t-1}$ are significantly incorrect, it may still result in wrong refinement for the entire local region. Therefore, similar to local correction, we refine global correspondences $\mathcal{G}^t \leftarrow \Omega_g(\mathcal{G}^t)$, where Ω_g represents the operation of global correspondence correction. The key difference is that, in the global stage, we use second-order consistency [6] to identify correct correspondences. By leveraging correspondences that satisfy maximum consistency, we rectify anomalous correspondences globally, completing one iteration of correspondence regeneration $\Theta(\cdot)$.

4. Experiment

We evaluate our algorithm using the indoor 3DMatch dataset [46] and the outdoor KITTI dataset [16], demonstrating its superior performance. Furthermore, we introduce the 3DMatch-EOR benchmark and conduct experiments to demonstrate the effectiveness of our method in scenarios with very few correct initial correspondences. Additionally, we also conduct robustness test experiments to verify that our algorithm can achieve robust registration using only weak features, thereby demonstrating its robustness across different descriptors. Finally, we conduct ablation studies to validate the effectiveness of each module.

Evaluation Metric. Following [6, 47], we use registration recall (RR), rotation error (RE), and translation error (TE) to evaluate registration performance. For outdoor scenes, we consider a registration successful if the error is within (5° , 60 cm), and for indoor scenes, within (15° , 30 cm). Additionally, we assess the quality of correspondences using feature matching recall (FMR) and inlier precision (IP) which represents the proportion of inliers to the total correspondences. To further demonstrate our method's ability to generate a higher number of accurate correspondences, we define two new metrics: inlier number ratio (INR) and inlier

Method	FPFH (Traditional Descriptor)						FCGF (Learning-based Descriptor)						Time(s)
	RR($\uparrow$)	RE($\downarrow$)	TE($\downarrow$)	IP($\uparrow$)	INR($\uparrow$)	IN($\uparrow$)	RR($\uparrow$)	RE($\downarrow$)	TE($\downarrow$)	IP($\uparrow$)	INR($\uparrow$)	IN($\uparrow$)	
PointDSC [3]	76.96	2.25	6.76	66.64	70.16	290.23	92.61	2.10	6.48	78.37	85.88	1174.62	0.22
VBReg [22]	78.01	2.28	7.34	67.23	70.93	166.56	92.77	2.24	6.82	79.92	86.30	690.38	1.02
Hunter [40]	84.70	1.80	6.46	74.39	79.27	309.28	94.05	1.86	6.54	82.55	88.01	1203.88	0.59
SM [25]	55.88	2.94	8.15	47.96	70.69	230.03	86.57	2.29	7.07	81.44	38.36	405.10	0.03
TEASER [38]	75.48	2.48	7.31	73.01	78.60	270.98	85.77	2.73	8.66	82.43	87.90	1182.44	0.07
RANSAC-1M [14]	64.20	4.05	11.35	63.96	69.03	264.09	88.42	3.05	9.42	77.96	83.08	672.22	0.93
SC ² -PCR [6]	83.24	2.21	6.70	73.64	78.78	289.17	93.16	2.11	6.44	80.17	87.24	1140.85	0.11
MAC [47]	83.90	2.11	6.80	-	-	-	93.72	2.07	6.52	-	-	-	0.95
FastMAC [48]	82.87	2.15	6.73	-	-	-	92.67	2.00	6.47	-	-	-	0.27
Regor (ours)	88.48	1.70	5.94	82.68	1253.74	2532.40	<u>93.96</u>	1.75	6.14	87.48	263.81	2620.30	0.36

Table 1. Quantitative comparison on the 3DMatch dataset with descriptors FPFH [32] and FCGF [9].




Figure 5. Correspondences before and after processing.

	RR(%)	RE($^{\circ}$)	TE(cm)	IP(%)	INR(%)	IN(/)
FPFH						
PointDSC [3]	27.91	4.27	10.45	23.80	25.76	46.52
VBReg [22]	30.83	4.38	10.92	26.61	29.70	48.41
Hunter [40]	36.90	3.89	10.05	<u>31.55</u>	<u>36.17</u>	<u>56.01</u>
RANSAC [14]	19.83	4.67	10.32	30.82	21.20	37.02
SC ² -PCR [6]	35.93	4.26	10.86	30.11	35.65	55.23
MAC [47]	40.88	3.66	<u>9.45</u>	-	-	-
FastMAC [48]	38.46	4.04	10.47	-	-	-
PCRegen (ours)	43.96	2.89	8.93	37.73	778.35	643.24
Predator						
PointDSC [3]	68.30	<u>3.45</u>	9.44	56.35	66.70	863.40
VBReg [22]	69.82	3.87	10.07	56.90	67.09	872.05
Hunter [40]	<u>71.10</u>	3.58	9.72	58.90	<u>69.06</u>	<u>892.07</u>
RANSAC [14]	64.85	4.28	11.04	56.44	66.81	869.21
SC ² -PCR [6]	69.46	3.48	9.63	58.07	68.63	883.59
MAC [47]	70.91	3.69	9.81	-	-	-
FastMAC [48]	68.77	3.90	10.32	-	-	-
PCRegen (ours)	72.04	3.17	<u>9.46</u>	63.95	220.46	2138.47

Table 2. Quantitative comparison on the 3DLoMatch Dataset.

numbers (IN). INR represents the ratio of the final number of inliers IN_{final} to the initial number of inliers $IN_{initial}$, defined explicitly in Appendix.

4.1. Evaluation on Indoor Scenes

Experimental Setup. For indoor scenes, we utilize the 3DMatch dataset. Our method is validated on both 3DMatch [46] and 3DLoMatch benchmarks [20]. Following [3, 6], we perform voxel downsampling on the raw point clouds, and use two descriptors, FPFH [32] and FCGF [10], to extract features as inputs for each method.

Result on 3DMatch. We compare our method against several baselines on the 3DMatch dataset, including PointDSC [3], VBReg [22], Hunter [40], SM [25]

TEASER [38], RANSAC [14], SC²-PCR [6], MAC [47], and FastMAC [48]. The first three baselines are learning-based methods, while others are geometry-based approaches. The results are summarized in Table 1. Our method outperforms the previous state-of-the-art [6, 47]. Using the FPFH [9], our method outperforms all others across all metrics. Notably, our method achieves 88.48% RR, surpassing SC²-PCR [6] by 5.2 percentage points (pp). It also achieves more accurate registration with the lowest RE and TE. One significant advantage of our method is obtaining more high-quality correspondences, as reflected in the INR and IN. Our Regor achieves a 1253.74% INR, which is much better than 78.78% of SC²-PCR. Previous methods [6, 47, 48] only perform top-down filtering, which inevitably limits the INR to below 100%. In contrast, our Regor regenerates more high-quality correspondences, producing around 10 times the original amount. This advantage allows our method to achieve robust registration even when there are few initial correct correspondences. Moreover, despite employing multiple iterations, it still maintains a comparable speed to the baselines [22, 40] due to its simple sampling strategy and efficient consistency computation.

Result on 3DLoMatch. Since Predator [20] is specifically designed for scenarios with low overlap, we use Predator [20] and FPFH [32] to extract features on 3DLoMatch. The comparative results are presented in Table 2. Our approach achieves the highest RR and RE on 3DLoMatch, regardless of the descriptors used. Due to the inherent weakness of FPFH, nearly all baselines perform poorly on 3DLoMatch. In contrast, our method shows significant performance improvement, surpassing MAC [47] by 3.08 pp.

Result on 3DMatch-EOR. The key to the effectiveness of our method lies in its ability to generate more accurate correspondences, particularly in scenarios with extremely high outlier ratios. To validate this, we select extreme outlier ratio cases from the 3DMatch/3DLoMatch datasets and propose the 3DMatch-EOR benchmark which is described in Appendix. We establish two benchmarks with outlier rates exceeding 99% and 90%, respectively. Our method is compared to SC²-PCR [6] and PointDSC [3], with results




Figure 6. Correspondences visualization in different stages.

	RR(%)	RE(°)	TE(cm)	FMR(%)	IP(%)	INR(%)	IN(/)	RR(%)	RE(°)	TE(cm)	FMR(%)	IP(%)	INR(%)	IN(/)
Outlier ratio above 90%														
	3DMatch							3DLoMatch						
PointDSC [3]	71.12	2.52	7.79	77.76	60.05	63.55	129.12	23.27	4.84	12.94	39.78	18.16	26.57	31.18
SC ² -PCR [6]	78.58	2.49	7.73	84.02	51.91	73.73	138.98	35.39	4.29	10.96	42.58	29.6	35.14	49.73
PCRegen (ours)	84.88	1.86	6.67	86.38	78.14	1457.54	2334.36	43.50	2.95	9.24	46.34	36.96	780.05	632.62
Outlier ratio above 99%														
	3DMatch							3DLoMatch						
PointDSC [3]	2.56	5.12	15.28	10.90	3.07	4.88	0.17	0.75	4.93	29.34	0.75	0.08	1.46	0.02
SC ² -PCR [6]	9.15	4.94	12.21	16.20	6.08	9.61	0.34	3.28	5.94	19.44	7.88	2.47	4.96	0.89
PCRegen (ours)	26.76	2.17	8.60	28.87	21.54	1545.12	666.94	9.27	3.36	12.39	11.84	6.30	291.40	118.45

Table 3. Results on the 3DMatch-EOR benchmark with extreme outlier ratios.

shown in Table 3. In cases with such high outlier rates, achieving robust registration is exceptionally challenging, as the number of initial correct correspondences is minimal. When the outlier rate exceeded 99%, the average inlier count of SC²-PCR is less than 1, leading to registration failure and achieving only 9.15%/3.28% in RR. In contrast, our method achieves an average of 667 inliers, an improvement of nearly 2000 times compared to SC²-PCR [6], resulting in a 26.76% RR. Our Regor enables success in scenarios where failure would otherwise be inevitable, allowing for relatively robust registration even with few initial inliers.

4.2. Evaluation on Outdoor Scenes

Experimental Setup. Following [7, 47], we validate our method’s performance on the outdoor KITTI dataset. We select sequences 8 to 10 for testing. Consistent with the settings [6], we perform downsampling using a 30 cm voxel grid and extract features using FPFH descriptors.

Result on KITTI. We conduct comparative experiments between our method and PointDSC [3], VBReg [22], Hunter [40], RANSAC [14], SC²-PCR [6], MAC [47], and FastMAC [48] on the KITTI dataset. The results are presented in Table 4. Our Regor achieves the best performance with the highest RR, outperforming RANSAC [14] by 25.41

	RR(%)	RE(°)	TE(cm)	IP(%)	INR(%)	IN(/)
PointDSC [3]	99.46	0.57	7.22	91.39	93.05	247.11
VBReg [22]	98.92	8.39	8.41	90.82	92.13	246.09
Hunter [40]	99.82	0.55	6.29	92.11	94.58	247.50
RANSAC [14]	74.41	1.55	30.20	85.29	87.20	189.70
SC ² -PCR [6]	99.82	0.61	8.16	93.59	95.94	248.46
MAC [47]	99.46	0.59	8.70	-	-	-
FastMAC [48]	98.02	0.55	8.24	-	-	-
Regor (ours)	99.82	0.50	5.36	79.57	2486.82	5353.22

Table 4. Quantitative comparison on the KITTI Dataset.

pp. It also demonstrates lower RE and TE, achieving the most accurate registration. Furthermore, our method shows superior performance in generating new correspondences, with INR over 25 times higher than those of other baselines [6, 40]. Despite IP is lower compared to SOTA [6], our method achieves more accurate pose estimation due to a larger number of high-quality inliers.

4.3. Robustness to Different Features

To further assess the robustness of various descriptors, we also test with traditional descriptors such as FPFH [32], SHOT [33], PFH [31], and PoRS [18], as well as learning-based descriptors like SpinNet [1], FCGF [9], Predator [20], and GeoTrans [30]. As shown in Table 5, our Regor con-




Figure 7. RR under different features.

Method	RR(%)↑	IP(%)↑	IN(%)↑
FPFH [32]	54.34	45.74	228.71
FPFH [32] + Regor	88.48 (+34.14)	82.68 (+36.94)	1253.74 (×5.48)
PFH [31]	69.38	57.28	286.38
PFH [31] + Regor	89.28 (+19.90)	81.72 (+24.44)	2477.39 (×8.65)
RoPS [18]	54.16	42.49	212.46
RoPS [18] + Regor	91.37 (+37.21)	84.29 (+41.80)	2542.32 (×11.96)
SHOT [33]	55.33	46.01	229.50
SHOT [33] + Regor	71.90 (+16.57)	66.41 (+20.40)	1981.10 (×8.63)
FCGF [9]	86.32	56.85	405.27
FCGF [9] + Regor	93.96 (+7.64)	87.48 (+30.63)	2620.30 (×6.46)
SpinNet [1]	88.60	47.53	444.84
SpinNet [1] + Regor	96.30 (+7.70)	88.24 (+40.71)	2799.74 (×6.29)
Predator [20]	89.00	58.01	444.38
Predator [20] + Regor	91.44 (+2.44)	86.84 (+28.83)	4608.34 (×10.37)
GeoTrans [30]	91.80	75.20	449.85
GeoTrans [30] + Regor	97.10 (+5.30)	91.01 (+15.81)	2649.30 (×5.88)

Table 5. Enhancement of feature-based methods.

sistently improves performance across all features, which demonstrates it is robust to different features. Notably, GeoTrans [30]+Regor achieves a SOTA RR of 97.10%.

Good Performance with Weak Features. Moreover, our method significantly enhance the performance of the weak descriptors, increasing the RR from 54.16% to 91.37% with RoPS [18]. As can be seen in Figure 7, with our Regor, the traditional descriptors [18, 31, 32] achieves 88%+ RR, surpassing the learning-based baseline [9]. When combined with our Regor, RoPS [18] exhibits performance that exceeds the learning-based methods [1, 9, 20], only second to GeoTrans [30]. This demonstrates that our method can achieve robust registration even with weak features.

4.4. Ablation Study

We conduct ablation experiments on the 3DMatch and 3DMatch-EOR benchmarks using FPFH descriptor.

Progressive or One-stage. We verify the effectiveness of our progressive strategy by only using the one-stage process. The results are shown in Table 6 (1) and (2). Our strategy improves RR by 1.65/2.82 pp compared to the one-stage approach. It is attributed to the heuristic setting: starting with a larger radius ensures global optimality and robustness, and gradually reducing the radius with each iter-

No.	Methods	3DMatch			3DMatch-EOR		
		RR	IP	IN	RR	IP	IN
1)	One-stage	86.81	76.34	290.32	23.94	15.88	9.13
2)	Progressive*	88.48	82.68	2532.40	26.76	21.54	666.94
3)	NNM	87.35	80.90	2487.23	25.06	19.82	643.45
4)	MM	86.81	81.03	2392.53	23.94	20.80	600.09
5)	GMM*	88.48	82.68	2532.40	26.76	21.54	666.94
6)	Only Local	87.35	78.40	2526.63	23.06	17.32	668.08
7)	Only Global	76.37	80.48	2502.21	24.65	16.57	641.42
8)	Local & Global*	88.48	82.68	2532.40	26.76	21.54	666.94
9)	SC	86.88	81.09	2491.54	22.54	18.18	574.74
10)	CTC*	88.48	82.68	2532.40	26.76	21.54	666.94

Table 6. Ablation study on 3DMatch and 3DMatch-EOR.

ation enhances accuracy. More importantly, IN drops by 1~2 orders of magnitude, suggesting that our progressive strategy is key to generating more correspondences.

Matching Strategy. We replace our GMM with mutual matching (MM) and nearest neighbor matching (NNM). Results are shown in Table 6 (3), (4) and (5). The model with GMM achieves the highest RR, IP, and IN. Compared with MM, using GMM results in a significant improvement in IN, which indicates our GMM is helpful to generate more inliers and addresses the issue of few initial inliers.

Local-Global Refinement. To demonstrate the importance of our correspondence refinement, we compare the complete model with two ablation models. The results, shown in Rows 6, 7, and 8 of Table 6, indicate a significant performance drop when either local or global refinement is removed, with RR falling to 87% and 86%. In addition, the IP of the ablation model decreases by over 4 pp, which underscores the substantial contribution of the local-global refinement to the quality of the correspondences.

Spatial Compatibility. We compared the SOTA second-order consistency (SC) [6, 7] with our CTC, and the results are shown at the bottom of Table 6. Our method with CTC obtains higher IP and IN, which indicates that CTC is more suitable for local refinement and achieves higher-quality correspondences. In addition, our CTC has less computational complexity than SC. Through it, we can implement an efficient iterative process.

5. Conclusion

In this paper, we introduce a new idea for optimizing correspondences. Unlike previous methods that focus solely on outlier removal, our Regor effectively addresses the challenge of limited correct initial correspondences. By generating more high-quality correspondences, our method enables more robust and accurate registration, achieving state-of-the-art performance across all datasets. Notably, it demonstrates significant improvements even under extreme outlier conditions. When combined with our Regor, using traditional descriptors can also achieve robust registration.

References

- [1] Sheng Ao, Qingyong Hu, Bo Yang, Andrew Markham, and Yulan Guo. Spinnet: Learning a general surface descriptor for 3d point cloud registration. In *CVPR*, pages 11753–11762, 2021. 1, 2, 7, 8
- [2] Sheng Ao, Qingyong Hu, Hanyun Wang, Kai Xu, and Yulan Guo. Buffer: Balancing accuracy, efficiency, and generalizability in point cloud registration. In *CVPR*, pages 1255–1264, 2023. 2
- [3] Xuyang Bai, Zixin Luo, Lei Zhou, Hongkai Chen, Lei Li, Zeyu Hu, Hongbo Fu, and Chiew-Lan Tai. Pointdsc: Robust point cloud registration using deep spatial consistency. In *CVPR*, pages 15859–15869, 2021. 1, 2, 6, 7, 3, 5
- [4] Daniel Barath and Jiří Matas. Graph-cut ransac. In *CVPR*, pages 6733–6741, 2018. 2
- [5] Daniel Barath and Gábor Valasek. Space-partitioning ransac. In *ECCV*, pages 721–737, 2022. 2
- [6] Zhi Chen, Kun Sun, Fan Yang, and Wenbing Tao. Sc2-pcr: A second order spatial compatibility for efficient and robust point cloud registration. In *CVPR*, pages 13221–13231, 2022. 1, 2, 5, 6, 7, 8, 3, 4
- [7] Zhi Chen, Kun Sun, Fan Yang, Lin Guo, and Wenbing Tao. Sc2-pcr+: Rethinking the generation and selection for efficient and robust point cloud registration. *IEEE TPAMI*, 2023. 2, 7, 8
- [8] Yuxin Cheng, Zhiqiang Huang, Siwen Quan, Xinyue Cao, Shikun Zhang, and Jiaqi Yang. Sampling locally, hypothesis globally: accurate 3d point cloud registration with a ransac variant. *Visual Intelligence*, 1(1):20, 2023. 1
- [9] Christopher Choy, Jaesik Park, and Vladlen Koltun. Fully convolutional geometric features. In *ICCV*, pages 8957–8965, 2019. 2, 6, 7, 8
- [10] Christopher Choy, Wei Dong, and Vladlen Koltun. Deep global registration. In *CVPR*, pages 2514–2523, 2020. 6
- [11] Chin Seng Chua and Ray Jarvis. Point signatures: A new representation for 3d object recognition. *IJCV*, 25:63–85, 1997. 4
- [12] Ondřej Chum and Jiří Matas. Optimal randomized ransac. *IEEE TPAMI*, 30(8):1472–1482, 2008. 2
- [13] David Eppstein, Maarten Löffler, and Darren Strash. Listing all maximal cliques in sparse graphs in near-optimal time. In *Algorithms and Computation: 21st International Symposium*, pages 403–414, 2010. 2
- [14] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981. 1, 2, 6, 7, 5
- [15] Isak Gath and Amir B. Geva. Unsupervised optimal fuzzy clustering. *IEEE TPAMI*, 11(7):773–780, 1989. 2
- [16] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, pages 3354–3361, 2012. 2, 5, 3, 4
- [17] Zan Gojcic, Caifa Zhou, Jan D. Wegner, and Andreas Wieser. The perfect match: 3d point cloud matching with smoothed densities. In *CVPR*, pages 5540–5549, 2019. 2
- [18] Yulan Guo, Ferdous Sohel, Mohammed Bennamoun, Min Lu, and Jianwei Wan. Rotational projection statistics for 3d local surface description and object recognition. *IJCV*, 105: 63–86, 2013. 2, 7, 8
- [19] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. Randla-net: Efficient semantic segmentation of large-scale point clouds. In *CVPR*, pages 11108–11117, 2020. 4
- [20] Shengyu Huang, Zan Gojcic, Mikhail Usvyatsov, Andreas Wieser, and Konrad Schindler. Predator: Registration of 3d point clouds with low overlap. In *CVPR*, pages 4267–4276, 2021. 1, 2, 6, 7, 8, 4
- [21] Xiaoshui Huang, Guofeng Mei, Jian Zhang, and Rana Abbas. A comprehensive survey on point cloud registration. *arXiv preprint arXiv:2103.02690*, 2021. 1
- [22] Haobo Jiang, Zheng Dang, Zhen Wei, Jin Xie, Jian Yang, and Mathieu Salzmann. Robust outlier rejection for 3d registration with variational bayes. In *CVPR*, pages 1148–1157, 2023. 1, 2, 6, 7
- [23] Shengze Jin, Iro Armeni, Marc Pollefeys, and Daniel Barath. Multiway point cloud mosaicking with diffusion and global optimization. In *CVPR*, pages 20838–20849, 2024. 2
- [24] Andrew E Johnson and Martial Hebert. Using spin images for efficient object recognition in cluttered 3d scenes. *TPAMI*, 21(5):433–449, 1999. 4
- [25] Marius Leordeanu and Martial Hebert. A spectral technique for correspondence problems using pairwise constraints. In *ICCV*, pages 1482–1489, 2005. 6, 4
- [26] Hongdong Li and Richard Hartley. The 3d-3d registration problem revisited. In *ICCV*, pages 1–8. IEEE, 2007. 4
- [27] Fan Lu, Guang Chen, Yinlong Liu, Lijun Zhang, Sanqing Qu, Shu Liu, and Rongqi Gu. Hregnet: A hierarchical network for large-scale outdoor lidar point cloud registration. In *ICCV*, pages 16014–16023, 2021. 1, 2
- [28] Carl Olsson, Fredrik Kahl, and Magnus Oskarsson. Branch-and-bound methods for euclidean registration problems. *IEEE TPAMI*, 31(5):783–794, 2008. 1
- [29] G. Dias Pais, Srikumar Ramalingam, Venu Madhav Govindu, Jacinto C. Nascimento, Rama Chellappa, and Pedro Miraldo. 3dregnet: A deep neural network for 3d point registration. In *CVPR*, pages 7191–7201, 2020. 2
- [30] Zheng Qin, Hao Yu, Changjian Wang, Yulan Guo, Yuxing Peng, and Kai Xu. Geometric transformer for fast and robust point cloud registration. In *CVPR*, pages 11143–11152, 2022. 1, 2, 7, 8
- [31] Radu Bogdan Rusu, Nico Blodow, Zoltan Csaba Marton, and Michael Beetz. Aligning point cloud views using persistent feature histograms. In *IROS*, pages 3384–3391, 2008. 2, 7, 8
- [32] Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast point feature histograms (fpfh) for 3d registration. In *ICRA*, pages 3212–3217, 2009. 2, 6, 7, 8, 5
- [33] Samuele Salti, Federico Tombari, and Luigi Di Stefano. Shot: Unique signatures of histograms for surface and texture description. *CVIU*, 125:251–264, 2014. 2, 7, 8
- [34] Ruwen Schnabel, Roland Wahl, and Reinhard Klein. Efficient ransac for point-cloud shape detection. In *Comput. Graph Forum*, pages 214–226, 2007. 2
- [35] Xinyu Xie, Yawen Cui, Tao Tan, Xubin Zheng, and Zitong Yu. Fusionmamba: Dynamic feature enhancement for mul-

timodal image fusion with mamba. *Visual Intelligence*, 2(1): 37, 2024. [1](#)

- [36] Kezheng Xiong, Maoji Zheng, Qingshan Xu, Chenglu Wen, Siqi Shen, and Cheng Wang. Speal: Skeletal prior embedded attention learning for cross-source point cloud registration. In *AAAI*, pages 6279–6287, 2024. [1](#)
- [37] Fan Yang, Lin Guo, Zhi Chen, and Wenbing Tao. One-inlier is first: Towards efficient position encoding for point cloud registration. *NeurIPS*, 35:6982–6995, 2022. [2](#)
- [38] Heng Yang, Jingnan Shi, and Luca Carlone. Teaser: Fast and certifiable point cloud registration. *IEEE TRO*, 37(2): 314–333, 2020. [2](#), [5](#), [6](#)
- [39] Jiaqi Yang, Xiyu Zhang, Shichao Fan, Chunlin Ren, and Yanning Zhang. Mutual voting for ranking 3d correspondences. *IEEE TPAMI*, 2023. [2](#)
- [40] Runzhao Yao, Shaoyi Du, Wenting Cui, Aixue Ye, Feng Wen, Hongbo Zhang, Zhiqiang Tian, and Yue Gao. Hunter: Exploring high-order consistency for point cloud registration with severe outliers. *IEEE TPAMI*, 2023. [2](#), [6](#), [7](#)
- [41] Kwang Moo Yi, Eduard Trulls, Yuki Ono, Vincent Lepetit, Mathieu Salzmann, and Pascal Fua. Learning to find good correspondences. In *CVPR*, pages 2666–2674, 2018. [2](#)
- [42] Huan Yin, Xuecheng Xu, Sha Lu, Xieyuanli Chen, Rong Xiong, Shaojie Shen, Cyrill Stachniss, and Yue Wang. A survey on global lidar localization: Challenges, advances and open problems. *IJCV*, pages 1–33, 2024. [1](#)
- [43] Hao Yu, Fu Li, Mahdi Saleh, Benjamin Busam, and Slobodan Ilic. Cofinet: Reliable coarse-to-fine correspondences for robust pointcloud registration. *NeurIPS*, 34:23872–23884, 2021. [1](#), [2](#)
- [44] Hao Yu, Zheng Qin, Ji Hou, Mahdi Saleh, Dongsheng Li, Benjamin Busam, and Slobodan Ilic. Rotation-invariant transformer for point cloud matching. In *CVPR*, pages 5384–5393, 2023. [2](#)
- [45] Junle Yu, Luwei Ren, Yu Zhang, Wenhui Zhou, Lili Lin, and Guojun Dai. Peal: Prior-embedded explicit attention learning for low-overlap point cloud registration. In *CVPR*, pages 17702–17711, 2023. [2](#)
- [46] Andy Zeng, Shuran Song, Matthias Nießner, Matthew Fisher, Jianxiong Xiao, and Thomas Funkhouser. 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. In *CVPR*, pages 1802–1811, 2017. [2](#), [5](#), [6](#), [3](#), [4](#)
- [47] Xiyu Zhang, Jiaqi Yang, Shikun Zhang, and Yanning Zhang. 3d registration with maximal cliques. In *CVPR*, pages 17745–17754, 2023. [1](#), [2](#), [5](#), [6](#), [7](#)
- [48] Yifei Zhang, Hao Zhao, Hongyang Li, and Siheng Chen. Fastmac: Stochastic spectral sampling of correspondence graph. In *CVPR*, pages 17857–17867, 2024. [6](#), [7](#)
- [49] Guiyu Zhao, Zhentao Guo, Xin Wang, and Hongbin Ma. Spherenet: Learning a noise-robust and general descriptor for point cloud registration. *IEEE TGRS*, 62:1–16, 2024. [2](#)

Progressive Correspondence Regenerator for Robust 3D Registration

Supplementary Material

6. Analysis and Proof

6.1. Proof of Theorem 1

Theorem 1. Assuming event Φ , where the consistency of correct correspondences is greater than the maximum consistency of incorrect correspondences, has a probability $P(\Phi) \approx 1$ (which means our method can correctly identify inliers), if the score of the local correspondences satisfies $\text{Score}(\mathcal{G}_i^t) \geq 1$, then the proportion of correct correspondences $b \geq a$.

Proof. For simplicity, we denote the consistency matrix $\mathbf{S}_i^t$ as $\mathbf{S}_{\text{CTC}}$. Thus, we can demonstrate that:

$$\begin{aligned}
 P(\text{Score}(\mathcal{G}_i^t) \geq 1) &= P\left(\frac{\|\mathbf{S}_{\text{CTC}}\|_1}{a \cdot N} \geq 1\right) = 1 \\
 &= P\left(\frac{\max_{1 \leq j \leq N} \sum_{i=1}^N |\mathbf{S}_{\text{CTC}ij}|}{a \cdot N} \geq 1\right) \\
 &= P\left(\frac{\max_{1 \leq j \leq N} \sum_{i=1}^N |\mathbf{S}_{\text{CTC}ij}|}{a \cdot N} \geq 1 | \Phi\right) P(\Phi) \\
 &+ P\left(\frac{\max_{1 \leq j \leq N} \sum_{i=1}^N |\mathbf{S}_{\text{CTC}ij}|}{a \cdot N} \geq 1 | \bar{\Phi}\right) P(\bar{\Phi}) \\
 &\stackrel{(1)}{\approx} P\left(\frac{\max_{1 \leq j \leq N} \sum_{i=1}^N |\mathbf{S}_{\text{CTC}ij}|}{a \cdot N} \geq 1 | \Phi\right) P(\Phi) \\
 &= \left[1 - P\left(\frac{\max_{1 \leq j \leq N} \sum_{i=1}^N |\mathbf{S}_{\text{CTC}ij}|}{a \cdot N} < 1 | \Phi\right)\right] P(\Phi) \\
 &= \left[1 - P\left(\frac{\sum_{i=1}^N |\mathbf{S}_{\text{CTC}il}|}{a \cdot N} < 1 | \Phi\right)^N\right] P(\Phi) \\
 &\stackrel{(2)}{\leq} \left[1 - P\left(\frac{b \cdot N}{a \cdot N} < 1\right)^N\right] P(\Phi) \\
 &= \left[1 - P\left(\frac{b}{a} < 1\right)^N\right] P(\Phi)
 \end{aligned} \tag{10}$$

where step (1) is based on our assumption that $P(\Phi) \approx 1$. Step (2) can be easily derived from the definition of consistency matrix. By definition, the number of correct correspondences is $b \cdot N$. According to the definition of our consistency matrix, When the geometric consistency of correct correspondences is greater than the maximum consistency of incorrect correspondences, that is, event Φ occurs, there is always:

$$\sum_{i=1}^n |\mathbf{S}_{\text{CTC}il}| \leq b \cdot N. \tag{11}$$

According to Eq. 11 and 10, we have already proved that $\left[1 - P\left(\frac{b}{a} < 1\right)^N\right] P(\Phi) \geq 1$. Below is our final derivation:

$$\begin{aligned}
 &\left[1 - P\left(\frac{b}{a} < 1\right)^N\right] P(\Phi) \geq 1 \\
 \Rightarrow &\left[1 - P\left(\frac{b}{a} < 1\right)^N\right] \geq 1 \\
 \Rightarrow &P\left(\frac{b}{a} < 1\right)^N \leq 0 \\
 \Rightarrow &P\left(\frac{b}{a} < 1\right) \leq 0 \\
 \Rightarrow &P\left(\frac{b}{a} \geq 1\right) = 1
 \end{aligned} \tag{12}$$

Thus, we have proved the proportion of correct correspondences b exceeds the threshold a . Therefore, given a , we can distinguish local regions where the inlier ratio is greater than a .

This theorem is crucial for our correspondence refinement, as it allows control over correspondences through hyperparameter a tuning. Additionally, It provides a theoretical guarantee for the quality of correspondence at each step. It is also valuable for other outlier removal methods.

6.2. Proof of Theorem 2

Theorem 2. Our generalized mutual matching (GMM) generates at least as many correct correspondences as those produced by mutual matching (MM), that is, $|\mathcal{G}_{\text{GMM}}| \geq |\mathcal{G}_{\text{MM}}|$ where $\mathcal{G}_{\text{GMM}}$ and $\mathcal{G}_{\text{MM}}$ represent the correspondences obtained from the GMM and MM methods, respectively.

Proof. According to the definition of the matching matrices $\mathbf{M}_1^{P \rightarrow Q}$ and $\mathbf{M}_1^{Q \rightarrow P}$ provided in Sec. 3.3, the mutual matching (MM) process can be expressed as follows:

$$\mathcal{G}_{\text{MM}} = \mathbb{G}(\mathbf{M}_1^{P \rightarrow Q} \odot \mathbf{M}_1^{Q \rightarrow P}) \tag{13}$$

where the function $\mathbb{G}(\mathbf{M})$ extracts correspondences from the matching matrix $\mathbf{M}$, which is defined as follows:

$$\mathbb{G}(\mathbf{M}) = \{(\mathbf{p}_i, \mathbf{q}_j) | \mathbf{M}(i, j) = 1, \mathbf{p}_i \in \mathbf{P}, \mathbf{q}_j \in \mathbf{Q}\} \tag{14}$$

In contrast, our GMM process can be expressed as:

$$\mathcal{G}_{\text{GMM}} = \mathbb{G}\left((\mathbf{M}_1^{P \rightarrow Q} \odot \mathbf{M}_2^{Q \rightarrow P}) \otimes (\mathbf{M}_1^{Q \rightarrow P} \odot \mathbf{M}_2^{P \rightarrow Q})\right) \tag{15}$$

Based on the definition and properties of function $\mathbb{G}$, we can

derive the following:

$$\begin{aligned}
& |\mathcal{G}_{GMM}| \\
&= |\mathbb{G}((\mathbf{M}_1^{\mathcal{P} \rightarrow \mathcal{Q}} \circ \mathbf{M}_2^{\mathcal{Q} \rightarrow \mathcal{P}}) \otimes (\mathbf{M}_1^{\mathcal{Q} \rightarrow \mathcal{P}} \circ \mathbf{M}_2^{\mathcal{P} \rightarrow \mathcal{Q}}))| \\
&\geq |\mathbb{G}((\mathbf{M}_1^{\mathcal{P} \rightarrow \mathcal{Q}} \circ \mathbf{M}_2^{\mathcal{Q} \rightarrow \mathcal{P}}))| \\
&\geq |\mathbb{G}((\mathbf{M}_1^{\mathcal{P} \rightarrow \mathcal{Q}} \circ \mathbf{M}_1^{\mathcal{Q} \rightarrow \mathcal{P}}))| \\
&= |\mathcal{G}_{MM}|
\end{aligned} \tag{16}$$

Thus, we complete the proof that $|\mathcal{G}_{GMM}| \geq |\mathcal{G}_{MM}|$.

6.3. Why our Regor perform well with few inliers?




Figure 8. Illustration of our progressive correspondence regenerator. The green and red lines indicate correct and incorrect correspondences at the current stage, respectively. The purple and yellow dots represent the seed correspondences (center points) and really correct corresponding points, respectively. The luminous part is the overlapping area. The dotted circles are the local corresponding regions.

The advantage of our algorithm lies in its ability to achieve robust registration even with very few inliers or under extreme outlier rates, which is unattainable for existing outlier removal methods. As shown in Fig. 8a, the initial correspondences are filled with numerous incorrect matches, with only one inlier. Methods relying on geometric consistency, such as [6, 47], would fail to extract this inlier. Moreover, it is impossible to solve the 6DoF pose with only a single correct correspondence.

In contrast, our method still performs well in such harsh situations. The reasons can be summarized in the following two aspects. **First**, the whole process is gradually generated and corrected, which does not depend on the quality of

the initial correspondence. **Second**, outliers with small errors conduce to regenerate new correct correspondences by local matching. For example, in Figure 8, our approach begins by sampling a seed correspondence (the purple point in the figure). For clarity, we illustrate a single seed point here. Although the seed correspondence is incorrect (the correct match is the yellow point in the figure), our next step, local matching, is crucial. If the seed correspondence has a relatively small distance error, it can still guide the matching of overlapping local regions. In these regions, despite the initial error, the overlap enables the generation of new correct correspondences. More importantly, performing local matching within these regions significantly reduces the search space compared to global matching across the entire point cloud, making it much easier to obtain correct correspondences.

Additionally, in the early iterations of our algorithm, we adopt a larger region radius for local matching. This increases the probability of overlap between local regions, mitigating the impact of erroneous correspondences during the initial stages. As the iterations progress, the region radius is gradually reduced to refine the matches and improve the overall quality of the correspondences. As a result, even with very few inliers, our method can utilize erroneous correspondences with small errors to generate new correct correspondences, making accurate pose estimation possible under such challenging conditions.

7. Implementation details

7.1. Point-level Pose Refinement

Pose refinement is crucial for achieving accurate pose estimation. In the final stage of correspondence regeneration, we perform pose refinement based on the established correspondences. Unlike correspondence-level pose refinement such as SC2-PCR [6] and MAC [47], our approach conducts pose refinement at the point level.

A key advantage of our approach is its ability to generate dense and accurate correspondences. However, traditional pose refinement evaluation metrics, such as inlier count (IC) [6], mean square error (MSE), truncated chamfer distance (TCD), and feature-and-spatial-consistency-constrained truncated chamfer distance (FS-TCD) [7], rely heavily on initial correspondences. Given the limited number of correct initial correspondences, these metrics fail to fully capture the benefits of our dense and accurate correspondences. To address this limitation, we propose a novel point-level evaluation metric, termed point-level truncated chamfer distance (PO-TCD):

$$\mathbf{R}^*, \mathbf{t}^* = \max_{\mathbf{R}, \mathbf{t}} \sum_{\mathbf{p}_i \in \mathbf{P}} \mathbb{1} \left(\min_{\mathbf{q}_j \in \mathbf{Q}} \|\mathbf{R}\mathbf{p}_i + \mathbf{t} - \mathbf{q}_j\| < \sigma_d \right). \tag{17}$$

It eliminates reliance on initial correspondences and instead computes the optimal truncation distance directly from dense point clouds, fully leveraging the benefits of our dense correspondences.

7.2. Hyperparameter Settings

Dataset	k_0	r_0	s_0	ω_k	ω_r	ω_s	k_{GMM}	a
3DMatch [46]	20	1	500	5	0.5	0.2	3	0.5
KITTI [16]	20	10	500	5	0.5	0.2	3	0.5

Table 7. Hyperparameter settings in different datasets.

The detailed hyperparameter settings are provided in Table 7. Due to the differing scales of indoor and outdoor datasets, we specify two distinct parameter configurations, differing only in the local region radius. Specifically, for indoor datasets, we set r_0 to 1m, while for outdoor datasets, r_0 is set to 10m. All other parameters remain consistent across both settings.

Here, we outline the implementation details of our approach. First, for local grouping, during the t th iteration, we perform random sampling with sampling quantity $s_0 \cdot (w_s)^t$ on the correspondences obtained from the $t - 1$ th iteration. Using a radius of $r_0 \cdot (r_s)^t$, we construct local point sets and standardize the number of points in each local set to $k_0 \cdot (w_k)^t$ by sampling to facilitate parallel computation. Next, we apply generalized mutual matching with $k_{GMM} = 3$ to obtain reliable correspondences. Finally, through a refinement process that combines local and global correspondence optimization, the global correspondences are updated. It is worth noting that before the first iteration of sampling, we utilize the SC2-PCR [6] module to ensure the initial quality of the correspondences. Additionally, after the final correspondence refinement stage, we perform point-level pose refinement to further enhance the results.

Throughout the entire process, point sets for each stage are stored and updated using index-based representations, significantly reducing memory overhead. Similarly, during Correspondence Merging, we employ an index-based hash $\mathbb{H}$ table to efficiently obtain unique global correspondences, ensuring both speed and memory efficiency.

8. Detailed Metrics

Inlier Number (IN). It represents the number of correct correspondences in the final correspondences, reflecting our ability to generate inliers.

Inlier Number Ratio (INR). It is the ratio of inliers IN_a in the final correspondences to the inliers IN_b in the initial correspondences, which can be calculated by:

$$INR = \frac{1}{H} \sum_{h=1}^H INR_h, \quad INR_h = \begin{cases} \frac{IN_a^h}{IN_b^h}, & IN_b^h \neq 0 \\ IN_a^h, & IN_b^h = 0 \end{cases} \quad (18)$$

where H represents the total number of point cloud pairs in the dataset, and INR_h denotes the INR for the h th point cloud pair.

Registration Recall (RR). Following [3, 6], we consider registration to be successful if the translation error (TE) and rotation error (RE) are within specified thresholds. For indoor scenes, the thresholds are 15° and 30cm, while for outdoor scenes, they are 5° and 60cm. The registration recall rate is calculated as:

$$RR = \frac{1}{H} \sum_{h=1}^H \mathbb{1}(\text{TE}_h < \sigma_d \wedge \text{RE}_h < \sigma_\theta), \quad (19)$$

where $\mathbb{1}$ represents the indicator function, and σ_d and σ_θ are the thresholds for rotation error and translation error, respectively. RE_h and TE_h denote the rotation error and translation error for the h th point cloud pair. They can be computed using the following formulas:

$$\begin{cases} \text{RE}_h = \arccos\left(\frac{\text{trace}(\hat{\mathbf{R}}_h^T \mathbf{R}_h) - 1}{2}\right) \\ \text{TE}_h = \|\hat{\mathbf{t}}_h - \mathbf{t}_h\| \end{cases} \quad (20)$$

where $\hat{\mathbf{R}}_h$, $\mathbf{R}_h$, $\hat{\mathbf{h}}_h$ and $\mathbf{h}_h$ are our predicted rotation matrix, the ground-truth rotation matrix, predicted translation vector and the ground-truth translation vector, respectively.

9. Additional Experiments

The information of all tested datasets is given in Table 8.

9.1. Results at different numbers of inliers

In Sec. 4.1, we propose the 3DMatch-EOR Benchmark, which includes two baselines: one with an outlier ratio exceeding 99% and another exceeding 90%. From these experiments, we conclude that our method demonstrates greater robustness under extreme outlier ratios. However, since these experiments are designed based on ratios, they do not verify the impact of the number of inliers. To address this, we conduct additional comparative experiments using FPFH as the feature extraction baseline, varying the number of inliers. Specifically, we designed seven scenarios with inlier counts ranging from 0~20, 20~40, 40~60, 60~80, 80~100, to 100~200, comparing our method against the outlier removal method SC2-PCR [6]. The experimental results are summarized in Table 9.

In all conditions, our method outperforms SC2-PCR [6] across all metrics. Even when the number of correct correspondences is extremely low (0~100), our method delivers excellent performance, thanks to its strong capability for generating new correspondences, as reflected in the INR and IN metrics. Furthermore, as the number of inliers decreases, feature matching becomes increasingly challenging, leading to a gradual decline in the RR of our method.

Benchmark	Indoor scene	Nuisances	Application scenario	# Matching pairs
3DMatch [46]	Indoor scene	Occlusion, real noise	Registration	1623
3DLoMatch [20]	Indoor scene	Limited overlap, occlusion, real noise	Registration	1781
3DMatch-EOR 90% (ours)	Indoor scene	High outlier ratio, occlusion, real noise	Registration	1270
3DMatch-EOR 99% (ours)	Indoor scene	Extreme high outlier ratio, occlusion, real noise	Registration	142
KITTI [16]	Outdoor scene	Clutter, occlusion, real noise	Detection, registration, segmentation	555

Table 8. Information of all tested datasets.

# Inlier number	# Pairs	Regor (ours)						Outlier removal [6]					
		RR($\uparrow$)	RE($\downarrow$)	TE($\downarrow$)	IP($\uparrow$)	INR($\uparrow$)	IN($\uparrow$)	RR($\uparrow$)	RE($\downarrow$)	TE($\downarrow$)	IP($\uparrow$)	INR($\uparrow$)	IN($\uparrow$)
0~20	55	3.87	3.57	10.36	3.59	273.01	38.44	3.64	4.82	23.9	1.49	3.86	0.02
20~40	95	36.84	3.98	9.22	29.65	1510.95	471.55	8.42	5.16	13.79	10.33	15.87	0.29
40~60	80	62.50	2.71	8.00	52.83	2536.91	1275.64	27.5	5.70	12.86	28.47	28.47	0.73
60~80	90	80.00	2.60	8.50	71.21	2434.90	1656.36	63.33	4.15	11.62	51.59	61.64	2.35
80~100	95	91.58	1.98	8.04	78.83	2288.78	2000.84	86.32	3.82	11.22	63.77	72.58	3.73
100~200	333	95.50	1.89	7.00	86.85	1625.94	2318.18	92.19	2.71	8.24	78.66	85.15	26.11

Table 9. Comparison between our Regor and outlier removal method [6] under different numbers of inliers.

However, the rate of decline for our method is significantly slower than that of SC2-PCR [6], indicating a certain degree of robustness in handling scenarios with very few inliers. However, when IN is less than 20, our method also cannot achieve robust registration.

9.2. Detailed Ablation Study




Figure 9. Ablation study on the progressive iteration.

Progressive Iteration. In Section 4.4, we conduct an ablation study to compare progressive and one-stage approaches. To further analyze the effect of iterations, we vary the number of iterations from 0 to 10 and evaluate the performance on the 3DMatch dataset. To balance registration success and speed, our study reports two metrics: registration recall (RR) and running time. The results, presented in Figure 9, show that RR increases significantly as the number of iterations rises from 1 to 4. Beyond 4th iterations, the improvement slows, eventually saturating, with occasional fluctuations. Meanwhile, runtime increases steadily with more iterations. To balance performance and efficiency, we ultimately select 4 iterations as the optimal configuration.

Prior-guided Local Grouping. We conduct a more de-

Methods	3DMatch			3DMatch-EOR		
	RR	IP	IN	RR	IP	IN
1) K nearest neighbor	88.32	82.55	2522.09	26.70	21.40	670.23
2) Radius nearest neighbor*	88.48	82.68	2532.40	26.76	21.54	666.94
3) Spectral technique [6, 25]	88.36	84.77	2501.41	26.19	24.33	630.92
4) Sampling with consistency score	88.06	84.09	2510.91	20.81	22.19	578.09
5) Random sampling*	88.48	82.68	2532.40	26.76	21.54	666.94

Table 10. Ablation study on prior-guided local grouping.

tailed ablation study on the prior-guided local grouping module, focusing on two aspects: sampling strategies and neighbor grouping strategies. For neighbor grouping, we evaluate two approaches: K-nearest neighbors (KNN) and radius-based neighbors (RNN). For sampling, we test three methods: spectral technique [25], sampling with consistency scores, and random sampling. As shown in Table 10, the RNN strategy outperforms KNN. This is because RNN defines local region sizes independently of the point cloud’s sparsity, while KNN is highly affected by variations in point cloud density. Surprisingly, as seen in Rows 3, 4, and 5 of Table 10, random sampling outperforms the other strategies, achieving the highest RR, albeit with relatively lower IP. This is because random sampling promotes a more uniform distribution of local regions across the point cloud, reducing the risk of getting trapped in local optima and improving robustness. Meanwhile, ablated models focus local regions on areas with higher-quality correspondences as determined by our method, resulting in higher IP.

Hyperparameter Setting. We conduct an ablation study on hyperparameter settings using the 3DMatch dataset. To simplify hyperparameter tuning, we design the required sampling number s_t , the local region radius r_t , and the number of points in the local region k_t to follow a geometric progression. This approach allows all parameters to be de-

Models	3DMatch			3DMatch-EOR		
	RR	IP	IN	RR	IP	IN
1) $k_0 = 50, \omega_k = 5$	88.42	80.48	2601.12	24.20	20.11	681.23
2) $k_0 = 20, \omega_k = 10$	88.48	81.82	2580.45	26.60	20.90	678.14
3) $k_0 = 20, \omega_k = 5^*$	88.48	82.68	2532.40	26.76	21.54	666.94
4) $r_0 = 2, \omega_r = 0.5$	88.42	84.09	2598.01	26.44	20.91	702.29
5) $r_0 = 0.5, \omega_r = 0.5$	87.21	80.92	2423.12	22.13	21.49	590.77
6) $r_0 = 1, \omega_r = 0.2$	87.43	82.56	2499.86	22.81	21.88	553.51
7) $r_0 = 1, \omega_r = 0.5^*$	88.48	82.68	2532.40	26.76	21.54	666.94

Table 11. Ablation study with different hyperparameters.

terminated by setting only the initial values s_0 , r_0 , k_0 , and the common ratios w_s , w_r , w_k . The results of the ablation study are presented in Table 11. As shown in Rows 2, 3, 4, and 7 of Table 11, increasing k_0 and r_0 leads to a higher number of inliers (IN), but does not necessarily result in a higher RR. Balancing accuracy and efficiency, we determine $k_0 = 20$ and $\omega_k = 5$ based on Rows 1, 2, and 3. Similarly, Rows 4, 5, 6, and 7 indicated that $r_0 = 1$ and $\omega_r = 0.5$ are the optimal radius hyperparameters.


9.3. Additional Qualitative Results

As demonstrated in Sec. 4.1 and 4.2, our Regor achieves superior quantitative performance on the 3DMatch and KITTI datasets. In this section, we provide additional visualizations to illustrate our qualitative results.

Additional Qualitative Results on 3DMatch. On the 3DMatch dataset, we compare the qualitative results of our method with those of SC2-PCR [6] and PointDSC [3]. The visualization of the correspondences is shown in Figure 10. As illustrated, when the initial correspondences are of poor quality, outlier removal methods [3, 6, 14] struggle to establish high-quality correspondences, which significantly hinders accurate registration. In contrast, our Regor progressively refines correspondences, generating dense and reliable inliers. This provides a robust foundation for accurate 6DoF pose estimation and demonstrates the adaptability of our approach to scenarios with extremely high outlier rates.

Figure 11 presents the qualitative results and registration visualizations. In some scenarios, SC2-PCR [6] and PointDSC [3] fail to achieve robust registration due to the limitations in correspondence quality. Compared to outlier removal methods [3, 6, 14], our method not only achieves robust registration but also yields lower pose errors.

Additional Qualitative Results on KITTI. On the KITTI dataset, we use FPFH [32] for feature extraction and the results of the qualitative experiments are shown in Figure 12. Our method achieves accurate and robust registration due to high-quality dense correspondences.




(a) SC2-PCR

(b) PointDSC

(c) Ours

Figure 10. Correspondence results on 3DMatch.



(a) GT

(b) SC2-PCR

(c) PointDSC

(d) Ours

Figure 11. Registration results on 3DMatch.




Figure 12. Qualitative results on KITTI odometry.