

Cross-PCR: A Robust Cross-Source Point Cloud Registration Framework

Guiyu Zhao, Zhentao Guo, Zewen Du, Hongbin Ma*

School of Automation, Beijing Institute of Technology
zhaoguiyu@bit.edu.cn, {zt_guo1230, dzw1114}@163.com, mathmhb@bit.edu.cn

Abstract

Due to the density inconsistency and distribution difference between cross-source point clouds, previous methods fail in cross-source point cloud registration. We propose a density-robust feature extraction and matching scheme to achieve robust and accurate cross-source registration. To address the density inconsistency between cross-source data, we introduce a density-robust encoder for extracting density-robust features. To tackle the issue of challenging feature matching and few correct correspondences, we adopt a loose-to-strict matching pipeline with a “loose generation, strict selection” idea. Under it, we employ a one-to-many strategy to loosely generate initial correspondences. Subsequently, high-quality correspondences are strictly selected to achieve robust registration through sparse matching and dense matching. On the challenging Kinect-LiDAR scene in the cross-source 3DCSR dataset, our method improves feature matching recall by 63.5 percentage points (pp) and registration recall by 57.6 pp. It also achieves the best performance on 3DMatch, while maintaining robustness under diverse downsampling densities.

Introduction

Cross-source point cloud registration (Huang, Mei, and Zhang 2023; Huang et al. 2021b) entails the alignment of point clouds obtained from diverse sensors into a shared coordinate system. It is the key to complete the calibration of different sensors for multi-sensor fusion. Currently, a variety of sensors and methods are employed for capturing point clouds, each offering distinct advantages in the data collection process. Consequently, harnessing the complementary strengths of different sensors for multi-sensor fusion holds great promise and finds applications in domains such as mapping, robotics, and remote sensing. In robotics, it is a cheap and feasible solution for localization. We build a map with high-precision LiDAR and then scan the environment with inexpensive sensors on the robot (Xiong et al. 2023; Ren et al. 2023). Regrettably, cross-source registration has seen limited advancements in recent years.

Cross-source point cloud registration is a notably challenging task, as it entails aligning point clouds originating from distinct sensors. These discrepancies in point cloud

collection methodologies result in significant variations in point cloud distribution, thereby introducing formidable obstacles to achieving robust registration. Therefore, cross-source point cloud registration is beset by two primary challenges (Huang, Mei, and Zhang 2023). **First**, as shown in Figure 1, there exist substantial disparities in the distribution and density between two point clouds from different sensors. For instance, LiDAR-generated point clouds manifest as sparse line-like structures, while those collected by depth cameras exhibit high density. This divergence poses a large hindrance to the extraction of similar features at corresponding points, leading to a failure in correct matching (Figure 1 right). Although DIF-PCR (Liu et al. 2023) proposes the density-robust feature, it primarily addresses the density inequality within one point cloud, rather than the density difference between two cross-source point clouds. **Second**, Different sensing mechanisms (Mallick, Das, and Majumdar 2014; Roriz, Cabral, and Gomes 2021) introduces lots of different noises and outliers into the cross-source point cloud data. These disturbances further exacerbate the difficulty of matching features. Consequently, only a limited number of accurate correspondences can be obtained, which also impedes the attainment of robust registration.

Currently, the majority of point cloud registration methods are primarily designed for homologous data. These methods encounter challenges in addressing two issues. Unfortunately, although some methods (Ao et al. 2021; Zhao et al. 2024) have good generalization ability, they perform badly in cross-source registration tasks with different densities between two point clouds. SPEAL (Xiong et al. 2023) proposes an outdoor cross-source registration method, but it faces challenges in addressing difficult matching issues in indoor scenes and fails to achieve robust registration. Due to the above two problems, it is difficult to achieve cross-source registration by previous feature-based methods. Consequently, a considerable proportion of cross-source point cloud registration algorithms rely on optimization-based methodologies (Zhao et al. 2023; Huang, Mei, and Zhang 2023). However, these techniques face hurdles in achieving robust registration across diverse scenarios.

We introduce a feature-based cross-source registration method named Cross-PCR. To deal with the challenges above, we present a density-robust feature extraction method and a loose-to-strict matching pipeline with the idea of

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

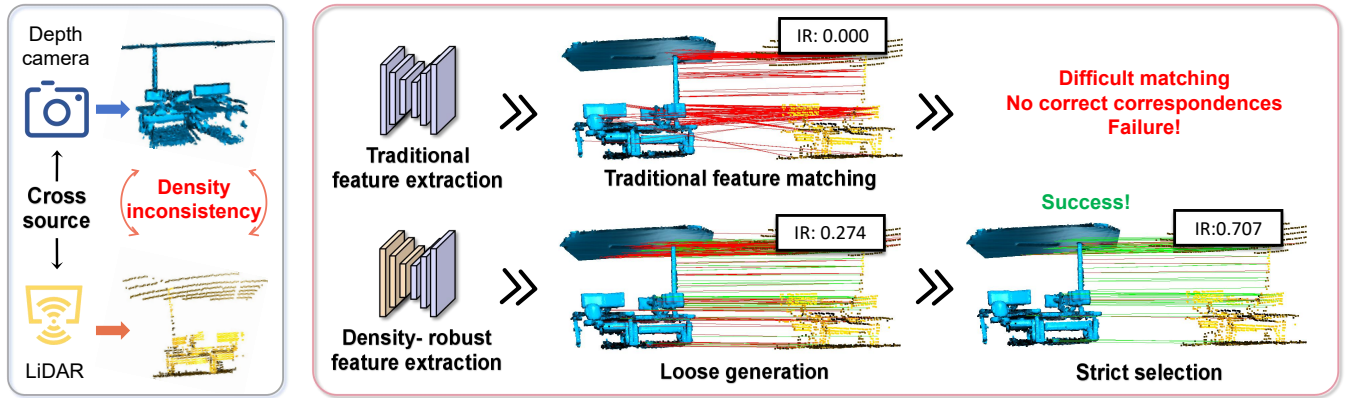


Figure 1: Existing methods often struggle with density differences and difficult matching, leading to limited correct correspondences. In contrast, our Cross-PCR introduces a density-robust feature extractor to extract density-robust features. Then, a loose-to-strict matching strategy, guided by the “loose generation, strict selection” idea, generates reliable correspondences.

“loose generation, strict selection”. **For the first challenge**, We propose a **density-robust feature extraction** network to address density inconsistencies between cross-source point clouds. Specifically, we design a multi-density fusion block to fuse the upsampled, unsampled, and downsampled point cloud features respectively. It is then integrated into each layer of the encoder to achieve deep fusion of density, capturing common structural features between point clouds of different densities. This greatly reduces the feature differences of cross-source point clouds in density. In addition, this simple and efficient design avoids large computational overhead and complex training strategy (Liu et al. 2023). **For the second one**, we devise a **loose-to-strict matching** framework involving sparse matching and dense matching. Despite extracting the density-robust feature, obtaining the most similar features at the corresponding points remains challenging in the context of cross-source registration. Our key idea is “loose generation, strict selection” which means loosely generating a large number of correspondences and strictly filtering out correct correspondences. To realize this, we propose a one-to-many matching strategy, which aims to capture as many correct correspondences as possible. Subsequently, the correspondences are systematically pruned, facilitating the completion of sparse matching through a robust consistency-based correspondence refinement mechanism. Despite these efforts, results after sparse matching may still contain spurious correspondences. We propose a prior-guided dense matching approach, which globally filters dense correspondences in the feature space. It finally results in achieving a more robust registration.

We evaluate our Cross-PCR on both cross-source and same-source datasets. On the 3DCSR dataset (Huang et al. 2021b), we achieve the best performance with great improvement. Our method successfully achieves the challenging registration from depth camera to LiDAR, with a 57.6 percentage point (pp) improvement in registration recall (RR) and 63.5 pp in feature matching recall (FMR). Notably, we also achieve the best results on 3DMatch (Zeng et al. 2017), obtaining a remarkable 94.5% RR and 88.7%

IR. Our main contributions are summarized as:

- We introduce a novel loose-to-strict matching pipeline with the idea of “loose generation, strict selection”, solving the problem of difficult feature matching in cross-source registration, even with an extremely low IR.
- We propose a density-robust feature extractor to extract the density-robust features, addressing differences in density across cross-source point clouds.
- Our method not only achieves robust cross-source registration, but it also exhibits effectiveness in same-source registration. Moreover, it is the first to achieve robust registration on the indoor Kinect-LiDAR benchmark.

Related Work

Same-Source Registration. Recently, same-source point cloud registration has advanced rapidly. Early methods primarily relied on ICP optimization (Segal, Haehnel, and Thrun 2009; Sharp, Lee, and Wehe 2002; Yang et al. 2015; Yang, Li, and Jia 2013) and feature matching with hand-crafted features (Rusu, Blodow, and Beetz 2009; Rusu et al. 2008; Salti, Tombari, and Di Stefano 2014). Recent advancements, however, have shown that learning-based methods (Zeng et al. 2017; Ao et al. 2021; Choy, Park, and Koltun 2019; Huang et al. 2021a; Yu et al. 2021; Qin et al. 2022) offer more accurate and robust registration. Among these, patch-based methods (Poiesi and Boscaini 2022; Ao et al. 2021, 2023) demonstrate strong generalization, while fragment-based methods (Choy, Park, and Koltun 2019; Huang et al. 2021a; Yu et al. 2021) provide better efficiency. A key trend is the adoption of coarse-to-fine strategies (Yu et al. 2021; Yang et al. 2022; Lin et al. 2023), improving registration accuracy. Transformer-based methods (Qin et al. 2022; Jin et al. 2024) have become state-of-the-art, leveraging geometric information (Qin et al. 2022; Yu et al. 2023) and positional encoding (Qin et al. 2022; Yang et al. 2022; Li and Harada 2022) to achieve outstanding results.

Cross-Source Registration. In contrast to same-source point cloud registration, the development of cross-source registration (Huang, Mei, and Zhang 2023; Zhao et al. 2023;

Huang et al. 2017) has been relatively slow, primarily due to variations in point cloud distribution and density. The challenge of extracting similar features at corresponding points has limited feature-based methods, leading to a reliance on optimization (Yang, Li, and Jia 2013; Tazir et al. 2018; Zhao et al. 2023) and model-based methods (Huang et al. 2017, 2019; Ling and Qin 2022). Huang et al. (Huang et al. 2017) employ a coarse-to-fine strategy using ESF descriptors for cross-source registration. Zhao et al. (Zhao et al. 2023) achieve superior results through fuzzy clustering (Gath and Geva 1989) and ICP (Segal, Haehnel, and Thrun 2009). SPEAL (Xiong et al. 2023) uses skeleton point features for outdoor cross-source registration but performs poorly with LiDAR and depth camera data. Due to the lack of large cross-source indoor datasets, there are currently no effective learning-based methods (Huang, Mei, and Zhang 2020) for achieving robust indoor cross-source registration.

Method

Problem Statement Cross-source point cloud registration aims to determine a rigid transformation $\mathbf{T} = \{\mathbf{R}, \mathbf{t}\}$, where $\mathbf{R} \in SO(3)$ represents a rotation matrix, and $\mathbf{t} \in \mathbb{R}^3$ represents a translation vector. The task is formulated as:

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

where \mathcal{G} represents the set of correspondences, and $g_i = (\mathbf{p}_{x_i}, \mathbf{q}_{y_i})$ denotes a pair of corresponding points. As the correspondences are initially unknown, it is imperative to establish these correspondences through feature extraction and matching before solving the optimization problem.

Density-Robust Feature Extraction

Density-Robust Encoder. Previously, the feature extraction network, employing KPConv-FPN (Thomas et al. 2019; Lin et al. 2017) as the backbone, leverages the advantageous characteristics of the FPN structure to ensure robust scale invariance in the features. However, in cross-source registration, the primary challenge lies in the density inconsistency between two point clouds. Unfortunately, the KPConv-FPN network is incapable of addressing this challenge, ultimately resulting in the failure to achieve robust cross-source registration. Therefore, we adopt a multi-density fusion block to devise a density-robust encoder (DRE), facilitating the capture of density-robust features.

Taking point cloud \mathbf{P} as an example, we get different levels of point clouds $\{\mathbf{P}, \tilde{\mathbf{P}}, \bar{\mathbf{P}}, \hat{\mathbf{P}}\}$ under the sampling strategy of KPConv (Thomas et al. 2019), where $\tilde{\mathbf{P}}$ and $\hat{\mathbf{P}}$ is the dense point cloud and the sparse point cloud, respectively. Then, as shown in Figure 2 (b), our density-robust encoder achieves hierarchical density-robust feature extraction on different levels. It is constructed by one KPConv block and three multi-density fusion blocks (MDB). For simplicity, we denote the l -level input feature and point cloud as \mathbf{F}^l and \mathbf{P}^l . In a MDB, by employing the farthest point sampling (FPS) (Qi et al. 2017) and upsampling method (Qi et al. 2017) on \mathbf{F}^l and \mathbf{P}^l , we generate three point clouds

$\{\mathbf{P}_{d_1}^l, \mathbf{P}_{d_2}^l, \mathbf{P}_{d_3}^l\}$ of varying densities along with their corresponding features $\{\mathbf{F}_{d_1}^l, \mathbf{F}_{d_2}^l, \mathbf{F}_{d_3}^l\}$ (see Figure 2 (c)),

$$\begin{aligned} \mathbf{F}_{d_1}^l &= \Psi(\text{FPS}(\Psi(\mathbf{F}^l; \mathbf{P}^l))), \\ \mathbf{F}_{d_2}^l &= \Psi^2(\mathbf{F}^l; \mathbf{P}^l), \\ \mathbf{F}_{d_3}^l &= \Psi(\text{UP}(\Psi(\mathbf{F}^l; \mathbf{P}^l))), \end{aligned} \quad (2)$$

where $\Psi(\mathbf{F}^l; \mathbf{P}^l)$ represents performing KPConv operator on \mathbf{F}^l within \mathbf{P}^l . FPS(\cdot) and UP(\cdot) denote downsampling and upsampling the point features, respectively. The feature fusion of different densities is accomplished through multi-layer perceptron (MLP), and the intermediate feature \mathbf{F}_m^l after fusion is obtained:

$$\mathbf{F}_m^l = \text{MLP}(\text{Cat}[\mathbf{F}_{d_1}^l, \mathbf{F}_{d_2}^l, \mathbf{F}_{d_3}^l]), \quad (3)$$

where concatenation operation and MLP are denoted as Cat and MLP(\cdot), respectively. Finally, the input feature \mathbf{F}^{l+1} of the next block is calculated as $\mathbf{F}^{l+1} = \Psi(\mathbf{F}_m^l; \mathbf{P}^{l+1})$. By our density-robust encoder, we extract the density-robust features $\mathbf{F}^{\hat{\mathbf{P}}}, \mathbf{F}^{\hat{\mathbf{Q}}} \in \mathbb{R}^{|\hat{\mathbf{P}}| \times d}$ of sparse points $\hat{\mathbf{P}}$ and $\hat{\mathbf{Q}}$. And features $\mathbf{F}^{\tilde{\mathbf{P}}}$ and $\mathbf{F}^{\tilde{\mathbf{Q}}}$ of the dense points $\tilde{\mathbf{P}}$ and $\tilde{\mathbf{Q}}$ is obtained by KPConv decoder (Bai et al. 2020).

Attention Backbone. We refine the sparse features $\mathbf{F}^{\hat{\mathbf{P}}}$ and $\mathbf{F}^{\hat{\mathbf{Q}}}$ by employing the attention backbone (Vaswani et al. 2017; Qin et al. 2022). Our approach effectively captures long-distance dependencies through the self-attention and cross-attention mechanisms. The refined sparse features are denoted as $\tilde{\mathbf{F}}^{\hat{\mathbf{P}}} \in \mathbb{R}^{|\hat{\mathbf{P}}| \times \tilde{d}}$ and $\tilde{\mathbf{F}}^{\hat{\mathbf{Q}}} \in \mathbb{R}^{|\hat{\mathbf{Q}}| \times \tilde{d}}$.

Sparse Matching

Despite our efforts to extract density-robust features, acquiring easily-matched features at corresponding points remains challenging. Figure 1 shows the differences between point clouds from depth cameras and LiDAR. These disparities hinder the extraction of features that satisfy the nearest-neighbor relationship at corresponding points. Consequently, traditional feature matching in cross-source registration often results in few or no correct correspondences, leading to registration failures. To address this, we introduce the ‘‘loose generation, strict selection’’ strategy. We propose a one-to-many correspondences generation method to produce broader and looser correspondences, maximizing the potential correct correspondences. Then, our correspondences filtering module selects the correct correspondences. **One-to-Many Correspondences Generation.** First, we normalize the sparse features $\tilde{\mathbf{F}}^{\hat{\mathbf{P}}}$ and $\tilde{\mathbf{F}}^{\hat{\mathbf{Q}}}$, followed by calculating the feature similarity matrix $\mathbf{D}^{\mathbf{F}} = \text{Norm}(\tilde{\mathbf{F}}^{\hat{\mathbf{P}}}) \cdot \text{Norm}(\tilde{\mathbf{F}}^{\hat{\mathbf{Q}}})^\top$ where Norm(\cdot) is a L2-normalization operation. Utilizing the one-to-many strategy, for each source sparse point $\hat{\mathbf{p}}_i$, we search the k nearest neighbors within point cloud $\hat{\mathbf{Q}}$ in the feature space, through Top-k selection:

$$\hat{\mathcal{G}}_i = \left\{ (\hat{\mathbf{p}}_i, \hat{\mathbf{q}}_{m_i}) \mid m_i \in \underset{j \in [1, |\hat{\mathbf{Q}}|]}{\text{Topk}} (-\mathbf{D}_{i,j}^{\mathbf{F}}) \right\}, \quad (4)$$

where Topk(\cdot) is the operation that finds the index of k maximum values. Finally, the correspondences $\hat{\mathcal{G}}_i$ of each

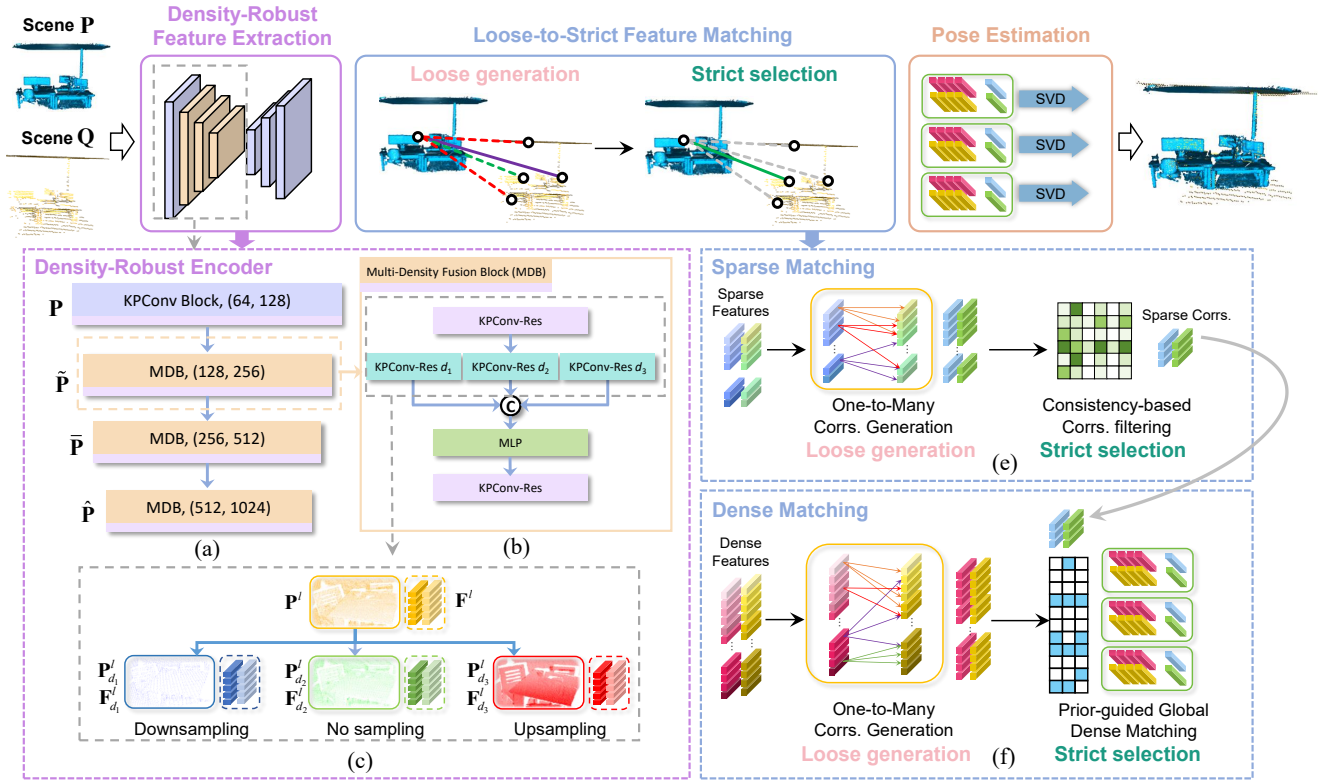


Figure 2: Our Cross-PCR consists of three parts: density-robust feature extraction, loose-to-strict feature matching, and pose estimation. First, in density-robust feature extraction, a pair of cross-source point clouds \mathbf{P} and \mathbf{Q} are fed into a density-robust encoder to obtain the density-robust features $\tilde{\mathbf{F}}^{\mathbf{P}}$ and $\tilde{\mathbf{F}}^{\mathbf{Q}}$. After that, in loose-to-strict feature matching, we propose a two-stage matching, which follows a “loose generation, strict selection” strategy to generate m sets of robust and reliable correspondences $\hat{\mathcal{G}}'_i$. Finally, in pose estimation, we propose a novel hypothesis selection method to select the best transformation \mathbf{T}^* .

point $\hat{\mathbf{p}}_i$ is combined into the initial sparse correspondences $\hat{\mathcal{G}} = \bigcup_{i=1}^n \hat{\mathcal{G}}_i$ where n is the number of points in $\hat{\mathbf{P}}$.

Consistency-based Correspondences Filtering. Though our one-to-many correspondences generation strategy effectively captures a greater number of correct correspondences, it also introduces a substantial number of outliers due to its looser conditions, resulting in a lower inlier ratio. However, in cases where the inlier ratio is extremely low, the estimators (Fischler and Bolles 1981; Besl and McKay 1992) become ineffective. Consequently, a strong correspondences filtering module is essential.

To improve efficiency, we use the spectral matching technique (Leordeanu and Hebert 2005) to screen key sparse correspondences $\hat{\mathcal{G}}^k$ on $\hat{\mathcal{G}}$ where $\hat{\mathcal{G}}^k \subseteq \hat{\mathcal{G}}$. We compute translation and rotation invariant measurements (TRIMs) (Yang, Shi, and Carlone 2020) to assess the consistency distance d_{ij} between the two correspondences $\hat{g}_i^k \in \hat{\mathcal{G}}^k$ and $\hat{g}_j \in \hat{\mathcal{G}}$

$$d_{ij} = \left| \|\hat{\mathbf{p}}_i - \hat{\mathbf{p}}_j\|_2 - \|\hat{\mathbf{q}}_i - \hat{\mathbf{q}}_j\|_2 \right| \quad (5)$$

where $\hat{g}_i^k = (\hat{\mathbf{p}}_i, \hat{\mathbf{q}}_i)$ and $\hat{g}_j = (\hat{\mathbf{p}}_j, \hat{\mathbf{q}}_j)$. A smaller distance d_{ij} indicates that the correspondence \hat{g}_i^k and \hat{g}_j satisfy a stronger consistency. Given a threshold σ_d , we get the consistency score between \hat{g}_i^k and \hat{g}_j , denoted as $s_{ij} =$

$\mathbb{1}(d_{ij} \leq \sigma_d)$ where $\mathbb{1}(\cdot)$ is the indicator function and consistency score matrix is denoted as $\mathbf{S}_k = [s_{ij}] \in \mathbb{R}^{|\hat{\mathcal{G}}^k| \times |\hat{\mathcal{G}}|}$. In the same way, we obtain consistency score matrix $\mathbf{S}_k \in \mathbb{R}^{|\hat{\mathcal{G}}^k| \times |\hat{\mathcal{G}}^k|}$ between key sparse correspondences $\hat{\mathcal{G}}^k$.

To enhance the robustness of correspondence selection, inspired by SC^2 measure (Chen et al. 2022), we calculate the weighted second-order consistency matrix $\mathbf{S}_s^* = \mathbf{S}_s \odot (\mathbf{S}_k \mathbf{W}_k \mathbf{S}_s)$ where operator \odot represents the element-wise product. The weight w_{ij} in matrix $\mathbf{W}_k \in \mathbb{R}^{|\hat{\mathcal{G}}^k| \times |\hat{\mathcal{G}}^k|}$, is defined as $w_{ij} = \exp(-\frac{d_{ij}^2}{2\sigma_d^2})$. The element s_{ij}^* of matrix \mathbf{S}_s^* represents the number of correspondences $\hat{\mathcal{G}}^k$ that simultaneously satisfy consistency with correspondence \hat{g}_i^k and \hat{g}_j . Then, we employ Top-k selection to choose k sparse correspondences $\hat{\mathcal{G}}'_i$ with the highest consistency scores:

$$\hat{\mathcal{G}}'_i = \left\{ (\hat{\mathbf{p}}_i, \hat{\mathbf{q}}_i) \mid l_i \in \text{Topk} \left((\mathbf{S}_s^*)_{ij}, (\hat{\mathbf{p}}_i, \hat{\mathbf{q}}_i) \in \hat{\mathcal{G}} \right), j \in [1, |\hat{\mathcal{G}}|] \right\}. \quad (6)$$

By merging each set of filtered correspondences, we get the refined sparse correspondences $\hat{\mathcal{G}}' = \bigcup_{i=1}^{|\hat{\mathcal{G}}^k|} \hat{\mathcal{G}}'_i$.

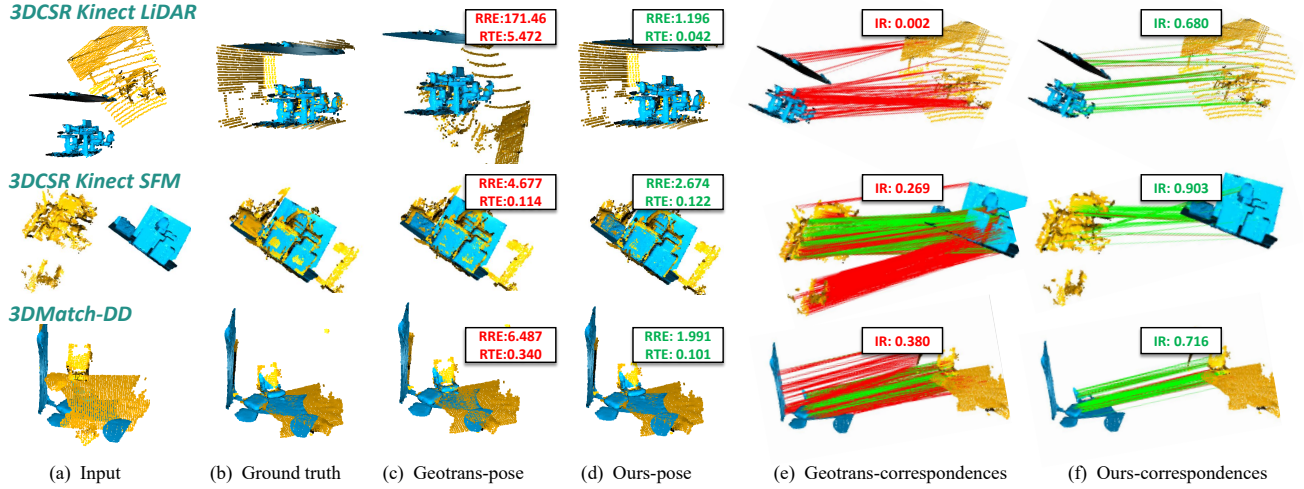


Figure 3: Correspondences and registration results compared with GeoTrans (Qin et al. 2022)

Dense Matching

Despite the refined sparse correspondences obtained through sparse matching, their limited number is insufficient for robust registration. However, their high quality allows us to use them as priors for guiding dense correspondences. Thus, we introduce prior-guided global dense matching based on these sparse correspondences. Notably, some incorrect correspondences may still exist in the refined sparse matches. If using coarse-to-fine methods (Qin et al. 2022; Yu et al. 2021), incorrect sparse matches can misguide local dense points via optimal transport (Sarlin et al. 2020; Yu et al. 2021), causing widespread errors in dense correspondences. To address this, we propose a more robust global-aware dense matching technique that transcends local matching constraints. Unlike coarse-to-fine strategies based on spatial neighbors, we adopt a global partitioning approach based on consistency.

Our general idea is to conduct a global search for every sparse correspondence by leveraging spatial consistency and aggregating dense correspondences that meet the consistency criteria. Each sparse correspondence \hat{g}_i in $\hat{\mathcal{G}}$ yields its set of dense correspondences $\tilde{\mathcal{G}}'_i$ by Top-k selection.

Prior-guided Global Dense Matching. Similar to sparse matching, we perform one-to-many correspondence generation on dense point features $\mathbf{F}^{\hat{\mathcal{P}}}$ and $\mathbf{F}^{\hat{\mathcal{Q}}}$, to establish loose dense correspondences $\tilde{\mathcal{G}}$. In the same way as Eq. 5, we calculate the consistency matrix $\mathbf{S}'_s \in \mathbb{R}^{|\hat{\mathcal{G}}'| \times |\tilde{\mathcal{G}}'|}$ of the refined sparse correspondences $\hat{\mathcal{G}}'$. Then, the sparse-to-dense consistency matrix $\mathbf{S}_{s2d} \in \mathbb{R}^{|\hat{\mathcal{G}}'| \times |\tilde{\mathcal{G}}'|}$ between sparse correspondences $\hat{\mathcal{G}}'$ and dense correspondences $\tilde{\mathcal{G}}$ is calculated. The consistency score s_{ij} in matrix \mathbf{S}_{s2d} is defined as:

$$s_{ij} = \mathbb{1} \left(\left| \|\hat{\mathbf{p}}_i - \tilde{\mathbf{p}}_j\|_2 - \|\hat{\mathbf{q}}_i - \tilde{\mathbf{q}}_j\|_2 \right| \leq \sigma_d \right), \quad (7)$$

where $(\hat{\mathbf{p}}_i, \hat{\mathbf{q}}_i) \in \hat{\mathcal{G}}'$, $(\tilde{\mathbf{p}}_j, \tilde{\mathbf{q}}_j) \in \tilde{\mathcal{G}}$. Then, we calculate the sparse-to-dense second-order consistency matrix \mathbf{S}_{s2d}^*

$$\mathbf{S}_{s2d}^* = \mathbf{S}_{s2d} \odot (\mathbf{S}'_s \mathbf{W}'_s \mathbf{S}_{s2d}) \quad (8)$$

where $\mathbf{W}'_s \in \mathbb{R}^{|\hat{\mathcal{G}}'| \times |\tilde{\mathcal{G}}'|}$ is the weight matrix between refined sparse correspondences $\hat{\mathcal{G}}'$, and its calculation is identical to that of \mathbf{W}_s . Top-k is used to select the top k dense correspondences $\tilde{\mathcal{G}}'_i \subseteq \tilde{\mathcal{G}}$ for each sparse correspondence \hat{g}_i :

$$\tilde{\mathcal{G}}'_i = \left\{ (\tilde{\mathbf{p}}_{l_i}, \tilde{\mathbf{q}}_{l_i}) \mid l_i \in \text{Topk} \left((\mathbf{S}_{s2d}^*)_{ij}, (\tilde{\mathbf{p}}_{l_i}, \tilde{\mathbf{q}}_{l_i}) \in \tilde{\mathcal{G}} \right) \right\} \quad (9)$$

This results in m groups of dense correspondences, completing the dense matching. Finally, m transformations $\mathbf{T}_i = \{\mathbf{R}_i, \mathbf{t}_i\}$ are obtained by weighted SVD (Besl and McKay 1992) where the weight is the consistency score within each group of dense correspondences $\tilde{\mathcal{G}}_i$. The details of weighted SVD are provided in the Appendix.

Hypothesis Selection. Due to the low inlier ratio of the initial correspondences generated from the cross-source point clouds, it may not be feasible to accurately determine the optimal transformation by utilizing initial correspondences for hypothesis selection. Consequently, different from the previous method (Chen et al. 2022; Qin et al. 2022), we choose to directly compute the chamfer-like truncated distance between the pair of sparse point clouds $\hat{\mathcal{P}}$ and $\hat{\mathcal{Q}}$. The selection of optimal transformation \mathbf{T}^* is completed by solving the following optimization problem:

$$\mathbf{R}^*, \mathbf{t}^* = \max_{\mathbf{R}_i, \mathbf{t}_i} \sum_{\hat{\mathbf{p}}_n \in \hat{\mathcal{P}}} \mathbb{1} \left(\min_{\hat{\mathbf{q}}_m \in \hat{\mathcal{Q}}} \|\hat{\mathbf{p}}_n - \hat{\mathbf{q}}_m\| < \tau_0 \right), \quad (10)$$

where τ_0 is a distance threshold (*i.e.* 0.1m).

Loss Function The loss function $\mathcal{L} = \mathcal{L}_s + \mathcal{L}_d$ is composed of sparse matching loss and dense matching loss. To improve robustness to low overlap, the sparse matching loss uses the overlap-aware circle loss (Qin et al. 2022). For efficiency, so we only use circle loss (Sun et al. 2020) to supervise metric learning of dense point features.

Experiments

We evaluate our Cross-PCR on both cross-source dataset 3DCSR (Huang et al. 2021b) and same-source dataset

Method	RR(\uparrow)	IR(\uparrow)	FMR(\uparrow)	RE(\downarrow)	TE(\downarrow)
<i>Kinect_sfm</i>					
GICP	12.5	-	-	1.90	0.03
JRMPC	0	3.3	14.8	-	-
GCTR ²	15.2	10.9	46.8	2.99	0.10
GCC ²	81.3	-	-	2.09	0.06
SpinNet	84.5	54.8	93.8	2.64	0.08
Predator	71.8	31.8	81.3	3.93	0.11
CoFiNet	37.5	25.3	90.6	4.19	0.10
GeoTransformer	84.4	40.2	90.6	1.95	0.06
Cross-PCR (<i>ours</i>)	93.8	89.4	96.9	2.20	0.06
<i>Kinect_lidar</i>					
GICP	0.6	-	-	12.8	0.24
JRMPC	0	0.5	3.1	-	-
GCTR ²	0	0.1	2.1	-	-
GCC ²	3.1	-	-	3.64	0.15
SpinNet	5.2	1.2	7.1	6.93	0.15
Predator	5.2	0.5	1.3	4.63	0.21
CoFiNet	5.2	1.2	7.1	4.65	0.15
GeoTransformer	9.1	3.4	9.7	4.60	0.17
Cross-PCR (<i>ours</i>)	66.7	35.5	73.2	2.83	0.14

Table 1: Quantitative results on 3DCSR.

3DMatch (Zeng et al. 2017). We also perform a robustness test to highlight the robustness to difference in density.

Following (Qin et al. 2022), we use five metrics to evaluate our performance: (1) *Inlier Ratio* (IR), represents the fraction of correspondences whose residuals is less than 0.1m. (2) *Feature Matching Recall* (FMR), represents the proportion of point cloud pairs whose IR is greater than 5% in the total point cloud pairs. (3) *Registration Recall* (RR)¹, represents the proportion of correctly registered point cloud pairs to the total point cloud pairs. (4) *Rotation Error* (RE) and *Translation Error* (TE), indicate registration errors.

Evaluation on Cross-Source Dataset

Dataset. 3DCSR dataset (Huang et al. 2021b) is an indoor cross-source dataset with two cross-source categories, containing 202 cross-source scenarios. The first category is *Kinect_sfm*, which consists of the point clouds collected by the *Kinect camera* and the point clouds generated by the *structure from motion* (SFM). The second one is *Kinect-LiDAR*, which consists of the point clouds collected by *LiDAR device* and the point clouds collected by *Kinect*.

***Kinect-SFM* results.** We conduct a comprehensive evaluation between our Cross-PCR and cross-source registration algorithms: GICP (Segal, Haehnel, and Thrun 2009), JRMPC (Huang et al. 2017), GCTR (Huang et al. 2019), GCC (Zhao et al. 2023) and learning-based same-source registration methods: SpinNet (Ao et al. 2021), Predator (Huang et al. 2021a), CoFiNet (Yu et al. 2021), GeoTransformer (Qin et al. 2022). The results on *Kinect-SFM* are shown in Table 1. We evaluate the registration results through RR, TE, and RE. Additionally, to underscore the superior feature matching capability of our method, we eval-

¹Following (Huang et al. 2021a), for same-source registration, RMSE<0.2m is considered as successful registration; Following (Zhao et al. 2023), for cross-source registration, RE<15° and TE<30cm is considered as successful registration.

²As the codes are not available, we reproduce results ourselves.

# Samples	3DMatch					3DLoMatch				
	5000	2500	1000	500	250	5000	2500	1000	500	250
<i>Feature Matching Recall (%)</i> \uparrow										
SpinNet	97.6	97.2	96.8	95.5	94.3	75.3	74.9	72.5	70.0	63.6
Predator	96.6	96.6	96.5	96.3	96.5	78.6	77.4	76.3	75.7	75.3
YOHO	98.2	97.6	97.5	97.7	96.0	79.4	78.1	76.3	73.8	69.1
CoFiNet	98.1	98.3	98.1	98.2	98.3	83.1	83.5	83.3	83.1	82.6
GeoTransformer	97.9	97.9	97.9	97.9	97.6	88.3	88.6	88.8	88.6	88.3
RoITr	98.0	98.0	97.9	98.0	97.9	89.6	89.6	89.5	89.4	89.3
Cross-PCR (<i>ours</i>)	97.9	97.7	97.9	97.7	97.7	83.5	83.3	83.3	83.1	83.4
<i>Inlier Ratio (%)</i> \uparrow										
SpinNet	47.5	44.7	39.4	33.9	27.6	20.5	19.0	16.3	13.8	11.1
Predator	58.0	58.4	57.1	54.1	49.3	26.7	28.1	28.3	27.5	25.8
YOHO	64.4	60.7	55.7	46.4	41.2	25.9	23.3	22.6	18.2	15.0
CoFiNet	49.8	51.2	51.9	52.2	52.2	24.4	25.9	26.7	26.8	26.9
GeoTransformer	71.9	75.2	76.0	82.2	85.1	43.5	45.3	46.2	52.9	57.7
RoITr	82.6	82.8	83.0	83.0	83.0	54.3	54.6	55.1	55.2	55.3
Cross-PCR (<i>ours</i>)	88.7	88.7	88.7	88.7	88.7	65.9	65.9	65.9	65.9	65.9
<i>Registration Recall (%)</i> \uparrow										
SpinNet	88.6	86.6	85.5	83.5	70.2	59.8	54.9	48.3	39.8	26.8
Predator	89.0	89.9	90.6	88.5	86.6	59.8	61.2	62.4	60.8	58.1
YOHO	90.8	90.3	89.1	88.6	84.5	65.2	65.5	62.4	56.5	48.0
CoFiNet	89.1	88.9	88.4	87.4	87.0	67.5	66.2	64.2	63.1	61.0
GeoTransformer	92.0	91.8	91.8	91.4	91.2	75.0	74.8	74.2	74.1	73.5
RoITr	91.9	91.7	91.8	91.4	91.0	74.7	74.8	74.8	74.2	73.6
Cross-PCR (<i>ours</i>)	94.5	94.2	94.2	94.3	94.0	73.7	73.9	74.1	74.2	74.1

Table 2: Quantitative results on 3DMatch and 3DLoMatch.

uate the correspondences results using IR and FMR. On the *Kinect-SFM* benchmark, our approach performs best on RR, IR, and FMR. Notably, benefiting from our loose-to-strict matching, our method demonstrates remarkable efficacy in feature matching, achieving a 34.6 pp improvement on IR, surpassing other algorithms by a considerable margin. Qualitative results are shown in Figure 3.

***Kinect-LiDAR* results.** We also assess the outcomes on *Kinect-LiDAR* in the same way, with the results presented in the right part of Table 1. Notably, cross-source data in the *Kinect-LiDAR* benchmark exhibits greater dissimilarity than that in *Kinect-SFM*. As indicated in Table 1, almost all methods fail to achieve correct registration on *Kinect-LiDAR* benchmark. Benefiting from the idea of “loose generation, strict selection” and density-robust features, our method is the first to achieve robust registration on the *Kinect-LiDAR* benchmark, with the lowest RE and TE. it achieves a 66.7% RR, 35.5% IR, and 73.2% FMR, marking respective improvements of 57.6 pp, 32.1 pp, and 63.5 pp compared to the previous SOTA method (Qin et al. 2022). Our method exhibits a remarkable generalization ability across diverse modal data, ensuring robust cross-source registration.

Evaluation on Same-Source Dataset

Dataset. The 3DMatch dataset (Zeng et al. 2017) is a large indoor dataset containing 62 scenarios, of which 46 are used for training, 8 for validation, and 8 for testing. Following (Huang et al. 2021a), point cloud pairs with overlap > 30% are split as 3DMatch, and those with 10% ~ 30% overlap are split as 3DLoMatch (Huang et al. 2021a).

Correspondence results. Following (Qin et al. 2022; Yu et al. 2023), we conduct experiments to assess the robustness of samples. Comparative evaluations are made with the state-of-the-art methods: SpinNet (Ao et al. 2021), Predator (Huang et al. 2021a), YOHO (Wang et al. 2022), CoFiNet (Yu et al. 2021), GeoTrans (Qin et al. 2022),

Benchmark	3DMatch									3DLoMatch								
	-			0.05			0.1			-			0.05			0.1		
Voxel length																		
Metrics	RR	IR	FMR	RR	IR	FMR	RR	IR	FMR	RR	IR	FMR	RR	IR	FMR	RR	IR	FMR
Predator	89.0	58.0	96.6	2.1	0.4	1.1	2.1	0.1	0.2	59.8	26.7	78.6	0.2	0	0	2.1	0.1	0.2
CoFiNet	89.1	49.8	98.1	2.0	0.4	1.5	0.4	0.1	0	67.5	24.4	83.1	0.1	0	0	0.1	0	0
GeoTrans	92.0	71.9	97.9	91.1	63.0	97.6	4.3	2.3	14.3	75.0	43.5	88.3	66.9	33.5	83.3	1.7	1.2	5.6
Cross-PCR (<i>ours</i>)	94.5	88.7	97.7	91.5	84.4	97.2	91.0	81.9	97.2	73.7	65.9	83.5	69.8	61.4	80.2	65.8	57.2	77.4

Table 3: Quantitative results on the 3DMatch-DD benchmark.

Voxel length	-			0.05			0.1		
Metrics	RR	IR	FMR	RR	IR	FMR	RR	IR	FMR
KPConv	92.2	87.4	96.5	90.7	85.5	96.2	57.6	40.9	69.2
DRE*	94.5	88.7	97.7	91.5	84.4	97.2	91.0	81.9	97.2

Table 4: Ablation study with DRE on 3DMatch-DD.

RoTr (Yu et al. 2023). Correspondence results are evaluated using FMR and IR, as shown in Table 2 (top and middle). For IR, our method exhibits significant enhancement, achieving 6.1 pp and 11.6 pp improvements over the state-of-the-art RoTr (Yu et al. 2023) on 3DMatch and 3DLoMatch, respectively. For FMR, our method also yields commendable results, but it is slightly worse than RoTr (Yu et al. 2023). The drop arises from the removal of ambiguous inliers through loose-to-strict matching, resulting in the correspondences with low IR (<5%) in certain scenarios.

Registration results. Following (Qin et al. 2022), we employ the RR to assess our registration results. For other methods, we employ a 50k RANSAC to estimate transformation. Compared to other methods, Cross-PCR yields the best outcomes in most cases, surpassing RoTr (Yu et al. 2023) by 2.6 pp. Our method also exhibits robustness to sampling, with only a 0.5 pp decrease in RR from 5000 to 250 samples. Moreover, instead of using time-expensive RANSAC, our method only uses weight SVD (Besl and McKay 1992) which has a low time cost. This underscores the high quality of the correspondences, enabling robust registration without relying on a robust estimator.

Robustness to Density Inconsistency

Experiments settings. Following (Huang et al. 2021a), we preprocess the original 3DMatch dataset. Then, we only downsample the target point cloud by using *voxel down-sample* to simulate the large density difference between the cross-source point cloud pair. By setting no voxel downsampling, voxel downsampling with a 0.05m side length, and with a 0.1m side length, we conduct three experiments.

Quantitative results. Table 3 presents quantitative results. On the 3DMatch-DD benchmark, our method consistently outperforms across diverse downsampling levels compared to other baselines (Huang et al. 2021a; Yu et al. 2021; Qin et al. 2022). In the challenging condition of voxel downsampling with a 0.1m side length, our method still achieves correct registration in certain scenarios. It only drops 3.5 pp and 7.9 pp in RR compared to the original, far less than GeoTrans (Qin et al. 2022), which drops 86.7 pp and 73.3 pp.

No.	Methods	<i>Kinect-sfm</i>			<i>Kinect-lidar</i>		
		RR	IR	FMR	RR	IR	FMR
1)	KPConv	87.5	85.2	93.8	57.7	30.1	61.7
2)	DRE*	93.8	89.4	96.9	66.7	35.5	73.2
3)	W/o loose generation	84.5	84.4	93.8	44.2	26.4	54.5
4)	W/ loose generation*	93.8	89.4	96.9	66.7	35.5	73.2
5)	W/o strict selection	81.1	77.9	90.6	27.3	17.5	34.4
6)	W/ strict selection*	93.8	89.4	96.9	66.7	35.5	73.2
7)	One-stage matching	68.8	29.5	93.8	56.2	10.3	50.9
8)	Two-stage matching*	93.8	89.4	96.9	66.7	35.5	73.2

Table 5: Ablation study on the 3DCSR dataset.

Ablation Study

Density-robust encoder. We substitute our density-robust encoder with KPConv-FPN (Qin et al. 2022). The results are shown in Table 5, entries (1) and (2). For the *Kinect-lidar* benchmark, the ablation model performance decreases significantly due to large density differences. Further ablation on the 3DMatch-DD benchmark (Table 4) shows that our method has much less performance degradation across various sampling levels compared to the ablation model.

W/o loose generation. We replace our one-to-many correspondences generation with traditional one-to-one matching. As shown in Table 5 (3) and (4), there is a substantial improvement on *Kinect-lidar*, with “loose generation”. This shows the technique contributes to achieving robust registration even when matching is extremely challenging.

W/o strict selection. As shown in Table 5 (5) and (6), we remove the consistency-based correspondences refinement module. The ablated model achieves low RR and IR, which demonstrates this module is critical to achieving robust registration. Entries (3) and (5) together prove the superiority of our “loose generation, strict selection” idea.

One-stage vs two-stage. We only reserve sparse matching for one-stage matching to compare with our two-stage matching. As shown in Table 5 (7) and (8), the one-stage matching method performs poorly on the *Kinect-sfm* benchmark, with RR decreasing by 25.0 pp. This proves the necessity of our prior-guided global dense matching and the superiority of two-stage matching.

Conclusion

We propose a novel cross-source point cloud registration framework. With the density-robust feature extractor, we address the issue of inconsistent data distribution among cross-source point clouds. To tackle the challenging matching problem in registration, we employ a “loose generation, strict selection” strategy. Utilizing a two-stage matching approach, we achieve reliable matching and robust registration.

References

- Ao, S.; Hu, Q.; Wang, H.; Xu, K.; and Guo, Y. 2023. BUFFER: Balancing Accuracy, Efficiency, and Generalizability in Point Cloud Registration. In *CVPR*, 1255–1264.
- Ao, S.; Hu, Q.; Yang, B.; Markham, A.; and Guo, Y. 2021. Spinnet: Learning a general surface descriptor for 3d point cloud registration. In *CVPR*, 11753–11762.
- Bai, X.; Luo, Z.; Zhou, L.; Fu, H.; Quan, L.; and Tai, C.-L. 2020. D3feat: Joint learning of dense detection and description of 3d local features. In *CVPR*, 6359–6367.
- Besl, P.; and McKay, N. D. 1992. A method for registration of 3-D shapes. *IEEE TPAMI*, 14(2): 239–256.
- Chen, Z.; Sun, K.; Yang, F.; and Tao, W. 2022. Sc2-pcr: A second order spatial compatibility for efficient and robust point cloud registration. In *CVPR*, 13221–13231.
- Choy, C.; Park, J.; and Koltun, V. 2019. Fully Convolutional Geometric Features. In *ICCV*, 8957–8965.
- Fischler, M. A.; and Bolles, R. C. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6): 381–395.
- Gath, I.; and Geva, A. B. 1989. Unsupervised optimal fuzzy clustering. *IEEE TPAMI*, 11(7): 773–780.
- Huang, S.; Gojcic, Z.; Usvyatsov, M.; Wieser, A.; and Schindler, K. 2021a. Predator: Registration of 3d point clouds with low overlap. In *CVPR*, 4267–4276.
- Huang, X.; Fan, L.; Wu, Q.; Zhang, J.; and Yuan, C. 2019. Fast registration for cross-source point clouds by using weak regional affinity and pixel-wise refinement. In *ICME*, 1552–1557.
- Huang, X.; Mei, G.; and Zhang, J. 2020. Feature-metric registration: A fast semi-supervised approach for robust point cloud registration without correspondences. In *CVPR*, 11366–11374.
- Huang, X.; Mei, G.; and Zhang, J. 2023. Cross-source point cloud registration: Challenges, progress and prospects. *Neurocomputing*, 126383.
- Huang, X.; Mei, G.; Zhang, J.; and Abbas, R. 2021b. A comprehensive survey on point cloud registration. *arXiv preprint arXiv:2103.02690*.
- Huang, X.; Zhang, J.; Wu, Q.; Fan, L.; and Yuan, C. 2017. A coarse-to-fine algorithm for matching and registration in 3D cross-source point clouds. *IEEE TCSVT*, 28(10): 2965–2977.
- Jin, S.; Armeni, I.; Pollefeys, M.; and Barath, D. 2024. Multiway Point Cloud Mosaicking with Diffusion and Global Optimization. In *CVPR*, 20838–20849.
- Leordeanu, M.; and Hebert, M. 2005. A spectral technique for correspondence problems using pairwise constraints. In *ICCV*, volume 2, 1482–1489.
- Li, Y.; and Harada, T. 2022. Leopard: Learning partial point cloud matching in rigid and deformable scenes. In *CVPR*, 5554–5564.
- Lin, C.-W.; Chen, T.-I.; Lee, H.-Y.; Chen, W.-C.; and Hsu, W. H. 2023. Coarse-to-fine point cloud registration with se(3)-equivariant representations. In *ICRA*, 2833–2840.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *CVPR*, 2117–2125.
- Ling, X.; and Qin, R. 2022. A graph-matching approach for cross-view registration of over-view and street-view based point clouds. *ISPRS J. Photogramm.*, 185: 2–15.
- Liu, Q.; Zhu, H.; Zhou, Y.; Li, H.; Chang, S.; and Guo, M. 2023. Density-invariant Features for Distant Point Cloud Registration. In *ICCV*, 18215–18225.
- Mallick, T.; Das, P. P.; and Majumdar, A. K. 2014. Characterizations of noise in Kinect depth images: A review. *IEEE Sens. J.*, 14(6): 1731–1740.
- Poiesi, F.; and Boscaini, D. 2022. Learning general and distinctive 3D local deep descriptors for point cloud registration. *IEEE TPAMI*, 45(3): 3979–3985.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 30.
- Qin, Z.; Yu, H.; Wang, C.; Guo, Y.; Peng, Y.; and Xu, K. 2022. Geometric transformer for fast and robust point cloud registration. In *CVPR*, 11143–11152.
- Ren, S.; Zeng, Y.; Hou, J.; and Chen, X. 2023. CorrI2P: Deep Image-to-Point Cloud Registration via Dense Correspondence. *IEEE TCSVT*, 33(3): 1198–1208.
- Roriz, R.; Cabral, J.; and Gomes, T. 2021. Automotive LiDAR technology: A survey. *IEEE TITS*, 23(7): 6282–6297.
- Rusu, R. B.; Blodow, N.; and Beetz, M. 2009. Fast Point Feature Histograms (FPFH) for 3D registration. In *ICRA*, 3212–3217.
- Rusu, R. B.; Blodow, N.; Marton, Z. C.; and Beetz, M. 2008. Aligning point cloud views using persistent feature histograms. In *IROS*, 3384–3391.
- Salti, S.; Tombari, F.; and Di Stefano, L. 2014. SHOT: Unique signatures of histograms for surface and texture description. *CVIU*, 125: 251–264.
- Sarlin, P.-E.; DeTone, D.; Malisiewicz, T.; and Rabinovich, A. 2020. Superglue: Learning feature matching with graph neural networks. In *CVPR*, 4938–4947.
- Segal, A.; Haehnel, D.; and Thrun, S. 2009. Generalized-icp. In *RSS*, 435. Seattle, WA.
- Sharp, G. C.; Lee, S. W.; and Wehe, D. K. 2002. ICP registration using invariant features. *IEEE TPAMI*, 24(1): 90–102.
- Sun, Y.; Cheng, C.; Zhang, Y.; Zhang, C.; Zheng, L.; Wang, Z.; and Wei, Y. 2020. Circle loss: A unified perspective of pair similarity optimization. In *CVPR*, 6398–6407.
- Tazir, M. L.; Gokhool, T.; Checchin, P.; Malaterre, L.; and Trassoudaine, L. 2018. CICIP: Cluster Iterative Closest Point for sparse–dense point cloud registration. *Rob. Auton. Syst.*, 108: 66–86.
- Thomas, H.; Qi, C. R.; Deschaud, J.-E.; Marcotegui, B.; Goulette, F.; and Guibas, L. J. 2019. Kpconv: Flexible and deformable convolution for point clouds. In *ICCV*, 6411–6420.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *NeurIPS*, 30.

Wang, H.; Liu, Y.; Dong, Z.; and Wang, W. 2022. You only hypothesize once: Point cloud registration with rotation-equivariant descriptors. In *ACM MM*, 1630–1641.

Xiong, K.; Zheng, M.; Xu, Q.; Wen, C.; Shen, S.; and Wang, C. 2023. SPEAL: Skeletal Prior Embedded Attention Learning for Cross-Source Point Cloud Registration. *arXiv preprint arXiv:2312.08664*.

Yang, F.; Guo, L.; Chen, Z.; and Tao, W. 2022. One-Inlier is First: Towards Efficient Position Encoding for Point Cloud Registration. *NeurIPS*, 35: 6982–6995.

Yang, H.; Shi, J.; and Carlone, L. 2020. Teaser: Fast and certifiable point cloud registration. *IEEE TRO*, 37(2): 314–333.

Yang, J.; Li, H.; Campbell, D.; and Jia, Y. 2015. Go-ICP: A globally optimal solution to 3D ICP point-set registration. *IEEE TPAMI*, 38(11): 2241–2254.

Yang, J.; Li, H.; and Jia, Y. 2013. Go-icp: Solving 3d registration efficiently and globally optimally. In *ICCV*, 1457–1464.

Yu, H.; Li, F.; Saleh, M.; Busam, B.; and Ilic, S. 2021. Cofinet: Reliable coarse-to-fine correspondences for robust pointcloud registration. *NeurIPS*, 34: 23872–23884.

Yu, H.; Qin, Z.; Hou, J.; Saleh, M.; Li, D.; Busam, B.; and Ilic, S. 2023. Rotation-invariant transformer for point cloud matching. In *CVPR*, 5384–5393.

Zeng, A.; Song, S.; Nießner, M.; Fisher, M.; Xiao, J.; and Funkhouser, T. 2017. 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. In *CVPR*, 1802–1811.

Zhao, G.; Guo, Z.; Wang, X.; and Ma, H. 2024. SphereNet: Learning a Noise-Robust and General Descriptor for Point Cloud Registration. *IEEE TGRS*, 62: 1–16.

Zhao, M.; Huang, X.; Jiang, J.; Mou, L.; Yan, D.-M.; and Ma, L. 2023. Accurate Registration of Cross-Modality Geometry via Consistent Clustering. *IEEE TVCG*.