

# Learning Instance-Aware Correspondences for Robust Multi-Instance Point Cloud Registration in Cluttered Scenes

Zhiyuan Yu<sup>1</sup> Zheng Qin<sup>2</sup> Lintao Zheng<sup>1</sup> Kai Xu<sup>1\*</sup>

<sup>1</sup> National University of Defense Technology

<sup>2</sup> Defense Innovation Institute, Academy of Military Sciences

## Abstract

Multi-instance point cloud registration estimates the poses of multiple instances of a model point cloud in a scene point cloud. Extracting accurate point correspondences is to the center of the problem. Existing approaches usually treat the scene point cloud as a whole, overlooking the separation of instances. Therefore, point features could be easily polluted by other points from the background or different instances, leading to inaccurate correspondences oblivious to separate instances, especially in cluttered scenes. In this work, we propose **MIRETR**, **Multi-Instance REgistration TRansformer**, a coarse-to-fine approach to the extraction of instance-aware correspondences. At the coarse level, it jointly learns instance-aware superpoint features and predicts per-instance masks. With instance masks, the influence from outside of the instance being concerned is minimized, such that highly reliable superpoint correspondences can be extracted. The superpoint correspondences are then extended to instance candidates at the fine level according to the instance masks. At last, an efficient candidate selection and refinement algorithm is devised to obtain the final registrations. Extensive experiments on three public benchmarks demonstrate the efficacy of our approach. In particular, MIRETR outperforms the state of the arts by 16.6 points on  $F_1$  score on the challenging ROBI benchmark. Code and models are available at <https://github.com/zhiyuanYU134/MIRETR>.

## 1. Introduction

Point cloud registration aims at estimating a rigid transformation that aligns two point clouds. In real-world application scenarios such as robotic bin picking, there is often the requirement of multi-instance registration where the point cloud of a model needs to be registered against multiple instances of the model in the target scene. *Multi-instance point cloud registration* is challenging due to the unknown number of instances and inter-instance occlusions

\*Corresponding author: Kai Xu (kevin.kai.xu@gmail.com).




Figure 1. MIRETR significantly improves the multi-instance registration results in cluttered scenes compared to the state-of-the-art GeoTransformer [38]. Benefiting from the instance-aware correspondences, our method can generate more accurate registrations (see the yellow boxes) and register the heavily-occluded instances with severe geometric deficiency (see the red boxes).

especially in a cluttered scene.

Existing approaches to multi-instance registration are mostly two-stage: They first extract point correspondences between the model and the scene point clouds and then recover per-instance transformations with multi-model fitting [28, 44, 55]. Clearly, the performance of such approach highly depends on the quality of the correspondences. Unlike pair-wise registration that can be pinned down by a small number of correspondences, multi-instance registration poses unique challenges. First, the correspondences should widely spread over the scene point cloud to cover as many instances as possible so that more instances can be found and registered. Second, the correspondences should

be accurately clustered into individual instances to estimate per-instance pose transformations.

Existing point correspondence methods can be broadly classified into *keypoint-based* and *keypoint-free* ones. Albeit successful in pairwise registration, these methods fail to effectively tackle the aforementioned challenges. On the one hand, keypoint-based methods [4, 14, 25] first detect keypoints and then match them to form correspondences. In the multi-instance setting, however, it is common that an instance is severely occluded so that too few keypoints can be detected to guarantee a successful registration.

On the other hand, keypoint-free methods [38, 54] bypass keypoint detection with a coarse-to-fine pipeline. They first extract correspondences between sparse superpoints and then refine them to dense point correspondences. To achieve reliable superpoint matching, they encode global context with feature correlation [54] or geometric structure [38] to learn distinctive superpoint features. When used in multi-instance registration, such approach may overlook the separation of individual instances and lead to the following two drawbacks. First, the superpoint features of an instance may be contaminated by the global context from the background or other instances, making the extracted features less informative. Thus, the features of a severely occluded instance could be overwhelmed by those more complete neighboring ones, making the occluded instances hard to register (see Fig. 1 (left)). Second, they cannot discriminate correspondences into different instances and have to rely on multi-model fitting [28, 44, 55] to cluster correspondences. Multi-model fitting needs to sample numerous hypotheses and finds difficulty in handling heavy occlusion.

We propose *MIRETR*, *Multi-Instance REgistration TRansformer*, which learns to extract instance-aware correspondences in a coarse-to-fine fashion. Our key motivation is to restrict the contextual feature encoding within the instance scope. An *Instance-aware Geometric Transformer* module is devised to jointly extract reliable *superpoint* features and predict instance masks in the scene. We resolve this chicken-and-egg problem with an iterative process. In each iteration, we first encode intra-instance geometric structure based on the instance mask of the previous iteration and then learn geometric consistency between the model and the scene. Next, the instance masks are refined based on the instance-aware features. As the iteration goes, the superpoint features gradually improve with increasingly relevant context and accurate instance masks, resulting in reliable registrations even for heavily-occluded instances (see Fig. 1 (right)).

Having obtained superpoint correspondences, we extend them to instance candidates based on instance masks. This allows us to directly extract instance-wise *point correspondences* and estimate a transformation for each instance candidate. An efficient candidate selection and refinement

method is proposed to eliminate duplicated instances and obtain the final registrations. Thanks to the instance-aware correspondences, our method bypasses the need of multi-model fitting, achieving superior efficiency and accuracy. Extensive experiments on three benchmarks [2, 10, 52] demonstrate the efficacy of our method. Our method surpasses the previous state-of-the-art methods [38, 55] by 16.6 points on  $F_1$  score on the cluttered ROBI benchmark. Our main contributions include:

- A multi-instance point cloud registration method which extracts instance-wise correspondences and estimates transformations without multi-model fitting.
- An instance-aware geometric transformer module which jointly learns instance-aware superpoint features and predicts instance masks.
- An efficient instance candidate selection and refinement method which removes duplicated candidates and generates the final registrations.

## 2. Related work

**Point cloud registration.** State-of-the-art point cloud registration methods can be categorized into direct registration methods and correspondence-based methods. Direct registration methods [1, 18, 26, 48, 49, 51, 53] use a neural network to estimate the transformation in an end-to-end manner. Correspondence-based methods [4, 14, 16, 17, 19, 25, 38, 54] extract point correspondences and estimate the rigid transformation with a robust pose estimator. Early point correspondence methods [4, 25] adopt a detect-then-match pipeline, where keypoints are first detected and then matched. Recent advances [38, 54] have bypassed keypoint detection by extracting correspondences in a coarse-to-fine fashion. However, existing methods mainly focus on pairwise registration and little attention has been paid to multi-instance registration, which faces new challenges such as unknown number of instances and heavy intra-instance occlusion. In this paper, we fill this gap by directly extracting instance-aware correspondences.

**Multi-model fitting.** Multi-model fitting aims at recovering multiple models from noisy data, which is an important step in multi-instance registration. Multi-model fitting methods can be classified into RANSAC-based methods and clustering-based methods. RANSAC-based methods [6, 7, 27, 28] adopt a hypothesize-and-verify manner to sequentially fit multiple models. Clustering-based methods [33–35, 44, 55] sample a huge set of hypotheses and cluster the correspondences according to their residuals under these hypotheses. However, these methods suffer from huge memory usage or long convergence time. In this work, we achieve accurate and efficient multi-instance registration without multi-model fitting.

**Open-set pose estimation.** Model retrieval [2, 3, 8, 9, 13, 20, 24, 29] and image-based pose estimation [11, 12,

22, 23, 31, 32, 37, 41–43, 47, 50] are two highly-related tasks to multi-instance registration. Model retrieval aims at estimating the 9-DoF pose parameters of the CAD models from a pre-built database that appears in a given scene. And image-based pose estimation recovers the 6-DoF poses for certain objects [12, 22, 23, 37, 41–43, 50] or categories [11, 31, 32, 47] in an RGB or RGB-D image. However, both tasks are based on the *close-set* assumption, where the training and testing data follow the same distribution. Model retrieval can only retrieve the CAD models in the database while image-based pose estimation can only handle known objects or categories in the training data. On the contrary, multi-instance registration is an *open-set* problem. It is based on geometric matching and thus can generalize well to novel objects and categories. We discuss the generality of MIRETR in Sec. 5.3.

### 3. Problem Statement

Given a model point cloud  $\mathcal{P} = \{\mathbf{p}_i \in \mathbb{R}^3 \mid i = 1, \dots, N\}$  and a scene point cloud  $\mathcal{Q} = \{\mathbf{q}_i \in \mathbb{R}^3 \mid i = 1, \dots, M\}$ , where  $\mathcal{Q}$  could contain multiple instances of  $\mathcal{P}$ , the goal of multi-instance registration is to register all instances of  $\mathcal{P}$  in  $\mathcal{Q}$ . Specifically,  $\mathcal{Q}$  is represented as  $\mathcal{Q} = \mathcal{Q}_0 \cup \mathcal{Q}_1 \cup \dots \cup \mathcal{Q}_K$ , where  $\mathcal{Q}_0$  is the background points and  $\{\mathcal{Q}_1, \dots, \mathcal{Q}_K\}$  are the  $K$  instances of  $\mathcal{P}$ . For each instance  $\mathcal{Q}_k$ , we estimate a 6-DoF pose  $\mathbf{T}_k \in \mathcal{SE}(3)$  aligning  $\mathcal{P}$  to  $\mathcal{Q}_k$ , which consists of a 3D rotation  $\mathbf{R}_k \in \mathcal{SO}(3)$  and a 3D translation  $\mathbf{t}_k \in \mathbb{R}^3$ . To solve this problem, one usually first extracts a set of correspondences  $\mathcal{C} = \{(\mathbf{p}_i, \mathbf{q}_i) \mid \mathbf{p}_i \in \mathcal{P}, \mathbf{q}_i \in \mathcal{Q}\}$  between the two point clouds. The correspondences are then clustered into different groups  $\{\mathcal{C}_1, \dots, \mathcal{C}_K\}$ , and each group recovers the pose of one instance by solving the following problem:

$$\min_{\mathbf{R}_k, \mathbf{t}_k} \sum_{(\mathbf{p}_i, \mathbf{q}_i) \in \mathcal{C}_k} \|\mathbf{R}_k \cdot \mathbf{p}_i + \mathbf{t}_k - \mathbf{q}_i\|_2^2. \quad (1)$$

This problem is very challenging as the number of instances is unknown and it is difficult to cluster the correspondences to the correct instances. In this work, we deal with this problem by extracting instance-aware correspondences.

### 4. Method

MIRETR extracts correspondences in a coarse-to-fine fashion similar to keypoint-free registration methods. The overall pipeline of our method is illustrated in Fig. 2. At the coarse level, we establish correspondences between down-sampled superpoints with an *Instance-aware Geometric Transformer* module (Sec. 4.1). At the fine level, each superpoint correspondence is extended to form an instance candidate, where *instance-wise point correspondences* are extracted for pose estimation (Sec. 4.2). At last, we merge similar instance candidates and refine the resultant poses to obtain the final registrations (Sec. 4.3).

Given the input point clouds  $\mathcal{P}$  and  $\mathcal{Q}$ , we progressively downsample them into sparse superpoints, denoted as  $\hat{\mathcal{P}}$  and  $\hat{\mathcal{Q}}$ , using grid subsampling [45]. KPFCNN [45] is adopted to extract multi-level point features. Let us denote the features of  $\mathcal{P}$ ,  $\mathcal{Q}$ ,  $\hat{\mathcal{P}}$  and  $\hat{\mathcal{Q}}$  as  $\mathbf{F}^{\mathcal{P}}$ ,  $\mathbf{F}^{\mathcal{Q}}$ ,  $\hat{\mathbf{F}}^{\mathcal{P}}$  and  $\hat{\mathbf{F}}^{\mathcal{Q}}$ , respectively. Following [38, 54], each superpoint  $\hat{\mathbf{p}}_i \in \hat{\mathcal{P}}$  ( $\hat{\mathbf{q}}_i \in \hat{\mathcal{Q}}$ ) is associated with a local patch  $\mathcal{G}_i^{\mathcal{P}}$  ( $\mathcal{G}_i^{\mathcal{Q}}$ ) by the point-to-node partition strategy [30] where each point is assigned to its nearest superpoint.

#### 4.1. Instance-aware Geometric Transformer

Given the superpoints  $\hat{\mathcal{P}}$  and  $\hat{\mathcal{Q}}$ , as well as their features  $\hat{\mathbf{F}}^{\mathcal{P}}$  and  $\hat{\mathbf{F}}^{\mathcal{Q}}$ , we first extract a set of superpoint correspondences whose local patches overlap with each other. Accurate superpoint matching relies on learning the geometric consistency between the two point clouds. To this end, previous works [38, 54] adopt transformer [46] to encode intra- and inter-point-cloud context in the global scope, as shown in Fig. 3 (a). Albeit successful in pairwise registration, this is problematic in multi-instance registration. In the scene point cloud, the superpoints from the background or other instances are de facto noises for modeling the geometric consistency and severely pollute the superpoint features. This becomes more severe in cluttered scenes. Considering an heavily-occluded instance, as its geometric structure is incomplete, the features of its superpoints could be easily overwhelmed by those from the background or nearby instances. As a result, it can hardly be registered.

To address these issues, we propose to make the superpoint features *instance-aware* with a novel *Instance-aware Geometric Transformer* module. The key insight is to restrict the intra-point-cloud context encoding in the scene point cloud within each individual instance. To this end, we first extract the  $k$ -nearest neighbors for each superpoint and conduct context aggregation within local regions. However, as shown in Fig. 3 (b), nearby superpoints do not necessarily reside in the same instance, and the feature pollution problem still exists. For this reason, we further design an instance masking mechanism by predicting an instance mask for each superpoint to select the neighbors in the same instance as it, and only aggregate the context among them. Fig. 3 (c) illustrates our instance masking mechanism.

As shown in Fig. 2 (bottom), the instance-aware geometric transformer contains three blocks. First, a *geometric encoding block* encodes intra-instance geometric structure based on the instance masks. Next, a *cross-attention block* enhances the superpoint features by modeling inter-point-cloud geometric consistency. At last, an *instance masking block* predicts new instance masks with the instance-aware superpoint features. We adopt  $N_t$  instance-aware geometric transformer modules to progressively refine the superpoint features and the instance masks.

**Geometric encoding block.** This block encodes the intra-

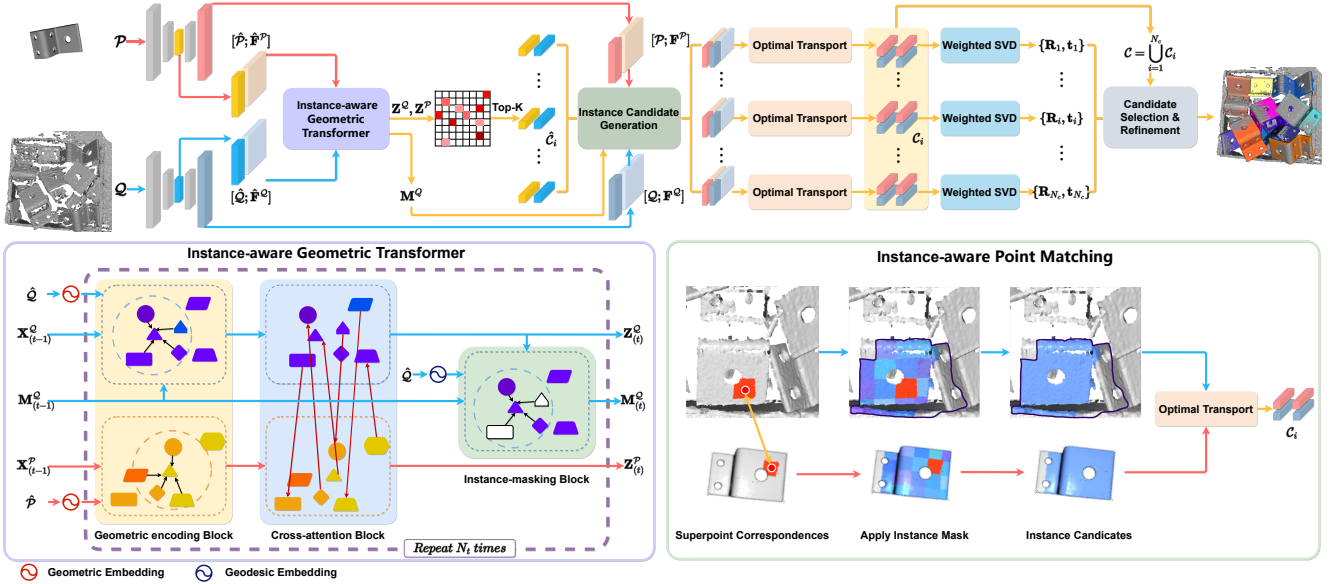


Figure 2. Overall pipeline of MIRETR. The backbone progressively downsamples two point clouds and extracts multi-level features. At the coarse level, the Instance-aware Geometric Transformer module extracts instance-aware superpoint features and establishes reliable superpoint correspondences. At the fine level, the superpoint correspondences are extended to instance candidates, where instance-wise point correspondences are extracted to estimate per-candidate poses. At last, a simple but effective candidate selection and refinement algorithm is adopted to generate the final registrations.

instance geometric context within each point cloud. Given an anchor superpoint  $\hat{\mathbf{q}}_i \in \hat{\mathcal{Q}}$ , its  $k$ -nearest superpoints  $\mathcal{N}_{i,k}^{\mathcal{Q}} = \{\hat{\mathbf{q}}_{i_1}, \dots, \hat{\mathbf{q}}_{i_k}\}$ , the input feature matrix  $\mathbf{X}^{\mathcal{Q}} \in \mathbb{R}^{|\hat{\mathcal{Q}}| \times d}$ , and the instance mask matrix  $\mathbf{M}^{\mathcal{Q}} \in \mathbb{R}^{|\hat{\mathcal{Q}}| \times k}$ , the output feature matrix  $\mathbf{Z}^{\mathcal{Q}} \in \mathbb{R}^{|\hat{\mathcal{Q}}| \times d}$  is computed as:

$$\mathbf{z}_i^{\mathcal{Q}} = \sum_{j=1}^k \frac{\exp(e_{i,j})}{\sum_{l=1}^k \exp(e_{i,l})} (\mathbf{x}_{i_j}^{\mathcal{Q}} \mathbf{W}^V), \quad (2)$$

where the attention score  $e_{i,j}$  is computed as:

$$e_{i,j} = \frac{(\mathbf{x}_i^{\mathcal{Q}} \mathbf{W}^Q)(\mathbf{x}_{i_j}^{\mathcal{Q}} \mathbf{W}^K + \mathbf{r}_{i,j} \mathbf{W}^R)^\top}{\sqrt{d}} + m_{i,j}^{\mathcal{Q}}, \quad (3)$$

where  $\mathbf{W}^Q, \mathbf{W}^K, \mathbf{W}^V, \mathbf{W}^R \in \mathbb{R}^{d \times d}$  are the projection weights for query, key, value and geometric embedding, respectively.  $\mathbf{r}_{i,j}$  is the geometric structure embedding [38] encoding the transformation-invariant geometric information among the superpoints. For the instance mask  $\mathbf{M}^{\mathcal{Q}}$ , we set  $m_{i,j}^{\mathcal{Q}} = 0$  if  $\hat{\mathbf{q}}_i$  and  $\hat{\mathbf{q}}_{i_j}$  are in the same instance, while  $m_{i,j}^{\mathcal{Q}} = -\infty$  otherwise.  $\mathbf{M}^{\mathcal{Q}}$  is initialized with all zeros and refined by the instance masking block described later. Benefiting from the instance masks, the superpoint features are instance-aware by encoding the feature correlation and the geometric context within their belonging instances. For the computation in  $\hat{\mathcal{P}}$ , we ignore the mask term in Eq. (3) and aggregate the features from all neighbors.

**Cross-attention block.** After encoding the intra-instance geometric context, we further learn the inter-point-cloud geometric consistency with a cross-attention block inspired

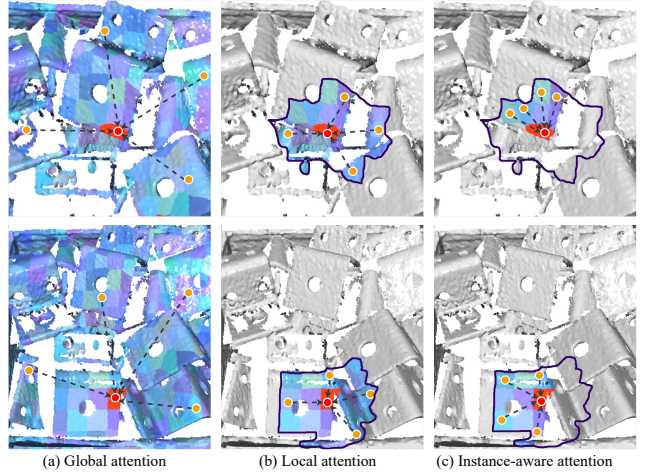


Figure 3. Comparison of (a) global attention, (b) local attention, and (c) instance-aware attention. The superpoints (patches) participating in the attention computation are color-coded. The anchor superpoints are in red. The  $k$ -nearest neighbors of the anchor are enclosed by the purple line.

by [38, 54]. Given the feature matrices  $\mathbf{X}^{\mathcal{P}} \in \mathbb{R}^{|\hat{\mathcal{P}}| \times d}$  of  $\hat{\mathcal{P}}$  and  $\mathbf{X}^{\mathcal{Q}} \in \mathbb{R}^{|\hat{\mathcal{Q}}| \times d}$  of  $\hat{\mathcal{Q}}$  from the geometric encoding block, the output feature matrix  $\mathbf{Z}^{\mathcal{P}} \in \mathbb{R}^{|\hat{\mathcal{P}}| \times d}$  of  $\hat{\mathcal{P}}$  is computed as:

$$\mathbf{z}_i^{\mathcal{P}} = \sum_{j=1}^{|\hat{\mathcal{Q}}|} \frac{\exp(e_{i,j})}{\sum_{k=1}^{|\hat{\mathcal{Q}}|} \exp(e_{i,k})} (\mathbf{x}_k^{\mathcal{Q}} \mathbf{W}^V), \quad (4)$$

where the attention score  $e_{i,j}$  is computed as:

$$e_{i,j} = \frac{(\mathbf{x}_i^{\mathcal{P}} \mathbf{W}^Q)(\mathbf{x}_j^{\mathcal{Q}} \mathbf{W}^K)^\top}{\sqrt{d}}, \quad (5)$$

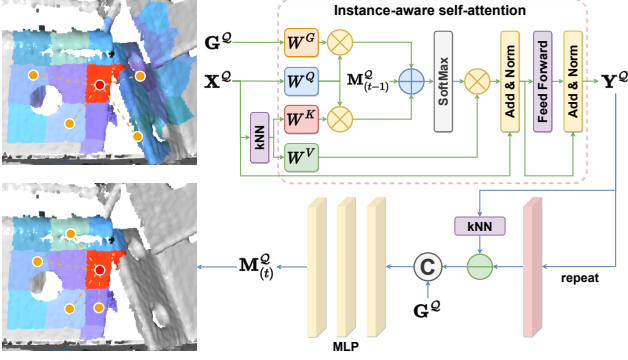


Figure 4. Structure of the Instance masking block.

where  $\mathbf{W}^Q, \mathbf{W}^K, \mathbf{W}^V \in \mathbb{R}^{d \times d}$  are the respective projection weights for query, key and value. The same computation goes for  $\mathbf{Z}^Q$ . Benefiting from the cross-attention block, the superpoint features in one point cloud are aware of the geometric structure of the other one, which facilitates modeling the geometric consistency between two point clouds.

**Instance masking block.** At last, this block consumes the cross-attention features  $\mathbf{X}^Q$  of  $\hat{Q}$  as well as the previous instance mask  $\mathbf{M}^Q_{(t-1)}$  to predict a refined instance mask  $\mathbf{M}^Q_{(t)}$ . We observe that two superpoints from different instances tend to be distant in the geodesic space. To this end, we design a geodesic self-attention mechanism by replacing the geometric structure embedding in Eq. (3) with a similar geodesic distance embedding to enhance the superpoint features. Please refer to the appendix for more details. The resultant superpoint features are denoted as  $\mathbf{Y}$ . At last, we adopt an MLP to predict a confidence score  $u_{i,j}$  for each neighbor  $\hat{\mathbf{q}}_{i,j} \in \mathcal{N}_i^Q$  indicating whether it belongs to the same instance as  $\hat{\mathbf{q}}_i$  based on the feature discrepancy between  $\mathbf{y}_i$  and  $\mathbf{y}_{i,j}$  and the geodesic distance embedding  $\mathbf{g}_{i,j}$  between  $\hat{\mathbf{q}}_i$  and  $\hat{\mathbf{q}}_{i,j}$ :

$$u_{i,j} = \sigma\left(\text{MLP}([\mathbf{y}_{i,j} - \mathbf{y}_i; \mathbf{g}_{i,j}])\right), \quad (6)$$

where  $\sigma(\cdot)$  is the sigmoid function and  $[\cdot]$  is the concatenation operator. At last, the confidence matrix  $\mathbf{U} \in \mathbb{R}^{|\hat{Q}| \times k}$  is converted into  $\mathbf{M}^Q_{(t)}$  with a thresholding function:

$$m_{i,j}^Q = \begin{cases} -\infty & u_{i,j} < \tau \\ 0 & \text{otherwise} \end{cases}, \quad (7)$$

where  $\tau$  is the confidence threshold. Thanks to the instance mask, our model can effectively learn instance-aware superpoint features and extract accurate superpoint correspondences covering more instances. Furthermore, the instance masks also help extract instance-aware dense point correspondences as described in Sec. 4.2.

**Superpoint matching.** We extract *superpoint correspondences*  $\hat{C}$  by matching the superpoint features from the last cross-attention block. As in [38], we select the top  $N_c$  su-

perpoint pairs with the highest cosine feature similarity as the superpoint correspondences.

## 4.2. Instance Candidate Generation

After obtaining the superpoint correspondences, we then refine them to dense point correspondences at the fine level. Previous methods [38, 54] opt to match the points within the local patches of two matched superpoints with an optimal transport layer [40]. Nonetheless, the local correspondences extracted in this manner often cluster closely, which could lead to unstable pose estimation as noted in [25]. This is aggravated in multi-instance registration, especially in cluttered scenes, as each instance is typically small and thus its pose cannot be globally optimized as in [38].

To address this issue, we propose to extract dense point correspondences *within the instance scope* by leveraging the instance masks from the coarse level. For each superpoint correspondence  $\hat{C}_k = (\hat{\mathbf{p}}_i, \hat{\mathbf{q}}_j)$ , we collect their neighboring superpoints  $\mathcal{N}_i^P$  and  $\mathcal{N}_j^Q$ . The superpoints from different instances of  $\hat{\mathbf{q}}_j$  are removed from  $\mathcal{N}_j^Q$  based on the instance mask  $\mathbf{M}_j$ . The points in the local patches of all superpoints in  $\mathcal{N}_i^P$  and  $\mathcal{N}_j^Q$  imply a potential occurrence of  $\mathcal{P}$  in  $\mathcal{Q}$ , i.e., an *instance candidate*, denoted as  $\mathcal{I}_k$ . Next, instance-wise *point correspondences* are extracted within  $\mathcal{I}_k$  with an optimal transport layer and mutual top- $k$  selection following [38], denoted as  $\mathcal{C}_k$ . At last, we estimate a pose  $\mathbf{T}_k = \{\mathbf{R}_k, \mathbf{t}_k\}$  for  $\mathcal{I}_k$  by solving Eq. (1) with weighted SVD [8]. Thanks to the powerful instance-aware geometric transformer, an instance candidate can cover a relatively large portion of an instance, with few points from the background or other instances, so the poses obtained in this step have already been very accurate.

## 4.3. Candidate Selection and Refinement

As several superpoint correspondences could belong to the same instance, there are commonly duplicated instances in the instance candidates. For this reason, we further design a simple but effective algorithm for duplicate removal and candidate refinement.

Inspired by non-maximum suppression (NMS) [39], we first sort the instance candidates by the inlier ratio on the global point correspondences  $\mathcal{C} = \bigcup_{i=1}^{N_c} \mathcal{C}_i$ . The inlier ratio of the instance candidate  $\mathcal{I}_k$  is:

$$\gamma_k = \frac{1}{|\mathcal{C}|} \sum_{(\mathbf{p}, \mathbf{q}) \in \mathcal{C}} [\|\mathbf{R}_k \mathbf{p} + \mathbf{t}_k - \mathbf{q}\| < \tau_2], \quad (8)$$

where  $[\cdot]$  is the Iverson bracket and  $\tau_2$  is the acceptance radius. Next, we select the instance candidate with the highest inlier ratio as the anchor candidate. All remaining candidates similar with the anchor are merged with it and removed from future computation. We define the similarity between two candidates  $\mathcal{I}_i$  and  $\mathcal{I}_j$  by the average distance

(ADD) [21] between their estimated poses:

$$s_{i,j} = 1 - \frac{\text{ADD}(\mathbf{T}_i, \mathbf{T}_j)}{r}, \quad (9)$$

where  $r$  is the diameter of  $\mathcal{P}$ . And  $\mathcal{I}_i$  and  $\mathcal{I}_j$  are considered similar if  $s_{i,j}$  is above  $\tau_s$ . To merge two candidates, we combine their point correspondences and recompute the pose by solving Eq. (1). At last, the preserved pose is iteratively refined with the surviving inliers as in [5, 38], leading to a final instance registration. The anchor is then removed from future computation. The above process repeats until there are no candidates left and we omit the instance registrations with too few inliers.

Compared to multi-model fitting methods [44, 55], the advantages of our MIRETR are two-fold. First, by learning instance-aware superpoint features, the influence from the background and other instances is minimized, thus the correspondences are more accurate and can cover more instances. Second, correspondence clustering-based methods tend to lose heavily-occluded instances as there are few correspondences on them, especially in cluttered scenes. On the contrary, our method can better recognize these hard instances as we bypass the clustering process by directly extracting instance candidates and merging similar ones.

#### 4.4. Loss Functions

We adopt three loss functions to train MIRETR: an overlap-aware circle loss [38]  $\mathcal{L}_{\text{circle}}$  to supervise the superpoint features, a negative log-likelihood loss [40]  $\mathcal{L}_{\text{nl}}$  to supervise the point matching, and a mask prediction loss [36]  $\mathcal{L}_{\text{mask}}$  to supervise the instance masks. The overall loss is computed as  $\mathcal{L} = \mathcal{L}_{\text{circle}} + \mathcal{L}_{\text{nl}} + \mathcal{L}_{\text{mask}}$ . Please refer to the supplementary material for more details.

### 5. Experiments

We evaluate MIRETR on indoor Scan2CAD [2, 44, 55], industrial ROBI [52] and synthetic ShapeNet [10] benchmarks. And we also introduce the implementation details and more experiments in the supplementary material.

**Baselines.** We compare with three state-of-the-art point cloud correspondence methods, FCGF [14], CoFiNet [54] and GeoTransformer [38], and four recent multi-model fitting methods, T-linkage [33], Ransacov [34], ECC [44], and PointCLM [55]. The point cloud correspondence methods and the multi-model fitting methods are pairwise integrated for a comprehensive comparison.

**Metrics.** Following [44, 55], we evaluate our method with three registration metrics: (1) *Mean Recall* (MR), the ratio of registered instances over all ground-truth instances, (2) *Mean Precision* (MP), the ratio of registered instances over all predicted instances, and (3) *Mean F<sub>1</sub> score* (MF), the harmonic mean of MP and MR. We also report *Inlier Ratio* (IR), the ratio of inliers over all extracted correspondences.

Model	IR(%)	MR(%)	MP(%)	MF(%)
FCGF [14] + T-Linkage [33]	36.27	54.75	22.76	32.15
FCGF [14] + Ransacov [34]		73.50	45.01	55.83
FCGF [14] + PointCLM [55]		86.99	70.05	77.60
FCGF [14] + ECC [44]		92.60	73.79	82.13
CoFiNet [54] + T-Linkage [33]	18.52	28.17	5.09	8.62
CoFiNet [54] + Ransacov [34]		38.29	13.69	20.16
CoFiNet [54] + PointCLM [55]		36.15	24.70	29.34
CoFiNet [54] + ECC [44]		57.26	17.24	26.50
GeoTransformer [38] + T-Linkage [33]	44.02	72.68	44.80	55.43
GeoTransformer [38] + Ransacov [34]		78.84	67.16	72.53
GeoTransformer [38] + PointCLM [55]		82.81	81.90	82.35
GeoTransformer [38] + ECC [44]		94.63	74.83	83.57
MIRETR (ours) + T-Linkage [33]	56.59	77.12	46.04	57.65
MIRETR (ours) + Ransacov [34]		84.78	71.34	77.48
MIRETR (ours) + PointCLM [55]		91.85	91.08	91.46
MIRETR (ours) + ECC [44]		96.52	89.03	92.62
MIRETR (ours, full pipeline)		95.70	91.21	93.40

Table 1. Evaluation results on Scan2CAD benchmark.

### 5.1. Evaluations on Scan2CAD

**Dataset.** Scan2CAD [2] is a scene-to-CAD alignment dataset build upon ScanNet [15] and ShapeNet [10]. It consists of 1506 scenes from ScanNet annotated with 14225 CAD models from ShapeNet and their poses in the scenes. Following [55], we replace the objects in the scenes with the corresponding aligned CAD models. We select the scene-model pairs which the scene contain multiple instances of the model for multi-instance registration. As last, we obtain 2184 point cloud pairs and use 70% pairs for training, 10% for validation and 20% for testing.

**Quantitative results.** As shown in Tab. 1, our MIRETR achieves consistent improvements over the baseline correspondence models. Notably, our method surpasses the previous best GeoTransformer on IR by 12 percentage points (pp). When coupled with PointCLM, it outperforms GeoTransformer by over 9 pp on the three registration metrics.

Moreover, our full pipeline surpasses all models with the baseline multi-model fitting methods. Although the ECC-based model performs slightly better than full MIRETR on MR, its MP is significantly worse. This indicates that ECC achieves high MR by predicting more registrations, which is impractical in real-world applications. And our method outperforms PointCLM on all three registration metrics, which demonstrates the strong effectiveness of our method.

### 5.2. Evaluations on ROBI

**Dataset.** ROBI [52] is a recent dataset for industrial bin-picking. It includes 7 reflective metallic industrial objects and 63 bin-picking scenes. Each scene contains cluttered instances of one industrial object. For each point cloud pair, the scene point cloud is backprojected from a depth image and the model point cloud is sampled from the CAD model of its corresponding industrial object. We obtain 4880 pairs in total, and we split the scenes for training, validation and testing according to the ratio of 6 : 1 : 2 for each object.

**Quantitative results.** As in Tab. 2, our method outper-

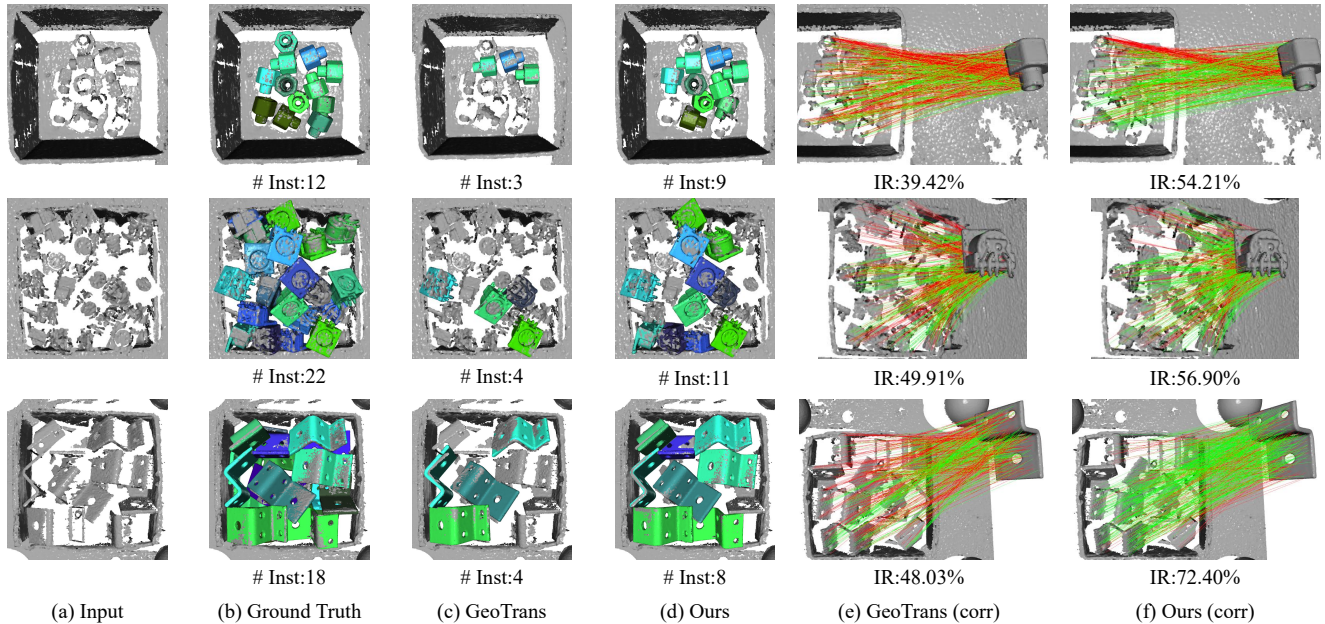


Figure 5. Registration results on ROBI benchmark. We visualize the successfully registered instances in (c) and (d). MIRETR registers more instances in the cluttered scenes (the 2<sup>nd</sup> and the 3<sup>rd</sup> rows) and with incomplete geometry (all three rows). And it extracts more accurate correspondences benefiting from the instance-aware correspondence learning mechanism.

Model	IR(%)	MR(%)	MP(%)	MF(%)
CoFiNet [54] + T-Linkage [33]	10.35	1.34	0.72	0.93
CoFiNet [54] + RansaCov [34]		1.25	1.39	1.31
CoFiNet [54] + PointCLM [55]		1.17	2.11	1.50
CoFiNet [54] + ECC [44]		3.21	8.50	4.66
GeoTransformer [38] + T-Linkage [33]	39.64	7.71	4.57	5.73
GeoTransformer [38] + RansaCov [34]		8.99	13.58	10.81
GeoTransformer [38] + PointCLM [55]		13.98	26.64	18.33
GeoTransformer [38] + ECC [44]		18.55	30.99	23.20
MIRETR (ours) + T-Linkage [33]	45.14	12.04	10.47	11.20
MIRETR (ours) + RansaCov [34]		14.14	26.29	18.38
MIRETR (ours) + PointCLM [55]		18.68	40.11	25.48
MIRETR (ours) + ECC [44]		24.65	34.85	28.91
MIRETR (ours, full pipeline)		38.51	41.19	39.80

Table 2. Evaluation results on ROBI benchmark.

forms the baseline correspondence methods by a large margin on all four metrics. Compared to Scan2CAD, ROBI is more challenging due to the more cluttered scenes, where instance-aware information is more important. Benefiting from the instance-aware geometric transformer module, our method extracts reliable instance-aware correspondences, which contributes to more accurate instance registrations.

Our full model achieves significant improvements over the ECC-based model on MR, MP and MF. The PointCLM-based model achieves comparable MP with the full model, but significantly worse MR, which means it can register only several easy instances but misses most heavily-occluded instances. Specifically, it can register on average only 6.2 instances per scene, while the number of our method is 13.7. Our method achieves both high MR and MP, showing strong robustness in cluttered scenes.

**Qualitative results.** Fig. 5 visualizes the registration results of GeoTransformer and MIRETR. Our method can




Figure 6. Visualizations of the predicted instance masks. The anchor superpoint (patch) is in red and the selected neighbors are color-coded. Our method can effectively reject the patches outside the interested instance while preserving most of those inside.

register more instances in cluttered scenes (see the 2<sup>nd</sup> and the 3<sup>rd</sup> rows) and with severe geometric deficiency (see all three rows). And MIRETR attains more high-quality correspondences thanks to our instance-aware design.

We further visualize the predicted instance masks Fig. 6. The instance masks effectively reject the patches out of the interested instance while preserving most of those inside. Specifically, our method achieves the instance mIoU of 69.26 pp while the result of the model without masking is merely 38.13 pp. The high-quality masks allow us to effectively encode intra-instance context and contribute to accurate registrations.

Fig. 7 compares the results under different overlap ratios.

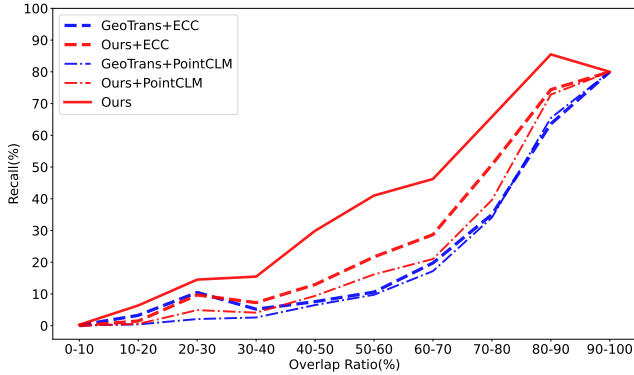


Figure 7. Results in different overlap levels on ROBI benchmark.

MIRETR consistently outperforms the baselines at all overlap levels, and the improvements are more significant when the overlap ratio is below 50%. This proves the strong capability of our method to handle low-overlap instances.

### 5.3. Evaluations on ShapeNet

**Dataset.** ShapeNet [10] is a large CAD model dataset which contains 51300 models from 55 categories. We use the models from first 30 categories for training and those from the rest 25 categories for testing to evaluate the generality to novel categories. We randomly sample at most 500 models from each category to avoid class imbalance. For each point cloud pair, the model point cloud is a random CAD model and the scene one is constructed by applying 4~16 random poses on the model. At last, we obtain 8634 pairs for training, 900 for validation, and 7683 for testing.

**Quantitative results.** As shown in Tab. 3, our MIRETR achieves the best results on all four metrics, surpassing the baselines by a large margin. GeoTransformer significantly outperforms CoFiNet, which proves the importance of geometric structure to achieving generality. And our method achieves improvements of more than 10 pp over GeoTransformer on all registration metrics with PointCLM and ECC. Thanks to the instance-aware geometric transformer, the correspondences extracted by our method provide strong instance information, facilitating the correspondence clustering process. At last, our full model surpasses the previous best ECC-based baseline by around 2 pp on MR, 11 pp on MP and 7 pp on MF, demonstrating the strong generality of MIRETR to unseen objects.

### 5.4. Ablation Studies

We conduct extensive ablation studies on ROBI benchmark to better understand our design choices in Tab. 4. We first ablate the instance-aware point matching and extract point correspondences within only the patches of the superpoint correspondences in Tab. 4 (b), which is significantly worse than the full model. The poses estimated by this method is unstable due to clustered correspondences, thus the duplicated candidates cannot be effectively removed by NMS.

Model	IR(%)	MR(%)	MP(%)	MF(%)
CoFiNet [54] + T-Linkage [33]	39.53	4.59	3.24	3.79
CoFiNet [54] + Ransacov [34]		7.35	5.33	6.17
CoFiNet [54] + PointCLM [55]		23.40	19.66	21.36
CoFiNet [54] + ECC [44]		41.49	26.60	32.41
GeoTransformer [38] + T-Linkage [33]	69.65	27.68	26.40	27.02
GeoTransformer [38] + Ransacov [34]		41.86	48.16	44.78
GeoTransformer [38] + PointCLM [55]		68.09	69.09	68.58
GeoTransformer [38] + ECC [44]		78.03	64.08	70.37
MIRETR (ours) + T-Linkage [33]	77.13	29.40	28.67	29.03
MIRETR (ours) + Ransacov [34]		45.63	46.85	46.23
MIRETR (ours) + PointCLM [55]		85.63	85.81	85.71
MIRETR (ours) + ECC [44]		92.56	82.08	87.01
MIRETR (ours, full pipeline)		94.95	93.94	94.44

Table 3. Evaluation results on ShapeNet benchmark.

Model	IR(%)	MR(%)	MP(%)	MF(%)
(a) full MIRETR	45.14	38.51	41.19	39.80
(b) w/o instance-aware point matching	40.52	18.15	27.54	21.88
(c) w/o instance masking block	36.78	17.07	45.37	24.80
(d) w/o both	38.81	14.13	27.44	18.65
(e) w/ global attention	25.98	7.72	18.86	10.96

Table 4. Ablation studies on ROBI benchmark.

Next, we remove the instance masking blocks in Tab. 4 (c) and the instance-aware point matching gathers the points from all neighbor superpoints. This model achieves a good MP, but a significantly worse MR, which means it can register only few instances (2.9 vs. 13.7 average instances per scene). We then ablate both modules in Tab. 4 (d), which further degrades the performance. At last, we replace the geometric encoding block with the global geometric self-attention [38] in Tab. 4 (e), which achieves the worst results. This indicates that the superpoint features can be polluted if global context is encoded. These results have proven the strong efficacy of our designs. More ablation studies are presented in the supplementary material.

## 6. Conclusion

We have proposed MIRETR, a coarse-to-fine method to extract instance-aware correspondences for multi-instance registration. At the coarse level, an instance-aware geometric transformer module jointly learns instance-aware superpoint features and predicts per-instance masks. At the fine level, the superpoint correspondences are extended to instance candidates according to the instance masks, where instance-wise point correspondences are extracted. At last, we devise a simple but effective candidate selection and refinement algorithm to obtain the final registrations, bypassing the need of multi-model fitting methods. Extensive experiments on three public benchmarks have demonstrated the efficacy of our method. For future work, we would like to extend MIRETR to multi-modal multi-instance registration to boost more applications.

**Acknowledgement.** This work is in part supported by the NSFC (62325211, 62132021, 62102435), the Major Program of Xiangjiang Laboratory (23XJ01009), and the National Key R&D Program of China (2018AAA0102200).



## References

- [1] Yasuhiro Aoki, Hunter Goforth, Rangaprasad Arun Srivatsan, and Simon Lucey. Pointnetlk: Robust & efficient point cloud registration using pointnet. In *CVPR*, pages 7163–7172, 2019. 2
- [2] Armen Avetisyan, Manuel Dahnert, Angela Dai, Manolis Savva, Angel X Chang, and Matthias Nießner. Scan2cad: Learning cad model alignment in rgb-d scans. In *CVPR*, pages 2614–2623, 2019. 2, 6
- [3] Armen Avetisyan, Tatiana Khanova, Christopher Choy, Denver Dash, Angela Dai, and Matthias Nießner. Scenecad: Predicting object alignments and layouts in rgb-d scans. In *ECCV*, pages 596–612, 2020. 2
- [4] Xuyang Bai, Zixin Luo, Lei Zhou, Hongbo Fu, Long Quan, and Chiew-Lan Tai. D3feat: Joint learning of dense detection and description of 3d local features. In *CVPR*, pages 6359–6367, 2020. 2
- [5] 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. 6
- [6] Daniel Barath and Jiri Matas. Progressive-x: Efficient, anytime, multi-model fitting algorithm. In *ICCV*, pages 3780–3788, 2019. 2
- [7] Daniel Barath, Denys Rozumny, Ivan Eichhardt, Levente Hajder, and Jiri Matas. Progressive-x+: Clustering in the consensus space. *arXiv preprint arXiv:2103.13875*, 2021. 2
- [8] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor fusion IV: control paradigms and data structures*, pages 586–606, 1992. 2, 5
- [9] Thomas O Binford. Survey of model-based image analysis systems. *The International Journal of Robotics Research*, 1(1):18–64, 1982. 2
- [10] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 2, 6, 8
- [11] Dengsheng Chen, Jun Li, Zheng Wang, and Kai Xu. Learning canonical shape space for category-level 6d object pose and size estimation. In *CVPR*, pages 11973–11982, 2020. 2, 3
- [12] Kai Chen, Rui Cao, Stephen James, Yichuan Li, Yun-Hui Liu, Pieter Abbeel, and Qi Dou. Sim-to-real 6d object pose estimation via iterative self-training for robotic bin picking. In *ECCV*, pages 533–550, 2022. 2, 3
- [13] Roland T Chin and Charles R Dyer. Model-based recognition in robot vision. *ACM Computing Surveys (CSUR)*, 18(1):67–108, 1986. 2
- [14] Christopher Choy, Jaesik Park, and Vladlen Koltun. Fully convolutional geometric features. In *CVPR*, pages 8958–8966, 2019. 2, 6
- [15] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, pages 5828–5839, 2017. 6
- [16] Haowen Deng, Tolga Birdal, and Slobodan Ilic. Ppf-foldnet: Unsupervised learning of rotation invariant 3d local descriptors. In *ECCV*, pages 602–618, 2018. 2
- [17] Haowen Deng, Tolga Birdal, and Slobodan Ilic. Ppfnet: Global context aware local features for robust 3d point matching. In *CVPR*, pages 195–205, 2018. 2
- [18] Kexue Fu, Shaolei Liu, Xiaoyuan Luo, and Manning Wang. Robust point cloud registration framework based on deep graph matching. In *CVPR*, pages 8893–8902, 2021. 2
- [19] Zan Gojcic, Caifa Zhou, Jan D Wegner, and Andreas Wieser. The perfect match: 3d point cloud matching with smoothed densities. In *CVPR*, pages 5545–5554, 2019. 2
- [20] Can Gümeli, Angela Dai, and Matthias Nießner. Roca: robust cad model retrieval and alignment from a single image. In *CVPR*, pages 4022–4031, 2022. 2
- [21] Stefan Hinterstoisser, Vincent Lepetit, Slobodan Ilic, Stefan Holzer, Gary Bradski, Kurt Konolige, and Nassir Navab. Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In *ACCV*, 2012. 6
- [22] Tomas Hodan, Frank Michel, Eric Brachmann, Wadim Kehl, Anders GlentBuch, Dirk Kraft, Bertram Drost, Joel Vidal, Stephan Ihrke, Xenophon Zabulis, et al. Bop: Benchmark for 6d object pose estimation. In *ECCV*, pages 19–34, 2018. 3
- [23] Yinlin Hu, Pascal Fua, Wei Wang, and Mathieu Salzmann. Single-stage 6d object pose estimation. In *CVPR*, pages 2930–2939, 2020. 3
- [24] Siyuan Huang, Siyuan Qi, Yixin Zhu, Yinxue Xiao, Yuanlu Xu, and Song-Chun Zhu. Holistic 3d scene parsing and reconstruction from a single rgb image. In *ECCV*, pages 187–203, 2018. 2
- [25] 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. 2, 5
- [26] Xiaoshui Huang, Guofeng Mei, and Jian Zhang. Feature-metric registration: A fast semi-supervised approach for robust point cloud registration without correspondences. In *CVPR*, pages 11366–11374, 2020. 2
- [27] Yasushi Kanazawa and Hiroshi Kawakami. Detection of planar regions with uncalibrated stereo using distributions of feature points. In *BMVC*, pages 1–10, 2004. 2
- [28] Florian Kluger, Eric Brachmann, Hanno Ackermann, Carsten Rother, Michael Ying Yang, and Bodo Rosenhahn. Consac: Robust multi-model fitting by conditional sample consensus. In *CVPR*, pages 4634–4643, 2020. 1, 2
- [29] Weicheng Kuo, Anelia Angelova, Tsung-Yi Lin, and Angela Dai. Mask2cad: 3d shape prediction by learning to segment and retrieve. In *ECCV*, pages 260–277, 2020. 2
- [30] Jiaxin Li, Ben M Chen, and Gim Hee Lee. So-net: Self-organizing network for point cloud analysis. In *CVPR*, pages 9397–9406, 2018. 3
- [31] Jiehong Lin, Zewei Wei, Zhihao Li, Songcen Xu, Kui Jia, and Yuanqing Li. Dualposenet: Category-level 6d object pose and size estimation using dual pose network with refined learning of pose consistency. In *ICCV*, pages 3560–3569, 2021. 3

- [32] Jiehong Lin, Zewei Wei, Changxing Ding, and Kui Jia. Category-level 6d object pose and size estimation using self-supervised deep prior deformation networks. In *ECCV*, pages 19–34, 2022. [3](#)
- [33] Luca Magri and Andrea Fusiello. T-linkage: A continuous relaxation of j-linkage for multi-model fitting. In *CVPR*, pages 3954–3961, 2014. [2](#), [6](#), [7](#), [8](#)
- [34] Luca Magri and Andrea Fusiello. Multiple model fitting as a set coverage problem. In *CVPR*, pages 3318–3326, 2016. [6](#), [7](#), [8](#)
- [35] Luca Magri, Fusiello Andrea, et al. Robust multiple model fitting with preference analysis and low-rank approximation. In *BMVC*, pages 20–1, 2015. [2](#)
- [36] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, pages 565–571, 2016. [6](#)
- [37] Georgy Ponimatkin, Yann Labbé, Bryan Russell, Mathieu Aubry, and Josef Sivic. Focal length and object pose estimation via render and compare. In *CVPR*, pages 3825–3834, 2022. [3](#)
- [38] 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](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [39] Rasmus Rothe, Matthieu Guillaumin, and Luc Van Gool. Non-maximum suppression for object detection by passing messages between windows. In *ACCV*, pages 290–306, 2015. [5](#)
- [40] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *CVPR*, pages 4938–4947, 2020. [5](#), [6](#)
- [41] Yifei Shi, Junwen Huang, Xin Xu, Yifan Zhang, and Kai Xu. Stablepose: Learning 6d object poses from geometrically stable patches. In *CVPR*, pages 15222–15231, 2021. [3](#)
- [42] Yongzhi Su, Mahdi Saleh, Torben Fetzner, Jason Rambach, Nassir Navab, Benjamin Busam, Didier Stricker, and Federico Tombari. Zebrapose: Coarse to fine surface encoding for 6dof object pose estimation. In *CVPR*, pages 6738–6748, 2022.
- [43] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3d orientation learning for 6d object detection from rgb images. In *ECCV*, pages 699–715, 2018. [3](#)
- [44] Weixuan Tang and Danping Zou. Multi-instance point cloud registration by efficient correspondence clustering. In *CVPR*, pages 6667–6676, 2022. [1](#), [2](#), [6](#), [7](#), [8](#)
- [45] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *ICCV*, pages 6411–6420, 2019. [3](#)
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017. [3](#)
- [47] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *CVPR*, pages 2642–2651, 2019. [3](#)
- [48] Yue Wang and Justin Solomon. Prnet: self-supervised learning for partial-to-partial registration. In *NeurIPS*, pages 8814–8826, 2019. [2](#)
- [49] Yue Wang and Justin M Solomon. Deep closest point: Learning representations for point cloud registration. In *ICCV*, pages 3523–3532, 2019. [2](#)
- [50] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*, 2017. [3](#)
- [51] Hao Xu, Shuaicheng Liu, Guangfu Wang, Guanghui Liu, and Bing Zeng. Omnet: Learning overlapping mask for partial-to-partial point cloud registration. In *ICCV*, pages 3132–3141, 2021. [2](#)
- [52] Jun Yang, Yizhou Gao, Dong Li, and Steven L Waslander. Robi: A multi-view dataset for reflective objects in robotic bin-picking. In *IROS*, pages 9788–9795, 2021. [2](#), [6](#)
- [53] Zi Jian Yew and Gim Hee Lee. Rpm-net: Robust point matching using learned features. In *CVPR*, pages 11824–11833, 2020. [2](#)
- [54] Hao Yu, Fu Li, Mahdi Saleh, Benjamin Busam, and Slobodan Ilic. Cofinet: Reliable coarse-to-fine correspondences for robust point cloud registration. *arXiv preprint arXiv:2110.14076*, 2021. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [55] Mingzhi Yuan, Zhihao Li, Qiuye Jin, Xinrong Chen, and Manning Wang. Pointclm: A contrastive learning-based framework for multi-instance point cloud registration. In *ECCV*, pages 595–611, 2022. [1](#), [2](#), [6](#), [7](#), [8](#)