

Coarse-to-Fine Point Cloud Registration with SE(3)-Equivariant Representations

Cheng-Wei Lin¹, Tung-I Chen¹, Hsin-Ying Lee¹, Wen-Chin Chen¹, and Winston H. Hsu^{1,2}

Abstract—Point cloud registration is a crucial problem in computer vision and robotics. Existing methods either rely on matching local geometric features, which are sensitive to the pose differences, or leverage global shapes, which leads to inconsistency when facing distribution variances such as partial overlapping. Combining the advantages of both types of methods, we adopt a coarse-to-fine pipeline that concurrently handles both issues. We first reduce the pose differences between input point clouds by aligning global features; then we match the local features to further refine the inaccurate alignments resulting from distribution variances. As global feature alignment requires the features to preserve the poses of input point clouds and local feature matching expects the features to be invariant to these poses, we propose an SE(3)-equivariant feature extractor to simultaneously generate two types of features. In this feature extractor, representations that preserve the poses are first encoded by our novel SE(3)-equivariant network and then converted into pose-invariant ones by a pose-detaching module. Experiments demonstrate that our proposed method increases the recall rate by 20% compared to state-of-the-art methods when facing both pose differences and distribution variances.

I. INTRODUCTION

Point cloud registration, the task of finding rigid transformations that align two input point clouds, is fundamental for several applications in computer vision and robotics. Depending on the application domain, point cloud registration methods need to fulfill the demand for various properties [1]. For example, some applications such as SLAM [2], [3] require the registration method to be real-time and accurate. Other applications such as scene reconstruction [4] expect the registration method to be robust to initial pose conditions.

The diversified requirements lead to a variety of registration approaches. Some prior approaches [5]–[7], namely *local* approaches, focused on time-efficiency and accuracy. However, due to the dependence on matching local geometric features, these approaches are sensitive to the magnitude changes in rigid transformations, and thus fail to handle large initial pose differences (Fig. 1-a). On the other hand, *global* approaches [8], [9] leverage global shape information to maintain robust against initial pose difference. Nonetheless, these global approaches usually produce alignment results inferior to those of local approaches when facing *distribution variances*, such as partial overlapping (Fig. 1-b). Distribution variances affect the overall shape that global methods rely on, while regional geometries remain unchanged.

To solve the registration problems with both large initial pose differences and distribution variances, we adopt a coarse-to-fine pipeline that takes advantage of both global

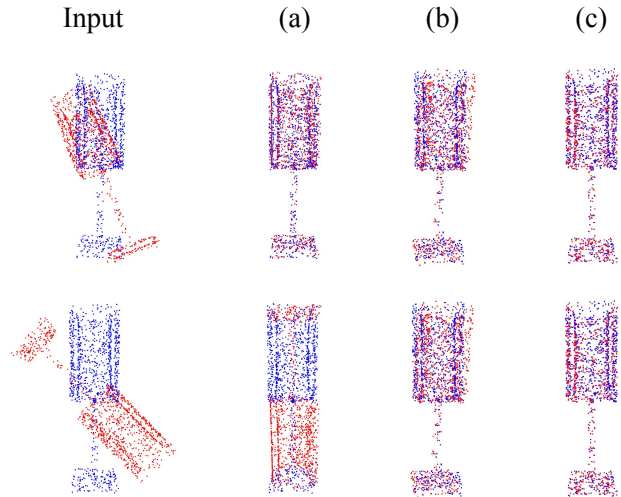


Fig. 1. The registration results of methods from each category. We denote source point clouds as red and target point clouds as blue. (a) Local methods (e.g. IDAM) only produce accurate alignments when initial errors are small enough. (b) Global methods (e.g. DeepGMR) are invariant to initial errors but their performance are unreliable. In contrast, (c) our coarse-to-fine method accurately aligns point clouds regardless of initial errors.

and local approaches (Fig. 2-a). A global register is applied as our coarse-grained register to reduce pose differences and roughly align input point clouds. Then, we utilize a local register to refine the inaccurate alignments caused by distribution variances.

On top of that, we employ a shared feature extractor to generate representations for both global and local registers. The feature alignment process in the global register requires representations to preserve the pose of input point clouds. On the contrary, the correspondence matching in the local register expects representations to get rid of the influences caused by the differences in the poses.

In our feature extractor, an SE(3)-equivariant neural network and a pose-detaching module are proposed to produce pose-preserving features and convert them into pose-invariant ones, respectively (Sec. III-A shows more details about SE(3)-equivariance). Unlike existing SE(3)-equivariant networks [10], [11], our SE(3)-equivariant neural network avoids time-intensive approximations and kernels that constrain the expressiveness of the representation. We preserve the translations by maintaining the center of feature embedding and preserve the rotations using Vector Neuron [12], a rotation-equivariant framework. Furthermore, our pose-detaching module normalizes the input point clouds

¹National Taiwan University, ²Mobile Drive Technology

to remove the translations and utilizes the orthogonality of rotation matrices to eliminate the rotations. Consequently, our novel SE(3)-equivariant feature extractor concurrently produces pose-preserving and pose-invariant representations, supporting both coarse- and fine-grained registers to effortlessly perform registrations.

We evaluate our method on ModelNet40 [13], which is composed of various object models. Following RPMNet [14], we pre-processed the dataset to simulate real-world situations, including sensor noises, independent scans, and partially overlapping point clouds. Furthermore, we evaluate the performance over multiple initial angle ranges to exhibit the influence of initial poses. Experimental results demonstrate that our method outperforms state-of-the-art methods and reaches a reliable performance under simulated real-world scenarios. In addition, ablations (Sec. IV-F) support that our feature extractor satisfies the feature requirements of both global and local registers yet remains time efficient.

To sum up, the overall contributions of this work can be summarized as follows:

- We apply a coarse-to-fine pipeline to resist impacts from both initial pose differences and distribution variances.
- We introduce a novel SE(3)-equivariant feature extractor, simultaneously obtaining representations for both global and local registers.
- Our method outperforms state-of-the-art methods on ModelNet40 under circumstances simulating the real world across different initial pose difference ranges.

II. RELATED WORK

A. Local Registration Methods

Local approaches are often used under circumstances where transformations are known to be small in magnitude. Iterative Closest Point (ICP) [5] iteratively matches the closest points as correspondences and minimizes the distance between these correspondences, which often causes the result to converge at a local minimum. To resolve this problem, a variety of strategies have been proposed to deal with outliers [15], handle the noises [16], or devise better distance metrics [17], [18]. However, the limitation of matching points on Euclidean space leads to recent work performing matching on feature space. PPFNet [19], 3DSmoothNet [20], SpinNet [21], and FCGF [22] follow this idea and solve the Procrustes problem based on the correspondences paired by their representations. Moreover, DCP [6], IDAM [7], RPMNet [14], DGR [23], ImLoveNet [24], and DetarNet [25] use the ground truth poses to supervise point matching and feature learning. Predator [26] and REGTR [27] further leverage the ground truth overlap regions. Another branch of work such as D3Feat [28], DeepVCP [29], PRNet [30], and GeoTransformer [31] leverages key points to enhance the time efficiency. The remaining challenge is that perfect correspondences rarely exist in real-world situations, and thereby recent work [14], [32] utilizes soft matching [33] to work under these conditions. Even so, these local methods still fail to handle large initial perturbations.

B. Global Registration Methods

Unlike local approaches, global approaches are designed to be invariant to the initial transformation error. Some methods such as GO-ICP [34], GOGMA [35], and GOSMA [36] search the SE(3) space using branch-and-bound techniques. Other methods [37]–[39] match the feature with robust optimization. However, these methods are unsuitable for real-time applications due to their large computation time. Fast Global Registration (FGR) [40] is presented to address this issue, achieving a similar speed to that of many local methods. To further improve the accuracy of the registration result, recent work handles the registration problem via learned global representations. DeepGMR [8] represents the global feature through GMM distributions and EquivReg [9] takes rotation-equivariant implicit feature embeddings as its global representations. Nevertheless, these learning-based methods often struggle with distribution differences.

C. Group Equivariant Neural Network

Some research concentrates on proposing group equivariant neural networks as a means of resisting group transformations. For instance, Convolution Neural Network (CNN) [41] are translation-equivariant, resulting in its performance consistency among the same images with different 2D translations. To prevent the effect of rotation, recent studies [11], [42], [43] construct the kernels by some steerable functions. However, these constrained kernels limit the flexibility of the network. Other studies [10], [44] obtain the equivariance property by lifting the input space to higher-dimensional spaces where the group is contained. These studies are time-intensive and cost more computational resources due to the integration of the entire group. Vector Neuron [12] presents a brand new SO(3)-equivariant framework. The major advantage of this framework is the capability of incorporating the SO(3)-equivariance property into existing networks, such as PointNet [45] or DGCNN [46]. We will later see how we design our SE(3)-equivariant feature extractor based on this simple idea, and use the extracted representation to cope with the registrations with notable initial transformations and distribution variances.

III. COARSE-TO-FINE REGISTRATION

Illustrated in Fig. 2, our coarse-to-fine registration pipeline begins with extracting global and local feature representations. These global representations are fed into the global register to estimate a rough alignment between input point clouds. Focusing on roughly aligned point clouds, the local register refines the alignment results given the correspondences formed by matching these local representations.

A. Preliminaries

A function $f: U \rightarrow V$ is equivariant to a set of transformations G , if for any $g \in G$, f and g commutes, i.e., $f(g \cdot u) = g \cdot f(u), \forall u \in U$. For instance, convolution layers are translation-equivariant because the outcome of applying a 2D translation to the input taken by convolution layers is identical to that of applying the 2D translation to the feature

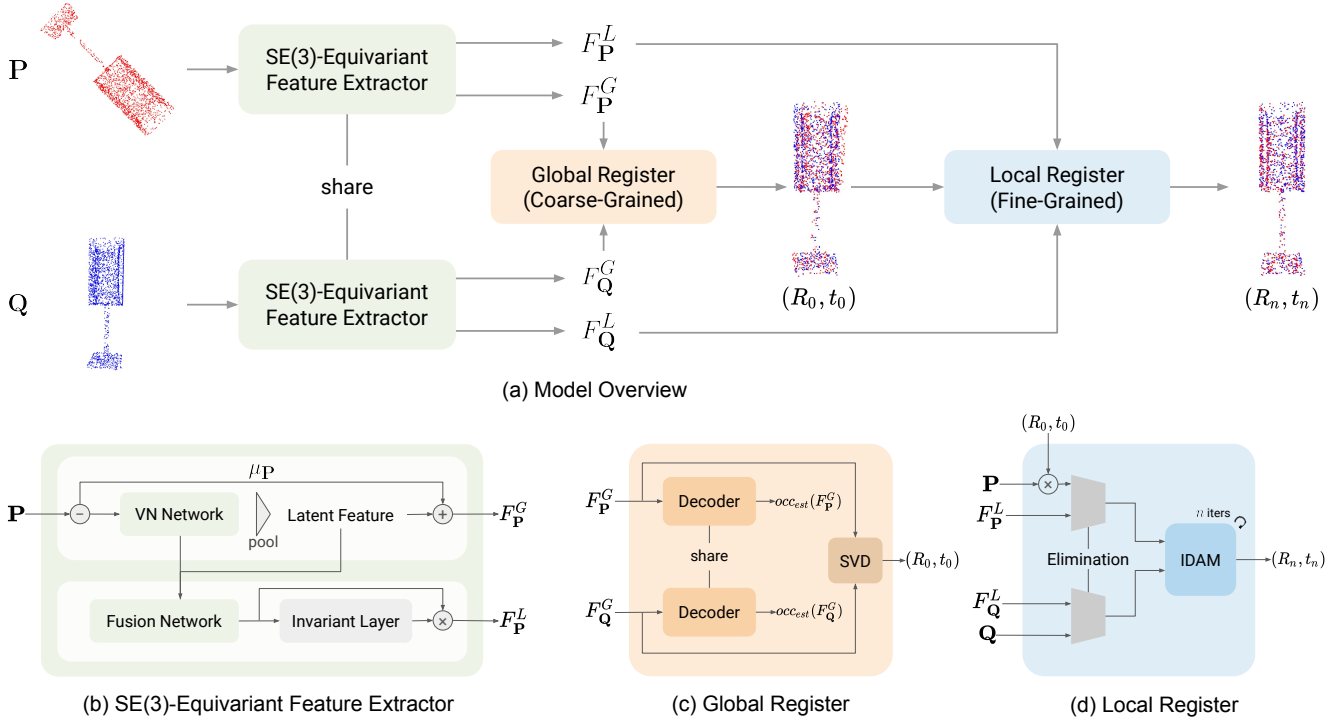


Fig. 2. (a) An overview of the coarse-to-fine pipeline. Given point clouds \mathbf{P} and \mathbf{Q} , (b) our SE(3)-equivariant feature extractor encodes SE(3)-equivariant global representations $F_{\mathbf{P}}^G$ and $F_{\mathbf{Q}}^G$, as well as SE(3)-invariant local representations $F_{\mathbf{P}}^L$ and $F_{\mathbf{Q}}^L$ (Sec. III-B). (c) As SE(3)-equivariant global representations preserve the poses of input point clouds, our global register produce rough alignments (R_0, t_0) by aligning these global representations (Sec. III-C). (d) Taken the roughly aligned point clouds and their SE(3)-invariant local representations, our local register further refines the alignment results by matching these local representations (Sec. III-D). Consequently, our coarse-to-fine pipeline accurately aligns input point clouds by (R_n, t_n) . (Best seen in color.)

map resulting from the convolution layers. Therefore, we can say a network is SE(3)-equivariant if and only if the network \mathbf{nn} fulfills that for any pair of rotation $\mathbf{R} \in SO(3)$ and translation $\mathbf{t} \in \mathbb{R}^3$,

$$\mathbf{nn}(\mathbf{R}\mathbf{x} + \mathbf{t}) = \mathbf{R} \cdot \mathbf{nn}(\mathbf{x}) + \mathbf{t}, \quad (1)$$

where \mathbf{x} is the input 3D representations. According to the above definition, our work is different from Neural Descriptor Fields [47], which removes the rotations and translations rather than preserves them.

B. SE(3)-Equivariant Feature Extractor

To prepare for registration, we extract global features that represent the overall shape and local features that summarize the regional geometries. The feature alignment process in the global register expects the global features to preserve the SE(3) transformations of input point clouds. In contrast, the correspondence matching process in the local register requires the local features to avoid the impact of these transformations. To this end, our feature extractor comprises two modules, one managing to produce SE(3)-equivariant global representations $F_{\mathbf{P}}^G \in \mathbb{R}^{C_0 \times 3}$ and the other aiming to generate SE(3)-invariant local representations $F_{\mathbf{P}}^L \in \mathbb{R}^{N \times C_1}$.

1) *SE(3)-Equivariant Global Features*: As the feature alignment process in the global register requires features to preserve the poses of input point clouds, we introduce a module to encode SE(3)-equivariant global representations.

Depicted in the upper row of Fig. 2-b, this module is composed of three operations: maintaining the mean of input point clouds to preserve translations, passing through an SO(3)-equivariant network to preserve rotations, and average-pooling point-wise features to respect the fact that point clouds are not ordered. We maintain the mean by subtracting it at the beginning of the network and adding it back to the final global representation; hence preserve the translations. Next, we create a SO(3)-equivariant network by plugging Vector Neuron (VN) [12], a rotation-equivariant framework, into a PointNet-based backbone, thus preserving the rotation and flexibly representing the attributes of each point. Lastly, the average-pooling approximates a permutation symmetry function, resulting in our global feature being permutation-invariant. With the above three operations, we encode feature representations that are SE(3)-equivariant and permutation-invariant. Denote the source point cloud as \mathbf{P} and the target point cloud as $\mathbf{Q} = \mathbf{R}\mathbf{M}\mathbf{P} + \mathbf{t}$, where \mathbf{M} is a permutation matrix. The relationship between their global representations $F_{\mathbf{P}}^G$ and $F_{\mathbf{Q}}^G$ is shown as follows:

$$F_{\mathbf{Q}}^G = \mathbf{R} \cdot F_{\mathbf{P}}^G + \mathbf{t}. \quad (2)$$

2) *SE(3)-Invariant Local Features*: Additionally, we introduce another module to generate SE(3)-invariant local features for the local register to match correspondence points according to these features. Displayed in the lower row of Fig. 2-b, we follow PointNet [45] and fuse the global features

into the local features by a VN-based fusion network. Then the fused feature embeddings FR (R denotes the preserved rotation) are made invariant to rotations by an invariant layer, which, based on VN, transforms FR into vectors VR . With the operation of multiplying FR by the transpose of VR , we obtain rotation-invariant local representations:

$$(FR)(VR)^T = FRR^T V^T = FV^T. \quad (3)$$

Moreover, these representations are also translation-invariant due to the mean subtraction mentioned in Sec. III-B.1. As a result, these representations are invariant to any transformation from SE(3) group. That is, for point clouds \mathbf{P} and $\mathbf{S} = \mathbf{R}\mathbf{P} + \mathbf{t}$, their local representations $F_{\mathbf{P}}^L$ and $F_{\mathbf{S}}^L$ are identical to each other. Extracting both SE(3)-equivariant global features and SE(3)-invariant local features, we fulfill the requirements of the following global and local registers.

C. Global Register (Coarse-Grained)

Given the SE(3)-equivariant global representations (Sec. III-B.1), our coarse-grained register aims to reduce the pose differences in registration problems. These global representations have the same poses as the input point clouds; therefore we can simply solve the registration task by aligning their global representations when the point clouds only differ in their pose (Sec. III-C.1). To further deal with the noisy point clouds from different scans, we follow [9] and employ the implicit representation loss (Sec. III-C.2) as well as the registration loss (Sec. III-C.3) to our training progress.

1) *Feature Alignment*: Our coarse-grained register operates registration on the global feature space since the global representations preserve the poses of input point clouds. Furthermore, these global representations are automatically paired due to the permutation-invariance property; thus the data association issue is removed, i.e., finding the correspondence point pairs between two point clouds is not necessary anymore. As a consequence, the registration problem of point clouds \mathbf{P} and $\mathbf{Q} = \mathbf{R}\mathbf{M}\mathbf{P} + \mathbf{t}$ can be solved by finding the rotation and translation between their global representations $F_{\mathbf{P}}^G = \{f_{\mathbf{P},1}^G, \dots, f_{\mathbf{P},c_0}^G\}$ and $F_{\mathbf{Q}}^G = \{f_{\mathbf{Q},1}^G, \dots, f_{\mathbf{Q},c_0}^G\}$, that is, optimizing the following objective function:

$$\operatorname{argmin}_{R, \mathbf{t}} \frac{1}{C_0} \sum_{i=1}^{C_0} \|\mathbf{R}f_{\mathbf{P},i}^G + \mathbf{t} - f_{\mathbf{Q},i}^G\|^2. \quad (4)$$

As validated in ICP [5], this least square optimization can be solved in closed form by using single value decomposition.

2) *Implicit Representation Loss*: The above feature alignment operation allows our coarse-grained register to align point clouds that differ in only permutations and rigid transformations, which is far from useful. In real-world applications, point clouds are captured from different scans, which do not result in identical sets of points. Inspired by EquivReg [9], we construct an encoder-decoder network, in which the encoder is the aforementioned SE(3)-equivariant feature extractor. The decoder estimates the occupancy value $occ_{est}(p; F_{\mathbf{P}}^G) \in [0, 1]$ of a queried position $p \in \mathbb{R}^3$ according to the encoded shape representation $F_{\mathbf{P}}^G \in \mathbb{R}^{C_0 \times 3}$

of point cloud \mathbf{P} . In practice, we query n sampled positions for the implicit shape reconstruction and defined the loss of the occupancy value prediction as:

$$L_{occ} = \sum_{i=1}^n \mathbf{cross_entropy}(occ_{est}(p_i; F_{\mathbf{P}}^G), occ_{gt}(p_i)), \quad (5)$$

where $occ_{gt}(p) \in \{0, 1\}$ is the ground truth occupancy value of the position p . With this network design, the global representations of point clouds captured from the same geometry are implicitly motivated to be similar.

3) *Registration Loss*: Since these global representations are similar but not exactly the same, the slight difference among these point clouds may lead to considerable errors in registration outcomes. Hence, we include the registration step in our training progress to further encourage precise registration. With the estimated rigid transformation ($\mathbf{R}_{est}, \mathbf{t}_{est}$), we follow DCP [6] and use the following registration loss to measure our coarse-grained register's agreement to the ground-truth rigid transformations ($\mathbf{R}_{gt}, \mathbf{t}_{gt}$):

$$L_{reg} = \|\mathbf{R}_{gt}^T \mathbf{R}_{est} - \mathbf{I}_3\|_F^2 + \|\mathbf{t}_{gt} - \mathbf{t}_{est}\|_2^2, \quad (6)$$

where $\|\cdot\|_F$ denotes the Frobenius norm and $\|\cdot\|_2$ denotes the Euclidean norm. Consequently, our coarse-grained register produces rough registration results ($\mathbf{R}_0, \mathbf{t}_0$), which will be later refined by our fine-grained register.

D. Local Register (Fine-Grained)

Since the original large-error registration tasks have been degraded into small-error ones with the global register (Sec. III-C), we here present a local register that focuses on refining the small residual perturbation. Considering speed and simplicity, we apply [7], a learning-based local method, as our fine-grained register. As illustrated in 2-d, the network majorly consists of a hard elimination (Sec. III-D.1) and an Iterative Distance-Aware Similarity Matrix Convolution Network (IDAM) (Sec. III-D.2).

1) *Hard Elimination*: To improve speed and filter out ambiguous regions, we employ an elimination network in the fine-grained register. According to the local representations, this elimination network gives every point a score and preserves the points along with their local representations of which the scores are in the top 1/6. As claimed in [7], this process leaves the prominent points, such as corner points, and thereby simplifies the matching procedure later in IDAM.

2) *IDAM*: Taking the eliminated points and representations, IDAM particularly learns two items, a similarity matrix and a weight vector. The similarity matrix presents the closeness of each source point and each target point, which can be used to find the corresponding points. The confidence in these correspondent points are approximated by the aforementioned weight vector. With the correspondent points and the confidence, a refined alignment result can be calculated by orthogonal Procrustes algorithms such as weighted SVD. Repeating this process n times, we eventually get the accurate registration results ($\mathbf{R}_n, \mathbf{t}_n$).

TABLE I
RECALL RATES ON CLEAN DATA (FIG. 3-A)

Methods	Recall rates across different angle ranges			
	[0, 45]	[0, 90]	[0, 135]	[0, 180]
IDAM [7]	99.3	98.2	96.9	97.0
RPMNet [14]	99.7	91.6	67.8	49.3
Predator [26]	99.8	91.2	53.9	33.6
DeepGMR [8]	100.0	100.0	100.0	100.0
EquivReg [9]	100.0	100.0	100.0	100.0
Ours	100.0	100.0	100.0	100.0

IV. EXPERIMENTS

A. Dataset: ModelNet40

Our network is trained on ModelNet40 [13], which consists of 12311 CAD models from 40 categories of objects. We follow the official split, taking 9843 models as the training set and the other 2468 models as the testing set. Besides, we use the tool provided by Stutz and Geiger [48] to make these models watertight. Given the watertight models, we determine the occupancy value of any 3D coordinate depending on whether it is inside or outside the model. As for the transformations between input point clouds, we follow the procedure in DCP [6]. That is, taking the specified maximum angle θ and maximum distance d as arguments, our initial rotation and translation are determined by three Euler angles sampled in the range $[0, \theta]$ and three coordinates sampled in the range $[-d, d]$, respectively.

B. Training and Testing Details

Our training process is composed of two stages. First, we trained our coarse-grained register together with our feature extractor. Then, we fixed the global feature module and trained the fine-grained register. For evaluation, we set the maximum distance for the initial translation to 0.5 across all settings. To demonstrate the robustness to initial pose differences, we tested the algorithms over initial rotations with maximum angles of 45° , 90° , 135° , and 180° .

C. Evaluation Metrics

We evaluate our registration result by computing the evaluation metric provided by RPMNet [14]:

$$\mathbf{Err}(\mathbf{R}) = \arccos\left(\frac{\text{tr}(\mathbf{R}_{gt}^T \mathbf{R}_{est}) - 1}{2}\right), \mathbf{Err}(\mathbf{t}) = \|\mathbf{t}_{gt} - \mathbf{t}_{est}\|_2, \quad (7)$$

where $\mathbf{Err}(\mathbf{R})$ denotes the rotation errors and $\mathbf{Err}(\mathbf{t})$ denotes the translation errors. In the following sections, we present the results as the recall rates of the rotation errors being less than 5° and the translation errors being less than 0.2.

D. Baselines

To give readers a better context of differences between local and global registration approaches, we compared our method with some representative approaches in each category. IDAM [7], RPMNet [14] and Predator [26] are the baselines representing local methods, while DeepGMR [8]

TABLE II
RECALL RATES ON NOISY DATA (FIG. 3-B)

Methods	Recall rates across different angle ranges			
	[0, 45]	[0, 90]	[0, 135]	[0, 180]
IDAM [7]	91.6	45.1	21.8	15.6
RPMNet [14]	91.4	62.0	33.8	25.5
Predator [26]	98.0	59.6	21.8	10.7
DeepGMR [8]	83.1	84.7	82.9	83.9
EquivReg [9]	89.0	89.0	88.9	89.3
Ours	98.8	98.9	98.8	98.9




Fig. 3. (a) Clean data: point clouds includes only permutation and pose difference. (b) Noisy data: point clouds are jittered by Gaussian noises, simulating sensor noises. (c) Independently sampled data: point clouds are separately sampled from the object surface, simulating independent scans. (d) Partially overlapping data: point clouds are cropped by farthest point sampling, simulating scans from different views.

and EquivReg [9] are the baselines representing global methods. It's worth mentioning that since normal estimations are usually inaccurate in the real world [8], we remove the input point pair features, which are calculated according to the normal, from RPMNet [14]. Besides, we did not translate the input of the EquivReg [9] because of its limited support for only rotational registration tasks. Its recall rates are hence formulated exclusively using rotation errors.

E. Performance on Multiple Point Cloud Settings

As illustrated in Fig. 3, we conducted our experiments with point clouds under multiple settings to simulate different real world scenarios, including clean, noisy, independently sampled, and partially overlapping point clouds.

1) *Clean Data*: We first test the registration of a point cloud and its permuted and transformed copy. The point clouds comprise 1024 points randomly sampled from the outer surface of each model. Under this setting, every point in the source point cloud has exactly one corresponding point in the target point cloud. Registration results are presented in Table I. We can see that global methods and our method achieve recall rates of 100% across different initial angle ranges. As for the local methods, IDAM performs appropriately while RPMNet barely handles small angle range cases.

2) *Noisy Data*: To simulate the noises caused by the sensor, we further jittered the points in both point clouds with noises sampled from $N(0, 0.01)$ and clipped to $[-0.05, 0.05]$. As shown in Table II, local methods share a similar trend of recall rate decrement along with the increment of initial angle ranges. On the other hand, global methods reach high performance across different angle ranges. Therefore, our

TABLE III

RECALL RATES ON INDEPENDENTLY SAMPLED DATA (FIG. 3-C)

Methods	Recall rates across different angle ranges			
	[0, 45]	[0, 90]	[0, 135]	[0, 180]
IDAM [7]	75.2	20.2	6.9	2.9
RPMNet [14]	90.5	58.4	22.1	11.2
Predator [26]	99.5	85.9	46.4	25.2
DeepGMR [8]	20.9	19.6	19.0	19.3
EquivReg [9]	43.8	44.0	44.4	42.6
Ours	90.6	89.5	89.9	90.2

TABLE IV

RECALL RATES ON PARTIALLY OVERLAPPING DATA (FIG. 3-D)

Methods	Recall rates across different angle ranges			
	[0, 45]	[0, 90]	[0, 135]	[0, 180]
IDAM [7]	98.8	69.4	37.6	27.1
RPMNet [14]	90.2	54.4	23.9	13.5
Predator [26]	96.4	53.2	16.7	8.8
DeepGMR [8]	46.7	48.2	48.2	46.6
EquivReg [9]	51.9	51.5	49.2	48.4
Ours	82.1	79.3	78.0	78.3

method can produce reliable coarse results and lead to even better results after the refinement.

3) *Independently Sampled Data*: Instead of repeating points, real world scans generally contain distinct point sets despite being captured from the same view. Simulating these real world scenarios, we *independently* sampled the 1024 points in source and target point clouds. The results in Table III reveal that global methods cannot handle these perturbations, and result in a significant performance drop. In contrast, local methods reach performances similar to those under conditions with Gaussian noises. On top of that, our method achieves consistent results regardless of the initial conditions, with only slightly poorer outcomes than those in Gaussian noise cases. Particularly, Predator reaches a better performance than ours on small initial angle ranges because the backbone of our fine-grained register is IDAM, whose results are inferior to those of Predator. We expect that changing the backbone of the fine-grained register will further enhance our performance.

4) *Partially Overlapping Data*: In most real-world applications, point clouds are captured from a single view, leading to incomplete shapes information on objects. To simulate these conditions, we followed RPMNet [14] and sub-sampled the point clouds using the farthest point sampling algorithm, retaining approximately 75% of the points. Table IV demonstrates that our method works properly under all initialization despite the slight degradation due to the information loss of the overall shape. Baseline methods either overfitted on cases with constrained angle ranges or failed to achieve reliable outcomes. Notably, local methods search for the answers in the constrained space rather than in the entire SE(3) space, which leads to their high performance in this constrained

TABLE V

ABLATION STUDY ON EQUIVARIANCE AND INVARIANCE

VN	Invariant Layer	Shared Encoder	Recall(%)	Speed(fps)
			17.9	17.8
	✓		21.4	17.8
✓			73.2	17.9
✓	✓		89.5	17.8
✓	✓	✓	90.2	36.4

space but poor performance in spaces other than that.

F. Ablations

To clarify the efficacy and efficiency of our feature extractor, we presented the ablation studies in Table V. These ablation studies are evaluated under the independently sampled setting using a single Titan RTX GPU.

Vector Neuron(VN) is used to generate SE(3)-equivariant global features, and the invariant layer is used to produce SE(3)-invariant local features. Results show that either non-equivariant global features or non-invariant local features will significantly decrease the recall rate. As a consequence, we confirm that global methods and local methods indeed expect their input feature representations to be SE(3)-equivariant and SE(3)-invariant, respectively.

Table V also indicates that using dual encoders to extract these features will be time-consuming, and thus cannot be used in real-time applications. Our shared SE(3)-equivariant encoder, on the other hand, produces over 30 fps without harming the efficacy. Moreover, we can further increase our speed by processing the local feature module (Sec. III-B.2) and the global register (Sec. III-C) in parallel.

V. CONCLUSION

In this paper, we proposed a coarse-to-fine point cloud registration method leveraging the representations extracted by an SE(3)-equivariant encoder. Our coarse-to-fine pipeline integrates the advantages of global and local registration methods, resulting in accurate alignment despite considerable initial transformations and severe distribution variances. Additionally, our SE(3)-equivariant encoder extracts the features fulfilling the necessity of the coarse-to-fine pipeline yet maintains the time efficiency required for real-time applications.

We conduct performance evaluations on data simulating the point clouds in real-world scenarios. Experiment results show that our network handles various input imperfections while still having room for improvement. For instance, when facing extreme noises, our coarse-grained register may fail to produce reliable results, which leaves substantial pose differences for our fine-grained register to handle. This problem is left for future work to solve.

VI. ACKNOWLEDGEMENT

This work was supported in part by National Science and Technology Council, Taiwan, under Grant NSTC 111-2634-F-002-022, and Mobile Drive Technology Co., Ltd (MobileDrive). We are grateful to the National Center for High-performance Computing.

REFERENCES

- [1] D. Bauer, T. Patten, and M. Vincze, "Reagent: Point cloud registration using imitation and reinforcement learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 14 586–14 594.
- [2] P. Kim, J. Chen, and Y. K. Cho, "Slam-driven robotic mapping and registration of 3d point clouds," *Automation in Construction*, vol. 89, pp. 38–48, 2018.
- [3] D. Cattaneo, M. Vaghi, and A. Valada, "Lcdnet: Deep loop closure detection and point cloud registration for lidar slam," *IEEE Transactions on Robotics*, 2022.
- [4] R. Y. Takimoto, M. d. S. G. Tsuzuki, R. Vogelaar, T. de Castro Martins, A. K. Sato, Y. Iwao, T. Gotoh, and S. Kagei, "3d reconstruction and multiple point cloud registration using a low precision rgb-d sensor," *Mechatronics*, vol. 35, pp. 11–22, 2016.
- [5] P. J. Besl and N. D. McKay, "Method for registration of 3-d shapes," in *Sensor fusion IV: control paradigms and data structures*, vol. 1611. Spie, 1992, pp. 586–606.
- [6] Y. Wang and J. M. Solomon, "Deep closest point: Learning representations for point cloud registration," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 3523–3532.
- [7] J. Li, C. Zhang, Z. Xu, H. Zhou, and C. Zhang, "Iterative distance-aware similarity matrix convolution with mutual-supervised point elimination for efficient point cloud registration," in *European conference on computer vision*. Springer, 2020, pp. 378–394.
- [8] W. Yuan, B. Eckart, K. Kim, V. Jampani, D. Fox, and J. Kautz, "Deepgmr: Learning latent gaussian mixture models for registration," in *European conference on computer vision*. Springer, 2020, pp. 733–750.
- [9] M. Zhu, M. Ghaffari, and H. Peng, "Correspondence-free point cloud registration with so(3)-equivariant implicit shape representations," in *Conference on Robot Learning*. PMLR, 2022, pp. 1412–1422.
- [10] M. Finzi, S. Stanton, P. Izmailov, and A. G. Wilson, "Generalizing convolutional neural networks for equivariance to lie groups on arbitrary continuous data," in *International Conference on Machine Learning*. PMLR, 2020, pp. 3165–3176.
- [11] F. Fuchs, D. Worrall, V. Fischer, and M. Welling, "Se (3)-transformers: 3d roto-translation equivariant attention networks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 1970–1981, 2020.
- [12] C. Deng, O. Litany, Y. Duan, A. Poulencard, A. Tagliasacchi, and L. J. Guibas, "Vector neurons: A general framework for so (3)-equivariant networks," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 12 200–12 209.
- [13] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao, "3d shapenets: A deep representation for volumetric shapes," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1912–1920.
- [14] Z. J. Yew and G. H. Lee, "Rpm-net: Robust point matching using learned features," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 824–11 833.
- [15] D. Chetverikov, D. Stepanov, and P. Krsek, "Robust euclidean alignment of 3d point sets: the trimmed iterative closest point algorithm," *Image and vision computing*, vol. 23, no. 3, pp. 299–309, 2005.
- [16] A. W. Fitzgibbon, "Robust registration of 2d and 3d point sets," *Image and vision computing*, vol. 21, no. 13-14, pp. 1145–1153, 2003.
- [17] K.-L. Low, "Linear least-squares optimization for point-to-plane icp surface registration," *Chapel Hill, University of North Carolina*, vol. 4, no. 10, pp. 1–3, 2004.
- [18] A. Segal, D. Haehnel, and S. Thrun, "Generalized-icp," in *Robotics: science and systems*, vol. 2, no. 4. Seattle, WA, 2009, p. 435.
- [19] H. Deng, T. Birdal, and S. Ilic, "Ppfnet: Global context aware local features for robust 3d point matching," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 195–205.
- [20] Z. Gojcic, C. Zhou, J. D. Wegner, and A. Wieser, "The perfect match: 3d point cloud matching with smoothed densities," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 5545–5554.
- [21] S. Ao, Q. Hu, B. Yang, A. Markham, and Y. Guo, "Spinnet: Learning a general surface descriptor for 3d point cloud registration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 11 753–11 762.
- [22] C. Choy, J. Park, and V. Koltun, "Fully convolutional geometric features," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 8958–8966.
- [23] C. Choy, W. Dong, and V. Koltun, "Deep global registration," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2514–2523.
- [24] H. Chen, Z. Wei, Y. Xu, M. Wei, and J. Wang, "Imlovenet: Misaligned image-supported registration network for low-overlap point cloud pairs," in *ACM SIGGRAPH 2022 Conference Proceedings*, 2022, pp. 1–9.
- [25] Z. Chen, F. Yang, and W. Tao, "Detarnet: Decoupling translation and rotation by siamese network for point cloud registration," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 1, 2022, pp. 401–409.
- [26] S. Huang, Z. Gojcic, M. Usvyatsov, A. Wieser, and K. Schindler, "Predator: Registration of 3d point clouds with low overlap," in *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 2021, pp. 4267–4276.
- [27] Z. J. Yew and G. H. Lee, "Regtr: End-to-end point cloud correspondences with transformers," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6677–6686.
- [28] X. Bai, Z. Luo, L. Zhou, H. Fu, L. Quan, and C.-L. Tai, "D3feat: Joint learning of dense detection and description of 3d local features," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 6359–6367.
- [29] W. Lu, G. Wan, Y. Zhou, X. Fu, P. Yuan, and S. Song, "Deepvcv: An end-to-end deep neural network for point cloud registration," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 12–21.
- [30] Y. Wang and J. M. Solomon, "Prnet: Self-supervised learning for partial-to-partial registration," *Advances in neural information processing systems*, vol. 32, 2019.
- [31] Z. Qin, H. Yu, C. Wang, Y. Guo, Y. Peng, and K. Xu, "Geometric transformer for fast and robust point cloud registration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 11 143–11 152.
- [32] A. Drory, T. Shomer, S. Avidan, and R. Giryes, "Best buddies registration for point clouds," in *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [33] S. Gold, C.-P. Lu, A. Rangarajan, S. Pappu, and E. Mjølness, "New algorithms for 2d and 3d point matching: Pose estimation and correspondence," *Advances in neural information processing systems*, vol. 7, 1994.
- [34] J. Yang, H. Li, D. Campbell, and Y. Jia, "Go-icp: A globally optimal solution to 3d icp point-set registration," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 11, pp. 2241–2254, 2015.
- [35] D. Campbell and L. Petersson, "Gogma: Globally-optimal gaussian mixture alignment," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 5685–5694.
- [36] D. Campbell, L. Petersson, L. Kneip, H. Li, and S. Gould, "The alignment of the spheres: Globally-optimal spherical mixture alignment for camera pose estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 11 796–11 806.
- [37] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [38] M. B. Horowitz, N. Matni, and J. W. Burdick, "Convex relaxations of se (2) and se (3) for visual pose estimation," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 1148–1154.
- [39] H. Maron, N. Dym, I. Kezurer, S. Kovalsky, and Y. Lipman, "Point registration via efficient convex relaxation," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 4, pp. 1–12, 2016.
- [40] Q.-Y. Zhou, J. Park, and V. Koltun, "Fast global registration," in *European conference on computer vision*. Springer, 2016, pp. 766–782.
- [41] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [42] T. S. Cohen and M. Welling, "Steerable cnns," *arXiv preprint arXiv:1612.08498*, 2016.
- [43] C. Esteves, C. Allen-Blanchette, A. Makadia, and K. Daniilidis, "Learning so (3) equivariant representations with spherical cnns," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 52–68.

- [44] R. Kondor and S. Trivedi, "On the generalization of equivariance and convolution in neural networks to the action of compact groups," in *International Conference on Machine Learning*. PMLR, 2018, pp. 2747–2755.
- [45] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 652–660.
- [46] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *Acm Transactions On Graphics (tog)*, vol. 38, no. 5, pp. 1–12, 2019.
- [47] A. Simeonov, Y. Du, A. Tagliasacchi, J. B. Tenenbaum, A. Rodriguez, P. Agrawal, and V. Sitzmann, "Neural descriptor fields: Se (3)-equivariant object representations for manipulation," *arXiv preprint arXiv:2112.05124*, 2021.
- [48] D. Stutz and A. Geiger, "Learning 3d shape completion under weak supervision," *CoRR*, vol. abs/1805.07290, 2018. [Online]. Available: <http://arxiv.org/abs/1805.07290>