

Slice Transformer and Self-supervised Learning for 6DoF Localization in 3D Point Cloud Maps

Muhammad Ibrahim¹, Naveed Akhtar¹, Saeed Anwar², Michael Wise¹ and Ajmal Mian¹




Fig. 1: The proposed self-supervised pretext task shuffles 3D point cloud quadrants to create pseudo-labels. Quadrant permutations result in 24 labels. 36 slices are generated from a single raw frame, covering 30° rotation that includes a 20° overlap with the neighboring slices - overlap not shown.

Abstract—Precise localization is critical for autonomous vehicles. We present a self-supervised learning method that employs transformers for the first time for the task of outdoor localization using LiDAR data. We propose a pre-text task that reorganizes the slices of a 360° LiDAR scan to leverage its axial properties. Our model, called Slice Transformer, employs multi-head attention while systematically processing the slices. To the best of our knowledge, this is the first instance of leveraging multi-head attention for outdoor point clouds. We additionally introduce the Perth-WA dataset, which provides a large-scale LiDAR map of Perth city in Western Australia, covering ~4km² area. Localization annotations are provided for Perth-WA. The proposed localization method is thoroughly evaluated on Perth-WA and Appollo-SouthBay datasets. We also establish the efficacy of our self-supervised learning approach for the common downstream task of object classification using ModelNet40 and ScanNN datasets. The code and Perth-WA data will be publicly released.

I. INTRODUCTION

Six degrees of freedom (6DoF) localization of vehicles is a key task for autonomous driving. Satellite-based localization lacks the required precision and does not work inside cities with tall structures, bridges and tunnels. To achieve the required level of precision in the mentioned scenarios, 3D LiDAR-based localization is an optimal choice.

A general approach towards LiDAR-based localization is to construct an offline 3D map and query the map with LiDAR frames during online navigation. Conventional techniques under this paradigm [1], [2] leverage frame registration. However, due to their large computational requirements, they are unable to provide a practical solution. More recently, deep learning based techniques have shown promising results. Among these methods, matching deep learning features

of an input frame with the deep features of a pre-computed map is a viable solution [3], [4]. Nevertheless, such methods also require post-processing with registration to achieve the desired precision level [5].

To contain the prediction latency within practical limits, the state-of-the-art methods compress the map into a neural model, and directly localize vehicles by regressing its 6DoF pose over the map [6]. For 3D point clouds, this demands high-fidelity representation learning by the model because of the unstructured nature of the data. Though effective, conventional deep learning methods based on multi-layer perceptrons [7] and convolutional networks [8] still fall short on the accuracy required for the critical task of 6DoF localization with LiDAR frames.

transformers [9] have recently surpassed conventional deep learning methods in performance. However, their extension to point cloud data is not straight-forward. In fact, currently, there is no widely known technique that can leverage the key strength of transformers, i.e., multi-head attention, for outdoor point cloud data. An additional challenge is that due to their low inductive bias, transformers also require an even larger amount of data than convolutional networks.

In this work, we make three major contributions to achieve highly precise 6DoF localization in 3D point clouds. First, we propose a pre-text task for self-supervised learning on point clouds, see Fig. 1. Our task exploits the intrinsic axial nature of the LiDAR data by systematic slicing of each frame. The slices are shuffled to allow pre-training on a very large amount of data in a self-supervised way, which is especially conducive for learning effective Transformer models. Second, we propose a first-of-its-kind based Transformer model that enables leveraging multi-head attention for outdoor point clouds. Specifically suited to our pre-text task, the model processes a LiDAR frame by slicing it, hence termed Slice Transformer. Third, we introduce a 3D LiDAR map of the Perth city in Western Australia that covers ~ 4 km² of the Central Business District (CBD), providing annotations for 6DoF localization as well as additional frames for self-supervised learning.

Professor Ajmal Mian is the recipient of an Australian Research Council Future Fellowship Award (project number FT210100268) funded by the Australian Government. Dr. Naveed Akhtar is recipient of an Office of National Intelligence National Intelligence Postdoctoral Grant (project number NIPG-2021-001) funded by the Australian Government.

¹Department of Computer Science, The University of Western Australia. muhammad.ibrahim@research, naveed.akhtar@, michael.wise@, ajmal.mian@) uwa.edu.au

²King Fahad University of Petroleum and Minerals (KFUPM), Dhahran, KSA, saeed.anwar@kfupm.edu.sa

We establish the baseline results for 6DoF localization on our dataset with the proposed method and PointLoc [6]. We also provide benchmarking of our technique on an existing Appollo-Southbay dataset [10]. Moreover, leveraging the backbone of our model, we demonstrate the effectiveness of our model for point cloud object classification task with ModelNet40 [11] and ScanObjectNN [12] datasets.

II. RELATED WORK

6DoF localization is an important task for self-driving vehicles [13], [14]. Whereas conventional methods for matching point cloud frames for such a task have used registration techniques [1], [15], more recent works focus on exploiting deep learning to relate an input frame to a 3D map [6], [10], [16], [17]. Among these methods, there are contributions that compress the map into a neural model and use that model as a 6DoF pose predictor for the vehicle [6], [10]. Using raw LiDAR frames, this prediction is particularly challenging due to the unstructured nature of the data, which conflicts with the high precision requirements of the task. Hence, it is imperative to continuously improve the neural models for this task along with the progress of deep learning.

Recent literature has seen Vision Transformer [9] and its variants, e.g., [18], [19], [20], [21], [22], [23], to outperform their convolutional and multi-layer perceptron based counterparts on a variety of tasks. Consequently, Transformer architectures have also started to emerge for point cloud processing. For instance, [24] and [25] employ transformers for 3D object detection, whereas [26] developed a method for point cloud segmentation. However, due to the complex nature of outdoor LiDAR data, Transformer models are yet to establish themselves on outdoor benchmarks.

In general, Transformer architectures require a huge amount of training data to learn effective models. In the domain of point clouds, this need is partially filled with self-supervised learning [27], [28], [29]. However, the pre-text tasks are designed keeping in mind the downstream tasks of object classification and segmentation. Representative examples of pre-text tasks in point cloud domain include self-shape correction of 3D CAD models [30], self-domain adaptation [31] and point cloud rotation [32]. In general, the issue with the self-supervised methods for point cloud processing is that the pre-training is performed either on images [29], 3D synthetic data [28], [31] or indoor point clouds [27], which is not particularly helpful for outdoor downstream applications like localization.

Contemporary self and un-supervised deep learning based point cloud techniques employ CNNs and MLPs. Nubert et al. [16] proposed a self-supervised method to predict poses from input LiDAR frames. Similarly, an un-supervised projection method was proposed by Cho et al. [33] to predict the poses from input. Both methods exploit structured 2D CNNs to process 3D point clouds indirectly. SelfVoxelO [34] exploits 3D CNNs to process the raw point cloud directly for pose estimation. Similarly, UnPWC-SVDLO [35] utilizes scene flow estimation network PointPWC [36] as a backbone for its un-supervised LiDAR odometry.

Leveraging the representation prowess of transformers for point cloud analysis is still under-explored. Along this nascent direction, encouraging results have started to emerge [37], [38]. However, the complexity of data, especially in the outdoor domain, has still not allowed researchers to exploit the key strength of transformers, i.e., multi-head attention, in their techniques. Current methods are limited to simple indoor scenes and synthetic objects. Our work fills this gap by enabling multi-head attention for outdoor tasks on challenging LiDAR scans. Moreover, we introduce a pre-text learning task especially suited to outdoor LiDAR data and present the first instance of outdoor localization with transformers in the domain of point clouds.

III. SELF-SUPERVISED LEARNING

A. Pre-text Task

Due to the complex nature of outdoor LiDAR data, it is imperative to induce models with a large amount of training samples. However, annotating a large number of LiDAR frames for any downstream task can become prohibitive. To side-step the issue, we propose a pre-text task to automate the labeling process with pseudo-labels. The central idea of the proposed pre-text task is illustrated in Fig. 1.

We first divide a LiDAR frame axially into 36 slices. These slices are split into four quadrants of 9 slices each. We generate a new class label by shuffling the quadrants, resulting in 24 pseudo-labels. Each slice contains 10° of new region and a 20° overlap with its neighboring slices. Our treatment of a frame in slices is governed by the tokenization of input required by the Transformer architecture. The slices naturally result in well-defined tokens. We provide the architectural details of our network in Section III-B. Using the pre-text task, we used 250K raw LiDAR frames to pre-train our model using the standard cross-entropy loss. These frames are taken from the proposed Perth-WA dataset (see Sec. V). The model is trained for 100 epochs on 240K raw frames and 10K frames are used for validation. Our model achieved around 95% classification accuracy on the validation set.

B. Proposed Slice Transformer Network

We propose a Transformer architecture that leverages multi-head attention to process outdoor LiDAR frames. Illustrated in Fig. 2, our network consumes slices of a frame, which is inline with the pre-text task discussed in III-A. We describe the major processes of the network below.

1) *Slice Extraction*: A point cloud frame of size $\mathbb{R}^{N \times 3}$ is transformed to $\mathbb{R}^{S \times D \times 3}$ where N , S and D are the number of input points, number of slices and dimension of each slice respectively. For slice extraction, we first transform X, Y, Z values of in Cartesian coordinated to Azimuth, Elevation, and Radius. Along the Azimuth, points falling in 30° slices are extracted. Then, a 10° rotation is applied to the point cloud to extract a new slice. This allows slice extraction with a 20° overlap. In our implementation, we filter out the less dense regions beyond 70m radius in a slice.

2) *Input Embeddings Generation*: We transform 3D slices of size $\mathbb{R}^{S \times D \times 3}$ into input emdeddings in $\mathbb{R}^{S \times D}$. To that end, we employ three 2D CNN layers with the input channel




Fig. 2. Schematics of the proposed method. A 3D point cloud frame of N points is divided into 36 slices with overlap. Embedding for each slice is created for tokenization and positional encoding is added. Local Feature Block (LFB) is proposed to enable multi-headed attention (MHA) over multi-layer perceptron (MLP) embeddings of input features. This process forms the core of Slice Transformer. Processing the features with multiple Slice Transformer blocks, MLPs and skip connections, the output is processed with a softmax layer for the pre-text classification task. The + sign stands for addition.

sizes 36, 64 and 128. These layers are followed by four 1D CNN layers, with output channel sizes 128, 128, 64, and 36. These layers are packed into a block, which is called an Input MLP block in Fig. 2 due to the dominant use of 1D CNN layers for projection purpose.

3) *Positional Embedding*: The positional information to the input embeddings in transformers is known to boost the network performance. Unlike the conventional ‘cosine’ and ‘sine’ positional encoding, we add learnable slice-wise positional encodings for improved performance. This learned positional encoding is added to the input data.

4) *Local Features Block (LFB)*: Local features play a key role in neural representation. We exploit both multi-headed attention (MHA) and CNNs to extract powerful features from the input point cloud. Local Feature Block (LFB) is the central component of our network that extracts local features of a raw point cloud frame. It incorporates pairs of the proposed Slice Transformer and an MLP in series as shown in Fig. 2. Between the pairs, skip-connections are used. We empirically selected three pairs of slice transformers and MLP in our network.

Slice Transformer: This module processes a point cloud as a sequence of slices, hence called Slice Transformer. As shown in Fig. 2, it comprises four Multi-head Attentions (MHAs) and three MLPs. The input and output of this block are in $\mathbb{R}^{S \times D}$. An MHA is based on a Scaled Dot-Product Attention defined as
$$\text{Atten}(q, k, v) = v \text{softmax}\left(\frac{q \cdot k^T}{\sqrt{dk}}\right), \quad (1)$$

where Atten is the attention function of a single head; q, k, v are query, key and value; dk is key’s dimension.

In the case of MHA, keys, queries and values are projected multiple times to different learned linear transformations. The attention function is computed in parallel over all projections separately. The outputs of these computations are concatenated. For efficiency, we use 16 heads with 128 dimensions each. Thus, the input to a single head attention function is $\mathbb{R}^{S \times 128}$. MHA permits the network to attend useful features from different representation subspaces at various positions. In Slice Transformer, we place MLPs before MHA modules at various locations to extract useful and diverse local features of the input data. Each MLP consists of six 1D CNNs with Batch

Normalization layers and ReLU activation functions. The input and output channel size for the CNNs are (36, 64), (64, 64), (64, 128), (128, 128), (128, 128), (128, 36). Mathematically, a single head performs the computation as shown in Eq. (1) while MHA performs concatenation operation over all heads computations, which can be expressed as

$$\text{MHA}(q, k, v) = \text{concat}(h_1, h_2, h_3 \dots h_i)W^O, \quad (2)$$

where h_i denote attention heads and $W^O \in \mathbb{R}^{dk \times 2048}$ is the projection parameter matrix.

5) *Output Classification*: Within the context of pre-text learning, this process is responsible for the final prediction. It is implemented with 5 1D CNNs with kernel size (1, 1), a maxpooling layer, and two linear layers. The 1D CNNs map the input features $\mathbb{R}^{S \times D}$ to $\mathbb{R}^{1024 \times 2048}$ and then maxpooling is employed channel-wise. Finally, linear layers and softmax is applied to generate the final output. Using the conventional cross-entropy loss, the network is trained to predict pseudo-labels for the pre-text task or the true labels for classification-see Sec. VI-B.

IV. LOCALIZATION

Self localization is a critical autonomous navigation task. We present a method for localization in 3D point cloud map with LiDAR input by building on the self-supervised model presented in the preceding section. Our technique is illustrated in Fig. 3. It divides a query frame into slices following the method of Sec. III-A, and extracts features for slices using a pre-trained backbone obtained using the technique in Sec. III. The output of this module is further processed using the components described below.

A. Features Filtering

Features extracted by the backbone network may also contain unwanted and noisy features. Inspired by [39], [40], we devise a neural modeling based filter to cleanse the unwanted features. To that end, we first estimate a mask of size $\mathbb{R}^{1 \times D}$ for the input slice features in $\mathbb{R}^{S \times D}$ with an MLP employing sigmoid activation. The mask is then broadcasted to the features by applying dot product between the mask and input features followed by normalization. This enables extraction of similar features from the input features map.

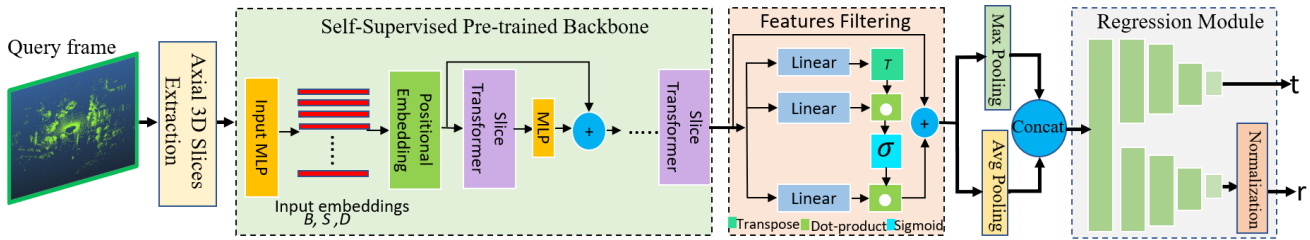


Fig. 3. Proposed end-to-end 6DoF localization method that leverages the self-supervised pre-trained backbone followed by a regression stage that employs two-headed fully-connected sub-network to predict translation and rotation parameters for the query point cloud frame.




Fig. 4. Constructed 3D point cloud map of Perth CBD, WA. The map shown binary annotations of roads and other objects.

B. Max-pooling and Average Pooling

The main responsibility of these layers is to identify slice-wise maximum and average features. The max-pooling and average pooling generate outputs with size $\mathbb{R}^{D \times 1}$, expecting inputs of size $\mathbb{R}^{D \times S}$. Outputs from these layers are concatenated for further processing.

C. Regression Module

The pre-trained model from Sec. III is a classifier whereas 6DoF prediction is better formulated as a regression problem. Therefore, we replace classification related modules with a Regression Module, while keeping the pre-trained backbone. The employed Regression Module consists of a common fully-connected (FC) layer at the initial stage, followed by two branches of four FC layers - see Fig. 3. The output channel sizes in each branch of the FC layers are 1024, 512, 256, 3. The layers are initialized with Xavier_uniform distribution, and use ReLU activations. We also normalize the rotations branch of the sub-network to account for the relatively smaller variations.

D. Training Loss Function

Unlike other localization methods which are predominantly restricted to a single loss for both rotation and translation, we use L_1 -loss for translation vector and Cosine Similarity Embedding loss for the rotation vector. We combine these losses with learnable balancing factors α and β , as the net batch loss in Eq. (3).

$$\mathcal{L} = \|t - t'\|e^\alpha + \alpha + (1 - \cos(r - r'))e^\beta + \beta, \quad (3)$$

where (t', r') , (t, r) are the ground-truth and predicted poses i.e., translations t (along X, Y, Z directions) and rotations r (in yaw, roll, pitch).

V. PERTH-WA DATASET

Another major contribution of this work is Perth-WA dataset that provides 6DoF annotations for localization. The

data comprises a LiDAR map of 4km² region of Perth Central Business District (CBD) in Western Australia as shown in Fig. 4. The scenes contain commercial structures, residential areas, food streets, complex routes, and hospital building etc. The data was collected in three different two-hour sessions under day/night conditions with sunny and cloudy weather. Unlike the existing related dataset, Apollo-SouthBay [10] and Oxford Radar RobotCar Dataset [41], Perth-WA dataset annotations do not rely on Inertial Measurement Unit (IMU). Instead, the labeling comes directly from the LiDAR frames themselves, as explained shortly.

In constructing a large 3D point cloud map, accumulation of registration errors is a major problem. This causes a drift in the map. To overcome that, we collected the data in closed loops. It is known that registration of closed loop LiDAR frames reduces error accumulation drastically [42]. Our map is constructed in an offline process, which involves extraction of loops from the scanned data, filtering out non-static objects, registration and merging the consecutive frames in a loop, followed by merging all the sub-maps into a single map. We further post-process the map to remove redundant points. In map construction, we leverage 3D normal-distribution transform (3D-NDT) algorithm [43] for registering the LiDAR frames. There are total nineteen sub-maps created from raw LiDAR frames which are merged to create a single 3D point cloud map. We allow two types of regions with dense and sparse point clouds in the map, the latter to create a more challenging scenario for localization. For the sparse region of the map, we skip three consecutive frames in the collected data.

To extract the ground-truth poses for Perth-WA dataset, we exploit the map creation process itself. Within a loop, a moving LiDAR frame is registered with a static point cloud, which generates a transformation matrix for the moving frame. To compute the transformation matrix of a frame, we multiply the previous frame's transformation to the transformation matrix of that frame. For instance, when the frames in a loop are $f_1, f_2, f_3, \dots, f_n$, and their corresponding transformation matrices are $T_1, T_2, T_3, \dots, T_n$, then the ground truth transformation matrices for each frame can be computed as $(T_0 \cdot T_1) = T_{g1}, (T_{g1} \cdot T_2) = T_{g2}, \dots, (T_{gn-1} \cdot T_n) = T_{gn}$. The T_0 is a 4×4 identity matrix for the first loop, however, it is the ground-truth transformation matrix of the last frame in a previous loop for any other loop. To compute the ground-truth transformation matrices for the next loop, we initialise the registration process from T_0 to compute the ground-truth of the first frame in that loop. After merging two loops, we

TABLE I
CLASSIFICATION RESULTS (ACC %) ON MODELNET40 AND
SCANOBJECTNN. PC = POINT CLOUD, N = NORMAL.

Method	Input	ModelNet40	ScanNN
PointNet++[44]	PC	90.7	77.9
DGCNN [45]	PC	89.0	78.4
SpiderCNN (4-layer) [46]	PC+N	92.4	73.7
PointNet[7]	PC	89.2	68.2
PointCNN[47]	PC+N	92.2	78.5
3DmFV [48]	PC	91.4	63.0
Transformer [27]	PC	91.4	77.2
Transformer+OcCo [49]	PC	92.1	80.4
NPCT [50]	PC	91.0	-
POS-BERT [28]	PC	92.1	83.2
Ours	PC	92.4	84.5

refined the values of T_0 with rotation and translation values to ensure that the frames align perfectly with the map. We use this strategy for all the loops to generate the ground-truth annotations with high precision. Finally, we convert the ground-truth transformation metrics to rotation angles (degrees) and translation vectors (meters) for all the frames to provide annotations for training and testing the model. The Perth-WA dataset contains 30K frames with 6DoFs poses.

VI. EXPERIMENTS

We evaluate the proposed method for localization and object classification tasks on benchmark datasets and compare with the state-of-the-art methods. Moreover, we also establish baseline results on the proposed Perth-WA dataset. For comparison, we choose existing localization and classification methods based on their popularity and availability of the author-provided code for a fair evaluation.

A. Setup

All experiments are performed on Ubuntu 18.04 operating system with PyTorch version 1.9 and using a single GeForce RTX 3090 GPU with 24GB memory. For the proposed method, each point cloud frame is divided into 36 slices, 30° each with 2048 points per slice and 20° overlap with the neighboring slices.

B. Classification Task

For classification, we directly fine-tune our self-supervised pre-trained model for the task. We present results on the popular ModelNet40 [11] and ScanObjectNN [12] datasets.

1) *Datasets*: ModelNet40 [11] contains 3D CAD models of 40 object classes. We use the standard training and test set in our experiments, comprising 9,843 and 2,468 samples, respectively. The ScanObjectNN dataset [12] comprises nearly 15K real-world object scans with occlusions and background. It is a challenging point cloud dataset for classification task. It has 15 categories with 2,902 distinct object instances. Each object in the dataset is defined by a list of 3D points (x, y, z values), normal vector and color at each point, and class label. We perform experiments on the perturbed (PB-T50-RS) variant of the dataset.

2) *Implementation Details*: In classification, we set the same training parameters for both datasets. We employ batch sizes of 16 and 4 for training and testing, respectively; and use Adam optimizer with learning rate 0.001 and weight

decay 0.00001. We use the standard Cross Entropy loss and train the model for 60 epochs each on both datasets.

3) Results on ModelNet40 and ScanObjectNN Datasets:

We compare our method with the state-of-the-art. Among the MLP and CNN based classification methods, we compare with the PointNet[7], PointNet++[44], DGCNN [45], SpiderCNN [46], PointCNN[47], and 3DmFV [48]. We also compare with transformer based methods, which include Transformer [27], Transformer+OcCo [49], POS-BERT [28] and NPCT [50]. Table I summarizes the results on ModelNet40 and ScanObjectNN datasets. Our approach outperforms all methods on ScanObjectNN dataset. On the ModelNet40 (synthetic dataset), our method shares the same highest accuracy with SpiderCNN, however, our method significantly outperforms SpiderCNN on ScanObjectNN (real dataset).

From the results, we can draw a few conclusions. Primarily, a pre-trained model on real 3D point cloud data provides useful prior knowledge to help in the tasks of classification and localization for self-driving vehicles. Secondly, most of the methods fail to perform well on real-world 3D point clouds even though they have higher accuracy on the synthetic datasets. Most of the methods, except Transformer+OcCo [49] and POS-BERT [28], require fully supervised learning, without leveraging pre-training, thereby requiring large amount of labelled data. Transformer+OcCo [49] is an unsupervised Transformer approach. However, it also utilizes ModelNet40 samples for pre-training. Synthetic data for pre-training is not an optimal option for outdoor downstream tasks, e.g., localization and object detection. Similarly, POS-BERT [28] is a self-supervised approach, however, it is pre-trained on indoor scenes. Unlike these methods, we performed pre-training on real-world LiDAR point clouds which proved beneficial for both indoor and outdoor point cloud processing tasks.

C. Localization Task

We present localization results on the proposed Perth-WA dataset and the Apollo-SouthBay dataset [10]. We did not consider KITTI [53] dataset for experiments due to large errors in ground-truth values.

1) *Implementation Details*: For fair benchmarking, we use the same implementation setup for localization in the Perth-WA dataset and the Apollo-SouthBay dataset [10]. We employ batch sizes 8 and 2 for training and testing, respectively. We use Adam optimizer with learning rate 0.001 and β values (0.9, 0.999). The Exponential scheduler with $\gamma = 0.9$ is used during training. The model is trained for 50 epochs on both datasets.

2) *Results on Perth-WA dataset*: Out of the 30K labeled frames, we choose 20K frames for training including sparse and dense parts of the map. For testing, an exclusive set of 10K frames is used which is not a part of the training set. These frames were originally skipped during the construction of the training map. The test set is challenging, in that a part of it consists of 2,200 frames that are taken from widespread regions of Perth CBD. We evaluate the performance of a recent point cloud based localization approach PointLoc [6]

TABLE II

RESULTS ON PERTH-WA DATASET. THE VALUES REPRESENT ABSOLUTE MEAN ERROR FOR ROTATION (IN DEGREES) AND ADDITIONALLY, THE MAXIMUM ERROR (2ND VALUE) FOR TRANSLATION (IN METERS). OUR METHOD HAS THE LEAST ERROR IN ALL CASES.

Method	Yaw	Roll	Pitch	Rot	X	Y	Z	Trans
PointLoc [6]	0.26°	1.96°	0.153°	0.75°, 1.51°	29.70, 83.48	37.49, 65.98	7.80, 19.03	25.00, 59.89
Ours (baseline)	0.32°	2.42°	0.27°	1.00°, 2.30°	14.20, 123.63	17.05, 45.15	8.50, 23.56	13.25, 74.10
Ours (pretrained)	0.17°	1.52°	0.096°	0.59°, 1.48°	6.26, 21.44	6.55, 18.95	2.86, 16.64	5.23, 21.36

TABLE III

RESULTS ON APOLLO-SOUTH BAY DATASET. THE VALUES REPRESENT RMSE FOR ROTATION (IN RAD) AND TRANSLATION (IN METERS). OUR METHOD ACHIEVES THE LOWEST AVERAGE ERRORS.

Route	Method	Yaw	Roll	Pitch	Rot	X	Y	Z	Trans
BaylandsToSeafood	Levinson et al.[51]	-	-	-	-	0.148	0.115	0.074	0.112
	Wan et al.[52]	0.054	-	-	-	0.036	0.026	0.019	0.027
	Ours	0.066	0.016	0.05	0.133	0.011	0.033	0.020	0.021
ColumbiaPark	Levinson et al.[51]	-	-	-	-	0.063	0.045	0.034	0.047
	Wan et al.[52]	0.081	-	-	-	0.046	0.034	0.024	0.035
	Ours	0.037	0.25	0.025	0.104	0.026	0.013	0.019	0.020
Hightway237	Levinson et al.[51]	-	-	-	-	0.161	0.138	0.061	0.120
	Wan et al.[52]	0.069	-	-	-	0.049	0.038	0.022	0.036
	Ours	0.048	0.23	0.118	0.132	0.013	0.039	0.014	0.022
MathildaAVE	Levinson et al.[51]	-	-	-	-	0.106	0.086	0.044	0.078
	Wan et al.[52]	0.060	-	-	-	0.040	0.030	0.020	0.030
	Ours	0.033	0.034	0.190	0.085	0.019	0.027	0.039	0.029
SanJoseDowntown	Levinson et al.[51]	-	-	-	-	0.103	0.075	0.055	0.077
	Wan et al.[52]	0.052	-	-	-	0.058	0.039	0.034	0.044
	Ours	0.061	0.088	0.147	0.099	0.054	0.044	0.029	0.043
SunnyvaleBigLoop	Levinson et al.[51]	-	-	-	-	0.132	0.097	0.070	0.099
	Wan et al.[52]	0.081	-	-	-	0.069	0.050	0.038	0.052
	Ours	0.084	0.053	0.241	0.126	0.022	0.043	0.069	0.045
Average	Levinson et al.[51]	-	-	-	-	0.119	0.093	0.046	0.089
	Wan et al.[52]	0.066	-	-	-	0.050	0.036	0.026	0.037
	Ours	0.055	0.117	0.128	0.113	0.024	0.033	0.031	0.030

for comparison. We also compute baseline results that uses our method without self-supervised pretraining. Note that, our method is a complete localization approach even without the self-supervised pre-trained backbone. Following [6], we use Mean Absolute Error and Max error values of poses for analysing the performance. Table II summarizes the results of our experiments. Our self-supervised localization approach consistently outperforms the baseline model and PointLoc for angular and translation mean error values. We can conclude from these results that our proposed self-supervised approach enables more effective point cloud feature learning, which makes it suitable for localization using complex outdoor LiDAR frames. These results also demonstrate that Perth-WA dataset is comparatively more challenging than the Appollo-SouthBay dataset (see Table III) for localization.

3) *Results on the Apollo-SouthBay Dataset:* The ApolloSouthBay [10] is a large scale localization dataset collected in San Francisco, USA. An IMU based system is utilized to collect the ground-truth poses for the LiDAR frames. The dataset covers six routes, BaylandsToSeafood, ColumbiaPark, Highway237, MathildaAVE, SanJoseDowntown, and SunnyvaleBigLoop. For each route, TrainData, TestData and MapData are provided. Results of our experiments on this dataset are summarized in Table III. We follow [10] and use RMSE as the evaluating metric. To establish benchmark

results, we train our model on training sets, and test our model on all the routes as shown in Table III. We compare our approach with the state-of-the-art localization methods, Levinson et al. [51] and Wan et al. [52]. Our method outperforms both by achieving the lowest average errors. This is mainly due to the ability of our method to exploit transformers with pre-training on large real-world data.

VII. CONCLUSIONS

This paper made three major contributions to outdoor localization using point cloud maps. First, it introduced a pre-text task that allows self-supervised learning, keeping in view state-of-the-art Transformer architectures. Second, it proposed a first-of-its-kind Transformer network that enables the use of multi-head attention to process outdoor LiDAR data. Third, it provided a large-scale point cloud map of Perth (Western Australia), covering nearly 4km² area. The dataset provides annotations for 6DoF localization problem and is more challenging compared to the existing Apollo-SouthBay dataset. We established the baseline on the proposed dataset and benchmarked our method on an existing Apollo-SouthBay dataset, showing highly competitive localization results. We also established the effectiveness of our pre-text task and the model by directly fine tuning the later for classification on ModelNet40 and ScanNN datasets.

REFERENCES

- [1] A. Segal, D. Haehnel, and S. Thrun, "Generalized-icp." in *Robotics: science and systems*, vol. 2, no. 4. Seattle, WA, 2009, p. 435.
- [2] T. Stoyanov, M. Magnusson, H. Andreasson, and A. J. Lilienthal, "Fast and accurate scan registration through minimization of the distance between compact 3d ndt representations," *The International Journal of Robotics Research*, vol. 31, no. 12, pp. 1377–1393, 2012.
- [3] Z. J. Yew and G. H. Lee, "3dfeat-net: Weakly supervised local 3d features for point cloud registration," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 607–623.
- [4] L. Li, M. Yang, L. Weng, and C. Wang, "Robust localization for intelligent vehicles based on pole-like features using the point cloud," *IEEE Transactions on Automation Science and Engineering*, vol. 19, no. 2, pp. 1095–1108, 2021.
- [5] G. Elbaz, T. Avraham, and A. Fischer, "3d point cloud registration for localization using a deep neural network auto-encoder," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [6] W. Wang, B. Wang, P. Zhao, C. Chen, R. Clark, B. Yang, A. Markham, and N. Trigoni, "Pointloc: Deep pose regressor for lidar point cloud localization," *IEEE Sensors Journal*, vol. 22, no. 1, pp. 959–968, 2021.
- [7] 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.
- [8] A. Komarichev, Z. Zhong, and J. Hua, "A-cnn: Annularly convolutional neural networks on point clouds," in *CVPR*, 19, pp. 7421–7430.
- [9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *arXiv preprint arXiv:1706.03762*, 2017.
- [10] W. Lu, Y. Zhou, G. Wan, S. Hou, and S. Song, "L3-net: Towards learning based lidar localization for autonomous driving," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6389–6398.
- [11] 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.
- [12] M. A. Uy, Q.-H. Pham, B.-S. Hua, T. Nguyen, and S.-K. Yeung, "Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [13] M. Elhousni and X. Huang, "A survey on 3d lidar localization for autonomous vehicles," in *2020 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2020, pp. 1879–1884.
- [14] X. Gao, Q. Wang, H. Gu, F. Zhang, G. Peng, Y. Si, and X. Li, "Fully automatic large-scale point cloud mapping for low-speed self-driving vehicles in unstructured environments," in *2021 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2021, pp. 881–888.
- [15] D. Kovalenko, M. Korobkin, and A. Minin, "Sensor aware lidar odometry," in *2019 European Conference on Mobile Robots (ECMR)*. IEEE, 2019, pp. 1–6.
- [16] J. Nubert, S. Khattak, and M. Hutter, "Self-supervised learning of lidar odometry for robotic applications," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 9601–9607.
- [17] R. Dubé, A. Cramariuc, D. Dugas, J. Nieto, R. Siegwart, and C. Cadena, "Segmap: 3d segment mapping using data-driven descriptors," *arXiv preprint arXiv:1804.09557*, 2018.
- [18] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," *arXiv preprint arXiv:2103.14030*, 2021.
- [19] C.-F. Chen, Q. Fan, and R. Panda, "Crossvit: Cross-attention multi-scale vision transformer for image classification," *arXiv preprint arXiv:2103.14899*, 2021.
- [20] L. Yuan, Y. Chen, T. Wang, W. Yu, Y. Shi, Z. Jiang, F. E. Tay, J. Feng, and S. Yan, "Tokens-to-token vit: Training vision transformers from scratch on imagenet," *arXiv preprint arXiv:2101.11986*, 2021.
- [21] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, "Training data-efficient image transformers & distillation through attention," in *International Conference on Machine Learning*. PMLR, 2021, pp. 10 347–10 357.
- [22] R. Strudel, R. Garcia, I. Laptev, and C. Schmid, "Segmenter: Transformer for semantic segmentation," *arXiv preprint arXiv:2105.05633*, 2021.
- [23] H. Chen, Y. Wang, T. Guo, C. Xu, Y. Deng, Z. Liu, S. Ma, C. Xu, C. Xu, and W. Gao, "Pre-trained image processing transformer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 12 299–12 310.
- [24] T. Guan, J. Wang, S. Lan, R. Chandra, Z. Wu, L. Davis, and D. Manocha, "M3det: Multi-representation, multi-scale, mutual-relation 3d object detection with transformers," *arXiv preprint arXiv:2104.11896*, 2021.
- [25] I. Misra, R. Girdhar, and A. Joulin, "An end-to-end transformer model for 3d object detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 2906–2917.
- [26] H. Zhao, L. Jiang, J. Jia, P. H. Torr, and V. Koltun, "Point transformer," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 16 259–16 268.
- [27] X. Yu, L. Tang, Y. Rao, T. Huang, J. Zhou, and J. Lu, "Point-bert: Pre-training 3d point cloud transformers with masked point modeling," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 19 313–19 322.
- [28] K. Fu, P. Gao, S. Liu, R. Zhang, Y. Qiao, and M. Wang, "Pos-bert: Point cloud one-stage bert pre-training," *arXiv preprint arXiv:2204.00989*, 2022.
- [29] Z. Wang, X. Yu, Y. Rao, J. Zhou, and J. Lu, "P2p: Tuning pre-trained image models for point cloud analysis with point-to-pixel prompting," *arXiv preprint arXiv:2208.02812*, 2022.
- [30] Y. Chen, J. Liu, B. Ni, H. Wang, J. Yang, N. Liu, T. Li, and Q. Tian, "Shape self-correction for unsupervised point cloud understanding," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 8382–8391.
- [31] I. Achituve, H. Maron, and G. Chechik, "Self-supervised learning for domain adaptation on point clouds," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2021, pp. 123–133.
- [32] O. Poursaeed, T. Jiang, H. Qiao, N. Xu, and V. G. Kim, "Self-supervised learning of point clouds via orientation estimation," in *2020 International Conference on 3D Vision (3DV)*, 2020, pp. 1018–1028.
- [33] Y. Cho, G. Kim, and A. Kim, "Unsupervised geometry-aware deep lidar odometry," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2145–2152.
- [34] Y. Xu, Z. Huang, K.-Y. Lin, X. Zhu, J. Shi, H. Bao, G. Zhang, and H. Li, "Selfvoxelo: Self-supervised lidar odometry with voxel-based deep neural networks," *arXiv preprint arXiv:2010.09343*, 2020.
- [35] Y. Tu, "Unpwc-svdlo: Multi-svd on pointpwc for unsupervised lidar odometry," *arXiv preprint arXiv:2205.08150*, 2022.
- [36] W. Wu, Z. Wang, Z. Li, W. Liu, and L. Fuxin, "Pointpwc-net: A coarse-to-fine network for supervised and self-supervised scene flow estimation on 3d point clouds," *arXiv preprint arXiv:1911.12408*, 2019.
- [37] Z. Zhou, C. Zhao, D. Adolphsson, S. Su, Y. Gao, T. Duckett, and L. Sun, "Ndt-transformer: Large-scale 3d point cloud localisation using the normal distribution transform representation," *arXiv preprint arXiv:2103.12292*, 2021.
- [38] J. Shan, S. Zhou, Z. Fang, and Y. Cui, "Ptt: Point-track-transformer module for 3d single object tracking in point clouds," *arXiv preprint arXiv:2108.06455*, 2021.
- [39] B. Wang, C. Chen, C. X. Lu, P. Zhao, N. Trigoni, and A. Markham, "Atloc: Attention guided camera localization," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 06, 2020, pp. 10 393–10 401.
- [40] Z. Huang, Y. Xu, J. Shi, X. Zhou, H. Bao, and G. Zhang, "Prior guided dropout for robust visual localization in dynamic environments," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 2791–2800.
- [41] D. Barnes, M. Gadd, P. Murcutt, P. Newman, and I. Posner, "The oxford radar robotcar dataset: A radar extension to the oxford robotcar dataset," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Paris, 2020. [Online]. Available: <https://arxiv.org/abs/1909.01300>
- [42] C.-K. Chng, A. Parra, T.-J. Chin, and Y. Latif, "Monocular rotational odometry with incremental rotation averaging and loop closure," in *2020 Digital Image Computing: Techniques and Applications (DICTA)*. IEEE, 2020, pp. 1–8.

- [43] M. Magnusson, "The three-dimensional normal-distributions transform: an efficient representation for registration, surface analysis, and loop detection," Ph.D. dissertation, Örebro universitet, 2009.
- [44] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," in *Advances in neural information processing systems*, 17, pp. 5099–5108.
- [45] S. S. Mohammadi, Y. Wang, and A. Del Bue, "Pointview-gcn: 3d shape classification with multi-view point clouds," in *2021 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2021, pp. 3103–3107.
- [46] Y. Xu, T. Fan, M. Xu, L. Zeng, and Y. Qiao, "Spidernn: Deep learning on point sets with parameterized convolutional filters," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 87–102.
- [47] Y. Li, R. Bu, M. Sun, W. Wu, X. Di, and B. Chen, "Pointcnn: Convolution on x-transformed points," *Advances in neural information processing systems*, vol. 31, 2018.
- [48] Y. Ben-Shabat, M. Lindenbaum, and A. Fischer, "3dmfv: Three-dimensional point cloud classification in real-time using convolutional neural networks," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3145–3152, 2018.
- [49] H. Wang, Q. Liu, X. Yue, J. Lasenby, and M. J. Kusner, "Unsupervised point cloud pre-training via occlusion completion," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 9782–9792.
- [50] M.-H. Guo, J.-X. Cai, Z.-N. Liu, T.-J. Mu, R. R. Martin, and S.-M. Hu, "Pct: Point cloud transformer," *Computational Visual Media*, vol. 7, no. 2, pp. 187–199, 2021.
- [51] J. Levinson and S. Thrun, "Robust vehicle localization in urban environments using probabilistic maps," in *2010 IEEE international conference on robotics and automation*. IEEE, 2010, pp. 4372–4378.
- [52] G. Wan, X. Yang, R. Cai, H. Li, Y. Zhou, H. Wang, and S. Song, "Robust and precise vehicle localization based on multi-sensor fusion in diverse city scenes," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 4670–4677.
- [53] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.