

Pose-Graph Attentional Graph Neural Network for Lidar Place Recognition

Milad Ramezani^{1,*†}, Liang Wang^{1,2†}, Joshua Knights^{1,3}, Zhibin Li¹, Pauline Pounds², Peyman Moghadam^{1,3}

Abstract—This paper proposes a pose-graph attentional graph neural network, called P-GAT, which compares (key) nodes between sequential and non-sequential sub-graphs for place recognition tasks as opposed to a common frame-to-frame retrieval problem formulation currently implemented in SOTA place recognition methods. P-GAT uses the maximum spatial and temporal information between neighbour cloud descriptors — generated by an existing encoder — utilising the concept of pose-graph SLAM. Leveraging intra- and inter-attention and graph neural network, P-GAT relates point clouds captured in nearby locations in Euclidean space and their embeddings in feature space. Experimental results on the large-scale publically available datasets demonstrate the effectiveness of our approach in scenes lacking distinct features and when training and testing environments have different distributions (domain adaptation). Further, an exhaustive comparison with the state-of-the-art shows improvements in performance gains. Code is available at <https://github.com/csiro-robotics/P-GAT>.

Index Terms—place recognition, spatiotemporal attention, SLAM

I. INTRODUCTION

ACCURATE and drift-free (re)-localisation is critical for many robotic and computer vision applications, such as autonomous navigation [1] and augmented reality [2]. Achieving reliable (re)-localisation is challenging, particularly in GPS-denied environments, such as indoor, subterranean or dense vegetated environments [3], [4], due to occlusion, complex geometry, and dynamic objects in scenes.

One promising direction for addressing the challenges associated with reliable (re)-localisation is to utilise a Place Recognition (PR) method to predict the coarse location of an agent within a database of previously visited places. Place Recognition is commonly framed as a retrieval task in computer vision and robotics, either vision-based [5]–[8] or lidar-based [9]–[13]. Given a query, the method involves retrieving the most similar key in the database by first encoding the input frame (image/point cloud) as a global descriptor and matching it against the global descriptors of previously visited places.

Despite remarkable improvements, visual PR is less robust against appearance, season, illumination and viewpoint variations in large-scale (*i.e.*, city-scale) areas. In this paper, we

Manuscript received: August, 12, 2023; Revised October, 23, 2023; Accepted November, 21, 2023.

This paper was recommended for publication by Editor Cesar Cadena Lerma upon evaluation of the Associate Editor and Reviewers' comments.

This work was funded by CSIRO's Machine Learning and Artificial Intelligence Future Science Platform (MLAI FSP). (*P.P. and P.M. share senior authorship.*) (*Corresponding author.) (†Equally contributed authors.)

¹ Robotics and Autonomous Systems, DATA61, CSIRO, Australia. E-mails: firstname.lastname@csiro.au

² The University of Queensland, Brisbane, Australia. E-mails: firstname.lastname@uq.edu.au

³ Queensland University of Technology (QUT), Brisbane, Australia. E-mails: firstname.lastname@qut.edu.au

Digital Object Identifier (DOI): see top of this page.




Fig. 1: P-GAT aims to optimise spatiotemporal information by relating point clouds within and across subgraphs leveraging an attentional graph neural network. If point clouds are captured in nearby locations (similar point clouds), our intra- (edges in black shades) and inter- (edges in blue shades) attention mechanism reweights their embeddings to bring them closer in feature space. Intra-attention enhances distant point cloud communication within subgraphs, while inter-attention facilitates potential place recognition in revisit areas.

consider the problem of lidar place recognition for large-scale environments. Despite all the progress in the field of lidar place recognition, most existing methods only encode a single lidar frame into a global descriptor, and hence, topological scene-level understanding is often neglected. There are few prior works [14], [15] that directly aggregate a sequence of lidar descriptors to generate one single global descriptor for each lidar sequence. However, these methods do not take advantage of the topological relationship between a sequence of point clouds in the context of a graph containing sets of nodes and edges.

To exploit the spatiotemporal information between neighbouring point clouds, we propose an attentional graph neural network called P-GAT, that uses topological information obtained by a pose-graph lidar SLAM system, between a set of point clouds to maximise the receptive field for training. Nodes and edges generated in a pose graph (as the robot explores the environment and pose-graph SLAM optimises robot poses) are further used to generate fully connected graphs (subgraphs) that contain positional information for a given robot travel distance. These subgraphs are further fed into our P-GAT model for place recognition by comparing a pair of subgraphs rather than just a pair of point clouds (Fig. 1). The communication between the nodes of the subgraph pairs is performed leveraging an attentional graph neural network. To address the dynamic properties of subgraphs (varying number of nodes) and the presence of nodes in multiple subgraphs (in contrast to comparing between pairs of point clouds/images), we develop a customised layer normalisation module using

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

a boolean mask for padding nodes and develop an averaging scheme for efficiently computing similarity scores during inference. We demonstrate that P-GAT can be integrated into any global point cloud descriptors (*i.e.*, cloud encoder agnostic), which greatly improves their robustness and generalisability. We extensively analyse and compare the performance of our network with the state-of-the-art over multiple large-scale public datasets. To characterise the properties of P-GAT in detail, we demonstrate the role of each component using numerous ablation studies. The proposed P-GAT can achieve the state-of-the-art on various benchmark datasets for lidar place recognition tasks.

II. RELATED WORK

This section reviews hand-crafted and learning-based lidar PR methods before introducing the attentional graph neural network and its applications.

A. Handcrafted Lidar Descriptors

The purpose of lidar PR methods is to create distinctive descriptors, based on a Local Reference Frame (LRF), along the robot's path to recognise revisited places regardless of its pose. Two types of descriptors are used: signatures and histograms. Histogram-based methods such as PFH [16] and FPFH [17], describe the 3D surface neighbourhood of a point by encoding a few geometric features obtained individually at each point according to local coordinates. DELIGHT [18], instead of geometric properties, computes intensity histograms. Signature-based algorithms like Scan Context [19] or *Segmatch* [20] use descriptor-based features to improve place recognition. Additionally, algorithms such as SHOT [21] combine both histograms and signatures for robust place recognition. However, handcrafted lidar descriptors require careful tuning of feature extraction and matching parameters depending on the operating environment.

B. Learning-based Lidar Descriptors

Recent advances in learning-based approaches have shown promising results in addressing the challenges mentioned earlier. CNN-based methods such as *Segmap* [22] or Efficient Segment Matching (ESM) [23] encode local patches of a point cloud into local embeddings. Local descriptors can later be used for localisation [24], [25]. These methods, however, are not defined in an end-to-end fashion to form a scene-level global descriptor for point clouds.

Using a convolutional bottom-up and top-down backbone, MinkLoc3D [10] and its variations [26], [27] extract local features and aggregate them into a global descriptor by a Generalised-Mean pooling (GeM) [28]. *LoGG3D-Net* [29] employs a sparse convolutional U-Net to encode point clouds into local features. During training, it ensures the maximum similarity of corresponding local features on a pair of neighbour point clouds by defining a local consistency loss. Unlike MinkLoc3D, LoGG3D-Net uses second-order pooling to aggregate local features to create global descriptors.

In contrast to convolutional models, methods have been proposed which are based on PointNet [30], an encoder

which works directly on an unordered point cloud due to permutation invariance to points utilising a symmetry function. PointNetVLAD [9] is a seminal lidar PR work with a PointNet-based backbone design. It uses NetVLAD [31] to aggregate local descriptors for the generation of global descriptors. To capture local contextual information, PCAN [11] adds an attention map mechanism for predicting the significance of point features. Re-weighted local features are further aggregated into a discriminative global descriptor using NetVLAD. Similarly, LPD-Net [32] aims to cover the limitation of PointNet in the extraction of local contextual information by aggregating neighbour features using a graph neural network. Global descriptors are finally generated using NetVLAD. In another effort, SOE-Net [12] adds an orientation-encoding unit in local descriptor extraction and a self-attention unit before the aggregation of local features through NetVLAD to improve the point-wise feature representation of PointNet. Recently, PPT-Net [13], inspired by SOE-Net [12] and pyramid structure of PointNet++ [33], proposed a pyramid point transformer design to learn the regional contextual information of a point cloud at multiple levels leveraging a grouped self-attention mechanism for the extraction of discriminative local embeddings. Local embeddings are also aggregated using NetVLAD.

In all the lidar PR methods mentioned above, point clouds are compared pairwise and thus, the topological and sequential relationship between a set of point clouds are not explored. The work that attempts to leverage temporal information between point clouds is Locus [15] which aims to relate the local features of the current point cloud to their correspondences in the previous point clouds using second-order temporal feature pooling; however, within a short time window consisting of only three frames. This limits the generalisation of Locus [15] under test-time distribution shifts. SeqOT [14] was also proposed to benefit from temporal information, although using a sequence of range images. Moreover, it generates only a single descriptor for each sequence.

C. Attentional Graph Neural Network

Attentional Graph Neural Network (AGNN) is designed to process data represented as a graph, using attention mechanisms [34] to selectively focus on nodes and edges. The attention mechanism helps reduce the graph's complexity and improves the effectiveness of local feature matching across nodes. AGNNs have been successfully applied to various tasks, such as image matching [35], object detection [36], and program re-identification [37]. To the best of our knowledge, AGNN has never been used in lidar PR tasks. We believe AGNN is well-suited for lidar PR (due to its efficiency in the aggregation of contextual information) when exploring spatiotemporal information existing between nearby point clouds.

III. THE PROPOSED METHOD

Our goal is to enhance PR performance by increasing the spatiotemporal information on scenarios where a graph with nodes and edges represents robot poses and spatial constraints in between. Robot poses serve as positional information.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

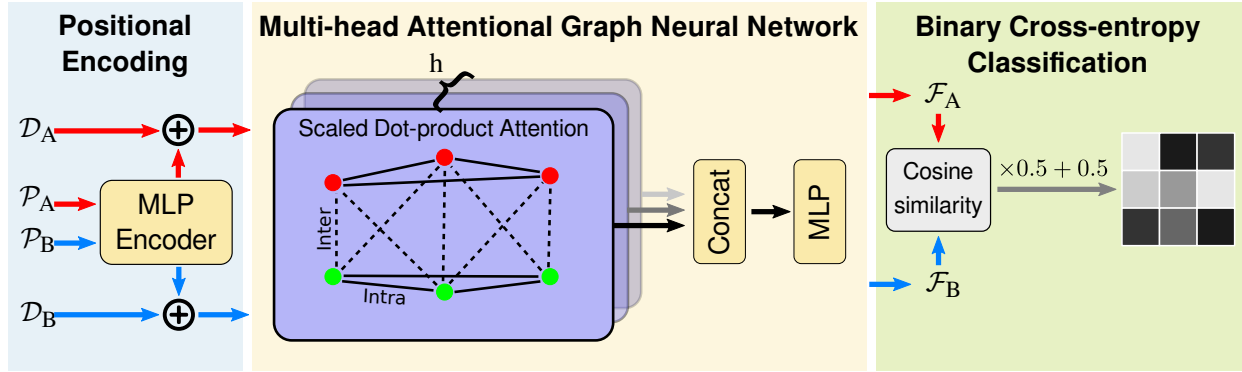


Fig. 2: P-GAT employs a multi-head attentional graph neural network for contextual and viewpoint information aggregation of sequential point clouds, refining each point cloud’s descriptor. The network uses an MLP encoder that supplements descriptors with normalised positions. Intra- and inter-attention mechanisms infer relationships for all nodes simultaneously, updating the descriptors accordingly. P-GAT predicts whether the pair of two point clouds are captured from the same place by comparing their final descriptors with shifted cosine similarity.

A. Overview

P-GAT is defined based on a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} and \mathcal{E} denote nodes and edges, respectively. The graph is inherently created by a SLAM system (pose-graph). For example, a lidar SLAM system registers consecutive lidar scans to create nodes and edges of the pose graph for further use in a back-end optimisation problem. Utilising pose-graph SLAM, we assume that accurate localisation is achievable within a local window of robot traversal, *e.g.*, 200 m, for a vehicle travelling in a large-scale environment. However, no global localisation is provided to P-GAT.

Our network comprises three major blocks: positional encoding, graph-based attention mechanism and binary classification. Provided local robot poses by SLAM, we encode relative positional information into embeddings computed by a 3D backbone to maintain the topological relationship between a sequence of keynodes (selected robot poses associated with a point cloud and its corresponding descriptor). P-GAT utilises both graph structure and attention mechanism, described in Sec. III-C, to increase the distinctiveness of positional-aware embeddings. Inspired by SuperGlue [35], which uses an AGNN for image matching, we aim to relate the descriptors of point clouds captured in nearby locations by aggregating contextual and viewpoint information into point clouds’ descriptors using an AGNN. Figure 2 depicts the overall architecture of P-GAT and its components.

B. Problem Formulation

To formulate the problem, we consider two sets of point clouds as two subgraphs $\mathcal{S}_A = \{\mathcal{V}_i^A\}_{i=1}^N$, and $\mathcal{S}_B = \{\mathcal{V}_j^B\}_{j=1}^M$, where M and N denote the number of keynodes in \mathcal{S}_A and \mathcal{S}_B , respectively. Our PR problem is now defined between pairs of subgraphs instead of two individual point clouds, as is common in the literature. We also formulate the PR problem as a binary classification problem using a similarity measure computed from descriptors between every pair of $(\mathcal{V}_i^A, \mathcal{V}_j^B)$ in subgraphs \mathcal{S}_A and \mathcal{S}_B .

All the pairs of keynodes from subgraphs \mathcal{S}_A and \mathcal{S}_B that are captured in the same place must be classified. As noted earlier, each keynode is associated with the robot position \mathbf{t}

and a descriptor \mathbf{d} , *i.e.*, $\mathcal{V} = \{\mathbf{t}, \mathbf{d}\}$. The position $\mathbf{t} \in \mathbb{R}^3$ is relative to the first node of the graph and normalised, *i.e.*, $\mathbf{p} = \frac{1}{\sigma}(\mathbf{t} - \mathbf{c})$, where \mathbf{c} is the centroid of the nodes’ positions in a subgraph and σ is a scalar showing the scatteredness of the keynodes. This normalisation is essential to maintain the geometry consistency between keynodes and network generalisation when training and testing data have different distributions. The descriptor $\mathbf{d} \in \mathbb{R}^E$ represents the input point cloud embedding as a fixed-size vector global descriptor at the time of keynode \mathcal{V} , and E denotes the descriptor dimension.

C. Attentional Graph Neural Network

To exploit the spatiotemporal information existing between keynodes in the pose graph created by SLAM, we build a fully-connected graph between the subgraph pairs $(\mathcal{S}_A, \mathcal{S}_B)$, *i.e.*, $\mathcal{G}_{AB} = (\mathcal{V}_{AB}, \mathcal{E}_{AB})$ with $N + M$ keynodes and $\frac{1}{2}((N + M)^2 - (N + M))$ undirected edges. Now, we can aggregate the relative positional and contextual constraints between the keynodes in \mathcal{G}_{AB} utilising an AGNN to enhance the descriptors’ representation. Since \mathcal{G}_{AB} has two types of connections (multiplexity [38]), *i.e.*, intra- and inter- subgraph edges, we effectively train the network to push the embeddings of the point clouds captured in nearby locations together. This allows embeddings to be invariant to dynamic and viewpoint changes resulting in more effective place recognition.

Positional Encoding: To embed the keynode relative position into the descriptor with a higher dimension E in each subgraph, we use Multi-Layer Perceptron (MLP) encoder to increase the dimension of the position \mathbf{p} to E and integrate it with the descriptor \mathbf{d} by element-wise addition \oplus :

$$\mathcal{X} = \mathcal{D} \oplus \text{MLP}_{\text{enc}}(\mathcal{P}), \quad (1)$$

where $\mathcal{X} \in \mathbb{R}^{E \times N}$ is the matrix of positional-aware descriptors and $\mathcal{D} \in \mathbb{R}^{E \times N}$ is the matrix consisting of original descriptors \mathbf{d}_i , $i \in \{1, \dots, N\}$. Position embedding is common in sequence-to-sequence learning problems [34], [39], enhancing the model’s ability to capture contextual cues. In our case, it improves the embeddings’ distinctiveness by allowing the GNN to distinguish between different embeddings of point clouds at different positions in the subgraph, thus capturing

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

the contextual and viewpoint information more effectively (See Sec. IV-E).

Multi-head Attentional GNN: As stated before, our graph between pairs of subgraphs is fully connected and multi-plex, *i.e.*, it comprises two types of edges: intra edges $\mathcal{E}_{\text{intra}}$ which are the edges between the nodes within one subgraph and inter edges $\mathcal{E}_{\text{inter}}$ which are the edges across the subgraphs in the input pair. Borrowing the terminology of *message passing* in GNN from [40], we aggregate the messages carried through edges $e_{ij} \in \mathcal{E}$ to keynode i , *i.e.*, $\mathbf{m}_{e_i} := \{\mathbf{m}_{e_{ij}} \mid \forall j : e_{ij} \in \mathcal{E}\}$. The AGNN is composed of multiple layers. At each layer $\ell \in \{1, \dots, L\}$, the message passing update, *i.e.*, updating the intermediate feature ${}^{(\ell)}\mathbf{x}_i$ of keynode \mathcal{V}_i , is concurrently conducted by aggregating messages across all the edges for all the keynodes once for subgraph \mathcal{S}_A and once for \mathcal{S}_B using:

$${}^{(\ell+1)}\mathbf{x}_i = {}^{(\ell)}\mathbf{x}_i + \mathcal{M}_t({}^{(\ell)}\mathbf{x}_i, \mathbf{m}_{e_i}), \quad (2)$$

where \mathcal{M}_t is the message function. Following [35], we use a MLP, albeit modified (See Sec. IV-C), and concatenate messages with intermediate representations, *i.e.*, $\mathcal{M}_t(\mathbf{x}, \mathbf{m}_e) = \text{MLP}([\mathbf{x} \parallel \mathbf{m}_e])$, where $[\cdot \parallel \cdot]$ denotes concatenation operator.

Following the attention mechanism described in [34], we aggregate messages received by each keynode through edges within (intra-attention) and across (inter-attention) subgraphs \mathcal{S}_A and \mathcal{S}_B . To this end, we consider the receiver keynode to be in subgraph \mathcal{S}_R and the entire sender keynodes to be in \mathcal{S}_S such that $(\mathcal{S}_R, \mathcal{S}_S) \in \{\mathcal{S}_A, \mathcal{S}_B\}^2$.

$$\begin{aligned} \mathbf{Q} &= \mathbf{W}_Q \mathcal{X}_R + \mathbf{b}_Q, \\ \begin{bmatrix} \mathbf{K} \\ \mathbf{V} \end{bmatrix} &= \begin{bmatrix} \mathbf{W}_K \\ \mathbf{W}_V \end{bmatrix} \mathcal{X}_S + \begin{bmatrix} \mathbf{b}_K \\ \mathbf{b}_V \end{bmatrix}, \end{aligned} \quad (3)$$

where matrices \mathbf{Q} , \mathbf{K} and \mathbf{V} consist of query vectors (features of keynodes that are being attended to) in \mathcal{S}_R , \mathbf{K} keys (features that are used to compute the attention scores) and \mathbf{V} values (features that are weighted by the attention scores to compute the output) in \mathcal{S}_S , respectively. \mathcal{X}_R and \mathcal{X}_S are the matrices packing the intermediate features \mathbf{x}_i from \mathcal{S}_R and \mathbf{x}_j from \mathcal{S}_S , respectively. Matrices \mathbf{W}_Q , \mathbf{W}_K and \mathbf{W}_V are linear transformation weights computed within training along with biases \mathbf{b}_Q , \mathbf{b}_K and \mathbf{b}_V . The entire messages can now be computed by scaled-dot product attention as follows:

$$\mathbf{m}_e = \text{softmax} \left(\frac{\mathbf{Q}^\top \mathbf{K}}{\sqrt{d_k}} \right) \mathbf{V}, \quad (4)$$

where d_k is the keys dimension. For large d_k values, the scaling is essential to avoid large magnitudes likely to cause minimal gradients when using a softmax function [34]. We use multi-head attention to capture different types of information and relationships within the input more effectively.

D. Classification Layer

The output of the AGNN block is two tensors representing the final descriptors of subgraph \mathcal{S}_A and \mathcal{S}_B , $\mathcal{F}^A = \{\mathbf{f}_i^A\}_{i=1}^N \in \mathbb{R}^{E \times N}$ and $\mathcal{F}^B = \{\mathbf{f}_j^B\}_{j=1}^M \in \mathbb{R}^{E \times M}$. The final prediction in P-GAT is performed using cosine similarity to produce a similarity matrix $\mathbf{S} = \{s_{ij}\} \in \mathbb{R}^{N \times M}$. Element s_{ij} is obtained from:

$$s_{ij} = \frac{\langle \mathbf{f}_i^A, \mathbf{f}_j^B \rangle}{\|\mathbf{f}_i^A\| \|\mathbf{f}_j^B\|}, \quad (5)$$

where $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ denote the inner product and the L2 norm, respectively.

Based on the definition of PR, the descriptors of two nearby point clouds should be similar, *i.e.*, s_{ij} close to 1, while the descriptors of two dissimilar point clouds should be distinct, *i.e.*, s_{ij} close to -1 . Since there can be multiple pairs of nearby point clouds on a given pair of two subgraphs, PR based on two subgraphs becomes a multi-label binary classification problem. Additionally, we map the scores in \mathbf{S} with the range $[-1, 1]$ to $\mathbf{P} = \{p_{ij}\}$ with the range $[0, 1]$ using:

$$p_{ij} = s_{ij} \times 0.5 + 0.5, \quad (6)$$

and interpret it as the probability that point cloud \mathcal{P}_i from subgraph \mathcal{S}_A and point cloud \mathcal{P}_j from subgraph \mathcal{S}_B represent the same place. Our classification problem can now be defined as a stochastic optimisation problem measuring the similarity between two probabilities.

E. Loss

Reformulating the binary classification problem as stochastic optimisation, we minimise the Kullback-Leibler (KL) divergence $D_{KL}(\mathbf{y} \parallel \mathbf{P}) = -\sum_{i=1}^N p_i \log(\frac{p_i}{y_i})$ between the predicted probabilities p_i and the true probabilities y_i (ground truth). The true probabilities follow the Bernoulli distribution, *i.e.*, the probability that random variable $x \in \{0, 1\}$ belongs to a class ($x = 1$) is $p(x = 1) = p$, otherwise $p(x = 0) = 1 - p$. Using the Bernoulli distribution properties, the KL divergence is converted into Binary Cross Entropy (BCE) [41]. Since we have multiple separate classifications to perform (between the keynodes in pairs of subgraphs), our final BCE loss is defined as follows:

$$\mathcal{L}(\mathbf{y}, \mathbf{P}) = -\sum_{i=1}^N \sum_{j=1}^M \omega_{ij} (y_{ij} \cdot \log p_{ij} + (1 - y_{ij}) \cdot \log(1 - p_{ij})), \quad (7)$$

where $y_{ij} \in \{0, 1\}$ is the ground truth label. If point cloud \mathcal{P}_i and point cloud \mathcal{P}_j represent the same place $y_{ij} = 1$, otherwise, $y_{ij} = 0$. ω_{ij} is a scalar hyperparameter to indicate whether the point clouds pair $(\mathcal{P}_i, \mathcal{P}_j)$ contributes in the loss function based on the conditions we follow to select positive and negative pairs (See Sec. IV-C).

IV. EXPERIMENTS

We briefly introduce the datasets used, followed by evaluation settings and the implementation details. Comparisons are made between our proposed network and three state-of-the-art architectures with different backbones, serving as baselines. Additionally, we assess our network's performance in comparison to existing lidar PR approaches. We then provide detailed ablation studies to verify the network design and the impact of each component on performance gains.

A. Datasets

We use three publicly-available large-scale datasets for evaluation. We detail the characteristics of each dataset below. **Oxford RobotCar dataset** [42] has been widely used for lidar PR, which is a processed subset of the overall Oxford RobotCar dataset [42]. It consists of point clouds captured

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

by travelling a route (~ 10 km) 44 times across Oxford, UK, over a year. To assess performance on the Oxford RobotCar dataset, point clouds from one trip are taken as queries and matched against point clouds from other trips in an iterative process [9]. The training and testing dataset split introduced by Uy *et al.* [9] is followed in this work. In total, ~ 24.7 k point clouds were used for training and testing.

In-house dataset [9] consists of data from three regions in Singapore – a University Sector (U.S.), a Residential Area (R.A.), and a Business District (B.D.). Similar to the evaluation on the Oxford dataset, we use the standard test dataset split described in [9]. Point clouds were collected on a car travelling a path in U.S. (~ 10 km), R.A. (~ 8 km) and B.D. (~ 5 km) 5 times at different times. Point clouds from a single trip are used as queries and evaluated iteratively with point clouds from the remaining trips as database [9]. In-house dataset used at the test time only to demonstrate generalisability. In total, ~ 4.5 k point clouds were used for testing.

MulRan dataset [43] consists of traversals of several environments in South Korea – the Daejeon Convention Center (DCC) (3 runs each ~ 5 km), the Riverside (3 runs each ~ 6 km) of Daejeon city, the Korea Advanced Institute of Science and Technology (KAIST) (3 runs each ~ 7 km) and Sejong city (Sejong) (3 runs each ~ 23 km). We use the DCC and Riverside environments, training with DCC sequences 1 and 2 and testing on sequence 3, and training with Riverside sequences 1 and 3 and testing on sequence 2. KAIST three sequences are used as unseen only for evaluation. Because the average distance of point clouds in the MulRan dataset is ~ 1 m, we only use the point clouds with a minimum of 20 m apart, resulting in a total of ~ 2.5 k point clouds used for training and testing.

B. Evaluation Criteria

The datasets described above include UTM coordinates, *i.e.*, ground truth obtained from IMU/GPS for Oxford/in-house and IMU/GPS/SLAM for MulRan, for each point cloud. Using this ground truth, we select a 25 m threshold to classify successful place recognition events (retrievals). We compare the cosine similarity between the refined global descriptors of each query in a run with the refined global descriptors of the remaining point clouds covering the same region in the database. For comparison, we use AR@N (and its varieties, *i.e.*, AR@1 and AR@1%), commonly used for lidar PR performance. This metric measures the percentage of correctly localised queries where at least one of the top-N database predictions matches the query. A perfect AR@N score would be 100%, meaning all the possible revisits are correctly identified.

C. Implementation Details

We implemented P-GAT in PyTorch, and Adam optimiser was used. The learning rate was set to $1e^{-4}$ without learning rate decay. Our model was trained with batch size 256 for 1500 epochs on Oxford and MulRan. The number of attention layers was set to 9 with four heads for multi-head AGNN, resulting in ~ 12 million parameters (in total) for learning. We used this single configuration across all the experiments.

We create an adjacency matrix using a travel distance threshold to generate fully connected subgraphs. Using this matrix, we create subgraphs that include (key)nodes, each encompassing both the pose and its corresponding feature information. The adjacency matrix plays a crucial role in the training process for pairing subgraphs, ensuring that nodes within each subgraph are fully interconnected. Upon pairing subgraphs, we merge the nodes and features from both subgraphs, forming a fully connected graph. Nodes define edges in between, and features are populated with descriptors obtained from an existing PR method. We randomly pair subgraphs from the database to select two subgraphs as input. Because most of the pairs of subgraphs do not overlap, the ground truth of the similarity matrix is almost a zero matrix, not allowing the model to update the parameters properly. Therefore, we defined a positive rate and set it to 30%, forcing the model to have a 30% probability of having subgraph pairs with at least one pair of positive nodes.

Subgraph generation is based on the robot’s travelling distance. A subgraph starts from node \mathcal{V}_i and stops at node \mathcal{V}_{i+n} when the travelling distance exceeds a threshold. The stride of the subgraph generation is one, *i.e.*, the next subgraph starts from node \mathcal{V}_{i+1} . Since the number of subgraph nodes can vary, we use a boolean mask to deal with padding nodes. Also, we implemented a customised version of layer normalisation (compatible with the dynamic behaviour of subgraphs) in the MLP blocks, calculating all nodes’ mean and standard deviation except padding nodes. This approach effectively handles the dynamic node number within subgraphs.

During testing, because one query node can appear in multiple subgraphs (unlike point cloud pair-wise comparison), we implemented an average scheme in which the similarity scores related to a query node are averaged when the query node is compared against the entire nodes from the database, *i.e.*, $\bar{s}_{ij} = \frac{1}{C_1 C_2} \sum_{n_1=1}^{C_1} \sum_{n_2=1}^{C_2} s_{ij}(\mathbf{d}_i^{(n_1)}, \mathbf{d}_j^{(n_2)})$, where C_1 and C_2 are the maximum numbers of subgraphs from which the i -th query node and j -th database node are seen, respectively.

D. Comparison to State-of-the-Art

To show that P-GAT is backbone agnostic, we separately trained the model on Oxford RobotCar [42] and MulRan datasets [43] using descriptors $\mathbf{d} \in \mathbb{R}^{256}$ obtained from the backbone models PointNetVLAD [9], MinkLoc3D [10], and PPT-Net [13]. Descriptors of these baselines were generated after training the models on the training splits described earlier.

Table I summarises the average recall@1% (AR@1%) and average recall@1 (AR@1) of the results of our P-GAT and three other baseline models. Across all the experiments, P-GAT displays substantial performance gains (on average above 10%) compared to the baselines, indicating that if P-GAT is integrated with any backbone, it improves the performance due to spatiotemporal information. Additionally, following the trend of models’ performance when trained on Oxford and tested on MulRan/in-house or trained on MulRan and tested on Oxford/in-house, P-GAT demonstrates more consistent improvement in performance, showing that P-GAT enhances

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

TABLE I: Across all the baseline and all dataset combinations, our P-GAT exhibits performance gains in place recognition (measured by Recall@1). The best performance is indicated in bold.

| Trained on: | | Tested on: | | | | | | | | | | | | | | | | |
|------------------|--------------|------------|-------------|-------------|-------------|--------------|-------------|-------------|-------------|-------------|-------------|--------------|-------------|--------------|-------------|-------------|-------------|-------------|
| | | Oxford | | DCC | | Riverside | | B.D. | | R.A. | | U.S. | | KAIST | | Average | | |
| | | AR@1 | AR@1% | AR@1 | AR@1% | AR@1 | AR@1% | AR@1 | AR@1% | AR@1 | AR@1% | AR@1 | AR@1% | AR@1 | AR@1% | AR@1 | AR@1% | |
| Oxford | PointNetVLAD | Baseline | 62.8 | 83.8 | 69.3 | 81.6 | 37.0 | 56.6 | 56.2 | 65.3 | 52.3 | 68.4 | 63.1 | 79.0 | 57.2 | 73.1 | 56.8 | 72.5 |
| | | P-GAT | 96.5 | 99.9 | 80.8 | 92.1 | 57.0 | 83.5 | 89.2 | 98.6 | 76.3 | 99.6 | 79.0 | 98.2 | 73.8 | 82.3 | 78.9 | 93.5 |
| | MinkLoc3D | Baseline | 93.0 | 97.9 | 77.7 | 91.2 | 46.3 | 81.1 | 81.5 | 88.5 | 80.4 | 91.2 | 86.7 | 95.0 | 78.8 | 91.6 | 77.8 | 90.9 |
| | | P-GAT | 98.0 | 99.9 | 93.9 | 94.7 | 75.6 | 86.4 | 98.0 | 99.8 | 94.3 | 100.0 | 98.0 | 100.0 | 89.1 | 94.2 | 92.4 | 96.4 |
| | PPT-Net | Baseline | 93.5 | 98.1 | 76.0 | 89.6 | 39.4 | 84.3 | 84.6 | 90.0 | 84.1 | 93.3 | 90.1 | 97.5 | 77.4 | 88.3 | 77.9 | 91.6 |
| | | P-GAT | 97.4 | 99.9 | 93.6 | 99.2 | 73.1 | 94.0 | 89.0 | 98.1 | 82.9 | 99.8 | 90.9 | 99.9 | 82.3 | 91.4 | 87.4 | 97.5 |
| DCC Riverside | PointNetVLAD | Baseline | 40.9 | 57.6 | 79.3 | 90.9 | 52.5 | 81.6 | 52.2 | 60.5 | 45.0 | 58.5 | 49.1 | 65.5 | 64.9 | 78.9 | 54.8 | 70.5 |
| | | P-GAT | 65.6 | 93.9 | 96.9 | 99.9 | 83.2 | 96.5 | 52.5 | 76.9 | 61.9 | 95.9 | 58.5 | 96.4 | 79.0 | 91.4 | 71.1 | 93.0 |
| | MinkLoc3D | Baseline | 68.9 | 81.5 | 95.4 | 99.2 | 78.1 | 92.7 | 79.4 | 85.7 | 79.5 | 88.7 | 84.0 | 92.5 | 85.3 | 93.5 | 81.5 | 90.5 |
| | | P-GAT | 78.9 | 97.3 | 93.2 | 98.9 | 87.3 | 95.6 | 95.4 | 99.9 | 81.3 | 99.2 | 89.5 | 99.9 | 82.2 | 92.8 | 86.8 | 97.7 |
| | PPT-Net | Baseline | 63.3 | 78.8 | 91.7 | 99.6 | 66.1 | 94.6 | 66.4 | 75.5 | 65.7 | 81.8 | 77.2 | 89.4 | 83.6 | 93.6 | 73.4 | 87.6 |
| | | P-GAT | 69.1 | 96.3 | 98.9 | 100.0 | 87.3 | 99.4 | 77.3 | 94.6 | 74.9 | 97.9 | 80.2 | 99.8 | 78.8 | 89.8 | 80.9 | 96.8 |




Fig. 3: Variance analysis: Following the training and testing split protocol of [9], we re-trained P-GAT, MinkLoc3D, PointNetVLAD and PPT-Net 5 times, with a fixed parameter setting, on the Oxford training set and evaluated on the Oxford/in-house testing split.

TABLE II: Average recall (%) at top 1% (AR@1%) and top 1 (AR@1) for the state-of-the-art lidar-based PR models trained on the Oxford RobotCar. Our P-GAT (trained on top of MinkLoc3D) performs best on all benchmarks.

| | Oxford | | U.S. | | R.A. | | B.D. | |
|----------------------|-------------|-------------|-------------|--------------|-------------|--------------|-------------|-------------|
| | AR@1 | AR@1% | AR@1 | AR@1% | AR@1 | AR@1% | AR@1 | AR@1% |
| PointNetVLAD [9] | 62.8 | 80.3 | 63.2 | 72.6 | 56.1 | 60.3 | 57.2 | 65.3 |
| PCAN [11] | 69.1 | 83.8 | 62.4 | 79.1 | 56.9 | 71.2 | 58.1 | 66.8 |
| LPD-Net [32] | 86.3 | 94.9 | 87.0 | 96.0 | 83.1 | 90.5 | 82.5 | 89.1 |
| EPC-Net [44] | 86.2 | 94.7 | — | 96.5 | — | 88.6 | — | 84.9 |
| HiTPR [45] | 86.6 | 93.7 | 80.9 | 90.2 | 78.2 | 87.2 | 74.3 | 79.8 |
| SOE-Net [12] | 89.4 | 96.4 | 82.5 | 93.2 | 82.9 | 91.5 | 83.3 | 88.5 |
| MinkLoc3D [10] | 93.0 | 97.9 | 86.7 | 95.0 | 80.4 | 91.2 | 81.5 | 88.5 |
| NDT-Transformer [46] | 93.8 | 97.7 | — | — | — | — | — | — |
| PPT-Net [13] | 93.5 | 98.1 | 90.1 | 97.5 | 84.1 | 93.3 | 84.6 | 90.0 |
| SVT-Net [47] | 93.7 | 97.8 | 90.1 | 96.5 | 84.3 | 92.7 | 85.5 | 90.7 |
| MinkLoc3D-S [26] | 92.8 | 81.7 | 83.1 | 67.7 | 72.6 | 57.1 | 70.4 | 62.2 |
| PVT3D [48] | 95.6 | 98.5 | 92.9 | 97.9 | 89.5 | 94.8 | 87.9 | 92.1 |
| P-GAT (Ours) | 98.0 | 99.9 | 98.0 | 100.0 | 94.3 | 100.0 | 98.0 | 99.8 |

the generalisation ability of the network when train and test data are from different distributions.

We further compare our P-GAT with the state-of-the-art methods listed in Table II. For this, P-GAT was trained using descriptors extracted from MinkLoc3D trained on the Oxford training subset. Comparing P-GAT and MinkLoc3D baseline, we observe that AR@1 for Oxford, U.S., R.A., and B.D. improves by 5%, 11.3%, 13.9%, and 16.5%, respectively. Additionally, although we used MinkLoc3D descriptors, the substantial gain obtained by P-GAT allows us to outperform the other state-of-the-art methods, such as PPT-Net, SVT-Net and PVT3D, which reported higher performance than that of vanilla MinkLoc3D. AR@1% scores for all datasets are all higher than 99%. High AR@1% indicates that P-GAT enables the identification of all possible revisits in a small subset of top candidates (1% of point clouds in the database).

To evaluate the performance variance of our model compared to the baselines, we trained P-GAT, PointNetVLAD, MinkLoc3D and PPT-Net 5 times on the Oxford training sub-

TABLE III: Impact of the positional encoding in the subgraph.

| | With position | | Without position | |
|--------|---------------|-------|------------------|-------|
| | AR@1 | AR@1% | AR@1 | AR@1% |
| Oxford | 98.0 | 99.9 | 95.8 | 99.8 |
| R.A. | 94.3 | 100.0 | 53.3 | 85.1 |
| U.S. | 98.0 | 100.0 | 58.8 | 89.8 |
| B.D. | 98.0 | 99.8 | 59.5 | 82.2 |

set and tested on the Oxford/in-house testing subset using fixed settings. Fig. 3 shows mean AR@N recalls (average of AR@N over 5 experiments) and their standard deviations for all four models. The variance of our model is comparable with that of MinkLoc3D and PPT-Net, and it decreases by increasing the number of top candidates. PointNet-VLAD shows greater variance than the other methods indicating the network has more nondeterministic behaviour. The recall curves provide a visual representation of the superior performance of our P-GAT model, demonstrating the potential of our attention-based approach in improving the accuracy and stability of place recognition.

E. Ablation Studies

We conducted ablation studies to validate our proposed method and the relative contribution of each component. For this purpose, we trained and evaluated P-GAT utilising the MinkLoc3D model trained on the Oxford dataset. This ensures a fair comparison with the results reported in Table II.

Effects of Positional Encoding: Table III compares the performance of our model with and without positional information. The model without positional information shows a negligible drop (2.2% in AR@1) in performance on the Oxford testing split. However, the model lacking positional information exhibited a considerably inferior performance on the in-house datasets. As seen, AR@1 and AR@1% decrease more than 40% and 10%, respectively, over R.A., U.S. and

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

TABLE IV: Impact of intra- and inter-attention mechanisms.

| Intra-Attn | Inter-Attn | Oxford | | U.S. | | R.A. | | B.D. | |
|------------|------------|-------------|-------------|------|--------------|-------------|--------------|-------------|-------------|
| | | AR@1 | AR@1% | AR@1 | AR@1% | AR@1 | AR@1% | AR@1 | AR@1% |
| - | - | 95.6 | 99.3 | 85.9 | 95.2 | 92.1 | 97.7 | 84.8 | 93.9 |
| - | ✓ | 96.4 | 99.7 | 94.7 | 99.9 | 92.1 | 99.9 | 97.6 | 99.9 |
| ✓ | - | 97.0 | 99.9 | 78.7 | 98.4 | 87.8 | 98.8 | 73.1 | 89.5 |
| ✓ | ✓ | 98.0 | 99.9 | 98.0 | 100.0 | 94.3 | 100.0 | 98.0 | 99.8 |

TABLE V: Impact of the travel distance in the subgraphs.

| Distance | Oxford | | U.S. | | R.A. | | B.D. | | | | | |
|----------|--------|-------------|-------------|-------|-------------|--------------|-------|-------------|--------------|----|-------------|-------------|
| | avg.N | AR@1 | AR@1% | avg.N | AR@1 | AR@1% | avg.N | AR@1 | AR@1% | | | |
| 50 | 3 | 74.3 | 99.2 | 2 | 57.1 | 89.3 | 2 | 52.3 | 87.6 | 2 | 42.1 | 73.8 |
| 100 | 5 | 91.1 | 99.9 | 4 | 78.7 | 98.4 | 4 | 87.8 | 98.8 | 4 | 73.1 | 89.5 |
| 200 | 10 | 98.0 | 99.9 | 8 | 98.0 | 100.0 | 8 | 94.3 | 100.0 | 8 | 98.0 | 99.8 |
| 300 | 15 | 88.6 | 97.1 | 12 | 83.1 | 99.3 | 12 | 69.7 | 98.8 | 12 | 90.4 | 99.6 |

*avg.N denotes the averaged number of nodes (rounded to nearest integer) in the subgraphs

B.D. These results demonstrate that positional information is critical for our model’s performance, especially when testing in unseen environments.

Intra- and Inter-Attention Mechanisms: In this ablation study, we examine the impact of inter- and intra-attention mechanisms on our model’s performance. Table IV displays the results, indicating that disabling intra-attention produces a slight decrease in performance across all datasets. On the other hand, disabling inter-attention results in a slight improvement in the AR@1 and AR@1% scores on the Oxford dataset but a significant decrease in AR@1 scores on the in-house datasets. The U.S. and R.A. datasets show a decrease of around 1% in AR@1% scores, whereas B.D. experiences a decrease of 10%. Notably, the model without the full attention mechanism demonstrates the lowest AR@1 and AR@1% scores on the Oxford dataset, while its performance on the in-house datasets is comparable and sometimes even better than the model without inter-attention. These findings suggest that both inter- and intra- attention mechanisms enhance place recognition accuracy. Moreover, disabling inter-attention may lead to overfitting on the training dataset, resulting in poor generalisation at test-time with distribution shifts.

Travel Distance in Subgraphs: We conducted an ablation study to investigate the impact of the robot’s travelling distance within subgraphs on the performance of our model. Table V tabulates the performance results for varying travelling distances (50 m, 100 m, 200 m, and 300 m) in subgraphs. The model trained on subgraphs with a 50 m length exhibits the lowest AR@1 and AR@1% scores on all datasets. As travelling distances in subgraphs increase, the model performance also improves. The optimal results were obtained from the subgraphs with a travelling distance of 200 m. However, the model trained on subgraphs with a travelling distance of 300 m exhibits a slight performance drop. This ablation study demonstrates the importance of the subgraph length in aggregating information within and across subgraphs. Short subgraphs avoid proper receptive fields between similar and dissimilar point clouds. On the other hand, large subgraphs contain irrelevant information, biasing the model performance.

F. Runtime Analysis and Memory Usage

We evaluated the computation time taken on average for each keynode in a subgraph and memory consumption to demonstrate that our presented system can run online. The timing results are collected by running the pre-trained models on a single NVIDIA RTX A3000 Mobile GPU with an Intel(R)

TABLE VI: Comparison of inference speed and memory consumption.

| Model | Memory Usage | Runtime |
|----------------------|---------------------|--------------|
| PointNetVLAD + P-GAT | (2.95+1.13=4.08)GiB | (22+18=40)ms |
| LPD-Net + P-GAT | (1.94+1.13=3.07)GiB | (35+18=53)ms |
| MinkLoc3D + P-GAT | (0.85+1.13=1.98)GiB | (29+18=47)ms |
| PPT-Net + P-GAT | (0.96+1.13=2.09)GiB | (30+18=48)ms |

Xeon(R) W-11855M CPU @ 3.20GHz CPU. Table VI reports P-GAT’s runtime per frame using embeddings extracted by the baselines listed. Overall, P-GAT adds a constant memory of ~ 1.1 GiB and ~ 20 ms to the inference time, demonstrating that, in an end-to-end fashion, the total computation time allows online operation.

V. CONCLUSION

We proposed P-GAT for large-scale place recognition tasks. Through the attention mechanism and graph neural network, we increase the descriptors’ distinctiveness by relating the entire point clouds collected in a sequence and in nearby locations in revisit areas. This design allows for exploiting spatiotemporal information. Extensive experiments on the Oxford dataset, the in-house dataset and the MulRan dataset demonstrate the effectiveness of the proposed method and its superiority compared to the state-of-the-art. Our proposed P-GAT’s average performance across key benchmarks is superior by above 10% over the original baselines demonstrating that P-GAT can be incorporated into any global descriptors, substantially improving their robustness and generalisation ability. In future work, we would like to extend the P-GAT for point cloud registration task to accurately determine the position and orientation of retrieved places.

REFERENCES

- [1] J. Huang, B. Zhou, Z. Fan, Y. Zhu, Y. Jie, L. Li, and H. Cheng, “FAEL: Fast Autonomous Exploration for Large-Scale Environments with a Mobile Robot,” *IEEE Robotics and Automation Letters*, 2023.
- [2] P.-E. Sarlin, M. Dusmanu, J. L. Schönberger, P. Speciale, L. Gruber, V. Larsson, O. Miksik, and M. Pollefeys, “LaMAR: Benchmarking localization and mapping for augmented reality,” *European Conference on Computer Vision*, 2022, pp. 686–704.
- [3] J. Knights, K. Vidanapathirana, M. Ramezani, S. Sridharan, C. Fookes, and P. Moghadam, “Wild-Places: A Large-Scale Dataset for Lidar Place Recognition in Unstructured Natural Environments,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 11 322–11 328.
- [4] M. Ramezani, E. Griffiths, M. Haghghat, A. Pitt, and P. Moghadam, “Deep Robust Multi-Robot Re-localisation in Natural Environments,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [5] R. Wang, Y. Shen, W. Zuo, S. Zhou, and N. Zheng, “TransVPR: Transformer-Based Place Recognition with Multi-Level Attention Aggregation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 13 648–13 657.
- [6] A. Ali-bey, B. Chaib-draa, and P. Giguère, “MixVPR: Feature Mixing for Visual Place Recognition,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 2998–3007.
- [7] H. Zhang, X. Chen, H. Jing, Y. Zheng, Y. Wu, and C. Jin, “ETR: An Efficient Transformer for Re-Ranking in Visual Place Recognition,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 5665–5674.
- [8] Z. Li, C. D. W. Lee, B. X. L. Tung, Z. Huang, D. Rus, and M. H. Ang, “Hot-NetVLAD: Learning Discriminatory Key Points for Visual Place Recognition,” *IEEE Robotics and Automation Letters*, 2023.
- [9] M. A. Uy and G. H. Lee, “PointNetVLAD: Deep Point Cloud Based Retrieval for Large-Scale Place Recognition,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

- [10] J. Komorowski, "MinkLoc3D: Point Cloud Based Large-Scale Place Recognition," in *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2021, pp. 1789–1798.
- [11] Zhang, Wenxiao and Xiao, Chunxia, "PCAN: 3D Attention Map Learning Using Contextual Information for Point Cloud Based Retrieval," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 436–12 445.
- [12] Xia, Yan and Xu, Yusheng and Li, Shuang and Wang, Rui and Du, Juan and Cremers, Daniel and Stilla, Uwe, "SOE-Net: A Self-Attention and Orientation Encoding Network for Point Cloud based Place Recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 11 348–11 357.
- [13] Hui, Le and Yang, Hang and Cheng, Mingmei and Xie, Jin and Yang, Jian, "Pyramid Point Cloud Transformer for Large-Scale Place Recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6098–6107.
- [14] J. Ma, X. Chen, J. Xu, and G. Xiong, "SeqOT: A Spatial-Temporal Transformer Network for Place Recognition Using Sequential LiDAR Data," *IEEE Transactions on Industrial Electronics*, 2022.
- [15] K. Vidanapathirana, P. Moghadam, B. Harwood, M. Zhao, S. Sridharan, and C. Fookes, "Locus: LiDAR-based Place Recognition using Spatiotemporal Higher-Order Pooling," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 5075–5081.
- [16] R. B. Rusu, N. Blodow, Z. C. Marton, and M. Beetz, "Aligning Point Cloud Views Using Persistent Feature Histograms," in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2008, pp. 3384–3391.
- [17] R. B. Rusu, N. Blodow, and M. Beetz, "Fast Point Feature Histograms (FPFH) for 3D Registration," in *2009 IEEE International Conference on Robotics and Automation*, 2009, pp. 3212–3217.
- [18] K. P. Cop, P. V. Borges, and R. Dubé, "DELIGHT: An Efficient Descriptor for Global Localisation using LiDAR Intensities," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 3653–3660.
- [19] G. Kim and A. Kim, "Scan Context: Egocentric Spatial Descriptor for Place Recognition within 3D Point Cloud Map," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 4802–4809.
- [20] R. Dubé, D. Dugas, E. Stumm, J. Nieto, R. Siegwart, and C. Cadena, "SegMatch: Segment Based Place Recognition in 3D Point Clouds," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 5266–5272.
- [21] S. Salti, F. Tombari, and L. Di Stefano, "SHOT: Unique signatures of histograms for surface and texture description," *Computer Vision and Image Understanding*, vol. 125, pp. 251–264, 2014.
- [22] Dubé, Renaud and Cramariuc, Andrei and Dugas, Daniel and Nieto, Juan and Siegwart, Roland and Cadena, Cesar, "SegMap: 3D Segment Mapping using Data-Driven Descriptors," *Robotics: Science and Systems Online Proceedings*, vol. 14, 2018.
- [23] Tinchev, Georgi and Penate-Sanchez, Adrian and Fallon, Maurice, "Learning to See the Wood for the Trees: Deep Laser Localization in Urban and Natural Environments on a CPU," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1327–1334, 2019.
- [24] Renaud Dubé and Andrei Cramariuc and Daniel Dugas and Hannes Sommer and Marcin Dymczyk and Juan Nieto and Roland Siegwart and Cesar Cadena, "SegMap: Segment-based mapping and localization using data-driven descriptors," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 339–355, 2020.
- [25] Ramezani, Milad and Tinchev, Georgi and Iuganov, Egor and Fallon, Maurice, "Online LiDAR-SLAM for Legged Robots with Robust Registration and Deep-Learned Loop Closure," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 4158–4164.
- [26] K. Żywanowski, A. Banaszczyk, M. R. Nowicki, and J. Komorowski, "MinkLoc3D-SI: 3D LiDAR Place Recognition With Sparse Convolutions, Spherical Coordinates, and Intensity," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1079–1086, 2021.
- [27] J. Komorowski, M. Wysoczanska, and T. Trzcinski, "EgoNN: Egocentric Neural Network for Point Cloud Based 6DoF Relocalization at the City Scale," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 722–729, 2021.
- [28] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature Pyramid Networks for Object Detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2117–2125.
- [29] Vidanapathirana, Kavisha and Ramezani, Milad and Moghadam, Peyman and Sridharan, Sridha and Fookes, Clinton, "LoGG3D-Net: Locally Guided Global Descriptor Learning for 3D Place Recognition," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2215–2221.
- [30] Qi, Charles R and Su, Hao and Mo, Kaichun and Guibas, Leonidas J, "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.
- [31] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "NetVLAD: CNN architecture for weakly supervised place recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 5297–5307.
- [32] Liu, Zhe and Zhou, Shunbo and Suo, Chuanzhe and Yin, Peng and Chen, Wen and Wang, Hesheng and Li, Haoang and Liu, Yun-Hui, "LPD-Net: 3D Point Cloud Learning for Large-Scale Place Recognition and Environment Analysis," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 2831–2840.
- [33] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [34] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention Is All You Need," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [35] P.-E. Sarlin, D. DeTone, T. Malisiewicz, and A. Rabinovich, "SuperGlue: Learning Feature Matching with Graph Neural Networks," in *CVPR*, 2020.
- [36] C. Min, J. Xu, L. Xiao, D. Zhao, Y. Nie, and B. Dai, "Attentional Graph Neural Network for Parking-Slot Detection," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3445–3450, 2021.
- [37] S. Wang, Z. Chen, D. Li, Z. Li, L.-A. Tang, J. Ni, J. Rhee, H. Chen, and P. S. Yu, "Attentional Heterogeneous Graph Neural Network: Application to Program Reidentification," in *Proceedings of the 2019 SIAM International Conference on Data Mining*. SIAM, 2019, pp. 693–701.
- [38] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela, "Community Structure in Time-Dependent, Multiscale, and Multiplex Networks," *science*, vol. 328, no. 5980, pp. 876–878, 2010.
- [39] J. Gehring, M. Auli, D. Grangier, D. Yarats, and Y. N. Dauphin, "Convolutional Sequence to Sequence Learning," in *International conference on machine learning*. PMLR, 2017, pp. 1243–1252.
- [40] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner *et al.*, "Relational inductive biases, deep learning, and graph networks," *arXiv preprint arXiv:1806.01261*, 2018.
- [41] P.-T. De Boer, D. P. Kroese, S. Mannor, and R. Y. Rubinstein, "A Tutorial on the Cross-Entropy Method," *Annals of Operations Research*, vol. 134, pp. 19–67, 2005.
- [42] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 Year, 1000km: The Oxford RobotCar Dataset," *The International Journal of Robotics Research (IJRR)*, vol. 36.
- [43] G. Kim, Y.-S. Park, Y. Cho, J. Jeong, and A. Kim, "MulRan: Multimodal Range Dataset for Urban Place Recognition," *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6246–6253, 2020.
- [44] L. Hui, M. Cheng, J. Xie, J. Yang, and M.-M. Cheng, "Efficient 3D Point Cloud Feature Learning for Large-Scale Place Recognition," *IEEE Transactions on Image Processing*, vol. 31, pp. 1258–1270, 2022.
- [45] Z. Hou, Y. Yan, C. Xu, and H. Kong, "HiTPR: Hierarchical Transformer for Place Recognition in Point Cloud," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2612–2618.
- [46] Z. Zhou, C. Zhao, D. Adolphson, S. Su, Y. Gao, T. Duckett, and L. Sun, "NDT-Transformer: Large-Scale 3D Point Cloud Localisation using the Normal Distribution Transform Representation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 5654–5660.
- [47] Z. Fan, Z. Song, H. Liu, Z. Lu, J. He, and X. Du, "SVT-Net: Super Light-Weight Sparse Voxel Transformer for Large Scale Place Recognition," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 1, 2022, pp. 551–560.
- [48] Y. Xia, M. Gladkova, R. Wang, J. F. Henriques, D. Cremers, and U. Stilla, "PVT3D: Point Voxel Transformers for Place Recognition from Sparse Lidar Scans," *arXiv preprint arXiv:2211.12542*, 2022.