MinkUNeXt: Point Cloud-based Large-scale Place Recognition using 3D Sparse Convolutions

J.J. Cabrera¹, A. Santo^{1,2}, A. Gil¹, C. Viegas³ and L. Payá¹

Abstract—This paper presents MinkUNeXt, an effective and efficient architecture for place-recognition from point clouds entirely based on the new 3D MinkNeXt Block, a residual block composed of 3D sparse convolutions that follows the philosophy established by recent Transformers but purely using simple 3D convolutions. Feature extraction is performed at different scales by a U-Net encoder-decoder network and the feature aggregation of those features into a single descriptor is carried out by a Generalized Mean Pooling (GeM). The proposed architecture demonstrates that it is possible to surpass the current state-ofthe-art by only relying on conventional 3D sparse convolutions without making use of more complex and sophisticated proposals such as Transformers, Attention-Layers or Deformable Convolutions. A thorough assessment of the proposal has been carried out using the Oxford RobotCar and the In-house datasets. As a result, MinkUNeXt proves to outperform other methods in the state-of-the-art.

Index Terms—Place Recognition, LiDAR, Point cloud embedding, 3D Sparse Convolutions

I. INTRODUCTION

I N many applications, mobile robots must perform autonomous navigation in a specific environment. As it moves, the robot should be able to recognize or identify different areas of the environment. This action is equivalent to finding a correspondence between its current sensor observations and a part of the stored database. This ability is commonly denoted as place recognition. In order to speed this process, frequently, authors have concentrated on describing some parts of the environment by means of an invariant descriptor. In this way, the robot should be able to recognize a part of the environment by finding the descriptor in the database that most ressembles the descriptor associated to its current observations. The concept of place recognition is of uttermost importance in tasks such as localization, mapping and navigation.

Place recognition and robot localization are two closely related concepts. Place recognition centers on the description of the current robot observations in a way that allows the robot to identify different locations in the map. Thus, place recognition focuses on the extraction and codification of relevant features found in the robot query observation in

³Univ of Coimbra, ADAI, Department of Mechanical Engineering, Rua Luís Reis Santos, Pólo II, 3030-788 Coimbra, Portugal. Email: carlos.viegas@uc.pt

979-8-3503-0704-7/23/\$31.00 ©2023 European Union



Fig. 1. Point cloud-based place recognition. Each query point cloud (red) is embedded into a global descriptor which is compared with the descriptors from the database point clouds (blue) by means of a Nearest Neighbour Search.

such a way that they can be compared to previously stored data (Fig. 1). Similarly, robot localization refers to the act of estimating the position and orientation of the robot within a known map. In this way, given a map of the robot, conformed by a series of submaps, a common process to carry out the global localization of the robot could consist of two phases [1]: a) rapidly finding a submap within the global database using the feature descriptors (place recognition) and b) performing a fine estimation of the position and orientation of the robot in that submap (robot localization). A similar technique is proposed in [2], where the descriptor is computed from LiDAR measurements. Next, a handcrafted descriptor is employed to rapidly retrieve some areas of interest in the map. The final localization step, based on the ICP (Iterative Closest Point) algorithm is able to compute the position and orientation within the submap.

To date, place recognition has been performed with different types of sensors: visual cameras [3], laser [4], LiDAR [5] and Radar [6] using different types of techniques. For example, place recognition has been extensively solved by means of techniques based on the Bag of Words algorithm using images [7], [8].

During the last few years LiDAR sensors have lowered in

¹Institute for Engineering Research (I3E) Email: {juan.cabreram, a.santo, arturo.gil, lpaya}@umh.es

²Valencian Graduate School and Research Network for Artificial Intelligence.

price and weight, while increasing in resolution. Therefore, LiDAR sensors permit obtaining a large number of precise measurements from the environment that define its shape and structure. Being a self-illuminated sensor, it is insensitive to changes in natural light, thus directly applicable to outdoor applications. In consequence, several new potential applications of LiDAR sensors in the area of mobile robotics have emerged and it is therefore necessary to focus on methods that achieve a robust description of the scene. In the literature, so far, we can find: a) Classical techniques based on a handcrafted description of LiDAR data to generate rotationally invariant representations [2], [9] and b) Descriptions based on the use of Deep Neural Networks, either operating directly on the coordinates of the points [10] or on the projection of the points to image coordinates [11].

In this manuscript a technique for the robust description of scenes captured by a LiDAR sensor based on the use of a Deep Neural Network is presented. Several improvements and modifications are proposed starting from the basis of several recent architectures. As a result, the proposed network is able to outperform all other existing methods in the context of place recognition. In summary, the main contributions of this paper are:

- MinkUNeXt: A new 3D Sparse Convolutional Neural Network for Place-Recognition. It is the first approach of a U-Net architecture for point cloud embedding and place-recognition. The architecture has been specifically developed and optimized for this problem. In addition, substantial improvements have been achieved both in terms of macro and micro design.
- The definition of a new residual block: the 3D MinkNext Block, which is entirely composed of 3D sparse convolutions and surpasses the performance of ResNet Blocks. It follows the philosophy proposed by ConvNeXt [12], which uses standard convolutions and was originally proposed for image classification, semantic segmentation and object detection.

As a result, the proposed topology is able to surpass significantly the current state of the art of point cloud place-recognition in terms of average recall at 1 (AR@1) and average recall at 1% (AR@1%), when compared to the most relevant methods in the literature.

The rest of the paper is organized as follows: Next, Section II deepens in the state of the art in relation with the use of Deep Neural Networks for the description of the structure of point clouds. After that, Section III defines in detail the proposed architecture to describe the point clouds. Next, Section IV describes the datasets, experiments and results. Finally, Section V presents the main conclusions.

II. STATE OF THE ART

This section offers a comprehensive overview of the current state-of-the-art in place recognition, specifically exploring the utilization of Deep Neural Networks with point cloud data. Many applications have emerged that concentrate on place recognition based on point clouds. In this section the methods are presented chronologically. In addition, in this manuscript a comparison of the main results achieved by the most relevant architectures is provided. In this context, the first approach to this task was tackled in [10] with PointNetVLAD, a network model based on PointNet [13] for feature extraction followed by a NetVLAD layer for feature aggregation. The point clouds taken as input by this type of architectures do not need to be sorted, as they use symmetric functions such as Multi Layer Perceptron (MLP) or Fully Connected layers. Next, a similar approach, named LPD-Net [14] improved the state of the art by incorporating a local feature extraction block at the beginning of the network and a subsequent graph-based neighbourhood aggregation.

After that, the MinkLoc3D architecture emerged [15]. It is based on a Feature Pyramid Network (FPN) with Sparse Convolutions for feature extraction [16], followed by a Generalized Mean Pooling (GeM) for the aggregation of the features into a single vector [17]. At that time, the MinkLoc3D architecture marked a significant milestone, as it significantly surpassed the existing state-of-the-art methods and also demonstrated that the use of 3D convolutional layers was a good choice for feature extraction from point clouds. Unlike previous network typologies, when using 3D convolutions, they do require a sorted point cloud as input, where the spatial relationships between points are preserved. The same situation occurs in an analogous way with images, where 2D convolutions have proven to be very efficient in feature extraction thanks to the neighbourhood relationships between pixels. In this sense, some 2D architectures have also emerged taking as input the projected point cloud into a spherical image (OverlapNet [11]). Other works, such as [1] propose creating a rotation-invariant handcrafted image: from a polar coordinate representation of the point cloud, the 2D distance between consecutive points belonging to the same elevation angle (ring) is computed and then, a histogram per ring is obtained generating a 2D handcrafted codification of the point cloud.

In addition, both monocular images and point clouds are used simultaneously by some architectures (MinkLoc++ [18], PIC-Net [19]). In this case, both architectures are formed by two branches, processing independently the image and the point cloud. Each branch results in a feature vector and both vectors are finally aggregated into a single vector by a pooling process. Alternatively, each point can be associated with a feature corresponding to the RGB value of the image [20]. This requires a precise calibration of the camera-LiDAR system. Otherwise, some authors propose to use the relative intensity returned by each LiDAR ray, referred to as MinkLoc-SI [21].

The DAGC architecture [22] was the first to introduce self-attention layers [23] for point cloud feature extraction to perform place recognition. Later, other authors continued the use of attention layers, obtaining results close to the state of the art. In this sense, NDT-Transformer was presented





Fig. 2. This diagram shows the architecture of the proposed MinkUNeXt, which is based on a semantic segmentation network (U-Net) modified and enhanced to perform place-recognition from point clouds.

[24], a network model based on 3 Transformer Encoders that takes as input a modified point cloud by using a Normal Distribution Transform (NDT). This approach preserves the geometrical shape of the point cloud while decreasing the memory complexity.

Simultaneously, PPT-Net [25], a Transformer with a pyramidal distribution followed by a NetVLAD layer, emerged. Based on a similar idea, SOE-Net [26] extracts local features using a series of MLPs and subsequently, it applies attention layers in the aggregation of those features. In addition, the Retriever [27] network also introduces self-attention layers within an autoencoder to perform local feature aggregation. Besides, looking for efficiency and the use of these architectures in real localization systems (which must work in real time), SVT-Net, an efficient Sparse Voxel Transformer based on sparse convolutional layers for feature extraction, was presented in [28].

Furthermore, HiTPR [29] employs Farthest Point Sampling [30] to reduce the dimensionality of the input cloud while preserving its original topological information. In addition, this work introduces a Transformer block for short-range local feature extraction and an additional Transformer block for extracting global information over long distances. The mentioned Transformer-based approaches presented similar results to those found in the state of the art. However, the presentation of TransLoc3D [31] constituted a significant advance. It is a network model also based on sparse convolutions but unlike other proposals, it extracts features at different scales in parallel by means of convolutional layers with different kernel size. In addition, it also employs ECA (Efficient Channel Attention) layers [32] in order to interact local features from different channels. This type of layers are also used by MinkLoc3Dv2 [5], an architecture based on MinkLoc3D [15]. MinkLoc3Dv2

includes the use of ECAs with an increased number of planes or channels (depth of the convolution matrices). To date, this network architecture shows the best results in terms of average recall at 1 (AR@1) in the Oxford RobotCar Dataset [33], partly due to the loss function they introduce in the training process and the high batch size with which they train.

Finally, the best result in terms of average recall at 1% (AR@1%) was obtained by KPPR [34], a network model based on Flexible and Deformable Convolutions (KPConv [35]). However, Minkloc3Dv2 is still ahead in terms of average recall at 1 (AR@1), which is a more demanding metric. Additional architectures have been proposed to date, making other types of contributions such as rotation invariance E²PN-GeM [36] and RPR-Net [37] or inference efficiency EPC-Net [38] and BPT [39].

This paper presents MinkUNeXt, an architecture based on MinkUNet [16] modified and enhanced to perform placerecognition from point clouds. It is an encoder-decoder architecture entirely based on the proposed 3D MinkNeXt Block, a residual block composed of 3D sparse convolutions that follows the philosophy proposed by ConvNeXt [12]. The feature extraction is performed by the U-Net encoder-decoder and the feature aggregation of those features into a single descriptor is carried out by a Generalized Mean Pooling (GeM) [40]. The proposed architecture demonstrates that it is possible to surpass the current state of the art by only relying on conventional 3D sparse convolutions without making use of more complex and sophisticated frameworks such as Transformers, Attention-Layers or Deformable Convolutions. In this way, this paper shows that the proposed architecture outputs results which are superior to those found in the literature while maintaining the efficiency, scalability and performance.

III. MINKUNEXT: GLOBAL POINT CLOUD DESCRIPTOR FOR PLACE RECOGNITION

Place recognition from point clouds can be approached as an embedding task. For this purpose, it is desirable to have an architecture capable of extracting the more descriptive features of the scene and, in addition, aggregating them into a single vector that most generally describes the information present in the scene. The present work presents a pioneering solution that employs a U-Net architecture [41] in the context of place recognition. Most architectures resembling U-Net were originally designed for semantic segmentation, where the goal is to assign a category to each pixel of an input image, or in this case, to each point of the input point cloud. However, the encoder-decoder topology of a U-Net is also capable to extract and fuse relevant features from the scene as will be shown in the experimental section.

A. Global Architecture

The proposed model is fed by a point cloud given as an unordered set of 3D coordinates $P = \{(x_i, y_i, z_i)\}$. This point cloud is quantized into a sparse tensor, which is a high-dimensional extension of a sparse matrix where non-zero elements are represented as a set of indices C (coordinates) and associated values (or features) F. Some papers [14], [24] propose to employ as feature some handcrafted attributes such as the vertical component of the normal vector, height variance, change of curvature or just the value of the coordinates. Others [5], [15] prefer initializating each coordinate's feature to one, i.e., the first convolution (stem) will only take as input features 'ones' for the non-empty voxels. This idea is also taken in the present paper, where the input data $\hat{P} = \{(\hat{x}_i, \hat{y}_i, \hat{z}_i, 1)\}$ is conformed by C, a set of 3D quantized coordinates and F, a vector of 'ones' whose length is equal to the number of quantized points.

The global architecture is represented in Fig. 2. The encoder of the network consists of five 3D Sparse Convolutions (coloured in yellow). Among them, the stem is the first convolution and it preserves the input dimension of the point cloud since its stride is fixed to 1 and the kernel size is 5. While each of the following four convolutions gradually decrease the spatial dimension, the receptive field increases since successive convolutional layers capture larger and larger patterns by combining information from previous layers. Each of those convolutions downsample its input dimension by 2 since they employ both kernel size and stride of 2. After the encoder, the dimension of the input point cloud is downsampled by 32.

In a common U-Net the decoder is composed of four 3D Sparse Transpose Convolutions that upsample the spatial dimension by 2, progressively reconstructing the input cloud. However, in this architecture it is proposed to partially reconstruct the input point cloud by only applying three transpose convolutions (coloured in orange), since our purpose is point cloud embedding and not semantic segmentation. Subsection IV-E will justify that features extracted with only three transpose convolutions are more robust for understanding the overall context of the scene. Furthermore, a Batch Normalization and a ReLU activation function (coloured in red) are applied after all the convolutions, which helps in stabilizing the training process. In addition, in this architecture it is proposed to employ the presented Residual MinkNeXt Block (coloured in blue) instead of the common ResNet Block after each ReLU (without taking into account the one corresponding to the stem). This kind of residual blocks provide a direct path for gradients to flow through the network, reducing overfitting and boosting the generalization capabilities on unseen data. In this architecture, it is also used to increase the number of features maps as it will further detailed in the following subsection III-B.

The U-Net architecture is characterized for having skip connections between the encoder and the decoder. On the one hand, the encoder would capture features at different spatial scales, from fine details (low-level) to more global structures (high-level) present in point clouds. On the other hand, thanks to the skip connections, the decoder would fuse the low-level and high-level features. After that, a Fully Connected Layer is added since it outputs features have been proven to perform robustly against viewpoint changes in visual place recognition [42]. Furthermore, this Fully Connected Layer is also employed to extend the feature maps up to a dimensionality of 512. Subsequently, the points descriptors that conform that feature map are aggregated into a single global descriptor by a Generalized Mean Pooling (GeM) [40].

B. Residual Block Architecture

As mentioned before, in this paper both a global and a residual block architecture are proposed. In this sense, a new residual block is designed (Fig. 3) which is entirely composed of 3D Sparse Convolutions and follows the philosophy proposed by ConvNeXt [12], surpassing the performance of ResNet Blocks. We have named this block MinkNeXt, since it takes advantage of the ResNet blocks and is fully implemented in Minkowski Engine [16].

In the global architecture (Fig. 2), the proposed residual block appears in blue colour after each ReLU activation function (except for the one corresponding to the stem). Since the residual block is generally employed to increase the number of features maps, the stem of the residual block is formed by a 1x1x1 convolution that widens the input dimension to the output channels size. After that, an inverted bottleneck is applied by expanding the dimension four times and then reducing it again to the output dimension through two 3D Sparse Convolutions. This inverted bottleneck was originally proposed by MobileNetV2 [43] and nowadays, it is an important design in every Transformer block. In addition, a 1x1x1 Convolution in the residual connection is also applied when the input and the output dimensions differ.

The activation function employed in this block is the Gaussian Error Linear Unit (GeLU) [44], which is smoother than ReLU and is utilized in the most advanced Transformers.

Finally, the normalization is carried out by LayerNorms [45] in the main stream of the block and by BatchNorms [46] in the residual connection.



Fig. 3. This diagram shows the proposed MinkNeXt Block. This residual block is an essential part of the global network, since increases the number of feature maps through an inverted bottleneck.

IV. EXPERIMENTS

This section describes the datasets (Subsection IV-A), the labelling (Subsection IV-B) and the training and evaluation of the proposed architecture (Subsection IV-C). Later, the implementation details are described in Subsection IV-D. Subsequently, in Subsection IV-E, we present an ablation study of the designing steps carried out to obtain the final architecture. Finally, the main results are compared with other approaches in the literature in Subsection IV-F.

A. Datasets

In order to train and evaluate the proposed method, the datasets and evaluation protocols introduced in [10] have been used. This is a common framework employed and respected by a large number of studies that is used to compare different proposals that address the place recognition task using point clouds. The benchmark consists of 2 datasets and 4 different environments:

- The Oxford RobotCar Dataset [33]. This dataset is generated using some SICK LMS-151 2D sensors mounted on a car. The dataset covers a 10 km trajectory along the city of Oxford. In total, 44 sequences of the same trajectory which are geographically divided into training (70%) and test (30%) are used. This results in 21,711 training submaps and 3,030 test submaps.
- The In-house Dataset [10]. This dataset consists of three different environments: a University Sector (U.S.), a Residential Area (R.A.), and a Business District (B.D.). These datasets are captured using a Velodyne-64 LiDAR mounted on a motorized vehicle that covers each of the three regions. The paths lengths are 10 km, 8 km and 5 km respectively. It is conformed by 5 different sequences from each of the U.S., R.A. and B.D. regions, which were captured at different times. In addition, each U.S. and R.A. sequence are geographically divided into train

and test. While the B.D. environment is only used for testing.

In both datasets, the LiDAR scans are taken at regular intervals of 12.5 m and 25 m for the training and test set, respectively. Also, both datasets are formed by a number of submaps. Each submap is constructed by capturing LiDAR scans consecutively along 20 m. Next, the scans are registered in a common frame and further processed to create a consistent submap. Each of these training and test submaps are filtered by removing the ground plane and also regularly sampled by a voxel grid filter in order to reduce its size to 4096 points. The XYZ coordinates of the points that constitute each point cloud are then shifted and scaled in order to obtain a point distribution with zero mean in the [-1, 1] range for each coordinate.

B. Labelling and similarity

Each submap in the dataset is tagged with the UTM coordinates of its respective centroid. This constitutes the identifier of each submap and is later used during the training and evaluation of the network. Next, we define the similarity between the submaps in the datasets. This concept is generally denoted as labelling in the literature and it is important because it is necessary to feed the model with structurally similar submaps captured from the same place and structurally dissimilar submaps from different places. In this sense, most of the proposed labelling protocols are based on the Euclidean distance of the UTM coordinates from which point clouds are captured (two point clouds are considered structurally similar if they are captured within a distance p and structurally different if they are taken from a distance greater than n where p < n). This procedure, of course, is a coarse approximation that assumes that submaps captured from the same area will possess a similar structure. However, it is a simple but effective manner of labelling the training data. In this paper, this method is adopted with p = 10m and n = 50m as in the majority of the referred manuscripts. Authors, have also proposed other methods for similarity labelling in the context of placerecognition. For example, [11] proposes to use the overlap between point clouds as an alternative method for labelling similar and dissimilar point clouds. In order to compute the overlap between two point clouds (i.e. submaps) a precise registration must be carried out, which hinders the application of this technique to large datasets.

C. Training and evaluation

As for the training and evaluation of the proposed method, the two evaluation protocols established in [10] have been followed:

- The first, baseline protocol, consists in training the model only with the Oxford training data and evaluating with the Oxford and In-house (U.S., R.A. and B.D.) test data.
- The second, refined protocol, consists in training with Oxford and In-house (U.S., R.A.) training data and evaluating with the Oxford and In-house (U.S., R.A. and B.D.) test data.

TABLE I NUMBER OF TRAINING AND TESTING SUBMAPS FOR THE BASELINE AND REFINED PROTOCOLS.

	Baseline	Protocol	Refined Protocol			
	Training	Test	Training	Test		
Oxford	21.7k	3.0k	21.7k	3.0k		
In-house	-	4.5k	6.7k	1.7k		

Table I summarizes the number of training and testing submaps corresponding to each dataset and each of the protocols defined above. The assessment of the LiDAR-based place recognition descriptors is carried out by means of the recall rate at top-K candidates. Following the most common evaluation methods (as in the manuscripts cited in Section II), the average recall at 1 (AR@1) and average recall at 1% (AR@1%) are used in order to ease the comparison with other techniques. We start with a "query submap" formed by a point cloud which is taken from the test dataset and point clouds submaps from different traversals that cover the same region from the map. Each query submap is processed by the network and it outputs, as a result, a descriptor vector that codifies its appearance. This descriptor is referred to as the "query descriptor". Next, the query descriptor is compared to all the descriptors in the map. The point cloud in the database that minimizes the distance is selected. Finally, the place recognition is considered to be successful if the query and the retrieved point cloud are within an Euclidean distance of 25 m.

D. Implementation details

In the present work the proposed model is trained following the procedure established in [5]. In this regard, the Truncated Smooth-AP (TSAP) loss function is employed, which tries to maximize the ranking of the positive top-k candidates:

$$\mathcal{L}_{TSAP} = \frac{1}{b} \sum_{q=1}^{b} (1 - AP_q) \tag{1}$$

Where b is the batch size and AP_q is the smooth average precision:

$$AP_{q} = \frac{1}{|P|} \sum_{i \in P} \frac{1 + \sum_{j \in P, j \neq i} G(d(q, i) - d(q, j); \tau)}{1 + \sum_{j \in \Omega, j \neq i} G(d(q, i) - d(q, j); \tau)} \quad (2)$$

Given a query point cloud q, the average precision AP_q is computed from the set of k closest candidates P (positives) and the set of all positives and negatives Ω . Also, the function G constitutes a Sigmoid function $G(x;\tau) = (1 + \exp(-\frac{x}{\tau}))^{-1}$ with a parameter τ that controls the sharpness. The term d(q, i) represents the Euclidean distance between the descriptor of a query point cloud q and the *i*th point cloud. The numerator represents a soft ranking of a positive point *i* among the top k positives (where k = 4), while the denominator represents a soft ranking of a positive point *i* among all other positives and negatives.

For the correct performance of this type of loss function, it is necessary to train with a high batch size, specifically a size

 TABLE II

 TRAINING PARAMETERS IN BASELINE AND REFINED PROTOCOLS

Parameter	Baseline	Refined
Batch Size (b)	2048	2048
Number of Epochs	400	500
Initial Learning Rate	1×10^{-3}	1×10^{-3}
LR Scheduler Steps	250, 350	350, 450
L2 Weight Decay	1×10^{-4}	1×10^{-4}
Sigmoid Temperature (τ)	0.01	0.01
Positives per Query (k)	4	4
Quantization Scale (qs)	0.01	0.01

of 2048 has been used with 400 and 500 training epochs for the baseline and refined protocol, respectively. The optimizer used to minimize the loss function is Adam with an Initial Learning Rate of 1e-3 and it is divided by 10 in the epochs given by the LR scheduler steps, which are epochs 250 and 350 for baseline protocol and epochs 350 and 450 for refined protocol. Table II summarizes all the parameter values described above.

Additionally, when working with sparse convolutions, the input point clouds need to be quantized by a factor of qs, which is set to 0.01. Since these clouds are already normalized to [-1, 1], spatial resolutions of 200 voxels are obtained on each coordinate axis. To increase the number of training instances and reduce model overfitting, a data augmentation has been carried out by applying a random jitter of a value between [0, 0.001] individually to each point of the point cloud, a random transformation to the global point cloud with a value between [0, 0.01] and a random removal of 10% of the points.

All experiments are carried out on a NVIDIA GeForce RTX 3090 GPU with 24 GB. Our code is publicly available on the project website https://github.com/juanjocabrera/MinkUNeXt.git.

E. Ablation study: From MinkUNet to MinkUNeXt

The design departs from the MinkUNet34C architecture [16] as a baseline. Next, the series of design decisions are described. Each design step is summarized in two main subsections: (1) global design and (2) residual block design, which are included next. For every step, both the procedure and the results are presented, starting from the MinkUNet34C until obtaining the MinkUNeXt architecture. The evolution of the network and results is presented in Fig. 4. Table III summarizes and describes the main design steps.

1) Global Design: As mentioned above, the starting point is the MinkUnet34C [16] architecture and it is first modified by adding a GeM pool layer. This step is marked as G1.1 in Fig. 4. The rest of the roadmap followed towards the final design is described next. Each of the design steps is classified in one of the following points: evaluating the cardinality, evaluating the number of channels, changing the number of skip connections and changing the stem to "Patchify".



Fig. 4. This diagram illustrates the design progress of the proposed architecture from MinkUNet up to MinkUNeXt. All the proposed modifications are summarized in Table III.

- **G1.** Evaluating the cardinality. The cardinality is defined as the number of parallel blocks, that enables the network to learn various input representations. In this sense, different cardinality configurations per residual block are assessed: (2, 3, 4, 6, 2, 2, 2, 2), (2, 2, 2, 2, 2, 2, 2, 2) and (1, 1, 1, 1, 1, 1, 1), corresponding to MinkUNet34, MinkUNet18 and MinkUNet14, respectively. These cardinality values represent the number of instances of each Residual Block that appear in blue colour in Fig. 2, but at this point still with ResNet Blocks. In addition, these cardinality configurations are summarized respectively in steps G1.1, G1.2 and G1.3 in Fig. 4. As ilustrated in the diagram, reducing the cardinality to the minimum, with no parallel blocks, shows a better performance and allows an improvement from 92.2% to 92.8% in terms of average recall at 1 (AR@1). From now on, it will be used 1 as cardinality of each residual block.
- **G2.** Evaluating the number of channels. The number of channels or filters correspond to the number of feature maps that the convolutional layer can learn. The number of filters corresponding to the convolutional layers of the encoder are fixed to (32, 64, 128, 256), but the number of channels of the decoder takes the following values (128, 128, 96, 96), (128, 128, 128, 128) and (192, 192, 128, 128) corresponding to MinkUNet14A, MinkUNet14B and MinkUNet14C. This number of filters of the decoder is summarized respectively in steps G1.3, G2.1 and G2.2 in Fig. 4. The best result is obtained with MinkUnet14C (G2.2) with an AR@1 of 93.3%. Thus, the number of

filters of the transposed convolutions that will be adopted in the subsequent architecture variations is (192, 192, 128, 128).

- **G3.** Changing the number of skip connections. The original U-Net is characterized by the presence of 4 skip connections between the encoder and the decoder networks. In this sense, in the present paper the performance of the network is studied when reducing the number of skip connections and removing the transpose convolutions after the last connection. Apart from the 4 skip connections already implemented in the above configurations, we have evaluated 2 and 3 skip connections corresponding to G3.1 and G3.2 in Fig. 4. By reducing the number of skip connections to 3 and removing the layers after the last connection, the model shows, by far, the greatest improvement on the AR@1, increasing the results from 93.3% to 96.3%. As a result, only 3 skip connections will be included between the encoder and decoder.
- **G4.** Changing the stem to "Patchify". The stem refers to the first layer in the network, which performs the initial processing. In this case, the first processing is carried out by a 3D Sparse Convolution with kernel size 5 and stride 1. The term "Patchify" refers to the act of splitting the input data into a sequence of patches. Visual Transformers [47] introduced this concept, originally inspired by NLP Transformers [23]. The Swin Transformer [48] uses as stem a non-overlapping convolution with kernel size 4 and stride 4. In this sense, these parameters are adopted

for the stem in G4, but the performance of the network has decreased from 96.3% to 92.8% so "Patchify" is discarded.

2) Residual Block Design: This section describes each design step from ResNet Block to the proposed MinkNeXt Block. The roadmap of the design of this residual block is divided in the following points: creating a Bottleneck in the residual block, creating an Inverted Bottleneck in the residual block, replacing ReLUs with GeLUs, substituting BN with LN and evaluating different kernel sizes.

- **R1. Creating a Bottleneck in the residual block.** A Bottleneck consists in reducing the dimensionality of the hidden layer and then expanding it to its original size using 1x1 convolutions. This modification led to worse results in the performance of the proposed architecture.
- **R2.** Creating an Inverted Bottleneck in the residual block. Every Transformer block is characterized by an inverted bottleneck, which consists in expanding the dimensionality of the feature map of the hidden layer and then reducing it to its original size by 1x1 convolutions. In this case, 3D sparse convolutions with kernel size 3 and stride 1 are employed to create the inverted bottleneck with a hidden dimension four times wider than the input dimension. Fig. 4 shows that this inverted bottleneck block produces better results compared to the previous ResNet block when analysed jointly with the following modification (R3).
- **R3.** Replacing ReLUs with GeLUs. The Rectified Linear Unit (ReLU) [49] is the most employed activation function over time due to its simplicity and efficiency. However, recent advanced Transformers such as Google's BERT [50] or OpenAI's GPT-4 [51] employ Gaussian Error Linear Units (GeLUs) [44], which is a smoother variant of ReLUs. Following the same philosophy, Re-LUs are replaced with GeLUs in both the ResNet Block and the inverted bottleneck block, steps R3.1 and R3.2 in Fig. 4, respectively. In both cases, the performance of the architecture improves, but better results are obtained with the proposed inverted bottleneck block, achieving an AR@1 of 97.0%. In consequence, an inverted bottleneck with GeLUs will be used as residual block.
- **R4.** Substituting BN with LN. Batch Normalization (BN) [46] plays a critical role in convolutional networks by enhancing convergence and mitigating overfitting. However, BN may introduce complexities that may negatively impact the model's performance. Recently, the simpler Layer Normalization [45] (LN) has been successfully implemented in Transformers. Thus, BN is replaced with LN in the proposed residual block, obtaining an improvement of the model performance up to 97.4%. As a result, Layer Normalization will be

employed instead of Batch Normalization in the residual block.

R5. Evaluating different kernel sizes. Vision Transformers are characterized by employing large kernel sizes with a minimum dimension of 7. However, as shown in Fig. 4 (R5), the usage of smaller kernel sizes is beneficial in the present place recognition task, both in the input, hidden and last layers of the residual block. In this sense, we find the best parameter configuration with a kernel size of 1 in the first convolution and kernel sizes of 3 in the hidden and last convolutions. This leads to the final model and residual block architectures, which we have named MinkUNeXt and MinkNeXt block, respectively.

TABLE III This table summarizes all modifications proposed in the architecture design progress from MinkUNet up to MinkUNeXt.

ID	Design modifications
G1.1	Cardinality: $(2, 3, 4, 6, 2, 2, 2, 2) \rightarrow (2, 2, 2, 2, 2, 2, 2, 2)$
G1.2	Cardinality: $(2, 2, 2, 2, 2, 2, 2, 2, 2) \rightarrow (1, 1, 1, 1, 1, 1, 1, 1)$
G2.1	Decoder channels: $(128, 128, 96, 96) \rightarrow (128, 128, 128, 128)$
G2.2	Decoder channels: $(128, 128, 96, 96) \rightarrow (192, 192, 128, 128)$
G3.1	4 skip connections \rightarrow 2 skip connections
G3.2	4 skip connections \rightarrow 3 skip connections
G4	Stem (k=5, s=1 \rightarrow k=4, s=4)
R1	ResNet Block \rightarrow Bottleneck
R2	ResNet Block \rightarrow Inv. Bottleneck
R3.1	ResNet Block with ReLUs \rightarrow ResNet Block with GeLUs
R3.2	Inv. Bottleneck with ReLUs \rightarrow Inv. Bottleneck with GeLUs
R4	Inv. Bottleneck with BNs \rightarrow Inv. Bottleneck with LNs
R5.1	Inv. Bottleneck 1st convolution (k=3 \rightarrow k=5)
R5.2	Inv. Bottleneck 1st convolution (k=3 \rightarrow k=7)
R5.3	Inv. Bottleneck 1st convolution (k=3 \rightarrow k=1)

F. Comparison with the state of the art

As defined in Subsection IV-C, the two training and evaluation protocols established in [10] have been followed for place recognition with the Oxford RobotCar and In-house datasets. The baseline protocol consists in training the model only with the Oxford training data and evaluating with the Oxford and In-house (U.S., R.A. and B.D.) test data. In contrast, the refined protocol consists in training with Oxford and In-house (U.S., R.A.) training data and evaluating with the Oxford and In-house (U.S., R.A. and B.D.) test data. These protocols are widely used in the literature, so that the comparison is performed on the same terms and conditions. Additionally, the comparative results shown here have been obtained from the same works that are referenced.

Tables IV and V present an overview of the results with different techniques proposed in the state of the art compared to the one proposed in this paper under the same training and evaluation protocols (baseline and refined), in terms of average recall at 1 (AR@1) and average recall at 1% (AR@1%). Each column presents the results obtained with each of the datasets, whereas the last two columns present the mean results.

1) Results with the Baseline Protocol: Table IV presents the results of several methods in terms of average recall at 1 (AR@1) and average recall at 1% (AR@1%). It can be observed that, PointNetVLAD established the starting point for place-recognition from point clouds with the Oxford Robotcar and the In-house dataset. PCAN slightly outperforms PointNetVLAD on most datasets. BPT stands out with really competitive results, especially in Oxford and U.S. RPR-Net outperforms BPT in U.S, R.A and B.D., showing better generalization capabilities. Some works, such as DAGC and Retriever, do not provide AR@1 results for all datasets. However, they presented AR@1% results which show a performance better than PCAN, but worse than BPT. Futhermore, LPD-Net, HiTPR, EPC-Net and E²PN-GeM show similar, but good results across multiple scenarios. SOE-Net, only provides AR@1% results which are really promising as they are close to MinkLoc3D, the first architecture to exceed 90% in AR@1 with the Oxford dataset. Moreover, HiBi-Net, PPT-Net and SVT-Net show slightly higher performance, specifically for the In-house dataset. TransLoc3D takes a step forward with the best result so far in Oxford and solid performance in the other scenarios, and its improved version MinkLoc3Dv2 outperforms the rest of the architectures. In addition, KPPR also shows a remarkable performance, but only presented average recall at 1% results in the case of U.S., R.A., B.D.

Finally, the proposed architecture, MinkUNeXt, demonstrates superior performance in terms of AR@1 and AR@1% on Oxford. It outperforms all of the existing methods with a 97.5% in AR@1 and 99.3% in AR@1%. However, the performance slightly decreases when the model is tested in U.S., R.A. and B.D. It should be highlighted that the Oxford dataset and the three in-house datasets were obtained using LiDARs that exhibit different characteristics, such as number of channels or spatial resolution. The Oxford dataset is captured with various SICK LMS-151 2D and the Inhouse dataset with a 64 channel Velodyne. Moreover, the submaps within the Oxford dataset contain scenes that are entirely urban, characterized by densely built environments with a more compact structure. In contrast, the scenarios present in the in-house dataset are considerably more open, with fewer obstructions and a more dispersed arrangement of urban elements. This difference in the nature of the captured scenes can significantly influence the results and performance of the model on each dataset.

2) Results with the Refined Protocol: As for the performances of the models when training with the refined protocol (Table V), PointNetVLAD introduced the starting reference point as well, surprisingly achieving a good performance in U.S. R.A. and B.D. despite the simplicity of the network architecture. PCAN and DAGC presented similar results to Point-NetVLAD for the In-house dataset, but especially better in Oxford. In contrast, LPD-Net and SOE-Net show substantially better performance in all metrics and datasets. MinkLoc3D also manages to exceed 90% on average recall at 1 (AR@1) in Oxford and generally performs well in all metrics and sets. PPT-Net does not provide values for average recall at 1 (AR@1), but shows a promising performance on average recall at 1% (AR@1%). Furthermore, SVT-Net stands out especially in U.S., R.A. and B.D. In addition, TransLoc3D achieves good results in all metrics, being one of the best methods overall. MinkLoc3Dv2 boasted the best results in the state of the art so far, showing improvements over MinkLoc3D.

Finally, the proposed MinkUNeXt model shows considerable improvements in average recall at 1 (AR@1) and average recall at 1% (AR@1%) for all the scenarios obtaining the best results of the state of the art so far. The average recall at 1 (AR@1) metric on Oxford dataset is 97.7% and outperforms the runner-up MinkLoc3Dv2 by 0.8 p.p. On the R.A., B.D. scenarios it surpasses MinkLoc3Dv2 by 0.1 to 1.1 p.p. Nevertheless, slightly worse results (0.3 p.p.) are obtained with this metric in the U.S. dataset. Regarding the results in terms of AR@1% for the refined protocol, there was little room for improvement. However, the results on Oxford are improved by 0.2 p.p. to reach 99.3%, on R.A. by 0.5 p.p. to reach 99.9% and on B.D. by 0.1 p.p. to reach 97.7%. In addition, although the model previously output slightly worse results for U.S. in terms of AR@1, the performance of the network in the AR@1% metric is equal to the best previous result in the state-of-the-art with a value of 99.9%. The mean AR@1 and AR@1% over all 4 datasets improves by 0.4% and 0.2%, respectively. To conclude, training the MinkUNeXt with the refined protocol overcomes the generalization difficulties presented when training with the baseline protocol, since the model adapts to both LiDAR characteristics.

V. CONCLUSION

This paper presents MinkUNeXt, an architecture based on MinkUNet [16] exhaustively modified and enhanced to perform place-recognition based on point clouds. It is an encoder-decoder architecture entirely based on the proposed 3D MinkNeXt Block: a residual block composed of 3D sparse convolutions that follows the philosophy proposed by ConvNeXt [12]. The feature extraction step is performed by a U-Net encoder-decoder. The feature aggregation of those features into a single descriptor is carried out by a Generalized Mean Pooling (GeM) [40]. The designed architecture demonstrates that it is possible to surpass the current state of the art by only relying on conventional 3D sparse convolutions without making use of more complex and sophisticated proposals such as Transformers, Attention-Layers or Deformable Convolutions.

The proposed network shows that the usage of a U-Net architecture for point cloud-based place recognition is beneficial, since it is able to capture both detailed and contextual information of the three-dimensional environment. The fusion of features from multiple spatial scales improves the robustness of the place recognition model, allowing it to adapt to variations in point cloud geometry and density, as well as to different scenarios.

TABLE IV
EVALUATION RESULTS IN TERMS OF AVERAGE RECALL AT 1 (AR@1) AND AT 1% (AR@1%) OF PLACE RECOGNITION METHODS
TRAINED USING THE BASELINE PROTOCOL

	Oxford		U.S.		R.A.		B.D.		Mean	
Method	AR@1	AR@1%	AR@1	AR@1%	AR@1	AR@1%	AR@1	AR@1%	AR@1	AR@1%
PointNetVLAD [10]	62.8	80.3	63.2	72.6	56.1	60.3	57.2	65.3	59.8	69.6
PCAN [52]	69.1	83.8	62.4	79.1	56.9	71.2	58.1	66.8	61.6	75.2
DAGC [22]	-	87.5	-	83.5	-	75.7	-	71.2	-	79.5
BPT [39]	85.7	93.3	80.5	89.3	77.4	86.6	74.1	78.5	79.4	86.9
Retriever [27]	-	91.9	-	91.9	-	87.4	-	85.5	-	89.2
RPR-Net [37]	81.0	92.2	83.2	94.5	83.3	91.3	80.4	86.4	82.0	91.1
LPD-Net [14]	86.3	94.9	87.0	96.0	83.1	90.5	82.5	89.1	84.7	92.6
HiTPR [29]	87.8	94.6	86.0	94.0	81.3	89.1	81.8	88.3	84.2	91.5
EPC-Net [38]	86.2	94.7	-	96.5	-	88.6	-	84.9	-	91.2
E ² PN-GeM [36]	84.8	93.2	88.1	95.3	83.7	90.5	83.3	87.7	85.0	91.7
SOE-Net [26]	-	96.4	-	93.2	-	91.5	-	88.5	-	92.4
MinkLoc3D [15]	93.0	97.9	86.7	95.0	80.4	91.2	81.5	88.5	85.4	93.2
HiBi-Net [53]	87.5	95.1	87.8	-	85.8	-	83.0	-	86.0	-
NDT-Transformer [24]	93.8	97.7	-	-	-	-	-	-	-	-
PPT-Net [25]	93.5	98.1	90.1	97.5	84.1	93.3	84.6	90.0	88.1	94.7
SVT-Net [28]	93.7	97.8	90.1	96.5	84.3	92.7	85.5	90.7	88.4	94.4
TransLoc3D [31]	95.0	98.5	-	94.9	-	91.5	-	88.4	-	93.3
MinkLoc3Dv2 [5]	96.3	98.9	90.9	96.7	86.5	93.8	86.3	91.2	90.0	95.1
KPPR [34]	91.5	97.1	-	98.0	-	95.1	-	92.1	-	95.6
MinkUNeXt (ours)	97.5	99.3	88.9	96.5	85.0	91.3	85.2	90.1	89.1	94.3

TABLE V EVALUATION RESULTS IN TERMS OF AVERAGE RECALL AT 1 (AR@1) AND AT 1% (AR@1%) OF PLACE RECOGNITION METHODS TRAINED USING THE REFINED PROTOCOL.

	Oxford		U.S.		R.A.		B.D.		Mean	
Method	AR@1	AR@1%	AR@1	AR@1%	AR@1	AR@1%	AR@1	AR@1%	AR@1	AR@1%
PointNetVLAD [10]	63.3	80.1	86.1	94.5	82.7	93.1	80.1	86.5	78.0	88.6
PCAN [52]	70.7	86.4	83.7	94.1	82.5	92.5	80.3	87.0	79.3	90.0
DAGC [22]	71.5	87.8	86.3	94.3	82.8	93.4	81.3	88.5	80.5	91.0
LPD-Net [14]	86.6	94.9	94.4	98.9	90.8	96.4	90.8	94.4	90.7	96.2
SOE-Net [26]	89.3	96.4	91.8	97.7	90.2	95.9	89.0	92.6	90.1	95.7
MinkLoc3D [15]	94.8	98.5	97.2	99.7	96.7	99.3	94.0	96.7	95.7	98.6
PPT-Net [25]	-	98.4	-	99.7	-	99.5	-	95.3	-	98.2
SVT-Net [28]	94.7	98.4	97.0	99.9	95.2	99.5	94.4	97.2	95.3	98.8
TransLoc3D [31]	95.0	98.5	97.5	99.8	97.3	99.7	94.8	97.4	96.2	98.9
MinkLoc3Dv2 [5]	96.9	99.1	99.0	99.7	98.3	99.4	97.6	99.1	97.9	99.3
MinkUNeXt (ours)	97.7	99.3	98.7	99.9	99.4	99.9	97.7	99.0	98.3	99.5

It should be also noted that the proposed method outputs results outperforming an already saturated state-of-the-art. In particular, the network achieved an AR@1 of 97.5% and an AR@1% of 99.3% when trained with the refined protocol. Thus, there is little room for improvement and larger and more diverse scenarios are needed in order to stimulate further progress.

Future work will consider the inclusion of visual information into the place recognition system. In this sense, we consider that it would result in a richer representation of the environment compared to the use of LiDAR with pure distance data. However, visual information is hindered by changing lighting conditions, weather and seasonal changes, which pose a great challenge.

ACKNOWLEDGMENTS

The Ministry of Science, Innovation and Universities (Spain) has supported this work through "Ayudas para la

Formación de Profesorado Universitario" (FPU21/04969). This work is also part of the project TED2021-130901B-100, funded by MCIN/AEI/10.13039501100011033 and the European Union "NextGenerationEU"/PRTR, and of the project PROMETEO/2021/075 funded by Generalitat Valenciana (Spain). In addition, this research was sponsored by national funds (Portugal) through FCT - Fundação para a Ciência e a Tecnologia, under project LA/P/0079/2020, DOI: 10.54499/LA/P/0079/2020. This work was further funded in the scope of the projects E-Forest—Multi-agent Autonomous Electric Robotic Forest Management Framework, ref. POCI-01-0247-FEDER-047104 and F4F—Forest for Future, ref. CENTRO-08-5864-FSE-000031, co-financed by European Funds through the programs Compete 2020 and Portugal 2020.

REFERENCES

- H. Yin, L. Tang, X. Ding, Y. Wang, and R. Xiong, "LocNet: Global localization in 3D point clouds for mobile vehicles," in 2018 IEEE Intelligent Vehicles Symposium (IV). IEEE, 2018, pp. 728–733.
- [2] 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.
- [3] M. Leyva-Vallina, N. Strisciuglio, and N. Petkov, "Place recognition in gardens by learning visual representations: data set and benchmark analysis," in *International Conference on Computer Analysis of Images* and Patterns. Springer, 2019, pp. 324–335.
- [4] M. Himstedt, J. Frost, S. Hellbach, H.-J. Böhme, and E. Maehle, "Large scale place recognition in 2D LiDAR scans using geometrical landmark relations," in 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE, 2014, pp. 5030–5035.
- [5] J. Komorowski, "Improving point cloud based place recognition with ranking-based loss and large batch training," in 2022 26th International Conference on Pattern Recognition (ICPR). IEEE, 2022, pp. 3699– 3705.
- [6] J. Komorowski, M. Wysoczanska, and T. Trzcinski, "Large-scale topological radar localization using learned descriptors," in *Neural Information Processing: 28th International Conference, ICONIP 2021, Sanur, Bali, Indonesia, December 8–12, 2021, Proceedings, Part II 28.* Springer, 2021, pp. 451–462.
- [7] X. Tang, W. Fu, M. Jiang, G. Peng, Z. Wu, Y. Yue, and D. Wang, "Place recognition using line-junction-lines in urban environments," in 2019 IEEE International Conference on Cybernetics and Intelligent Systems (CIS) and IEEE Conference on Robotics, Automation and Mechatronics (RAM), 2019, pp. 530–535.
- [8] S. Arshad and G.-W. Kim, "A robust feature matching strategy for fast and effective visual place recognition in challenging environmental conditions," *International Journal of Control, Automation and Systems*, vol. 21, no. 3, pp. 948–962, 2023.
- [9] Y. Wang, Z. Sun, C.-Z. Xu, S. E. Sarma, J. Yang, and H. Kong, "LiDAR iris for loop-closure detection," in 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2020, pp. 5769–5775.
- [10] M. A. Uy and G. H. Lee, "PointNetVLAD: Deep point cloud based retrieval for large-scale place recognition," in *Proceedings of the IEEE* conference on computer vision and pattern recognition, 2018, pp. 4470– 4479.
- [11] X. Chen, T. Läbe, A. Milioto, T. Röhling, J. Behley, and C. Stachniss, "Overlapnet: A siamese network for computing LiDAR scan similarity with applications to loop closing and localization," *Autonomous Robots*, pp. 1–21, 2022.
- [12] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, "A convnet for the 2020s," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 11976–11986.
- [13] 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.
- [14] Z. Liu, S. Zhou, C. Suo, P. Yin, W. Chen, H. Wang, H. Li, and Y.-H. Liu, "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.
- [15] J. Komorowski, "Minkloc3D: Point cloud based large-scale place recognition," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021, pp. 1790–1799.
- [16] C. Choy, J. Gwak, and S. Savarese, "4D spatio-temporal convnets: Minkowski convolutional neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 3075–3084.
- [17] F. Radenović, G. Tolias, and O. Chum, "Fine-tuning CNN image retrieval with no human annotation," *IEEE transactions on pattern* analysis and machine intelligence, vol. 41, no. 7, pp. 1655–1668, 2018.
- [18] J. Komorowski, M. Wysoczańska, and T. Trzcinski, "MinkLoc++: LiDAR and monocular image fusion for place recognition," in 2021 International Joint Conference on Neural Networks (IJCNN). IEEE, 2021, pp. 1–8.
- [19] Y. Lu, F. Yang, F. Chen, and D. Xie, "Pic-net: Point cloud and image collaboration network for large-scale place recognition," *arXiv preprint* arXiv:2008.00658, 2020.

- [20] H. Song, W. Choi, and H. Kim, "Robust vision-based relativelocalization approach using an RGB-depth camera and LiDAR sensor fusion," *IEEE Transactions on Industrial Electronics*, vol. 63, no. 6, pp. 3725–3736, 2016.
- [21] 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.
- [22] Q. Sun, H. Liu, J. He, Z. Fan, and X. Du, "Dagc: Employing dual attention and graph convolution for point cloud based place recognition," in *Proceedings of the 2020 International Conference on Multimedia Retrieval*, 2020, pp. 224–232.
- [23] 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.
- [24] Z. Zhou, C. Zhao, D. Adolfsson, 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.
- [25] L. Hui, H. Yang, M. Cheng, J. Xie, and J. Yang, "Pyramid point cloud transformer for large-scale place recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6098–6107.
- [26] Y. Xia, Y. Xu, S. Li, R. Wang, J. Du, D. Cremers, and U. Stilla, "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. 11348– 11357.
- [27] L. Wiesmann, R. Marcuzzi, C. Stachniss, and J. Behley, "Retriever: Point cloud retrieval in compressed 3D maps," in 2022 International Conference on Robotics and Automation (ICRA). IEEE, 2022, pp. 10925–10932.
- [28] 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.
- [29] Z. Hou, Y. Yan, C. Xu, and H. Kong, in 2022 International Conference on Robotics and Automation (ICRA). IEEE, 2022, pp. 2612–2618.
- [30] 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.
- [31] T.-X. Xu, Y.-C. Guo, Z. Li, G. Yu, Y.-K. Lai, and S.-H. Zhang, "TransLoc3D: Point cloud based large-scale place recognition using adaptive receptive fields," arXiv preprint arXiv:2105.11605, 2021.
- [32] Q. Wang, B. Wu, P. Zhu, P. Li, W. Zuo, and Q. Hu, "ECA-Net: Efficient channel attention for deep convolutional neural networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11534–11542.
- [33] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The Oxford robotcar dataset," *The International Journal of Robotics Research*, vol. 36, no. 1, pp. 3–15, 2017.
- [34] L. Wiesmann, L. Nunes, J. Behley, and C. Stachniss, "KPPR: Exploiting momentum contrast for point cloud-based place recognition," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 592–599, 2022.
- [35] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. J. Guibas, "Kpconv: Flexible and deformable convolution for point clouds," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 6411–6420.
- [36] C. E. Lin, J. Song, R. Zhang, M. Zhu, and M. Ghaffari, "Se (3)equivariant point cloud-based place recognition," in *Conference on Robot Learning*. PMLR, 2023, pp. 1520–1530.
- [37] Z. Fan, Z. Song, W. Zhang, H. Liu, J. He, and X. Du, "RPR-Net: A point cloud-based rotation-aware large scale place recognition network," in *European Conference on Computer Vision*. Springer, 2022, pp. 709– 725.
- [38] 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.
- [39] Z. Hou, Y. Shang, T. Gao, and Y. Yan, "BPT: binary point cloud transformer for place recognition," arXiv preprint arXiv:2303.01166, 2023.
- [40] F. Radenović, G. Tolias, and O. Chum, "Fine-tuning cnn image retrieval with no human annotation," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 7, pp. 1655–1668, 2018.
- [41] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing*

and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18. Springer, 2015, pp. 234–241.

- [42] N. Sünderhauf, F. Dayoub, S. McMahon, B. Talbot, R. Schulz, P. Corke, G. Wyeth, B. Upcroft, and M. Milford, "Place categorization and semantic mapping on a mobile robot," in 2016 IEEE international conference on robotics and automation (ICRA). IEEE, 2016, pp. 5729– 5736.
- [43] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *Proceedings* of the IEEE conference on computer vision and pattern recognition, 2018, pp. 4510–4520.
- [44] D. Hendrycks and K. Gimpel, "Gaussian error linear units (gelus)," arXiv preprint arXiv:1606.08415, 2016.
- [45] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," arXiv preprint arXiv:1607.06450, 2016.
- [46] S. Ioffe, "Batch renormalization: Towards reducing minibatch dependence in batch-normalized models," Advances in neural information processing systems, vol. 30, 2017.
- [47] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [48] 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," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10012–10022.
- [49] V. Nair and G. E. Hinton, "Rectified linear units improve restricted boltzmann machines," in *Proceedings of the 27th international conference on machine learning (ICML-10)*, 2010, pp. 807–814.
- [50] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," arXiv preprint arXiv:1810.04805, 2018.
- [51] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [52] W. Zhang and C. Xiao, "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. 12436–12445.
- [53] D. W. Shu and J. Kwon, "Hierarchical bidirected graph convolutions for large-scale 3D point cloud place recognition," *IEEE Transactions on Neural Networks and Learning Systems*, 2023.