

RecNet: An Invertible Point Cloud Encoding through Range Image Embeddings for Multi-Robot Map Sharing and Reconstruction

Nikolaos Stathoulopoulos, Mario A.V. Saucedo, Anton Koval and George Nikolakopoulos

Abstract—In the field of resource-constrained robots and the need for effective place recognition in multi-robotic systems, this article introduces RecNet, a novel approach that concurrently addresses both challenges. The core of RecNet’s methodology involves a transformative process: it projects 3D point clouds into range images, compresses them using an encoder-decoder framework, and subsequently reconstructs the range image, restoring the original point cloud. Additionally, RecNet utilizes the latent vector extracted from this process for efficient place recognition tasks. This approach not only achieves comparable place recognition results but also maintains a compact representation, suitable for sharing among robots to reconstruct their collective maps. The evaluation of RecNet encompasses an array of metrics, including place recognition performance, the structural similarity of the reconstructed point clouds, and the bandwidth transmission advantages, derived from sharing only the latent vectors. Our proposed approach is assessed using both a publicly available dataset and field experiments¹, confirming its efficacy and potential for real-world applications.

I. INTRODUCTION

Place recognition plays a fundamental role in both single and multi-robot Simultaneous Localization and Mapping (SLAM) systems, as it allows robots to identify previously visited locations, while mitigating the effects of drift. In these systems, the ability of robots to recognize previously visited locations, allows them to anchor their position estimates and correct accumulated errors over time [1]. Image-based place recognition has been widely explored [2], but it suffers from certain limitations, such as sensitivity to lighting conditions and occlusions, which can restrict its reliability in challenging real-world environments. In contrast, LiDAR-based place recognition has gained popularity thanks to its ability to provide robust and precise 3D representations of the environment. Despite the advantages of LiDAR-based methods, they come with their own set of challenges. Direct computations in the 3D point cloud data can be computationally intensive, limiting real-time performance on resource-constrained robotic platforms. Additionally, in multi-robot systems, a major challenge lies in efficiently sharing individual maps and point cloud representations among agents. Centralized communication approaches suffer from limitations in scalability and are prone to single points of failure, hindering the potential benefits of collaborative multi-robot mapping

The Authors are with the Robotics and AI Group, Department of Computer, Electrical and Space Engineering, Luleå University of Technology, 971 87 Luleå, Sweden. Corresponding Author’s Email: niksta@ltu.se

This work has been partially funded by the European Union’s Horizon 2020 Research and Innovation Programme, under the Grant Agreement No. 101003591 NEX-GEN SIMS, and partially by the Sustainable Underground Mining (SUM) Academy Programme, under project SP14.

¹<https://www.youtube.com/watch?v=f9BnK34XkuQ>

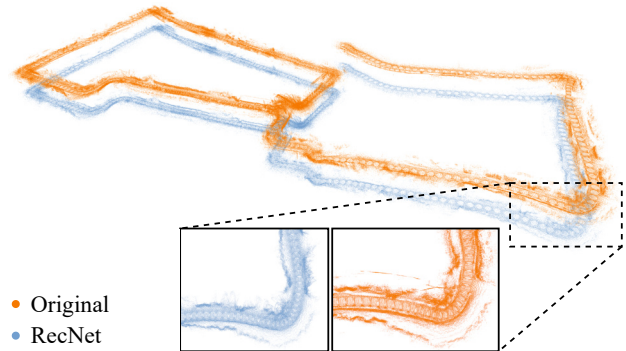


Fig. 1: The original and the reconstructed map using RecNet on an experiment around Luleå University of Technology.

and map-merging [3]. Consequently, the need arises for a more decentralized and efficient communication mechanism that allows agents to share their point cloud representations directly and selectively. With multiple robots operating over extended periods of time, the accumulation of vast amounts of point cloud data becomes inevitable [4]. This accumulation not only strains the onboard computational resources of the robots but also leads to significant communication overhead, when sharing maps among agents for collaborative mapping [5]. Additionally, as the number of robots in a multi-robot system grows, the complexity of map-merging and coordination escalates exponentially, demanding more sophisticated and scalable solutions.

In response to the challenges posed by resource-constrained robots and the need for effective place recognition in multi-robot systems, this article presents a novel approach that tackles both aspects simultaneously. Our method integrates LiDAR-based place recognition with lightweight descriptors, enabling efficient point cloud representation and sharing with other agents. By transforming 3D point clouds into compact range images and then lightweight descriptors, we significantly reduce computational burden and communication overhead. The central focus of this approach is to facilitate seamless data sharing [6] and to promote multi-robot collaboration in exploration missions [7]. This not only enhances the capabilities of individual robots but also empowers the entire multi-robot system to collectively build and maintain a comprehensive map. The efficiency of our method is crucial in resource-constrained scenarios, where computational and communication limitations can hinder the effectiveness of traditional approaches [8].

The contributions of the proposed framework can be summarized as follows: (a) Introduction of a novel descriptor pipeline that not only matches the performance

of state-of-the-art approaches like those discussed in [9] and [10], but also distinguishes itself by offering a reverse conversion to the original point cloud representation. This unique feature sets it apart from conventional methods; (b) Unlike the majority of existing approaches, RecNet provides point cloud reconstruction, while simultaneously achieves faster descriptor extraction when compared to [11], a feature that opens up possibilities for real-time global localization. Furthermore, the reconstruction quality was evaluated using publicly available structural similarity metrics [12], on a well-established dataset. (c) The descriptors generated by our framework provide a concise representation, reducing data transmission load by over 10 times. Additionally, the reconstruction process achieves a geometrical similarity of approximately 70%. The proposed framework was evaluated in both the KITTI dataset [13] and a real-world outdoor experiment. These evaluations showcase the robustness and versatility of our approach across various environments and scenarios, offering compelling evidence of its effectiveness in addressing real-world challenges in point cloud recognition.

II. RELATED WORK

In this section, we first present point cloud compression methods, followed by a discussion on place recognition techniques, with an emphasis on learning-based approaches for 3D point clouds. A comprehensive review of related literature on global localization in 3D point clouds is available in [14].

Point cloud compression: Several approaches have been proposed for compressing LiDAR point cloud data. The RIDDLE compression pipeline [15], employs quantization followed by deep delta encoding and entropy encoding to achieve efficient compression of range image data. Feng et al. [16] present a system that exploits spatial and temporal redundancies in LiDAR point cloud sequences. By identifying key frames and utilizing overlaps between consecutive frames, their approach achieves compression ratios ranging from $\times 40$ -90 while maintaining high accuracy and real-time processing. Sun et al. [17] propose a compression algorithm that segments 3D range data into ground and main objects, leveraging shape-based prediction methods inspired by 3D video coding techniques. By efficiently compressing prediction residuals, their method achieves compression ratios of nearly 5% without distance distortion. Wang et al. [18] propose a three-stage framework focusing on utilizing entropy modeling and a novel attention convolutional layer for feature extraction in scanning LiDAR point clouds. Their framework, employing RICNet_{stage1} and RICNet_{stage2} for compression and decompression, demonstrates promising results in terms of compression ratio and reconstruction quality, especially in diverse LiDAR line scenarios.

Place recognition: The most common step in 3D point cloud place recognition is to extract discriminative features. PointNetVLAD [2] introduced an end-to-end trainable global descriptor by extracting local features leveraging PointNet [19] and then using the NetVLAD aggregator [20], leading to global descriptors for LiDAR-based place recognition. However, the limited descriptive power of PointNet

prompted the proposal of LPD-Net [21] as an alternative approach. More recently, MinkLoc3D [22] demonstrated superior efficiency and performance compared to previous methods by employing sparse convolutions to effectively capture point-level features. On the other hand, LoGG3D-Net [23] improved global descriptors with a local consistency loss, ensuring feature consistency within point clouds. SegMatch [24] and SegMap [11] introduced a high-level perception technique that segments point clouds into distinct and discriminative elements at the object level. To achieve this, they utilize a 3D CNN to encode segment features and identify candidate correspondences through k-nearest neighbours (kNN) in feature space. This approach provides both the ability for place recognition and reconstructing the segments back to point clouds, inspiring the development of our approach. However, segmentation approaches may suffer from limitations in various environments, where complex or dynamic scenes can lead to inaccuracies in segment boundaries, affecting recognition performance.

Recognizing the limitations of both segment-based and direct point cloud computations, the research community emerged with alternative solutions that focus on converting the 3D point cloud representation into the image space. The work in [25] introduced the LORAX algorithm for point cloud registration in localization, involving division into super-points with Random Sphere Cover Set, projection onto 2D depth maps, dimension reduction using a Deep Neural Network Auto-Encoder, and coarse registration with iterative fine-tuning. Nevertheless, the authors acknowledged that their implementation is not optimized for real-time operation, and no timings were provided. To improve computational efficiency, many methods dealing with larger scans opt to convert them into an intermediary representation prior to processing with a deep neural network. The work in [10] presented OREOS, designed for place recognition, which additionally estimates the yaw discrepancy between scans. Their approach involves projecting the input data into a 2D range image using a spherical projection model. Taking the concept of using range images further, OverlapNet [9] takes advantage of multiple cues such as range, normals, intensity, and semantic classes. These cues are projected as spherical images, extracted from the point cloud, thereby enhancing the framework's performance. A more advanced version, OverlapTransformer [26], offers a rotation-invariant representation and faster inference capabilities, although it lacks the ability to estimate yaw angles. Another notable approach, LiDAR Iris [27], encodes the height information of a 3D point cloud into a binary image, transforming it into the Fourier domain to achieve rotation invariance.

III. THE PROPOSED APPROACH

The scope of this article is to develop a novel pipeline for compressing 3D point cloud information with the purpose of: (a) performing place recognition related tasks, like loop closure [1] or map-merging [3], and (b) being able to reconstruct the 3D point cloud to its original form. A critical constraint imposed on this pipeline is the requirement for




Fig. 2: The overall RecNet pipeline is composed of two identical *encoder* legs that share weights, a single *decoder* leg and a *tail* network, responsible for estimating the similarity between the two latent bottleneck vectors.

the compressed data to have a substantially reduced size compared to the original, ensuring efficient transmission between robots, while minimizing communication network congestion. Therefore, given a point cloud $\mathcal{P} \in \mathbb{R}^3$, we are looking for a series of functions that will compress the information into a compact representation $\beta \in \mathbb{R}^N$. The vector β needs to be able to provide enough information in order to recover most of the spatial information from \mathcal{P} into the new reconstructed point cloud $\hat{\mathcal{P}} \in \mathbb{R}^3$. In this article, we assume that the fourth dimension of the original point cloud, the intensity, is not recovered, as we keep the focus on the spatial information. The proposed approach is split into multiple modules, as described in the following subsections.

A. 3D Point cloud to Range Image

We utilize spherical projections of LiDAR scans as our input data. These projections are commonly employed to accelerate computations [9], [26]. Our approach involves transforming the point cloud \mathcal{P} into what is usually called a *vertex map* $\mathcal{V} : \mathbb{R}^2 \rightarrow \mathbb{R}^3$. In this map, each pixel corresponds to the nearest 3D point. To achieve this, every point $p_i = (x, y, z) \in \mathbb{R}^3$ undergoes a conversion process. Initially, it is transformed using the function $\Pi : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ into spherical coordinates and subsequently mapped to image coordinates $I = (u, v) \in \mathbb{R}^2$. This process can be summarized as follows:

$$\begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} \frac{1}{2} [1 - \arctan(yx^{-1})\pi^{-1}] w \\ [1 - (\arcsin(zr^{-1}) + f_{\text{up}})f^{-1}] h \end{pmatrix}, \quad (1)$$

where $r = \|p\|_2$ is the range, $f = f_{\text{up}} + f_{\text{down}}$ corresponds to the vertical field-of-view and w, h denote the width and height dimensions of the resultant vertex map \mathcal{V} .

B. The Encoder-Decoder pair

Following the initial compression of three-dimensional data onto a two-dimensional plane, we apply a subsequent compression step utilizing an *encoder*. This *encoder* takes a range image I as input and generates a multidimensional descriptive vector $\mathcal{E}(I) = \beta \in \mathbb{R}^{256 \times 64}$. Immediately thereafter, the bottleneck β is passed to a *decoder* with the objective of reconstructing the image I . The *decoder* can be conceptualized as the approximate inverse function of the *encoder*, aiming to produce a reconstructed image, $\hat{I} \in \mathbb{R}^2$.

This can be formally defined as $\mathcal{D}(\beta) = \tilde{\mathcal{E}}^{-1}(\beta) = \hat{I}$. The role of β is to encapsulate sufficient descriptive information to facilitate image restoration. As a final step, in line with our earlier objective (a) (performing place recognition tasks), we identify potential candidates, by assessing the similarity between each query vector β_1 , and the corresponding map vectors β_2 . We achieve this by passing the bottlenecks through a smaller Convolutional Neural Network (CNN) to get an estimate on the similarity of the two latent vectors, defined as $\mathcal{T}(\beta_1, \beta_2) = \hat{c} \in [0, 1]$. For a more comprehensive understanding of the network architecture and the loss functions, we provide detailed information in sub-section III-D.

C. Range Image to 3D Point cloud

The final stage of the process involves the conversion of the predicted image \hat{I} back into a three-dimensional representation, denoted as $\hat{\mathcal{P}} \in \mathbb{R}^3$. This conversion can be understood as the approximate inverse mapping of the transformation, detailed in Section III-A, represented by the mapping function $\tilde{\Pi}^{-1} : \mathbb{R}^2 \rightarrow \mathbb{R}^3$. Solving Eq. (1) for x, y , and z given r from the pixel intensity, is straightforward. It is worth noting that the original transformation Π operates from a three-dimensional space to a two-dimensional one, resulting in a loss of information due to the quantization of the image grid and is not a fully invertible and 1-1 mapping. Consequently, it is expected that the resulting point cloud will serve as an approximation of the original point cloud due to this inherent loss of information.

D. The Network Architecture

The proposed network architecture is illustrated in Fig. 2. To optimize the network for lightweight and rapid processing, we exclusively leverage the depth information from the point cloud to generate input range images, which also serve as the target images for our reconstruction task. This choice is motivated by our desire to maintain the network's efficiency, as employing normals images would imply additional computation time for computing neighboring relations. Furthermore, in line with findings from [9], it has been observed that the incorporation of intensity, normal, and semantic images has a negligible impact on the network's performance, except for the yaw estimation aspect, which falls beyond the scope of this article. As a result, the input

for the *encoder* is a tensor with dimensions $1 \times 64 \times 900$, with the channel, height, and width, respectively. In this tensor, the single channel represents the range information, as previously mentioned. RecNet adopts a Siamese network design, comprising two *encoder* legs, a single *tail*, and a single *decoder* leg. Notably, the *encoders* are constructed to be identical, sharing weights across their respective counterparts. Further insights into the trainable parameters are available in Table I.

Encoder: The *encoder* is composed of two legs, all possessing an identical architecture and weight sharing. Each leg represents a fully convolutional network (FCN) featuring ten convolutional layers. We abstain from incorporating max-pooling layers, as they tend to erode vital positional information, a critical element in the image reconstruction process. This architecture maintains a degree of lightness, generating feature volumes sized at $256 \times 1 \times 64$, denoted as β . The inherent translation-equivariance of the FCN permits the future application of yaw discrepancy estimation through a correlation head [9].

Decoder: The *decoder*, featuring a single leg, mirrors the architecture of the *encoder*. However, 2D convolutions are replaced by ten transpose 2D convolutions, facilitating the restoration of the original height and width of the image, while reducing the number of channels. This module operates akin to the gradient of a 2D convolution concerning its input. While commonly referred to as a fractionally-strided convolution or deconvolution, it does not perform a true inverse of convolution. The *decoder's* output directly yields the reconstructed range image in the original size.

Tail: The *tail* is a straightforward CNN comprising three convolutional layers and a fully connected layer. It takes two bottleneck vectors as input, and prior to entering the convolutional layers, these vectors undergo preprocessing through a *delta* layer that involves subtracting one from the other, resulting in a new vector denoted as $\Delta(\beta_1, \beta_2) = \beta_1 - \beta_2 \in \mathbb{R}^{256 \times 64}$. Subsequently, the convolutional layers follow to extract prominent features that aid in distinguishing between the latent vectors. Lastly, the fully connected layer quantifies the similarity between the two vectors, offering a confidence level for their similarity.

E. Loss functions

We adopt an end-to-end training approach to achieve two key objectives: reconstructing the original range image and learning compact representations for tasks related to place recognition. These objectives are governed by two distinct loss terms: the reconstruction loss (L_{rec}) and the place recognition loss (L_{pr}), with a weighting factor α used to combine them effectively. This combination is expressed as:

$$L(I, \hat{I}, \beta_1, \beta_2) = L_{rec}(I, \hat{I}) + \alpha L_{pr}(\beta_1, \beta_2) \quad (2)$$

To begin with, the reconstruction loss (L_{rec}) employs a custom loss function called Weighted Pixel Gradient MSE. This function blends the conventional pixel-wise Mean Squared Error (MSE) loss (L_{MSE}) with a weighted loss component

TABLE I: The layers of the RecNet architecture. Every operator is followed by a PReLU activation function, except the last layers of the *encoder* and *decoder*. A batch normalization is added every other Conv2D and ConvTran2D.

	Operator	Stride	Features	Kernel	Output Shape
Encoder	Conv2D	(2, 2)	16	(5, 15)	$16 \times 30 \times 443$
	Conv2D	(2, 2)	16	(3, 15)	$16 \times 14 \times 215$
	Conv2D	(2, 2)	32	(3, 13)	$32 \times 6 \times 102$
	Conv2D	(2, 1)	32	(2, 12)	$32 \times 3 \times 90$
	Conv2D	(2, 1)	64	(2, 9)	$64 \times 1 \times 82$
	Conv2D	(1, 1)	64	(1, 7)	$64 \times 1 \times 76$
	Conv2D	(1, 1)	128	(1, 5)	$128 \times 1 \times 72$
	Conv2D	(1, 1)	128	(1, 5)	$128 \times 1 \times 68$
	Conv2D	(1, 1)	256	(1, 3)	$256 \times 1 \times 66$
	Conv2D	(1, 1)	256	(1, 3)	$256 \times 1 \times 64$
Decoder	ConvTran2D	(1, 1)	256	(1, 3)	$256 \times 1 \times 66$
	ConvTran2D	(1, 1)	128	(1, 3)	$128 \times 1 \times 68$
	ConvTran2D	(1, 1)	128	(1, 5)	$128 \times 1 \times 70$
	ConvTran2D	(1, 1)	64	(1, 5)	$64 \times 1 \times 74$
	ConvTran2D	(1, 1)	64	(1, 7)	$64 \times 1 \times 82$
	ConvTran2D	(2, 1)	32	(2, 9)	$32 \times 2 \times 90$
	ConvTran2D	(2, 1)	32	(2, 12)	$32 \times 5 \times 102$
	ConvTran2D	(2, 2)	16	(3, 13)	$16 \times 13 \times 215$
	ConvTran2D	(2, 2)	16	(3, 15)	$16 \times 30 \times 443$
	ConvTran2D	(2, 2)	1	(5, 15)	$1 \times 64 \times 900$
Tail	Conv2D	(5, 5)	32	(9, 9)	$32 \times 50 \times 50$
	Conv2D	(3, 3)	64	(5, 5)	$64 \times 16 \times 16$
	Conv2D	(1, 1)	128	(3, 3)	$128 \times 14 \times 14$
	Dense	-	-	-	1

accounting for the gradient differences between corresponding pixels in both images (L_{∂}). The MSE is defined as:

$$L_{MSE}(I, \hat{I}) = \frac{1}{N} \sum_{i=0}^N (I_i - \hat{I}_i)^2 \quad (3)$$

The Gradient MSE loss is a component of the custom loss function designed to measure the difference between the gradients of the reconstructed image and the gradients of the target image. It serves as a way to encourage the model to generate images with similar local edge and texture features as the target image. The Gradient MSE loss is defined as:

$$L_{\partial}(I, \hat{I}) = \frac{1}{N} \left(\sum_{i=0}^N \left| \frac{\partial I_i}{\partial u} - \frac{\partial \hat{I}_i}{\partial u} \right| + \sum_{i=0}^N \left| \frac{\partial I_i}{\partial v} - \frac{\partial \hat{I}_i}{\partial v} \right| \right) \quad (4)$$

Initially, the gradients of both the reconstructed image and the target image are determined. These gradients reflect how pixel values change in both horizontal and vertical directions. Subsequently, the absolute disparities between the gradients of the reconstructed image and the target image are calculated. These calculations are performed separately for the u -direction and the v -direction. The absolute gradient differences are summed across all pixels in the image, and finally the total gradient difference is divided by the total

IV. EXPERIMENTAL EVALUATION

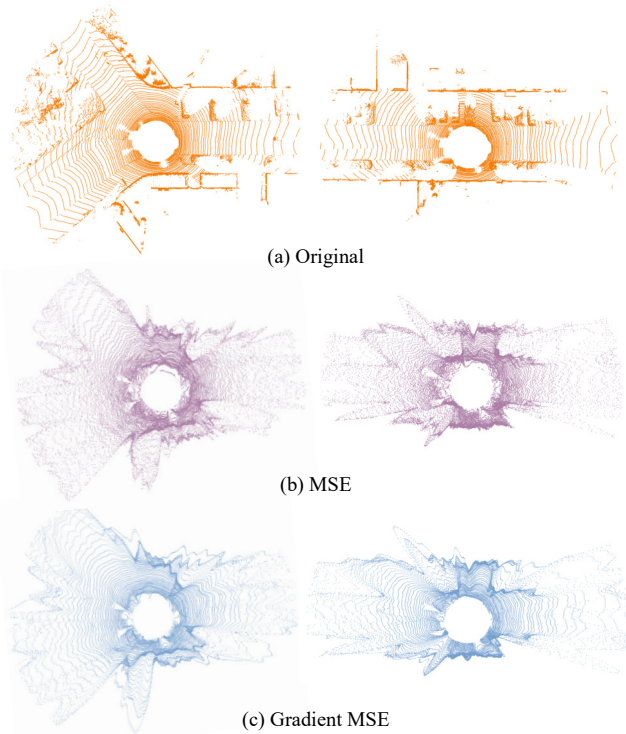


Fig. 3: Two LiDAR scans from the KITTI dataset, where (a) depicts the original scan, (b) is the reconstructed using the MSE and (c) is the reconstructed using the Gradient MSE.

number of pixels (N) to compute the average absolute gradient difference. The gradient difference loss encourages the model to pay attention to fine details and textures in the image, as it penalizes discrepancies in how pixel values change from one location to another compared to the target image. This is particularly useful in image generation tasks where maintaining fine-grained structures and edges is important, such as denoising, super-resolution, and image reconstruction. Moving to the place recognition loss, we have devised a custom loss function with the primary objective of training the model to predict a similarity score between the bottleneck vectors corresponding to the query point clouds and those of the map. This similarity score is calculated as the Euclidean distance between the two corresponding poses but it is transformed using an exponential function to constrain its range between 0 and 1. We approach this as a simple regression problem, where our loss function aims to minimize the mean square error between the network’s predicted similarity score \hat{c} , and the target similarity score, defined as follows:

$$L_{pr}(p_1, p_2, \hat{c}) = (C(p_1, p_2) - \hat{c})^2, \quad (5)$$

where, $C(p_1, p_2) = (\exp(-d(p_1, p_2)/m))^{-1}$, represents the target similarity score, and $d(\cdot)$ calculates the Euclidean distance between the two poses, p_1 and p_2 . The parameter m serves as a threshold that defines the desired scale for the similarity measurement.

The experimental evaluation validates our approach, focusing on two main aspects: place recognition and point cloud reconstruction. Initially, we assess our method using KITTI data [13], which includes LiDAR scans of urban areas in Germany. Our setup follows the previous works [9], [10], with sequences 03-10 for training, sequence 02 for validation, and sequence 00 for evaluation. To test our method’s generalization, we collected LiDAR scans using an Ouster OS1-32 while driving around Luleå University of Technology, referred to as LTU. This scanner has fewer beams, resulting in smaller projected images ($1 \times 32 \times 450$). To adapt, we adjusted the network, setting the bottleneck size to $\beta_{mini} \in \mathbb{R}^{128 \times 32}$ and removing the last two convolutional layers of the encoder and decoder. RecNet is implemented in PyTorch, and training was conducted on a workstation with an NVIDIA GeForce RTX 4090 graphics card and a 13th Gen Intel Core i9-13900K. All experiments were performed on a laptop with an 11th Gen Intel Core i7-1165G7, even without GPU support. The preprocessing time needed for a single point cloud scan is approximately 78ms, while the encoding for the bottleneck vector takes just 2ms.

A. Place Recognition Performance

This section evaluates our approach’s effectiveness in place recognition and loop closure candidate identification. We utilize a 10-meter threshold for the loss function, resulting in a similarity score of about 0.6 for scans at this distance. For candidate classification, we employ a 0.75 threshold, that corresponds to approximately a 3-meter radius. To assess candidate retrieval, we use sequence 00 from the KITTI dataset. We designate the initial 170 seconds, along with the corresponding predicted bottleneck vectors, as map vectors. The remaining data is employed for querying potential candidates. The obtained results, depicted in Fig. 4, indicate performance comparable to or surpassing state-of-the-art lightweight approaches such as LocNet [28], OREOS [10] and OverlapNet [9]. Notably, our bottleneck vector β is not only smaller than OverlapNet’s (approximately 5.5 times smaller) but also retains the ability to reconstruct the point cloud scan. Better performance than OREOS is expected due to the similarity in the encoding process, however RecNet has a greater number of convolutional layers and a larger latent vector, yielding a deeper feature extraction process.




Fig. 4: The place recognition performance of RecNet compared to others, in the 00 sequence of KITTI.

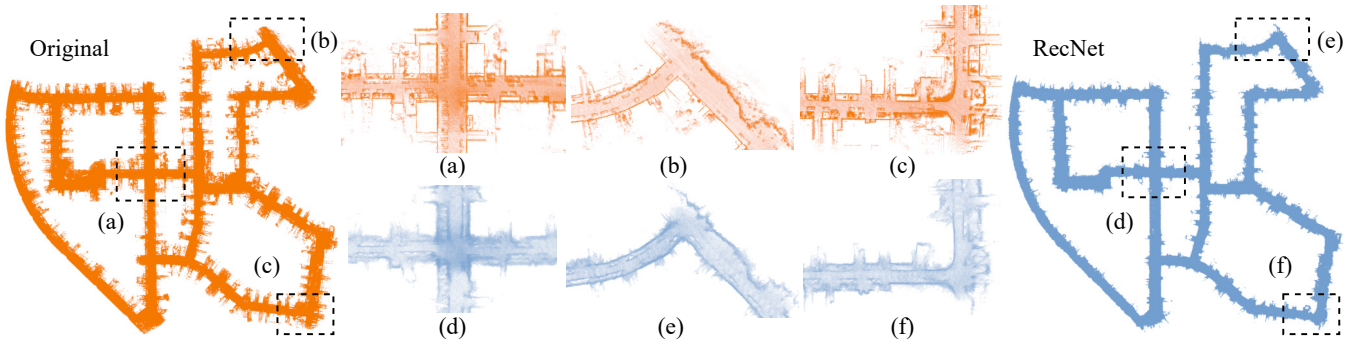


Fig. 5: The original and reconstructed point cloud maps from the 00 sequence of KITTI. The highlighted segments (a) - (c) provide a closer look at the original map, while the segments (d) - (f) demonstrate the corresponding reconstructed parts. RecNet facilitates a reconstruction that resembles the original map by only transmitting the latent vectors of the network.

B. Point cloud Reconstruction

The second phase of the evaluation examines how well RecNet reconstructs the original point cloud, gauging its fidelity and precision in preserving the essential details. In order to evaluate the reconstruction, we use four metrics. The first is a simple correspondence percentage at 0.5 meters. This metric is not very descriptive of the ability of the scan to maintain certain aspects like geometrical shape, curvatures and so on. For that we present the three metric scores proposed in [12], which measure the geometric, curvature and normal vectors' similarity. In addition to the reconstructed scans, we include to the comparison a sub-sampled version of the scans. We downsample with a voxel size of 0.75, in order to match the bottleneck vectors' flattened size. The scores presented in Table II correspond to the mean and standard deviation of the average of all the scans in the sequence 00 of KITTI and all the LiDAR scans attained from the LTU drive. In Fig. 3, we present the original and reconstructed point clouds using our proposed pixel-wise Gradient MSE and the simple MSE. Clearly, the gradient approach preserves line clarity and closely resembles the original point cloud. Fig. 5 and 1 illustrate the results of registering individual point cloud scans based on odometry, for the KITTI dataset and the LTU drive, respectively. Our reconstructed map retains the overall shape and information of the original map. This promising outcome paves the way for evaluating its applicability in real-world multi-robot missions, showcasing its potential for practical use.

TABLE II: Comparison of the PointSSIM [12] similarity score, for different reconstruction loss functions.

	Method	Corr@0.5m [%]	GeomSim [%]	NormSim [%]	CurvSim [%]
KITTI	Downsampled	76.2 ± 2.1	23.4 ± 1.6	80.3 ± 1.3	17.7 ± 2.8
	\mathcal{L}_{MSE}	68.8 ± 2.7	64.2 ± 2.5	80.7 ± 1.1	25.0 ± 2.5
	$\mathcal{L}_{MSE} + \mathcal{L}_{\partial}$	69.2 ± 2.4	65.9 ± 2.5	81.4 ± 1.3	27.8 ± 1.9
LTU	Downsampled	75.6 ± 2.7	25.9 ± 1.7	65.9 ± 2.4	19.5 ± 2.6
	\mathcal{L}_{MSE}	65.8 ± 3.0	68.7 ± 2.6	67.0 ± 2.2	36.4 ± 2.8
	$\mathcal{L}_{MSE} + \mathcal{L}_{\partial}$	67.3 ± 2.9	71.1 ± 2.6	67.8 ± 1.9	38.9 ± 2.4

TABLE III: Statistics for the transmission of the point clouds and the corresponding bottleneck vectors.

Statistics:	KITTI	LTU
Mission Duration (s)	470	368
Bandwidth for transmitting original clouds (kB/s)	2,560	1,220
Bandwidth for bottleneck vectors (kB/s)	131	88
Final map size from original clouds (MB)	897	204
Final size of bottleneck vectors (MB)	62	33

C. Point cloud Transmission

A key objective is to achieve an efficient representation for inter-robot communication, therefore the third aspect measures the reduction in point cloud transmission bandwidth afforded by our approach. This metric serves as a tactile measure of the real-world viability and applicability of our methodology. In Table III, we present a comparative analysis of the bandwidth requirements for transmitting downsampled original point clouds (where the downsampling matches established SLAM practices [29]) against the bandwidth necessary for transmitting only the bottleneck vectors. The results demonstrate the substantial reduction in bandwidth demands achieved through RecNet, with reductions exceeding tenfold in both the KITTI and the LTU drive experiments. This reduction emphasizes the efficiency and potential resource-saving benefits that our method brings to the front in the context of multi-robot systems.

V. CONCLUSIONS

In conclusion, RecNet is a novel framework for extracting descriptive information from point clouds, with place recognition performance akin to state-of-the-art approaches. Moreover, RecNet has the capacity for point cloud reconstruction, yielding substantial bandwidth reductions, particularly beneficial in the context of multi-robot systems. The evaluation of structural similarity between original and reconstructed point clouds, using a state-of-the-art metric on an openly accessible dataset, reinforces its reliability. At last, RecNet paves the way for extensive exploration of its utility in autonomous multi-robot systems, spanning from navigation, localization, map-merging, and beyond.

REFERENCES

- [1] C. Denniston, Y. Chang, A. Reinke, K. Ebadi, G. S. Sukhatme, L. Carlone, B. Morrell, and A. Akbar Agha-mohammadi, "Loop Closure Prioritization for Efficient and Scalable Multi-Robot SLAM," *IEEE Robotics and Automation Letters*, vol. 7, pp. 9651–9658, 2022.
- [2] M. A. Uy and G. H. Lee, "PointNetVLAD: Deep Point Cloud Based Retrieval for Large-Scale Place Recognition," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4470–4479.
- [3] N. Stathouloupoulos, A. Koval, A.-a. Agha-mohammadi, and G. Nikolakopoulos, "FRAME: Fast and Robust Autonomous 3D Point Cloud Map-Merging for Egocentric Multi-Robot Exploration," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 3483–3489.
- [4] A. Agha, K. Otsu, B. Morrell, D. D. Fan, R. Thakker, A. Santamaria-Navarro, S.-K. Kim, A. Bouman, X. Lei, J. Edlund, et al., "NeBula: Quest for Robotic Autonomy in Challenging Environments; TEAM CoSTAR at the DARPA Subterranean Challenge," *Field Robotics*, vol. 2, pp. 1432–1506, 2021.
- [5] J. Gielis, A. Shankar, and A. Prorok, "A Critical Review of Communications in Multi-robot Systems," *Current Robotics Reports*, vol. 3, no. 4, pp. 213–225, Dec 2022.
- [6] G. Damigos, N. Stathouloupoulos, A. Koval, T. Lindgren, and G. Nikolakopoulos, "Communication-aware control of large data transmissions via centralized cognition and 5g networks for multi-robot map merging," p. 22, Jan 2024.
- [7] Y. Chang, L. Ballotta, and L. Carlone, "D-Lite: Navigation-Oriented Compression of 3D Scene Graphs for Multi-Robot Collaboration," *IEEE Robotics and Automation Letters*, vol. 8, no. 11, pp. 7527–7534, 2023.
- [8] B. Lindqvist, N. Stathouloupoulos, S. Karlsson, Y. Bai, A. Patel, J. Haluska, I. Tevetzidis, A. Koval, C. Kanellakis, and G. Nikolakopoulos, "Deployment of Autonomous Uavs in Underground Mines: Field Evaluations and Use-Case Demonstrations," *SSRN*, 2023.
- [9] 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*, vol. 46, 01 2022.
- [10] L. Schaupp, M. Bürki, R. Dubé, R. Siegwart, and C. Cadena, "OREOS: Oriented Recognition of 3D Point Clouds in Outdoor Scenarios," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 3255–3261.
- [11] R. Dubé, A. Cramariuc, D. Dugas, H. Sommer, M. Dymczyk, J. Nieto, R. Siegwart, and C. 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.
- [12] E. Alexiou and T. Ebrahimi, "Towards a Point Cloud Structural Similarity Metric," in *2020 IEEE International Conference On Multimedia And Expo Workshops (ICMEW)*. IEEE, 2020, pp. 6. 1–6. [Online]. Available: <https://github.com/mmspg/pointssim>
- [13] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [14] H. Yin, X. Xu, S. Lu, X. Chen, R. Xiong, S. Shen, C. Stachniss, and Y. Wang, "A Survey on Global LiDAR Localization: Challenges, Advances and Open Problems," 2023.
- [15] L. Theis and W. Shi, "RIDDLE: Lidar Data Compression with Range Image Deep Delta Encoding," in *IEEE / CVF Computer Vision and Pattern Recognition Conference 2022*, 2022, pp. 1–19.
- [16] Y. Feng, S. Liu, and Y. Zhu, "Real-time spatio-temporal LiDAR point cloud compression," *IEEE International Conference on Intelligent Robots and Systems*, pp. 10766–10773, 2020.
- [17] X. Sun, H. Ma, Y. Sun, and M. Liu, "A Novel Point Cloud Compression Algorithm Based on Clustering," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2132–2139, 2019.
- [18] S. Wang and M. Liu, "Point Cloud Compression with Range Image-Based Entropy Model for Autonomous Driving," in *Computer Vision ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXII*. Berlin, Heidelberg: Springer-Verlag, 2022, p. 323340.
- [19] R. Q. Charles, H. Su, M. Kaichun, and L. J. Guibas, "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 77–85.
- [20] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "NetVLAD: CNN Architecture for Weakly Supervised Place Recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 5297–5307.
- [21] Z. Liu, S. Zhou, C. Suo, P. Yin, W. Chen, H. Wang, H. Li, and Y. Liu, "LPD-Net: 3D Point Cloud Learning for Large-Scale Place Recognition and Environment Analysis," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 2831–2840.
- [22] 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.
- [23] K. Vidanapathirana, M. Ramezani, P. Moghadam, S. Sridharan, and C. Fookes, "LoGG3D-Net: Locally Guided Global Descriptor Learning for 3D Place Recognition," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 2215–2221.
- [24] 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.
- [25] G. Elbaz, T. Avraham, and A. Fischer, "3D Point Cloud Registration for Localization Using a Deep Neural Network Auto-Encoder," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2472–2481.
- [26] J. Ma, J. Zhang, J. Xu, R. Ai, W. Gu, and X. Chen, "OverlapTransformer: An Efficient and Yaw-Angle-Invariant Transformer Network for LiDAR-Based Place Recognition," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6958–6965, 2022.
- [27] 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)*, 2020, pp. 5769–5775.
- [28] S. Gidaris and N. Komodakis, "LocNet: Improving Localization Accuracy for Object Detection," in *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*. IEEE Computer Society, 2016, pp. 789–798.
- [29] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and R. Daniela, "LIO-SAM: Tightly-coupled Lidar Inertial Odometry via Smoothing and Mapping," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 5135–5142.