

HypLiLoc: Towards Effective LiDAR Pose Regression with Hyperbolic Fusion

Sijie Wang^{1*} Qiyu Kang^{1*} Rui She¹ Wei Wang² Kai Zhao¹ Yang Song³ Wee Peng Tay¹
¹Nanyang Technological University ²University of Oxford ³C3.AI
 {wang1679;qiyu.kang;rui.she;kai.zhao;wptay}@ntu.edu.sg
 peter.wang1221@gmail.com yang.song@c3.ai

Abstract

LiDAR relocalization plays a crucial role in many fields, including robotics, autonomous driving, and computer vision. LiDAR-based retrieval from a database typically incurs high computation storage costs and can lead to globally inaccurate pose estimations if the database is too sparse. On the other hand, pose regression methods take images or point clouds as inputs and directly regress global poses in an end-to-end manner. They do not perform database matching and are more computationally efficient than retrieval techniques. We propose HypLiLoc, a new model for LiDAR pose regression. We use two branched backbones to extract 3D features and 2D projection features, respectively. We consider multi-modal feature fusion in both Euclidean and hyperbolic spaces to obtain more effective feature representations. Experimental results indicate that HypLiLoc achieves state-of-the-art performance in both outdoor and indoor datasets. We also conduct extensive ablation studies on the framework design, which demonstrate the effectiveness of multi-modal feature extraction and multi-space embedding. Our code is released at: <https://github.com/sijieaaa/HypLiLoc>

1. Introduction

Visual relocalization aims at estimating the 6-degree of freedom (DoF) pose of an agent using perception sensors, such as LiDARs and cameras. It plays a crucial role in many fields that include robot navigation [12], autonomous driving [23], and scene recognition [22]. Image-based relocalization methods have achieved good performance in various applications [15, 33, 36]. However, images taken from cameras can only capture RGB color information and are easily influenced by environmental conditions, including low illumination and light reflections. By contrast, LiDARs, which cast active beams to estimate the depth of surrounding objects, are more robust against those changes.

In recent years, the LiDAR has become an important sensor in smart robots, autonomous vehicles, and mobile devices. LiDAR-based relocalization, which is a basic and important module impacting other perception tasks, has attracted more attention [8, 19, 20, 27, 37, 46]. One of the classical approaches, LiDAR odometry, estimates the relative poses among successive LiDAR frames to obtain locally accurate pose estimation. However, errors accumulate over the trajectory, resulting in unsatisfactory global pose estimation. To compensate for the error, LiDAR odometry is usually treated as a component in a complete simultaneous localization and mapping system (SLAM), where the global pose estimated by a global positioning method or detected loop closure is used to correct the accumulated error in the LiDAR odometry [32, 43].

LiDAR-based retrieval is also used for relocalization [34]. It first constructs a database of LiDAR features learned from all candidate LiDAR frames. During inference, given a query LiDAR scan, the similarities between the query feature and all features stored in the database are computed so that the top-matched poses can be obtained. Although this approach provides accurate global pose estimation, it inherently suffers from high computation cost and storage burden [38]. Therefore, it is more appropriate for offline scenarios rather than for real-time mobile applications.

Pose regression is favored as a relocalization method due to its lower computation and storage cost during inference. The pose regression network is still trained on a database containing LiDAR frames in an end-to-end manner to obtain a regression model. During inference, taking the LiDAR scan as input, the pose regression network directly regresses the global pose without any pre-constructed candidate database or map. It can mitigate the high computation and storage burden that occurs in the LiDAR-based retrieval methods. As a result, pose regression can be operated in real-time to satisfy various relocalization requirements in robotics, unmanned aerial vehicles (UAVs), mobile relocalization APPs, autonomous vehicles, and SLAM systems.

In this paper, we propose a relocalization method called HypLiLoc, which is a pose regression network with LiDAR

*These authors contribute equally.

data as input. HypLiLoc uses a parallel feature extraction design, in which 3D features and 2D spherical projection features are obtained in two backbone branches simultaneously. The paper [24] leverages hyperbolic embeddings for 3D point clouds that can be viewed as hierarchical compositions of small parts. We thus follow this motivation to design our pipeline with hyperbolic learning. Specifically, we conduct feature fusion in both Euclidean and hyperbolic spaces to enhance the information representation and to achieve more effective multi-modal feature interaction. We test HypLiLoc in both outdoor and indoor datasets. Experiments indicate that HypLiLoc surpasses current approaches and achieves state-of-the-art (SOTA) performance.

Our main contributions are summarized as follows:

1. We propose a novel LiDAR-based pose regression network HypLiLoc. It has one backbone that learns 3D features directly from the 3D point cloud and another backbone that learns features from a 2D projection of the point cloud onto a spherical surface. To achieve effective multi-modal feature interaction, the features are embedded in both Euclidean and hyperbolic spaces using multi-space learning. An attention mechanism is then used to fuse the features from different spaces together.
2. We test our network in both outdoor and indoor datasets, where it outperforms current LiDAR pose regression counterparts and achieves SOTA performance. We also conduct extensive ablation studies on the effectiveness of each design component.

2. Related Work

In this section, we shall introduce more relocalization works, including LiDAR odometry, point cloud retrieval, and pose regression. Besides, we shall also provide more details of hyperbolic learning that are related to our method design.

2.1. LiDAR Odometry

LiDAR odometry methods address the relocalization problem under local views. Given several close or nearby LiDAR scans, they estimate the relative poses among them. The iterative closest point (ICP) method [6] solves the relative pose by iteratively searching correspondence between source points and target points and optimizing the least square error. Besides, another popular method is LOAM [46], which classifies the keypoints into edges and planes and uses KD-trees to search the neighborhood of the keypoints. DCP [39] uses PointNet [28] and DGCNN [41] as backbones to extract point cloud features, and then it predicts the pose using the Transformer module.

2.2. Point Cloud Retrieval

Point cloud retrieval approaches treat the relocalization task as the place recognition problem [34]. The core of these approaches is query-database matching. A database needs to be constructed to store the features of all candidate LiDAR scans with corresponding poses. In the inference process, for a given query LiDAR scan, its feature is extracted by the neural network, and then the query-database feature matching is performed for every possible pair. The final pose estimation is obtained from the top-matched pairs.

2.3. Pose Regression

Given the query sensor data, the pose regression models directly outputs the pose using the trained neural network. They do not depend on the query-database matching procedure, which speeds up the inference stage significantly when compared to point cloud retrieval methods. These models are still trained on a database of training samples that include sensor scans and the ground truth sensor poses. However, during the inference stage, the database is no longer required, in contrast to retrieval methods.

PoseNet [17] proposes simultaneous learning for location and orientation by integrating balance parameters. MapNet [15] uses visual odometry as the post-processing technique to optimize the regressed poses. AD-PoseNet [16] leverages semantic masks to drop out the dynamic area in the image. AtLoc [36] introduces global attention to guide the network to learn better representations. MS-Transformer [33] focuses on simultaneous pose regression for multiple scenes using a single network.

A LiDAR actively casts beams to estimate the sparse depth of its surrounding environment. Since LiDARs are less likely to be influenced by illumination changes than cameras, they have become core sensors in many applications. PointLoc [38] uses LiDAR point clouds to achieve pose regression. Its model consists of PointNet++ [29] followed by self-attention modules to generate point cloud features. The paper [45] studies the memory-friendly pose regression learning scheme and proposes four LiDAR pose regression models.

2.4. Hyperbolic Learning

Hyperbolic embedding for features has been proposed for datasets that have some underlying tree structure [31]. The paper [13] derives hyperbolic versions of several deep learning tools, including multinomial logistic regression, feed-forward networks, and recurrent networks. In the field of natural language processing (NLP), [25] and [26] introduce feature embeddings with hyperbolic models. In the field of deep graph learning, HGCN [7] considers hyperbolic node embeddings in Graph Convolutional Neural Networks (GCNs). GIL [47] proposes to use weighted embedding features in

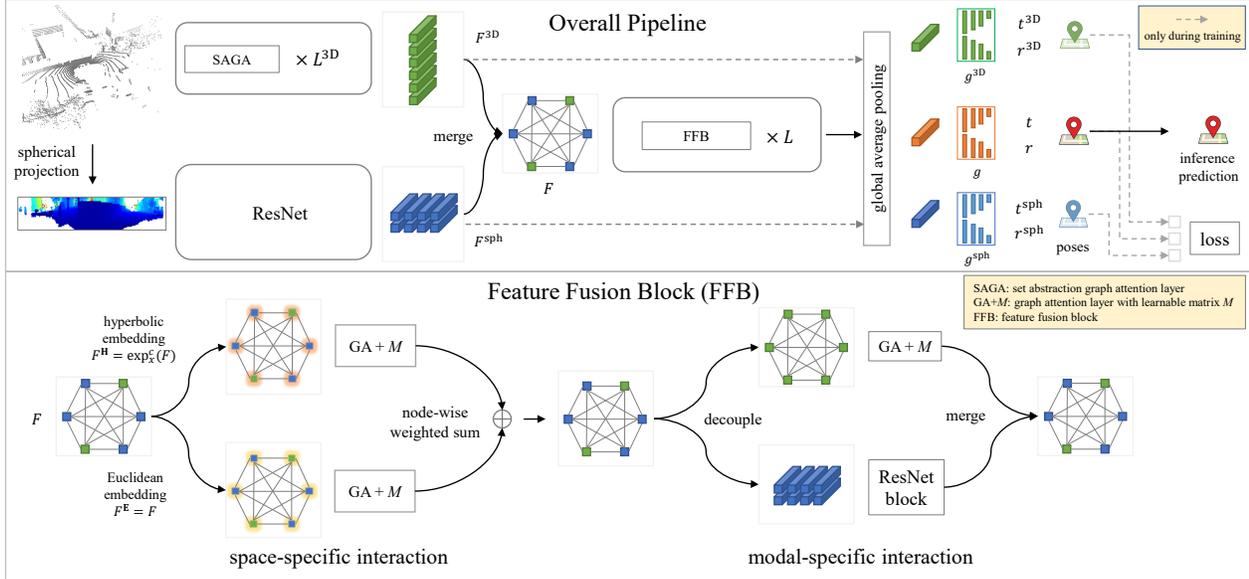


Figure 1. The overall architecture of our proposed HypLiLoc. We use two backbone branches to perform feature extraction. In the 3D backbone, we consider both local set abstraction and global attention aggregation. In the feature fusion block, the extracted multi-modal features are embedded into both Euclidean and hyperbolic spaces to achieve space-specific interaction. The fusion features are then decoupled to their own modality to perform modal-specific interaction. The final training loss is applied on both the 3D/projection level and the final fusion level.

both Euclidean and hyperbolic spaces. In the computer vision community, [11] uses pair-wise cross-entropy loss with hyperbolic distances to train the vision transformer [10], and [4] considers hyperbolic embeddings in the semantic segmentation task. More recently, hyperbolic embeddings have also been studied for the 3D point cloud [24], where the 3D point cloud is treated as nature compositions of small parts that follow the hierarchical architecture. This motivates us to introduce hyperbolic embeddings in our pipeline for better feature representations.

3. Proposed Model

In this section, we provide a detailed description of our proposed approach. We first summarize the HypLiLoc pipeline as follows.

1. Given a LiDAR point cloud scan, in addition to the traditional backbone of extracting 3D features from the point cloud, we additionally project the 3D points into a sphere to generate a 2D projection image. These two types of features are extracted by separate backbones.
2. We merge the two modal features together as the fusion features. The fusion features are then embedded in both Euclidean and hyperbolic spaces to achieve more effective representations.
3. After features interact in different spaces and modalities,

the global feature vector is obtained by applying the global average pooling operation on the fusion features. The final pose prediction is generated using the global feature vector with the pose regression head.

3.1. Modal-Specific Backbones

Projection Feature Extraction. Multi-modal feature extraction has shown promising performance in various tasks [1, 30, 42, 44]. The point cloud generated by LiDARs is convertible into multiple modalities by projecting 3D points into specific 2D spaces. Each projection provides us with a different way to define the neighbors of a point so that the point can aggregate feature representations from different definitions of its “neighborhood”. To this end, we consider two typical projection methods, including the spherical projection and the bird’s-eye view (BEV) projection. For the currently most commonly used multi-line spinning LiDAR, we visualize the point cloud projection in Fig. 2.

Suppose we have as the input a set of N LiDAR points $\{p_i\}_{i=1}^N$, each represented by 3D Cartesian coordinates $p_i = (x_i, y_i, z_i) \in \mathbb{R}^3$. For the spherical projection, we first convert the Cartesian coordinates into polar coordinates as:

$$\phi_i = \arctan\left(\frac{z_i}{\sqrt{x_i^2 + y_i^2}}\right), \quad (1)$$

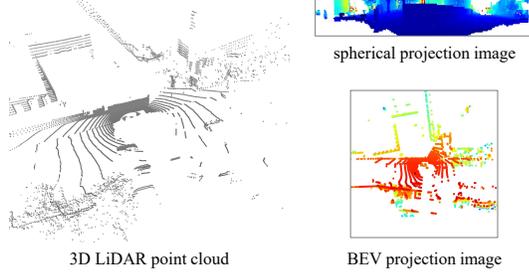


Figure 2. Visualization of the spherical and BEV projection methods.

$$\theta_i = \arctan\left(\frac{y_i}{x_i}\right), \quad (2)$$

$$r_i = \sqrt{x_i^2 + y_i^2 + z_i^2}. \quad (3)$$

Each point projected on the sphere is then denoted as $p_i^{\text{sph}} = (\phi_i, \theta_i) \in \mathbb{R}^2$. Projecting all the N points, an image $I^{\text{sph}} \in \mathbb{R}^{H \times W}$ is obtained, in which each pixel contains the radius value:

$$I^{\text{sph}}\left(\left\lfloor \frac{\phi_i H}{2\pi} \right\rfloor, \left\lfloor \frac{\theta_i W}{2\pi} \right\rfloor\right) = r_i. \quad (4)$$

On the other hand, the BEV projects each point p_i as (x_i, y_i) and generates the image $I^{\text{BEV}} \in \mathbb{R}^{H' \times W'}$, in which each pixel contains the height value:

$$I^{\text{BEV}}\left(\left\lfloor \frac{y_i H'}{2y_{\max}} \right\rfloor, \left\lfloor \frac{x_i W'}{2x_{\max}} \right\rfloor\right) = z_i, \quad (5)$$

where x_{\max} and y_{\max} are the maximum LiDAR range in the x and y directions, respectively.

We test the two projection counterparts in Section 4.4, where the spherical projection performs better than the BEV. The reason could be that LiDARs operate with the spinning mechanism, which is better modeled by the spherical projection. By contrast, the BEV projection loses information as some points are stacked on the same pixel and is thus not a bijective mapping.

In the following discussion, we use the spherical projection strategy in our pipeline. We treat the spherical projection points as the image modality input, while the 3D features are extracted directly from the point cloud in a separate backbone that we will introduce later. Following the golden rule for 2D image processing, we use ResNet [14] as the backbone for the image modality. Denoting the ResNet backbone as $f^{\text{sph}}(\cdot)$, the final spherical projection features $F^{\text{sph}} \in \mathbb{R}^{H^{\text{sph}} \times W^{\text{sph}} \times C}$ are obtained as:

$$F^{\text{sph}} = f^{\text{sph}}(I^{\text{sph}}). \quad (6)$$

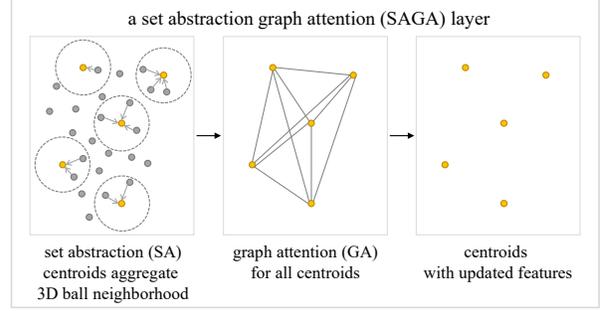


Figure 3. The SAGA layer consists of a SA layer and a GA layer.

3D Feature Extraction. Effective 3D point feature extraction is critical in the model design. PointNet++ [29] has shown promising performance in various tasks [38, 39]. In our pipeline, we use PointNet++ as the backbone branch for 3D feature extraction.

PointNet++ only considers the neighboring information within a determined range, i.e., in the set abstraction (SA) layer, each centroid uses the maximum neighboring feature value as its updated feature as shown in Fig. 3. In pose regression, the estimation accuracy of the pose benefits from an effective global representation. Thus to enable PointNet++ to additionally aggregate more global information, we introduce an additional graph attention (GA) layer after each SA layer to build the set abstraction graph attention (SAGA) layer as shown in Fig. 3. In this GA layer, we construct a complete graph whose node set contains all centroids from the SA layer. We denote the output of the SA layer as $P \in \mathbb{R}^{N^{\text{SA}} \times C^{\text{SA}}}$ with N^{SA} centroids, each with a C^{SA} -dimensional feature vector. We first use a Fully-Connected (FC) layer to generate the multi-head features:

$$P_k^{\text{FC}} = PW_k + b_k, \quad (7)$$

where W_k and b_k are a linear operation and an additive bias, respectively. These are learnable parameters of the k -th head. Then the attention weight matrix $A_k \in \mathbb{R}^{N^{\text{SA}} \times N^{\text{SA}}}$ can be obtained by computing the dot product among all the neighboring nodes:

$$A_k = \text{Softmax}\left(P_k^{\text{FC}} P_k^{\text{FC}\top}\right), \quad (8)$$

where $\text{Softmax}(\cdot)$ denotes the row-wise softmax function. The output features of the GA layer $P^{\text{GA}} \in \mathbb{R}^{N^{\text{GA}} \times C^{\text{GA}}}$ are generated by concatenating the weighted features from all heads as:

$$P^{\text{GA}} = \left\| \left\| A_k P_k^{\text{FC}} \right\|_k \right\|, \quad (9)$$

where $\|$ denotes the concatenation operation. We then stack L^{3D} such SAGA layers to build the 3D features extraction

backbone. We denote the final output features as $F^{3D} \in \mathbb{R}^{N^{3D} \times C}$ shown in Fig. 1.

3.2. Hyperbolic Feature Learning

In this subsection, we first state the motivation for such hyperbolic feature learning. We then introduce the hyperbolic embedding operators that will be used in Section 3.3 to fuse features extracted from the 3D LiDAR point cloud and spherical LiDAR projection in Section 3.1.

Motivation. After feature extraction using the two backbone branches, we need an effective fusion strategy to consider both point features and projection features. Embedding in a hyperbolic space has recently gained increased interest and shown promising performance in various fields [4, 7, 11, 13, 47]. The paper [24] argues that 3D point cloud objects possess inherent hierarchies due to their nature as compositions of small parts, which can be embedded in the hyperbolic space. Following this motivation, we consider leveraging the hyperbolic embedding method in our pipeline, such that features can be equipped with more various representations that come from different embedding spaces. Our ablation study (cf. Table 3 of Section 4.4) also indicates that hyperbolic embedding can lead to improvements in the pose estimation accuracy.

Hyperbolic Embedding. Different from the common Euclidean space, the hyperbolic space is equipped with constant negative curvature and has a different metric rather than the Euclidean ℓ_2 norm $\|\cdot\|$. This special metric renders a ball in hyperbolic embedding space to have exponentially increased volume with respect to its radius rather than polynomially as in Euclidean spaces.

Similar to [11], we use the n -dimensional *Poincaré ball* ($\mathbb{D}_c^n, g^{\mathbb{D}}$) for our hyperbolic embedding with the parameter c indicating constant negative curvature $-c^2$. More specifically, $g^{\mathbb{D}}$ is the Riemannian metric, and \mathbb{D}_c^n is defined as:

$$\mathbb{D}_c^n = \{x \in \mathbb{R}^n : c\|x\|^2 < 1, c \geq 0\}, \quad (10)$$

where the distance d between two points x and y on \mathbb{D}_c^n is defined as:

$$d(x, y) = \frac{2}{\sqrt{c}} \operatorname{arctanh}(\sqrt{c}\| -x \oplus_c y \|), \quad (11)$$

where \oplus_c is the *Mobius addition* defined as follows:

$$x \oplus_c y = \frac{(1 + 2c\langle x, y \rangle + c\|y\|^2)x + (1 - c\|x\|^2)y}{1 + 2c\langle x, y \rangle + c^2\|x\|^2\|y\|^2}. \quad (12)$$

To support features transferring from Euclidean spaces to hyperbolic spaces, the differentiable bijective operator named *exponential map* is induced. For a fixed base point $x \in \mathbb{D}_c^n$, where the tangent space at x is a Euclidean space, the exponential map $\exp_x^c : \mathbb{R}^n \rightarrow \mathbb{D}_c^n$ establishes the connection

between the tangent Euclidean space and the hyperbolic space at x as:

$$\exp_x^c(v) = x \oplus_c \left(\tanh\left(\sqrt{c} \frac{\lambda_x^c \|v\|}{2}\right) \frac{v}{\sqrt{c}\|v\|} \right), \quad (13)$$

where λ_x^c is the *conformal factor*. In our pipeline, we assume the input features are in this tangent space and would like to embed them in the hyperbolic space \mathbb{D}_c^n .

3.3. Feature Fusion Block

Based on the above-mentioned motivation, we propose the feature fusion block (FFB) to achieve effective feature interaction. Each FFB conducts both space-specific interaction and modal-specific interaction alternatively, which is similar to the commonly used cross-self-attention operation. We stack L such FFBs in our pipeline.

Feature Merging. Given the extracted 3D features F^{3D} and spherical projection features F^{sph} , we first pass them through an ℓ_2 normalization layer such that all features are constrained on a sphere. This is a common way to process multi-modal data [30]. We then formulate a fusion graph with complete edge connections, in which each node contains features from either F^{3D} or F^{sph} . In addition, to enable features to interact directly with the global representation, we add two extra node features that are processed by the global average pooling module. The fusion graph node features are collected in the following set:

$$F = \{F^{3D}, F^{\text{sph}}, \text{Pooling}(F^{3D}), \text{Pooling}(F^{\text{sph}})\}. \quad (14)$$

Space-specific Interaction. We embed the fusion features F into the Euclidean and hyperbolic spaces as $F^{\mathbb{H}} = \exp_x^c(F)$ (where the exp operator is applied node-wise) and $F^{\mathbb{E}} = F$, respectively, to perform feature interaction using GA layers. Specifically, in the same way as (7), we obtain the k -th head FC features for $F^{\mathbb{H}}$ or $F^{\mathbb{E}}$, denoted as F_k^{FC} . We additionally leverage a learnable matrix M regarded as a feature relationship metric such that the attention weights are computed as:

$$A_k^M = \text{Softmax}\left(F_k^{\text{FC}} M F_k^{\text{FC}\top}\right). \quad (15)$$

In Riemannian geometry, a Riemannian metric on a smooth manifold is a smooth symmetric covariant 2-tensor field that is positive definite at each point. The learnable matrix M can be viewed as a more general extension of the Riemannian metric, where we do not impose any constraint on it, leaving it to update freely. The effectiveness of this design can be seen in Table 5, where the free metric surpasses other counterparts.

The learned feature embeddings from the Euclidean and hyperbolic spaces are then passed into two different GA

layers and finally fused together using element-wise adding:

$$F^{\text{space}} = w^{\text{E}} F^{\text{E}} + w^{\text{H}} F^{\text{H}}, \quad (16)$$

where w^{E} and w^{H} denote learnable weights for the Euclidean and hyperbolic embeddings F^{E} and F^{H} , respectively. Each node feature has thus aggregated information from both Euclidean and hyperbolic spaces, which can be viewed as an adaptive combination of linearity and non-linearity that can contribute to a more effective feature representation.

Modal-specific Interaction. We next decouple the merged features back to 3D features and projection features again, enabling them to turn around and learn information within their own modality. This is similar to the self-attention operation in the cross-self-attention pipeline. Specifically, for the 3D features, we pass them through a GA layer (with the learnable matrix) with preceding and succeeding MLP layers, while for the 2D projection features, we pass them through a basic ResNet block. After the modal-specific interaction, 3D features and projection features are merged together again using (14) to reconstruct the fusion features.

3.4. Pose Regression Head and Loss Function

The task of LiDAR pose regression requires predicting a 6-DoF pose. However, since the translation and rotation elements do not scale compatibly, the regression converges in different basins. To deal with this problem, previous methods [38, 45] consider the regression head with two parallel MLPs for translation and rotation regression, respectively. We thus use the same decoding head design as [38], which consists of two MLP layers for translation and rotation regression as shown in Fig. 1.

During training, to provide sufficient supervision to the whole pipeline, we use not only the fusion features but also the 3D and projection features at lower levels. As shown in Fig. 1, for the three features $F^{3\text{D}}$, F^{sph} , and F , we use three different regression heads $g^{3\text{D}}$, g^{sph} , g respectively to predict their corresponding 6-DoF poses $(t^{3\text{D}}, r^{3\text{D}})$, $(t^{\text{sph}}, r^{\text{sph}})$, (t, r) . Specifically, we first perform global average pooling and then regression to obtain the predicted poses, which can be described as follows:

$$(t^{3\text{D}}, r^{3\text{D}}) = (g^{3\text{D}} \circ \text{Pooling})(F^{3\text{D}}), \quad (17)$$

$$(t^{\text{sph}}, r^{\text{sph}}) = (g^{\text{sph}} \circ \text{Pooling})(F^{\text{sph}}), \quad (18)$$

$$(t, r) = (g \circ \text{Pooling})(F), \quad (19)$$

where \circ denotes the composition operation, and $\text{Pooling}(\cdot)$ denotes the global average pooling operation. As for the rotation, we use the logarithmic format of the quaternion [38, 45]. Denoting the translation and rotation targets as t^*

and r^* , the final loss function is computed as:

$$\begin{aligned} \mathcal{L} = & (\|t^{3\text{D}} - t^*\| + \|t^{\text{sph}} - t^*\| + \|t - t^*\|)e^{-\lambda} + \lambda \\ & + (\|r^{3\text{D}} - r^*\| + \|r^{\text{sph}} - r^*\| + \|r - r^*\|)e^{-\gamma} + \gamma \end{aligned} \quad (20)$$

where λ and γ are learnable parameters. During inference, the pose (t, r) predicted by the fusion features F is treated as the final prediction.

4. Experiments

In this section, we first evaluate our proposed model on datasets collected from outdoors and indoors. We next present ablation studies to demonstrate the effectiveness of our model design.

4.1. Implementation Details

We use ResNet34 [14] pre-trained on ImageNet [9] as the backbone for projection features extraction. We use a batch size of 32. The number of attention heads is set as 8. Following [11], we set the base point x as 0 for hyperbolic embedding. We set $L^{3\text{D}} = 2$ and $L = 2$. The Adam [18] optimizer with the initial learning rate 1×10^{-3} and weight decay 5×10^{-4} is used for training. We train our network for 150 epochs. All the experiments are conducted on either an NVIDIA RTX 3090 GPU or an NVIDIA RTX A5000 GPU.

4.2. Datasets

Oxford Radar is a large-scale outdoor autonomous driving dataset [5]. It provides data from multi-modal sensors, including LiDARs, cameras, Radars, and GPS, but in our experiments, we use only LiDAR information. It contains sensor data in the time span of 1 year and a length span of 1000 km. In addition, it covers various seasons and weather conditions, which thus allows a comprehensive evaluation of the models. Following [38, 45], we use the same benchmark data split setting, and we also report the mean translation rotation error.

vReLoc is an indoor robot dataset [38]. It consists of data from LiDARs, cameras, depth cameras, and motion trackers. In our experiments, we use only LiDAR information. It contains both static and dynamic scenarios with people walking around. Following [38, 45], we use the same benchmark data split setting, and we also report the median translation and rotation error.

4.3. Main Results

We first compare HypLiLoc with other baselines on the Oxford Radar dataset. From Table 1, we observe that HypLiLoc achieves SOTA performance in all metrics. Especially on the route Full-9, HypLiLoc obtains 3.45 m mean translation error in the city-wise relocalization task compared to the second best performer PoseSOE with 7.27 m,

	Model	Full-6	Full-7	Full-8	Full-9
<i>LiDAR Retrieval</i>	PointNetVLAD [34]	28.48 / 5.19	17.62 / 3.95	23.59 / 5.87	13.71 / 2.57
<i>LiDAR Odometry</i>	DCP [40]	18.45 / 2.08	14.84 / 2.17	16.39 / 2.26	13.60 / 1.86
<i>Image-based PR</i>	PoseLSTM [35]	26.36 / 6.54	74.00 / 9.85	128.25 / 18.59	19.12 / 3.05
	MapNet [15]	48.21 / 6.06	61.01 / 5.85	75.35 / 9.67	44.34 / 4.54
	AD-MapNet [16]	18.43 / 3.28	19.18 / 3.95	66.21 / 9.42	15.10 / 1.82
	AtLoc+ [36]	17.92 / 4.73	34.03 / 4.01	71.51 / 9.91	10.53 / 1.97
	MS-Transformer [33]	11.69 / 5.66	65.38 / 9.01	88.63 / 19.80	7.62 / 2.53
<i>LiDAR-based PR</i>	PointLoc [38]	13.81 / <u>1.53</u>	9.81 / <u>1.27</u>	11.51 / <u>1.34</u>	9.51 / <u>1.07</u>
	PosePN [45]	16.32 / 2.43	14.32 / 3.06	13.48 / 2.60	9.14 / 1.78
	PosePN++ [45]	10.64 / 1.78	9.59 / 1.92	<u>9.01</u> / 1.51	8.44 / 1.71
	PoseSOE [45]	<u>8.81</u> / 2.04	<u>7.59</u> / 1.94	9.21 / 2.12	<u>7.27</u> / 1.87
	PoseMinkLoc [45]	11.20 / 2.62	14.69 / 2.90	12.35 / 2.46	10.06 / 2.15
	HypLiLoc (ours)	6.00 / 1.31	6.88 / 1.09	5.82 / 0.97	3.45 / 0.84

Table 1. Mean translation and rotation error (m°) on the Oxford Radar dataset. The best and the second-best results in each metric are highlighted in **bold** and underlined, respectively. PR stands for pose regression. HypLiLoc achieves the best performance in all metrics.

	Model	Seq-05	Seq-06	Seq-07	Seq-14
<i>Image-based PR</i>	PoseLSTM [35]	0.16 / 4.23	0.18 / 5.28	0.24 / 7.05	0.13 / 4.81
	MapNet [15]	0.26 / 6.67	0.28 / 6.91	0.39 / 9.17	0.25 / 6.85
	AD-MapNet [16]	0.17 / 3.33	0.21 / 3.37	0.24 / 4.38	0.14 / 4.12
	AtLoc+ [36]	0.18 / 4.32	0.24 / 5.14	0.26 / 6.04	0.16 / 4.61
	MS-Transformer [33]	0.16 / 3.98	0.15 / 3.56	0.18 / 5.32	0.13 / 4.83
<i>LiDAR-based PR</i>	PointLoc [38]	<u>0.12</u> / 3.00	0.10 / 2.97	0.13 / 3.47	0.11 / 2.84
	PosePN [45]	<u>0.12</u> / 4.38	<u>0.09</u> / 3.16	0.17 / 3.94	0.08 / 3.27
	PosePN++ [45]	0.15 / 3.12	0.10 / 3.31	<u>0.15</u> / <u>2.92</u>	0.10 / <u>2.80</u>
	PoseSOE [45]	0.14 / 3.15	0.11 / <u>2.90</u>	<u>0.15</u> / 3.06	0.11 / 3.20
	PoseMinkLoc [45]	0.16 / 5.17	0.11 / 3.74	0.21 / 5.74	0.12 / 3.64
	HypLiLoc (ours)	0.09 / 2.52	0.08 / 2.58	0.13 / 2.55	<u>0.09</u> / 2.34

Table 2. Median translation and rotation error (m°) on the vReLoc dataset. The best and the second-best results in each metric are highlighted in **bold** and underlined, respectively. PR stands for pose regression. HypLiLoc achieves the best performance in 7 out of 8 metrics.

which demonstrates the effectiveness of HypLiLoc. In addition, compared with camera pose regression approaches that take images as inputs, LiDAR-based ones are generally more accurate. This verifies that point clouds generated by LiDARs are a more effective data modality for the re-localization task. Table 1 also indicates that for large-scale pose estimation, LiDAR pose regression approaches surpass both retrieval-based and odometry-based ones, and thus this approach is promising for many applications.

Note that pose regression approaches can be integrated into SLAM systems [2, 3, 21] to achieve even better accuracy and to perform fast global pose estimating, especially in cases where a global navigation satellite system is not available (e.g., indoors and urban areas with dense skyscrapers).

We next test HypLiLoc on the indoor vReLoc dataset in Table 2, where it achieves SOTA performance in 7 out of 8 metrics and shows strong competitiveness. We note in the indoor environment, the LiDAR-based approaches also generally outperform image-based ones.

Method	Mean Error (m°) on Full-8
base model	9.78 / 1.99
+ global graph attention	8.91 / 1.74
+ spherical-projection backbone	7.26 / 1.36
+ feature fusion block	6.19 / 1.13
+ learnable metric (full model)	5.82 / 0.97
full model w/o hyperbolic branch	6.57 / 1.19
full model w/o Euclidean branch	6.24 / 1.16

Table 3. Ablation study for different modules on Full-8 route of the Oxford Radar dataset.

4.4. Analysis

Ablation Study. We provide insights into our design choices for HypLiLoc by ablating each module. From Table 3, every module in our design contributes to the final improved estimation accuracy. Making use of information from the projected point cloud image and hyperbolic-Euclidean feature fusion strategy both contribute to more accurate pose regression outputs.

#Modalities	3D	Sph.	BEV	Mean Error (m/°) on Full-8
1	✓			8.91 / 1.74
		✓		8.94 / 2.18
			✓	9.46 / 2.33
2	✓	✓		5.82 / 0.97
	✓		✓	9.44 / 1.80
		✓	✓	6.77 / 1.02
3	✓	✓	✓	6.32 / 1.01

Table 4. Comparison of different projection methods on Full-8 route of the Oxford Radar dataset. For the single modality, we do not use the feature fusion block.

Method	Mean Error (m/°) on Full-8
w/o M	6.19 / 1.13
Riemannian	6.34 / 1.28
positive definite	6.18 / 1.13
symmetric	6.02 / 1.24
no constraint	5.82 / 0.97

Table 5. Comparison of different constraints on Full-8 route of the Oxford Radar dataset.

Different Projection Strategies. We next compare different modality strategies. As shown in Table 4, we first test the performance using the single modality input, including the 3D point cloud, the spherical projection, and the BEV projection. Among them, the 3D and the spherical projection show similar performances, while the BEV performance is worse. This verifies our insight that the BEV projection is not a bijective mapping, and thus less information is retained.

When feeding two modalities, the combination *3D + spherical* surpasses the other two counterparts, which is our final model choice for HypLiLoc. When we further add the BEV input, the performance drops instead.

Learnable Matrix M Design. We test the performance by applying different constraints on the learnable matrix M in (15). The Riemannian metric, which formulates a positive definite and symmetric matrix, has the strictest constraints. However, as shown in Table 5, the Riemannian metric does not provide performance improvements compared with the setting without any metric. If we only impose either the positive definite constraint or the symmetric constraint, the performance improves. Furthermore, if we exclude all constraints and enable the metric to evolve freely, we can achieve optimal performance.

Computational Time and Storage. We compare LiDAR-based models that belong to different relocalization pipelines. As observed in Table 6, regression-based models can operate at least 2 times as fast as both the retrieval-based and odometry-based approaches. For the runtime memory, regression-based models need only 1/3 that of retrieval and odometry methods (less than 7 GB). HypLiLoc can perform inference at a speed of 48 FPS, which is 4 times as fast as

	Model	Runtime Speed	Runtime Total Memory
<i>Retrieval</i>	PointNetVLAD [34]	11FPS	26GB
<i>Odometry</i>	DCP [40]	10FPS	22GB
<i>Regression</i>	PointLoc [38]	22FPS	7GB
	HypLiLoc (ours)	48FPS	6GB

Table 6. Comparison of the runtime speed and the runtime total memory of different models.

than the retrieval model and over 2 times as fast as PointLoc applications.

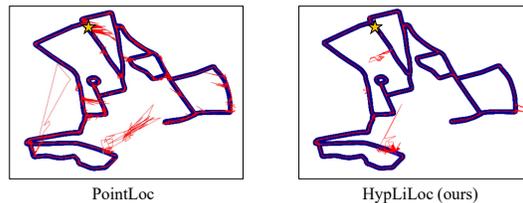


Figure 4. Trajectory visualization on the Oxford Radar dataset. The ground truth trajectories are shown in bold blue lines, and the estimated trajectories are shown in thin red lines.

Visualization. We visualize a typical output pose trajectory of HypLiLoc and PointLoc in Fig. 4. HypLiLoc outputs a smoother and more accurate pose trajectory compared with PointLoc.

5. Limitations

Although we have tested the proposed model in a city-wise dataset, verification of HypLiLoc’s performance in challenging scenarios (e.g. with noise perturbations and adversarial attacks) is necessary for practical implementations.

6. Conclusion

In this work, we propose HypLiLoc, a novel network for LiDAR-based pose regression. It achieves effective feature extraction with global graph attention, hyperbolic-Euclidean interaction, and modal-specific learning. It achieves SOTA performance in both outdoor and indoor datasets.

7. Acknowledgment

This research is supported by A*STAR under its RIE2020 Advanced Manufacturing and Engineering (AME) Industry Alignment Fund – Pre Positioning (IAF-PP) (Grant No. A19D6a0053) and the National Research Foundation, Singapore and Infocomm Media Development Authority under its Future Communications Research and Development Programme.

References

- [1] Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text. *Advances in Neural Information Processing Systems*, 34:24206–24221, 2021. [3](#)
- [2] Thangarajah Akilan, Edna Johnson, Gaurav Taluja, Japneet Sandhu, and Ritika Chadha. Multimodality weight and score fusion for slam. In *IEEE Canadian Conference on Electrical and Computer Engineering (CCECE)*, pages 1–4. IEEE, 2020. [7](#)
- [3] Ihtisham Ali, Sari Peltonen, and Atanas Gotchev. Bi-directional loop closure for visual slam. *arXiv preprint arXiv:2204.01524*, 2022. [7](#)
- [4] Mina Ghadimi Atigh, Julian Schoep, Erman Acar, Nanne van Noord, and Pascal Mettes. Hyperbolic image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4453–4462, 2022. [3](#), [5](#)
- [5] Dan Barnes, Matthew Gadd, Paul Murcutt, Paul Newman, and Ingmar Posner. The oxford radar robotcar dataset: A radar extension to the oxford robotcar dataset. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6433–6438. IEEE, 2020. [6](#)
- [6] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor fusion IV: control paradigms and data structures*, volume 1611, pages 586–606. Spie, 1992. [2](#)
- [7] Ines Chami, Zhitao Ying, Christopher Ré, and Jure Leskovec. Hyperbolic graph convolutional neural networks. *Advances in neural information processing systems*, 32, 2019. [2](#), [5](#)
- [8] Pierre Dellenbach, Jean-Emmanuel Deschaud, Bastien Jacquet, and François Goulette. Ct-icp: Real-time elastic lidar odometry with loop closure. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 5580–5586. IEEE, 2022. [1](#)
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [6](#)
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [3](#)
- [11] Aleksandr Ermolov, Leyla Mirvakhabova, Valentin Khrukov, Nicu Sebe, and Ivan Oseledets. Hyperbolic vision transformers: Combining improvements in metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7409–7419, 2022. [3](#), [5](#), [6](#)
- [12] Baofu Fang, Gaofei Mei, Xiaohui Yuan, Le Wang, Zaijun Wang, and Junyang Wang. Visual slam for robot navigation in healthcare facility. *Pattern Recognition*, 113:107822, 2021. [1](#)
- [13] Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. *Advances in neural information processing systems*, 31, 2018. [2](#), [5](#)
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [4](#), [6](#)
- [15] Joao F Henriques and Andrea Vedaldi. Mapnet: An allocentric spatial memory for mapping environments. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8476–8484, 2018. [1](#), [2](#), [7](#)
- [16] Zhaoyang Huang, Yan Xu, Jianping Shi, Xiaowei Zhou, Hujun Bao, and Guofeng Zhang. Prior guided dropout for robust visual localization in dynamic environments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2791–2800, 2019. [2](#), [7](#)
- [17] Alex Kendall, Matthew Grimes, and Roberto Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *Proceedings of the IEEE international conference on computer vision*, pages 2938–2946, 2015. [2](#)
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [19] Kenji Koide, Masashi Yokozuka, Shuji Oishi, and Atsuhiko Banno. Voxelized gicp for fast and accurate 3d point cloud registration. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11054–11059. IEEE, 2021. [1](#)
- [20] Bin Li, Mu Hu, Shuling Wang, Lianghao Wang, and Xiaojin Gong. Self-supervised visual-lidar odometry with flip consistency. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3844–3852, 2021. [1](#)
- [21] Ruihao Li, Sen Wang, and Dongbing Gu. Ongoing evolution of visual slam from geometry to deep learning: Challenges and opportunities. *Cognitive Computation*, 10(6):875–889, 2018. [7](#)
- [22] Huei-Yung Lin, Chia-Wei Yao, Kai-Sheng Cheng, and Van Luan Tran. Topological map construction and scene recognition for vehicle localization. *Autonomous Robots*, 42(1):65–81, 2018. [1](#)
- [23] Weixin Lu, Yao Zhou, Guowei Wan, Shenhua Hou, and Shiyu Song. L3-net: Towards learning based lidar localization for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6389–6398, 2019. [1](#)
- [24] Antonio Montanaro, Diego Valsesia, and Enrico Magli. Rethinking the compositionality of point clouds through regularization in the hyperbolic space. *arXiv preprint arXiv:2209.10318*, 2022. [2](#), [3](#), [5](#)
- [25] Maximillian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. *Advances in neural information processing systems*, 30, 2017. [2](#)
- [26] Maximillian Nickel and Douwe Kiela. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. In *International Conference on Machine Learning*, pages 3779–3788. PMLR, 2018. [2](#)
- [27] Julian Nubert, Shehryar Khattak, and Marco Hutter. Self-supervised learning of lidar odometry for robotic applications. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9601–9607. IEEE, 2021. [1](#)

- [28] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas 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*, pages 652–660, 2017. [2](#)
- [29] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. [2](#), [4](#)
- [30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. [3](#), [5](#)
- [31] Rik Sarkar. Low distortion delaunay embedding of trees in hyperbolic plane. In *International Symposium on Graph Drawing*, pages 355–366. Springer, 2011. [2](#)
- [32] Tixiao Shan, Brendan Englot, Drew Meyers, Wei Wang, Carlo Ratti, and Daniela Rus. Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping. In *2020 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 5135–5142. IEEE, 2020. [1](#)
- [33] Yoli Shavit, Ron Ferens, and Yosi Keller. Learning multi-scene absolute pose regression with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2733–2742, 2021. [1](#), [2](#), [7](#)
- [34] Mikaela Angelina Uy and Gim Hee Lee. Pointnetvlad: Deep point cloud based retrieval for large-scale place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4470–4479, 2018. [1](#), [2](#), [7](#), [8](#)
- [35] Florian Walch, Caner Hazirbas, Laura Leal-Taixe, Torsten Sattler, Sebastian Hilsenbeck, and Daniel Cremers. Image-based localization using lstms for structured feature correlation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 627–637, 2017. [7](#)
- [36] Bing Wang, Changhao Chen, Chris Xiaoxuan Lu, Peijun Zhao, Niki Trigoni, and Andrew Markham. Atloc: Attention guided camera localization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10393–10401, 2020. [1](#), [2](#), [7](#)
- [37] Han Wang, Chen Wang, Chun-Lin Chen, and Lihua Xie. F-loam: Fast lidar odometry and mapping. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4390–4396. IEEE, 2021. [1](#)
- [38] Wei Wang, Bing Wang, Peijun Zhao, Changhao Chen, Ronald Clark, Bo Yang, Andrew Markham, and Niki Trigoni. Pointloc: Deep pose regressor for lidar point cloud localization. *IEEE Sensors Journal*, 22(1):959–968, 2021. [1](#), [2](#), [4](#), [6](#), [7](#), [8](#)
- [39] Yue Wang and Justin M. Solomon. Deep closest point: Learning representations for point cloud registration. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. [2](#), [4](#)
- [40] Yue Wang and Justin M Solomon. Deep closest point: Learning representations for point cloud registration. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3523–3532, 2019. [7](#), [8](#)
- [41] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5):1–12, 2019. [2](#)
- [42] Jianyun Xu, Ruixiang Zhang, Jian Dou, Yushi Zhu, Jie Sun, and Shiliang Pu. Rpvnet: A deep and efficient range-point-voxel fusion network for lidar point cloud segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16024–16033, 2021. [3](#)
- [43] Wei Xu and Fu Zhang. Fast-lío: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter. *IEEE Robotics and Automation Letters*, 6(2):3317–3324, 2021. [1](#)
- [44] Xu Yan, Jiantao Gao, Chaoda Zheng, Chao Zheng, Ruimao Zhang, Shuguang Cui, and Zhen Li. 2dpass: 2d priors assisted semantic segmentation on lidar point clouds. In *European Conference on Computer Vision*, pages 677–695. Springer, 2022. [3](#)
- [45] Shangshu Yu, Cheng Wang, Chenglu Wen, Ming Cheng, Minghao Liu, Zhihong Zhang, and Xin Li. Lidar-based localization using universal encoding and memory-aware regression. *Pattern Recognition*, 128:108685, 2022. [2](#), [6](#), [7](#)
- [46] Ji Zhang and Sanjiv Singh. Low-drift and real-time lidar odometry and mapping. *Autonomous Robots*, 41(2):401–416, 2017. [1](#), [2](#)
- [47] Shichao Zhu, Shirui Pan, Chuan Zhou, Jia Wu, Yanan Cao, and Bin Wang. Graph geometry interaction learning. *Advances in Neural Information Processing Systems*, 33:7548–7558, 2020. [2](#), [5](#)