

# Learning Off-Road Terrain Traversability with Self-Supervisions Only

Junwon Seo, Sungdae Sim, and Inwook Shim

**Abstract**—Estimating the traversability of terrain should be reliable and accurate in diverse conditions for autonomous driving in off-road environments. However, learning-based approaches often yield unreliable results when confronted with unfamiliar contexts, and it is challenging to obtain manual annotations frequently for new circumstances. In this paper, we introduce a method for learning traversability from images that utilizes only self-supervision and no manual labels, enabling it to easily learn traversability in new circumstances. To this end, we first generate self-supervised traversability labels from past driving trajectories by labeling regions traversed by the vehicle as highly traversable. Using the self-supervised labels, we then train a neural network that identifies terrains that are safe to traverse from an image using a one-class classification algorithm. Additionally, we supplement the limitations of self-supervised labels by incorporating methods of self-supervised learning of visual representations. To conduct a comprehensive evaluation, we collect data in a variety of driving environments and perceptual conditions and show that our method produces reliable estimations in various environments. In addition, the experimental results validate that our method outperforms other self-supervised traversability estimation methods and achieves comparable performances with supervised learning methods trained on manually labeled data.

**Index Terms**—Semantic scene understanding, deep learning for visual perception, vision-based navigation, field robots.

## I. INTRODUCTION

RECENT advancements in visual perception enabled the success of fast-moving autonomous off-road vehicles. Estimating the traversability of the terrain with visual sensors is a crucial component of off-road driving. Numerous studies have made significant improvements in traversability estimation using large-scale datasets with human annotations and RGB images that provide semantically rich information about complex environments [1], [2]. However, the datasets

Manuscript received 7 February 2023; accepted 29 May 2023. Date of publication 8 June 2023; date of current version 23 June 2023. This letter was recommended for publication by Associate Editor D. Cattaneo and Editor M. Vincze upon evaluation of the reviewers' comments. This work was supported in part by the Agency For Defense Development grant funded by Korean Government in 2023, in part by the Korea Institute for Advancement of Technology (KIAT) grant funded by Korea Government (MOTIE) under Grant P0020536, HRD Program for Industrial Innovation, and in part by the National Research Foundation of Korea (NRF) grant funded by the Korea Government (MSIT) under Grant NRF-2022R1F1A1073505. (Corresponding author: Inwook Shim.)

Junwon Seo and Sungdae Sim are with the Agency for Defense Development, Daejeon 34186, Republic of Korea {junwon.vision, sungdae.sim}@gmail.com

Inwook Shim was with the Agency for Defense Development, Daejeon 34186, Republic of Korea. He is now with the Department of Smart Mobility Engineering, Inha University, Republic of Korea iwshim@inha.ac.kr

The multimedia material is available at <https://bit.ly/3YdKanw>.

Digital Object Identifier (DOI): 10.1109/LRA.2023.3284356



Fig. 1. We present a traversability estimation method that can be trained without human annotation in various environments. *Top*: Self-supervised traversability data gathered on diverse environments. There exists a large variance in visual appearances. *Bottom*: Traversability estimation results of the model learned solely with the self-supervised traversability labels.

only contain observations for a specific and limited context, resulting in unreliable estimates for unobserved conditions.

To successfully adapt to new circumstances, frequent manual annotation is required, which is not only unsustainable but also erroneous. Due to the high cost of the data-labeling procedure, obtaining sufficient labeled data regarding the various environments would be challenging. The labels produced by human experts often provide inadequate information for learning traversability in complex environments since the ground truth regarding traversable regions can not be clearly defined in off-road environments. In addition, the domain-specific annotations would lose their relevance in unfamiliar environments. For instance, various conditions, including places, seasons, weather, lighting, and camera settings, can significantly affect the visual appearance of an outdoor environment and the performance of estimations. Consequently, the vehicle cannot accurately predict traversability from images in a variety of situations if only static and constrained datasets are utilized.

While it is impractical to manually annotate images of every single environment, labels on traversable regions can be automatically generated by exploiting the vehicle trajectories in a self-supervised fashion [3]–[8]. Various works present the self-supervised approaches to learning traversability, which leverage self-supervised traversability labels instead of human-provided annotations [3]–[9]. However, they focus primarily on traversal cost analysis or terrain categorization in confined contexts rather than identifying traversable regions reliably in diverse environments. In order for a vehicle to operate

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

successfully and sustainably in a range of environments, it would be desirable that traversability is learned solely by utilizing the self-supervised traversability data.

Nonetheless, learning traversability from self-supervised traversability data is challenging due to the following reasons. First, since the vehicle has no experience traversing non-traversable regions, no labels for non-traversable regions can be obtained. While supervised learning-based methods [10] learn to differentiate between regions with distinct labels, the self-supervised traversability data only contains labels for a single class. It lacks supervision for discriminating between traversable and non-traversable regions, leading to overconfident predictions [3]. Second, the self-supervised traversability label is incomplete. Only a small portion of the traversable regions are labeled by trajectories, leaving the remainder unlabeled, and some of the labels may be inaccurate due to occlusions of trajectories [8]. Lastly, while some methods leverage one-class classification methods for learning with the self-supervised label, they not only fail to produce a dependable prediction for off-road images but also do not conduct experiments in various environments.

In this work, we propose a self-supervised traversability estimation method that learns traversability only from self-supervisions without explicit labels. We present an automated labeling process that can produce reliable self-supervised traversability labels on images by utilizing past vehicle trajectories. With the self-supervised labels, our algorithm learns traversability in off-road environments with complex distributions by leveraging Positive-Unlabeled(PU) [11] learning method and the 2D normalizing flow [12]. Moreover, to complement for insufficient supervision of the self-supervised labels, we employ approaches for self-supervised learning of visual representations to obtain discriminative representations from images [13].

To demonstrate the efficacy of our method, we collect large-scale driving data under a variety of conditions, including terrain types, places, weather, seasons, and lighting conditions. We conduct extensive experiments using our dataset along with the public dataset, RELIS-3D [2]. Our comprehensive quantitative and qualitative evaluations demonstrate that our method can effectively learn traversability in a wide range of unstructured and unknown environments.

## II. RELATED WORKS

### A. Traversability Estimation

With developments in learning traversability, autonomous driving has made significant progress in urban and off-road environments [14]. Most early works on estimating traversability concentrate on analyzing simple geometric and visual features [15], [16]. With the development of deep neural networks, semantic segmentation is widely utilized to classify terrains into predefined terrain classes leveraging large datasets [17]. In addition, numerous approaches have been developed to identify traversable regions in unstructured environments [17], [18]. Fully convolutional networks for image segmentation [10] have significantly improved the off-road traversability estimation performance since images

contain semantically rich and dense information about the off-road environments [19]. However, such methods heavily rely on training data, which leads to incorrect estimations when confronted with data from distributions not included in the training data [20]. The supervised learning methods may not generalize well to changing and unknown environmental circumstances [4]. For the widespread deployment of autonomous vehicles off-road, where the likelihood of encountering an unfamiliar context is considerable, the model should be capable of working reliably in various environments.

### B. Self-Supervised Learning of Traversability

For reliable traversability estimation in a wide range of environments, self-supervised approaches are proposed, which exploit a vehicle's driving experience to learn the traversability of a terrain [3], [4], [9], [21], [22]. These methods enable automated procedures to self-label visual data for learning traversability. For example, measurements from proprioceptive sensors are used to assess the traversal cost of terrain or to classify terrains. However, they either rely on manual labels or are oblivious to the fact that estimations in unseen environments can be unreliable. As labels on non-traversable regions cannot be acquired via self-labeling, the estimations are prone to over-confident predictions, which might lead to navigational failure [5]. Consequently, identifying traversable regions with reliability is a crucial problem. While our previous work [8] has shown that traversability can be learned using point clouds, the method for identifying the traversable region using images has not been exhaustively examined in a variety of environments.

One-class classification algorithms can be employed to distinguish traversable and non-traversable regions [5]–[7], [23]. For example, normalizing flow [24] shows a great performance for traversability classification on multi-modal images [5]. However, it freezes the feature encoder after pretraining and encodes local patches, which do not incorporate global scene information. The features would simply capture low-level meanings, such as the color and texture of terrains, without their semantic information required to discriminate between traversable and non-traversable terrain [25]. In addition, an autoencoder is used to identify high-risk terrains based on the reconstruction error [6]. The simple autoencoder-based reconstruction focuses on low-frequency details, resulting in the estimation that high-frequency details are simply classified as non-traversable. The autoencoder is also known to generalize well to unseen data, resulting in large false-positive predictions in which well-reconstructed non-traversable regions are assigned a low anomaly score [26]. The one-class classification algorithm should be capable of extracting more discriminative features from images with self-supervised labels in order to reliably identify traversable regions in diverse off-road contexts.

### C. Self-Supervised Learning of Visual Representations

Instead of predicting human-annotated labels, approaches on self-supervised visual pre-training learn without labels by solving pretext tasks. The pretext tasks include the reconstruction

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

of inputs, instance discriminations, and clustering with pseudo-labels [27], [28]. Most state-of-the-art methods are contrastive learning, in which the network is trained to attract positive sample pairs and repel negative sample pairs [13]. Due to the incomplete labeling of the self-supervised traversability data, the acquisition of highly discriminative features from the data is challenging. By complementing the short supervision of self-supervised traversability data with self-supervised learning of visual representations, the visual representation for learning traversability could be more discriminative.

### III. METHODS

Our goal is to train a network that can successfully embed the complex data distribution of environments, allowing for precise and reliable estimation in a wide range of unstructured contexts. Given image  $\mathbf{x}$ , we generate a self-supervised label  $\hat{\mathbf{y}}$  that does not require manual annotation. Then, we learn a model that estimates a pixel-wise traversability  $\mathbf{y}_i = f(\mathbf{x}_i)$ , where  $\mathbf{x}_i$  is an image pixel and  $\mathbf{y}_i$  is terrain traversability representing whether or not the terrain is traversable.

#### A. Self-Supervised Traversability Label

From images gathered while driving, the self-supervised label  $\hat{\mathbf{y}}$  is generated by an automated procedure, as illustrated in Fig. 2. Since the regions traversed by a vehicle during data collection can be considered safe to traverse, we can designate such regions as traversable. The wheel-terrain contact points are calculated using the trajectories recovered by the SLAM [29]. The trajectories of horizon  $\alpha$  from time  $t_i$  are converted to contact points, denoted as  $\mathbf{T}(t_i, t_{i+\alpha})$ .

Prior to the labeling, the contact points are filtered to eliminate false-positive labels. Since the past trajectory is projected onto the 2D images, parts of the contact points can be occluded due to obstacles or rotations of the vehicle. Without filtering, the obstacles would be labeled as traversable, leading to a large number of false positives in estimations. Although these false-positive labels can be avoided by shortening the horizons, this shortens supervision for learning. The contact points are therefore filtered using LiDAR points captured simultaneously with the images. Similar to an occlusion filtering algorithm [6], a contact point is filtered as occluded if it has a longer radial distance than the nearest LiDAR points in spherical coordinates.

However, numerous undesirable noises exist in LiDAR point clouds acquired under a variety of off-road conditions (e.g., dust, rain, snow). The noises can be regarded as an obstacle during filtering, which may hinder the effectiveness of filtering. For robust labeling in a variety of unstructured environments, unsupervised LiDAR denoising [30] is performed prior to occlusion filtering. With the denoised point cloud, the contact points are filtered, denoted as  $\mathbf{T}'(t_i, t_{i+\alpha})$ .

Finally, the contact points are projected into the camera coordinates to generate the self-supervised label of the PU type by the following equations:

$$\hat{\mathbf{y}} = \mathbf{K} \cdot [\mathbf{R}|\mathbf{t}] \cdot \mathbf{T}'(t_i, t_{i+\alpha}) \quad (1)$$

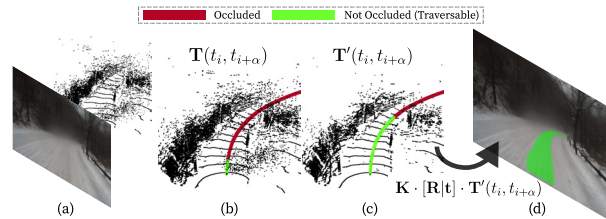


Fig. 2. Overview of our proposed automated procedure for self-supervised traversability label generation. (a) On captured data, the wheel-contact points are considered traversable. The occluded points are filtered before projection to image coordinates to eliminate false positive labels. (b) Noises in LiDAR points are regarded as obstacles, leading to erroneous filtering. (c) Unsupervised LiDAR denoising is performed to improve the efficacy of filtering in off-road conditions. (d) Trajectories are then projected to image coordinates and labeled as traversable.

where  $\mathbf{K}$  and  $[\mathbf{R}|\mathbf{t}]$  represent the intrinsic camera calibration matrix and the world-to-camera transformation matrix respectively. On the image coordinates, pixels between the left and right wheel-terrain contact points are labeled as *positive*, while all other pixels are left *unlabeled*. Note that the label consists of a relatively small number of positive pixels and unlabeled pixels are a combination of traversable and non-traversable regions.

#### B. Learning Traversability

In this section, we propose a method for learning traversability with self-supervised labels only. The overall architecture of our learning method is illustrated in Fig. 3.

1) *One-Class Classification*: The image is forwarded to a feature encoder for image segmentation, denoted as  $f$ , that maps images  $\mathbf{x} \in \mathbb{R}^{h \times w \times 3}$  into pixel-wise features  $\mathbf{z} \in \mathbb{R}^{h \times w \times d}$ . Our backbone encoder is PSPNet [10] which can capture global context through the pyramid pooling module.

The feature embedding space can be trained to minimize the volume of a positive-data-enclosing hypersphere. Then, the similarity metric between a feature and the center of the hypersphere,  $p(\mathbf{z}_i) \in [0, 1]$ , can be used to determine the traversability of the  $i^{\text{th}}$  pixel. The simple one-class classification loss can be used for positive pixels:

$$\mathcal{L}_{\text{OCC}} = 1 - p(\mathbf{z}_i). \quad (2)$$

However, the representations of positive pixels have high intra-class variation as there exist various representations of traversable regions in off-road environments. The loss function Eq. (2) assumes a single center and pushes dissimilar features towards a single center, thereby diminishing the discriminative power of the representations. Not only is the solution susceptible to a hypersphere-collapse solution, in which the majority of data can be trivially mapped to the hypersphere center, but it is also incapable of effectively capturing multimodal distributions.

Normalizing flow [24] can be used to project complex distributions of features to a simple distribution while avoiding a trivial solution. The flow model, denoted as  $g$ , transforms a pixel-wise feature  $\mathbf{z}_i \in \mathbb{R}^d$  into a flow feature  $\mathbf{z}_i^f \in \mathbb{R}^{d^f}$  with a tractable distribution using a bijective invertible mapping. We

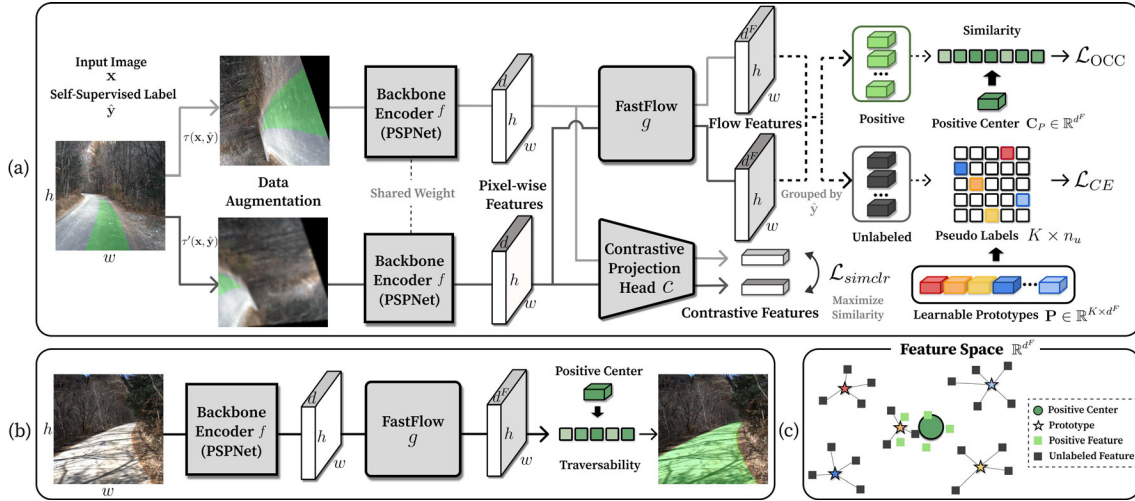


Fig. 3. High-level structure of the proposed method. (a) Training. (b) Inference. (c) Illustration of the embedding space of flow features at training. The feature space is learned with the clustering pretext task, enhancing the discriminative power of features for the one-class classification.

adopt the 2D normalizing flow model with affine coupling layers, Fastflow [12], which produces more accurate features for segmentation. The likelihood of a flow feature can be simply defined as  $p(\mathbf{z}_i^F) = \mathbf{z}_i^F \cdot \mathbf{C}_p$ , which represents cosine similarity with the hypersphere center of positive data,  $\mathbf{C}_p \in \mathbb{R}^{d^F}$ . Then, the traversability of an image pixel can be easily calculated using the change of the variable formula:

$$\log p(\mathbf{z}_i) = \log p(\mathbf{z}_i^F) + \log \left| \det \left( \frac{\partial \mathbf{z}_i^F}{\partial \mathbf{z}_i} \right) \right|, \quad (3)$$

where the determinant of the Jacobian  $\frac{\partial \mathbf{z}_i^F}{\partial \mathbf{z}_i}$  can be calculated with affine coupling layers [24]. Intuitively, the Jacobian penalizes the trivial solutions that have constant mappings.

However, the normalizing flow model cannot be end-to-end trained with a feature encoder because it would generate a trivial backbone encoder while preventing the trivial flow model. In addition, the network does not utilize unlabeled data. The network would be overfitted to the distributions of traversable regions, limiting its ability to discriminate between traversable and non-traversable regions.

2) *Self-Supervised Clustering with Unlabeled Data*: We use unlabeled data in a self-supervised manner to train the network end-to-end so that the network can learn better embeddings while avoiding trivial solutions. Motivated by the clustering-based self-supervised learning of visual representations [27], [28], our methodology solves the clustering pretext task. The flow features of unlabeled pixels are jointly self-labeled and clustered with a set of  $K$  learnable prototypes,  $\mathbf{P} \in \mathbb{R}^{K \times d^F}$ , which functions as cluster centers.

By taking the softmax of the similarity between prototypes and unlabeled features, the posterior distribution of unlabeled pixels to prototypes,  $\mathbf{Q} \in \mathbb{R}^{K \times n_u}$ , is computed, where  $n_u$  is the number of unlabeled pixels within a batch. Then, soft cluster assignment  $\mathbf{A} \in \mathbb{R}^{K \times n_u}$  from features to prototypes is computed by optimizing the following equation with an equipartition constraint:

$$\max_{\mathbf{A}} \text{Tr}(\mathbf{A}^T \mathbf{Q}) \quad \text{s.t.} \quad \mathbf{A} \cdot \mathbf{1}^{n_u} = \frac{n_u}{K} \cdot \mathbf{1}^K \quad (4)$$

The constraints ensure that the prototypes equally partition the assignments, thereby preventing trivial solutions in which features are collapsed into equal representations [27]. This optimization problem can be efficiently solved with a few iterations of the *Sinkhorn-Knopp* algorithm [31]. By minimizing the cross-entropy loss between the posterior distribution and the optimized cluster assignment, the features and prototypes are simultaneously updated:

$$\mathcal{L}_{CE} = -\frac{1}{n_u} \sum_k^K \sum_j^{n_u} \mathbf{A}_{kj} \log(\mathbf{Q}_{kj}). \quad (5)$$

However, the learned representations are still insufficient because the supervision of the self-supervised labels is restricted to a small portion of the entire traversable regions. In contrast to supervised learning methods, which are trained to explicitly distinguish traversable and non-traversable regions with full labels, the unsupervised clustering objective may attempt to learn simplistic features.

3) *Self-Supervised Contrastive Learning*: To supplement the representation power of backbone features, the encoder  $f$  is simultaneously learned with the contrastive pretext task of self-supervised visual representation learning [13] alongside other objectives in order to generate a more powerful visual representation for a given data distribution.

Given  $N$  images in a minibatch, two random *views* are generated for each image as a positive pair by random data augmentations,  $\tau$  and  $\tau'$ . The remaining  $2N - 2$  augmented views of images within a minibatch are regarded as negative pairs. The augmentation comprises low-level image transformations. The contrastive feature of each view,  $\mathbf{z}^c \in \mathbb{R}^{d^c}$ , is produced by forwarding the pixel-wise features into the contrastive projection head, denoted as  $c$ .

Then, we minimize the contrastive loss function for each data with the cosine similarity as follows:

$$\mathcal{L}_{simclr} = -\log \frac{\exp\left(\frac{\mathbf{z}^c \cdot \mathbf{z}_+^c}{\lambda}\right)}{\exp\left(\frac{\mathbf{z}^c \cdot \mathbf{z}_+^c}{\lambda}\right) + \sum_{\mathbf{z}_-^c}^{2N-2} \exp\left(\frac{\mathbf{z}^c \cdot \mathbf{z}_-^c}{\lambda}\right)}. \quad (6)$$

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

$\lambda$  is a temperature hyperparameter, and  $\mathbf{z}_+^C$  and  $\mathbf{z}_-^C$  denotes features of positive and negative pairs, respectively. By minimizing the loss, features of positive pairs are pulled while pushing those of negative pairs away. In a self-supervised manner, it regularizes the features to be more semantically meaningful and discriminative for traversability estimation.

#### IV. EXPERIMENTS

In this section, we validate that our self-supervised traversability estimation method can effectively learn traversability in a wide range of environments. We first describe the dataset used for the evaluation, followed by the experimental setup as well as implementation details. Then, we present both quantitative and qualitative results of our traversability estimation method. Lastly, we present detailed ablation studies demonstrating that our self-supervised traversability estimation method is capable of learning traversability in a variety of environments and under appearance changes without human annotations.

##### A. Datasets

1) *Driving Data Under Adverse Conditions*: We collected driving data in a variety of environments using our platform [32] equipped with an RGB camera, VLP-32 LiDAR. It comprises about 20,000 images gathered under a wide range of conditions, including varying places, seasons, weather, lighting, terrain types, obstacles, and lens conditions.

According to the characteristics of the environments, we divide our data into five categories: *paved*, *unpaved*, *snowy*, *rainy*, and *night*. The *paved* contains images of urban and rural areas with paved roads. The *unpaved* includes images acquired while driving on unpaved off-roads, where obstacles, dust, and smoke are captured on camera and the drivable regions are less clear. The *snowy* and *rainy* categories consist of images taken in the context of snow and rain, with snowed surfaces and puddles, as well as frost and raindrops on the lens. The *night* is composed of images obtained in dark areas with headlights on. For evaluation, 300 images per category are manually annotated by an expert. They are chosen from a subset of sequential images and excluded from training.

2) *RELLIS-3D*: We also present experimental results using the publicly available RELLIS-3D off-road dataset [2], which contains RGB camera images with pixel-level annotation. The self-supervised traversability labels are generated from the raw data using LiDAR-based SLAM [29]. The data is divided into a training set containing 4,827 images and a validation set with the remaining images. Although the annotation does not indicate which points are traversable, we define the *grass*, *dirt*, *asphalt*, *concrete*, and *mud* classes as traversable and the *tree*, *pole*, *vehicle*, *object*, *person*, *fence*, *barrier*, *rubble*, and *bush* classes as non-traversable.

##### B. Experimental Setup

First, we demonstrate that the model learned with our self-supervised traversability estimation methods is more effective than models trained in a fully supervised manner using datasets

with human annotations. For the comparison, the PSPNet is trained in a supervised manner with the following datasets: *KITTI* road detection, *RELLIS-3D*, and our labeled dataset of outdoor driving scenes (*Outdoor*). For the outdoor dataset, about 1K images of the paved and unpaved categories are randomly selected and manually labeled for supervised learning. The model is also trained with the aforementioned three datasets altogether (*ALL*).

Then, we show that our method yields higher performances than other self-supervised traversability estimation algorithms. Our approach is compared with the method that uses normalizing flow [24] on top of the pre-trained backbone (*Real-NVP*) [5] and the method based on reconstruction-based anomaly detection with autoencoder (*AE Based*) [6].

##### C. Evaluation Metrics

The Area Under the Receiver Operating Characteristic (AUROC) is used to quantitatively evaluate the methods. AUROC quantifies the likelihood of a positive sample having a higher normal score than a negative sample, and therefore evaluates the one-class classification algorithms regardless of the threshold. In addition, we evaluate our methods using the standard evaluation metrics of the KITTI road detection system. Maximum F1-measure (MaxF), average precision (AP), precision rate (PRE), recall rate (REC), false positive rate (FPR), and false negative rate (FNR) are the metrics included. Note that the four latter measures are obtained at the threshold of the maximum F1 measure.

##### D. Implementation Details

We use PSPNet with ResNet50 [33] as a backbone embedding network for every method for fair comparisons. We use the flow model with eight transformation blocks composed of affine coupling layers. The contrastive head consists of adaptive average pooling and two MLP layers with a ReLU in the middle. Both the flow model and the contrastive head produce 128 dimensional vectors that are  $l_2$  normalized.

Our models are trained for 60 epochs with a mini-batch size of 64, using the Adam optimizer with a learning rate of  $1e^{-3}$  and a polynomial learning rate decay. The sum of the means of the three losses (Eq. 2,5, and 6) is used as the objective of the optimization. The random data augmentation pipeline includes  $256 \times 256$  pixel random cropping, flipping, random color jittering, random gray-scale conversion, gaussian blurring, rotation, and random perspective transformation. The number of learnable prototypes is set to 256 and they are randomly initialized by the normal distribution. We execute three iterations of the Sinkhorn-Knopp algorithm and set the temperature parameters  $\lambda$  as 0.1. For data labeling, we set horizons  $\alpha$  for self-supervised labels in Section III-A to 100, indicating that we utilized trajectories 10 seconds ahead of the image acquisition.

##### E. Experimental Results

The Table I and Fig. 4 provide the quantitative and qualitative results for our dataset. In most of the categories,

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

TABLE I  
QUANTITATIVE RESULTS OF TRAVERSABILITY ESTIMATION METHODS FOR OUR DATASET COLLECTED UNDER VARIOUS CONDITIONS.

	Paved			Unpaved			Snowy			Rainy			Night		
	AUROC	MaxF	AvgPrec	AUROC	MaxF	AvgPrec	AUROC	MaxF	AvgPrec	AUROC	MaxF	AvgPrec	AUROC	MaxF	AvgPrec
PSPNet (KITTI [14])	0.8732	0.7820	0.8570	0.9100	0.8039	0.8486	0.6669	0.7467	0.8245	0.9095	0.8034	0.8605	0.9323	0.8778	0.9256
PSPNet (RELLIS-3D [2])	0.9143	0.7728	0.7562	0.9401	0.8398	0.7690	0.6187	0.7670	0.7223	0.9143	0.7728	0.7562	0.9669	0.9168	0.9059
PSPNet (Outdoor)	0.9861	<b>0.9414</b>	<b>0.9386</b>	0.9729	0.8829	0.8828	0.7592	0.8593	0.8945	0.9757	0.8899	0.9094	0.9673	0.9139	0.9372
PSPNet (ALL)	<b>0.9870</b>	0.9396	0.9383	0.9720	0.9051	0.8809	0.9459	0.8337	0.8454	0.9672	0.8839	0.9019	0.9565	0.9105	0.9338
Real-NVP [5]	0.5878	0.6101	0.4863	0.6352	0.5075	0.4313	0.5459	0.5245	0.3719	0.6712	0.5175	0.4878	0.5145	0.6286	0.5344
AE Based [6]	0.7923	0.7102	0.6905	0.8443	0.6795	0.6712	0.8120	0.6683	0.6544	0.8321	0.6600	0.6677	0.8609	0.7851	0.8235
<b>Ours</b>	0.9660	0.8888	0.9306	<b>0.9813</b>	<b>0.9143</b>	<b>0.9168</b>	<b>0.9770</b>	<b>0.8930</b>	<b>0.9052</b>	<b>0.9815</b>	<b>0.9044</b>	<b>0.9160</b>	<b>0.9750</b>	<b>0.9169</b>	<b>0.9387</b>

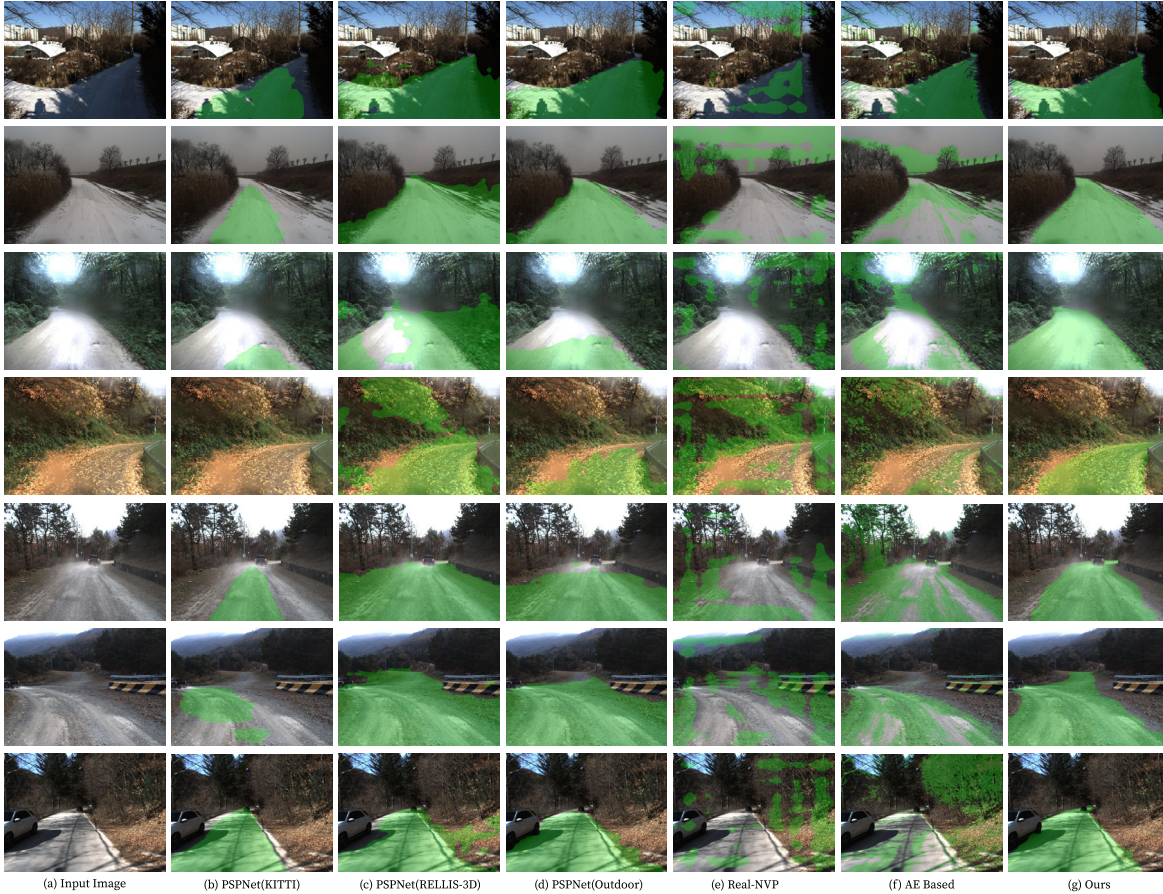


Fig. 4. Qualitative traversability estimation results for our dataset. The thresholds of the maximum F1 scores are used for the visualizations. The pixels where the estimated traversability exceeds the thresholds are colored green. More results are available in the multimedia material.

our method outperforms models trained with manual labels in a supervised way, implying that our methods can yield comparable or even greater results than a model trained with laborious manual annotations. In addition, combining multiple datasets for supervised learning does not seem to improve the traversability estimation performance on target distributions. These results demonstrate that distribution shifts have a significant impact on the performance of traversability estimation, suggesting the necessity of self-supervised traversability estimation methods for autonomous vehicles to operate effectively in widespread environments. Due to the fact that our method exploits self-supervised labels of the target distributions, the model can estimate traversability reliably in the presence of distributional shifts.

Our method shows a significant margin compared to other self-supervised methods based on one-class classification. Note that the methods based on autoencoders produce a large number of false positives, meaning that the simple reconstruction-based anomaly detections fail to distinguish non-traversable regions. Ours produces fewer false negative occurrences than others, indicating that it obtains discriminative features for identifying traversable regions in complex off-road environments.

The Table II and Fig. 5 illustrate results for RELLIS-3D. Our method shows better performance than others and even yields comparable performance with models overfit with manual annotations of the RELLIS-3D.

TABLE II  
QUANTITATIVE RESULTS FOR RELIS-3D.

	AUROC	MaxF	AvgPrec	PRE	REC	FPR	FNR
PSPNet (KITTI [14])	0.7565	0.7345	0.7431	0.6241	0.8922	0.6248	0.1078
PSPNet (RELLIS-3D [2])	<b>0.9814</b>	<b>0.9400</b>	<b>0.9457</b>	0.9378	0.9423	0.0702	0.0578
PSPNet (Outdoor)	0.8160	0.7860	0.7895	0.7104	0.8798	0.4030	0.1202
PSPNet (ALL)	0.9814	0.9396	0.9454	<b>0.9388</b>	0.9404	<b>0.0690</b>	0.0607
Real-NVP [5]	0.5625	0.7001	0.5710	0.5464	<b>0.9742</b>	0.9280	<b>0.0258</b>
AE Based [6]	0.7348	0.7437	0.7079	0.6250	0.9181	0.6323	0.0819
<b>Ours</b>	0.9036	0.8622	0.9164	0.8738	0.8508	0.1466	0.1492

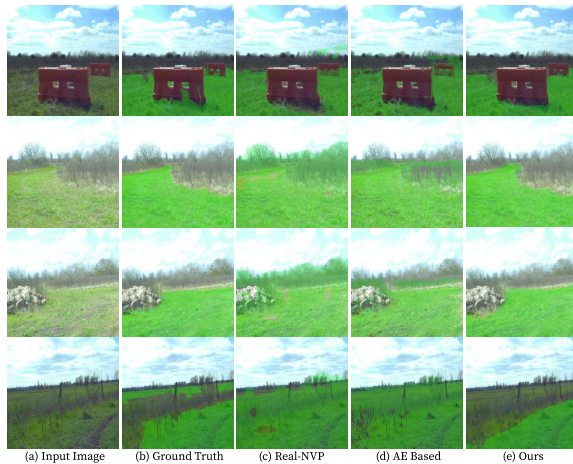


Fig. 5. Qualitative results for RELIS-3D. Note that the ground truths are sometimes inaccurate and labeled for terrain classification rather than traversability. For example, some pixels on obstacles are labeled as traversable due to thin overhanging vegetation, and distant grass across fences is also labeled as traversable. Instead of merely assessing terrain type, our method identifies traversable regions that are contextually significant.

### F. Ablation Studies

We present comprehensive ablation studies to examine the validity of each component of our methodology. We quantitatively verify the efficacy of the self-supervised traversability labeling, PU learning algorithm with the normalizing flow, self-supervised clustering, and self-supervised contrastive learning. The ablations are trained using data from all categories of our dataset. The results are shown in Table III and Fig. 6 shows the ROC curves of the results.

1) *Self-Supervised Labels*: First, to validate the efficacy of our self-supervised labeling algorithm, we compare the impacts of labels obtained without occlusion filtering, without LiDAR denoising, and with varying horizons. The labels generated without occlusion filtering result in a high FPR, indicating that non-navigable regions, such as obstacles, are incorrectly estimated as traversable due to unreliable labels. The labels created without LiDAR denoising lead to a higher FNR, implying that the model is trained with fewer supervisions because self-labels for traversable regions are misclassified as occluded due to LiDAR noise. Similarly, lowering the horizons of trajectories increased the FNR because it reduces the number of positive pixels of the labels.

2) *PU Learning*: Second, we replace the flow model with 1D flow (Real-NVP) and simple MLP layers to evaluate the efficacy of flow models. We observe that 2D normalizing flow produces better results compared to 1D flow. It confirms that

TABLE III  
QUANTITATIVE RESULTS OF THE ABLATION STUDIES.

	AUROC	MaxF	AvgPrec	PRE	REC	FPR	FNR
PSPNet (KITTI [14])	0.9092	0.8044	0.8504	0.8451	0.7674	0.0643	0.2326
PSPNet (RELLIS-3D [2])	0.9354	0.8126	0.7644	0.7521	0.8836	0.1331	0.1164
PSPNet (Outdoor)	0.9718	0.8765	0.8927	0.8532	0.9011	0.0708	0.0989
PSPNet (ALL)	0.9732	0.9013	0.8909	0.8806	<b>0.9230</b>	0.0572	<b>0.0770</b>
Real-NVP [5]	0.6539	0.5326	0.4480	0.3838	0.8698	0.6856	0.1302
AE Based [6]	0.7261	0.5847	0.5342	0.4779	0.7530	0.4039	0.2470
<b>Ours</b>	<b>0.9741</b>	<b>0.9082</b>	<b>0.9143</b>	<b>0.9014</b>	0.9151	<b>0.0456</b>	0.0849
1) Self-Supervised Labels							
w.o. Occlusion Filtering	0.9272	0.8144	0.8682	0.8201	0.8087	0.0871	0.1913
w.o. Denoising	0.8861	0.8019	0.8414	0.8679	0.7452	0.0557	0.2548
w. horizon $\alpha = 30$	0.9497	0.8839	0.9125	0.8846	0.8833	0.0530	0.1167
w. horizon $\alpha = 50$	0.9632	0.8861	0.90951	0.8841	0.8880	0.0535	0.1120
2) PU Learning							
w.o. Fastflow(w. Real-NVP)	0.9498	0.8690	0.8994	0.8852	0.8534	0.0544	0.1466
w.o. Fastflow(w. MLP)	0.5948	0.5236	0.5049	0.3555	0.9934	0.8843	0.0066
w.o. Unlabeled Pixels	0.5015	0.4954	0.3879	0.3293	1.0000	1.0000	0.0000
3) Self-Supervised Learning							
w.o. Contrastive Loss	0.9266	0.8449	0.8829	0.8831	0.8098	0.0526	0.1902
w. Reconstruction Loss	0.9348	0.8191	0.8684	0.8083	0.8301	0.0966	0.1699

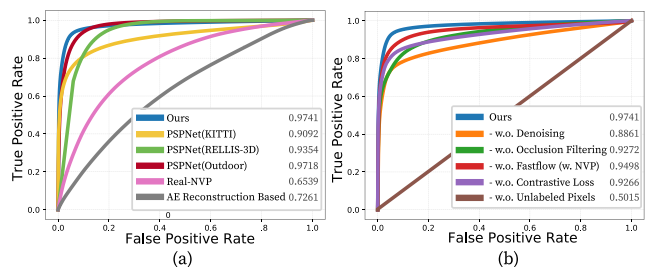


Fig. 6. ROC curves for (a) traversability estimation methods and (b) the ablation studies. The values indicate the AUROC of the results. Notably, our method yields low false positive rates compared to other methods, which is essential for safe navigation in unstructured environments.

using 2D normalizing flow is more effective for localizing traversable regions with pixel-wise features while avoiding a trivial solution. The model trained by replacing the flow model with MLPs results in low AUROC and high FPR, implying that the models produce a trivial solution without the normalizing flow. Also, the model is trained without clustering loss in order to highlight the effectiveness of using unlabeled data. Without the loss for unlabeled pixels, Eq. (5), the performance severely diminishes as AUROC approaches 0.5. It indicates that the use of unlabeled data is essential for avoiding trivial solutions and learning more discriminative features about environments.

3) *Self-Supervised Learning for Visual Representation*: Then, we verify the efficacy of our self-supervised contrastive learning. The models are trained without contrastive learning and with the reconstruction pretext task. The model trained without contrastive learning exhibits a low REC, denoting that the image features are less relevant for learning traversability, as the self-labels cannot provide guidance regarding the distinction between traversable and non-traversable regions. The performance of the model learned concurrently with the reconstruction pretext task is improved, but still inferior to the model with contrastive loss. The reason for this is that the naive reconstruction objective tends to focus on texture and color rather than semantic meanings.

## V. CONCLUSIONS

This paper introduces a self-supervised traversability estimation method that can learn traversability in varied environments without manual annotations. Using the self-supervised labels only, the network is trained to predict traversable regions using a one-class classification method. Self-supervised learning of visual representation is incorporated into the learning in order to improve the network. Extensive experiments demonstrate that the proposed method is capable of learning traversability more effectively than others. We believe that our method can be leveraged for the wider deployment of autonomous vehicles since it is capable of easily adapting to a variety of contexts by precisely embedding the data distribution of target environments.

Future work includes using labeled data with domain adaptation and semi-supervised learning. Also, we are investigating incremental learning and online learning for a more general model that can be used in various environments.

## REFERENCES

- [1] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [2] P. Jiang, P. Osteen, M. Wigness, and S. Saripalli, "Relis-3d dataset: Data, benchmarks and analysis," in *IEEE international conference on robotics and automation (ICRA)*, 2021, pp. 1110–1116.
- [3] L. Wellhausen, A. Dosovitskiy, R. Ranftl, K. Walas, C. Cadena, and M. Hutter, "Where should i walk? predicting terrain properties from images via self-supervised learning," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1509–1516, 2019.
- [4] J. Zürn, W. Burgard, and A. Valada, "Self-supervised visual terrain classification from unsupervised acoustic feature learning," *IEEE Transactions on Robotics*, vol. 37, no. 2, pp. 466–481, 2020.
- [5] L. Wellhausen, R. Ranftl, and M. Hutter, "Safe robot navigation via multi-modal anomaly detection," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1326–1333, 2020.
- [6] R. Schmid, D. Atha, F. Schöller, S. Dey, S. Fakoorian, K. Otsu, B. Ridge, M. Bjelonic, L. Wellhausen, M. Hutter *et al.*, "Self-supervised traversability prediction by learning to reconstruct safe terrain," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 12419–12425.
- [7] N. Hirose, A. Sadeghian, M. Vázquez, P. Goebel, and S. Savarese, "Gonet: A semi-supervised deep learning approach for traversability estimation," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 3044–3051.
- [8] J. Seo, T. Kim, K. Kwak, J. Min, and I. Shim, "Scate: A scalable framework for self-supervised traversability estimation in unstructured environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 888–895, 2023.
- [9] H. Dahlkamp, A. Kaehler, D. Stavens, S. Thrun, and G. Bradski, "Self-supervised monocular road detection in desert terrain," in *Proceedings of Robotics: Science and Systems (RSS)*, Philadelphia, USA, August 2006.
- [10] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2017, pp. 2881–2890.
- [11] J. Bekker and J. Davis, "Learning from positive and unlabeled data: A survey," *Machine Learning*, vol. 109, no. 4, pp. 719–760, 2020.
- [12] J. Yu, Y. Zheng, X. Wang, W. Li, Y. Wu, R. Zhao, and L. Wu, "Fastflow: Unsupervised anomaly detection and localization via 2d normalizing flows," *arXiv preprint arXiv:2111.07677*, 2021.
- [13] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International Conference on Machine Learning (ICML)*. PMLR, 2020, pp. 1597–1607.
- [14] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2012, pp. 3354–3361.
- [15] M. Wermelinger, P. Fankhauser, R. Diethelm, P. Krüsi, R. Siegwart, and M. Hutter, "Navigation planning for legged robots in challenging terrain," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 1184–1189.
- [16] J. Sock, J. Kim, J. Min, and K. Kwak, "Probabilistic traversability map generation using 3d-lidar and camera," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 5631–5637.
- [17] P. Borges, T. Peynot, S. Liang, B. Arain, M. Wildie, M. Minareci, S. Lichman, G. Samvedi, I. Sa, N. Hudson *et al.*, "A survey on terrain traversability analysis for autonomous ground vehicles: Methods, sensors, and challenges," *Field Robotics*, vol. 2, no. 1, pp. 1567–1627, 2022.
- [18] A. Shaban, X. Meng, J. Lee, B. Boots, and D. Fox, "Semantic terrain classification for off-road autonomous driving," in *Conference on Robot Learning (CoRL)*. PMLR, 2022, pp. 619–629.
- [19] T. Guan, D. Kothandaraman, R. Chandra, A. J. Sathymoorthy, K. Weerakoon, and D. Manocha, "Ga-nav: Efficient terrain segmentation for robot navigation in unstructured outdoor environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8138–8145, 2022.
- [20] D. Barnes, W. Maddern, and I. Posner, "Find your own way: Weakly-supervised segmentation of path proposals for urban autonomy," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 203–210.
- [21] D. Kim, J. Sun, S. M. Oh, J. M. Rehg, and A. F. Bobick, "Traversability classification using unsupervised on-line visual learning for outdoor robot navigation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2006, pp. 518–525.
- [22] M. Guaman Castro, S. Triest, W. Wang, J. M. Gregory, F. Sanchez, J. G. Rogers III, and S. Scherer, "How does it feel? self-supervised costmap learning for off-road vehicle traversability," in *2023 International Conference on Robotics and Automation (ICRA)*. IEEE, 2023.
- [23] T. Ji, S. T. Vuppala, G. Chowdhary, and K. Driggs-Campbell, "Multi-modal anomaly detection for unstructured and uncertain environments," in *Conference on Robot Learning (CoRL)*, vol. 155. PMLR, 16–18 Nov 2021, pp. 1443–1455.
- [24] L. Dinh, J. Sohl-Dickstein, and S. Bengio, "Density estimation using real NVP," in *International Conference on Learning Representations (ICLR)*, 2017.
- [25] D. Gudovskiy, S. Ishizaka, and K. Kozuka, "Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2022, pp. 98–107.
- [26] D. Gong, L. Liu, V. Le, B. Saha, M. R. Mansour, S. Venkatesh, and A. v. d. Hengel, "Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (CVPR)*, 2019, pp. 1705–1714.
- [27] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 9912–9924, 2020.
- [28] Y. M. Asano, C. Rupprecht, and A. Vedaldi, "Self-labelling via simultaneous clustering and representation learning," in *International Conference on Learning Representations (ICLR)*, 2020.
- [29] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, "Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping," in *2020 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2020, pp. 5135–5142.
- [30] G. Bae, B. Kim, S. Ahn, J. Min, and I. Shim, "Slide: Self-supervised lidar de-snowing through reconstruction difficulty," in *European Conference on Computer Vision (ECCV)*, 2022.
- [31] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport," in *Advances in Neural Information Processing Systems (NeurIPS)*, C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger, Eds., vol. 26, 2013, pp. 2292–2300.
- [32] S. Sim, J. Min, S. Ahn, J. Lee, J. S. Lee, G. Bae, B. Kim, J. Seo, and T. S. Choe, "Build a multi-sensor dataset for autonomous driving in adverse weather conditions," *The Journal of Korea Robotics Society*, vol. 17, no. 3, pp. 245–254, 2022.
- [33] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2016, pp. 770–778.