

The Devil is in the Details: Simple Remedies for Image-to-LiDAR Representation Learning

Wonjun Jo¹, Kwon Byung-Ki², Kim Ji-Yeon³, Hawook Jeong⁴, Kyungdon Joo⁵, and Tae-Hyun Oh^{1,2,3,6}

¹ Department of Electrical Engineering, POSTECH, South Korea

² Graduate School of AI, POSTECH, South Korea

³ Department of Convergence IT Engineering, POSTECH, South Korea

⁴ RideFlux Inc., South Korea

⁵ Artificial Intelligence Graduate School, UNIST, South Korea

⁶ Institute for Convergence Research and Education in Advanced Technology, Yonsei University, South Korea

{jo1jun, byungki.kwon, jiyeon.kim, taehyun}@postech.ac.kr,
hawook@rideflux.com, kyungdon@unist.ac.kr

Abstract. LiDAR is a crucial sensor in autonomous driving, commonly used alongside cameras. By exploiting this camera-LiDAR setup and recent advances in image representation learning, prior studies have shown the promising potential of image-to-LiDAR distillation. These prior arts focus on the designs of their own losses to effectively distill the pre-trained 2D image representations into a 3D model. However, the other parts of the designs have been surprisingly unexplored. We find that fundamental design elements, *e.g.*, the LiDAR coordinate system, quantization according to the existing input interface, and data utilization, are more critical than developing loss functions, which have been overlooked in prior works. In this work, we show that simple fixes to these designs notably outperform existing methods by 16% in 3D semantic segmentation on the nuScenes dataset and 13% in 3D object detection on the KITTI dataset in downstream task performance. We focus on overlooked design choices along the spatial and temporal axes. Spatially, prior work has used cylindrical coordinate and voxel sizes without considering their side effects yielded with a commonly deployed sparse convolution layer input interface, leading to spatial quantization errors in 3D models. Temporally, existing work has avoided cumbersome data curation by discarding unsynced data, limiting the use to only the small portion of data that is temporally synced across sensors. We analyze these effects and propose simple solutions for each overlooked aspect. [Project page](#)

1 Introduction

Understanding 3D scenes with LiDAR is crucial for autonomous driving. With the advance of the neural network, the fundamental 3D scene understanding tasks, *e.g.*, 3D semantic segmentation or 3D object detection [31,41,8], have shown promising results with the 3D point cloud annotations [4,30,3]. However,

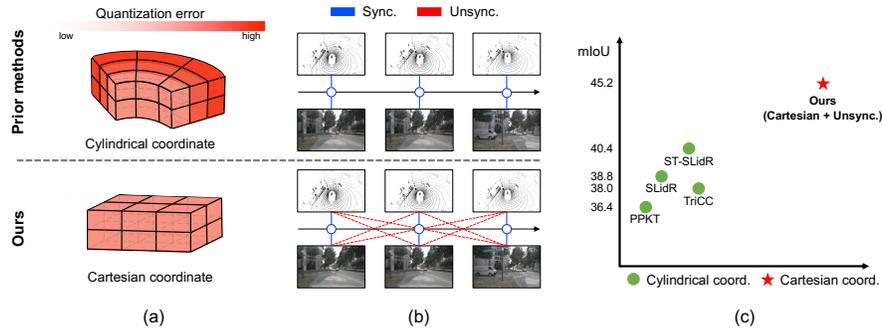


Fig. 1. Our simple treatments. (a) The quantization error of cylindrical coordinate increases with distance, while Cartesian coordinate has uniform quantization error regardless of distance. (b) Utilization of unsynced data collected at different times boosts up the combinations of image-LiDAR paired data. (c) Compared to the existing methods, we draw attention to underexplored treatments; quantization errors in coordinate system and unsynced data utilization. With these simple treatments, we achieve the state-of-the-art performance in 3D semantic segmentation task. The mIoU is measured by linear probing evaluation protocol in the nuScenes dataset.

annotating large-scale 3D LiDAR datasets requires time-consuming efforts with intensive labor [36,40]. As a workaround, by leveraging image representation learning, *image-to-LiDAR distillation* methods [29,20,26,18] have been developed and demonstrated their effectiveness in an annotation-efficient way. These methods transfer 2D image representation pre-trained on large data to the data-scarce 3D LiDAR domain. This facilitates learning rich and transferable 3D representations.

A series of recent image-to-LiDAR distillation studies [20,26,18] have focused on improving loss function designs. While some have reported good results, they depend very strongly on the standard setting used in SLidR by Sautier *et al.* [29], where the SLidR models are used as backbones without modification, and their proposed data processing remains the same. Following this convention without careful consideration introduces two bottlenecks.

First, previous work has primarily used the cylindrical coordinate as a standard 3D input space for quantizing raw LiDAR point clouds by considering how LiDAR data are acquired. Zhu *et al.* [41] show the effectiveness of the cylindrical coordinate in LiDAR-only tasks, but we found that this is detrimental in our image-to-LiDAR distillation scenario. The cylindrical coordinate can maintain *uniform density* across varying distances by a trade-off between increasing the voxel size and decreasing the density of LiDAR points according to distance [41]. While this appears to be beneficial, when using the standard sparse convolution input layer used in SLidR, the cylindrical coordinate inherently increases quantization errors as the distance increases from the origin (see Fig. 1a). This causes a degradation in the spatial domain, especially at far distances. Second, recent studies [29,20,26,18] have only used synced pairs of image-LiDAR data in a dataset to simply guarantee accurate 2D-3D matching without bells and

whistles. However, synced data account for only a small portion of the acquired data and discard the majority of the unsynced data. Thus, a significant portion of temporal information in the existing dataset is under-utilized.

Motivated by these, we propose simple treatments for image-to-LiDAR distillation that effectively handle two unexplored directions: quantization domain and unsynced data. The first treatment is to change the quantization domain from cylindrical to Cartesian, ensuring *uniform quantization errors* regardless of distances, and to set a much smaller voxel size. This simple strategy prevents the degradation of LiDAR points’ spatial resolution, ensuring more accurate representation association in image-to-LiDAR distillation. The second treatment is the use of unsynced data acquired at different times (see Fig. 1b). It should be noted that abundant unsynced data offers additional information for self-supervised training, but accurately estimating point-pixel correspondence may be considered necessary, which is often challenging. We propose a simple point-pixel matching method called Positive Pair Mining (PPM). The PPM module corrects inaccurate point-pixel matching caused by moving objects in a stratified manner, when image and point cloud data are fetched from different times. This PPM module enables the utilization of numerous unsynced LiDAR-image pair data.

We show that these simple treatments significantly improve the performance in two downstream tasks: 3D semantic segmentation and 3D object detection (see Fig. 1c). Our method outperforms all the records the previous state-of-the-art achieved under the benchmark setting. This hints at several issues with the existing evaluation protocol, *e.g.*, the loose tuning and the dissonance of the design choices of the baselines, which prevents understanding the true performance of the prior work. This cannot be inferred from each previous work alone. We summarize our main contributions as follows:

- We find that prior work underexplored and overlooked fundamental designs, such as the LiDAR coordinate system, quantization, and data utilization. Our simple remedies for spatial and temporal designs exhibit significant performance improvements, *i.e.*, state-of-the-art in downstream tasks.
- We argue that changing the coordinate from cylinder to Cartesian prevents LiDAR spatial resolution degradation and ensures uniform quantization error.
- We propose a compact module, *i.e.*, PPM, that corrects incorrect pixel-point matching and enables the model to utilize unsynced image-LiDAR paired data.
- We demonstrate consistent improvements, implying that our method not only improves LiDAR representation learning fundamentally but also sets the new baseline and protocol in the field that paves the way for further development.

2 Related Work

Our work is related to self-supervised learning that aims to learn useful representations without relying on labeled datasets, which can be categorized based on the data modalities. We briefly review the related lines of work.

2D Image-based Self-Supervised Learning. Recent approaches have shown remarkable results in learning 2D image representations with contrastive learning. These methods [5,11,12,14,21,24,32,35] contrast instance-level representation by augmenting different views of the same instance. Since contrastive learning is easily affected by the number of negative samples, several methods [2,38] propose objective functions that minimize the redundancy of the features and maximize the similarity between the same instance features. Such carefully designed contrast-based pre-training regime enables the image representations to be well-transferred to the downstream tasks, *e.g.*, image segmentation or image detection. As a teacher network, we leverage the powerful pre-trained 2D model, MoCov2 [6], and distill the 2D image knowledge into a 3D student network.

3D Point Cloud-based Self-Supervised Learning. Analogous to the image domain, 3D self-supervised learning approaches have focused on learning useful representations from the 3D point cloud. Reconstruction-based methods [7,28,33] generate object-level paired representations by perturbing point clouds, *e.g.*, distortion, random arrangement, or occlusion in pre-training. Most recent work [15,36,22,40,37] allows object-level 3D representations to transfer into 3D object detection or 3D segmentation while utilizing a single modality dataset, a 3D point cloud. Temporal consistency-based methods [34,23] find the corresponding point cloud across different timestamps by tracking the objects. STRL [16] exploits spatio-temporal cues from a 3D point cloud and generates temporally-correlated frames to learn invariant representations. While STRL proposes spatial/temporal data alignments for 3D self-supervised learning, single-modality representations suffer from a lack of appearance information in images.

2D-to-3D Distillation-based Self-Supervised Learning. Multi-modal self-supervised learning leverages pre-trained 2D image representations and distills them into 3D representations by applying contrastive loss to maximize the similarity between 2D-3D paired features. PPKT [19] proposes a pixel-to-point knowledge transfer learning method. They use the back-projection function to align 2D and 3D features and transfer knowledge between heterogeneous networks. SLidR [29] contrasts 2D image superpixels and corresponding 3D point cloud following SLIC [1]. Based on SLidR, ST-SLidR [20] proposes the semantically-tolerant and class-agnostic balanced loss. Due to the heterogeneous data acquisition frequencies of LiDAR and camera sensors, the aforementioned methods utilize only *synced* images-LiDAR pairs for multi-modal contrastive learning. Recent work, TriCC [26], propose the triplet loss function enabling the utilization of extended timestamps while they only fetch *synced* frames. Moreover, they obtain a pixel-point matching table by computing the similarity between pixel features and point features. Instead, we directly compute a 3D transformation matrix to obtain accurate correspondence between points and pixels. Another line of work, Seal [18] replace the superpixel-driven matching [1] with the foundation models, *e.g.* SAM [17], X-Decoder [42], OpenSeeD [39], and SEEM [43]. ScaLR [27] improves image-to-LiDAR distillation performance by scaling up the dataset, 2D backbone, and 3D backbone. Though not distillation, GPC [25] present a point cloud colorization method as a pretext task loss to

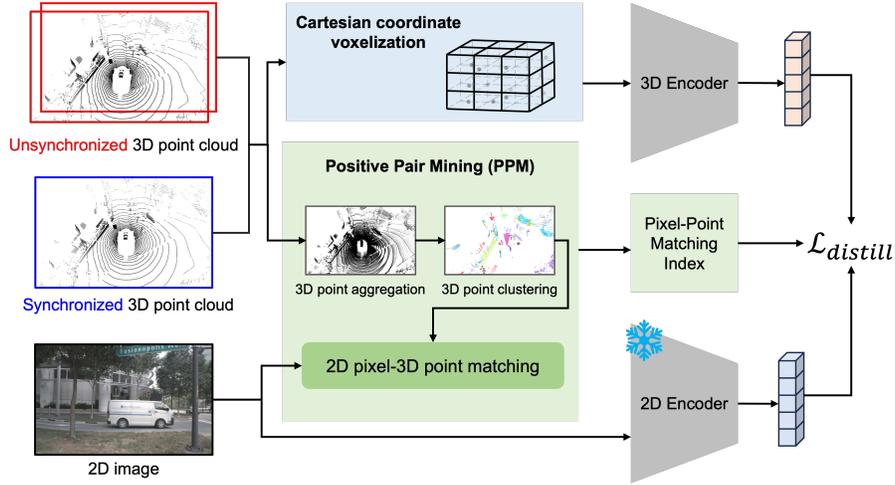


Fig. 2. Overall pipeline. Given 3D LiDAR point cloud and 2D image, our goal is to distill pre-trained 2D image representations into a 3D backbone model. Based on our observation that Cartesian coordinate prevents degradation in spatial resolution of LiDAR and ensures uniform quantization error, synced/unsynced points are quantized in Cartesian domain and then 3D encoder extracts the point-wise features. Unlike the prior work, we utilize unsynced point cloud for pixel-point matching. Our PPM module corrects inaccurate pixel-point matching for moving objects. Note that our pre-trained 2D encoder is frozen during training. We contrast the matched pixel-point features.

match synchronized images. The aforementioned methods focus on the designs of their own losses to effectively distill the pre-trained 2D image representations into the 3D model. Instead, we carefully examine the fundamental designs of the input interface and data utilization, which were overlooked by the prior art.

3 Simple Treatments

In this section, we first explore the changing coordinates for proper quantization (Sec. 3.1). Next, we provide the details of the PPM module for addressing unsynced data (Sec. 3.2). Finally, we outline a pipeline for integrating these treatments into image-to-LiDAR distillation (Sec. 3.3).

3.1 Treatment 1: Cartesian Coordinate

In the realm of LiDAR sensor data, cylindrical coordinate mitigates the issue of increasing sparsity in LiDAR point clouds with distance [41]. This approach allows for voxel sizes to expand with distance, enhancing the likelihood of points being included within a voxel. As a result, the density of voxels remains uniformly distributed, irrespective of the distance from the origin. However, as the distance increases, so does the voxel size, leading to a larger average distance between

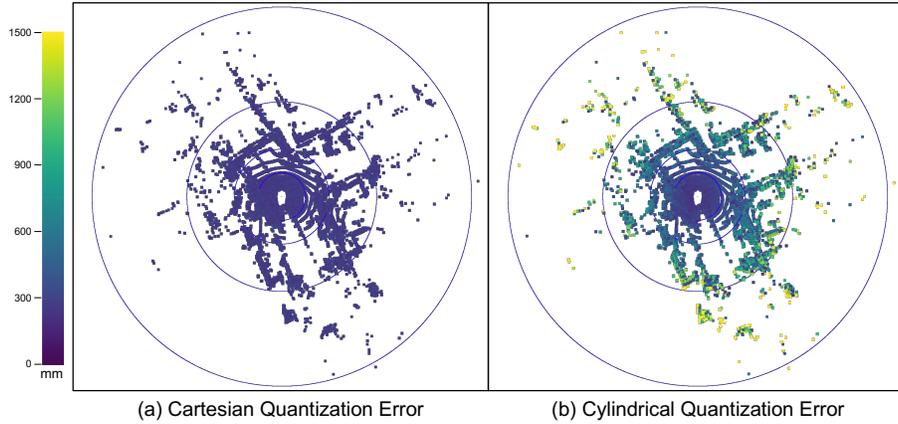


Fig. 3. Quantization error characteristics differ by coordinate system. (a) In the case of Cartesian quantization, the voxel size remains constant regardless of the distance, resulting in a uniform quantization error that does not vary with distance. On the other hand, (b) with Cylindrical quantization, the voxel size increases as the distance increases, leading to a rise in quantization error due to its characteristic of expanding voxel size with distance.

the original and quantized positions of points within a voxel. This signifies an increase in the quantization error as one moves further from the origin (see Fig. 3a). Consequently, with increasing distance from the origin, the enlarged voxel volume results in a lower data resolution, preserving details at close range while significantly losing details at far distances. This ultimately leads to a degradation in overall data resolution. Conversely, using Cartesian coordinates maintains the consistent voxel size throughout the space, resulting in the uniform quantization error across all distances. The average error determined by voxel size remains constant, regardless of position (see Fig. 3b). Thus, while some information may be lost at close range, the loss of detail at larger distances is considerably lessened. This ultimately results in the preservation of overall data resolution.

We postulate that small and uniform quantization error is more critical than the emphasis on uniform density [41] in performance. The degradation of overall data resolution would confine image-to-LiDAR distillation, causing a decline in performance. To alleviate this, we propose the adoption of cartesian coordinates in place of cylindrical coordinates within the image-to-LiDAR distillation scheme. Our first simple treatment involves a modification to the conventional process shared by existing image-to-LiDAR distillation methods before the point cloud feed-forward to the 3D network, which is as follows:

1. Transforming the raw 3D points from cartesian to cylindrical coordinate.
2. Voxelizing of the input with 3D cylindrical partitioning, with voxel sizes of $\{\delta\rho = 10cm, \delta\phi = 1^\circ, \delta z = 10cm\}$ and quantizing

We propose a modification to this process:

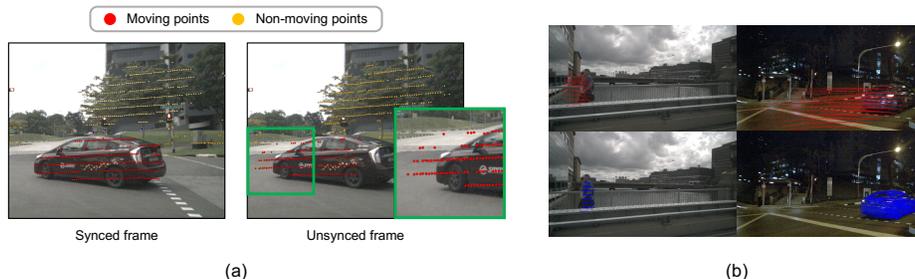


Fig. 4. Visualization of the effects of PPM module. (a) Compared to the synced frame, unsynced LiDAR points are misaligned with pixels in moving object regions. (b) With the PPM module, we observe that moving points are well matched with pixels in the corresponding 2D images.

1. **Not** transforming the 3D points from cartesian to cylindrical coordinate.
2. Voxelizing of the input using 3D cartesian partitioning, with voxel sizes of $\{\delta x = 10cm, \delta y = 10cm, \delta z = 10cm\}$ and quantizing.

By cartesian coordinates, we could alleviate overall resolution degradation, ensuring that image-to-LiDAR distillation is not constrained by degraded resolution.

3.2 Treatment 2: Utilization of Unsynced Data

Utilizing only synced data for image-to-LiDAR distillation learning significantly limits data utilization. Intuitively, given N number of subsequent point clouds $[P^1, \dots, P^t, \dots, P^N]$ and images $[I^1, \dots, I^t, \dots, I^N]$, where t denotes time, the available data is limited to N instances. Thanks to various combinations of LiDAR-image data, we can increase the quantity of data by utilizing unsynced data from different timestamps instead of only using synced data from the same timestamp. However, additional data processing on moving points is required to leverage unsynced data. Figure 4a demonstrates the reason. In the unsynced case, the non-moving points are accurately aligned with objects in the 2D images, but the moving points are mismatched because obtaining projection from moving points to the 2D image is intractable. To utilize unsynced data, we propose a simple treatment, *i.e.*, the Positive Pair Mining (PPM) module. The PPM module dynamically registers point clouds from the source to the target timestamp, regardless of whether the objects are moving or stationary. Using the calibration matrix provided with the synced data, the dynamically registered point clouds can be more accurately matched to the images corresponding to the target time.

Specifically, let P^t and I^s denotes the point cloud and image, which are acquired at time t and s from different sensors, respectively, which we represent $P^t = [\mathbf{p}_1^t, \dots, \mathbf{p}_{N_t}^t]$ and $I^s \in \mathbb{R}^{H \times W \times 3}$, where $\mathbf{p}_i^t \in \mathbb{R}^3$ denotes the i -th 3D point at the time t in cartesian coordinate, and H , W and N_t the height and width of the image, and the number of points at time t , respectively. Given a known

pre-calibrated relative pose between the LiDAR and the camera, the 3D-to-2D projection $T_{t,s} : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ outputs the projected 2D coordinate of I^s , *i.e.*, $\mathbf{x}_i^s = T_{t,s}(\mathbf{p}_i^t)$. If $t = s$, $T_{t,s}$ ensures the accurate 3D-to-2D correspondence between \mathbf{p}_i^t and \mathbf{x}_i^s . However, for $t \neq s$, it does not guarantee the correspondence. To address this potentially unreliable matching case, we design the PPM module. Given point clouds at an unsynced frame, we first inspect whether they are likely moving because the projection is unreliable with moving 3D points. we subsequently find the more accurate 3D-to-2D matching by the mining of a 3D transformation matrix $Z_{t,s,i} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ for the moving points. To be concise, we omit the subscript indices. Z is combined with T to create the pixel-point matching index. As shown in Fig. 4b, moving points are better matched to the image when utilizing the PPM module than when it is not.

Our proposed PPM module consists of unsupervised methods. Notably, it involves aggregating unsynced and synced 3D point clouds from consecutive frames and applying 3D point clustering to detect objects. Additionally, it includes ground removal to separate objects, moving cluster tracking to identify moving objects, and cluster-wise Iterative Closest Point (ICP) to register points for moving objects specifically. See supplementary material for details.

3.3 Image-to-LiDAR Distillation with Simple Treatments

We combined the two simple treatments to construct the overall pipeline. Figure 2 demonstrates the overall pipeline. Compared to the baseline [29], we employ cartesian coordinate and enable 3D representation learning using unsynced data by introducing the PPM module.

As shown in Fig. 2, we sample the 3D point clouds P^t and P^s , where t and s denote synced and unsynced frames. We also sample the 2D image I^t at synced frame. The point clouds quantized in cartesian coordinate and image pass through the 3D and 2D networks, resulting in corresponding 3D point-wise features $\mathbf{F}^t \in \mathbb{R}^{N'_t \times D}$ and $\mathbf{F}^s \in \mathbb{R}^{N'_s \times D}$, and 2D pixel-wise features $\mathbf{G}^t \in \mathbb{R}^{H \times W \times D}$. Because of quantization, N' is lower than N and the number of pixel-point matching index is reduced. Then, we find the matching between \mathbf{F} and \mathbf{G} using the pixel-point matching index. At the time t , the pair of 3D point and 2D pixel features $\{\mathbf{f}_i^t, \mathbf{g}_i^t\}$ are follows:

$$\{\mathbf{f}_i^t, \mathbf{g}(\mathbf{x}_i^t) \mid \mathbf{x}_i^t = T(\mathbf{p}_i^t)\}, \quad (1)$$

where \mathbf{g} is function that exploits 2D pixel feature \mathbf{g}_i^t from \mathbf{G}^t at 2D coordinate \mathbf{x}_i^t . If the 3D point comes from the unsynced frame s , we employ Z for the pairing of $\{\mathbf{f}_j^s, \mathbf{g}_j^t\}$, which defined as:

$$\{\mathbf{f}_j^s, \mathbf{g}(\mathbf{x}_j^t) \mid \mathbf{x}_j^t = T(Z(\mathbf{p}_j^s))\}. \quad (2)$$

Note that Z is the identity matrix when applied to non-moving points, including ground points. We then gather $M = N'_t + N'_s$ number of paired features $\{\mathbf{f}_k, \mathbf{g}_k\}_{k=1}^M$ from the synced and unsynced data. Then, we apply the image-to-LiDAR contrastive distillation loss to enforce the 3D feature to be similar to the

paired 2D feature. We use the same self-supervised loss with SLiDR [29] to see the effects of our changes without any fancy loss design, which is defined as

$$\mathcal{L}_{distill} = -\frac{1}{M} \sum_{k=1}^M \log\left[\frac{\exp(\langle \mathbf{f}_k, \mathbf{g}_k \rangle / \tau)}{\sum_{d=1}^M \exp(\langle \mathbf{f}_d, \mathbf{g}_k \rangle / \tau)}\right], \quad (3)$$

where $\langle \cdot, \cdot \rangle$ denotes the normalized inner product in the l_2 -normalized \mathbb{R}^D , and $\tau > 0$ is the temperature term. The positive and negative pairs for the contrastive learning are determined following the implementation of Sautier *et al.* [29], which groups the pixel-wise features by super-pixel segmentation [1] and their corresponding point-wise features. We simplify the above expressions by using pixel-wise representation instead of superpixel-wise.

4 Experiments

In this section, we describe the pre-training details of 3D network (Sec. 4.1). In Sec. 4.2, we present the fine-tuning results with varying amounts of annotation and ablation studies of the proposed PPM module and coordinates. Then, we provide the experiment results on the downstream tasks (Sec. 4.3 and Sec. 4.4).

4.1 Experiment Setup

Networks. The 3D backbone network is composed of sparse residual U-Net architecture [9] and a linear layer, which projects 3D features to D -dimensional space. The sparse residual U-Net architecture is composed of $3 \times 3 \times 3$ with sparse convolutions. The input to the 3D backbone network is voxelized in either cylindrical or Cartesian coordinates, with voxel sizes altered to $10cm$ or $5cm$ for experimentation. This aspect differs from traditional image-to-LiDAR distillation methods. The 2D network is composed of 1) MoCov2 [6] that is a frozen network pre-trained by a self-supervised method, whose architecture is ResNet 50 [13], 2) a trainable linear layer that projects 2D features to D -dimensional space, and 3) the bilinear upsampling layer to obtain a pixel-wise feature of input resolution from a reduced output resolution.

Pre-training Datasets. We pre-train the 3D network on the nuScenes dataset [4], which has 700 training scenes, and split into 600 remaining training scenes and 100 mini-validation scenes for choosing the training hyperparameters. The sampling rates of LiDAR and the camera are 20Hz and 12Hz, respectively, and synced data has a sampling rate of 2Hz. We refer to the timestamps of synced data at 2Hz as keyframes and the timestamps of the remaining frame data as inter-frames. The LiDAR and image data belonging to inter-frames are inherently unsynced due to the difference in data collection frequencies. We train the 3D network with keyframe and inter-frame data using the proposed method.

Data augmentation. We follow the data augmentation strategy used in Sautier *et al.* [29]. For the 3D point clouds, we apply the composition of random transformations, including a z-axis rotation, a 50% probability random flip against

Table 1. Impact of coordinate and voxel size. When using cylindrical coordinates, average quantization error is larger, leading to a reduction in data resolution. In contrast, using Cartesian coordinates results in a smaller average quantization error, maintaining data resolution. Reducing voxel size decreases the quantization error and preserves data resolution better. As data resolution is well preserved, performance improves.

Method	Coordinate	Voxel size (<i>cm</i>)	Quan. error (<i>mm</i>)	nuScenes
				Lin. Prob. (100%)
SLidR [29]	Cylindrical	10	229.2	38.8
Ours	Cartesian	10	96.5	40.8
Ours	Cartesian	5	48.2	41.2

the x and y-axes, and a random cuboid dropping. For the images, we apply the random horizontal flip and random crop-resizing transformations.

Hyperparameters. The 3D backbone network is pre-trained using 4 A100 GPUs, employing a batch size 16 and spanning 50 epochs. We adopt the Stochastic Gradient Descent (SGD) for the optimization algorithm with an initial learning rate set at 0.5, momentum of 0.9, weight decay of 0.0001, and dampening of 0.1. We use the cosine annealing scheduler that progressively reduces the learning rate from its initial value to 0 by the conclusion of the 50th epoch.

4.2 Ablation Studies

Coordinate system and voxel size. Table 1 ablates the effect of the coordinate system and voxel size. As shown in Fig. 3, cylindrical coordinate increases the spatial quantization errors as the distance increases, which introduces performance degradation to 3D models. We thus convert the input coordinate system from cylindrical to Cartesian, which ensures the uniform quantization error regardless of the distances. This approach reduces the overall quantization error and preserves the data resolution well. Additionally, by decreasing the voxel size, we further reduce the overall quantization error, thereby preserving the data resolution even more. Our method with the Cartesian coordinate interface enjoys a clear benefit of +2.0% in the 3D semantic segmentation task with nuScenes dataset. Moreover, we obtain an additional +0.4% gain with the smaller voxel sizes. We postulate that preserving data resolution effects improves image-to-LiDAR distillation learning and results in performance improvement. Therefore, our joint treatment of the spatial aspect, *i.e.*, coordinate change and voxel size adjustment, is more impactful than the prior art that only focuses on loss designs.

Unsynced data utilization. In this ablation study, the cartesian coordinate system and a voxel size of 5*cm* were applied as the default settings. In Table 2, we evaluate the impact of our utilization of unsynced frames and our proposed Positive Pair Matching (PPM) module. Since unsynced data is collected at different times, the misalignments between 2D pixels and 3D points occur in moving objects. Without dedicated adjustment, the sole use of unsynced frames hinders the distillation and causes performance degradation. Our proposed PPM module

Table 2. Impact of data utilization and PPM. To assess the effect of the proposed PPM module, we build two baseline methods, (A) and (B). Method (A) is the reproduced SLiDR with the Cartesian coordinate. For a fair comparison, we have increased the batch size of SLiDR, since (A) is trained on synced data only. Method (B) also utilizes the unsynced data and applies the nearest alignment to minimize misalignment between image and point cloud. With utilizing unsynced data, the proposed PPM module leads to significant performance improvement.

Methods	Data utilization	Alignment method	nuScenes
			Lin. Prob. (100%)
(A)	Synced	None	41.6
(B)	Synced + Unsynced	Nearest	41.5
Ours	Synced + Unsynced	PPM	45.2

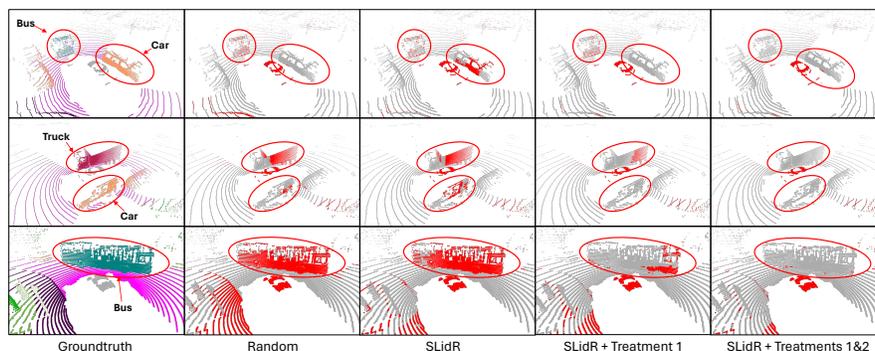


Fig. 5. Qualitative results of error map. Error map among the Random, SLiDR [29], SLiDR w. Cartesian, and Ours when fine-tuned with 1% labeled data. Gray and red colors denote correct and incorrect points, respectively. We visualize the results of the validation set in nuScenes [4]. When our treatments were applied sequentially, we observe the qualitative improvement. Treatment1 and treatment2 denote Cartesian coordinate and utilization of unsynced data, respectively.

resolves the issue by correcting inaccurate pixel-point matching and enables the utilization of numerous unsynced data pairs. Our methodology involves utilizing inter-frame data in addition to key-frame data, which effectively doubles the batch size. For a fair comparison, we increase the batch size of SLiDR twofold. Additionally, to validate the effectiveness of the PPM’s misalignment correction capability, we compare it with a baseline that minimizes misalignment, called the nearest alignment. The nearest alignment baseline utilizes unsynced data by employing both inter-frame LiDAR and inter-frame images but minimizes potential misalignment between inter-frame data by sampling inter-frame LiDAR and matching it with the closest timestamp’s inter-frame image. We find that our PPM module brings significant performance improvement (+3.7%) compared to other baselines. As shown in Fig. 5, we observed that the qualitative results improve as our treatment is applied sequentially.

Table 3. 3D semantic segmentation results on nuScenes and SemanticKITTI validation sets. We compare our method with the existing 3D representation learning methods using the nuScenes and SemanticKITTI datasets. Our method surpasses the existing methods across all metrics. The table is horizontally partitioned based on whether (Top) fully-unsupervised methods or (Bottom) SAM were used.

Method	nuScenes						SemanticKITTI		
	Lin. Prob.	1%	5%	10%	25%	100%	1%	5%	10%
Random	8.1	30.3	47.7	56.6	64.8	74.2	39.5	52.1	55.6
PointContrast [36]	21.9	32.5 (+2.2)	-	57.1 (+0.5)	-	74.3 (+0.1)	41.1 (+1.6)	-	-
DepthContrast [40]	22.1	31.7 (+1.4)	-	57.3 (+0.7)	-	74.1 (-0.1)	41.5 (+2.0)	-	-
PPKT [19]	36.4	37.8 (+7.5)	51.7(+4.0)	59.2 (+2.6)	66.8 (+2.0)	73.8 (-0.4)	43.9 (+4.4)	53.1 (+1.0)	57.3 (+1.7)
SLidR [29]	38.8	38.2 (+7.9)	52.2 (+4.5)	58.8 (+2.2)	66.2 (+1.4)	74.6 (+0.4)	44.6 (+5.1)	52.6 (+0.5)	56.0 (+0.4)
ST-SLidR [20]	40.4	40.7 (+10.4)	54.6 (+6.9)	60.7 (+4.1)	67.7 (+2.9)	75.1 (+0.9)	44.7 (+5.2)	-	-
TriCC [26]	38.0	41.2 (+10.9)	54.1 (+6.4)	60.4 (+3.8)	67.6 (+2.8)	75.6 (+1.4)	45.9 (+6.4)	55.9 (+3.8)	59.0 (+3.4)
Ours	45.2	42.7 (+12.4)	56.8 (+9.1)	63.3 (+6.7)	69.8 (+5.0)	75.7 (+1.5)	50.3 (+10.8)	61.1 (+9.0)	63.3 (+7.7)
Seal [18]	44.9	45.8 (+15.5)	55.6 (+7.9)	62.9 (+6.3)	68.4 (+3.6)	75.6 (+1.4)	46.6 (+7.1)	-	-
Ours-SAM	48.0	43.0 (+12.7)	57.1 (+9.4)	64.7 (+8.1)	70.0 (+5.2)	75.9 (+1.7)	52.1 (+12.6)	61.5 (+9.4)	63.6 (+8.0)

Table 4. Few-shot 3D object detection results on KITTI. Using the KITTI dataset, we fine-tune our model with few labeled dataset and evaluate the mAP@R40 metric. Our method outperforms the recent Image-to-LiDAR distillation methods, *i.e.*, PPKT [19], SLidR [29], and TriCC [26] by a large margin and achieves on average more than a two-fold performance increase. † denotes KITTI pretraining, ‡ denotes Waymo [30] pretraining. The table is horizontally partitioned by different epoch settings.

Method	Epoch	Fine-tune (5%)			Fine-tune (10%)			Fine-tune (20%)		
		Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard
Random	50	73.7	56.6	50.7	74.6	58.8	53.9	77.9	63.7	59.2
PPKT [19]	50	75.7 (+2.0)	59.6 (+3.0)	54.4 (+3.7)	78.3 (+3.7)	63.7 (+4.9)	58.4 (+4.5)	78.9 (+1.0)	64.8 (+1.1)	59.9 (+0.7)
SLidR [29]	50	74.5 (+0.8)	58.8 (+2.2)	52.9 (+2.2)	78.1 (+3.5)	63.5 (+4.7)	58.3 (+4.4)	77.6 (-0.3)	63.8 (+0.1)	59.2
Ours	50	84.3 (+10.6)	69.3 (+12.7)	64.0 (+13.3)	84.0 (+9.4)	70.9 (+12.1)	66.1 (+12.2)	84.0 (+6.1)	70.9 (+7.2)	66.1 (+6.9)
TriCC [26]	20	77.9 (+4.2)	61.3 (+4.7)	56.2 (+5.5)	79.6 (+5.0)	64.6 (+5.8)	59.3 (+5.4)	80.0 (+2.1)	65.9 (+2.2)	60.7 (+1.5)
GPC† [25]	80	-	62.7 (+6.1)	-	-	66.4 (+7.6)	-	-	70.7 (+7.0)	-
GPC‡ [25]	12	-	63.9 (+7.3)	-	-	67.0 (+8.2)	-	-	70.1 (+6.4)	-
Ours	20	83.4 (+9.7)	69.3 (+12.7)	63.7 (+13.0)	82.5 (+7.9)	70.0 (+11.2)	65.5 (+11.6)	82.5 (+4.6)	70.0 (+6.3)	65.5 (+6.3)

4.3 Transfer to Semantic Segmentation

We evaluate our pre-trained 3D backbone by fine-tuning it on a 3D semantic segmentation task. We choose the nuScenes [4] dataset with 16 semantic categories and the SemanticKITTI [3] dataset with 19 semantic categories. We measure the Mean Intersection over Union (mIoU) metric for evaluation. Following common fine-tuning settings [29,20] we evaluate the performance on the mini-validation set of each dataset. For nuScenes dataset, we utilize two evaluation protocols: 1) linear probing with fully exploiting annotations and 2) fine-tuning with different portions of annotations, *i.e.* {1%, 5%, 10%, 25%, 100%}. For the SemanticKITTI dataset, we report the fine-tuning results with few-shot {1%, 5%, 10%} settings.

In Table 3, we evaluate our learned representations on the 3D semantic segmentation task. To demonstrate, we establish a Random model by randomly initializing it without using the pre-trained weight from the 3D network. For the nuScenes dataset [4], we find that all the self-supervised models initialized with pre-trained weights show improvements over the model initialized with random weight. Interestingly, we observe that our method outperforms all the competing

Table 5. Per-class fine-tuning performance on nuScenes with 1% of annotation. We evaluate the per-class IoU on nuScenes. Compared to competing methods, our method achieves the best performance in the average class. Specifically, it significantly boosts the performance for classes that are challenging to learn due to their smaller resolution, *e.g.*, bicycle, pedestrian, traffic cone.

Method	barrier	bicycle	bus	car	const. veh.	motorcycle	pedestrian	traffic cone	trailer	truck	driv. surf.	other flat	sidewalk	terrain	manmade	vegetation	mIoU
Random	0.0	0.0	8.1	65.0	0.1	6.6	21.0	9.0	9.3	25.8	89.5	14.8	41.7	48.7	72.4	73.3	30.3
PointContrast [36]	0.0	1.0	5.6	67.4	0.0	3.3	31.6	5.6	12.1	30.8	91.7	21.9	48.4	50.8	75.0	74.6	32.5
DepthContrast [40]	0.0	0.6	6.5	64.7	0.2	5.1	29.0	9.5	12.1	29.9	90.3	17.8	44.4	49.5	73.5	74.0	31.7
PPKT [19]	0.0	2.2	20.7	75.4	1.2	13.2	45.6	8.5	17.5	38.4	92.5	19.2	52.3	56.8	80.1	80.9	37.8
SLiDR [29]	0.0	3.1	15.2	72.0	0.9	18.8	43.2	12.5	14.7	33.3	92.8	29.4	54.0	61.0	80.2	81.9	38.3
ST-SLiDR [20]	0.0	2.7	16.0	74.5	3.2	25.4	50.9	20.0	17.7	40.2	92.0	30.7	54.2	61.1	80.5	82.9	40.7
TriCC [26]	0.0	2.6	20.7	73.6	0.3	18.9	49.2	22.0	16.9	33.4	94.5	43.1	57.2	62.1	82.3	82.6	41.2
Ours	0.0	3.4	26.0	77.4	2.4	21.4	61.4	30.9	17.5	44.4	92.2	28.3	54.1	62.1	79.6	81.6	42.7

methods by a large margin, *e.g.*, +37.1% for linear probing and +12.4% for fine-tuning with 1% annotations. We also evaluate the fine-tuning performance on the SemanticKITTI dataset [3] to verify the effectiveness. Although our model is only pre-trained on the nuScenes dataset, our model surpasses all the competing methods across varying annotation settings. This demonstrates that our method can generalize to unseen datasets and different volumes of annotations. These results reveal the dissonance of the existing baselines’ design choices, and our simple treatments of input interface and data utilization effectively resolve the issues. Additionally, we compare our methodology with Seal [18], which utilizes Vision Foundation Models (VFM) with a superpixel segmenter instead of SLIC [1]. Despite using SLIC as is, our method shows superior performance in all metrics except one compared to Seal. Furthermore, when we replace SLIC with one of the VFMs, Segment Anything Models (SAM) [17], in our methodology, dubbed Ours-SAM, we observe a consistent overall improvement in performance. This suggests that our methodology, proposing a way to utilize raw data itself, can be applied to other methods to achieve consistently improved performance.

4.4 Transfer to Object Detection

To validate the usefulness of the 3D representation learned by the proposed method, we report performance on another challenging downstream task, few-shot 3D object detection. For the few-shot 3D object detection dataset, we adopt the KITTI dataset [10]. For a fair comparison, we exactly follow the fine-tuning implementation details of SLiDR [29]. Table 4 shows the mAP@R40 performance of networks fine-tuned on 5%, 10%, and 20% annotations on the difficulty levels defined by the KITTI dataset, *e.g.*, easy, moderate, and hard. For a fair comparison with TriCC [26], we provide the performance of the pre-trained network for 20 epochs as well. Our method archives the best performance across all metrics and surpasses the other Image-to-LiDAR distillation methods, *e.g.*, PPKT [19], SLiDR [29], TriCC [26] by a large margin regardless of the number of epochs. Compared to the improvements achieved by previous works at random, we achieve, on average, more than a two-fold performance increase with simple

treatments, and in the fine-tune 5% moderate setting, we achieve a four-fold performance improvement. Our method also shows a significant performance improvement compared to GPC [25] pre-trained on KITTI and Waymo, which have similar or identical LiDAR characteristics to downstream dataset *i.e.* KITTI.

4.5 Per-class IoU

We report the 3D semantic segmentation performance of per-class by fine-tuning 1% on the nuScenes dataset (see Table 5). Compared to the competing methods, Ours achieves the best performance at the average of classes and is ranked first in 16 classes. By balancing the overall quantization error to increase data resolution, performance across all classes improves. Notably, this leads to a significant enhancement in performance for classes that are difficult to learn due to their small resolution, such as bicycle, pedestrian, and traffic cone. Especially with treatments focused on data utilization, we show good performance compared to ST-SLiDAR [20], which is designed with a loss aware of rare classes (e.g., bicycle, motorcycle, pedestrian, and traffic cone) that have a small number of points.

5 Conclusion

We delve into 3D representation learning, especially in image-to-LiDAR distillation, and find that the prior arts overlook the fundamental designs, *e.g.*, the LiDAR coordinate system, quantization, and data under-utilization. Compared to the existing works focusing on their own design of loss functions, we propose simple remedies for spatial and temporal designs. Spatially, we change the domain of quantization from cylindrical to Cartesian, which ensures uniform quantization error regardless of distance and prevents the degradation of spatial resolution. Temporally, we propose a PPM module that enables the model to utilize the unsynced data by aligning mismatched pixel-point paired data. Our spatial and temporal treatments facilitate the synergies and exhibit significant performance improvements, *i.e.*, state-of-the-art performance under the same benchmark setting in two downstream tasks: 3D semantic segmentation and 3D object detection. We hope that our work provides common ground by offering our new baseline and input protocol to contribute to future image-to-LiDAR distillation research and help pave the way for further development.

Acknowledgement. This project was supported by RideFlux and also supported by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.RS-2022-II220124, Development of Artificial Intelligence Technology for Self-Improving Competency-Aware Learning Capabilities; No. 2020-0-00004, Development of Previsional Intelligence based on Long-term Visual Memory Network; No.RS-2020-II201336, Artificial Intelligence Graduate School Program(UNIST))

References

1. Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Süsstrunk, S.: Slic superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **34**(11), 2274–2282 (2012)
2. Bardes, A., Ponce, J., LeCun, Y.: Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906* (2021)
3. Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J.: Semantickitti: A dataset for semantic scene understanding of lidar sequences. In: *IEEE International Conference on Computer Vision (ICCV)*. pp. 9297–9307 (2019)
4. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 11621–11631 (2020)
5. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International Conference on Machine Learning (ICML)*. pp. 1597–1607. PMLR (2020)
6. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297* (2020)
7. Chen, Y., Liu, J., Ni, B., Wang, H., Yang, J., Liu, N., Li, T., Tian, Q.: Shape self-correction for unsupervised point cloud understanding. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8382–8391 (2021)
8. Cheng, R., Razani, R., Taghavi, E., Li, E., Liu, B.: 2-s3net: Attentive feature fusion with adaptive feature selection for sparse semantic segmentation network. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 12547–12556 (2021)
9. Choy, C., Gwak, J., Savarese, S.: 4d spatio-temporal convnets: Minkowski convolutional neural networks. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3075–3084 (2019)
10. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2012)
11. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: *2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06)*. vol. 2, pp. 1735–1742. IEEE (2006)
12. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9729–9738 (2020)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
14. Henaff, O.: Data-efficient image recognition with contrastive predictive coding. In: *International conference on machine learning*. pp. 4182–4192. PMLR (2020)
15. Hou, J., Graham, B., Nießner, M., Xie, S.: Exploring data-efficient 3d scene understanding with contrastive scene contexts. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15587–15597 (2021)
16. Huang, S., Xie, Y., Zhu, S.C., Zhu, Y.: Spatio-temporal self-supervised representation learning for 3d point clouds. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 6535–6545 (2021)

17. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. arXiv preprint arXiv:2304.02643 (2023)
18. Liu, Y., Kong, L., Cen, J., Chen, R., Zhang, W., Pan, L., Chen, K., Liu, Z.: Segment any point cloud sequences by distilling vision foundation models. arXiv preprint arXiv:2306.09347 (2023)
19. Liu, Y.C., Huang, Y.K., Chiang, H.Y., Su, H.T., Liu, Z.Y., Chen, C.T., Tseng, C.Y., Hsu, W.H.: Learning from 2d: Contrastive pixel-to-point knowledge transfer for 3d pretraining. arXiv preprint arXiv:2104.04687 (2021)
20. Mahmoud, A., Hu, J.S., Kuai, T., Harakeh, A., Paull, L., Waslander, S.L.: Self-supervised image-to-point distillation via semantically tolerant contrastive loss. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7102–7110 (2023)
21. Misra, I., Maaten, L.v.d.: Self-supervised learning of pretext-invariant representations. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6707–6717 (2020)
22. Nunes, L., Marcuzzi, R., Chen, X., Behley, J., Stachniss, C.: Segcontrast: 3d point cloud feature representation learning through self-supervised segment discrimination. IEEE Robotics and Automation Letters **7**(2), 2116–2123 (2022)
23. Nunes, L., Wiesmann, L., Marcuzzi, R., Chen, X., Behley, J., Stachniss, C.: Temporal consistent 3d lidar representation learning for semantic perception in autonomous driving. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5217–5228 (2023)
24. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
25. Pan, T.Y., Ma, C., Chen, T., Phoo, C.P., Luo, K.Z., You, Y., Campbell, M., Weinberger, K.Q., Hariharan, B., Chao, W.L.: Pre-training lidar-based 3d object detectors through colorization (2024)
26. Pang, B., Xia, H., Lu, C.: Unsupervised 3d point cloud representation learning by triangle constrained contrast for autonomous driving. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5229–5239 (2023)
27. Puy, G., Gidaris, S., Boulch, A., Siméoni, O., Sautier, C., Pérez, P., Bursuc, A., Marlet, R.: Three pillars improving vision foundation model distillation for lidar. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21519–21529 (2024)
28. Sauder, J., Sievers, B.: Self-supervised deep learning on point clouds by reconstructing space. Advances in Neural Information Processing Systems **32** (2019)
29. Sautier, C., Puy, G., Gidaris, S., Boulch, A., Bursuc, A., Marlet, R.: Image-to-lidar self-supervised distillation for autonomous driving data. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9891–9901 (2022)
30. Sun, P., Kretschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2446–2454 (2020)
31. Tang, H., Liu, Z., Zhao, S., Lin, Y., Lin, J., Wang, H., Han, S.: Searching efficient 3d architectures with sparse point-voxel convolution. In: European Conference on Computer Vision (ECCV). pp. 685–702. Springer (2020)
32. Tian, Y., Krishnan, D., Isola, P.: Contrastive representation distillation. arXiv preprint arXiv:1910.10699 (2019)

33. Wang, H., Liu, Q., Yue, X., Lasenby, J., Kusner, M.J.: Unsupervised point cloud pre-training via occlusion completion. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9782–9792 (2021)
34. Wu, Y., Zhang, T., Ke, W., Süssstrunk, S., Salzmann, M.: Spatiotemporal self-supervised learning for point clouds in the wild. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5251–5260 (2023)
35. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3733–3742 (2018)
36. Xie, S., Gu, J., Guo, D., Qi, C.R., Guibas, L., Litany, O.: Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In: European Conference on Computer Vision (ECCV). pp. 574–591. Springer (2020)
37. Yin, J., Zhou, D., Zhang, L., Fang, J., Xu, C.Z., Shen, J., Wang, W.: Proposalcontrast: Unsupervised pre-training for lidar-based 3d object detection. In: European Conference on Computer Vision. pp. 17–33. Springer (2022)
38. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow twins: Self-supervised learning via redundancy reduction. In: International Conference on Machine Learning. pp. 12310–12320. PMLR (2021)
39. Zhang, H., Li, F., Zou, X., Liu, S., Li, C., Yang, J., Zhang, L.: A simple framework for open-vocabulary segmentation and detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1020–1031 (2023)
40. Zhang, Z., Girdhar, R., Joulin, A., Misra, I.: Self-supervised pretraining of 3d features on any point-cloud. In: IEEE International Conference on Computer Vision (ICCV). pp. 10252–10263 (2021)
41. Zhu, X., Zhou, H., Wang, T., Hong, F., Ma, Y., Li, W., Li, H., Lin, D.: Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9939–9948 (2021)
42. Zou, X., Dou, Z.Y., Yang, J., Gan, Z., Li, L., Li, C., Dai, X., Behl, H., Wang, J., Yuan, L., et al.: Generalized decoding for pixel, image, and language. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15116–15127 (2023)
43. Zou, X., Yang, J., Zhang, H., Li, F., Li, L., Wang, J., Wang, L., Gao, J., Lee, Y.J.: Segment everything everywhere all at once. *Advances in Neural Information Processing Systems* **36** (2024)