

DOS: Distilling Observable Softmaps of Zipfian Prototypes for Self-Supervised Point Representation

Mohamed Abdelsamad^{1,2}, Michael Ulrich¹, Bin Yang¹, Miao Zhang¹,
Yakov Miron¹, Abhinav Valada²

¹Bosch Center for Artificial Intelligence

²University of Freiburg

Abstract

Recent advances in self-supervised learning (SSL) have shown tremendous potential for learning 3D point cloud representations without human annotations. However, SSL for 3D point clouds still faces critical challenges due to irregular geometry, shortcut-prone reconstruction, and unbalanced semantics distribution. In this work, we propose *DOS* (Distilling Observable Softmaps), a novel SSL framework that self-distills semantic relevance softmaps only at observable (unmasked) points. This strategy prevents information leakage from masked regions and provides richer supervision than discrete token-to-prototype assignments. To address the challenge of unbalanced semantics in an unsupervised setting, we introduce Zipfian prototypes and incorporate them using a modified Sinkhorn-Knopp algorithm, *Zipf-Sinkhorn*, which enforces a power-law prior over prototype usage and modulates the sharpness of the target softmap during training. *DOS* outperforms current state-of-the-art methods on semantic segmentation and 3D object detection across multiple benchmarks, including nuScenes, Waymo, SemanticKITTI, ScanNet, and ScanNet200, without relying on extra data or annotations. Our results demonstrate that observable-point softmaps distillation offers a scalable and effective paradigm for learning robust 3D representations.

Introduction

SSL has emerged as a compelling alternative to supervised pretraining for 3D point cloud representation learning, enabling large-scale training without the need for expensive, labor-intensive annotations. While recent methods (Wu et al. 2025; Abdelsamad et al. 2025; Abou Zeid et al. 2025; Hermosilla, Stippel, and Sick 2025; Lang et al. 2024b) have demonstrated promising results across 3D perception tasks, several fundamental challenges persist. The irregular and unstructured nature of point clouds, their susceptibility to shortcut learning, and the inherently long-tailed distribution of semantic categories all contribute to a persistent performance gap between supervised and self-supervised representations on downstream tasks.

Most existing SSL methods for 3D point clouds follow either reconstruction or distillation paradigms. Reconstruction-based approaches (Abdelsamad et al. 2025; Zhang et al. 2023) often rely on low-level geometry,

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

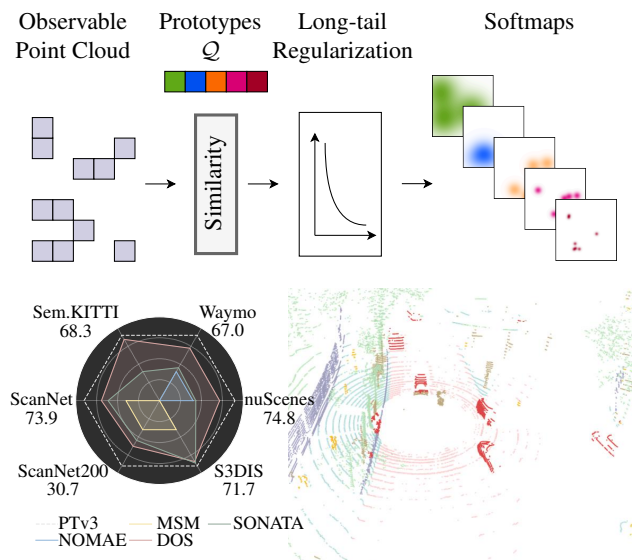


Figure 1: *Overview of DOS*. *Top*: *DOS* learns label-free 3D representations by distilling semantic softmaps at visible points, with prototype similarity and long-tail regularization. *Left*: *DOS* under linear probing outperforms prior SSL methods on six 3D segmentation benchmarks, learning strong representation in both indoor and outdoor scenes. *Right*: Provided with five samples only, a frozen *DOS* successfully segments the objects in a completely new domain.

encouraging shortcut learning rather than semantic understanding. Distillation methods that leverage 2D vision models (Abou Zeid et al. 2025) inherit architectural biases and struggle to exploit 3D-specific structure. Recent masked self-distillation methods (Wu et al. 2025; Hermosilla, Stippel, and Sick 2025) aim to learn semantic representations, but suffer from *positional leakage*, where masked regions inadvertently influence supervision. Clustering-based objectives introduce additional instability, especially in ambiguous or imbalanced regions, while feature regression treats all points equally, ignoring semantic salience.

To address these limitations, we propose *DOS* (Distilling Observable Softmaps), a novel SSL framework that combines *observable self-distillation* with *semantic softmaps*. Ob-

servable self-distillation supervises only visible (unmasked) points, mitigating leakage and ensuring features are deduced solely from observable geometry. Softmaps encode prototype activations across space, encouraging point-wise semantic competition and providing richer gradients than clustering or feature regression. This formulation allows the model to learn nuanced spatial variation and better capture rare semantic concepts. Despite improved stability, softmaps can suffer from degraded semantic diversity. While centering strategies help regularize and mitigate this, they also enforce uniformity, which misaligns with the naturally long-tailed distribution of 3D semantics. To address this, we introduce *Zipf-Sinkhorn*, a frequency-aware assignment strategy that replaces the uniform prior with a Zipfian distribution, better reflecting real-world semantics and yielding more robust, diverse representations.

We evaluate DOS on semantic segmentation and object detection tasks across six diverse benchmarks, covering both indoor and outdoor domains, and demonstrate that it consistently outperforms prior methods, with strong few-shot and cross-domain transfer capabilities. To summarize, our main contributions are as follows:

- We introduce *observable self-distillation*, which eliminates positional leakage by restricting supervision to unmasked points.
- We propose *semantic softmaps* as a distillation target, promoting inter-point competition and rich spatial reasoning.
- We develop *Zipf-Sinkhorn*, a prototype assignment scheme aligned with the long-tailed nature of real-world semantics.
- We release a general-purpose LiDAR backbone pretrained across multiple datasets, enabling strong transfer across 3D domains, and systematically study its generalization capabilities through cross-domain and few-shot evaluations.

DOS achieves state-of-the-art results across five 3D benchmarks—reaching up to *95% of supervised performance* under linear probing, *surpassing supervised baselines* under full fine-tuning, and significantly improving 3D object detection, especially in low-label settings. It transfers effectively to unseen domains without adaptation. Additionally we conduct a thorough ablations that confirm the effectiveness of each proposed component.

Related Work

SSL for 3D Scene Understanding. Self-supervised learning (SSL) has become a key technique for learning 3D representations without manual annotations, particularly for scene understanding with LiDAR and RGB-D data. However, the irregular structure of point clouds demands SSL methods distinct from those in 2D vision. Early 3D SSL methods adopt masked modeling (Yu et al. 2022; Zhang et al. 2023), where models reconstruct occluded regions from visible geometry. These approaches often leak positional cues through masked tokens, enabling shortcut learning that limits semantic abstraction. NOMAE (Abdelsamad et al. 2025) mitigates leakage by reconstructing local neighborhoods, but remains

sensitive to density and susceptible to geometric shortcuts. Recent methods (Wu et al. 2025; Hermosilla, Stippel, and Sick 2025) adopt self-distillation for semantic abstraction, yet face two key limitations: supervision still relies on masked tokens, leaving positional leakage unresolved; and clustering-based objectives assume uniform prototype usage, hindering long-tail semantic modeling. We address both challenges in this work.

Structured Relevance and Spatial Supervision. A longstanding line of research in vision explores the use of spatially structured signals to improve representation learning. In 2D settings, saliency maps and class activation techniques (Selvaraju et al. 2017) highlight task-relevant regions to guide attention. Other approaches (Wang and Gupta 2017; Zhang et al. 2020; Lang et al. 2023, 2024a) leverage spatial relevance maps to weight supervision, encouraging the network to focus on semantically informative areas. Relevance modeling remains underexplored in self-supervised 3D learning despite its success in supervised settings.

Long-Tail Distribution in 3D Data. Real-world 3D scenes exhibit imbalanced semantics, with a few frequent and many rare categories (Çanakçı et al. 2025; Hindel et al. 2025). While long-tail learning has been studied in supervised settings (Lin et al. 2017; Tao and et al. 2020; Mohan et al. 2025), it remains underexplored in self-supervised learning, where most methods assume uniform class or prototype distributions. SwAV (Caron et al. 2020) mitigates collapse by enforcing balanced prototype usage via the Sinkhorn algorithm, but this clashes with the natural frequency imbalance. Zipf’s law (Clauset, Shalizi, and Newman 2009), which captures power-law distributions common in vision and language, offers a more realistic semantic prior. In this work, we explicitly integrate such a prior into a self-supervised 3D framework.

Method: Distilling Observable Softmaps

In this section, we begin by outlining the general student-teacher framework shared with recent masked distillation methods. We then introduce our key contributions: *observable*-point distillation to avoid masked token leakage, *Softmap* loss for a richer learning signal, and *Zipf-Sinkhorn* regularized softmaps to handle long-tail semantics.

General Framework

We adopt a standard masked self-distillation framework based on a student-teacher architecture, commonly used in recent 3D SSL approaches (Wu et al. 2025; Hermosilla, Stippel, and Sick 2025). As illustrated in Figure 2, a point cloud $\mathcal{P} = \{(\mathbf{x}_i, \mathbf{f}_i)\}_{i=1}^N$, where $\mathbf{x}_i \in \mathbb{R}^3$ denotes the spatial coordinates and $\mathbf{f}_i \in \mathbb{R}^d$ the feature of point i , is randomly cropped and duplicated into two different views, $\mathcal{P}^{(1)}$ and $\mathcal{P}^{(2)}$. Each view undergoes separate spatial and photometric augmentations while preserving the original point coordinates to enable cross-view correspondence. For each augmented view, we apply a random point-wise mask by selecting a visible indices subset $\mathcal{I}_v \subset \{1, \dots, N\}$, yielding the visible input $\mathcal{P}_v = \{(\mathbf{x}_i, \mathbf{f}_i)\}_{i \in \mathcal{I}_v}$. The student network processes only \mathcal{P}_v , while

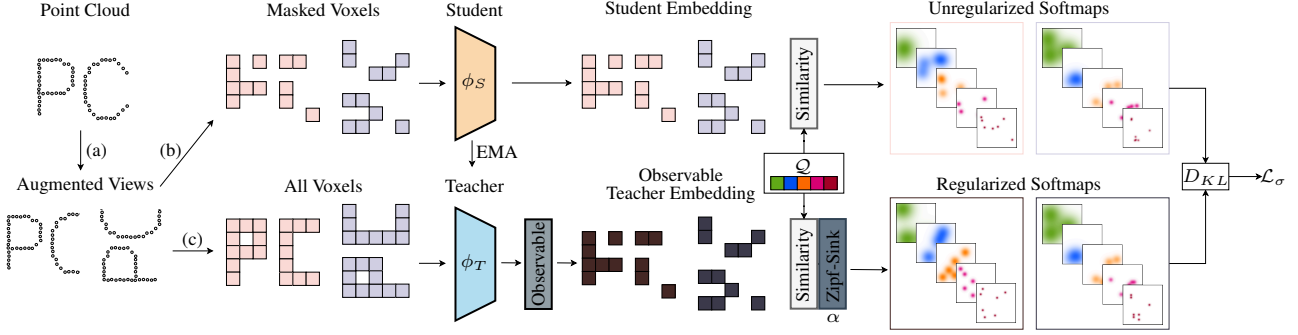


Figure 2: Overview of the proposed *DOS* framework. (a) A point cloud is augmented into two views. (b) The views are masked, voxelize and passed through the student; (c) the corresponding unmasked views go through the teacher, and outputs are filtered to retain only observable voxels. For clarity, only same-view supervision is shown. Both student and teacher compute similarity to prototypes \mathcal{Q} , followed by spatial normalization into softmaps. Teacher softmaps are regularized with Zipf-Sinkhorn to reflect long-tail semantics. The student learns to match these via KL divergence.

the teacher processes the full point cloud \mathcal{P} . Both teacher and student share the same network architecture, and the teacher’s weights are updated via an exponential moving average (EMA) of the student parameters.

This architectural design provides a modular setup for masked distillation and is shared with recent methods. However, it also inherits several limitations: reliance on positional cues from masked tokens, supervision focusing on point similarity, and unawareness of long-tail distributions. We address these issues with three core changes: (1) we distill only from observable (unmasked) points to avoid information leakage, (2) we introduce *Softmap*, a soft relational objective that distills prototype-specific softmaps across points, and (3) regularizing with *Zipf-Sinkhorn*, an optimal transport method aligned with real-world frequency distributions.

Observable Point Distillation

Standard masked distillation methods supervise the student at masked tokens using feature regression (Hermosilla, Stippel, and Sick 2025) or prototype contrastive losses (Wu et al. 2025). However, recovering these masked tokens often relies on their positional embeddings to infer features, leading to information leakage and shortcut learning. To avoid this, we discard masked tokens entirely and apply supervision only at observable (unmasked) points, the subset actually seen by the student. The teacher, in contrast, processes the full input \mathcal{P} and produces targets over the visible subset \mathcal{I}_v . Importantly, since the teacher’s output reflects the full context while the student only sees partial input, the student is implicitly encouraged to reason about the missing (masked) regions. This encourages context-aware learning without access to positional cues from the masked areas.

Softmap: Semantic Softmaps Distillation

Most existing distillation methods, whether based on feature regression (Hermosilla, Stippel, and Sick 2025; Abou Zeid et al. 2025) or soft clustering (Wu et al. 2025), align teacher and student outputs at each point independently. While ef-

fective, such objectives overlook a key aspect of semantic understanding: the relative importance of different points for a given concept. Instead of comparing distributions across features at fixed locations, we supervise the student based on how each prototype’s activation is distributed across points. Specifically, the student learns to predict a *semantic softmap*, a normalized relevance map indicating where a prototype is active. The target is the teacher’s normalized activation across the visible points, reframing distillation as a distribution-matching problem across space rather than features.

Formally, let $\phi_T(\mathcal{P})$ and $\phi_S(\mathcal{P}_v)$ be the teacher and student embeddings, and $\mathcal{Q} = \{q_k\}_{k=1}^K$ the learnable prototype set. We compute cosine similarity between point embeddings and prototypes via:

$$s_{ik}^T = \exp\left(\frac{\cos(\phi_T(\mathcal{P})_i, q_k)}{\tau_T}\right),$$

$$s_{ik}^S = \exp\left(\frac{\cos(\phi_S(\mathcal{P}_v)_i, q_k)}{\tau_S}\right),$$

where τ_T, τ_S are temperature parameters, $i \in \mathcal{I}_v$ indexes visible points, and k indexes prototypes. To construct softmaps, these similarities are normalized across visible points:

$$S_T(i, k) = \frac{s_{ik}^T}{\sum_{j \in \mathcal{I}_v} s_{jk}^T}, \quad S_S(i, k) = \frac{s_{ik}^S}{\sum_{j \in \mathcal{I}_v} s_{jk}^S}.$$

This contrasts with clustering-based soft targets normalization $S_T^c(i, k) = s_{ik}^T / \sum_{k'=1}^K s_{ik'}^T$, where scores are normalized across prototypes for each point independently. Next, we apply Zipf-Sinkhorn to regularize the teacher’s softmaps using a long-tailed prior, yielding $\tilde{S}_T(i, k)$ (see Section Zipf-Sinkhorn Regularization). The Softmap loss is defined as a KL divergence:

$$\mathcal{L}_\sigma(\mathcal{P}, \mathcal{P}_v, \mathcal{Q}) = \frac{1}{K} \sum_{k=1}^K \text{KL}(\tilde{S}_T(:, k) \| S_S(:, k))$$

$$= -\frac{1}{K} \sum_{k=1}^K \sum_{i \in \mathcal{I}_v} \tilde{S}_T(i, k) \log S_S(i, k)$$

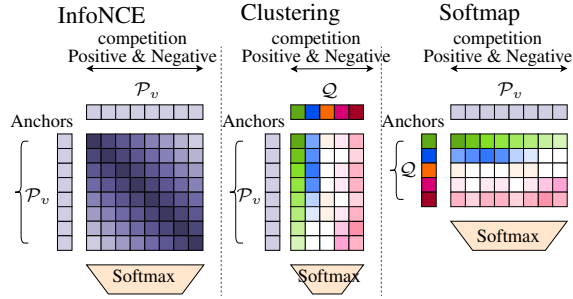


Figure 3: Conceptual comparison of contrastive strategies. *Left*: InfoNCE contrasts positive and negative pairs via pairwise similarity. *Middle*: Clustering aligns points to prototypes but lacks inter-point competition. *Right*: Softmap (DOS) distills dense soft targets and applies softmax across points, encouraging spatially structured, concept-aware representations. Softmap distillation yields stronger features empirically.

While Softmap loss does not rely on explicit positive-negative pairs, it shares a key property with InfoNCE (Oord, Li, and Vinyals 2018): each prototype induces competition across points, treating them as soft positives and negatives with respect to a semantic concept. Unlike clustering-based methods that normalize over prototypes, Softmap distillation applies softmax across points (see Fig. 3), shifting the objective from point-to-prototype matching to spatial reasoning. This formulation yields richer gradients, encourages spatially grounded representations, and allows even weakly activated prototypes to influence learning.

For each student view $a \in \{1, 2\}$, we apply the Softmap loss to align the student’s predicted softmaps with Zipf-regularized teacher softmaps from both the same-view and cross-view teachers:

$$\mathcal{L}_a = \frac{1}{2} \left(\mathcal{L}_\sigma(\mathcal{P}^{(a)}, \mathcal{P}_v^{(a)}, \mathcal{Q}) + \mathcal{L}_\sigma(\mathcal{P}^{(\bar{a})}, \mathcal{P}_v^{(a)}, \mathcal{Q}) \right),$$

where $\mathcal{P}_v^{(a)}$ is the masked student input, $\mathcal{P}^{(a)}$ and $\mathcal{P}^{(\bar{a})}$ are full inputs to the same-view and cross-view teachers, and \bar{a} denotes the opposite view index. The total training loss is $\mathcal{L}_{\text{total}} = \mathcal{L}_1 + \mathcal{L}_2$. This setup enforces semantic consistency between views by aligning their softmaps at observable student points. Which promotes high-level, spatially structured representations without explicit reconstruction of masked tokens.

Zipf-Sinkhorn Regularization

Softmaps provide richer gradients and a stronger inductive bias than clustering, enabling stable training even without explicit centering. However, performance remains suboptimal due to two failure modes: (1) semantic prototype collapse, where multiple prototypes redundantly activate on the same regions; and (2) point-wise collapse, where parts of the scene are left unrepresented. To address this, we aim for prototypes that (i) activate in distinct regions and (ii) collectively provide full scene coverage. While standard Sinkhorn-Knopp promotes these goals, it additionally enforces uniform prototype usage, a rigid constraint misaligned with real-world 3D

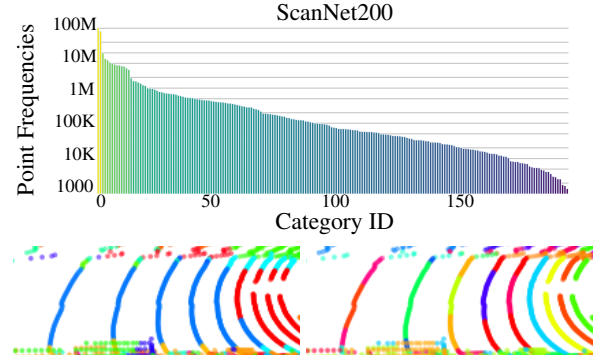


Figure 4: Zipfian Distribution and Prototype Assignment. *Top*: Category frequencies in ScanNet200 follow a Zipf-like distribution, illustrating the long-tail nature of 3D semantics. *Bottom*: Prototype activation maps after pretraining. *Left*: Zipf prior yields consistent activations in frequent regions (e.g., road) and reduces fragmentation. *Right*: Uniform prior enforces equal usage, leading to over-segmentation and learning biased toward low-level geometry (e.g., ego distance).

Algorithm 1: Zipf-Sinkhorn for Softmaps Regularization

- Require:** Teacher similarities matrix $F \in \mathbb{R}^{N_v \times K}$, Zipf exponent α , iterations T
- Ensure:** Zipf-aware softmap $\tilde{S}_T \in \mathbb{R}^{N_v \times K}$
- 1: **Compute Zipf Prior:**
 - 2: $w_k \propto \frac{1}{k^\alpha}$ for $k = 1, 2, \dots, K$, then normalize $\mathbf{w} \leftarrow \mathbf{w} / \sum_k w_k$
 - 3: **Normalize Similarities:**
 - 4: Normalize: $F \leftarrow F / \sum_{i,k} F_{i,k}$
 - 5: **Iterative Normalization (Sinkhorn):**
 - 6: **for** $t = 1$ to T **do**
 - 7: Normalize rows: $F_{i,:} \leftarrow F_{i,:} / \sum_k F_{i,k} \quad \forall i$
 - 8: Normalize columns to match \mathbf{w} : $F_{:,k} \leftarrow F_{:,k} \cdot \mathbf{w}_k / \sum_i F_{i,k} \quad \forall k$
 - 9: **end for**
 - 10: **Column Normalize for Prototype activation:**
 - 11: $\tilde{S}_T(i, k) \leftarrow F_{i,k} / \sum_j F_{j,k}$
 - 12: **return** \tilde{S}_T
-

semantics, which follow a Zipfian distribution (Figure 4, top). As a result, uniform balancing tends to over-segment frequent structures and assign prototypes based on low-level cues like distance, reducing semantic fidelity (Figure 4, bottom right).

To address the mismatch between uniform prototype usage and real-world semantics, we propose *Zipf-Sinkhorn*, a modified optimal transport step that incorporates a power-law prior over prototypes. Instead of enforcing a uniform marginal, we assign a Zipfian prior $\pi \in \mathbb{R}^K$, where $\pi_k \propto \frac{1}{k^\alpha}$ and $\alpha > 0$ controls the sharpness. This reflects the semantic frequency imbalance in 3D data where frequent prototypes receive broader activation, while rare ones remain sharp and selective.

Let $F \in \mathbb{R}^{N_v \times K}$ be the similarity matrix between N_v visi-

Method	nuScenes		Waymo		SemKITTI		ScanNet		ScanNet200		S3DIS Area 5	
	mIoU	mAcc	mIoU	mAcc	mIoU	mAcc	mIoU	mAcc	mIoU	mAcc	mIoU	mAcc
PTv3 (2023)	80.4	87.2	71.3	80.5	69.1	76.1	77.6	85.0	35.3	46.0	73.4	78.9
MSM (2025)	-	-	-	-	-	-	68.7	-	26.8	-	59.5	-
NOMAE (2025)	65.1	77.9	59.2	69.9	-	-	-	-	-	-	-	-
Sonata* (2025)	66.1	77.2	60.5	72.5	62.0	72.5	72.5	83.1	29.3	41.6	72.3	81.2
DOS (ours)	74.1	84.8	66.1	77.1	67.5	78.1	72.8	83.3	29.1	41.1	70.6	79.2
DOS* (ours)	74.8	84.2	67.0	77.7	68.3	78.4	73.9	83.5	30.7	41.7	71.7	81.4
D-DITR (2025)	80.7	-	72.1	-	69.8	-	79.2	-	37.7	-	75.0	-
MSM (2025)	-	-	-	-	-	-	78.5	-	35.7	-	73.2	-
NOMAE (2025)	81.8	87.7	72.3	82.5	-	-	-	-	-	-	-	-
Sonata* (2025)	<u>81.7</u>	87.9	72.9	81.9	72.6	77.9	<u>79.4</u>	86.1	36.8	46.5	76.0	81.6
DOS (ours)	81.5	87.6	<u>73.3</u>	<u>83.8</u>	<u>73.1</u>	<u>81.0</u>	78.7	<u>86.2</u>	36.7	<u>46.6</u>	74.2	83.6
DOS* (ours)	81.8	<u>87.8</u>	73.9	83.9	73.5	81.3	79.7	86.8	<u>37.1</u>	46.8	<u>75.1</u>	83.8

Table 1: Comparison of semantic segmentation performance on the validation split of multiple indoor and outdoor point cloud datasets. Methods marked with * use additional data. For finetuning, we **bold** the best result and underline the second-best.

ble points and K prototypes. We apply an entropy-regularized Sinkhorn procedure (Algorithm 1) to iteratively normalize F such that its columns match π and its rows remain uniform. After convergence, columns are renormalized to yield the final Zipf-aware softmap \tilde{S}_T . This simple modification does not introduce additional computational cost beyond computing the Zipf prior once at initialization. However, by aligning supervision with long-tailed semantics, Zipf-Sinkhorn improves representational diversity during pretraining.

Implementation Details

We use PTv3 (Wu et al. 2023) with layer normalization as the default encoder. Pretraining is done on $2 \times A100$ GPUs with a batch size of 16 and completes in ~ 20 hours, depending on dataset size. A 70% masking ratio is applied uniformly for semantic segmentation experiments, while for object detection we use a 60% masking ratio. We use mask block sizes of 40 cm for indoor and 1 m for outdoor datasets. We apply the pretraining supervision at a voxel size of 0.08 m for indoor scenes and 0.2 m for outdoor scenes. To obtain voxel features at these resolutions, we follow Sonata (Wu et al. 2025) and upcast features from deeper, coarser layers to the target voxel grid, then concatenate them with the features computed at that resolution. We use 1024 prototypes for all experiments. For semantic segmentation, we append a lightweight PTv3 decoder; for object detection, we use the pretrained encoder as a backbone for a CenterPoint detector (Yin, Zhou, and Krahenbuhl 2021). During fine-tuning, we follow the standard settings used in the respective supervised counterparts. Full details are provided in the supplementary material.

Experimental Evaluation

Semantic Segmentation Results

Table 1 presents semantic segmentation results on six 3D benchmarks, spanning diverse environments and sensors. For autonomous driving, we evaluate on nuScenes (Fong et al. 2022), Waymo (Sun et al. 2020), and SemanticKITTI (Behley et al. 2019); for indoor scene understanding, we use ScanNet (Dai et al. 2017), ScanNet200 (Hu, Dai, and Nießner

2020), and S3DIS (Armeni et al. 2016). We report performance using two protocols: (i) linear probing (LP) with a frozen encoder and linear head, and (ii) fine-tuning (FT) with full encoder-decoder optimization. Additional dataset details are provided in the supplementary material.

DOS achieves consistently strong results across all benchmarks, outperforming prior SSL approaches under both LP and FT. Notably, *DOS reaches up to 95% of the supervised performance under LP* and surpasses the supervised PTv3 baseline under fine-tuning on every dataset. Compared to NOMAE (Abdelsamad et al. 2025), DOS achieves substantially higher LP performance on all outdoor datasets, indicating that self-distillation yields more semantic features than geometry-centric reconstruction. Additionally, DOS outperforms D-DITR (Abou Zeid et al. 2025) across nearly all metrics, despite D-DITR using multi-view images and vision-based distillation.

Compared to Sonata (Wu et al. 2025) and MSM (Hermosilla, Stippel, and Sick 2025), both of which use masked self-distillation, DOS offers consistent improvements across LP and FT. For example, on Waymo and nuScenes, DOS improves over Sonata by +6.5 and +8.0 mIoU (LP), and by +0.7 and +0.1 mIoU (FT). These improvements stem from DOS’s key innovations: *observable self-distillation*, *softmap supervision*, and *Zipf-Sinkhorn regularization*, which together mitigate information leakage, reduce prototype collapse, and promote semantically diverse learning. With additional training data (denoted as DOS*), the performance further improves across all settings, achieving state-of-the-art SSL results on all six benchmarks.

Object Detection Results

This section presents the object detection performance on the nuScenes validation dataset. Table 2 evaluates the nuScenes detection score (NDS) and mean average precision (mAP) under two settings: full-data, where both SSL pretraining and fine-tuning use the complete training set, and a commonly adopted data-efficient regime, where fine-tuning is conducted with only 20% of the labeled frames and rele-

Methods	NDS	mAP
<i>20% Labeled Frames (Data-Efficient Setting)</i>		
ALSO (Boulch et al. 2023)	48.2	41.2
GD-MAE (Yang et al. 2023a)	48.8	42.6
Learning from 2D (Liu et al. 2021)	49.2	48.8
UniPAD (Yang et al. 2023b)	55.8	48.1
NOMAE (Abdelsamad et al. 2025)	60.9	54.4
DOS (ours)	62.1	57.1
<i>Full Data (100% Labeled)</i>		
CenterPoint (no pre-training)	65.4	57.6
CenterPoint + DOS (ours)	69.7	65.5

Table 2: Object detection on nuScenes val set under full-data (complete training set) and 20%-label (fine-tuning with 20% labeled frames). DOS improves performance in both regimes.

Training dataset	nuScenes		Waymo		SemKITTI	
	mIoU	mAcc	mIoU	mAcc	mIoU	mAcc
nuScenes	74.1	84.8	55.5	68.7	57.3	68.5
Waymo	66.2	78.1	<u>66.1</u>	<u>77.1</u>	64.1	74.8
SemKITTI	59.5	72.5	58.5	71.4	<u>67.5</u>	<u>78.1</u>
Nu-SK-Wa	74.8	84.2	67.0	77.7	68.3	78.4

Table 3: Cross-domain segmentation performance of DOS when pretrained on individual datasets (rows) and evaluated on different domains (columns).

vant baselines are reported. DOS-pretrained model improves 3D object detection on nuScenes in both low-resource and full-data regimes. In the 20%-label setting, DOS surpasses all baselines on mAP and NDS. In the full-data regime, integrating DOS with CenterPoint significantly boosts performance, yielding a +4.3 NDS and +7.9 mAP improvement over training from scratch, highlighting the transferability of DOS to detection tasks.

Cross-Domain Transfer Analysis

We evaluate cross-domain generalization by pretraining DOS on nuScenes, Waymo, or SemanticKITTI, then evaluating it on the other two. Table 3 shows that even without joint training, DOS exhibits strong transfer. A model trained solely on Waymo surpasses all other methods, including Sonata trained on all datasets, when evaluated on nuScenes and SemanticKITTI, highlighting the robustness of the learned features. We then train a single DOS model on all three datasets combined. This unified model performs consistently well across domains, confirming DOS’s suitability as a general-purpose 3D backbone.

To assess generalization to unseen domains, we test on the ParisLuco (Sanchez et al. 2024) dataset. In the zero-shot setting, we fine-tune on nuScenes or SemanticKITTI and evaluate directly on ParisLuco. In Table 4, we observe that DOS significantly outperforms a supervised-from-scratch baseline using the same labels. In the few-shot case, fine-tuning with only five labeled scenes yields state-of-the-art performance. Figure 5 compares zero-shot transfer from SemanticKITTI to ParisLuco3D. Despite no adaptation, DOS yields notably

# Shots	setting	nuScenes	SemKITTI	CoLa
0	supervised	6.9	18.0	37.5
0	DOS(ours)	50.8	43.8	63.9
5	DOS(ours)	51.5	41.9	85.9

Table 4: ParisLuco evaluation under zero-shot and few-shot settings. Zero-shot: model fine-tuned on nuScenes or SemanticKITTI and evaluated directly on ParisLuco. Few-shot: five annotated ParisLuco scenes used.

Model	mIoU	mACC	ACC
Masked self-distillation			
naive	54.7	66.7	88.9
token jitter	55.1	68.3	89.2
observable (ours)	69.3	81.7	93.4
Training target			
clustering	69.3	81.7	93.4
feature regression	63.0	76.6	92.2
softmax (ours)	72.3	83.4	94.0
+ Zipf prior			
uniform (alpha = 0)	72.3	83.4	94.0
alpha = 1.3	74.1	84.8	93.7

Table 5: Ablation on DOS components, evaluated via LP on the nuScenes val set (semantic segmentation).

improved segmentation over the supervised PTv3 baseline, demonstrating stronger generalization to out-of-distribution urban scenes. To support downstream adoption, we release our multi-dataset pretrained weights as a general-purpose LiDAR backbone for cross-domain transfer.

Ablation Study

We conduct ablation studies to evaluate the design choices behind DOS and justify each component’s contribution. In addition to our default PTv3 backbone, we assess DOS with an alternative architecture to test generality. Further analysis on centering strategies, label-efficient settings, and prototype count is presented in the supplementary material.

Evaluating Components of DOS. We evaluate the contribution of each component in DOS via an incremental ablation study on the nuScenes dataset, using LP for semantic segmentation after 50 epochs of pretraining. All models use a PTv3 backbone with consistent hyperparameters, including a 70% masking ratio and fixed block size. We begin with a standard *masked self-distillation* baseline that supervises the student via an online clustering loss. This naive setup yields 54.7 mIoU and suffers from positional leakage through masked tokens, as also observed by prior works (Wu et al. 2025; Hermosilla, Stippel, and Sick 2025). Token jittering (Wu et al. 2025), proposed as a mitigation, improves this only slightly to 55.1 mIoU. Replacing this with our proposed *observable self-distillation*, which discards masked tokens and supervises only visible voxels, leads to a substantial improvement of +14.6 mIoU, achieving 69.3 mIoU. Unlike Sonata, our ap-

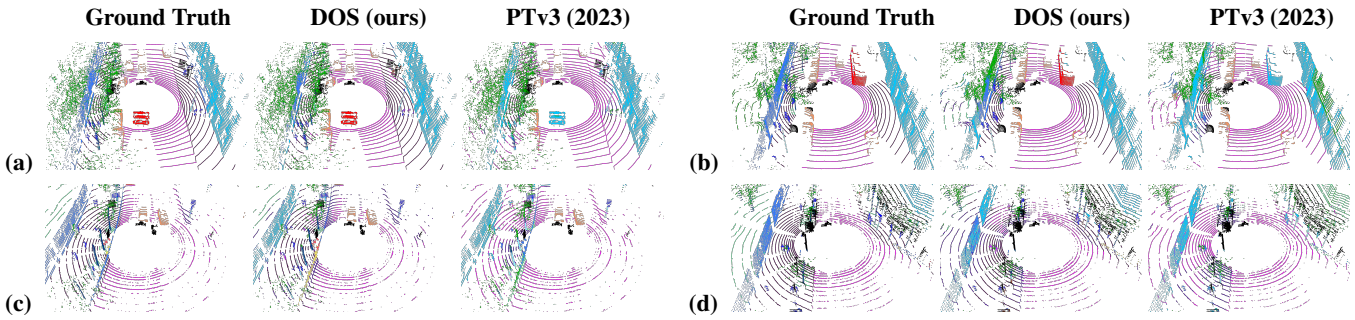


Figure 5: Zero-shot transfer to ParisLuco3D from SemanticKITTI using the same PTv3 (2023) backbone. DOS shows superior generalization, recovering buses (a, b), reducing wall/building confusion, and separating road from sidewalk. Remaining issues include railing–bike confusion (c) and similar classes swapping (d).

Method	0.1%	1%	10%	50%	100%
Supervised	28.2	41.6	68.7	78.7	80.4
NOMAE (2025)(ft)	35.8	48.1	69.9	80.1	81.8
DOS (ours, ft)	38.8	53.1	71.5	80.5	81.8

Table 6: Results with varying amounts of annotated data, evaluated for semantic segmentation mIoU on the nuScenes val set.

Method	Limited Scenes					Limited Annotations (pts)				
	1%	5%	10%	20%	100%	20	50	100	200	full
Supervised	25.8	48.9	61.0	67.0	77.2	60.1	67.9	71.4	72.7	77.2
Sonata (2025)(lin)	43.6	62.5	68.6	69.8	72.5	69.0	70.5	71.1	71.5	72.5
DOS (ours, lin)	45.8	63.4	69.7	70.9	72.8	69.4	70.8	71.4	71.7	72.8

Table 7: Results with varying amounts of annotated data, evaluated for semantic segmentation mIoU on the ScanNet val set.

Zipf Exponent α	0.0	0.1	0.3	0.6	0.9	1.3	1.6	2.0	3.0
ScanNet (mIoU)	71.9	72.2	71.9	72.3	72.5	72.8	72.4	71.9	69.8
ScanNet200 (mIoU)	27.9	28.1	28.2	28.6	28.7	29.1	29.2	29.0	27.7

Table 8: Ablation of Zipf exponent α on ScanNet and ScanNet200. Reported scores are mIoU under LP.

α	Head (66)	Common (68)	Tail (66)
0.0	50.4	20.5	10.6
1.3	50.8	23.5	13.2

Table 9: Effect of Zipf exponent α on ScanNet200 head/common/tail splits (66/68/66 classes). Reported scores are mIoU under LP.

proach does not rely on auxiliary heads or additional labeled data, yet achieves significantly higher performance under this restricted distillation setup.

Next, we assess the impact of the distillation target. Keeping the observable setup fixed, we replace clustering with a *feature regression* loss (e.g., cosine similarity), which degrades performance to 63.0 mIoU. By contrast, our proposed *softmax distillation* boosts results to 72.3 mIoU, a +3.0 improvement over clustering, demonstrating the benefits of spatially-aware supervision and inter-point competition. Finally, we integrate a Zipfian prior over prototype usage via *Zipf-Sinkhorn*, yielding a further improvement from 72.3 to 74.1 mIoU. This suggests that aligning prototype assignments with a Zipfian prior, by better capturing the long-tail distribution of 3D semantics, enhances semantic diversity and representation quality. We adopt this final setup in all main experiments.

Label-Efficiency We evaluate DOS under limited supervision settings to assess its robustness in data-scarce regimes. Table 6 and Table 7 report semantic segmentation results on nuScenes and ScanNet across various label budgets. On both benchmarks, DOS consistently outperforms supervised training from scratch and prior SSL methods such as Sonata

and NOMAE. In the linear probing setting, DOS achieves strong performance even with only 1% of labeled scenes, and reaches parity with full supervision at just 20% labeling. In the fine-tuning setting, DOS-pretrained models yield the highest accuracy across all label fractions, including extremely low-resource cases (e.g., 0.1% on nuScenes). These results demonstrate that DOS learns semantically rich and transferable representations that enable effective downstream adaptation with minimal annotation.

Effect of Zipf Exponent. We evaluate the impact of the Zipf exponent α in Zipf-Sinkhorn regularization on ScanNet and ScanNet200, which share geometry but differ in label granularity (20 vs. 200 classes). This setup helps us assess how prior sharpness influences representation quality across coarse and fine semantic spaces. As shown in Table 8, moderate values of α yield the best performance: mIoU peaks at $\alpha = 1.3$ on ScanNet (72.8) and $\alpha = 1.6$ on ScanNet200 (29.2). Lower α values (≈ 0) enforce overly uniform usage, while higher values (≥ 2) degrade performance and cause training instability due to prototype under-utilization and weakened gradients. To better understand where the gains on ScanNet200 come from, we further break down performance

using the standard head/common/tail split of 66/68/66 classes based on class frequency (Table 9). When increasing α from 0.0 to 1.3, performance on head classes remains nearly unchanged (50.4 \rightarrow 50.8 mIoU), whereas common and tail classes improve substantially (20.5 \rightarrow 23.5 and 10.6 \rightarrow 13.2 mIoU, respectively). This indicates that mild Zipf priors mainly help by strengthening representations for medium-frequency and rare categories rather than further boosting already well-represented head classes. These results align with empirical Zipfian statistics observed in real-world 3D datasets, and support the use of mild Zipf priors to balance semantic coverage, specialization, and stability.

Effect of Prototype Count. We study how the number of prototypes affects representation quality for three objectives: clustering, softmap distillation, and softmap with Zipf-Sinkhorn. As shown in Table 10, softmap-based objectives consistently outperform clustering across all settings, with clear gains even at low prototype counts. This highlights the benefit of modeling contextual relevance and inter-point competition in softmaps compared to the point-wise, independent clustering assignments. Softmap variants also saturate early, typically between 1024 and 4096 prototypes, indicating better training efficiency and stability. In contrast, clustering keeps improving with more prototypes but remains worse overall, and further gains would likely require impractically large cluster counts due to memory constraints.

Prototype Count	Clustering	Softmap	Softmap + Zipf
<i>nuScenes</i>			
32	65.1	70.9	70.2
128	69.5	69.8	73.3
1024	68.7	72.1	74.1
4096	69.6	72.3	74.0
<i>ScanNet</i>			
32	59.6	66.1	65.0
128	62.7	70.8	70.6
1024	66.7	71.5	72.8
4096	69.3	71.7	72.9

Table 10: LP mIoU (%) on nuScenes and ScanNet for varying prototype counts and objectives.

Dataset	w/o Cross-View	Full DOS (Ours)
nuScenes	74.3	74.1
Waymo	63.5	66.1
ScanNet	61.8	72.8

Table 11: Effect of cross-view supervision on linear probing performance (mIoU %) across datasets.

Effect of Cross-View Supervision. We ablate cross-view supervision in DOS by training a variant without the view alignment objective. For a fair comparison, we double the batch size and number of training epochs so that the total compute matches the default setup, and evaluate on nuScenes,

Fine-tuning		Linear Probing	
Method	mIoU	SPUNet	mIoU
scratch	73.3	DOS	55.6
UniPAD	79.4	+ double channels	57.4
NOMAE	80.1	+ distill from PTv3	71.5
DOS	80.2	PTv3 + DOS (ref.)	74.8
PTv3 + DOS (ref.)	81.8		

Table 12: Transferability of DOS across backbone architectures. We report mIoU on nuScenes under LP and FT.

Waymo, and ScanNet with linear probing. As shown in Table 11, the effect of cross-view supervision is dataset dependent: on ScanNet, with dense indoor point clouds, removing cross-view alignment causes a substantial performance drop, while on sparser outdoor datasets such as nuScenes and Waymo, the effect is minimal or slightly negative. This suggests that cross-view supervision is most beneficial when point cloud density supports meaningful spatial alignment across views.

Transferability Across Architectures. We evaluate DOS on SPUNet to test generality beyond PTv3, using the same pretraining and evaluation. As shown in Table 12, DOS significantly improves SPUNet finetuning over training from scratch, confirming its value as a general purpose initialization. However, LP yields weaker representations than PTv3, suggesting sensitivity to backbone capacity. To investigate further, we double SPUNet channel width, which improves LP scores but still falls short of PTv3. When PTv3 is unavailable or too heavy, we recommend distilling from a pretrained DOS encoder. With just 10 epochs of distillation from a frozen DOS model, SPUNet outperforms its self distilled version, which highlights an efficient path to lightweight 3D backbones.

Conclusion

We introduced *DOS*, an SSL framework for 3D point clouds that addresses key limitations of existing methods. DOS builds on three core innovations: (1) *observable self-distillation*, which avoids positional leakage by supervising only visible points, (2) *semantic softmaps*, which enhance representation learning through point-wise competition and relevance modeling, and (3) *Zipf-Sinkhorn*, a prototype assignment strategy aligned with the naturally long-tailed distribution of real-world semantics. Together, these components enable DOS to learn high-quality 3D features without labels. DOS outperforms SOTA across five challenging 3D benchmarks and shows strong generalization in cross-domain and few-shot settings. We release a general-purpose LiDAR backbone pretrained across multiple datasets to accelerate progress in 3D scene understanding.

Limitations and Future Work. While DOS currently handles indoor and outdoor domains separately, unifying them within a single pretraining framework remains a promising direction. Moreover, although DOS closely approaches supervised performance under LP, surpassing this baseline is a challenge for future research in 3D representation learning.

References

- Abdelsamad, M.; Ulrich, M.; Gläser, C.; and Valada, A. 2025. Multi-Scale Neighborhood Occupancy Masked Autoencoder for Self-Supervised Learning in LiDAR Point Clouds. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- Abou Zeid, K.; Yilmaz, K.; de Geus, D.; Hermans, A.; Adrian, D.; Linder, T.; and Leibe, B. 2025. DINO in the Room: Leveraging 2D Foundation Models for 3D Segmentation. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- Armeni, I.; Sener, O.; Zamir, A. R.; Jiang, H.; Brilakis, I.; Fischer, M.; and Savarese, S. 2016. 3D Semantic Parsing of Large-Scale Indoor Spaces. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- Behley, J.; Garbade, M.; Milioto, A.; Quenzel, J.; Behnke, S.; Stachniss, C.; and Gall, J. 2019. SemanticKITTI: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Boulch, A.; Sautier, C.; Michele, B.; Puy, G.; and Marlet, R. 2023. ALSO: Automotive Lidar Self-supervision by Occupancy estimation. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- Çanakçı, A. S.; Vödisch, N.; Petek, K.; Burgard, W.; and Valada, A. 2025. Label-Efficient LiDAR Panoptic Segmentation. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. In *NeurIPS*.
- Clauset, A.; Shalizi, C. R.; and Newman, M. E. J. 2009. Power-Law Distributions in Empirical Data. *SIAM Review*, 51(4): 661–703.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- Fong, W. K.; Mohan, R.; Hurtado, J. V.; Zhou, L.; Caesar, H.; Beijbom, O.; and Valada, A. 2022. Panoptic nusenes: A large-scale benchmark for lidar panoptic segmentation and tracking. *IEEE Robotics and Automation Letters*, 7(2): 3795–3802.
- Hermosilla, P.; Stippel, C.; and Sick, L. 2025. Masked Scene Modeling: Narrowing the Gap Between Supervised and Self-Supervised Learning in 3D Scene Understanding. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- Hindel, J.; Mohan, R.; Bratulic, J.; Cattaneo, D.; Brox, T.; and Valada, A. 2025. Label-Efficient LiDAR Semantic Segmentation with 2D-3D Vision Transformer Adapters. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Hu, R.; Dai, A.; and Nießner, M. 2020. ScanNet200: 3D Semantic Scene Understanding for 200 Classes. *arXiv preprint arXiv:2011.12336*.
- Lang, C.; Braun, A.; Schillingmann, L.; Haug, K.; and Valada, A. 2024a. Self-supervised representation learning from temporal ordering of automated driving sequences. *IEEE Robotics and Automation Letters*, 9(3): 2582–2589.
- Lang, C.; Braun, A.; Schillingmann, L.; and Valada, A. 2023. Self-supervised multi-object tracking for autonomous driving from consistency across timescales. *IEEE Robotics and Automation Letters*, 8(11): 7711–7718.
- Lang, C.; Braun, A.; Schillingmann, L.; and Valada, A. 2024b. A point-based approach to efficient LiDAR multi-task perception. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 13261–13268.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Doll’ar, P. 2017. Focal Loss for Dense Object Detection. In *ICCV*.
- Liu, Y.; Huang, Y.; Chiang, H.; Su, H.; Liu, Z. Y.; Chen, C.; Tseng, C.; and Hsu, W. H. 2021. Learning from 2D: Pixel-to-Point Knowledge Transfer for 3D Pretraining. *arXiv preprint arXiv:2104.04687*.
- Mohan, R.; Hindel, J.; Drews, F.; Gläser, C.; Cattaneo, D.; and Valada, A. 2025. Open-Set LiDAR Panoptic Segmentation Guided by Uncertainty-Aware Learning. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31.
- Sanchez, J.; Soum-Fontez, L.; Deschaud, J.-E.; and Goulette, F. 2024. ParisLuco3D: A High-Quality Target Dataset for Domain Generalization of LiDAR Perception. *IEEE Rob. and Auto. Letters*, 9(6): 5052–5059.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 618–626.
- Sun, P.; Kretzschmar, H.; d’Arcy, X.; et al. 2020. Scalability in perception for autonomous driving: Waymo open dataset. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- Tao, Z.; and et al. 2020. Few-Shot Object Detection via Feature Reweighting. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- Wang, X.; and Gupta, A. 2017. Learning to Discover Cross-Modal Semantic Correspondence. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 3627–3635.
- Wu, X.; DeTone, D.; Frost, D.; Shen, T.; Xie, C.; Yang, N.; Engel, J.; Newcombe, R.; Zhao, H.; and Straub, J. 2025. Sonata: Self-Supervised Learning of Reliable Point Representations. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- Wu, X.; Jiang, L.; Wang, P.-S.; Liu, Z.; Liu, X.; Qiao, Y.; Ouyang, W.; He, T.; and Zhao, H. 2023. Point Transformer V3: Simpler, Faster, Stronger. *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 4840–4851.

- Yang, H.; He, T.; Liu, J.; Chen, H.; Wu, B.; Lin, B.; He, X.; and Ouyang, W. 2023a. GD-MAE: Generative Decoder for MAE Pre-training on LiDAR Point Clouds. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- Yang, H.; Zhang, S.; Huang, D.; Wu, X.; Zhu, H.; He, T.; Tang, S.; Zhao, H.; Qiu, Q.; Lin, B.; He, X.; and Ouyang, W. 2023b. UniPAD: A Universal Pre-Training Paradigm for Autonomous Driving. *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 15238–15250.
- Yin, T.; Zhou, X.; and Krahenbuhl, P. 2021. Center-based 3D Object Detection and Tracking. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 11784–11793.
- Yu, Q.; Tang, Y.; Xu, Y.; Zhao, T.; Liu, Y.; Susskind, J.; and Wang, H. 2022. Point-BERT: Pre-Training 3D Point Cloud Transformers with Masked Point Modeling. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- Zhang, K.; Qi, H.; Jia, J.; Chen, R.; and Yu, G. 2023. GeoMAE: Geometry-Aware Masked Autoencoder for Self-Supervised Point Cloud Learning. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*.
- Zhang, X.; Yang, Y.; Qi, Q.; Yuille, A. L.; and Shen, W. 2020. Revisiting Weakly Supervised Pre-training of Visual Transformers. In *European Conference on Computer Vision (ECCV)*, 292–308.