

Exploring Single Domain Generalization of LiDAR-based Semantic Segmentation under Imperfect Labels

Weitong Kong^{1,*}, Zichao Zeng^{2,*}, Di Wen¹, Jiale Wei¹,
Kunyu Peng^{1,†}, June Moh Goo², Jan Boehm², and Rainer Stiefelhagen¹

Abstract—Accurate perception is critical for vehicle safety, with LiDAR as a key enabler in autonomous driving. To ensure robust performance across environments, sensor types, and weather conditions without costly re-annotation, domain generalization in LiDAR-based 3D semantic segmentation is essential. However, LiDAR annotations are often noisy due to sensor imperfections, occlusions, and human errors. Such noise degrades segmentation accuracy and is further amplified under domain shifts, threatening system reliability. While noisy-label learning is well-studied in images, its extension to 3D LiDAR segmentation under domain generalization remains largely unexplored, as the sparse and irregular structure of point clouds limits direct use of 2D methods. To address this gap, we introduce the novel task Domain Generalization for LiDAR Semantic Segmentation under Noisy Labels (DGLSS-NL) and establish the first benchmark by adapting three representative noisy-label learning strategies from image classification to 3D segmentation. However, we find that existing noisy-label learning approaches adapt poorly to LiDAR data. We therefore propose DuNe, a dual-view framework with strong and weak branches that enforce feature-level consistency and apply cross-entropy loss based on confidence-aware filtering of predictions. Our approach shows state-of-the-art performance by achieving 56.86% mIoU on SemanticKITTI, 42.28% on nuScenes, and 52.58% on SemanticPOSS under 10% symmetric label noise, with an overall Arithmetic Mean (AM) of 49.57% and Harmonic Mean (HM) of 48.50%, thereby demonstrating robust domain generalization in DGLSS-NL tasks. The code is available at <https://github.com/MKong17/DGLSS-NL.git>.

I. INTRODUCTION

In autonomous driving, system safety and reliability are paramount. Accurate perception is essential to this goal. LiDAR provides precise 3D geometry and is central to reliable perception [1], [2]. However, the quality of LiDAR data varies across sensors, environments, and operating conditions, which challenges consistent interpretation. Domain generalization in 3D LiDAR semantic segmentation trains

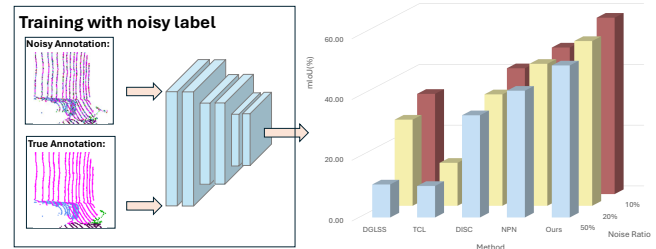


Fig. 1. We inject symmetric label noise into the training set according to predefined noise ratios. The figure reports the segmentation performance on the test set, where the vertical axis shows mIoU(%) (↑), the horizontal axis corresponds to the applied methods, and the third dimension indicates the results under different ratios of symmetric noise. DuNe consistently achieves the best results across all noise levels.

models that must perform in unseen domains without target data [3]. This capability is critical for robust perception and safe decision making in diverse real-world scenarios [4], [5].

Most existing domain generalization methods assume perfect annotations. In practice, 3D LiDAR labels are often imperfect because dense and irregular point sets are difficult to annotate consistently [6]–[8]. Label noise degrades segmentation performance, and the degradation is amplified under domain shift [9], which threatens reliability once deployed. These observations motivate a setting that treats noisy supervision and domain shift jointly.

In the image domain, noisy-label learning has been extensively studied, including loss correction, sample reweighting, and semi-supervised or contrastive strategies. TCL [10], a representative method, applies Gaussian modeling and pseudo-labeling to enable contrastive learning. On the other hand, DISC [11] dynamically selects clean samples while simultaneously exploiting informative signals from noisy ones. NPN [12] accumulates prediction statistics to construct partial and negative label sets, thereby mitigating the effect of overconfident noise. Although these approaches demonstrate strong performance on images, directly transferring them to point clouds is non-trivial. Unlike images, point clouds are sparse, irregular, and orderless [13], [14], which necessitates adaptations that respect the unique characteristics of 3D geometry.

This work presents the first systematic study of domain generalization for LiDAR semantic segmentation under noisy labels. We construct a controlled benchmark by corrupting source-domain labels with symmetric noise at predefined ratios, and evaluate both in-domain and cross-domain without accessing target samples or statistics. To ground the study, we adapt three representative noisy-label strategies from images

* Equal contribution

† Corresponding Author (email:kunyu.peng@kit.edu)

¹ The authors are with Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Karlsruhe 76131, Germany.

² The authors are with Department of Civil, Environmental and Geomatic Engineering, University College London, London WC1E 6BT, U.K.

The project is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB-1574 – 471687386. This work was supported in part by the SmartAge project sponsored by the Carl Zeiss Stiftung (P2019-01-003; 2021-2026). The authors gratefully acknowledge the computing time provided on the high-performance computer HoreKa by the National High-Performance Computing Center at KIT (NHR@KIT). This center is jointly supported by the Federal Ministry of Education and Research and the Ministry of Science, Research and the Arts of Baden-Württemberg, as part of the National High-Performance Computing (NHR) joint funding program (<https://www.nhr-verein.de/en/our-partners>). HoreKa is partly funded by the German Research Foundation (DFG).

to large-scale point clouds, namely TCL [10], DISC [11], and NPN [12], on a unified LiDAR backbone with matched optimization and augmentation. Building on empirical findings, we introduce **DuNe**, a **Dual-view** framework for learning with Noisy labels in 3D point clouds. It couples a strong geometry-aware view with a weak view, aligns them via bottleneck consistency, and employs confidence-filtered partial and negative supervision.

- **DGLSS-NL benchmark.** We establish the DGLSS-NL benchmark for single-source LiDAR semantic segmentation under controlled symmetric label noise, with rigorous in-domain and cross-domain evaluation to enable reproducible and fair comparison.
- **Standardized transfers and diagnostic insights.** We adapt three representative noisy-label strategies from images (TCL, DISC, NPN) to large-scale point clouds on a unified LiDAR backbone with matched training recipes, and conduct a controlled diagnostic analysis that disentangles the roles of sample selection, contrastive objectives, and negative learning under domain shift and varying noise severities, yielding actionable insights for future 3D methods.
- **Dual-view noise-robust generalization.** We introduce **DuNe**, a dual-view framework that fuses a geometry-aware strong view with a complementary weak view and integrates noise-aware supervision to resist label corruption and shift. The resulting model consistently surpasses the strongest transferred baseline across datasets and noise levels, improving Arithmetic Mean (AM) and Harmonic Mean (HM) of overall mIoU percentages by +4.6 / +4.6 at 10% noise, +6.2 / +6.9 at 20% noise, and +8.6 / +8.2 at 50% noise on average. These results are obtained by training on SemanticKITTI and evaluating on SemanticKITTI, nuScenes, and SemanticPOSS (denoted as $K \rightarrow K/N/P$).

II. RELATED WORK

A. LiDAR Semantic Segmentation

LiDAR semantic segmentation (LSS) assigns semantic categories to 3D points, with existing methods being point- [13], [15], projection- [16], or voxel- [17] based. Point-based methods directly process raw points and capture fine-grained local geometry but scale poorly to large outdoor scenes [13], [18]. Projection-based methods transform point clouds into 2D views, enabling efficient processing with standard CNNs and even off-the-shelf vision transformers, though inevitably at the cost of geometric distortion and information loss [16], [19], [20]. Voxel-based methods discretize the 3D space and apply sparse convolutions to balance accuracy and efficiency [17], [21]. Hybrid designs combine multiple representations for improved robustness [22]. Despite these advances, current models still rely heavily on costly manual annotations [6], [23], [24], which are often imperfect, imbalance, prone to annotation noise, and amplify sensitivity to dataset bias [8].

B. Domain Generalization in LiDAR Perception

While LSS models have achieved notable performance, they often degrade severely when applied to unseen environments, highlighting the need for 3D domain generalization (3DDG). Early studies such as MetaSets [25] addressed 3DDG for classification, where meta-learning with geometry-based transformations was introduced to bridge the gap between synthetic and real point clouds. Similarly, MAL [26] expanded the source domain with adversarially generated intermediate domains on transformation manifolds, improving generalization to unseen target sets. Beyond classification, Domain Generalization for LiDAR Semantic Segmentation (DGLSS) module [3] formalized DG in large-scale outdoor LiDAR. DGLSS addressed sparsity variations across sensors and scene distribution shifts challenges and proposed sparsity augmentation together with consistency regularizations. This established the first benchmark for LiDAR DG. However, existing DGLSS approaches assume clean supervision. In practice, large-scale LiDAR annotations are costly, imperfect, and inevitably noisy. This motivates our work on DGLSS under noisy labels (DGLSS-NL), where robustness to annotation noise is jointly addressed with domain generalization.

C. Noisy-label Learning

Although domain generalization alleviates dataset bias, its assumption of clean annotations rarely holds in real-world LiDAR datasets, where supervision is frequently corrupted by occlusion, sparsity, long-range effects, or human errors [6], [23], [24]. Noisy-label learning addresses this issue by enabling reliable training under imperfect annotations. In the image domain, research has focused on loss modification [10], [27], [28], regularization strategies [29], [30], and dynamic sample selection [11], [12], [31]. Despite progress, transferring existing techniques to LiDAR segmentation is challenging due to irregular points, costly annotations, and geometric constraints. LSS faces noisy labels, overlooked by domain generalization methods, while noise-robust approaches rarely address domain shifts. This motivates our noise-robust domain generalization framework for large-scale 3D point clouds.

III. BENCHMARK AND BASELINES

A. Noisy Labels

In our benchmark, we introduce noisy labels into existing domain generalization models for 3D point cloud semantic segmentation by adapting strategies from the image domain, aiming to evaluate model robustness under different noise levels. To simulate noise, we replace the ground-truth labels of certain samples with other categories in the dataset, following standard noisy-label protocols [27].

We distinguish between two types of noise: symmetric and asymmetric. Symmetric noise randomly flips each label to another class with equal probability, while asymmetric noise better reflects real-world annotation errors, where mislabeling is biased toward semantically similar categories (e.g., “truck” \rightarrow “bus”). Since asymmetric noise requires prior knowledge of class-level confusion statistics, we focus

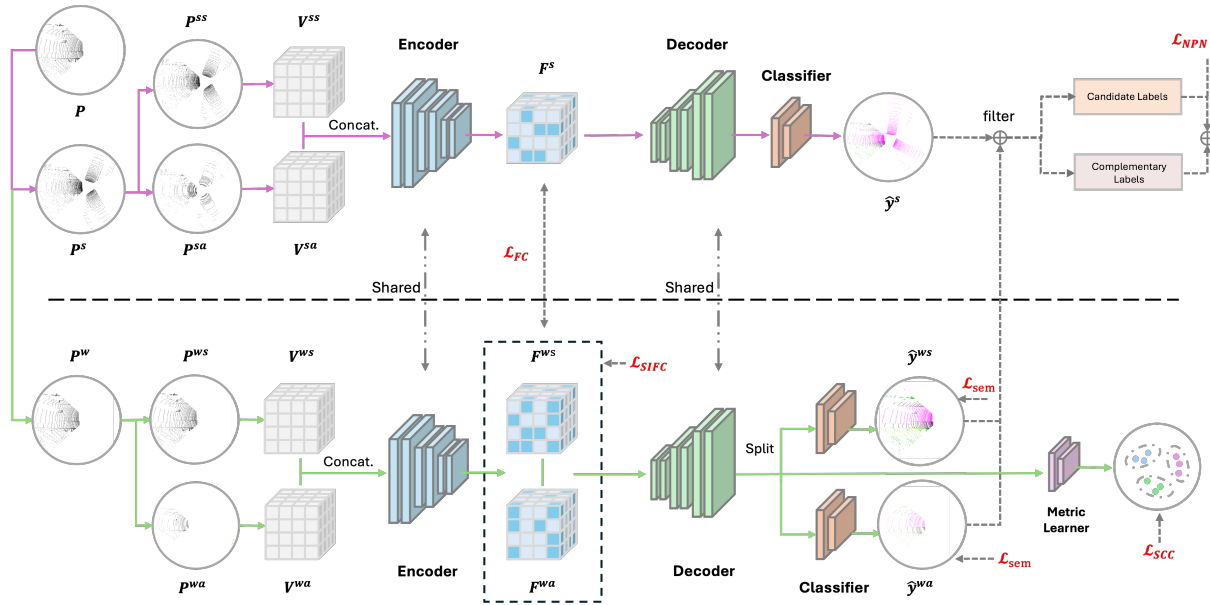


Fig. 2. **Overview of our proposed dual-view training pipeline - DuNe.** Each input LiDAR scan is first augmented by PolarMix [32] to generate a *strong* view and a *weak* view. Both views are further processed with sparsity augmentation. For the strong view, the entire augmented scan is used to construct candidate labels and complementary labels [12], which provide noise-robust supervision. For the weak view, we explicitly split the original and sparsity-augmented versions to form paired inputs, and enforce consistency loss and semantic correlation loss between them. This design allows the framework to jointly exploit label-level robustness and view-level consistency for improved noisy-label learning.

only on symmetric noise as a principled starting point. In practice, small noise (e.g., 2% or 5%) levels cause little disturbance due to the inherent robustness of the backbone, while very high noise ratios (e.g., above 50%) prevent the model from learning meaningful representations [28]. Therefore, we select three representative noise settings of 10%, 20%, and 50% for our experiments, while keeping the test set noise-free to ensure fair evaluation [33].

B. Baseline

We adopt a ResNet-based MinkowskiEngine [34] as the evaluation backbone, where the performance on unseen datasets reflects both segmentation quality and generalization ability. Building on this backbone, we are the first, to the best of our knowledge, to successfully overcome the challenges of transferring three representative noise-robust methods from the image domain to 3D LiDAR semantic segmentation and establish them as baselines: **TCL** [10] applies Gaussian modeling and pseudo-labeling to enable contrastive learning. In the 3D point cloud domain, it suffers from two issues: (1) clustering-based clean sample selection is prohibitively expensive for large-scale point cloud data and significantly slows down inference; and (2) stochastic augmentations and varying point counts hinder stable prototype construction for pseudo-labels. We mitigate these by restricting contrastive learning to a small subset of aligned points and a limited number of mixed-view pairs. **DISC** [11] dynamically selects clean samples while exploiting informative signals from noisy ones. However, it also faces challenges in large-scale point cloud segmentation: (1) the threshold-based comparison for distinguishing clean and hard labels is computationally expensive; and (2) points are not strictly aligned across batches, as newly appearing and disappearing points can

drastically affect the statistics. To mitigate these issues, we only maintain statistics for points that reappear across scans, which stabilizes the estimation while reducing computational cost. **NPN** [12] accumulates prediction statistics to construct partial and negative label sets, thereby mitigating the effect of overconfident noise. Its major challenge in 3D point cloud segmentation is that the varying number of points prevents the construction of consistent candidate label sets for individual samples, which makes it difficult to continuously update the candidates. To overcome this limitation, we instead take the maximum prediction within each batch as the candidate label, while still preserving complementary labels and the associated penalty function.

C. Dataset

SemanticKITTI [6] is a large-scale LiDAR semantic segmentation benchmark with fine-grained point-wise annotations across diverse driving scenarios. It contains over 40k scans with approximately 100k points per frame, making it a comprehensive source domain for studying segmentation performance. **NuScenes** [23] is a multi-sensor autonomous driving dataset covering 1,000 driving scenes in diverse urban conditions with varying weather, traffic density, and viewpoints. Its LiDAR scans introduce significant domain shifts compared to the source and are used as an unseen target domain for evaluating cross-dataset generalization. **SemanticPOSS** [24] is collected in Chinese urban environments and provides dense semantic annotations for street scenes under different geographic and environmental conditions. It serves as another target domain, complementing nuScenes by adding geographic and annotation diversity.

This setup enables a joint study of in-domain segmentation on SemanticKITTI and cross-domain generalization on

D. Evaluation Metric

We follow the evaluation protocol of the Single Domain Generalization for LiDAR Semantic Segmentation (DGLSS) [3] baseline. Specifically, mean Intersection-over-Union (mIoU) is used to measure semantic segmentation performance on individual datasets, while the Arithmetic Mean (AM) and Harmonic Mean (HM) are employed to summarize results across multiple target domains. mIoU reflects segmentation accuracy, AM indicates the overall level of cross-domain generalization, and HM provides a stricter criterion by emphasizing performance balance across domains. Under this protocol, we first evaluate the DGLSS baseline at different noise ratios and observe that its mIoU drops significantly as noise increases, indicating degraded segmentation accuracy and generalization capability. We then compare representative noise-robust methods transferred from the image domain with our proposed dual-view framework - **DuNe** under the same setting. Higher mIoU corresponds to better segmentation performance, while larger AM and HM indicate stronger cross-domain generalization and robustness.

IV. METHODOLOGY

A. Problem Definition

In this work, we study Domain Generalization for 3D LiDAR Semantic Segmentation under Noisy Labels, *i.e.*, DGLSS-NL. Given a 3D point cloud $P = \{p_i \in \mathbb{R}^3\}_{i=1}^N$, the goal is to learn a segmentation model $f(\cdot)$

$$y_i = f(p_i), \quad f: \mathbb{R}^3 \rightarrow \{1, \dots, C\} \quad (1)$$

that assigns each point p_i a true semantic label $y_i \in 1, \dots, C$, where C denotes the number of classes.

To enable DGLSS-NL, training is conducted on a noisy source domain \mathcal{D}_S and evaluation is performed on an unseen target domain \mathcal{D}_T :

$$\mathcal{D}_S = \{(p_i, \tilde{y}_i)\}_{i=1}^{N_S}, \quad \mathcal{D}_T = \{p_j\}_{j=1}^{N_T}, \quad \mathcal{D}_S \cap \mathcal{D}_T = \emptyset, \quad (2)$$

where N_S, N_T are the number of points in the source and target domains, respectively and \tilde{y}_i is defined as:

$$\tilde{y}_i = \begin{cases} y_i, & \text{with probability } (1 - \eta) \\ c \sim \mathcal{U}(C \setminus \{y_i\}), & \text{with probability } \eta \end{cases} \quad (3)$$

where $\eta \in \{0.1, 0.2, 0.5\}$ is the symmetric noise ratio and \mathcal{U} denotes uniform distribution over classes excluding y_i . The noise is synthetically injected to approximate real-world annotation errors (see Section III-A).

The training objective is to minimize the segmentation loss \mathcal{L}_{seg} on \mathcal{D}_S under noisy supervision,

$$\min_{\theta} \mathbb{E}_{(p, \tilde{y}) \sim \mathcal{D}_S} [\mathcal{L}_{\text{seg}}(f(p), \tilde{y})], \quad (4)$$

while ensuring that $f(\cdot)$ is robust to label noise and can generalize effectively to the unseen domain \mathcal{D}_T .

B. Dual-View Framework (DuNe)

As illustrated in Fig. 2, each point cloud $P = \{p_i\}_{i=1}^N$ is first augmented into two complementary views using the PolarMix strategy [32]:

$$P^w = \{p_i^w\}_{i=1}^N, \quad P^s = \{p_i^s\}_{i=1}^{N'}, \quad (5)$$

where P^w (weak view) preserves the structural fidelity of the original scan with the same cardinality N , and P^s (strong view) may have $N' \geq N$ points due to additional rotated-pasted instances introduced by PolarMix. In practice, PolarMix applies a scene-level swapping $Sw(\cdot)$ and an instance-level rotate-paste $Rp(\cdot)$ operation, yielding

$$P^s = Sw(P, P') \oplus Rp(P, P'), \quad (6)$$

where P' denotes another scan, and \oplus indicates concatenation.

Subsequently, both views are further processed by the baseline DGLSS module [3] to enhance sparsity. Specifically, the 3D point cloud is projected into a range view, and a randomly selected row is removed to simulate beam-missing artifacts:

$$\mathcal{A}(P) = \text{RowDrop}(\text{RangeProj}(P)), \quad (7)$$

which reduces the number of points to $\tilde{N} < N$ (or $\tilde{N} < N'$ for strong views). This yields four derived views:

$$P^{ss} = \mathcal{A}(P^s) \in \mathbb{R}^{\tilde{N}' \times 3}, \quad P^{sa} = P^s \in \mathbb{R}^{N' \times 3}, \quad (8)$$

$$P^{ws} = \mathcal{A}(P^w) \in \mathbb{R}^{\tilde{N} \times 3}, \quad P^{wa} = P^w \in \mathbb{R}^{N \times 3}. \quad (9)$$

Each view is voxelized and encoded with a ResNet-based sparse convolutional network implemented via MinkowskiEngine [34]:

$$F^v = \phi(P^v), \quad v \in \{ss, sa, ws, wa\}, \quad (10)$$

where $\phi(\cdot)$ denotes the sparse 3D encoder that maps voxelized point clouds into high-dimensional feature tensors.

A lightweight decoder with transposed sparse convolutions is then applied to upsample the encoded features back to the original voxel resolution for point-wise prediction. In the strong-view branch, the decoded features are concatenated and directly classified as

$$\hat{y}^s = \psi(\text{Concat}(F^{ss}, F^{sa})), \quad (11)$$

while in the weak-view branch, the concatenated features are decoded and split back into two parts:

$$(\hat{y}^{ws}, \hat{y}^{wa}) = \psi(\text{Concat}(F^{ws}, F^{wa})), \quad (12)$$

where $\psi(\cdot)$ denotes the task-specific decoder and classifier. During inference, only the strong branch is utilised, while the weak branch and consistency losses are disabled for computational efficiency.

C. Loss Function

To improve robustness against sparsity variation, semantic ambiguity, and noisy supervision, we adopt two complementary objectives: DGLSS module and NPN module.

DGLSS Loss: Following the DGLSS framework [3], we introduce two consistency terms in addition to the standard cross-entropy (CE) loss. First, the Sparsity-Invariant Feature Consistency (SIFC) encourages feature alignment across scans with different sparsity levels:

$$\mathcal{L}_{\text{SIFC}} = \frac{1}{N} \sum_{i=1}^N \|F_i^{ws} - F_i^{wa}\|_1. \quad (13)$$

Second, the Semantic Correlation Consistency (SCC) enforces stable inter-class relationships across domains. For each scan, decoded features are passed through a metric learner to obtain class-wise prototypes $Z_i \in \mathbb{R}^{C \times d}$, where each row is the average embedding of a class. The SCC loss aligns pairwise prototype correlations across scans:

$$\mathcal{L}_{\text{SCC}} = \frac{1}{L} \sum_i \sum_{j \neq i} (Z_i^\top Z_j - Z_j^\top Z_i), \quad (14)$$

Finally, besides these two consistency terms, we also adopt a weighted CE loss \mathcal{L}_{sem} to handle class imbalance. The overall DGLSS objective is

$$\mathcal{L}_{\text{DGLSS}} = \mathcal{L}_{\text{sem}} + \alpha \mathcal{L}_{\text{SIFC}} + \beta \mathcal{L}_{\text{SCC}}, \quad (15)$$

where α and β balance the contributions.

NPN Loss: To further handle noisy labels, we adopt NPN [12]. Given the predicted label \hat{y}^s from the strong branch, we decompose the label space into a candidate label set \hat{Y} (including \hat{y}^{ws} , \hat{y}^{wa} and \hat{y}^s) and its complementary label set (all remaining classes). The Partial Label Learning (PLL) term encourages the model to predict one of the candidate labels, while the Negative Learning (NL) term explicitly penalizes complementary labels:

$$L_{\text{NL}} = -\frac{1}{N} \sum_{i=1}^N \sum_{c \neq \hat{y}_i} \log(1 - p_\theta(\hat{y} = c | x_i)), \quad \hat{y} \in \hat{Y}. \quad (16)$$

The final NPN objective integrates PLL, NL, and a confidence regularization penalty:

$$\mathcal{L}_{\text{NPN}} = \mu L_{\text{NL}} + \nu L_{\text{CE}} + L_{\text{pen}}, \quad (17)$$

where μ and ν are balancing coefficients.

Overall Objective: Finally, we integrate the DGLSS and NPN losses within the dual-view framework. In addition to the consistency and semantic correlation regularizations, we introduce a dual-view feature consistency loss \mathcal{L}_{FC} to encourage similarity between the strong and weak view representations. The overall training objective is defined as

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{DGLSS}} + \mathcal{L}_{\text{NPN}} + \lambda \mathcal{L}_{\text{FC}}, \quad (18)$$

where λ balances the dual-view feature consistency term. This unified design jointly enforces feature-level consistency, semantic correlation, and noise-robust supervision, enabling the network to learn reliable representations from noisy labels while maintaining strong generalization across domains.

TABLE I

COMPARISON OF 3D SEMANTIC SEGMENTATION RESULTS TRAINED ON KITTI UNDER DIFFERENT NOISE RATIOS. ALL VALUES ARE mIoU(%) (↑). K, N, AND P DENOTE SEMANTICKITTI, nuSCENES, AND SEMANTICPOSS. K→N INDICATE MODELS TRAINED ON K AND EVALUATED ON N.

Dataset	Symmetric Noise					
	0%	2%	5%	10%	20%	50%
K → K	58.06	55.70	54.75	32.99	28.35	10.86
K → N	42.28	34.96	36.42	21.76	16.48	7.53
K → P	49.09	41.81	44.28	22.29	18.86	9.38
AM	49.81	45.33	44.28	27.38	22.42	9.19
HM	48.99	42.96	43.74	26.23	20.85	9.25

V. EXPERIMENTS

A. Experimental Setup

All models are implemented in PyTorch with sparse convolution operations from MinkowskiEngine, and trained on NVIDIA A100 GPUs with 40GB memory.

Backbone. We adopt the DGLSS [3] model as the backbone baseline and evaluate its performance under four supervision settings: (i) clean labels (noise-free baseline), (ii) noisy labels without correction, (iii) three noise-robust strategies migrated from the image domain (TCL [10], DISC [11], and NPN [12]), and (iv) our proposed dual-view learning framework - **DuNe**, which incorporates bottleneck consistency and adopts partial/negative label supervision. For in-domain evaluation, models are trained and tested on the SemanticKITTI [6] dataset (K → K). For cross-domain evaluation, models are trained on SemanticKITTI and directly tested on nuScenes [23] (K → N) and SemanticPOSS [24] (K → P) without target-domain fine-tuning.

Hyperparameters. Stochastic gradient descent (SGD) with momentum 0.9 and weight decay $1e^{-4}$ is used as the optimizer. The initial learning rate is set to 0.01 and decayed by a cosine annealing schedule. Unless otherwise specified, the loss balancing weights are set to $\mu = 1.0$ for the consistency loss, $\nu = 1.0$ for the negative cross-entropy loss, and $\lambda = 2.0$ for the prototype contrastive loss.

B. Results

Effect of Label Noise on the Baseline. Table I shows that the baseline DGLSS is highly vulnerable to noisy supervision. With clean labels, it achieves 58.06% mIoU on SemanticKITTI, 42.28% on nuScenes, and 49.09% on SemanticPOSS. However, performance deteriorates rapidly with increasing noise: on SemanticKITTI, mIoU drops to 32.99% at 10% noise and to only 10.86% at 50%. Similar trends appear in cross-dataset evaluations, where nuScenes decreases from 42.28% to 7.53% and SemanticPOSS from 49.09% to 9.38%. While low noise levels (2%–5%) have only a marginal effect due to the inherent stability of the backbone, moderate to high noise ratios (20%–50%) substantially impair both in-domain accuracy and cross-domain robustness, confirming the necessity of noise-robust strategies in LiDAR semantic segmentation.

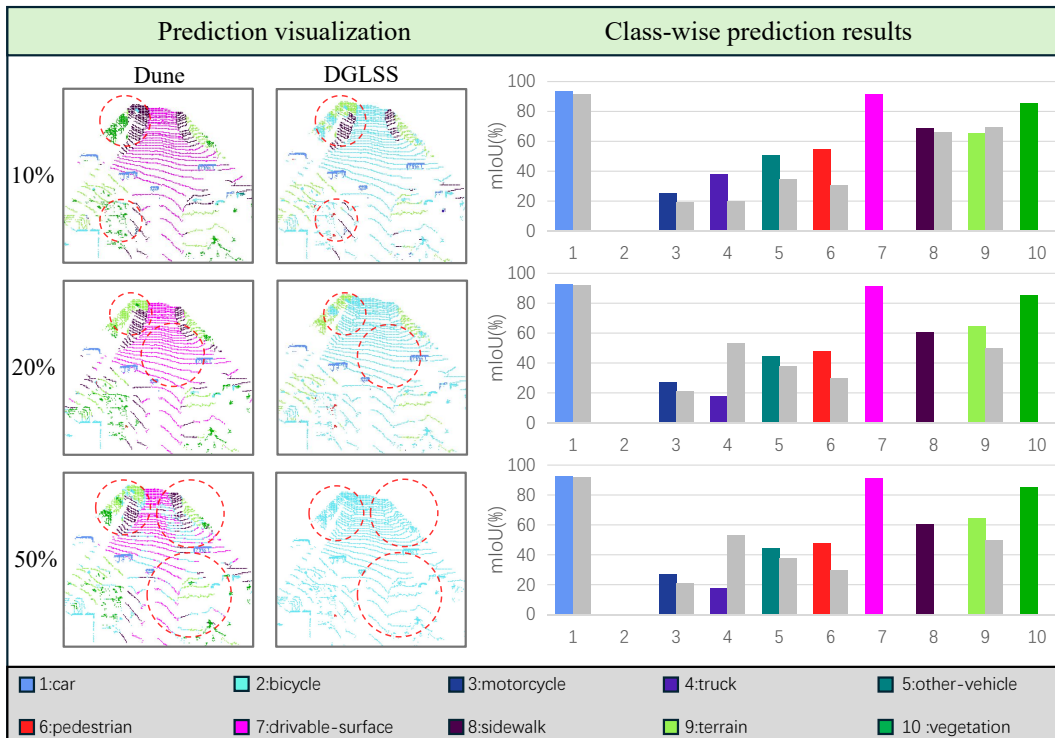


Fig. 3. Qualitative results under 10%, 20%, and 50% label noise. For each noise ratio, we compare segmentation predictions of our framework DuNe with the baseline DGLSS. Corresponding class-wise prediction histograms (10 categories) are also shown. The left column visualizes predictions, where circles highlight differences between methods. The right column presents class-wise prediction distributions: colored bars indicate results from our method, while gray bars correspond to predictions from the DGLSS baseline. Our method produces more coherent segmentations and maintains balanced class distributions, even under severe noise.

TABLE II

COMPARISON OF METHODS UNDER DIFFERENT NOISE RATIOS. ALL VALUES ARE mIoU (%) (↑). K, N, AND P DENOTE SEMANTICKITTI, NUSCENES, AND SEMANTICPOSS, RESPECTIVELY. ALL RESULTS ARE OBTAINED BY TRAINING ON SEMANTICKITTI AND EVALUATING ON THE DATASETS INDICATED BELOW.

Noise Ratio	Method	K	N	P	AM	HM
10%	TCL [10]	19.59	18.98	22.11	19.28	19.28
	DISC [11]	41.42	36.79	31.03	36.23	35.48
	NPN [12]	52.05	38.00	42.88	45.02	43.93
	DuNe	56.86	42.28	52.58	49.57	48.50
20%	TCL	14.14	10.33	8.86	10.82	10.80
	DISC	36.68	34.44	25.85	31.26	30.33
	NPN	48.72	31.84	36.63	40.28	38.53
	DuNe	53.20	39.69	48.05	46.44	45.46
50%	TCL	10.37	11.28	10.19	10.82	10.82
	DISC	33.64	28.22	24.98	29.31	28.67
	NPN	41.78	30.59	31.07	36.19	35.32
	DuNe	52.37	37.18	43.07	44.78	43.49

Comparison with Transferred Noise-Robust Methods.

Table II compares DuNe with three representative noise-robust strategies adapted from the image domain: TCL [10], DISC [11], and NPN [12]. TCL consistently fails to improve robustness and in many cases even underperforms the noisy baseline, showing that this method does not transfer effectively to 3D LiDAR segmentation. DISC achieves moderate gains by dynamically selecting cleaner samples, e.g., on

SemanticKITTI it improves performance from 32.99% to 41.42% at 10% noise, with similar trends on nuScenes and SemanticPOSS. NPN further enhances robustness by leveraging partial and negative labels, reaching 41.78% on SemanticKITTI, 30.59% on nuScenes, and 31.07% on SemanticPOSS at 50% noise.

Performance of the Proposed Framework. Our method DuNe consistently achieves the best results across all datasets and noise levels according to Table III. It maintains 52.37% on SemanticKITTI, 37.18% on nuScenes, and 43.07% on SemanticPOSS under 50% noise, substantially outperforming TCL and DISC and clearly surpassing NPN. These results demonstrate that while generic 2D noise-robust methods provide partial benefits, principled designs tailored to 3D LiDAR are essential for mitigating the impact of noisy supervision. More importantly, our approach enhances generalization to unseen domains. Under 10% symmetric noise, when trained on SemanticKITTI, our method achieves 56.86% mIoU on SemanticKITTI, 42.28% mIoU on nuScenes and 52.58% on SemanticPOSS, resulting in AM and HM of 49.57% and 48.50%, respectively. These results consistently surpass the baseline and approach the performance of DGLSS trained with clean labels, demonstrating strong robustness against both label noise and domain shift. Overall, the proposed dual-view learning framework not only restores segmentation accuracy under noisy supervision but also substantially improves cross-domain generalization,

TABLE III

PER-CLASS IOU(%)(\uparrow) UNDER DIFFERENT NOISE RATIOS. HERE, **K**, **N**, AND **P** DENOTE SEMANTICKITTI, nuSCENES, AND SEMANTICPOSS, RESPECTIVELY. **K**→**N** INDICATE MODELS TRAINED ON **K** AND EVALUATED ON **N**. WE OMIT THE NON-EXISTENT COMMON CLASSES IN SEMANTICPOSS AND MARK THEM WITH '-'.

Noise Ratio	Method	Dataset	mIoU	car	bicycle	motorcycle	truck	other-vehicle	pedestrian	drivable-surface	sidewalk	terrain	vegetation	AM	HM
10%	DGLSS	K → K	32.99	91.44	0.08	19.38	19.63	34.53	30.13	0.01	65.72	69.00	0.00	27.38	26.23
		K → N	21.76	76.06	0.02	17.51	11.94	17.65	24.46	0.01	35.48	34.50	0.00		
		K → P	22.29	59.66	4.60	-	-	-	45.68	0.01	0.00	1.49	-		
	DuNe	K → K	56.86	92.92	0.58	13.24	54.59	51.14	50.78	91.70	62.19	68.25	83.19	49.57	48.50
		K → N	42.28	78.06	0.11	33.15	26.76	12.33	44.34	84.98	34.48	41.56	67.01		
		K → P	52.58	64.53	7.52	-	-	-	50.61	66.21	-	74.03	-		
20%	DGLSS	K → K	28.35	92.24	0.06	20.82	53.12	37.57	29.85	0.00	0.00	49.84	0.00	22.42	20.85
		K → N	16.48	73.26	0.02	15.11	18.07	7.79	29.91	0.00	0.00	20.67	0.00		
		K → P	18.66	51.84	4.58	-	-	-	35.38	0.00	0.00	1.52	-		
	DuNe	K → K	53.20	92.76	0.28	27.19	17.46	44.79	47.92	91.32	60.83	64.45	84.99	46.44	45.46
		K → N	39.69	79.61	0.08	36.30	15.20	15.83	38.08	83.70	36.58	35.48	56.05		
		K → P	48.05	61.86	6.64	-	-	-	43.25	60.43	-	68.05	-		
50%	DGLSS	K → K	10.86	0.38	0.06	8.24	21.47	45.74	32.68	0.00	0.00	0.00	0.00	9.19	8.89
		K → N	7.53	0.13	0.02	13.82	23.84	8.02	29.44	0.00	0.00	0.00	0.00		
		K → P	9.38	0.05	4.38	-	-	-	42.47	0.00	0.00	0.00	-		
	DuNe	K → K	52.37	90.05	0.22	19.33	41.72	36.48	42.78	90.84	61.24	59.62	81.46	44.78	43.49
		K → N	37.18	72.60	0.06	31.09	17.45	12.24	35.07	78.33	38.55	25.22	61.21		
		K → P	43.07	54.58	6.65	-	-	-	40.46	40.17	-	73.48	-		

TABLE IV

ABLATION STUDY ON THE CONTRIBUTIONS OF NPN AND POLARMIX UNDER THE DGLSS SETTING. ALL VALUES ARE REPORTED AS MIOU(%)(\uparrow). HERE, **K**, **N**, AND **P** DENOTE SEMANTICKITTI, nuSCENES, AND SEMANTICPOSS, RESPECTIVELY. **K**→**N** INDICATE MODELS TRAINED ON SEMANTICKITTI AND EVALUATED ON nuSCENES.

Method	K → K	K → N	K → P	AM	HM
10% DGLSS	32.99	21.76	22.29	27.38	26.23
10% DGLSS + PolarMix	44.12	33.48	43.94	38.80	38.07
10% DGLSS + NPN	52.05	38.00	42.88	45.02	43.93
10% DGLSS + NPN + PolarMix	55.75	35.92	38.67	45.83	43.69
50% Noise strong branch	30.71	19.84	24.90	25.27	24.11
20% Noise weak branch	51.03	33.74	42.86	42.38	40.62
10% Noise weak branch	54.13	36.96	48.53	45.55	43.93
10% DuNe	56.86	42.28	52.58	49.57	48.50

establishing a solid benchmark for noise-robust LiDAR semantic segmentation.

C. Ablation Study

Table IV reports the ablation results on the contributions of NPN and PolarMix under the DGLSS setting. Several important observations can be made.

Effect of PolarMix. Compared to the plain DGLSS baseline (32.99%, 21.76%, and 22.29% on SemanticKITTI, nuScenes, and SemanticPOSS), adding PolarMix under 10% label noise already yields clear gains, reaching 44.12%, 33.48%, and 43.94%. This indicates that geometric mixing can effectively enhance data diversity and support cross-domain transfer, although the in-domain improvement remains moderate.

Effect of NPN. Applying NPN instead of PolarMix results in more substantial robustness improvements, with performance rising to 52.05%, 38.00%, and 42.88% on the three benchmarks. This demonstrates that partial-label and negative learning significantly mitigate the negative effect of corrupted annotations and stabilize the optimization.

Combination of NPN and PolarMix. When both are applied simultaneously, the model achieves 55.75%, 35.92%, and 38.67%. Although the SemanticPOSS result is slightly lower than NPN-only, the SemanticKITTI accuracy surpasses both single variants. These results suggest that noise modeling and geometric augmentation are complementary, each contributing to robustness and generalization under noisy supervision.

Effect of dual-branch consistency. Although combining NPN and PolarMix improves performance, the overall generalization remains limited, with AM and HM reaching only 45.83% and 43.69%. To further bridge the semantic gap between different representations, we introduce a consistency loss on the bottleneck features of the dual-branch design. Our experiments show that under 10% and 20% noise, using the strong achieves noticeably performance of generalization with higher AM and HM. This suggests that the denser augmented point clouds provide richer cues and promote the learning of transferable knowledge. However, under 50% noise, the dense but corrupted strong set amplifies the adverse effects of label noise, causing performance to collapse. Therefore, we adopt a selective strategy: using the strong set for 10% and 20% noise, and the weak set for 50% noise, enabling more robust performance across different noise levels.

Full framework. Our method **DuNe** achieves the best results across all datasets, with 56.86%, 42.28%, and 52.58% on SemanticKITTI, nuScenes, and SemanticPOSS under 10% noise, respectively. These improvements confirm that integrating NPN and PolarMix within our dual-view training strategy is crucial for achieving robustness and generalization under noisy supervision.

VI. CONCLUSION

We studied robust domain-generalized LiDAR semantic segmentation under noisy supervision. To establish a benchmark, we introduced symmetric label noise into SemanticKITTI, nuScenes, and SemanticPOSS, and adapted three representative noise-robust strategies from the image domain. This revealed both the sensitivity of standard training to corrupted labels and the limited transferability of 2D methods to 3D point clouds. To address this, we proposed **DuNe**, a dual-view learning framework combining PolarMix augmentation, bottleneck consistency, and partial/negative label supervision. Experiments showed that **DuNe** not only recovers segmentation accuracy under noise but also improves cross-dataset generalization. We hope this benchmark and framework foster future research on noise-robust LiDAR perception for autonomous driving.

REFERENCES

- [1] F. Pomerleau, F. Colas, and R. Siegwart, "A review of point cloud registration algorithms for mobile robotics," *Foundations and Trends® in Robotics*, vol. 4, no. 1, pp. 1–104, 2015.
- [2] K. Peng, J. Fei, K. Yang, A. Roitberg, J. Zhang, F. Bieder, P. Heidenreich, C. Stiller, and R. Stiefelwagen, "Mass: Multi-attentional semantic segmentation of lidar data for dense top-view understanding," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 9, pp. 15 824–15 840, 2022.
- [3] H. Kim, Y. Kang, C. Oh, and K.-J. Yoon, "Single domain generalization for lidar semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 17 587–17 598.
- [4] Y. Fu, Y. Xie, Y. Fu, and Y.-G. Jiang, "Styleadv: Meta style adversarial training for cross-domain few-shot learning," in *CVPR*, 2023.
- [5] Y. Fu, Y. Wang, Y. Pan, L. Huai, X. Qiu, Z. Shangguan, T. Liu, Y. Fu, L. Van Gool, and X. Jiang, "Cross-domain few-shot object detection via enhanced open-set object detector," in *ECCV*, 2024.
- [6] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, "Semantickitti: A dataset for semantic scene understanding of lidar sequences," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019.
- [7] P. Sun, H. Kretschmar, X. Dotiwalla, A. Chouard *et al.*, "Scalability in perception for autonomous driving: Waymo open dataset," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 2443–2451.
- [8] D. Griffiths and J. Boehm, "Weighted point cloud augmentation for neural network training data class-imbalance," *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 42, pp. 981–987, 2019.
- [9] B. Frenay and M. Verleysen, "Classification in the presence of label noise: A survey," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 25, no. 5, pp. 845–869, 2014.
- [10] Z. Huang, J. Zhang, and H. Shan, "Twin contrastive learning with noisy labels," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2023.
- [11] Y. Li, H. Han, S. Shan, and X. Chen, "Disc: Learning from noisy labels via dynamic instance-specific selection and correction," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, June 2023, pp. 24 070–24 079.
- [12] M. Sheng, Z. Sun, Z. Cai, T. Chen, Y. Zhou, and Y. Yao, "Adaptive integration of partial label learning and negative learning for enhanced noisy label learning," in *Proc. AAAI Conf. Artif. Intell.*, vol. 38, no. 5, Mar. 2024, pp. 4820–4828.
- [13] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 652–660.
- [14] C. Wu, K. Wang, Z. Zhong, H. Fu, J. Zheng, J. Zhang, J. Pfommer, and J. Beyerer, "Rethinking attention module design for point cloud analysis," in *Proc. Int. Conf. Pattern Recognit.* Springer, 2024, pp. 249–267.
- [15] J. Fei, K. Peng, P. Heidenreich, F. Bieder, and C. Stiller, "Pillarsegnet: Pillar-based semantic grid map estimation using sparse lidar data," in *2021 IEEE intelligent vehicles symposium (IV)*. IEEE, 2021, pp. 838–844.
- [16] B. Wu, A. Wan, X. Yue, and K. Keutzer, "Squeezeseg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3d lidar point cloud," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2018.
- [17] Y. Zhou and O. Tuzel, "Voxelnet: End-to-end learning for point cloud based 3d object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018.
- [18] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," in *NeurIPS*, 2017.
- [19] A. Milioto, I. Vizzo, J. Behley, and C. Stachniss, "Rangenet++: Fast and accurate lidar semantic segmentation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2019.
- [20] J. Moh Goo, Z. Zeng, and J. Boehm, "Zero-shot detection of buildings in mobile lidar using language vision model," *ISPRS-International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 48, pp. 107–113, 2024.
- [21] B. Graham, M. Engelcke, and L. van der Maaten, "3d semantic segmentation with submanifold sparse convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018.
- [22] Q. Hu, B. Yang, L. Xie, S. Rosa, Y. Guo, Z. Wang, N. Trigoni, and A. Markham, "Randla-net: Efficient semantic segmentation of large-scale point clouds," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020.
- [23] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020.
- [24] Y. Pan, B. Gao, J. Mei, S. Geng, C. Li, and H. Zhao, "Semanticpos: A point cloud dataset with large quantity of dynamic instances," in *2020 IEEE Intelligent Vehicles Symposium (IV)*, 2020, pp. 687–693.
- [25] C. Huang, Z. Cao, Y. Wang, J. Wang, and M. Long, "Metasets: Meta-learning on point sets for generalizable representations," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 8863–8872.
- [26] H. Huang, C. Chen, and Y. Fang, "Manifold adversarial learning for cross-domain 3d shape representation," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2022, pp. 272–289.
- [27] G. Patrini, A. Rozza, A. Krishna Menon, R. Nock, and L. Qu, "Making deep neural networks robust to label noise: A loss correction approach," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017.
- [28] L. Jiang, Z. Zhou, T. Leung, L.-J. Li, and L. Fei-Fei, "MentorNet: Learning data-driven curriculum for very deep neural networks on corrupted labels," in *Int. Conf. Mach. Learn.*, 2018, pp. 2304–2313.
- [29] A. Ghosh, H. Kumar, and P. S. Sastry, "Robust loss functions under label noise for deep neural networks," in *Proc. AAAI Conf. Artif. Intell.*, vol. 31, no. 1, 2017.
- [30] X. Ma, H. Huang, Y. Wang, S. Romano, S. Erfani, and J. Bailey, "Normalized loss functions for deep learning with noisy labels," in *Int. Conf. Mach. Learn.*, 2020, pp. 6543–6553.
- [31] B. Han, Q. Yao, X. Yu, G. Niu, M. Xu, W. Hu, I. Tsang, and M. Sugiyama, "Co-teaching: Robust training of deep neural networks with extremely noisy labels," in *Adv. Neural Inf. Process. Syst.*, 2018, pp. 8527–8537.
- [32] A. Xiao, J. Huang, D. Guan, K. Cui, S. Lu, and L. Shao, "Polarmix: A general data augmentation technique for lidar point clouds," *Advances in Neural Information Processing Systems*, vol. 35, pp. 11 035–11 048, 2022.
- [33] M. Li and C. Zhu, "Noisy label processing for classification: A survey," *arXiv preprint arXiv:2404.04159*, 2024.
- [34] C. Choy, J. Gwak, and S. Savarese, "4d spatio-temporal convnets: Minkowski convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 3075–3084.