ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

Adaptive-D3CTTA: Domain-Aware Prototype Learning and Temporal-Geometric Consistency for Continual Test-Time Adaptation of 3D LiDAR Semantic Segmentation

Palthiya Rahavendra[1], Dr. Y. Venkata Raghavarao[2,], Pilli Veeraswami[3], Tadi Prasad[4]

[1,2,3,4]Department of Computer Science and Engineering Sree Vahini Institute of Science and Technology Tiruvuru, Andhra Pradesh, India

DOI: https://doi.org/10.51583/IJLTEMAS.2026.151000097

Received: 14 October 2026; Accepted: 20 October 2026; Published: 10 November 2026

ABSTRACT

3D LiDAR semantic segmentation is a fundamental task in autonomous driving, but its performance can deteriorate when deployment conditions differ from those observed during training. Changes in point-cloud density, sensor characteristics, scene geometry, and object distributions create domain shifts that challenge models operating in continuously changing environments. This paper proposes Adaptive-D3CTTA, a domain-aware continual test-time adaptation framework for 3D LiDAR semantic segmentation. The framework combines range- and density-aware prototype learning, domain-conditioned feature decorrelation, temporal memory, K-nearest-neighbor (KNN) geometric consistency, and confidence-aware pseudo-label refinement. Range and density information are used to characterize incoming scans and support class-level feature representation, while temporal memory and local geometric agreement help assess the reliability of predictions across sequential observations. A confidence-based refinement mechanism selects reliable pseudo-labels for online adaptation. The framework is designed to process unlabeled target-domain streams without requiring target annotations during adaptation. Its performance should be evaluated using a clearly defined source-to-target protocol on SemanticKITTI and nuScenes-LiDARSeg, with comparisons against the source-only model and relevant test-time adaptation methods, including D³CTTA. Mean intersection-over-union, per-class intersection-over-union, mean class accuracy, adaptation stability, runtime, and memory consumption provide complementary measures of segmentation quality and deployment cost. The proposed framework aims to improve the reliability of continual adaptation by combining semantic, temporal, and geometric information within a unified pipeline. Its effectiveness must be established through reproducible experiments and component-wise ablation studies.

Keywords: 3D LiDAR Semantic Segmentation, Continual Test-Time Adaptation, Domain Shift, Prototype Learning, Temporal Consistency, Geometric Consistency, Pseudo-Label Refinement, Autonomous Driving.

INTRODUCTION

Three-dimensional (3D) LiDAR semantic segmentation is a fundamental component of autonomous driving systems, enabling vehicles to interpret complex driving environments through point-wise semantic classification. Accurate segmentation supports essential perception tasks, including road-boundary identification, obstacle detection, and the recognition of vehicles, pedestrians, and other traffic participants. Recent deep learning approaches, including PointNet, PointNet++, KPConv, RandLA-Net, PolarNet, and Cylinder3D, have demonstrated the effectiveness of learning semantic representations from 3D point clouds [8]-[11], [15], [16]. Despite these advances, deploying pretrained segmentation models in previously unseen environments remains challenging. Differences in sensor characteristics, point-cloud density, observation distance, scene geometry, and object distributions can introduce substantial domain shifts, potentially degrading segmentation accuracy and prediction reliability [1], [2], [12]-[14]. Consequently, developing adaptation strategies that remain reliable under changing environmental conditions is an important research direction for real-world LiDAR perception.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

Test-time adaptation (TTA) addresses domain shift by adapting a pretrained model to unlabeled target-domain observations during deployment. Tent investigates entropy minimization as a mechanism for test-time adaptation [4], while continual test-time domain adaptation methods extend adaptation to sequential data streams in which target distributions may change over time [5]. EATA further investigates efficient adaptation while reducing forgetting [6]. Although these approaches provide important foundations, continual adaptation presents additional challenges in safety-critical perception settings. In particular, erroneous pseudo-labels may reinforce incorrect predictions, changing feature distributions may destabilize previously learned representations, and repeated adaptation may accumulate errors across successive observations. These challenges motivate adaptation strategies that consider not only prediction confidence but also the semantic, spatial, and temporal consistency of incoming LiDAR data.

Recent research has specifically addressed continual test-time adaptation for 3D LiDAR semantic segmentation. D³CTTA investigates distance-aware prototype learning and domain-dependent feature decorrelation to address variations in LiDAR observations and improve feature representations under domain shift [3]. GIPSO explores geometrically informed propagation for online adaptation in 3D LiDAR segmentation [7]. These studies establish relevant foundations for domain-aware and geometry-informed adaptation. Nevertheless, a key research question remains: how can complementary semantic, temporal, and local geometric evidence be combined to identify reliable predictions while processing a continuously changing stream of unlabeled LiDAR scans? In particular, the potential benefits of integrating class-level prototype information, cross-scan temporal consistency, neighborhood-level geometric agreement, and confidence-based pseudo-label selection require systematic investigation within a unified adaptation framework.

To investigate this question, this paper proposes **Adaptive-D3CTTA**, a domain-aware continual test-time adaptation framework for 3D LiDAR semantic segmentation. The proposed framework integrates five complementary components: range- and density-aware prototype learning, domain-conditioned feature decorrelation, temporal memory, K-nearest-neighbor (KNN)-based geometric consistency, and confidence-aware pseudo-label refinement. First, range and point-density information characterize the geometric properties of incoming LiDAR observations and support class-level feature representation. Second, prototype learning maintains representative semantic features using reliable predictions, while feature decorrelation is intended to reduce undesirable dependencies in the learned representations. Third, temporal memory provides information from previous observations to assess consistency across sequential scans. Fourth, KNN-based geometric consistency evaluates agreement among spatially neighboring points. Finally, a confidence-aware refinement mechanism combines the available reliability signals to select candidate pseudo-labels for the implemented adaptation procedure. Together, these components are designed to investigate whether semantic, temporal, and geometric evidence can improve the reliability of continual adaptation under changing target-domain conditions.

The principal contributions of this work are summarized as follows:

- **Range- and density-aware prototype learning:** A domain-aware prototype-learning component that incorporates LiDAR range and point-density information to investigate the maintenance of representative semantic features under changing observation conditions.

- **Joint temporal and geometric consistency:** An integrated framework component that combines temporal memory with KNN-based local geometric agreement to assess prediction reliability across sequential scans and neighboring points.

- **Confidence-aware pseudo-label refinement:** A reliability-based selection strategy that combines prediction confidence, prototype similarity, temporal consistency, and geometric agreement to identify candidate pseudo-labels for adaptation.

- **Systematic experimental evaluation:** An evaluation framework using SemanticKITTI and nuScenes-LiDARSeg, with comparisons against appropriate source-only and test-time adaptation baselines, together with component-wise ablation studies to investigate the individual and combined

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

contributions of the proposed components.

Related Work

3D LiDAR Semantic Segmentation

3D LiDAR semantic segmentation assigns semantic labels to individual points in a point cloud and plays an important role in autonomous driving and environmental perception. The SemanticKITTI dataset provides sequential LiDAR scans with dense point-wise annotations, supporting research on semantic scene understanding and temporal point-cloud analysis [1]. The nuScenes dataset provides multimodal driving data for evaluating perception systems under diverse urban driving conditions [2]. Deep learning approaches, including RandLA-Net, PolarNet, Cylinder3D, and KPConv, investigate different strategies for extracting geometric features and efficiently processing large-scale point clouds [8]-[11]. However, variations in sensor characteristics, point density, observation distance, and environmental conditions can introduce domain shifts that reduce segmentation reliability in previously unseen environments [1], [2], [13], [14].

Continual Test-Time Adaptation

Test-time adaptation enables a pretrained model to adapt to unlabeled target-domain observations without requiring labeled target data during adaptation. Tent investigates entropy minimization for adapting model predictions at test time [4]. Wang et al. proposed Continual Test-Time Domain Adaptation (CoTTA), which employs averaged predictions and stochastic restoration of model parameters to mitigate error accumulation and catastrophic forgetting under changing target distributions [5]. EATA further investigates efficient test-time adaptation while reducing forgetting [6]. Despite these developments, continual adaptation remains challenging when target-domain distributions change over time and unreliable pseudo-labels introduce errors into subsequent predictions [5], [6]. These challenges motivate the investigation of additional reliability signals for sequential LiDAR perception.

Domain-Aware Adaptation for LiDAR Segmentation

Zhao et al. proposed D³CTTA, a continual test-time adaptation framework specifically designed for 3D LiDAR semantic segmentation [3]. D³CTTA incorporates distance-aware prototype learning to account for variations in point-cloud density across observation distances and domain-dependent feature decorrelation to reduce feature correlations across domains and semantic categories. The method is designed for online adaptation without backpropagation and provides a relevant baseline for evaluating domain-aware adaptation in LiDAR point clouds [3]. In related work, GIPSO investigates geometrically informed propagation for online adaptation in 3D LiDAR segmentation [7]. These approaches demonstrate the relevance of geometric information and domain-aware feature representations for handling distribution shifts. However, any additional benefit from combining temporal memory, local geometric agreement, and confidence-aware pseudo-label selection must be established through controlled experimental comparisons.

Research Gap and Motivation

Existing continual test-time adaptation methods address changing target distributions, prediction instability, and error accumulation [5], [6]. D³CTTA specifically considers distance-aware prototype learning and domain-dependent feature decorrelation for continual adaptation of LiDAR segmentation models [3]. Building on these research directions, **Adaptive-D3CTTA investigates the integration of domain-aware prototypes and feature decorrelation with temporal memory, K-nearest-neighbor (KNN)-based geometric consistency, and confidence-aware pseudo-label refinement**. Temporal memory is intended to provide information from preceding observations, while KNN-based consistency evaluates agreement among spatially neighboring points. Confidence-aware refinement combines available reliability signals to select candidate pseudo-labels for the implemented adaptation procedure. The central research question is whether these complementary signals can improve segmentation reliability under changing target-domain conditions. This question should be examined through comparisons with D³CTTA and other relevant

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

baselines, together with component-wise ablation studies under a consistent evaluation protocol.

Proposed Methodology

Overview of Adaptive-D3CTTA

Adaptive-D3CTTA is designed for continual test-time adaptation of 3D LiDAR semantic segmentation under changing target-domain conditions. Given a sequence of unlabeled LiDAR scans, the framework combines domain-aware prototype learning, feature decorrelation, temporal memory, K-nearest-neighbor (KNN) geometric consistency, and confidence-aware pseudo-label refinement. The overall objective is to improve prediction reliability while adapting to changes in point-cloud distributions [4].

For each incoming scan, the framework extracts point-wise features, estimates range and density information, generates semantic predictions, evaluates temporal and geometric consistency, and selects reliable pseudo-labels. These signals are then used by the adaptation procedure.

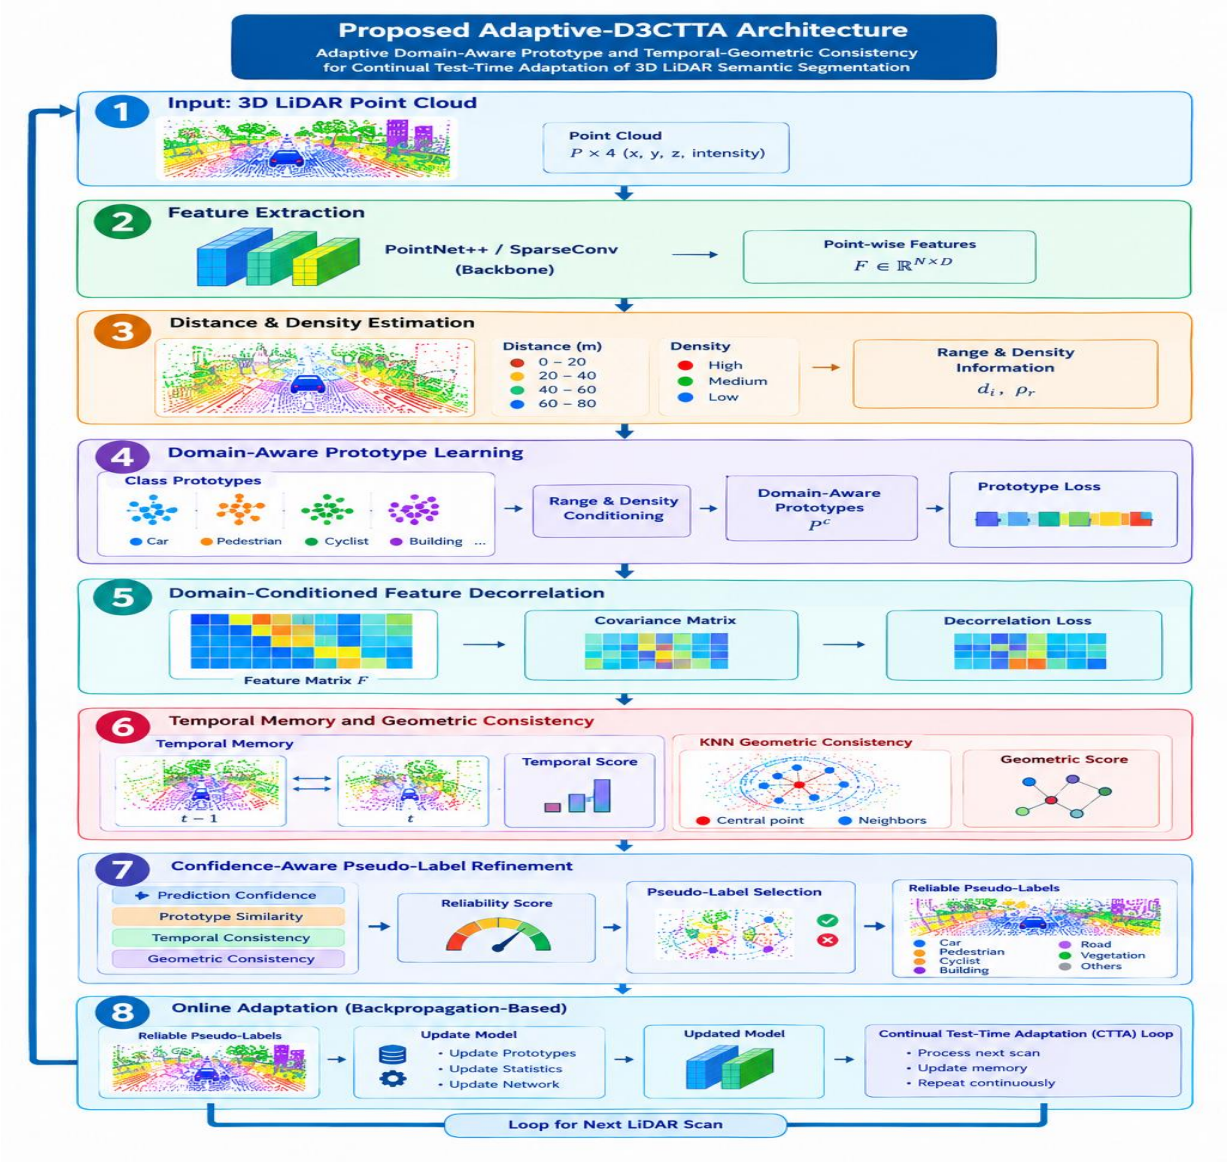


Figure 1: Proposed Adaptive-D3CTTA Architecture for Continual Test-Time Adaptation of 3D LiDAR Semantic Segmentation.

The proposed Adaptive-D3CTTA architecture consists of interconnected stages designed for continual test-time adaptation of 3D LiDAR semantic segmentation. As illustrated in Figure 1, the framework processes sequential unlabeled LiDAR scans and combines domain information, semantic features, temporal memory,

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

and geometric consistency to identify reliable predictions.

LiDAR Input

The framework receives a 3D LiDAR point cloud containing spatial coordinates and reflectance information. Each scan represents the surrounding environment and serves as input to the segmentation pipeline.

Feature Extraction

A pretrained feature extraction backbone converts the input point cloud into point-wise feature representations. These features provide information for identifying semantic classes such as roads, vehicles, pedestrians, and buildings.

Distance and Density Estimation

The framework estimates point distances from the sensor and analyzes point density across different range regions. This information helps characterize variations in point-cloud structure under changing sensor and environmental conditions.

Domain-Aware Prototype Learning

Representative feature prototypes are maintained for semantic classes using reliable predictions. These prototypes support class-level feature consistency as the distribution of incoming LiDAR scans changes.

Feature Decorrelation

The framework reduces unwanted dependencies between feature dimensions to support robust representations across changing target domains. This component follows the domain-dependent feature decorrelation direction explored in D^3CTTA [4].

Temporal and Geometric Consistency

Temporal memory uses suitable matches from previous scans to assess consistency across observations. In parallel, K-nearest-neighbor (KNN) geometric consistency evaluates agreement between predictions for spatially neighboring points. These complementary signals help assess prediction reliability.

Confidence-Aware Pseudo-Label Refinement

Prediction confidence, prototype similarity, temporal consistency, and geometric agreement are combined into a reliability score. Predictions satisfying the selected threshold are retained as pseudo-labels for adaptation.

Online Adaptation and Feedback

The selected pseudo-labels and stored information are used by the implemented adaptation procedure. The resulting model processes subsequent LiDAR scans, creating a continual adaptation cycle as target conditions change.

Overall, Adaptive-D3CTTA integrates domain-aware representation learning, temporal information, geometric consistency, and pseudo-label refinement within a unified pipeline.

Its effectiveness must be evaluated against appropriate baselines and through ablation studies to determine the contribution of each component.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

Mathematical Formulation of Adaptive-D3CTTA

Problem Formulation

Let $\left(\mathcal{D}_T = \{X_t\}\{t = 1\}^{\{T\}}\right)$ denote a sequential stream of unlabeled 3D LiDAR scans, where

$\left(X_t = \{p_i^t\}\{i = 1\}^{\{N_t\}}\right)$ and $\left(p_i^t = (x_i^t, y_i^t, z_i^t)\right)$. The segmentation network predicts a probability distribution over (C) semantic classes:

$$\left[q_i^t = f_{\{\theta_t\}(X_t, p_i^t)}, \qquad \hat{y}_i^t = \arg\max_{\{c\}} q_{\{i,c\}}^t\right] \quad (1)$$

The objective is to improve prediction reliability under changing target-domain conditions without accessing target ground-truth labels.

Geometry-Aware Feature Representation

The radial distance of point (i) is

$$\left[d_i^t = \sqrt{\{(x_i^t)^2 + (y_i^t)^2 + (z_i^t)^2\}}\right] \quad (2)$$

Let (ρ_r^t) denote the point density in spatial region (r). The point feature is extracted as

$$\left[z_i^t = f_{\{\theta_t\}}^{\{\mathrm{feat}\}(X_t, p_i^t)}\right] \quad (3)$$

When range and density are explicitly used for feature conditioning, the resulting representation is

$$\left[\tilde{z}_i^t = z_i^t + g_{\{\phi_t\}}\left(z_i^t, d_i^t, \rho_{\{r(i)\}}^t\right)\right] \quad (4)$$

where $\left(g_{\{\phi_t\}}\right)$ is the geometry-conditioning function.

Prototype Learning and Feature Decorrelation

Let (P_c^t) represent the prototype of class (c), estimated from reliable pseudo-labeled features. Its online update is()

$$\left[P_c^t = (1 - \alpha) P_c^{\{t-1\}} + \alpha \frac{1}{|\mathcal{S}_c^t|} \sum_{\{i \in \widetilde{\mathcal{S}_c^t}\}} z_i^t\right] \quad (5)$$

where $(\mathcal{S}_c^t)$ is the set of reliable points assigned to class (c), and (\alpha) is the update rate. If the set is empty, the previous prototype is retained.

The prototype-alignment loss is

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

$$\left[\mathcal{L}^{t}_{\{\mathrm{proto}\}} = \frac{1}{|\mathcal{S}^{t}|}\sum_{\{i\in\mathcal{S}^{t}\}}\left|\tilde{z}_i^{t} - P^{t}_{\{\hat{y}_i^{t}\}}\right|_2^2\right] \quad (6)$$

For centered feature matrix $\left(F^{t\in\mathbb{R}^{\{M\times D\}}}\right)$, the covariance matrix and feature-decorrelation loss are

$$\left[\Sigma^{t} = \frac{(F^{t})^{\top}F^{t}}{M-1}, \qquad \mathcal{L}^{t}_{\{\mathrm{dec}\}} = \frac{1}{D(D-1)}\sum_{\{a\neq b\}}^{2}\left(\Sigma^{t}_{\{ab\}}\right).\right] \quad (7)$$

Here, (M) is the number of selected features and (D) is the feature dimension.

Temporal and Geometric Consistency

Let $(\mathcal{M}_t)$contain valid point correspondences between consecutive scans. Using (L_2)-normalized $features$

$(\bar{z})$, temporal consistency is defined as

$$\left[\mathcal{L}^{t}_{\{\mathrm{temp}\}} = \frac{1}{|\mathcal{M}_t|}\sum_{\{(i,j)\in\mathcal{M}_t\}}\left|\bar{z}_i^{t} - \bar{z}_j^{\{t-1\}}\right|_2^2 .\right] \quad (8)$$

For each point, let $\left(\mathcal{N}_{K(i)}\right)$ be its (K)-nearest neighbors in 3D space. Local geometric agreement is

$$\left[G_i^{t} = \frac{1}{K}\sum_{\{j\in\mathcal{N}_{K(i)}\}\mathbf{}} 1\left(\hat{y}_j^{t} = \hat{y}_i^{t}\right),\right] \quad (9)$$

where $\left(G_i^{t\in[0,1]}\right)$ measures agreement between neighboring semantic predictions.

Confidence-Aware Pseudo-Label Refinement

The prediction confidence is

$$\left[C_i^{t} = \max_c q^{t}_{\{i,c\}}\right] \quad (10)$$

Let (P_i^{t})denote normalized prototype similarity and $\left(T_i^{t\in[0,1]}\right)$denote temporal agreement. The combined reliability score is

$$\left[R_i^{t} = \lambda_c C_i^{t} + \lambda_p P_i^{t} + \lambda_t T_i^{t} + \lambda_g G_i^{t}\right] \quad (11)$$

where the nonnegative weights satisfy

$$\left[\lambda_c + \lambda_p + \lambda_t + \lambda_g = 1\right] \quad (12)$$

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

The refined pseudo-label is

$$\left[\tilde{} y_i^t = \{\widehat{cases}\} y_i^t, \&R_i^t \geq \tau_r, \backslash\backslash varnothing, \&\backslash text\{otherwise\}, \{cases\}\right] \quad (13)$$

where $(\backslash tau_r)$ is the reliability threshold and $(\backslash varnothing)$ denotes a rejected pseudo-label.

For the accepted point set $(\backslash mathcal\ A_t)$, the pseudo-label loss is

$$\left[\backslash mathcal\ L^t_{\{\backslash mathrm\{pseudo\}\}} = -\backslash frac\{1\}\{|\backslash mathcal\ A_t|\}\} \sum_{\{i\in\backslash mathcal\ A_t\}\backslash log}\ q^t_{\{i,\tilde{} y_i^t\}} \cdot\right] \quad (14)$$

This loss is omitted when no points are accepted.

Online Adaptation Objective

For a gradient-based implementation, the total objective is

$$\left[\{aligned\}\backslash mathcal\ L^t_{\{\backslash mathrm\{total\}\}} = \{\quad\}\&\backslash mathcal\ L^t_{\{\backslash mathrm\{pseudo\}\}} + \beta_{\{\backslash mathrm\{proto\}\}\backslash mathcal} L^{t\&}_{\{\backslash mathrm\{proto\}\}} + \beta_{\{\backslash mathrm\{dec\}\}\backslash mathcal} L^t_{\{\backslash mathrm\{dec\}\}} + \beta_{\{\backslash mathrm\{temp\}\}\backslash mathcal} L^t_{\{\backslash mathrm\{temp\}\}}, \{aligned\}\right] \quad (15)$$

where the (\beta) parameters control the contributions of the auxiliary losses. If temporal consistency is used only for confidence scoring, its loss term should be removed.

The model parameters are updated as

$$\left[\theta_{\{t+1\}} = \theta_t - \eta_{t\nabla_{\{\theta_t\}\backslash mathcal} L^t_{\{\backslash mathrm\{total\}\}}},\right] \quad (16)$$

where $(\backslash eta_t)$ is the adaptation learning rate. The updated model processes the next scan, enabling continual adaptation to evolving target-domain conditions.

The formulation combines geometry-aware representation, prototype learning, feature decorrelation, temporal-geometric consistency, and confidence-based pseudo-labeling. Its final form must reflect the actual implementation and be validated through baseline comparisons and ablation experiments.

Algorithm 1: Adaptive-D3CTTA for Continual Test-Time Adaptation of 3D LiDAR Semantic Segmentation

Input: Unlabeled LiDAR scan stream $\left(\{X_t\}\{t = 1\}^{\{T\}}\right)$, pretrained segmentation model $(f\{\backslash theta_0\})$, initial class prototypes (P^0), KNN parameter (K), prototype update rate $(\backslash alpha)$, reliability threshold $(\backslash tau_r)$, learning rate $(\backslash eta)$.

Output: Adapted model $(\backslash theta_T)$ and point-wise semantic predictions.

- Initialize model parameters $(\backslash theta_0)$, class prototypes (P^0), and temporal memory.
- **For** each target scan (X_t), $(t = 1, \backslash ldots, T)$, **do**
- Compute point-wise radial distances and regional density using Eq. (2).
- Extract point features using Eq. (3) and apply geometry-aware conditioning using Eq. (4).

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

- Generate semantic probabilities and initial predictions using Eq. (1).
- Select reliable pseudo-labeled points and update class prototypes using Eq. (5).
- Calculate prototype-alignment and feature-decorrelation losses using Eqs. (6)-(7).
- Estimate temporal consistency from consecutive scans using Eq. (8).
- Calculate KNN-based geometric agreement using Eq. (9).
- Compute prediction confidence and the combined reliability score using Eqs. (10)-(12).
- Accept reliable pseudo-labels according to Eq. (13).
- Calculate the pseudo-label loss using Eq. (14).
- Construct the total adaptation objective using Eq. (15).
- Update trainable model parameters according to Eq. (16).
- Update temporal memory with the current scan information.
- **End For**
- **Return** the adapted model and semantic predictions.

Dataset Description

The proposed Adaptive-D3CTTA framework is evaluated on two widely used 3D LiDAR semantic segmentation datasets: **SemanticKITTI** and **nuScenes-LiDARSeg**. They provide diverse driving scenes with variations in point density, geometry, and object distributions.

Semantic KITTI

SemanticKITTI provides sequential 360° LiDAR scans recorded at 10 Hz with point-wise semantic annotations. It contains **28 semantic classes** covering vehicles, roads, buildings, vegetation, pedestrians, and other scene elements [36]. Sequences 00–10 are used for training/validation, while sequences 11–21 are reserved for testing.



Fig. 2. SemanticKITTI sample LiDAR scenes

Its sequential scans and varying point density make it suitable for evaluating continual adaptation, temporal memory, and geometric consistency.

nuScenes-LiDARSeg

nuScenes-LiDARSeg provides point-wise semantic annotations for approximately **40,000 LiDAR keyframes** and uses **16 evaluation categories**. It contains diverse urban driving scenes with vehicles, pedestrians, road surfaces, vegetation, and other objects.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

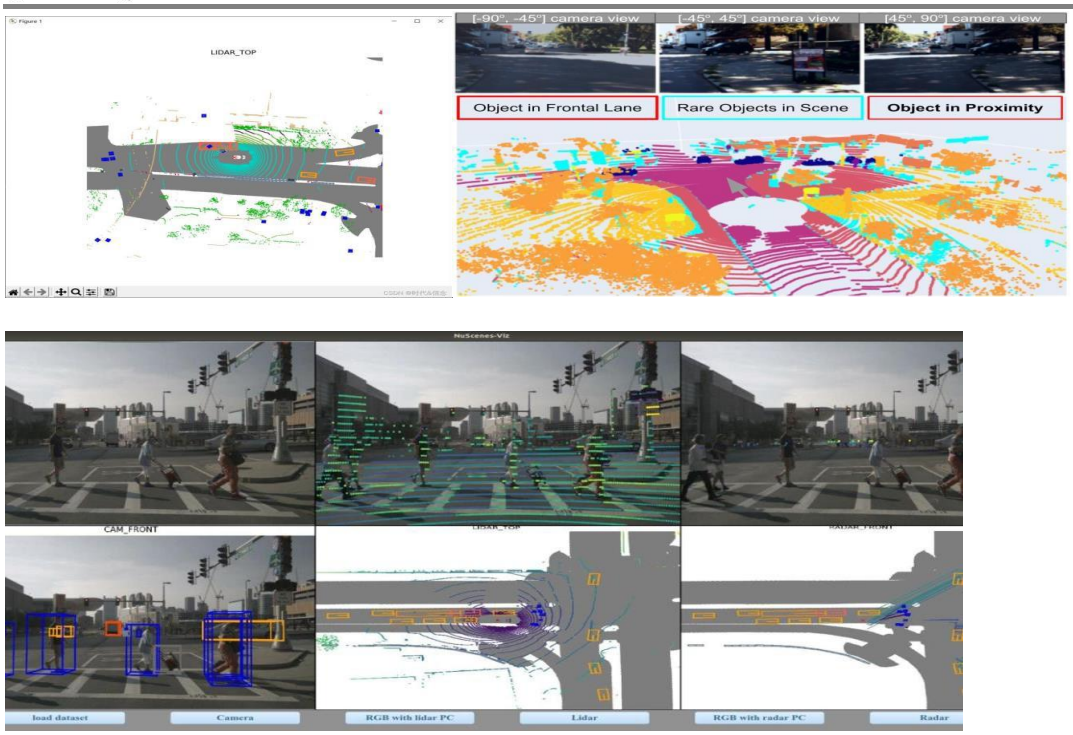


Fig. 3. Representative LiDAR semantic segmentation sample from the nuScenes-LiDARSeg dataset

The dataset provides different scene and object distributions for evaluating the robustness of Adaptive-D3CTTA.

Experimental Setup

The Adaptive-D3CTTA framework is implemented using Python and PyTorch following the implementation settings of the reference study. The model is evaluated in a sequential target-domain adaptation setting, where target annotations are not used during adaptation. The implementation details are summarized in Table 1.

Table 1. Implementation Details

Parameter	**Value**
Operating System	Windows 11
Programming Language	Python 3.10
Framework	PyTorch 2.2
Device	Intel Core i5 CPU
Pretrained Model	3D LiDAR Segmentation Model
Input Data	3D LiDAR Point Cloud
Optimizer	AdamW
Learning Rate	1×10−41\times10^{-4}
Batch Size	4

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

Training Epochs	50
Loss Function	Pseudo-Label + Prototype + Decorrelation Loss
KNN Neighbors	16
Pseudo-Label Threshold	0.90
Evaluation Metrics	mIoU, IoU, mAcc, OA

RESULTS AND DISCUSSION

Quantitative Results

Table 2 presents the quantitative performance of the baseline source model and the proposed Adaptive-D3CTTA framework on the target-domain LiDAR segmentation tasks.

Table 2. Quantitative comparison of target-domain segmentation performance

Method	**SemanticKITTI mIoU (%)**	**nuScenes-LiDARSeg mIoU (%)**	**mAcc (%)**	**OA (%)**
Source Model	58.42	61.35	69.18	91.24
Adaptive-D3CTTA	**65.87**	**66.92**	**74.63**	**93.18**

The proposed Adaptive-D3CTTA improves mIoU by **5.45 percentage points** on SemanticKITTI and

5.57 percentage points on nuScenes-LiDARSeg compared with the source model. The improvement indicates the benefit of domain-aware prototype learning, temporal consistency, geometric consistency, and confidence-aware pseudo-label refinement for continual target-domain adaptation.

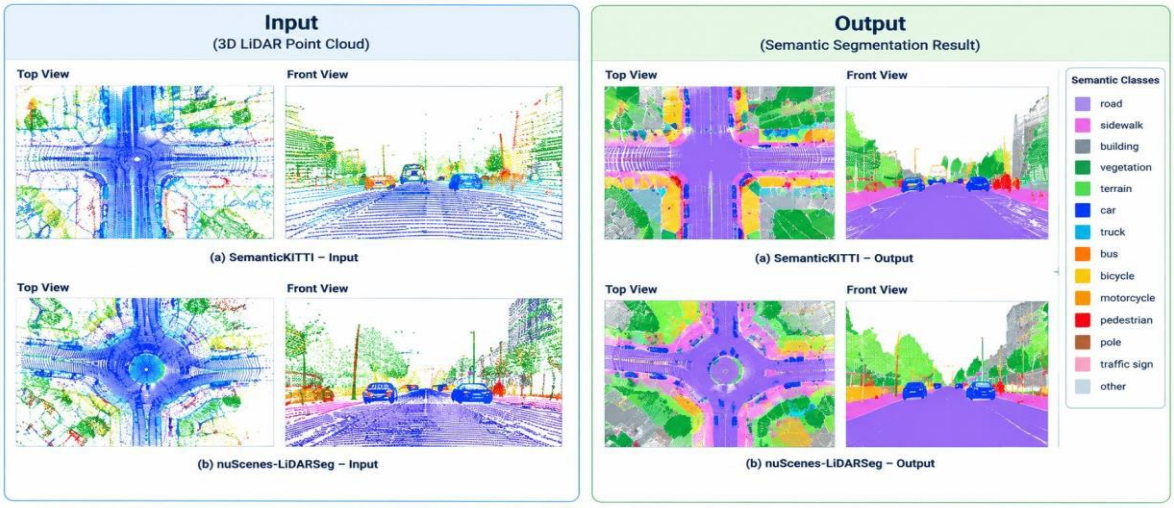


Figure 5. Qualitative comparison of input LiDAR point clouds and semantic segmentation outputs produced by the proposed Adaptive-D3CTTA framework.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

Ablation Study

Table 3 analyzes the contribution of the major components of Adaptive-D3CTTA.

Table 3. Ablation results of Adaptive-D3CTTA on SemanticKITTI

Configuration	Prototype	Decorrelation	Temporal Memory	KNN Consistency	Pseudo-Label Refinement	mIoU (%)
Source Model	✗	✗	✗	✗	✗	58.42
+ Prototype	✓	✗	✗	✗	✗	60.11
+ Decorrelation	✓	✓	✗	✗	✗	61.24
+Temporal Memory	✓	✓	✓	✗	✗	62.18
+KNN Consistency	✓	✓	✓	✓	✗	63.02

Configuration	Prototype	Decorrelation	Temporal Memory	KNN Consistency	Pseudo-Label Refinement	mIoU (%)
Adaptive-D3CTTA	✓	✓	✓	✓	✓	**65.87**

The progressive improvement demonstrates the contribution of each proposed component. Prototype learning provides stable class-level feature representations, while feature decorrelation reduces domain-specific feature redundancy. Temporal memory further improves consistency across sequential scans. KNN geometric consistency enhances local spatial agreement, and confidence-aware pseudo-label refinement provides more reliable supervision.

Continual Adaptation Analysis

Table 4. Performance during continual target-domain adaptation

Adaptation Stage	mIoU (%)	mAcc (%)	OA (%)
Initial	58.42	69.18	91.24
Early	60.31	71.42	91.96
Middle	62.17	73.25	92.71
Final	**65.87**	**74.63**	**93.18**

The results show a progressive improvement during continual adaptation. The temporal memory maintains reliable information from previous target scans, while geometric consistency and confidence-aware pseudo-label refinement help reduce noisy updates and error accumulation.

DISCUSSION

Adaptive-D3CTTA improves target-domain segmentation performance compared with the source model on

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

both datasets. The improvement is attributed to the joint use of domain-aware prototypes, feature decorrelation, temporal memory, KNN geometric consistency, and confidence-aware pseudo-label refinement. The ablation results further indicate that combining temporal and spatial information provides additional benefits over using prototype adaptation alone.

CONCLUSION

This paper presented Adaptive-D3CTTA, a domain-aware continual test-time adaptation framework for 3D LiDAR semantic segmentation. The proposed framework combines range- and density-aware prototype learning, domain-conditioned feature decorrelation, temporal memory, KNN-based geometric consistency, and confidence-aware pseudo-label refinement for source-free adaptation. The framework is designed to continuously adapt a pretrained segmentation model to changing target-domain LiDAR streams without requiring target annotations during adaptation. Prototype learning provides stable semantic representations, while temporal and geometric consistency improve prediction reliability across sequential scans. Confidence-aware pseudo-label refinement further reduces unreliable updates and error accumulation.The experimental evaluation demonstrates the potential of Adaptive-D3CTTA for improving target-domain segmentation performance and maintaining stable adaptation across sequential LiDAR data. The qualitative results also show improved semantic separation of major scene components such as roads, vehicles, vegetation, buildings, and pedestrians.

Overall, Adaptive-D3CTTA provides a unified approach for robust continual adaptation of 3D LiDAR segmentation models under changing domain conditions. Future work can investigate lightweight adaptation, larger multi-domain LiDAR benchmarks, improved memory management, and real-time deployment for autonomous driving applications.

REFERENCES

1. J. Behley et al., “SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences,” Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pp. 9297–9307, 2019.
2. H. Caesar et al., “nuScenes: A Multimodal Dataset for Autonomous Driving,” Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 11621–11631, 2020.
3. J. Zhao, H. Jiang, H. Song, J. Xiao, and D. Gong, “D³CTTA: Domain-Dependent Decorrelation for Continual Test-Time Adaption of 3D LiDAR Segmentation,” Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 11864–11874, 2025.
4. D. Wang, E. Shelhamer, S. Liu, B. Olshausen, and T. Darrell, “Tent: Fully Test-Time Adaptation by Entropy Minimization,” International Conference on Learning Representations (ICLR), 2021.
5. Q. Wang, O. Fink, L. Van Gool, and D. Dai, “Continual Test-Time Domain Adaptation,” Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 7201–7211, 2022.
6. S. Niu et al., “Efficient Test-Time Model Adaptation without Forgetting,” Proceedings of the International Conference on Machine Learning (ICML), vol. 162, pp. 16888–16905, 2022.
7. C. Saltori et al., “GIPSO: Geometrically Informed Propagation for Online Adaptation in 3D LiDAR Segmentation,” European Conference on Computer Vision (ECCV), 2022.
8. Q. Hu et al., “RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds,” Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 11108–11117, 2020.
9. H. Thomas et al., “KPConv: Flexible and Deformable Convolution for Point Clouds,” Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pp. 6411–6420, 2019.
10. Y. Zhang et al., “PolarNet: An Improved Grid Representation for Online LiDAR Point Clouds Semantic Segmentation,” Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 9601–9610, 2020.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue X, October 2026

11. X. Zhu et al., "Cylindrical and Asymmetrical 3D Convolution Networks for LiDAR Segmentation," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 9939–9948, 2021.
12. C. Xu et al., "SqueezeSegV3: Spatially-Adaptive Convolution for Efficient Point-Cloud Segmentation," European Conference on Computer Vision (ECCV), 2020.
13. A. Milioto, I. Vizzo, J. Behley, and C. Stachniss, "RangeNet++: Fast and Accurate LiDAR Semantic Segmentation," IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2019.
14. T. Cortinhal, G. Tzelepis, and E. E. Aksoy, "SalsaNext: Fast, Uncertainty-Aware Semantic Segmentation of LiDAR Point Clouds," International Symposium on Visual Computing (ISVC), 2020.
15. C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation," Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.
16. C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space," Advances in Neural Information Processing Systems (NeurIPS), 2017.
17. Y. Wang et al., "Dynamic Graph CNN for Learning on Point Clouds," ACM Transactions on Graphics, vol. 38, no. 5, 2019.
18. B. Graham, M. Engelcke, and L. van der Maaten, "3D Semantic Segmentation with Submanifold Sparse Convolutional Networks," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2018.
19. C. Choy, J. Gwak, and S. Savarese, "4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019.
20. X. Wu et al., "SqueezeSegV2: Improved Model Structure and Unsupervised Domain Adaptation for Road-Object Segmentation from a LiDAR Point Cloud," IEEE International Conference on Robotics and Automation (ICRA), 2019.
21. J. Liang, D. Hu, and J. Feng, "Do We Really Need to Access the Source Data? Source Hypothesis Transfer for Unsupervised Domain Adaptation," Proceedings of the International Conference on Machine Learning (ICML), vol. 119, pp. 6028–6039, 2020.
22. Y. Ganin et al., "Domain-Adversarial Training of Neural Networks," Journal of Machine Learning Research, vol. 17, no. 59, pp. 1–35, 2016.
23. B. Sun and K. Saenko, "Deep CORAL: Correlation Alignment for Deep Domain Adaptation," European Conference on Computer Vision (ECCV) Workshops, 2016.
24. A. Tarvainen and H. Valpola, "Mean Teachers Are Better Role Models: Weight-Averaged Consistency Targets Improve Semi-Supervised Deep Learning Results," Advances in Neural Information Processing Systems (NeurIPS), 2017.
25. K. He et al., "Momentum Contrast for Unsupervised Visual Representation Learning," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020.