
How Do Images Align and Complement LiDAR? Towards a Harmonized Multi-modal 3D Panoptic Segmentation

Yining Pan¹ Qiongjie Cui¹ Xulei Yang² Na Zhao¹

Abstract

LiDAR-based 3D panoptic segmentation often struggles with the inherent sparsity of data from LiDAR sensors, which makes it challenging to accurately recognize distant or small objects. Recently, a few studies have sought to overcome this challenge by integrating LiDAR inputs with camera images, leveraging the rich and dense texture information provided by the latter. While these approaches have shown promising results, they still face challenges, such as misalignment during data augmentation and the reliance on post-processing steps. To address these issues, we propose **Image-Assists-LiDAR (IAL)**, a novel multi-modal 3D panoptic segmentation framework. In IAL, we first introduce a modality-synchronized data augmentation strategy, **PieAug**, to ensure alignment between LiDAR and image inputs from the start. Next, we adopt a transformer decoder to directly predict panoptic segmentation results. To effectively fuse LiDAR and image features into tokens for the decoder, we design a **Geometric-guided Token Fusion (GTF)** module. Additionally, we leverage the complementary strengths of each modality as priors for query initialization through a **Prior-based Query Generation (PQG)** module, enhancing the decoder’s ability to generate accurate instance masks. Our IAL framework achieves state-of-the-art performance compared to previous multi-modal 3D panoptic segmentation methods on two widely used benchmarks. Code and models are publicly available at <https://github.com/IMPL-Lab/IAL.git>.

1. Introduction

3D Panoptic segmentation simultaneously assigns semantic labels and identifies distinct instances, effectively unifying 3D semantic (Zhao et al., 2021; Xu et al., 2023b) and instance (Li & Zhao, 2024) segmentation to provide a holistic understanding of the scene. This task is particularly crucial for real-world applications, such as dynamic object tracking (Yang et al., 2023) and autonomous driving (Hong et al., 2021; Cao et al., 2024). LiDAR is an indispensable sensor for perceiving the 3D world, with its LiDAR point cloud typically serving as the sole input for 3D panoptic segmentation (Razani et al., 2021; Zhou et al., 2021; Li et al., 2022a). However, LiDAR inherently faces limitations in detecting small or distant objects due to its radial emission pattern, which results in sparse returns along each laser ray (Li et al., 2022b). Consequently, entities that are small or located at a distance may not receive sufficient information. In contrast, camera images provide denser and more detailed representations, effectively compensating for the sparsity in LiDAR data, particularly in these challenging scenarios.

This complementary nature has motivated the use of multi-modal information for enhanced panoptic segmentation. Recently, LCPS (Zhang et al., 2023) and Panoptic-FusionNet (Song et al., 2024) have pioneered LiDAR-and-image fusion methods for multi-modal 3D panoptic segmentation. However, these methods only perform augmentation on the LiDAR side, leading to misalignment between the two modalities. This misalignment hinders effective integration of information from both modalities, causing the models to rely predominantly on LiDAR data, rather than fully utilizing both LiDAR and image data. Moreover, the prediction heads in LCPS and Panoptic-FusionNet do not directly predict 3D panoptic segmentation results. Instead, they use a post-processing strategy that involves clustering instances after semantic segmentation (Zhou et al., 2021). This strategy presents two issues: 1) The post-processing step is inefficient and limits the effectiveness of segmentation by relying on preliminary results; 2) Their convolution-based prediction heads rely on local context, which may be suboptimal for panoptic segmentation as it requires global context for accurate predictions.

To address the first limitation, we propose a modality-

¹Singapore University of Technology and Design (SUTD)
²Institute for Infocomm Research (I2R), A*STAR, Singapore. Correspondence to: Na Zhao <na.zhao@sutd.edu.sg>.

synchronized augmentation strategy – **PieAug**. PieAug ensures that the augmentation of multi-view images is synchronized with the augmentation of their corresponding LiDAR pairs, enabling well-aligned and enriched LiDAR and image inputs. Notably, PieAug is a general multi-modal data augmentation strategy designed for outdoor segmentation tasks. Its LiDAR-specific augmentation can be seen as a generalization of existing point cloud augmentation techniques, e.g., instance augmentation (Zhou et al., 2021), PolarMix (Xiao et al., 2022), LaserMix (Kong et al., 2023b).

To overcome the second post-processing limitation, we propose adopting a transformer decoder for multi-modal 3D panoptic segmentation, inspired by its success in 3D panoptic segmentation (Xiao et al., 2025) and multi-modal 3D object detection (Bai et al., 2022; Yan et al., 2023). By leveraging global context and directly predicting class labels and mask outputs, the transformer decoder eliminates the inefficiencies and constraints associated with post-processing. Despite its promise, adopting a transformer decoder introduces new challenges, particularly in *designing effective queries and tokens as inputs*. To overcome these challenges, we introduce a **Geometric-guided Token Fusion** (GTF) module and a **Prior-based Query Generation** (PQG) module. Combined with the PieAug strategy, these components form our proposed solution: a novel **Image-Assist-LiDAR** transformer-based framework, named **IAL**, for multi-modal 3D panoptic segmentation in autonomous driving scenarios.

GTF module integrates the sparse, cylinder-shaped LiDAR features with the compact, grid-shaped image features to create input tokens. Specifically, we adopt Cylinder3D (Zhu et al., 2020) to extract LiDAR features and use raw LiDAR points as geometric guidance to locate corresponding image patches for each cylindrical voxel. Additionally, we design a scale-aware positional embedding to encode the cylindrical voxels’ locations and their receptive fields, facilitating the fusion of image patches and cylindrical voxels. This approach enhances feature fusion while mitigating projection errors caused by variations in cylindrical voxel shapes.

Our PQG module leverages prior knowledge from LiDAR and image inputs, which provide complementary strengths for object perception, to improve query initialization. Specifically, we generate two groups of instance queries – *geometric-prior* and *texture-prior* instance queries – derived from LiDAR and image modalities, respectively. Geometric-prior queries exploit LiDAR’s geometric features, which are well-suited for detecting nearby or large objects rich in geometric information. In contrast, texture-prior queries leverage images by applying state-of-the-art detection and segmentation models, such as Grounding-DINO (Liu et al., 2023a) and SAM (Kirillov et al., 2023), to better identify distant and small objects. To handle challenging scenarios where both LiDAR and images fail to provide

reliable instance queries, we introduce a set of learnable parameters as *no-prior instance queries*. Consequently, the three groups of instance queries, combined with a set of semantic queries, are input into the transformer decoder to predict instance masks and semantic labels.

Our contributions can be summarized as: **1)** We present IAL, a novel transformer-based multi-modal framework for multi-modal 3D panoptic segmentation, eliminating the cumbersome post-processing steps required by previous methods. **2)** We propose PieAug, a multi-modal augmentation technique that not only addresses the asynchronization issue but also serves as a generalized formulation of existing LiDAR augmentation methods. **3)** We design the GTF and PQG modules that can effectively fuse image and LiDAR features as tokens and queries for the transformer decoder. **4)** Our IAL achieves state-of-the-art performance in outdoor panoptic segmentation, surpassing previous methods by 2.5% and 4.1% in PQ on the nuScenes and SemanticKITTI benchmarks, respectively.

2. Related Work

3D Panoptic Segmentation. Most advanced approaches for LiDAR-based 3D panoptic segmentation can be categorized into three main groups: top-down, bottom-up, and single-path methods (Li & Chen, 2022). The top-down method typically follows a detection-first principle, where bounding boxes are predicted initially, followed by the generation of instance masks from points within those boxes (Sirohi et al., 2021; Xu et al., 2023a; Ye et al., 2023). In contrast, bottom-up methods (Zhou et al., 2021; Hong et al., 2021; Razani et al., 2021; Xu et al., 2022; Li et al., 2022a) begin with semantic segmentation predictions and then generate instance masks through operations such as grouping and clustering. Both top-down and bottom-up approaches are limited by the performance of their preliminary predictions (object detection or semantic segmentation, respectively), which hinders their ability to achieve holistic perception. On the other hand, single-path methods treat panoptic segmentation as a unified task, simultaneously addressing the segmentation of both “stuff” and “thing” classes. For example, MaskRange (Gu et al., 2022), MaskPLS (Marcuzzi et al., 2023), and P3Former (Xiao et al., 2025) utilize learnable queries to predict masks and classes for both “thing” and “stuff” object types. Additionally, PUPS (Su et al., 2023) and DQFormer (Yang et al., 2025) explore the use of auxiliary categorical and positional embeddings for query initialization. However, these methods mainly rely on LiDAR inputs and often struggle with recognizing small or distant objects, where geometric information degrades. This limitation inspires the use of camera images, which contain rich texture information, to enhance panoptic segmentation.

Multi-modal 3D Scene Understanding. Multi-modal

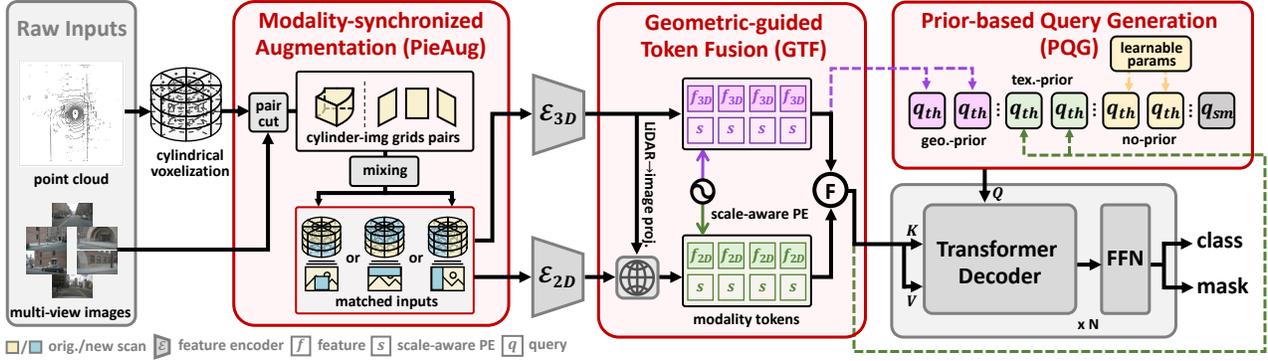


Figure 1. The architecture overview of our Image-Assists-LiDAR (IAL) framework. We first voxelize the point cloud into cylindrical voxels. In **PieAug**, we synchronize augmentation by pairing cylindrical and image grids, mixing original and new scans in different modes. Using a transformer-based structure, we design token fusion (GTF) and query initialization (PQG) modules to align and complement both modalities. In **GTF**, features from both modalities are bridged via LiDAR-to-image projection and scale-aware positional embeddings. We generate instance queries for “thing” classes using modality knowledge (specifically, geometric and texture priors from LiDAR and image features) and learnable parameters (for instances without advanced modality priors) in **PQG**. These instance queries, along with semantic queries, are iteratively updated through multiple transformer decoder layers to produce the final panoptic predictions.

learning with LiDAR and images has been extensively studied for 3D semantic segmentation (Zhuang et al., 2021; Krispel et al., 2020; Yan et al., 2022; Li et al., 2023b; Liu et al., 2023b; Man et al., 2023; Wu et al., 2024; An et al., 2025) and object detection (Liang et al., 2022; Yin et al., 2024; Bai et al., 2022; Li et al., 2024). A key challenge in these tasks is to effectively fuse the features of both modalities to leverage their complementary strengths. To address this, existing methods design fusion modules to align data from different sensors. Additionally, some object detection studies (Bai et al., 2022; Chen et al., 2023; Li et al., 2024; Zhang et al., 2024) have shown that image-driven queries improve the detection of challenging objects. While these approaches are not directly applicable to panoptic segmentation, they inspire our design of the Geometric-guided Token Fusion (GTF) and Prior-based Query Generation (PQG). For multi-modal 3D panoptic segmentation, LCPS (Zhang et al., 2023) and Panoptic-FusionNet (Song et al., 2024) are two pioneering works. LCPS designs a point-to-mask mapping for LiDAR to image fusion, while Panoptic-FusionNet directly applies point-to-pixel mapping by geometric information. However, these methods primarily apply augmentations only to LiDAR, leading to misalignment between modalities and an over-reliance on LiDAR features. Additionally, their reliance on post-processing makes panoptic inference inefficient and limits the effectiveness of true multi-modal perception.

LiDAR-Based Data Augmentation. Studies on LiDAR-based data augmentation often rely on instance- or scene-level mixing. Instance mixing methods (Zhou et al., 2021; Xiao et al., 2022; Zhao et al., 2022) augment point clouds by copying instance points from one scan to another, while scene-wise mixing approaches, such as LaserMix (Kong et al., 2023b) and PolarMix (Xiao et al., 2022), divide the scene into multiple intervals along inclination or azimuth

angles and selectively swap these intervals between scans. RangeFormer (Kong et al., 2023a) represents the 3D scene as a range-view image and applies tailored augmentation strategies for range-view learning. Similarly, UniMix (Zhao et al., 2024) extends scene-wise mixing to different attributes, including intensity and semantic channels. In multi-modal 3D scene understanding, data augmentation for both modalities remains underexplored. LaserMix++ (Kong et al., 2024) extends LaserMix to multi-modal scenes but relies on a single augmentation strategy. MSeg3D (Li et al., 2023b) applies asymmetric augmentation to each modality, using only simple local transformations – reducing the risk of misalignment but at the cost of data diversity. To address these limitations, we propose a general multi-modal augmentation strategy that incorporates diverse instance- and scene-level mixing to enhance both cross-modal alignment and panoptic segmentation performance.

3. Methodology

In the multi-modal 3D panoptic segmentation, we are given a 3D point cloud consisting of N discrete sampling points, denoted as $\mathbf{P} = \{\mathbf{p}_j \in \mathbb{R}^{1 \times 4}\}_{j=1}^N$, where each point \mathbf{p}_j contains its Cartesian coordinates in Euclidean space and its reflection intensity. The point cloud is associated with K view images, represented as $\mathcal{I} = \{\mathbf{I}_k \in \mathbb{R}^{H \times W \times 3}\}_{k=1}^K$, H and W denote the height and width of images. The goal of this task is to effectively utilize \mathbf{P} and \mathcal{I} to predict both semantic and instance labels for each point.

Framework Overview. In this paper, we introduce Image-Assist-LiDAR (IAL), a novel transformer-based framework for multi-modal 3D panoptic segmentation, as illustrated in Fig. 1. To process the sparse and irregular LiDAR point cloud, we first apply *cylindrical voxelization*, converting points into cylindrical-shaped voxels based on their polar coordinates. As a result, each voxel \mathbf{v}_i contains a varying

number of points, with voxel shapes differing along the radial axis. The point cloud \mathbf{P} can then be represented as $\mathbf{V} = \{\mathbf{v}_i\}_{i=1}^M$, where M is the number of valid cylindrical voxels. We apply modality-synchronized augmentation through our proposed PieAug strategy (Sec. 3.1), ensuring consistency across LiDAR and image data by pairing each voxel with its corresponding image regions and employing a generalized augmentation operator for diverse effects.

The augmented 3D voxels and images are then processed by 3D encoder \mathcal{E}_{3D} and 2D encoder \mathcal{E}_{2D} , extracting voxel-wise features $\mathbf{F}^{3D} \in \mathbb{R}^{M \times D}$ and image features $\mathbf{F}^{2D} \in \mathbb{R}^{K \times H \times W \times D}$, where D is the feature dimension. The 3D encoder uses Cylinder3D (Zhu et al., 2020), known for its strong generalization in 3D panoptic segmentation (Xiao et al., 2025; Zhang et al., 2023), and 2D encoder is SwiftNet (Oršić & Šegvić, 2021) with a ResNet-18 backbone.

Next, we use \mathbf{F}^{3D} and \mathbf{F}^{2D} to create tokens and queries for a transformer decoder, enabling cross-modal interaction. Token features are formed by concatenating voxel features with their aligned image counterparts, guided by a unified, scale-aware positional embedding (Sec.3.2). Meanwhile, instance queries \mathbf{q}_{th} are initialized using modality-compensated priors (Sec.3.3). These queries, along with semantic queries \mathbf{q}_{sm} are fed into the transformer decoder to predict the instance masks and semantic labels.

3.1. Modality-Synchronized Augmentation

To mitigate modality misalignment and enhance diversity during data augmentation, we propose PieAug. The key idea is to extract a flexible number of cylindrical voxels along the height, angle, or radius axes – analogous to cutting a variable-sized pie slice from a cake – and swap it with a corresponding slice from another scene, as illustrated in Fig. 2. To maintain synchronization across modalities, each 3D “pie” is paired with its corresponding image patch, which is exchanged simultaneously.

For simplicity, we illustrate the process of finding the corresponding image patch for a single voxel from one camera view. This can be easily extended to a pie-shaped region with multiple image views. Given a voxel containing N_i points, denoted as $\mathbf{v}_i = \{\mathbf{p}_j\}_i^{N_i}$, where $\sum_{i=1}^M N_i = N$, we first project the LiDAR point \mathbf{p}_j from 3D coordinate system to its corresponding 2D coordinates in the image plane using the following transformation:

$$\pi(\mathbf{p}_j) = \mathbf{K} \times \mathbf{T} \times [\mathbf{p}_{j,1}, \mathbf{p}_{j,2}, \mathbf{p}_{j,3}, 1]^\top, \quad (1)$$

where $\mathbf{K} \in \mathbb{R}^{3 \times 3}$ is the camera intrinsic matrix and $\mathbf{T} \in \mathbb{R}^{4 \times 4}$ is the extrinsic transformation matrix. Next, we define \mathbf{g}_i as the bounding rectangle that encloses all projected points $\pi(\mathbf{p}_j)$. This ensures that each voxel corresponds to a specific region in the image, denoted as $\langle \mathbf{v}_i, \mathbf{g}_i \rangle$:

$$\mathbf{g}_i = \mathcal{B}(\{\pi(\mathbf{p}_j) \mid \mathbf{p}_j \in \mathbf{v}_i\}). \quad (2)$$

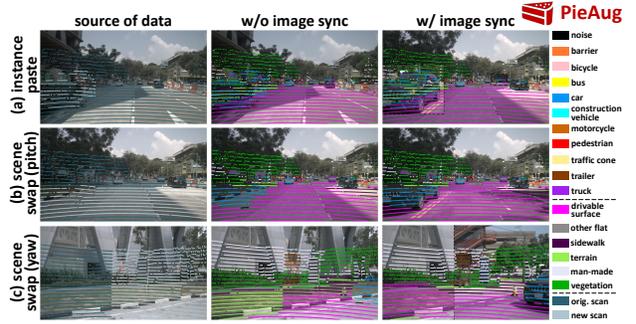


Figure 2. Motivation and implementation variants of PieAug. Each column illustrates the motivation for LiDAR-image synchronized augmentation. Each row displays a different *pie-cut* strategy. Point clouds are projected on camera images, with colors indicating semantic labels or data sources. Best viewed in color.

Here $\mathcal{B}(\cdot)$ is the operator that fits a bounding rectangle enclosing a set of pixels.

During augmentation, we determine the size and position of the *pie-cut* using a 3D binary mask $\mathbf{S} \in \{0, 1\}^{R \times \Theta \times Z}$, where R , Θ , and Z represent the binning resolution along the radial, angular, and height axes of the cylindrical voxelization, respectively. Each voxel is assigned a corresponding mask value $\mathbf{S}(r, \theta, z)$, where $\mathbf{S}(r, \theta, z) = 1$ indicates that the voxel is replaced with one from a new scan; otherwise, it remains unchanged from the original scene. Consequently, the augmented cylinder \mathbf{V}^{aug} is obtained as follows:

$$\mathbf{V}^{\text{aug}} = \mathbf{V}^{\text{org}} \otimes (1 - \mathbf{S}) + \mathbf{V}^{\text{new}} \otimes \mathbf{S}, \quad (3)$$

where \otimes denotes an element-wise masking operation, multiplying the voxel values by the corresponding mask. Since each voxel \mathbf{v}_i aligns with an image region \mathbf{g}_i , we apply image augmentation in parallel using the same mask \mathbf{S} .

Instance Pasting. We first illustrate how PieAug achieves instance-level augmentation (copy and paste) by selecting *pie-cut* voxels corresponding to s sampled instances from a new scene. We apply transformations such as translation, rotation, and scaling to each instance. Next, we identify the indices of voxels that overlap with the s transformed instances as \mathcal{C} . The mask \mathbf{S} is then constructed as:

$$\mathbf{S} = \bigcup_{r=1, \theta=1, z=1}^{R \times \Theta \times Z} \mathbb{1}[(r, \theta, z) \in \mathcal{C}]. \quad (4)$$

Scene Swapping. We then illustrate how PieAug achieves scene-level augmentation, including scene swapping, which involves dividing the voxels evenly along a chosen axis (height or angle) and swapping them alternately. We achieve this by selecting all voxels along the radial axis, as well as one of the height or angle axes, and then freely choosing a number of slices from the remaining axis. For example, the mask for selecting b slices from the angle axis is defined as:

$$\mathbf{S}(r, \theta, z) = \begin{cases} 1, & \text{if } \theta \in \mathcal{O} \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

Here, \mathcal{O} denotes the set of indices corresponding to the selected b angle slices. A similar masking strategy can be applied to height slices.

Remarks. With the instance-level and scene-level augmentation capabilities described above, PieAug generalizes most existing LiDAR-based augmentation techniques. For example: 1) Panoptic-PolarNet: perform instance-level augmentation by sampling instances based on their semantic labels and applying transformations using translation and XY-plane rotation. 2) PolarMix (instance branch): perform instance-level augmentation by selecting instances based on their labels and applying transformations by rotating duplicated instances multiple times along the Z-axis. 3) PolarMix (scene branch): perform scene-level augmentation by selecting half of the slices along the azimuth angle. 4) LaserMix: perform scene-level augmentation by selecting slices at different inclination angles. As a result, PieAug offers greater flexibility in combining instance-level and scene-level augmentations by performing multiple augmentation rounds with different voxel indices. Furthermore, due to the synchronized transformations applied to both modalities, PieAug ensures well-aligned LiDAR and image augmentations.

3.2. Geometric-Guided Token Fusion

Since cylindrical voxelization results in voxels of varying sizes (larger in regions farther from the central origin), this poses two key challenges for generating multi-modal tokens: 1) how to align image features with LiDAR features, and 2) how to effectively fuse. To address these issues, we propose **Geometric-guided Token Fusion (GTF)**, as illustrated in Fig. 3, which leverages the rich geometric information from LiDAR to guide alignment and enable effective fusion.

Specifically, we align features at the voxel level by projecting all physical points within a voxel \mathbf{v}_i onto the image plane and averaging their corresponding image features to create an aggregated representation:

$$\tilde{\mathbf{F}}_i^{2D} = \frac{1}{N_i} \sum_j^{N_i} \mathbf{F}^{2D}(\pi(\mathbf{p}_j)). \quad (6)$$

We refer to the paired voxel-wise LiDAR feature \mathbf{F}_i^{3D} and image feature $\tilde{\mathbf{F}}_i^{2D}$ as the *contents* of the i -th multi-modal token. Notably, aggregating image features by projecting all physical points within a voxel preserves feature validity. As shown in Fig. 3(a), using only the voxel centroid can lead to misalignment, as the projected location may not correspond accurately to the relevant image region.

Position embedding (PE) has proven effective in aligning features from different modalities (Yan et al., 2023). However, due to the varying sizes of cylindrical voxels, encoding only the voxel centroid is suboptimal, as a single point may not sufficiently represent the voxel’s perceptive field. Even when using physical points for PE, capturing the full perceptive field of a voxel or its corresponding image region

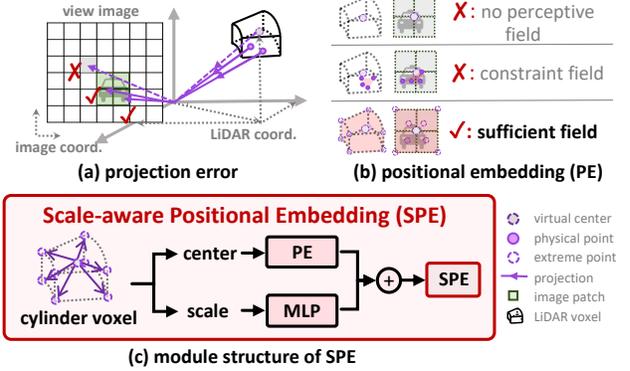


Figure 3. (a) and (b) illustrate two challenges in LiDAR–image fusion introduced by cylindrical voxelization. In (a), relying on a virtual voxel center can lead to projection errors, which become more pronounced for larger voxels and cause mismatches between LiDAR and image data. This motivates the use of physical points. In (b), canonical PE (row 1) overlooks the varying sizes of voxel and image features, while focusing solely on the region of physical points (row 2) limits the receptive field. Hence, we introduce a **scale-aware PE** that uses extreme points, as shown in (c).

remains inadequate, as illustrated by the red region in the second row of Fig. 3(b). This limitation arises because physical points provide only a partial representation of objects in both the LiDAR and image domains, especially at greater distances where LiDAR points are sparsely distributed. To address this issue, we propose a unified **Scale-aware Position Embedding (SPE)** for both LiDAR and image tokens. SPE ensures consistent perception of perceptive fields by incorporating shared scale embeddings to both modalities, as shown in the third row of Fig. 3(c).

Specifically, as illustrated in Fig. 3(c), we introduce an *extreme point set*, denoted as $\hat{\mathbf{v}}_i = \{\hat{\mathbf{p}}_j\}_{j=1}^8$, consisting of the eight corners of each cylindrical voxel to represent its scale. In the LiDAR space, these extreme points define the spatial partitioning of the scene. In the image space, they outline potential perceptive field regions, particularly in areas with sparse LiDAR points, allowing for a more adaptive perception range. In SPE, each cylindrical voxel is embedded with both its scale and centroid position:

$$\mathbf{s}_i = \psi(\text{Avg}(\hat{\mathbf{v}}_i)) + \phi(\|\hat{\mathbf{v}}_i - \text{Avg}(\hat{\mathbf{v}}_i)\|_2), \quad (7)$$

where $\text{Avg}(\cdot)$ denote the averaging method and $\|\cdot\|$ is L2 norm operation in Cartesian space. ψ represents a mixed-parameterized positional encoding function (Xiao et al., 2025), which embeds the centroid position in both Cartesian and polar spaces. ϕ is a multi-layer perceptron (MLP) that projects the scale feature into the same dimension as the centroid embedding. With SPE, the final fused multi-modal token $\mathbf{F}_i^{\text{fuse}}$ for i -th cylindrical voxel is obtained by:

$$\mathbf{F}_i^{\text{fuse}} = \text{Cat}[(\mathbf{F}_i^{3D} \oplus \mathbf{s}_i), (\tilde{\mathbf{F}}_i^{2D} \oplus \mathbf{s}_i)]. \quad (8)$$

3.3. Prior-Based Query Generation

Previous works, such as P3Former (Xiao et al., 2025), initialize queries by a set of learnable parameters. However,

Table 1. Preliminary study of positional embedding for objects of thing classes. We conduct the experiment on our LiDAR branch. “GT” denotes using the ground truth center position, while “Noise” denotes adding Gaussian noise with a kernel size of 3 to the GT center position. “th” and “st” is the thing and stuff classes.

Modality	GT	Noise	PQ	mIoU	PQ th	PQ st
LiDAR			77.0	75.9	77.8	75.7
LiDAR	✓		83.2	82.3	88.5	74.4
LiDAR	✓	✓	81.8	79.8	86.8	73.6

such queries tend to prioritize easier samples while neglecting more challenging ones. Additionally, preliminary results have validated that accurate positional embedding significantly enhances the model’s ability to locate objects. As demonstrated in Table 1, applying the ground truth center position for thing classes resulted in a 6.2% improvement in overall PQ and a 10.7% improvement in PQ for thing classes specifically. Even adding noise on the position, the improvement is still significant.

Inspired by this observation, we propose the **Prior-based Query Generation (PQG)** module to explicitly leverage texture features from the image domain, and geometric information from LiDAR domain as prior knowledge to generate well-informed initializations for instance queries. Specifically, we design three groups of queries: geometric-prior, texture-prior, and no-prior queries.

Geometric-Prior Query. Potential geometric-prior instances are those for which LiDAR features exhibit minimal degradation, allowing geometric characteristics to provide sufficient information for accurate positional hints. Compared to the rich texture features captured in images, LiDAR features offer more precise location predictions. Therefore, we generate location hints for geometric-prior queries by predicting a center heatmap and performing Non-Maximum Suppression (NMS) sampling according to confidence scores and range radius threshold. Specifically, we predict the class-agnostic heatmap in the polar Bird’s-Eye-View (BEV) space using a structure similar to (Yin et al., 2021; Bai et al., 2022). For each selected instance proposal (identified by its coordinates in the center heatmap), we lift it into 3D space by averaging all valid voxels across the height dimension.

Texture-Prior Query. For objects that are small or located far from the sensor, geometric information often becomes unreliable, making accurate location prediction difficult. To address this issue, we use texture information from images to discover potential texture-prior instances. First, we extract mask proposals using the pre-trained image segmentation models Grounding-DINO and SAM (Liu et al., 2023a; Kirillov et al., 2023). We then lift each 2D mask into a 3D frustum and collect all the 3D points that fall within it. To mitigate noise caused by overlapping along the depth dimension, we cluster these points into several groups using

the unsupervised DBSCAN (Ester et al., 1996) algorithm. Finally, the centroids of these clusters serve as location hints for texture-prior instances.

Given location hints from both modalities, we apply Farthest Point Sampling (FPS) to obtain a fixed number l_{pr} of location hints. It is worth noting that for large objects easily recognizable by both LiDAR and images, global sampling provides a holistic view and reduces redundant candidate proposals. Similar to the positional embedding used for multi-modal tokens, we apply SPE to embed both the query location and scale features. We then extract the query content by indexing into the corresponding voxel features \mathbf{F}^{fuse} , and finally add the SPE to this content to form the final query representation.

No-Prior Query. We hypothesize that instances without advanced priors exhibit a specific feature representation paradigm, allowing them to be recognized through a set of learnable parameters. We set the number of no-prior queries as l_{it} . This implicit paradigm learning enables the model to search within a smaller candidate pool and effectively identify potential instances.

All geometric-, texture-, and no-prior queries are concatenated and fed into the transformer decoder for “thing” class prediction, i.e., 3D instance segmentation. Semantic queries are initialized following the approach in P3Former, with auxiliary semantic supervision applied. These semantic queries are used to predict segmentation results. We follow the process in P3Former to combine the instance and semantic predictions into the final panoptic segmentation.

4. Experimental Results

4.1. Experimental Setting

Datasets. **nuScenes** (Caesar et al., 2020; Fong et al., 2022) is a large-scale, multi-modal dataset designed for autonomous driving, containing data from a 32-beam LiDAR, 5 radars, and 6 RGB cameras. It includes 40,157 frames of outdoor scenes, with 34,149 frames labeled for training and validation, and the remaining reserved for testing. The panoptic annotations cover 10 “thing” classes, 6 “stuff” classes, and 1 class for noisy labels. **SemanticKITTI** (Behley et al., 2019; 2021) is an outdoor dataset derived from KITTI Vision Benchmark (Geiger et al., 2012). It includes data from a 64-beam LiDAR sensor and two front-view cameras, including 8 “thing” classes and 11 “stuff” classes, comprising 19,130 frames for training, 4,071 frames for validation, and 20,351 frames for testing.

Evaluation Metrics. Consistent with the standard works (Kirillov et al., 2018; Zhang et al., 2023; Xiao et al., 2025), panoptic quality (PQ) is selected as the primary metric. PQ is defined as the product of segmentation quality

Table 2. Comparison of panoptic segmentation performance on the nuScenes validation set. Top results are shown in **bold**. “M.” indicates which modality (or modalities) each method uses. “P.-PCSCNet” is the LiDAR branch of “P.-FusionNet” (Song et al., 2024).

Method	M.	PQ	PQ [†]	RQ	SQ	PQ th	RQ th	SQ th	PQ st	RQ st	SQ st	mIoU
DS-Net (Hong et al., 2021)	L	42.5	51.0	50.3	83.6	32.5	38.3	83.1	59.2	70.3	84.4	70.7
EfficientLPS (Sirohi et al., 2021)	L	62.0	65.6	73.9	83.4	56.8	68.0	83.2	70.6	83.6	83.8	65.6
P.-PolarNet (Zhou et al., 2021)	L	67.7	71.0	78.1	86.0	65.2	74.0	87.2	71.9	84.9	83.9	69.3
P.-PHNet (Li et al., 2022a)	L	74.7	77.7	84.2	88.2	74.0	82.5	89.0	75.9	86.9	86.8	79.7
CFNet (Li et al., 2023c)	L	75.1	78.0	84.6	88.8	74.8	82.9	89.8	76.6	87.3	87.1	79.3
CenterLPS (Mei et al., 2023)	L	76.4	79.2	88.0	86.2	77.5	88.4	87.1	74.6	87.3	84.9	77.1
LCPS (Zhang et al., 2023)	L	72.9	77.6	82.0	88.4	72.8	80.5	90.1	73.0	84.5	85.5	75.1
P.-PCSCNet (Song et al., 2024)	L	72.7	75.4	84.8	86.4	71.2	82.9	86.6	75.1	84.2	84.2	69.8
P3Former (Xiao et al., 2025)	L	75.9	78.9	84.7	89.7	76.9	83.3	92.0	75.4	87.1	86.0	76.8
IAL (our LiDAR branch)	L	77.0	79.6	85.1	90.2	77.8	83.8	92.6	75.7	87.3	86.2	75.9
LCPS (Zhang et al., 2023)	L+C	79.8	84.0	88.5	89.8	82.3	89.6	91.7	75.6	86.5	86.7	80.5
P.-FusionNet (Song et al., 2024)	L+C	77.2	79.3	87.2	87.8	77.5	87.7	88.2	76.2	85.9	86.0	73.4
IAL (ours)	L+C	82.3	84.7	89.7	91.5	85.3	90.6	94.1	77.3	88.2	87.2	80.6

Table 3. Comparison on the nuScenes test set. Top and runner-up results are marked in **bold** and underline, respectively. “*” indicates the use of additional temporal frames and detection annotations. Our method is evaluated without test-time augmentation or ensembling.

Method	M.	PQ	PQ [†]	RQ	SQ	PQ th	RQ th	SQ th	PQ st	RQ st	SQ st	mIoU
EfficientLPS (Sirohi et al., 2021)	L	62.4	66.0	74.1	83.7	57.2	68.2	83.6	71.1	84.0	83.8	66.7
P.-PolarNet (Zhou et al., 2021)	L	63.6	67.1	75.1	84.3	59.0	69.8	84.3	71.3	83.9	84.2	67.0
P.-PHNet (Li et al., 2022a)	L	80.1	82.8	87.6	91.1	82.1	88.1	93.0	76.6	86.6	<u>87.9</u>	80.2
CPSeg (Li et al., 2023a)	L	73.2	76.3	82.7	88.1	72.9	81.3	89.2	74.0	85.0	<u>86.3</u>	73.7
MaskPLS (Marcuzzi et al., 2023)	L	61.1	64.3	68.5	86.8	54.3	58.8	87.8	72.4	84.5	85.1	74.8
LCPS (Zhang et al., 2023)	L	72.8	76.3	81.7	88.6	72.4	80.0	90.2	73.5	84.6	86.1	74.8
LidarMultiNet* (Ye et al., 2023)	L	81.4	84.0	88.9	91.3	83.9	89.9	93.1	77.3	87.1	88.2	82.2
IAL (our LiDAR branch)	L	75.1	77.7	83.0	90.1	75.0	80.9	92.4	75.2	86.5	86.4	73.3
4DFormer (Athar et al., 2023)	L+C	78.0	81.4	86.6	89.7	80.0	87.8	90.9	74.6	84.5	87.6	<u>80.4</u>
LCPS (Zhang et al., 2023)	L+C	79.5	82.3	87.7	90.3	81.7	88.6	92.2	75.9	86.3	87.3	78.9
IAL (ours)	L+C	82.0	84.3	89.3	91.6	84.8	90.2	93.8	77.5	87.8	87.8	79.9

(SQ) and recognition quality (RQ):

$$PQ = \underbrace{\frac{\sum_{TP} IoU}{|TP|}}_{SQ} \times \underbrace{\frac{|TP|}{|TP| + \frac{1}{2}|FP| + \frac{1}{2}|FN|}}_{RQ}, \quad (9)$$

where IoU denotes the Intersection over Union, TP denotes True Positives and so as for others. These metrics can be further extended to “thing” and “stuff” classes, denoted as PQth, PQst, RQth, RQst, SQth, and SQst. We also report PQ[†] (Porzi et al., 2019), which replaces PQ with mIoU for stuff classes.

Implementation Details. We follow standard practice (Zhou et al., 2021; Li et al., 2022a; Xiao et al., 2025) to represent point clouds by discretizing the 3D space into cylindrical voxels of size $[480 \times 360 \times 32]$. The LiDAR branch is built upon the architecture of P3former. For nuScenes, the polar coordinate range is defined as $[-50m, 50m] \times [0, 2\pi] \times [-5m, 3m]$, while for SemanticKITTI, the height range is adjusted to $[-4m, 2m]$. All images are resized to 640×360 . For augmentation, we employ the following strategies: instance pasting and scene-swapping (split the scene along the height and angle axes, with the number of splits randomly chosen from $[3, 4, 5]$ each time). We set the ratio of application for these three augmentation

strategies to 0.4:0.05:0.05, respectively. Additionally, we perform basic transformations including random rotation, flipping, and scaling. We set the number of prior-based and no-prior instance queries to $l_{pr} = l_{lt} = 128$. We use AdamW (Kingma & Ba, 2014) as the optimizer, with a default weight decay of 0.01. The entire model is trained from scratch with a batch size of 2, using 4 NVIDIA A40 GPUs. The training spans 80 epochs for nuScenes and 36 epochs for SemanticKITTI. The initial learning rate is set to 0.0008 and decays by half at epochs $[60, 75]$ for nuScenes and $[30, 32]$ for SemanticKITTI, respectively. All model results listed in the following sections are NOT employed with any test-time augmentation (TTA) method.

4.2. Benchmark Results

nuScenes. We present comprehensive comparison results for LiDAR panoptic segmentation performance on the nuScenes validation and test sets, as shown in Table 2 and Table 3. Due to the limited number of multi-modal methods, currently only LCPS (Zhang et al., 2023) and Panoptic-FusionNet (Song et al., 2024), we also include LiDAR-only methods for comparison. Notably, our method IAL achieves the best performance across all metrics on the validation set and ranks first or second on most metrics

Table 4. Comparison of panoptic segmentation performance on the SemanticKITTI validation set. Top results are shown in **bold**.

Method	M.	PQ	PQ [†]	RQ	SQ	mIoU
P.-PolarNet	L	59.1	64.1	70.2	78.3	64.5
DS-Net	L	57.7	63.4	68.0	77.6	63.5
EfficientLPS	L	59.2	65.1	69.8	75.0	64.9
P.-PHNet	L	61.7	–	–	–	65.7
CenterLPS	L	62.1	67.0	72.0	80.7	–
LCPS	L	55.7	65.2	65.8	74.0	61.1
P3Former	L	62.6	66.2	72.4	76.2	–
IAL (LiDAR)	L	62.0	65.1	71.9	76.0	64.9
LCPS	L+C	59.0	68.8	68.9	79.8	63.2
IAL (ours)	L+C	63.1	66.3	72.9	81.4	66.0

in the test set. Specifically, IAL outperforms LCPS and Panoptic-FusionNet by a significant margin of 2.5% and 5.1% in PQ on the validation set, as shown in Table 2. This improvement is attributed to superior performance in both recognition (surpassing the two previous works by 1.2% and 2.5% in RQ, respectively) and segmentation (surpassing them by 1.7% and 3.7% in SQ). Furthermore, our model demonstrates superior performance on both “thing” and “stuff” classes, achieving a 7.8% and 1.1% improvement in metrics compared to the latest work, Panoptic-FusionNet. Compared to the LiDAR-only baseline (using the same augmentation strategies as P3Former adopts), IAL achieves a 5.3% improvement, primarily due to a 7.5% increase from thing classes, demonstrating the effectiveness of image assistance in detecting and recognizing objects. In Table 3, IAL also demonstrates superior performance, achieving the highest scores across most metrics on the nuScenes leaderboard. These outstanding results highlight the effectiveness of our modules for modality alignment and compensation.

SemanticKITTI presents a significant challenge due to its use of only two front-view cameras, limiting the availability of image features to support LiDAR. As shown in Table 4, despite these constraints, our IAL achieves a 4.1% improvement in PQ over the state-of-the-art multi-modal baseline LCPS, demonstrating the robustness of our method even under limited image supervision.

4.3. Ablation Studies

To validate the effectiveness of our proposed components, we conduct comprehensive ablation studies on the overall proposal framework in Table 5 and provide detailed analyses for each individual module in Table 6. All experiments are conducted on the nuScenes validation set using the same hyper-parameters for fair comparison.

As shown in Table 5, compared to the baseline that uses only basic point cloud transformations (row 1), PieAug improves PQ by 2.7%, benefiting from better input alignment and enriched scene context. Building on this, GTF further boosts PQ by 2.7% and RQ by 2.1%, demonstrating that

Table 5. Ablation study of the proposed modules in our framework. “Pie” denotes the PieAug module.

Pie	GTF	PQG	PQ	PQ [†]	RQ	SQ	mIoU
			75.7	78.1	84.4	88.3	73.8
✓			78.4	81.0	86.9	90.0	78.2
✓	✓		81.1	83.5	89.0	90.9	80.2
✓	✓	✓	82.3	84.7	89.7	91.5	80.6

Table 6. Ablation study of PQG module. “Geo.”, “Tex.”, and “NP.” represent geometric prior, texture prior, and no-prior queries, respectively. We set the total number of queries to 256 for a fair comparison. In configurations combining prior (geometric or texture) and no-prior queries, 128 queries are allocated to each set.

Geo.	Tex.	NP.	PQ	PQ th	PQ st	mIoU
		✓	81.2	83.8	76.8	79.8
✓		✓	81.3	83.9	77.0	80.0
	✓	✓	81.1	83.4	77.2	80.0
✓	✓		80.7	83.0	77.0	80.0
✓	✓	✓	82.3	85.3	77.3	80.6

Table 7. Comparison of augmentation strategies. “Img” indicates whether image-synchronized augmentation is applied.

Method	Img	PQ	PQ [†]	RQ	SQ	mIoU
PolarMix		80.3	82.8	87.9	91.0	78.6
LaserMix		80.6	83.0	88.5	90.8	79.3
PieAug (ours)		81.4	83.7	89.1	91.0	80.1
PieAug (ours)	✓	82.3	84.7	89.7	91.5	80.6

unified scale embedding and accurate projection enhance multi-modal representations. Finally, incorporating the PQG module brings an additional 1.2% gain in PQ, validating our hypothesis that initializing queries with modality priors leads to more precise object predictions than using purely learnable parameters.

The effectiveness of the PQG module is validated in Table 6. Rows 1–3 show comparable performance, suggesting that purely learnable queries tend to overfit easy or redundant samples, even when uni-modal priors are available. Row 4 shows a slight drop, likely due to an excessive number of prior-based queries exceeding the number of ground-truth instances, resulting in more false positives. In contrast, our design (row 5) assigns strong geometric and texture priors to confident regions, while reserving learnable queries for harder, low-prior cases. This balanced allocation improves the model’s ability to handle both easy and difficult samples, leading to superior overall performance.

4.4. Augmentation Methods Comparison

As shown in Table 7, we compare PieAug with LiDAR-only augmentation methods, including PolarMix (instance pasting and scene mixing) and LaserMix (inclination angle splitting) by 2.0% and 1.7% in PQ. Even with LiDAR-only augmentation, PieAug achieves superior performance, demonstrating its effectiveness as a generalized framework.

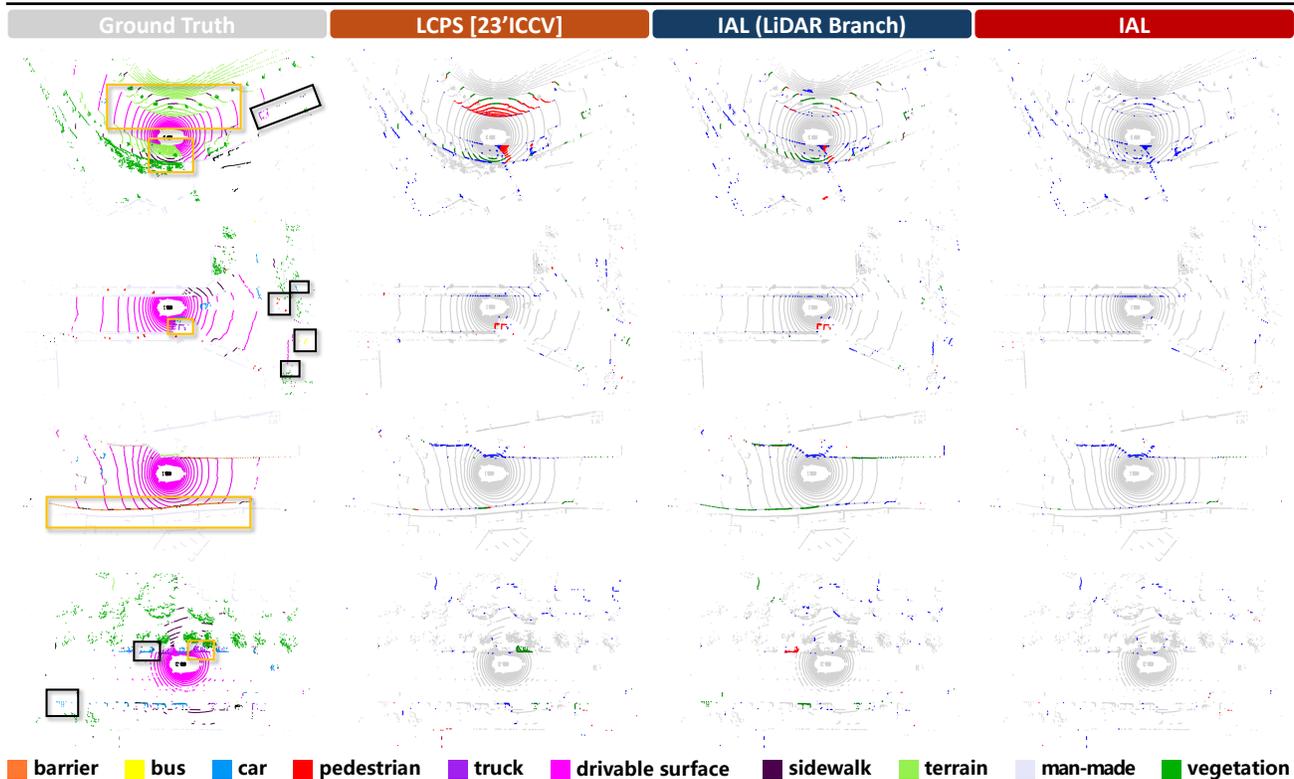


Figure 4. Qualitative comparison of our method with the preliminary multi-modal panoptic segmentation baseline, LCPS. To highlight the differences, we mark **false positive** and **false negative** predictions, which affect recognition quality, as well as **well-matched** and **mismatch points** in true positive predictions, which impact segmentation quality. GT is colored by semantic label. Best viewed in color.

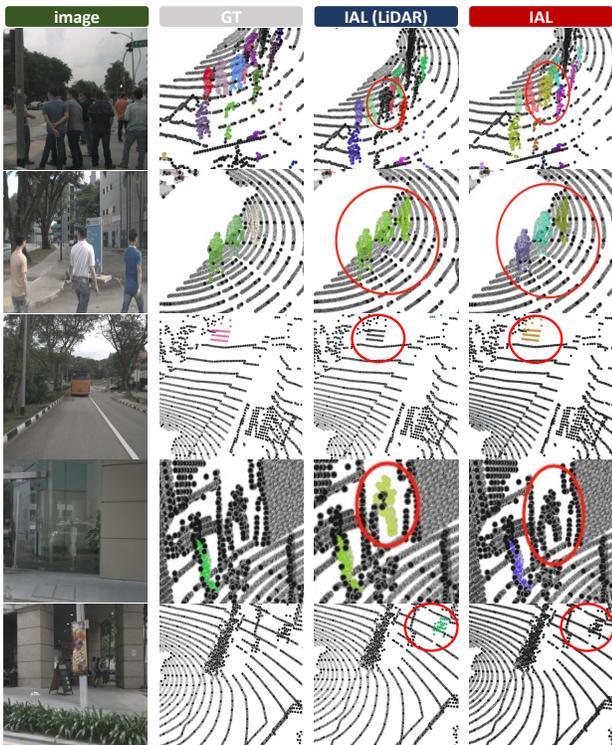


Figure 5. Visualization of instance predictions. Red circles highlight instances where the LiDAR branch fails to segment correctly, but our multi-modal method succeeds. Best viewed in color.

4.5. Qualitative Results and Discussion

We present qualitative evaluations on nuScenes validation set. As illustrated in the error maps in Fig. 4, our method notably reduces false positives (red points) and false negatives (green points) compared to LCPS. Furthermore, IAL outperforms its LiDAR branch in detecting remote objects (highlighted in the black boxes) and recognizing ambiguous classes (in yellow boxes), leveraging the assistance of image data. In Fig. 5, we compare instance prediction with GT, LiDAR branch, and IAL alongside the corresponding images. IAL showcases significant performance improvements in: (1) distinguishing multiple objects when they are clustered together (rows 1 and 2); (2) detecting distant objects (row 3); (3) recognizing false positive objects (rows 4 and 5).

5. Conclusion

This paper proposes IAL, a multi-modal 3D panoptic segmentation framework that harmonizes LiDAR and images through PieAug (synchronized augmentation), GTF (geometry-guided fusion), and PQG (prior-based queries). IAL directly predicts panoptic results via a transformer decoder, eliminating post-processing and achieving state-of-the-art performance on nuScenes (82.3% PQ) and SemanticKITTI (63.1% PQ). Texture-prior queries enhance small/distant object recognition, while geometric-prior queries improve large/nearby instance localization.

Impact Statement

This paper aims to enhance multi-modal 3D panoptic segmentation. There are minor potential societal consequences of our work, none of which we feel must be specifically highlighted here.

References

- An, Z., Sun, G., Liu, Y., Li, R., Wu, M., Cheng, M.-M., Konukoglu, E., and Belongie, S. Multimodality helps few-shot 3d point cloud semantic segmentation, 2025. URL <https://arxiv.org/abs/2410.22489>.
- Athar, A., Li, E., Casas, S., and Urtasun, R. 4d-former: Multimodal 4d panoptic segmentation. In Tan, J., Toussaint, M., and Darvish, K. (eds.), *Proceedings of The 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pp. 2151–2164. PMLR, 06–09 Nov 2023. URL <https://proceedings.mlr.press/v229/athar23a.html>.
- Bai, X., Hu, Z., Zhu, X., Huang, Q., Chen, Y., Fu, H., and Tai, C.-L. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1090–1099, 2022.
- Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., and Gall, J. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In *Proc. of the IEEE/CVF International Conf. on Computer Vision (ICCV)*, 2019.
- Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Gall, J., and Stachniss, C. Towards 3D LiDAR-based semantic scene understanding of 3D point cloud sequences: The SemanticKITTI Dataset. *The International Journal on Robotics Research*, 40(8-9):959–967, 2021. doi: 10.1177/02783649211006735.
- Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., and Beijbom, O. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, 2020.
- Cao, A.-Q., Dai, A., and de Charette, R. Pasco: Urban 3d panoptic scene completion with uncertainty awareness. In *CVPR*, 2024.
- Chen, X., Zhang, T., Wang, Y., Wang, Y., and Zhao, H. Futr3d: A unified sensor fusion framework for 3d detection. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 172–181, 2023.
- Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD’96, pp. 226–231. AAAI Press, 1996.
- Fong, W. K., Mohan, R., Hurtado, J. V., Zhou, L., Caesar, H., Beijbom, O., and Valada, A. Panoptic nuscenes: A large-scale benchmark for lidar panoptic segmentation and tracking. In *ICRA*, 2022.
- Geiger, A., Lenz, P., and Urtasun, R. Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 3354–3361, 2012.
- Gu, Y., Huang, Y., Xu, C., and Kong, H. Maskrange: A mask-classification model for range-view based lidar segmentation. *ArXiv*, abs/2206.12073, 2022. URL <https://api.semanticscholar.org/CorpusID:250048842>.
- Hong, F., Zhou, H., Zhu, X., Li, H., and Liu, Z. LiDAR-based Panoptic Segmentation via Dynamic Shifting Network. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13085–13094, Nashville, TN, USA, June 2021. IEEE. ISBN 978-1-66544-509-2. doi: 10.1109/CVPR46437.2021.01289. URL <https://ieeexplore.ieee.org/document/9577774/>.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. URL <https://api.semanticscholar.org/CorpusID:6628106>.
- Kirillov, A., He, K., Girshick, R. B., Rother, C., and Dollár, P. Panoptic segmentation. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9396–9405, 2018. URL <https://api.semanticscholar.org/CorpusID:4853375>.
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., Dollár, P., and Girshick, R. Segment anything. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 3992–4003, 2023. doi: 10.1109/ICCV51070.2023.00371.
- Kong, L., Liu, Y., Chen, R., Ma, Y., Zhu, X., Li, Y., Hou, Y., Qiao, Y., and Liu, Z. Rethinking range view representation for lidar segmentation. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 228–240, 2023a. doi: 10.1109/ICCV51070.2023.00028.
- Kong, L., Ren, J., Pan, L., and Liu, Z. Lasermix for semi-supervised lidar semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21705–21715, 2023b.

- Kong, L., Xu, X., Ren, J., Zhang, W., Pan, L., Chen, K., Ooi, W. T., and Liu, Z. Multi-modal data-efficient 3d scene understanding for autonomous driving, 2024. URL <https://arxiv.org/abs/2405.05258>.
- Krispel, G., Opitz, M., Waltner, G., Possegger, H., and Bischof, H. Fuseseg: Lidar point cloud segmentation fusing multi-modal data. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 1874–1883, 2020.
- Li, E., Razani, R., Xu, Y., and Liu, B. Cpseg: Cluster-free panoptic segmentation of 3d lidar point clouds. pp. 8239–8245, 2023a. URL <https://api.semanticscholar.org/CorpusID:240419848>.
- Li, J., He, X., Wen, Y., Gao, Y., Cheng, X., and Zhang, D. Panoptic-phnet: Towards real-time and high-precision lidar panoptic segmentation via clustering pseudo heatmap. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11799–11808, 2022a. URL <https://api.semanticscholar.org/CorpusID:248811224>.
- Li, J., Dai, H., Han, H., and Ding, Y. Mseg3d: Multi-modal 3d semantic segmentation for autonomous driving. In *CVPR*, pp. 21694–21704, 2023b.
- Li, L. and Zhao, N. End-to-end semi-supervised 3d instance segmentation with pteacher. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5352–5358. IEEE, 2024.
- Li, X. and Chen, D. A survey on deep learning-based panoptic segmentation. *Digital Signal Processing*, 120:103283, 2022. ISSN 1051-2004. doi: <https://doi.org/10.1016/j.dsp.2021.103283>. URL <https://www.sciencedirect.com/science/article/pii/S1051200421003225>.
- Li, X., Zhang, G., Wang, B., Hu, Y., and Yin, B. Center focusing network for real-time lidar panoptic segmentation. pp. 13425–13434, 06 2023c. doi: 10.1109/CVPR52729.2023.01290.
- Li, Y., Yu, A. W., Meng, T., Caine, B., Ngiam, J., Peng, D., Shen, J., Lu, Y., Zhou, D., Le, Q. V., et al. Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 17182–17191, 2022b.
- Li, Y., Fan, L., Liu, Y., Huang, Z., Chen, Y., Wang, N., and Zhang, Z. Fully sparse fusion for 3d object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- Liang, T., Xie, H., Yu, K., Xia, Z., Lin, Z., Wang, Y., Tang, T., Wang, B., and Tang, Z. Bevfusion: A simple and robust lidar-camera fusion framework. volume 35, pp. 10421–10434, 2022.
- Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Li, C., Yang, J., Su, H., Zhu, J., et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023a.
- Liu, Y., Runnan, C., Li, X., Kong, L., Yang, Y., Xia, Z., Bai, Y., Zhu, X., Ma, B., Li, Y., Qiao, Y., and Hou, Y. Uniseg: A unified multi-modal lidar segmentation network and the openpcseg codebase. pp. 21605–21616, 10 2023b. doi: 10.1109/ICCV51070.2023.01980.
- Man, Y., Gui, L.-Y., and Wang, Y.-X. BEV-Guided Multi-Modality Fusion for Driving Perception. In *CVPR*, 2023.
- Marcuzzi, R., Nunes, L., Wiesmann, L., Behley, J., and Stachniss, C. Mask-Based Panoptic LiDAR Segmentation for Autonomous Driving. 8(2):1141–1148, 2023. doi: 10.1109/LRA.2023.3236568. URL <https://www.ipb.uni-bonn.de/wp-content/papercite-data/pdf/marcuzzi2023ral.pdf>.
- Mei, J., Yang, Y., Wang, M., Li, Z., Hou, X., Ra, J., Li, L., and Liu, Y. CenterLPS: Segment Instances by Centers for LiDAR Panoptic Segmentation. In *Proceedings of the 31st ACM International Conference on Multimedia*, MM ’23, pp. 1884–1894, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701085. doi: 10.1145/3581783.3612080. URL <https://dl.acm.org/doi/10.1145/3581783.3612080>.
- Oršić, M. and Šegvić, S. Efficient semantic segmentation with pyramidal fusion. *Pattern Recognition*, 110:107611, 2021. ISSN 0031-3203. doi: <https://doi.org/10.1016/j.patcog.2020.107611>. URL <https://www.sciencedirect.com/science/article/pii/S0031320320304143>.
- Porzi, L., Bulò, S. R., Colovic, A., and Kotschieder, P. Seamless scene segmentation. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8269–8278, 2019. URL <https://api.semanticscholar.org/CorpusID:145047913>.
- Razani, R., Cheng, R., Li, E., Taghavi, E. M., Ren, Y., and Liu, B. Gp-s3net: Graph-based panoptic sparse semantic segmentation network. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 16056–16065, 2021. URL <https://api.semanticscholar.org/CorpusID:237213370>.

- Sirohi, K., Mohan, R., Buscher, D., Burgard, W., and Valada, A. Efficientlps: Efficient lidar panoptic segmentation. *IEEE Transactions on Robotics*, 38:1894–1914, 2021. URL <https://api.semanticscholar.org/CorpusID:231933783>.
- Song, H., Cho, J., Ha, J., Park, J., and Jo, K. Panoptic-fusionnet: Camera-lidar fusion-based point cloud panoptic segmentation for autonomous driving. *Expert Syst. Appl.*, 251(C), October 2024. ISSN 0957-4174. doi: 10.1016/j.eswa.2024.123950. URL <https://doi.org/10.1016/j.eswa.2024.123950>.
- Su, S., Xu, J., Wang, H., Miao, Z., Zhan, X., Hao, D., and Li, X. PUPS: Point Cloud Unified Panoptic Segmentation, February 2023. URL <http://arxiv.org/abs/2302.06185>. arXiv:2302.06185 [cs].
- Wu, Y., Liu, J., Gong, M., Miao, Q., Ma, W., and Xu, C. Joint semantic segmentation using representations of lidar point clouds and camera images. *Information Fusion*, 108:102370, 2024.
- Xiao, A., Huang, J., Guan, D., Cui, K., Lu, S., and Shao, L. Polarmix: a general data augmentation technique for lidar point clouds. 2022.
- Xiao, Z., Zhang, W., Wang, T., Loy, C. C., Lin, D., and Pang, J. Position-Guided Point Cloud Panoptic Segmentation Transformer. *International Journal of Computer Vision*, 133(1):275–290, January 2025. ISSN 1573-1405. doi: 10.1007/s11263-024-02162-z. URL <https://doi.org/10.1007/s11263-024-02162-z>.
- Xu, S., Wan, R., Ye, M., Zou, X., and Cao, T. Sparse cross-scale attention network for efficient lidar panoptic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 2920–2928, 2022.
- Xu, Y., Fazlali, H., Ren, Y., and Liu, B. Aop-net: All-in-one perception network for lidar-based joint 3d object detection and panoptic segmentation. In *2023 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1–7, 2023a. doi: 10.1109/IV55152.2023.10186743.
- Xu, Y., Hu, C., Zhao, N., and Lee, G. H. Generalized few-shot point cloud segmentation via geometric words. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 21506–21515, 2023b.
- Yan, J., Liu, Y., Sun, J., Jia, F., Li, S., Wang, T., and Zhang, X. Cross modal transformer: Towards fast and robust 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 18268–18278, 2023.
- Yan, X., Gao, J., Zheng, C., Zheng, C., Zhang, R., Cui, S., and Li, Z. 2dpass: 2d priors assisted semantic segmentation on lidar point clouds. In *European Conference on Computer Vision*, pp. 677–695. Springer, 2022.
- Yang, J., Cen, J., Peng, W., Liu, Shuai and Hong, F., Li, X., Zhou, K., Chen, Q., and Liu, Z. 4d panoptic scene graph generation. In *NeurIPS*, 2023.
- Yang, Y., Mei, J., Du, S., Xiao, Y., Wu, H., Xu, X., and Liu, Y. Dqformer: Toward unified lidar panoptic segmentation with decoupled queries for large-scale outdoor scenes. *IEEE Transactions on Geoscience and Remote Sensing*, 63:1–15, 2025. doi: 10.1109/TGRS.2025.3558951.
- Ye, D., Zhou, Z., Chen, W., Xie, Y., Wang, Y., Wang, P., and Foroosh, H. Lidarmultinet: towards a unified multi-task network for lidar perception. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence, AAAI’23/IAAI’23/EAAI’23*. AAAI Press, 2023. ISBN 978-1-57735-880-0. doi: 10.1609/aaai.v37i3.25429. URL <https://doi.org/10.1609/aaai.v37i3.25429>.
- Yin, J., Shen, J., Chen, R., Li, W., Yang, R., Frossard, P., and Wang, W. Is-fusion: Instance-scene collaborative fusion for multimodal 3d object detection. In *CVPR*, 2024.
- Yin, T., Zhou, X., and Krähenbühl, P. Center-based 3d object detection and tracking. *CVPR*, 2021.
- Zhang, H., Liang, L., Zeng, P., Song, X., and Wang, Z. *SparseLIF: High-Performance Sparse LiDAR-Camera Fusion for 3D Object Detection*, pp. 109–128. 09 2024. ISBN 978-3-031-72760-3. doi: 10.1007/978-3-031-72761-0_7.
- Zhang, Z., Zhang, Z., Yu, Q., Yi, R., Xie, Y., and Ma, L. Lidar-camera panoptic segmentation via geometry-consistent and semantic-aware alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 3662–3671, October 2023.
- Zhao, H., Zhang, J., Chen, Z., Zhao, S., and Tao, D. Unimix: Towards domain adaptive and generalizable lidar semantic segmentation in adverse weather. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14781–14791, June 2024.
- Zhao, N., Chua, T.-S., and Lee, G. H. Few-shot 3d point cloud semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8873–8882, 2021.

Zhao, Y., Zhao, N., and Lee, G. H. Synthetic-to-real domain generalized semantic segmentation for 3d indoor point clouds. *arXiv preprint arXiv:2212.04668*, 2022.

Zhou, Z., Zhang, Y., and Foroosh, H. Panoptic-PolarNet: Proposal-free LiDAR Point Cloud Panoptic Segmentation . In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13189–13198, Los Alamitos, CA, USA, June 2021. IEEE Computer Society. doi: 10.1109/CVPR46437.2021.01299. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR46437.2021.01299>.

Zhu, X., Zhou, H., Wang, T., Hong, F., Ma, Y., Li, W., Li, H., and Lin, D. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. *arXiv preprint arXiv:2011.10033*, 2020.

Zhuang, Z., Li, R., Jia, K., Wang, Q., Li, Y., and Tan, M. Perception-aware multi-sensor fusion for 3d lidar semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 16280–16290, 2021.

Supplementary Materials

The supplementary materials are organized as follows: Sec. A extends the ablation study with token fusion stage; Sec. B provides additional qualitative results to visualize the improvements brought by each module; Sec. C analyzes the efficiency of IAL; Sec. D investigates how image inputs enhance LiDAR under perturbations from lighting and weather conditions; Sec. E discusses the potential broader impact of IAL; and Sec. F outlines current limitations and future work.

A. Ablation Study of Token Fusion

To analyze the effectiveness of the Geometric-guided Token Fusion (GTF) module, we divide GTF into two components: Token Selection (Sel) and Token Positional Embedding (PE). The full version of GTF uses all physical points within a cylindrical voxel (denoted as “set”) for token selection and embeds the scale between all extreme points (“scl”) to indicate perception regions. Alternative designs degrade token selection to a virtual center (“ctr”) and positional embedding to the center or extreme points (“ext”) of the voxel. All results are evaluated using the same experimental setting as the ablation studies in the main manuscript.

Table 8 reveals that using a point set rather than the virtual center for every token to construct the image feature contributes up to a 0.9% increase in PQ performance, and an up to 0.7% improvement in RQ. This verifies that precise LiDAR-image projection helps LiDAR voxels find corresponding image patches, and image features assist in recognition. For positional embedding, using scale-aware embedding indicates the potential perception regions of both LiDAR and image tokens, achieving the highest performance. This advanced improvement diminishes when scaling is degraded to using extreme points or solely the center of the voxel.

Table 8. Ablation study of the GTF module. “Sel” and “PE” denote the designs for token selection and positional embedding, respectively. We evaluate different configurations for component ablation: “ctr” represents the voxel virtual center, “set” refers to physical points within a voxel, and “-”, “ctr”, “ext”, and “scl” indicate not implementing PE, the use of center, extreme, or scale embeddings.

Sel	PE	PQ	PQ [†]	RQ	SQ	mIoU
ctr	-	78.4	81.0	86.9	90.0	78.2
ctr	ctr	79.7	82.2	87.8	90.5	78.3
ctr	ext	79.7	82.3	87.7	90.6	78.5
ctr	scl	80.4	82.7	88.3	90.7	78.4
set	-	79.3	81.7	87.5	90.3	77.7
set	ctr	80.6	83.0	88.4	90.8	80.1
set	ext	80.0	82.4	88.0	90.6	77.9
set	scl	81.1	83.5	89.0	90.9	80.2

B. Qualitative Results for Modular Performance

We present qualitative evaluations of the ablation studies for the GTF and PQG modules on the nuScenes validation set, as illustrated in the error maps in Fig. 6. With the assistance of the GTF module (comparing column 3 with column 2), our model demonstrates improved performance in distinguishing ambiguous objects, such as recognizing barriers and pedestrians from the background (highlighted in yellow boxes in rows 2 and 5), as well as differentiating buses and trucks from cars (rows 3-5). These classes often share similar geometric appearances, especially when point clouds are sparse. However, images provide rich texture features that help distinguish each class, even in limited regions. The GTF module enhances LiDAR voxel data by embedding it with more accurate image features, allowing for better receptive field estimation for each cylindrical voxel and facilitating a strong alignment of LiDAR and image features. The PQG module further enhances object perception, as shown by the comparison between column 4 and column 3, especially for small-scale and remote objects highlighted in black boxes. Even for small objects like pedestrians, bicycles, cars, and trailers, which have few points, the PQG module succeeds in accurate detection, demonstrating its superior performance by utilizing both geometric and texture priors, as well as the learnable capability of no-prior queries.

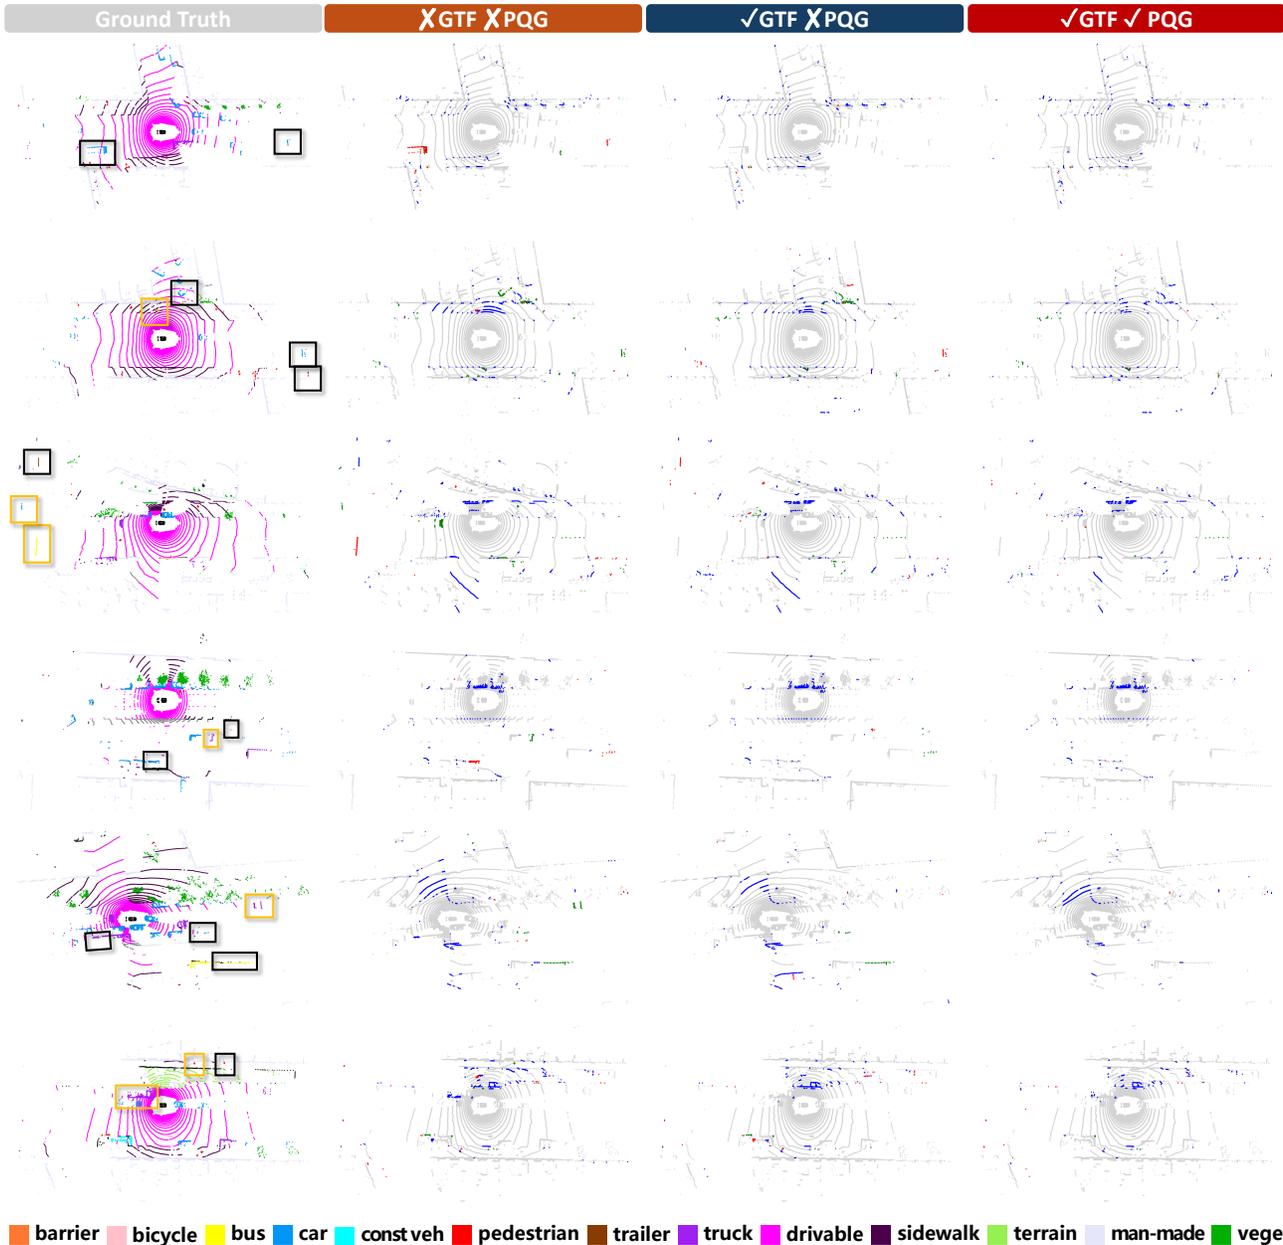


Figure 6. Qualitative comparison of the ablation study for GTF and PQG modules. To emphasize the differences, we mark **false positive** and **false negative** predictions, which affect recognition quality, as well as **well-matched** and **mismatch points** in true positive predictions, which impact segmentation quality. GT is colored by semantic label. Best viewed in color.

C. Time and Memory Cost

We compare inference speed, model size, and performance between our method and the main baseline, LCPS. We also report a lightweight variant (denoted by *), which excludes the 2D mask pre-processing step (Grounding-DINO and SAM) to highlight the efficiency of our framework’s major components. All latency measurements are conducted on a single NVIDIA A40 GPU with batch size 1. For a fair comparison, we measure LCPS latency using its official codebase on our hardware. As shown in Table 9, our method achieves over 2× faster inference and a +2.5% gain in PQ compared to LCPS. Even when including the mask generation time, our approach remains comparable in speed.

Table 9. Comparison of models in terms of inference speed (FPS), model size (#Params), and Panoptic Quality (PQ). * denotes the result of core components of our model. All latency measurements are conducted on the same device.

Model	FPS	#Params (M)	PQ
LCPS	1.7	77.7	79.8
IAL	4.0*	81.8*	82.3
	0.9	859.9	82.3

Table 10. Performance on the full nuScenes validation set and its nighttime/rain subsets. Best results are highlighted in **bold**.

Model	Full Val Set	Night Split	Rain Split
# of scan	6,019	602	1,088
LCPS	79.8	64.3	76.8
IAL (LiDAR branch)	77.0	63.2	73.1
IAL (full model)	82.3	70.5	81.2

D. Image Assists LiDAR Under Adverse Conditions

We evaluate the performance of our IAL on the **nighttime** and **rain** splits of the nuScenes val set. In the nighttime scenario, image quality is significantly degraded; in the rain scenario, both LiDAR and image encounter perturbations. As shown in Table 10, IAL (row 4) outperforms LCPS (row 2) not only on the full set but also under each adverse condition. We further ablate cross-modal interaction by comparing the full IAL to its LiDAR-only branch (row 3). Even under the degraded nighttime and rain splits, the full model gains +7.3% and +8.1% on PQ, respectively, confirming the image’s effective assistive role for LiDAR. This improvement can be attributed to two main factors: 1. Modality-synchronized augmentation (PieAug), which exposes the model to more diverse samples, including nighttime and rain scenarios, by mixing synchronized LiDAR and image data. This allows the model to generalize better to rare conditions like nighttime scenes. 2. The combination of three types of queries in our PQG module, where no-prior queries complement the texture-prior and geometric-prior queries, helping the model to effectively identify potential instances. Additionally, pre-trained Grounding-DINO and SAM further stabilise 2D mask generation under distribution shifts thanks to their large-scale training.

E. Potential Broader Impact of This Work

Based on the details in the paper, the broader impacts of this work can be highlighted as follows:

- The proposed multi-modal 3D panoptic segmentation framework (IAL) advances the field of autonomous driving by improving object detection and segmentation through the integration of LiDAR and image data. This technology has direct implications for the safety, accuracy, and efficiency of autonomous vehicles, particularly in complex, real-world environments. By addressing challenges such as the sparsity of LiDAR data and the difficulty of recognizing small or distant objects, this work enhances perception systems, enabling more reliable decision-making in autonomous driving.
- Furthermore, the advancements in modality-synchronized augmentation (PieAug) and geometric-guided token fusion (GTF) represent significant contributions to the broader field of sensor fusion in robotics and autonomous systems. These innovations could be adapted for other applications requiring high-precision environmental understanding, such as robotics, agriculture, and urban planning.
- As with all AI technologies, ethical considerations should be taken into account, particularly concerning privacy, security, and the potential for job displacement in industries such as transportation and logistics. However, the broader societal impact is generally positive, particularly in terms of enhancing public safety and reducing the risks associated with human error in driving.

F. Limitations

While our work achieves strong results across key metrics, it is important to acknowledge certain limitations. Specifically, the sampling method in our Query Initialization (PQG) module relies on a relatively simple and generic approach. While our work demonstrates strong performance across key metrics, it is important to note that the extraction of texture-prior queries relies on generic, large-scale pre-trained models rather than methods specifically designed for this task or benchmark. Although this approach ensures broad applicability, it may not fully leverage task-specific characteristics that could further enhance performance. Nevertheless, through our carefully designed PQG module, we still achieve competitive results. In future work, we plan to explore more specialized sampling methods tailored to the task, which could further improve the quality of texture-prior queries and overall performance.