

HD²-SSC: High-Dimension High-Density Semantic Scene Completion for Autonomous Driving

Zhiwen Yang, Yuxin Peng*

Wangxuan Institute of Computer Technology, Peking University
{yangzhiwen, pengyuxin}@pku.edu.cn

Abstract

Camera-based 3D semantic scene completion (SSC) plays a crucial role in autonomous driving, enabling voxelized 3D scene understanding for effective scene perception and decision-making. Existing SSC methods have shown efficacy in improving 3D scene representations, but suffer from the inherent input-output *dimension gap* and annotation-reality *density gap*, where the 2D planar view from input images with sparse annotated labels leads to inferior prediction of real-world dense occupancy with a 3D stereoscopic view. In light of this, we propose the corresponding **High-Dimension High-Density** Semantic Scene Completion (**HD²-SSC**) framework with expanded pixel semantics and refined voxel occupancies. To bridge the dimension gap, a High-dimension Semantic Decoupling module is designed to expand 2D image features along a pseudo third dimension, decoupling coarse pixel semantics from occlusions, and then identify focal regions with fine semantics to enrich image features. To mitigate the density gap, a High-density Occupancy Refinement module is devised with a “detect-and-refine” architecture to leverage contextual geometric and semantic structures for enhanced semantic density with the completion of missing voxels and correction of erroneous ones. Extensive experiments and analyses on the SemanticKITTI and SSCBench-KITTI-360 datasets validate the effectiveness of our HD²-SSC framework.

Code —

<https://github.com/PKU-ICST-MIPL/HD2-AAAI2026>

Extended version — <https://arxiv.org/abs/2511.07925>

Introduction

Accurate 3D geometry perception of the surrounding environment is essential for autonomous driving systems, enabling downstream tasks such as planning, navigation, and interaction with dynamic environments. 3D Semantic Scene Completion (SSC) (Roldao, De Charette, and Verroust-Blondet 2022) aims to infer both the occupancy and semantics of the 3D space, providing a voxelized understanding of the scene. Considering the complexity of outdoor environments, LiDAR-based methods (Roldao, de Charette, and

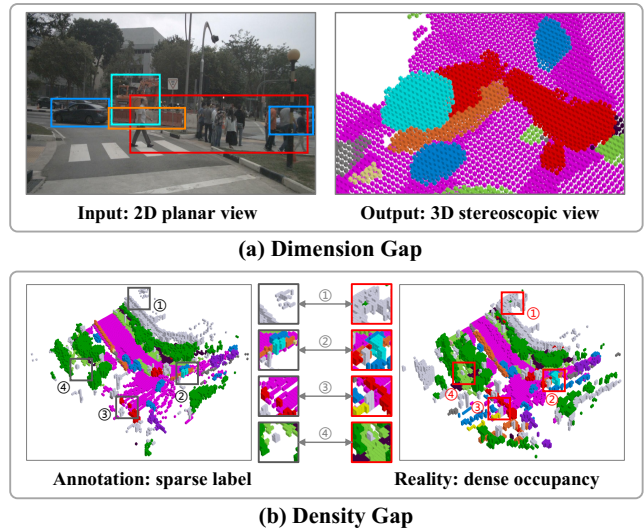


Figure 1: Illustrations of (a) Dimension Gap, highlighting the disparity between input coarse pixel semantics with occlusion and output distinct fine voxel semantics, (b) Density Gap, depicting the difference between annotated sparse labels and ground-truth dense occupancies.

Verroust-Blondet 2020; Cheng et al. 2021a; Xia et al. 2023; Li et al. 2023a; Lee et al. 2024; Wang et al. 2025) have been widely explored due to their ability to directly capture 3D structural information from input point clouds. However, they face the inherent challenges of high sensor cost and limited scalability for LiDAR data, restricting their applications to large-scale deployment in real-world scenarios.

Recently, camera-based semantic scene completion has garnered increasing attention as a cost-effective and scalable alternative utilizing 2D images as input. The pioneering work MonoScene (Cao and de Charette 2022) lifts 2D image features into dense 3D volumes with depth projection. Building on this, bird’s-eye-view (BEV) (Huang et al. 2021; Li et al. 2023b) and tri-perspective view (TVP) (Huang et al. 2023) representations are designed for improved computation efficiency. With these advances, transformer-based methods (Li et al. 2022; Wei et al. 2023) utilized voxel queries to aggregate image features into 3D scenes. Vox-

*Corresponding author.

Former (Li et al. 2023e) introduces a two-stage framework. Efforts have been made for improved performance through neighboring object structures and view spatial differences (Xue et al. 2024), and dense-sparse-dense design for hybrid guidance (Mei et al. 2024).

Despite significant progress, existing methods typically focus on refining voxel-based scene representations, but treat pixel features and voxel semantics indiscriminately during view transformation and occupancy prediction, facing two key challenges. (1) **Dimension Gap**: As illustrated in Figure 1 (a), input images are captured from a 2D planar view, leading to coarse pixel features confused by multiple object semantics with occlusions. In contrast, SSC requires fine-grained voxel semantics for distinct objects in a 3D stereoscopic view, necessitating the expansion and decoupling of coarse pixel features from occlusion. (2) **Density Gap**: As shown in Figure 1 (b), manual annotations derived from LiDAR sensors yield inherently sparse labels with interspace due to the finite resolution of collected point clouds. On the other hand, real-world scenes exhibit dense, consistent occupancy with rich contextual details, demanding comprehensive detection and meticulous refinement of voxel occupancy to enrich semantic density.

To address these challenges, we introduce the *High-Dimension High-Density Semantic Scene Completion (HD²-SSC)* framework. Specifically, the *High-dimension Semantic Decoupling (HSD)* module first expands and decouples coarse pixel semantics to alleviate the dimension gap. The process begins with expanding 2D image features into voxelized features through the Pseudo Voxelization block, with an orthogonal loss promoting distinct expanded semantics. The Semantic Aggregation block then extracts global semantics with pixel queries and integrates high-dimension voxelized features using semantic clustering with decoupling loss. Then, to mitigate the density gap, the *High-density Occupancy Refinement (HOR)* module operates in a “detect-and-refine” manner to enhance semantic density with contextual geometric and semantic structures. In the detection phase, binary predictions are generated for a comprehensive collection of occupied voxels, focusing on the identification of geometric critical voxels. In the subsequent refinement phase, we conduct class-wise prediction to identify semantic critical voxels, with respect to the prediction confidence score. Then the overall distributions of geometric and semantic critical voxels are aligned for consistent contextual structures, enhancing semantic density with completion of missing voxels and correction of erroneous ones. Extensive experiments and analyses on SemanticKITTI and SSCBench-KITTI-360 demonstrate the superiority of our HD²-SSC over SOTA methods.

The main contributions are summarized as follows:

- We propose HD²-SSC to address the challenges of dimension and density gap in camera-based semantic scene completion, improving SSC performance with completion of missing voxels and correction of erroneous ones.
- HSD bridges the dimension gap through expanding and decoupling coarse pixel semantics with orthogonal loss, then aggregating high-dimension voxelized semantics upon semantic clustering with decoupling loss.

- HOR addresses the density gap in a “detect-and-refine” paradigm to identify geometric and semantic critical voxels, aligning the overall critical distributions for consistent contextual details and improved semantic density.

Related Work

LiDAR-based SSC Methods

Due to the complexity of outdoor scenes and rich 3D structural information embedded in point clouds, LiDAR-based methods (Liong et al. 2020; Cheng et al. 2021b; Zhang et al. 2023; Ye et al. 2023) have long been the predominant solutions for SSC. UDNNet (Zou et al. 2021) pioneered the use of a 3D U-Net architecture with LiDAR data. SGCNet (Zhang et al. 2018) enhanced efficiency through spatial group convolutions for sparse 3D processing. Further advancements focused on refining 3D scene representations, such as multi-view fusion for completing sparse scenes (Cheng et al. 2021a) and local implicit functions for continuous scene modeling (Rist et al. 2021). SCPNet (Xia et al. 2023) transferred multi-frame information to single-frame models for lightweight architectures. SSA-SC (Yang et al. 2021) and SSC-RS (Mei et al. 2023) adopted multi-branch designs to hierarchically fuse semantic and geometric features. Despite the advancements, LiDAR-based methods remain computationally expensive due to large volume of LiDAR points, limiting their scalability to real-world applications.

Camera-based SSC Methods

In recent years, camera-based SSC methods (Miao et al. 2023; Tong et al. 2023; Yao et al. 2023; Wang et al. 2024; Yang and Peng 2025b,a) have gained increasing attention due to their cost-effectiveness and ease of deployment. MonoScene (Cao and de Charette 2022) pioneers a voxel-based approach by sampling image features along lines of sight. Another line of research leverages Bird’s-Eye-View (BEV) representations to aggregate multi-view image features. LSS (Phillion and Fidler 2020) introduces a frustum-based lifting mechanism, which is further enhanced by BEVDet (Huang et al. 2021) through aligning multi-frame features. BEVDepth (Li et al. 2023c) and BEVStereo (Li et al. 2023b) incorporate camera-aware depth estimation and dynamic temporal stereo information for more precise depth learning. Beyond BEV, TPVFormer (Huang et al. 2023) proposes a Tri-Perspective View (TPV) representation, decomposing voxel grids onto three orthogonal planes for efficient scene encoding. Transformer-based methods utilize BEV queries (Jiang et al. 2023) or voxel queries (Ma et al. 2024). Voxformer (Li et al. 2023e) adopts a two-stage framework with class-agnostic query proposals. Symphonize (Jiang et al. 2024) introduces a scene-from-instance paradigm to enhance feature interactions. SGN (Mei et al. 2024) presents a one-stage SSC framework with a dense-sparse-dense strategy, exploiting hybrid guidance to improve segmentation boundaries for more precise 3D semantic scene completion.

In general, existing camera-based SSC methods mainly focus on devising delicate architectures to refine 3D features, but overlook two critical challenges: the input-output

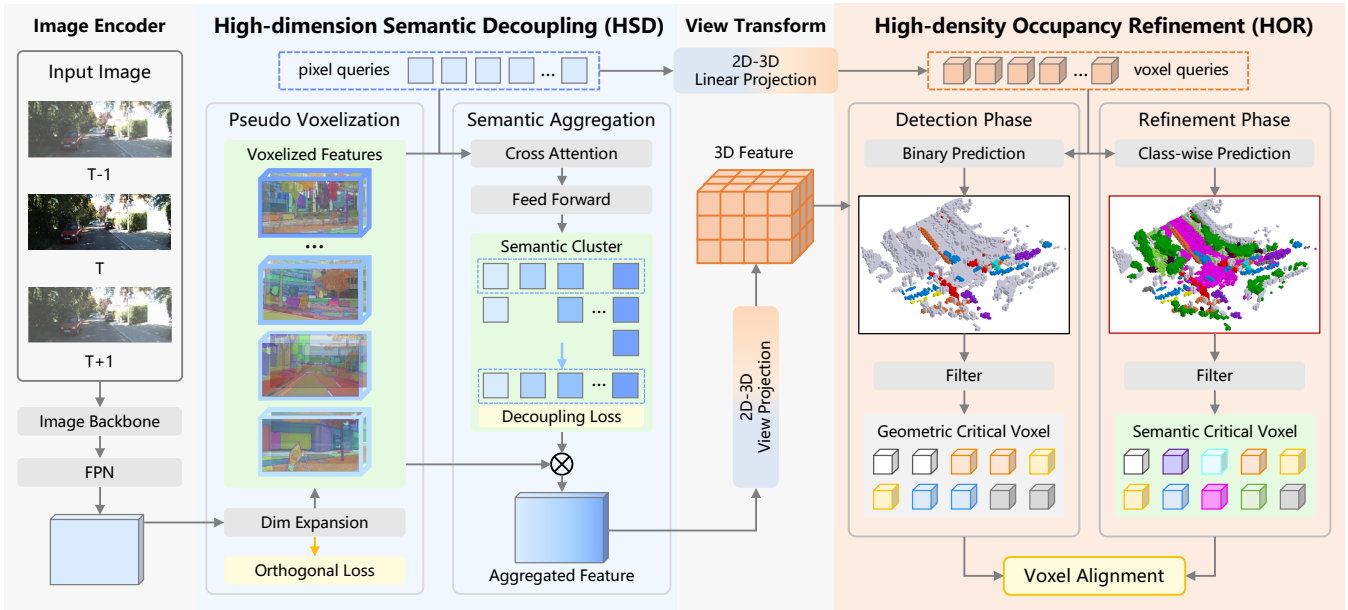


Figure 2: The overall architecture of our HD²-SSC. The High-dimension Semantic Decoupling (HSD) module expands and decouples coarse pixel semantics with orthogonal loss, then aggregates high-dimension voxelized semantics via semantic clustering with decoupling loss. The High-density Occupancy Refinement (HOR) module adopts a “detect-and-refine” architecture to identify geometric and semantic critical voxels, whose overall distributions are aligned for consistent contextual details.

dimension gap and annotation-reality density gap. In contrast, our HD²-SSC approach decouples coarse pixel semantics to bridge the dimension gap for aligned view transformation, and aligns critical voxel distributions for consistent contextual details to mitigate the density gap.

Approach

Figure 2 illustrates our HD²-SSC framework, consisting of four main modules: (1) an Image Encoder extracting 2D features, (2) an HSD module that addresses the dimension gap by decoupling coarse pixel semantics with orthogonal loss and aggregating high-dimension voxelized semantics upon semantic clustering, (3) a View Transform module projecting 2D pixel features to 3D voxel features, and (4) an HOR module that mitigates the density gap through aligning geometric and semantic critical voxel distributions in a “detect-and-refine” manner.

Preliminary

Problem Setup. Given input stereo images $I_l^{\text{rgb}}, I_r^{\text{rgb}}$, the objective of SSC is to predict the geometry and semantics of 3D scenes in front. The output is represented as a voxel grid $Y \in \mathbb{R}^{H \times W \times Z}$, where H, W, Z correspond to the grid’s length, width, and height, respectively. Each voxel is categorized as either empty denoted by c_0 or occupied by one of the semantic classes in $c \in \{c_1, \dots, c_N\}$.

Image Encoder. Consistent with previous methods (Mei et al. 2024), we employ the ResNet-50 (He et al. 2016) network with FPN (Lin et al. 2017) to extract 2D features $F^{2D} \in \mathbb{R}^{N_t \times C \times H \times W}$ from input RGB images, where N_t is

the image number of temporal inputs, C is the feature channel and (H, W) denotes the image resolution.

View Transformer. Following (Mei et al. 2024), we construct 3D features by sampling 2D features via 3D-2D projection mapping with camera parameters. Let $x \in \mathbb{R}^{H \times W \times Z \times 3}$ denote the centroid of $H \times W \times Z$ voxels in world coordinates. We establish the projection mapping $\pi(x)$ from pixel (u, v) using the camera intrinsic parameter K and extrinsic parameter $T = [R, t]$:

$$\begin{aligned} [h_c, w_c, z_c]^T &= R \cdot p + t \\ z_c \circ [u, v, 1]^T &= K \cdot [h_c, w_c, z_c]^T \end{aligned} \quad (1)$$

where \circ denotes element-wise product.

High-dimension Semantic Decoupling

HSD module consists of two blocks to address the dimension gap. The Pseudo Voxelization block expands and decouples coarse pixel semantics, and the Semantic Aggregation block aggregates high-dimension voxelized semantics.

Pseudo Voxelization. Given image feature F_{cam} extracted from the image encoder, a Dim Expansion (DE) layer is employed to lift it along a pseudo “semantic dimension” into pseudo voxelized features. This step is essential for addressing the dimension gap, generating multiple candidates for decoupling occluded objects within coarse pixel semantics:

$$F_{\text{pseudo}} = \text{DE}(F_{\text{cam}}, D_{\text{exp}}) \quad (2)$$

where $F_{\text{pseudo}} = \{F_{\text{pseudo}}^i \in \mathbb{R}^{N_{\text{cam}} \times C_{2D} \times H_{2D} \times W_{2D}}\}_{i=1}^{D_{\text{exp}}}$ represents pseudo voxelized features, and D_{exp} is the chan-

nel of expanded dimension, $DE(\cdot)$ represents the dim expansion operation with 2D convolution layers. However, the pseudo voxelized features may still exhibit similar semantics at the same 2D coordinates (h_{2D}, w_{2D}) , which limits their ability to generalize across different occluded objects. Therefore, we employ an orthogonal loss to further promote the distinct voxelized semantics:

$$\mathcal{L}_{\text{orth}} = \lambda |W_{\text{DE}}W_{\text{DE}}^T - I| \quad (3)$$

where W_{DE} represents the weight matrix of the dim expansion projection, I indicates the identity matrix, and λ is the regularization parameter for the orthogonal loss.

Semantic Aggregation. We first introduce pixel queries $Q_{\text{pixel}} \in \mathbb{R}^{N_{\text{query}} \times C_{2D}}$ to collect global semantics from the pseudo voxelized features through cross attention mechanism. Then, to further assign different pseudo voxelized feature slices F_{pseudo}^i with discrete semantics, we perform semantic clustering (DPC-kNN) on the collected global semantics in Q_{pixel} :

$$\text{DPC}(Q_{\text{pixel}}, D_{\text{exp}}) = \arg \min_C \sum_{i=1}^{D_{\text{exp}}} \sum_{v \in C_i} \|v - c_i\|^2 \quad (4)$$

where C represents the D_{exp} semantic clusters corresponding to the expanded dimension of the pseudo voxelized features, and c_i is the centroid of the i -th cluster C_i . Additionally, a decoupling loss is adopted to enhance the distinction among different semantic clusters:

$$\mathcal{L}_{\text{decouple}} = \sum_{i \neq j} \frac{C_i \cdot C_j}{|C_i| \cdot |C_j|} \quad (5)$$

Then, discriminative regions of the pseudo voxelized features are located with respect to the similarity against semantic clusters, thereby aggregating high-dimension voxelized semantics, as illustrated in Figure 3:

$$F_{\text{agg}} = \sum_{i=1}^{D_{\text{exp}}} F_{\text{pseudo}}^i \cdot \max_{C_j} \text{sim}(F_{\text{pseudo}}^i, C_j) \quad (6)$$

High-density Occupancy Refinement

HOR module is designed in a ‘‘detect-and-refine’’ architecture to identify geometric and semantic critical voxels, respectively. Then the critical voxels are aligned to ensure contextual geometric and semantic consistency.

Detection Phase. During the detection phase, the voxel queries $Q_{\text{voxel}} \in \mathbb{R}^{N_{\text{query}} \times C_{3D}}$ and voxel feature $F_{\text{voxel}} \in \mathbb{R}^{C_{3D} \times H_{3D} \times W_{3D} \times Z_{3D}}$ are passed through a binary classification head for a comprehensive detection of occupied voxels, together with a heuristic separation between foreground and background voxels:

$$M_{\text{o-f}}, M_{\text{f-b}} = H_{\text{bc}}(F_{\text{voxel}}, Q_{\text{voxel}}) \quad (7)$$

where H_{bc} denotes the binary classification head, $M_{\text{o-f}}, M_{\text{f-b}}$ represent the score maps distinguishing

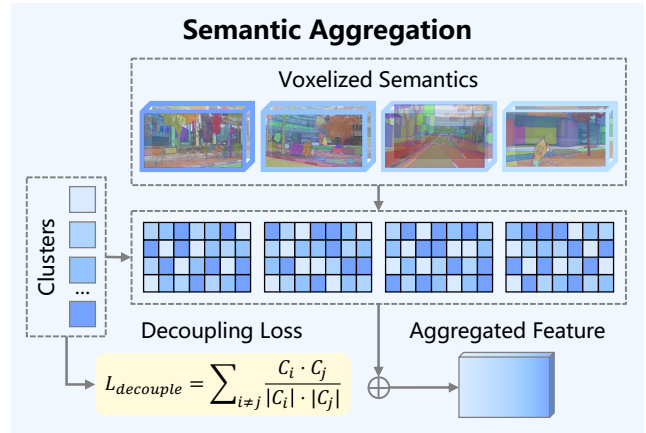


Figure 3: Illustration of aggregating the high-dimension voxelized semantics concerning the semantic clusters with decoupling loss.

occupied voxels from free ones, foreground voxels from background ones, respectively. The generated score maps are added to represent the geometric density score of each voxel, indicating its general structural significance. Then we select geometric critical voxels with the highest geometric density scores:

$$V_{\text{geo}} = \text{top}_k(M_{\text{o-f}} + M_{\text{f-b}}) \quad (8)$$

where k is the number of selected geometric critical voxels.

In the detection phase, we focus on the general classification between occupied and free, foreground and background voxels, and select geometric critical voxels with discriminative structural contextual information.

Refinement Phase. During the refinement phase, we first conduct class-wise prediction with the voxel queries and voxel features to generate the initial semantic scene completion results:

$$Y_{\text{init}} = H_{\text{cc}}(F_{\text{voxel}}, Q_{\text{voxel}}) \quad (9)$$

where H_{cc} denotes the class-wise classification head, Y_{init} represents the initial semantic scene completion results. Furthermore, we select semantic critical voxels with Y_{init} indicating the classification confidence and contextual semantic significance:

$$V_{\text{sem}} = \text{top}_k(\max(Y_{\text{init}})) \quad (10)$$

where k is the number of selected semantic critical voxels, aligned with geometric critical voxels.

Voxel Alignment. After selecting geometric and semantic critical voxels with discriminative geometric and semantic contextual information, respectively, we align the overall critical voxel distributions to promote consistent contextual semantic and geometry for further refinement:

$$\mathcal{L}_{\text{critical}} = \mathcal{L}_{\text{KL}}(V_{\text{geo}} || V_{\text{sem}}) + \mathcal{L}_{\text{KL}}(V_{\text{sem}} || V_{\text{geo}}) \quad (11)$$

where $\mathcal{L}_{\text{KL}}(\cdot || \cdot)$ is the Kullback-Leibler divergence loss measuring the difference between two distributions. With

Method	SC		SSC																			
	IoU	mIoU	car(3.92%)	bicycle(0.03%)	motorcycle(0.03%)	truck(0.16%)	other-veh.(0.20%)	person(0.07%)	bicyclist(0.07%)	motorcyclist(0.05%)	road(1.530%)	parking(1.12%)	sidewalk(1.13%)	other-grnd.(0.56%)	building(14.10%)	fence(3.90%)	vegetation(39.3%)	trunk(0.51%)	terrain(9.17%)	pole(0.29%)	traf.-sign(0.08%)	
<i>Mono-Input Methods</i>																						
MonoScene [†] (2022)	36.86	11.08	23.26	0.61	0.45	6.98	1.48	1.86	1.20	0.00	56.52	14.27	26.72	0.46	14.09	5.84	17.89	2.81	29.64	4.14	2.25	
TPVFormer [†] (2023)	35.61	11.36	23.81	0.36	0.05	8.08	4.35	0.51	0.89	0.00	56.50	20.60	25.87	0.85	13.88	5.94	16.92	2.26	30.38	3.14	1.52	
OccFormer [†] (2023)	36.50	13.46	25.09	0.81	1.19	25.53	8.52	2.78	2.82	0.00	58.85	19.61	26.88	0.31	14.40	5.61	19.63	3.93	32.62	4.26	2.86	
IAMSSC* (2024)	44.29	12.45	26.26	0.60	0.15	8.74	5.06	1.32	3.46	0.01	54.55	16.02	25.85	0.70	17.38	6.86	24.63	4.95	30.13	6.35	3.56	
<i>Temporal-Input Methods</i>																						
VoxFormer* (2023e)	44.15	13.35	26.54	1.28	0.56	7.26	7.81	1.93	1.97	0.00	53.57	19.69	26.52	0.42	19.54	7.31	26.10	6.10	33.06	9.15	4.94	
DepthSSC* (2023)	45.84	13.28	25.94	0.35	1.16	6.02	7.50	<u>2.58</u>	6.32	0.00	55.38	18.76	27.04	<u>0.92</u>	19.23	8.46	26.37	4.52	30.19	7.42	4.09	
Symphonies* (2024)	41.92	14.89	28.68	<u>2.54</u>	2.82	<u>20.44</u>	13.89	3.52	2.24	0.00	56.37	15.28	27.58	0.95	21.64	8.40	25.72	6.60	30.87	9.57	5.76	
HASSC* (2024)	44.58	14.74	27.33	1.07	1.14	17.06	8.83	2.25	<u>4.09</u>	0.00	57.23	19.89	29.08	1.26	20.19	7.95	27.01	7.71	33.95	9.20	4.81	
H2GFormer* (2024)	44.69	14.29	28.21	0.95	0.91	6.80	9.32	1.15	0.10	0.00	57.00	21.74	29.37	0.34	20.51	7.98	27.44	7.80	36.26	9.88	5.81	
CGFormer [†] (2024)	45.99	<u>16.87</u>	<u>34.32</u>	4.61	<u>2.71</u>	<u>19.44</u>	7.67	2.38	4.08	0.00	65.51	<u>20.82</u>	<u>32.31</u>	0.16	23.52	<u>9.20</u>	26.93	8.83	39.54	10.67	<u>7.84</u>	
SGN* (2024)	<u>46.21</u>	15.32	<u>33.31</u>	0.61	0.46	6.03	9.84	0.47	0.10	0.00	59.10	19.05	29.41	0.33	<u>25.17</u>	9.96	<u>28.93</u>	<u>9.58</u>	38.12	<u>13.25</u>	7.32	
HD²-SSC* (Ours)	47.59	17.44	32.68	1.72	0.54	24.76	<u>10.73</u>	1.61	1.73	0.00	<u>60.80</u>	23.72	33.12	0.11	28.55	8.83	30.90	11.08	<u>38.78</u>	13.41	8.34	

Table 1: Camera-based 3D semantic scene completion results on the SemanticKITTI (Behley et al. 2019) validation set. [†] denotes the methods employing EfficientNet-B7 (Tan and Le 2019) as image backbone, and * represents the methods utilizing ResNet-50 (He et al. 2016) as image backbone. Best results are highlights in bold, and second-best scores are underlined.

the aligned critical voxels serving as soft constraint on the complementary effect of contextual geometry and semantics, we utilize a refined multi-layer perceptron (MLP_{refine}) to enhance the initial SSC predictions:

$$Y_{\text{refine}} = Y_{\text{init}} + \text{MLP}_{\text{refine}}([V_{\text{geo}}, V_{\text{sem}}]) \quad (12)$$

where Y_{refine} represents the final refined SSC prediction, $[\cdot, \cdot]$ is the concatenation operation.

Experiments

In accordance with existing camera-based SSC methods (Mei et al. 2024; Li et al. 2023e), extensive experiments and analyses are conducted to validate our HD²-SSC on the SemanticKITTI and SSCBench-KITTI-360 datasets.

Experimental Setup

Dataset. (1) The **SemanticKITTI** (Behley et al. 2019) dataset, based on the KITTI Odometry Benchmark (Geiger, Lenz, and Urtasun 2012), contains 22 autonomous driving sequences with 20 semantic classes and annotations. We follow the official protocol to split the total sequences into (00-07, 09-10) / (08) / (11-21) for the training / validation / test sets respectively. (2) The **SSCBench-KITTI-360** (Li et al. 2023d) dataset consists of 9 densely annotated sequences of urban driving scenes. It is separated into a training set from sequences (00, 02-05, 07, 10), a validation set from sequence (06), and a test set from sequence (09). For both datasets, the semantic labels are within the range of $[0 \sim 51.2m, -25.6 \sim 25.6m, -2 \sim 4.4m]$, and the target

voxel grids are represented as $256 \times 256 \times 32$ voxel grids with the resolution of $0.2m$.

Evaluation Metrics. Following (Mei et al. 2024), we adopt the Intersection over Union (IoU) of occupied voxels as the evaluation metric for the class-agnostic scene completion (SC) task, and the mean Intersection over Union (mIoU) metric for the semantic scene completion (SSC) task.

Implementation Details. We crop the input RGB images of cam2 to size 1220×370 for SemanticKITTI and images of cam1 to size 1408×376 for SSCBench-KITTI-360, extracting 2D feature maps with $1/16$ of input resolutions. The feature dimension C is set to 128. The size $H \times W \times Z$ of the 3D feature volume is $128 \times 128 \times 16$, and the final predictions are upsampled to $256 \times 256 \times 32$. The channel dimensions for 2D and 3D features are $C_{2D} = 256$ and $C_{3D} = 32$. The expanded dimension in the pseudo voxelization block is $D_{\text{exp}} = 4$, and the number of query vectors is $N_{\text{query}} = 100$. The number of critical voxels is set as $k = 4096$. HD²-SSC is trained for 24 epochs on 4 A6000 GPUs with a total batchsize of 4. The AdamW (Loshchilov 2017) optimizer is utilized with an initial learning rate of $2e-4$ and a weight decay of $1e-2$.

Main Results

Quantitative comparison. Table 1 and Table 2 present the comparison results between our HD²-SSC approach and other state-of-the-art camera-based SSC methods on the SemanticKITTI validation set and SSCBench-KITTI-360 test

Method	SSC		IoU																		
	IoU	mIoU	car(2.85%)	bicycle(0.01%)	motorcycle(0.01%)	truck(0.16%)	other-veh.(5.75%)	person(0.02%)	road(14.98%)	parking(2.31%)	sidewalk(6.43%)	other-grnd.(2.05%)	building(15.67%)	fence(0.96%)	vegetation(41.99%)	terrain(7.10%)	pole(0.22%)	traf.-sign(0.06%)	other-struct.(4.33%)	other-obj.(0.28%)	
<i>Mono-Input Methods</i>																					
MonoScene [†] (2022)	37.87	12.31	19.34	0.43	0.58	8.02	2.03	0.86	48.35	11.38	28.13	3.32	32.89	3.53	26.15	16.75	6.92	5.67	4.20	3.09	
GaussianFormer2* (2024)	38.31	13.90	21.08	2.55	4.21	12.41	5.73	1.59	54.12	11.04	32.31	3.34	32.01	4.98	28.94	17.33	3.57	5.48	5.88	3.54	
TPVFormer [†] (2023)	40.22	13.64	21.56	1.09	1.37	8.06	2.57	2.38	52.99	11.99	31.07	3.78	34.83	4.80	30.08	17.52	7.46	5.86	5.48	2.70	
OccFormer [†] (2023)	40.27	13.81	22.58	0.66	0.26	9.89	3.82	2.77	54.30	13.44	31.53	3.55	36.42	4.80	31.00	19.51	7.77	8.51	6.95	4.60	
IAMSSC* (2024)	41.80	12.97	18.53	2.45	1.76	5.12	3.92	3.09	47.55	10.56	28.35	4.12	31.53	6.28	29.17	15.24	8.29	7.01	6.35	4.19	
<i>Temporal-Input Methods</i>																					
VoxFormer* (2023e)	38.76	11.91	17.84	1.16	0.89	4.56	2.06	1.63	47.01	9.67	27.21	2.89	31.18	4.97	28.99	14.69	6.51	6.92	3.79	2.43	
DepthSSC* (2023)	40.85	14.28	21.90	2.36	<u>4.30</u>	11.51	4.56	2.92	50.88	12.89	30.27	2.49	37.33	5.22	29.61	21.59	5.97	7.77	5.24	3.51	
Symphonies* (2024)	44.12	18.58	<u>30.02</u>	1.85	5.90	<u>25.07</u>	<u>12.06</u>	8.20	54.94	13.83	32.76	<u>6.93</u>	35.11	<u>8.58</u>	38.33	11.52	14.01	<u>9.57</u>	<u>14.44</u>	<u>11.28</u>	
CGFormer [†] (2024)	<u>48.07</u>	<u>20.05</u>	29.85	<u>3.42</u>	3.96	17.59	6.79	<u>6.63</u>	63.85	<u>17.15</u>	40.72	5.53	42.73	8.22	<u>38.80</u>	24.94	16.24	17.45	10.18	6.77	
SGN* (2024)	47.06	18.25	29.03	3.43	2.90	10.89	5.20	2.99	58.14	15.04	36.40	4.43	<u>42.02</u>	7.72	38.17	<u>23.22</u>	<u>16.73</u>	<u>16.38</u>	9.93	5.86	
HD²-SSC* (Ours)	48.58	20.62	33.00	0.00	2.98	26.90	12.38	4.75	<u>60.17</u>	17.58	<u>37.55</u>	8.7	40.72	9.51	41.69	14.69	18.7	12.35	16.65	12.75	

Table 2: Camera-based 3D semantic scene completion results on the SSCBench-KITTI-360 (Behley et al. 2019) test set. [†] denotes the methods employing EfficientNet-B7 (Tan and Le 2019) as image backbone, and * represents the methods utilizing ResNet-50 (He et al. 2016) as image backbone. Best results are highlights in bold, and second-best scores are underlined.

Variants	IoU	mIoU
Baseline	44.15	13.35
Baseline + HSD	46.45	15.58
Baseline + HOR	46.07	16.12
Baseline + HSD + HOR (Ours)	47.59	17.44

Table 3: Ablation study on the SemanticKITTI validation set dataset of different components of our HD²-SSC.

Variants	IoU	mIoU
HD²-SSC (Ours)	47.59	17.44
w/o L_{orth}	46.93 (-0.66)	16.64 (-0.8)
w/o $L_{decouple}$	46.85 (-0.74)	16.78 (-0.66)
w/o $L_{critical}$	46.49 (-1.1)	16.31 (-1.13)

Table 4: Ablation study on the SemanticKITTI validation set dataset of different loss terms.

set respectively, where the best results are highlighted in bold and the second-best results are underlined. It can be observed that our HD²-SSC approach achieves state-of-the-art performance in terms of both class-agnostic scene completion (IoU) and class-aware semantic scene completion (mIoU). Specifically, regarding the best comparison method SGN (Mei et al. 2024), our approach achieves superior performance with **1.38%** IoU, **2.12%** mIoU enhancement on the SemanticKITTI validation set and **1.52%** IoU, **2.37%** mIoU enhancement on the SSCBench-KITTI-360 test set,




Figure 4: Effect of the expanded dimension on the SSC performance, evaluated on the SemanticKITTI validation set.

respectively. Such performance improvements are attributed to the decoupled pixel semantics and aligned contextual details, which address the dimension and density gap. Although SGN employs a “dense-sparse-dense” architecture for dynamically selecting discriminative voxel features, it overlooks the coarse and confusing pixel semantics in image features, and thus suffers from misaligned geometric and semantic contextual details. In contrast, our HD²-SSC approach expands and decouples the coarse pixel semantics with orthogonal and decoupling loss, then aligns geometric and semantic critical voxels for consistent contextual details, enabling accurate SSC predictions with the completion of missing voxels and correction of erroneous ones.

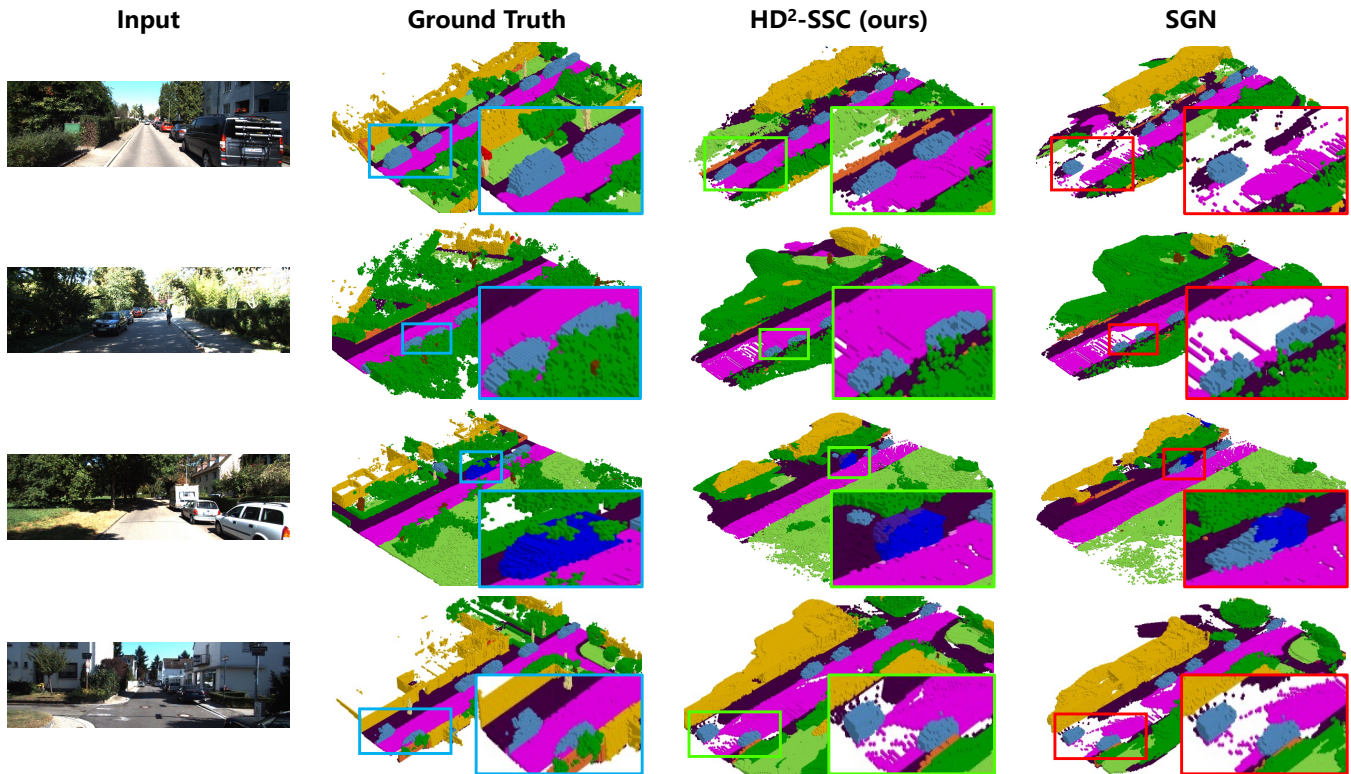


Figure 5: Visualization results of SSC prediction on the SemanticKITTI validation set. We highlight the occupancy ground truth with blue boxes, false SSC predictions of the best comparison method SGN with red boxes, and the improved SSC predictions from our HD²-SSC approach with green boxes. Better viewed when zoomed in.

Ablation on architectural components. To further assess the effectiveness of HD²-SSC, we conduct an ablation study to evaluate the impact of key components, as shown in Table 3. Incorporating HSD yields improvements of 2.30% IoU and 2.23% mIoU. Adding the HOR module enhances SSC performance by 1.92% IoU and 2.77% mIoU. These results underscore the benefits of decoupling coarse pixel semantics and aligning geometric and semantic distribution. Notably, HSD and HOR modules complement each other, with their combined application achieving superior results.

Ablation on different losses. Table 4 presents the ablation results on different losses, where disabling each loss causes performance degradation, validating their effectiveness: (1) L_{orth} improves semantic diversity for decoupling coarse pixel semantics. (2) L_{decouple} promotes semantic distinction in semantic clusters. (3) L_{critical} aligns critical voxel distributions for semantic and geometry consistency.

Effect of expanded dimension. Figure 4 demonstrates the effect of expanded dimension, where the best performance is achieved with $D_{\text{exp}} = 4$. When the expanded dimension keeps increasing, SSC performance begins to degrade. This phenomenon arises because, without explicit pixel semantic labels, too many expanded dimensions may bring additional confusion of “imaginary” semantics that do not correspond to real-world objects, leading to inferior performance.

Qualitative results. Figure 5 provides the visualization results from the SemanticKITTI validation set. We outline ground truth with blue boxes in the second row. The red boxes in the fourth row indicate false predictions where the best comparison method SGN (Mei et al. 2024) struggles, while the green boxes in the third row highlight the improved scene completion results from our HD²-SSC approach with completion of missing voxels and correction of erroneous ones. This improvement is attributed to HD²-SSC addressing the dimension and density gap with decoupled pixel semantics and refined voxel predictions.

Conclusion

In this paper, we identify the challenges of dimension and density gap in camera-based SSC, and introduce the HD²-SSC approach to address them. HSD expands and decouples coarse pixel semantics with orthogonal and decoupling loss, mitigating the dimension gap. HOR aligns geometric and semantic critical voxels to promote consistent contextual details, bridging the density gap. Experiments and analyses on the SemanticKITTI and SSCBench-KITTI-360 datasets validate the effectiveness of our HD²-SSC.

Limitation. Despite our effectiveness, failure cases still occur in occluded and distant areas with confusing features. Our future work aims to incorporate physical regularities to complement such low-quality semantic features.

Acknowledgments

This work was supported by the grants from the National Natural Science Foundation of China (62525201, 62132001, 62432001) and Beijing Natural Science Foundation (L247006, L257005).

References

- Behley, J.; Garbade, M.; Milioto, A.; Quenzel, J.; Behnke, S.; Stachniss, C.; and Gall, J. 2019. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9297–9307.
- Cao, A.-Q.; and de Charette, R. 2022. Monoscene: Monocular 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3991–4001.
- Cheng, R.; Agia, C.; Ren, Y.; Li, X.; and Bingbing, L. 2021a. S3cnet: A sparse semantic scene completion network for lidar point clouds. In *Conference on Robot Learning*, 2148–2161. PMLR.
- Cheng, R.; Razani, R.; Taghavi, E.; Li, E.; and Liu, B. 2021b. 2-s3net: Attentive feature fusion with adaptive feature selection for sparse semantic segmentation network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12547–12556.
- Geiger, A.; Lenz, P.; and Urtasun, R. 2012. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, 3354–3361. IEEE.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Huang, J.; Huang, G.; Zhu, Z.; Ye, Y.; and Du, D. 2021. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*.
- Huang, Y.; Thammatadatrakoon, A.; Zheng, W.; Zhang, Y.; Du, D.; and Lu, J. 2024. Probabilistic Gaussian Superposition for Efficient 3D Occupancy Prediction. *arXiv preprint arXiv:2412.04384*.
- Huang, Y.; Zheng, W.; Zhang, Y.; Zhou, J.; and Lu, J. 2023. Tri-perspective view for vision-based 3d semantic occupancy prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9223–9232.
- Jiang, H.; Cheng, T.; Gao, N.; Zhang, H.; Lin, T.; Liu, W.; and Wang, X. 2024. Symphonize 3d semantic scene completion with contextual instance queries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20258–20267.
- Jiang, Y.; Zhang, L.; Miao, Z.; Zhu, X.; Gao, J.; Hu, W.; and Jiang, Y.-G. 2023. Polarformer: Multi-camera 3d object detection with polar transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1042–1050.
- Lee, J.; Lee, S.; Jo, C.; Im, W.; Seon, J.; and Yoon, S.-E. 2024. Semicity: Semantic scene generation with triplane diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 28337–28347.
- Li, P.; Zhao, R.; Shi, Y.; Zhao, H.; Yuan, J.; Zhou, G.; and Zhang, Y.-Q. 2023a. Lode: Locally conditioned eikonal implicit scene completion from sparse lidar. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 8269–8276. IEEE.
- Li, Y.; Bao, H.; Ge, Z.; Yang, J.; Sun, J.; and Li, Z. 2023b. Bevstereo: Enhancing depth estimation in multi-view 3d object detection with temporal stereo. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1486–1494.
- Li, Y.; Ge, Z.; Yu, G.; Yang, J.; Wang, Z.; Shi, Y.; Sun, J.; and Li, Z. 2023c. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1477–1485.
- Li, Y.; Li, S.; Liu, X.; Gong, M.; Li, K.; Chen, N.; Wang, Z.; Li, Z.; Jiang, T.; Yu, F.; et al. 2023d. SSCBench: A Large-Scale 3D Semantic Scene Completion Benchmark for Autonomous Driving. *arXiv preprint arXiv:2306.09001*.
- Li, Y.; Yu, Z.; Choy, C.; Xiao, C.; Alvarez, J. M.; Fidler, S.; Feng, C.; and Anandkumar, A. 2023e. Voxformer: Sparse voxel transformer for camera-based 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9087–9098.
- Li, Z.; Wang, W.; Li, H.; Xie, E.; Sima, C.; Lu, T.; Qiao, Y.; and Dai, J. 2022. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European conference on computer vision*, 1–18. Springer.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Liong, V. E.; Nguyen, T. N. T.; Widjaja, S.; Sharma, D.; and Chong, Z. J. 2020. Amvnet: Assertion-based multi-view fusion network for lidar semantic segmentation. *arXiv preprint arXiv:2012.04934*.
- Loshchilov, I. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Ma, Q.; Tan, X.; Qu, Y.; Ma, L.; Zhang, Z.; and Xie, Y. 2024. Cotr: Compact occupancy transformer for vision-based 3d occupancy prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19936–19945.
- Mei, J.; Yang, Y.; Wang, M.; Huang, T.; Yang, X.; and Liu, Y. 2023. SSC-RS: Elevate LiDAR semantic scene completion with representation separation and BEV fusion. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1–8. IEEE.
- Mei, J.; Yang, Y.; Wang, M.; Zhu, J.; Ra, J.; Ma, Y.; Li, L.; and Liu, Y. 2024. Camera-based 3d semantic scene completion with sparse guidance network. *IEEE Transactions on Image Processing*.

- Miao, R.; Liu, W.; Chen, M.; Gong, Z.; Xu, W.; Hu, C.; and Zhou, S. 2023. Occdepth: A depth-aware method for 3d semantic scene completion. *arXiv preprint arXiv:2302.13540*.
- Phillion, J.; and Fidler, S. 2020. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, 194–210. Springer.
- Rist, C. B.; Emmerichs, D.;ENZweiler, M.; and Gavrilu, D. M. 2021. Semantic scene completion using local deep implicit functions on lidar data. *IEEE transactions on pattern analysis and machine intelligence*, 44(10): 7205–7218.
- Roldao, L.; de Charette, R.; and Verroust-Blondet, A. 2020. Lmscnet: Lightweight multiscale 3d semantic completion. In *2020 International Conference on 3D Vision (3DV)*, 111–119. IEEE.
- Roldao, L.; De Charette, R.; and Verroust-Blondet, A. 2022. 3D semantic scene completion: A survey. *International Journal of Computer Vision*, 130(8): 1978–2005.
- Tan, M.; and Le, Q. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, 6105–6114. PMLR.
- Tong, W.; Sima, C.; Wang, T.; Chen, L.; Wu, S.; Deng, H.; Gu, Y.; Lu, L.; Luo, P.; Lin, D.; et al. 2023. Scene as occupancy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8406–8415.
- Wang, R.; Ma, Y.; Yao, Y.; Tao, S.; Li, H.; Zhu, Z.; Liu, Y.; and Zuo, X. 2025. L2COcc: Lightweight Camera-Centric Semantic Scene Completion via Distillation of LiDAR Model. *arXiv preprint arXiv:2503.12369*.
- Wang, S.; Yu, J.; Li, W.; Liu, W.; Liu, X.; Chen, J.; and Zhu, J. 2024. Not all voxels are equal: Hardness-aware semantic scene completion with self-distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14792–14801.
- Wang, Y.; and Tong, C. 2024. H2gformer: Horizontal-to-global voxel transformer for 3d semantic scene completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 5722–5730.
- Wei, Y.; Zhao, L.; Zheng, W.; Zhu, Z.; Zhou, J.; and Lu, J. 2023. Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 21729–21740.
- Xia, Z.; Liu, Y.; Li, X.; Zhu, X.; Ma, Y.; Li, Y.; Hou, Y.; and Qiao, Y. 2023. Scpnet: Semantic scene completion on point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 17642–17651.
- Xiao, H.; Xu, H.; Kang, W.; and Li, Y. 2024. Instance-aware monocular 3D semantic scene completion. *IEEE Transactions on Intelligent Transportation Systems*, 25(7): 6543–6554.
- Xue, Y.; Li, R.; Wu, F.; Tang, Z.; Li, K.; and Duan, M. 2024. Bi-SSC: Geometric-Semantic Bidirectional Fusion for Camera-based 3D Semantic Scene Completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20124–20134.
- Yang, X.; Zou, H.; Kong, X.; Huang, T.; Liu, Y.; Li, W.; Wen, F.; and Zhang, H. 2021. Semantic segmentation-assisted scene completion for lidar point clouds. In *RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3555–3562.
- Yang, Z.; and Peng, Y. 2025a. GaLa-2.5 D: Global-Local Alignment with 2.5 D Semantic Guidance for Camera-based 3D Semantic Scene Completion in Autonomous Driving. *Chinese Journal of Electronics*.
- Yang, Z.; and Peng, Y. 2025b. SPHERE: Semantic-PHysical Engaged REpresentation for 3D Semantic Scene Completion. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 7681–7690.
- Yao, J.; Li, C.; Sun, K.; Cai, Y.; Li, H.; Ouyang, W.; and Li, H. 2023. Ndc-scene: Boost monocular 3d semantic scene completion in normalized device coordinates space. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 9421–9431. IEEE Computer Society.
- Yao, J.; and Zhang, J. 2023. Depthssc: Depth-spatial alignment and dynamic voxel resolution for monocular 3d semantic scene completion. *arXiv preprint arXiv:2311.17084*.
- Ye, D.; Zhou, Z.; Chen, W.; Xie, Y.; Wang, Y.; Wang, P.; and Foroosh, H. 2023. Lidarmultinet: Towards a unified multi-task network for lidar perception. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 3231–3240.
- Yu, Z.; Zhang, R.; Ying, J.; Yu, J.; Hu, X.; Luo, L.; Cao, S.-Y.; and Shen, H.-L. 2024. Context and geometry aware voxel transformer for semantic scene completion. *arXiv preprint arXiv:2405.13675*.
- Zhang, J.; Zhao, H.; Yao, A.; Chen, Y.; Zhang, L.; and Liao, H. 2018. Efficient semantic scene completion network with spatial group convolution. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 733–749.
- Zhang, Y.; Zhu, Z.; and Du, D. 2023. Occformer: Dual-path transformer for vision-based 3d semantic occupancy prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9433–9443.
- Zhang, Z.; Wang, B.; Yu, Z.; and Zhao, F. 2023. Attention guided enhancement network for weakly supervised semantic segmentation. *Chinese Journal of Electronics*, 32(4): 896–907.
- Zou, H.; Yang, X.; Huang, T.; Zhang, C.; Liu, Y.; Li, W.; Wen, F.; and Zhang, H. 2021. Up-to-down network: Fusing multi-scale context for 3d semantic scene completion. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 16–23. IEEE.