

Unleashing Semantic and Geometric Priors for 3D Scene Completion

Shiyuan Chen^{1*}, Wei Sui^{2*}, Bohao Zhang², Zeyd Boukhers³, John See⁴, Cong Yang^{1†}

¹School of Future Science and Engineering, Soochow University, Suzhou, China

²D-Robotics, Beijing, China

³Fraunhofer Institute for Applied Information Technology, Sankt Augustin, Germany

⁴School of Mathematical and Computing Sciences, Heriot-Watt University Malaysia, Putrajaya, Malaysia

sychen66@stu.suda.edu.cn, {wei.sui, bohao.zhang}@d-robotics.cc,

zeyd.boukhers@fit.fraunhofer.de, J.See@hw.ac.uk, cong.yang@suda.edu.cn

Abstract

Camera-based 3D semantic scene completion (SSC) provides dense geometric and semantic perception for autonomous driving and robotic navigation. However, existing methods rely on a coupled encoder to deliver both semantic and geometric priors, which forces the model to make a trade-off between conflicting demands and limits its overall performance. To tackle these challenges, we propose FoundationSSC, a novel framework that performs dual decoupling at both the source and pathway levels. At the source level, we introduce a foundation encoder that provides rich semantic feature priors for the semantic branch and high-fidelity stereo cost volumes for the geometric branch. At the pathway level, these priors are refined through specialised, decoupled pathways, yielding superior semantic context and depth distributions. Our dual-decoupling design produces disentangled and refined inputs, which are then utilised by a hybrid view transformation to generate complementary 3D features. Additionally, we introduce a novel Axis-Aware Fusion (AAF) module that addresses the often-overlooked challenge of fusing these features by anisotropically merging them into a unified representation. Extensive experiments demonstrate the advantages of FoundationSSC, achieving simultaneous improvements in both semantic and geometric metrics, surpassing prior bests by +0.23 mIoU and +2.03 IoU on SemanticKITTI. Additionally, we achieve state-of-the-art performance on SSCBench-KITTI-360, with 21.78 mIoU and 48.61 IoU.

Code —

<https://github.com/D-Robotics-AI-Lab/FoundationSSC>

Introduction

Semantic Scene Completion (SSC) infers the complete geometry and semantics of a 3D scene from partial observations. It is essential for high-level 3D perception applications, such as intelligent driving (Zheng et al. 2024) and robotic navigation (Wu et al. 2025). Early SSC methods primarily utilised LiDAR input (Song et al. 2017; Roldao, De Charette, and Verroust-Blondet 2020), but they faced

*Equal contribution. Work done during Shiyuan Chen’s internship at D-Robotics.

†Corresponding author.

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

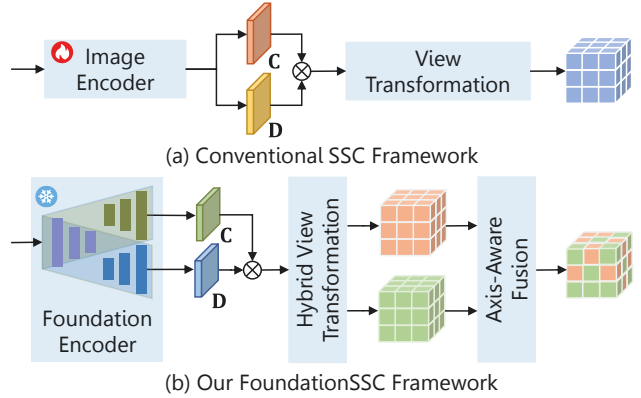


Figure 1: The conventional SSC framework (top) depends on a coupled image encoder, providing limited feature priors for both the semantic and geometric branches, resulting in an inherent trade-off. In contrast, our FoundationSSC (bottom) utilises a foundation encoder that provides robust decoupled semantic and geometric priors, effectively addressing conflicts at both the source and pathway levels.

challenges due to high hardware costs and limited resolution. As a result, there has been a recent shift towards camera-based SSC methods, which are gaining significant attention for their lower costs and richer texture information.

Despite significant advancements, existing camera-based SSC methods are fundamentally limited by their dependence on a single, coupled source of features, as shown in Fig. 1 (a). This restriction creates an inherent trade-off between the conflicting demands of semantic and geometric reasoning, ultimately hindering overall performance. Some research has attempted to address these issues by incorporating additional priors. For example, certain studies have aimed to enhance geometric understanding by utilising external stereo depth information (Li et al. 2024b). However, these approaches often function in a coarse and lossy manner, failing to fully capitalise on the rich probabilistic cues available in the cost volume. Conversely, other studies focus on integrating semantic priors from instance-level information (Xiao et al. 2024) or Vision-Language Models (VLMs) (Wang et al. 2025). While this approach can bene-

fit segmentation, it introduces a strong semantic bias into the coupled encoder, causing it to neglect the precise spatial features required for geometric reasoning and thus worsening the trade-off. Overall, these methods remain limited, as they merely add external inputs to a coupled framework without addressing the underlying feature conflict. Fortunately, recent advances in Vision Foundation Models (VFMs) (Oquab et al. 2024) present a promising opportunity. These models exhibit remarkable zero-shot generalisation on downstream tasks, such as semantic segmentation or depth estimation. Nonetheless, we believe that simply replacing the image encoder with VFMs is not enough. The primary challenge lies in designing a novel framework that can effectively leverage these robust priors to resolve the coupling problem.

In this paper, we propose FoundationSSC, a novel SSC framework designed to unleash the semantic and geometric priors from foundation models, as shown in Fig. 1 (b). Our approach operates on a core principle of dual decoupling. It begins with a unified foundation encoder that harnesses the power of VFMs to provide separate priors for the semantic and geometric branches, thereby addressing the trade-off at the feature source. These distinct priors are then processed through specialised, decoupled pathways aimed at maximising their potential. The resulting high-quality 2D representations are transformed into 3D space using a hybrid view transformation. Additionally, we introduce a novel Axis-Aware Fusion (AAF) module, which fuses the complementary 3D feature volumes in a way that respects their anisotropic characteristics, resulting in a unified 3D representation. Through these integrated designs, FoundationSSC effectively resolves the semantic-geometric coupling issues that have limited previous work, producing high-quality 2D and 3D feature representations. Consequently, our framework achieves state-of-the-art performance on SemanticKITTI (Behley et al. 2019) and SSCBench-KITTI-360 (Liao, Xie, and Geiger 2022) datasets, simultaneously enhancing both geometric (Intersection over Union, IoU) and semantic (mean IoU, mIoU) metrics. In summary, our contributions are as follows:

- We propose a novel framework that decouples semantics and geometric reasoning at both the source and pathway levels. This design fundamentally resolves the inherent trade-off of existing methods.
- We design an Axis-Aware Fusion (AAF) module that effectively fuses complementary 3D feature volumes from the hybrid view transformation, which are typically used in isolation or discarded by prior work, by respecting the anisotropic structure of driving scenes.
- The proposed FoundationSSC achieves state-of-the-art performance with a mIoU of 19.32 and an IoU of 48.12 on SemanticKITTI, as well as a mIoU of 21.78 and an IoU of 48.61 on SSCBench-KITTI-360, demonstrating the significant advantages of our decoupling design.

Related Work

Vision Foundation Models

VFMs are trained on large-scale web datasets, enabling them to learn highly generalisable visual representations

with impressive zero-shot capabilities. One of the early influential models, Contrastive Language-Image Pre-training (CLIP) (Radford et al. 2021), learns joint vision-language embeddings through contrastive learning. Subsequent models, like the Segment Anything Model (SAM) (Kirillov et al. 2023; Ravi et al. 2025), offer a universal segmentation framework; however, their features are not optimised for general-purpose representation learning. To address this limitation, the DINO series (Caron et al. 2021; Oquab et al. 2024) employs self-supervised learning to produce dense, high-quality features that transfer effectively to downstream tasks such as detection and segmentation. This success has paved the way for geometry-focused VFMs. The DepthAnything family (Yang et al. 2024a,b) fine-tunes DINOv2 for robust monocular depth estimation. More recently, FoundationStereo (Wen et al. 2025) has extended this approach to stereo matching, achieving strong zero-shot generalisation across diverse domains. Inspired by these advancements, some recent SSC works have leveraged the priors of VFMs to enhance semantic features (Cui et al. 2025; Wang et al. 2025). However, their exploration into geometry-focused foundation models remains limited.

Camera-Based Semantic Scene Completion

MonoScene (Cao and De Charette 2022) is a pioneering effort that extends the SSC task to cameras by employing a Features Line of Sight Projection (FLoSP) for projecting 2D features into 3D space. TPVFormer (Huang et al. 2023) introduces an efficient tri-perceptive view representation, created from image features. VoxFormer (Li et al. 2023a) employs a sparse-to-dense strategy, which propagates image features to dense voxels through deformable attention mechanism. OccFormer (Zhang, Zhu, and Du 2023) builds upon the Lift-Splat-Shoot (LSS) paradigm to facilitate feature lifting using both local and global transformer pathways.

Subsequent works have proposed more specialised solutions, often leveraging temporal cues from multiple frames (Li et al. 2024a; Lee et al. 2025), utilising more accurate geometric information obtained from stereo matching (Li et al. 2024b; Xue et al. 2024; Yu et al. 2024), exploiting spatial biases (Wang and Tong 2024; Bae, Ha, and Kim 2025), and integrating richer semantic knowledge derived from instance-level features (Jiang et al. 2024) or vision-language models (Wang et al. 2025; Cui et al. 2025). However, despite the diversity of these approaches, they are often limited by a tightly coupled architecture, which results in a trade-off between semantic and geometric features.

Our proposed FoundationSSC disentangles semantic and geometric features at both the source and processing pathway levels, effectively resolving the trade-off that has constrained earlier work.

Methodology

As illustrated in Fig. 2, FoundationSSC starts with a Foundation Encoder that achieves source-level decoupling, producing distinct semantic features and geometric cost volume features. These priors are then enhanced through Decoupled Semantic and Geometric Pathways: the semantic features

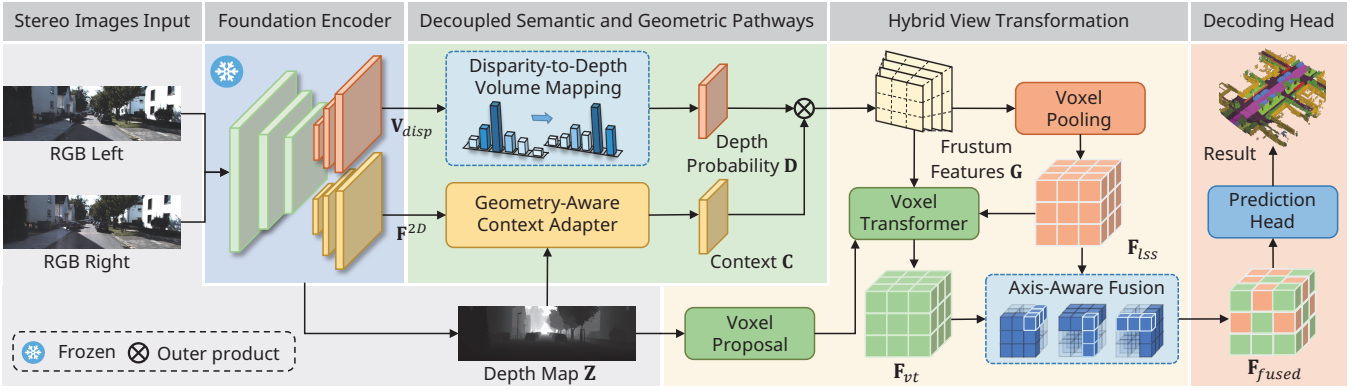


Figure 2: Overview of our proposed FoundationSSC framework. It begins with a Foundation Encoder producing decoupled priors, which are then enhanced in Decoupled Semantic Geometric Pathways. A Hybrid View Transformation subsequently lifts and fuses these priors into a unified 3D volume, which is processed by a Decoding Head to yield the final prediction.

are integrated with 3D structural awareness, while the cost volumes are transformed into rich probabilistic depth distributions. Following this, a Hybrid View Transformation elevates these refined 2D representations into two complementary 3D feature volumes. Our proposed Axis-Aware Fusion module then combines these volumes anisotropically into a unified 3D representation, which is ultimately passed to a decoding head to generate the final prediction.

Foundation Encoder

To address the balance between semantic and geometric information at the feature source, we utilise the pre-trained, frozen FoundationStereo model (Wen et al. 2025) as an integrated foundation encoder. This model effectively provides reliable, separate semantic and geometric priors. This choice is particularly advantageous because FoundationStereo inherits valuable semantic features from its DINOv2 lineage (Oquab et al. 2024; Yang et al. 2024b) and also offers strong geometric reasoning based on stereo input. Consequently, we have a robust encoder that generates intrinsically separated semantic and geometric features, displaying impressive zero-shot generalisation without the need for task-specific fine-tuning. From this frozen encoder, we extract three essential outputs for our decoupled pathways:

- **Monocular image features F^{2D} :** Multi-scale features from the frozen DepthAnythingV2 backbone are fused via a lightweight Feature Pyramid Network (FPN) to produce a single-scale, robust 2D feature map, which serves as the primary input for the semantic context branch.
- **Cost volume features V_{disp} :** The disparity probability distribution is generated by FoundationStereo’s attentive hybrid cost filtering module. The cost volume features provide rich geometric priors for our depth branch, avoiding the information loss of collapsed depth maps.
- **Dense depth map Z :** The final disparity map predicted by FoundationStereo, which is then converted to a metric depth map using camera intrinsics. This is primarily used for generating voxel proposals in the view transformation and geometric prior for the semantic context branch.

Decoupled Semantic and Geometric Pathways

To fully unleash the potential of the decoupled priors from our foundation encoder, we introduce decoupled semantic and geometric pathways. Each pathway is enhanced by a specialised module that is specifically designed to address its unique challenges and maximise its contributions.

Geometry-Aware Context Adapter. While VFM features exhibit impressive semantic capabilities, their 2D-centric pre-training limits understanding of 3D geometry (El Banani et al. 2024), leading to geometric inconsistencies when applied to 3D perception tasks. One possible solution is to fine-tune these models using 3D data (Yue et al. 2024); however, this approach can involve significant computational costs. To address this issue without expensive fine-tuning, we introduce the Geometry-Aware Context Adapter (GCA), a lightweight module designed to incorporate 3D structural awareness into the existing VFM priors.

The GCA begins by constructing a comprehensive geometry prior, M^g , from the depth map Z . This prior matrix fuses two distinct relationships between all pixel pairs: a depth relationship matrix, M^d , which encodes their metric distance in 3D space, and a spatial distance matrix, M^s , which represents their Manhattan distance on the 2D image plane. These matrices are balanced by a learnable parameter $\alpha \in [0, 1]$:

$$\begin{aligned} M_{i,j,i',j'}^d &= |Z_{ij} - Z_{i'j'}| \\ M_{i,j,i',j'}^s &= |i - i'| + |j - j'| \\ M^g &= \alpha M^d + (1 - \alpha) M^s \end{aligned} \quad (1)$$

where matrices are of shape $HW \times HW$ (H and W denote the height and width of the feature map). This matrix M^g serves as a structural bias that is integrated directly into the geometry self-attention mechanism.

Specifically, the standard attention map is multiplied element-wise by a term derived from M^g . This operation effectively directs the attention to concentrate on geometrically coherent pixel relationships. It achieves this by reducing the connections between pixels that are close in 2D but

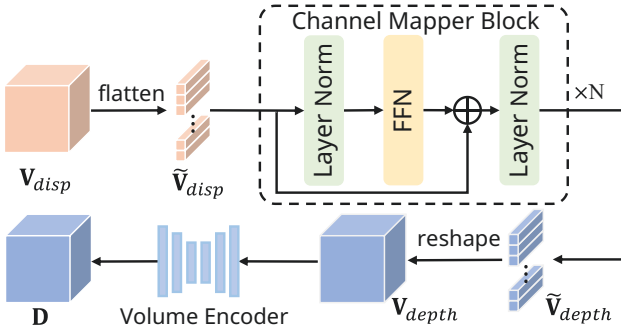


Figure 3: Illustration of the DDVM module, which transforms disparity cost volume to depth distribution.

distant in 3D. The operation can be expressed as:

$$\text{GeoAttn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}, \mathbf{M}^g) = (\text{Softmax}(\mathbf{Q}\mathbf{K}^T) \odot \beta^{\mathbf{M}^g}) \mathbf{V} \quad (2)$$

where \odot denotes element-wise multiplication and $\beta \in (0, 1)$ is a learnable decay rate. To achieve both representational diversity and computational efficiency, we follow the approach outlined in (Yin et al. 2025) by assigning distinct decay rates to different attention heads. Additionally, we implement the attention mechanism using an axes decomposition strategy. By incorporating explicit 3D structure into the attention modulation, GCA enhances the 2D VFM features \mathbf{F}^{2D} , improving their geometric consistency. These refined features then serve as enhanced context features \mathbf{C} for the process of lifting into the 3D voxel space.

Disparity-to-Depth Volume Mapping. The stereo cost volume, \mathbf{V}_{disp} , contains rich, probabilistic information about geometric correspondences, offering a far more detailed representation of 3D structure than a single depth map. Unlocking this potential is a core objective of our framework. However, its disparity-aligned structure is incompatible with the depth-binned format required by the LSS pipeline. The conventional solution is to first collapse the cost volume into a deterministic depth map and then build a one-hot depth distribution. This process acts as a significant information bottleneck by irretrievably discarding crucial details about matching ambiguities and uncertainty.

To address the issue of information loss, we propose the Disparity-to-Depth Volume Mapping (DDVM) module. This lightweight network learns a direct, non-linear mapping from a disparity-based volume to a depth-based volume. As illustrated in Fig. 3, the module transforms features from the disparity dimension into the depth dimension.

Given the cost volume features $\mathbf{V}_{disp} \in \mathbb{R}^{D_{disp} \times H \times W}$, the spatial dimensions are flattened to reshape the representation into $\tilde{\mathbf{V}}_{disp} \in \mathbb{R}^{D_{disp} \times HW}$. While $\mathbb{R}^{D_{disp} \times H \times W}$ preserves the 2D spatial layout, $\mathbb{R}^{D_{disp} \times HW}$ denotes its flattened form. A non-linear mapping function, $f(\cdot)$, transforms the feature channels directly from disparity to depth:

$$\tilde{\mathbf{V}}_{depth} = f(\tilde{\mathbf{V}}_{disp}) \quad (3)$$

where the output $\tilde{\mathbf{V}}_{depth}$ has the shape $\mathbb{R}^{D_{depth} \times HW}$. The mapping function $f(\cdot)$ comprises stacked channel-mapper

blocks, each consisting of a feed-forward network (FFN), layer normalisation, and a residual connection. This design enables the network to learn complex, non-linear correspondences between disparity and depth channels.

To restore spatial structure, $\tilde{\mathbf{V}}_{depth}$ is reshaped into 3D volumes $\mathbf{V}_{depth} \in \mathbb{R}^{D_{depth} \times H \times W}$. A refinement module then employs a shallow 3D CNN encoder $\phi(\cdot)$ to enhance spatial coherence and smoothness, followed by a softmax operation along the depth dimension to produce the probabilistic depth distribution, $\mathbf{D} \in \mathbb{R}^{D_{depth} \times H \times W}$:

$$\mathbf{D} = \text{Softmax}(\phi(\mathbf{V}_{depth})) \quad (4)$$

Unlike methods that depend on explicit analytical conversion, our learned mapping is more flexible and robust. It retains the detailed probabilistic characteristics of the original cost volume, offering a much richer geometric prior for the subsequent view transformation.

Hybrid View Transformation

We follow (Yu et al. 2024) for 2D-to-3D lifting, which integrates an LSS-based voxel initialisation with transformer-based voxel refinement. In the LSS view transformation, the context features $\mathbf{C} \in \mathbb{R}^{C \times H \times W}$ and the categorical depth distributions $\mathbf{D} \in \mathbb{R}^{D \times H \times W}$ are combined via the outer product $\mathbf{C} \otimes \mathbf{D}$ to construct the 3D frustum representation $\mathbf{G} \in \mathbb{R}^{C \times D \times H \times W}$. After this, voxel pooling (Phillion and Fidler 2020) is applied to splat frustum features \mathbf{G} into a 3D voxel grid, resulting in the 3D feature representation $\mathbf{F}_{lss} \in \mathbb{R}^{C \times X \times Y \times Z}$, where (X, Y, Z) denotes the spatial resolution of the voxel grid. In the voxel transformer, the LSS volume and a depth map are utilised to generate query proposals \mathbf{F}_q , where the depth map first identifies a sparse set of occupied voxel locations, and the LSS volume provides their initial features. Deformable cross-attention is employed to gather relevant 3D frustum features for each proposal, followed by deformable self-attention among voxels to produce the final 3D feature representation $\mathbf{F}_{vt} \in \mathbb{R}^{C \times X \times Y \times Z}$.

Inspired by (Li et al. 2023b, 2024c), we hypothesise that the two 3D feature volumes contain complementary information: \mathbf{F}_{lss} preserves precise geometric details for near-field objects, while \mathbf{F}_{vt} captures superior contextual information for distant and occluded regions. However, previous methods, such as those proposed by (Yu et al. 2024), often overlook the fusion of these features, treating \mathbf{F}_{lss} as a disposable intermediate representation and thereby risking the loss of valuable information. Effectively fusing these complementary representations is a critical challenge that is often neglected, and we will address this issue next.

Axis-Aware Fusion. Our initial fusion exploration started with a vanilla 3D channel attention method inspired by 2D techniques (Dai et al. 2021). Although this method demonstrated some performance improvements, we believe its effectiveness is constrained by the isotropic nature of its fusion mechanisms. Specifically, the 3D global pooling operation compresses the entire volume into a single descriptor, which results in the loss of crucial directional cues.

To enable spatially aware 3D feature fusion, we introduce the Axis-Aware Fusion (AAF) module. The core concept behind AAF is to independently capture and merge features




Figure 4: Illustration of the fusion unit for a single axis (e.g., the Z -axis) within the proposed AAF module.

based on contextual information along each of the three primary axes (X , Y , Z). To accomplish this, the AAF module utilises three parallel fusion units, with each unit dedicated to one axis and operating on the orthogonal plane to extract axis-specific context. For example, the unit for the Z -axis works on the XY -plane.

As illustrated in Fig. 4, within each unit, a dynamic, axis-specific attention map, denoted as σ_d , is created by synthesising both local and global information. Specifically, the input features (F_{lss} and F_{vt}), are processed through two simultaneous pathways: the local pathway employs 3D convolutions to capture fine-grained spatial details, while the global pathway utilises our anisotropic pooling strategy (pooling along the two axes of the plane) to extract a global context feature specific to the orthogonal axis.

The features from both pathways are then fused to generate the attention map σ_d . The final fused feature for the entire module, F_{fused} , is obtained by summing the outputs from the three parallel, axis-specific fusion units:

$$F_{fused} = \sum_{d \in XY, XZ, YZ} (\sigma_d F_{lss} + (1 - \sigma_d) F_{vt}) \quad . \quad (5)$$

Eq. 5 leverages both local and global information specific to each axis. As a result, the AAF module effectively merges complementary features in a highly adaptive and spatially aware way. The fusion operation produces a unified 3D feature representation, enhanced by a thorough understanding of the scene’s structure along all three orthogonal axes.

Training Loss

To train the FoundationSSC, we follow MonoScene (Cao and De Charette 2022), utilise affinity losses \mathcal{L}_{scal}^{geo} and \mathcal{L}_{scal}^{sem} to optimise the scene-wise and class-wise metrics, in conjunction with the cross-entropy loss \mathcal{L}_{ce} weighted by class frequencies. Additionally, we employ two auxiliary losses: a 2D semantic segmentation loss \mathcal{L}_s for feature enrichment and a depth loss \mathcal{L}_d (Li et al. 2024e) for geometric supervision. The total loss is a weighted sum of these components:

$$\mathcal{L} = \lambda_d \mathcal{L}_d + \lambda_s \mathcal{L}_s + \mathcal{L}_{ce} + \mathcal{L}_{scal}^{geo} + \mathcal{L}_{scal}^{sem} \quad . \quad (6)$$

where λ_d and λ_s are the weights for the auxiliary losses. Following preliminary experiments, we set $\lambda_d = 0.001$ and $\lambda_s = 1$, respectively.

Experiment

We evaluate our FoundationSSC on SemanticKITTI (Behley et al. 2019) and SSCBench-KITTI-360 (Li et al. 2024d). Following previous methods (Cao and De Charette 2022; Yu et al. 2024), we report the IoU and mIoU as evaluation metrics. Refer to the supplementary material for more information about datasets, metrics, and implementation details.

Quantitative Results

We present the quantitative results on SemanticKITTI hidden test set. As summarised in Table 1, FoundationSSC achieves state-of-the-art performance with a mIoU of 19.32 and an IoU of 48.12, demonstrating a +2.69 increase in mIoU and a +3.71 increase in IoU compared to the CGFormer baseline (Yu et al. 2024). The simultaneous and substantial gains across both metrics provide direct evidence that the dual-decoupling design effectively balances the inherent trade-off of coupled architectures. Furthermore, despite relying solely on stereo images as input, FoundationSSC outperforms all temporal methods, including HTCL (Li et al. 2024a) and SOAP (Lee et al. 2025). On a per-class level, our model achieves the best or second-best performance in most categories. Benefiting from the strong feature priors provided by the foundation encoder, our model demonstrates particularly robust results in long-tail classes, such as bicycle and motorcycle.

We further report results on SSCBench-KITTI-360 in Table 2, where FoundationSSC once again achieves top performance with 48.61 IoU and a 21.78 mIoU. The consistent superiority across two benchmarks underscores the robustness and generalizability of our proposed framework.

Ablation Study

We conduct ablation studies on SemanticKITTI validation set to evaluate the effectiveness of each component.

Ablation on the Architectural Components. Table 3 analyses the contribution of each component in FoundationSSC. Our baseline is a simplified CGFormer variant that excludes the TPV branch. First, we assess the impact of source-level decoupling (a). Replacing the baseline’s encoder with our Foundation Encoder (FE) results in substantial gains across both metrics (+2.06 mIoU, +1.33 IoU), validating the effectiveness of using high-quality, disentangled priors. Next, we ablate the components of our pathway-level decoupling based on the FE-only model. The GCA module (b) enhances both semantic and geometric features (+0.36 mIoU, +0.56 IoU) by injecting 3D awareness. The DDVM module (c) provides an even larger, geometry-focused boost (+0.52 mIoU, +1.32 IoU) by preserving probabilistic information. Moreover, combining both modules (d) results in a synergistic improvement (+0.97 mIoU, +1.23 IoU) over the FE-only model, demonstrating that our specialised pathways work together effectively. Lastly, adding the AAF module (e) to the full pathway model (d) contributes to further performance enhancements (+0.80 mIoU, +0.07 IoU), verifying our anisotropic fusion strategy. In total, our whole model demonstrates remarkable gains of +3.83 mIoU and +2.63 IoU over the baseline, confirming that each component of

Method	IoU	mIoU	road	sidewalk	parking	other-grnd.	building	car	truck	bicycle	motorcycle	other-veh.	vegetation	trunk	terrain	person	bicyclist	motorcyclist	fence	pole	traf.-sign
MonoScene	34.16	11.08	54.7	27.1	24.8	5.7	14.4	18.8	3.3	0.5	0.7	4.4	14.9	2.4	19.5	1.0	1.4	0.4	11.1	3.3	2.1
TPVFormer	34.25	11.26	55.1	27.2	27.4	6.5	14.8	19.2	3.7	1.0	0.5	2.3	13.9	2.6	20.4	1.1	2.4	0.3	11.0	2.9	1.5
VoxFormer	43.21	13.41	54.1	26.9	25.1	7.3	23.5	21.7	3.6	1.9	1.6	4.1	24.4	8.1	24.2	1.6	1.1	0.0	13.1	6.6	5.7
OccFormer	34.53	12.32	55.9	30.3	31.5	6.5	15.7	21.6	1.2	1.5	1.7	3.2	16.8	3.9	21.3	2.2	1.1	0.2	11.9	3.8	3.7
Symphonies	42.19	15.04	58.4	29.3	26.9	11.7	24.7	23.6	3.2	3.6	2.6	5.6	24.2	10.0	23.1	<u>3.2</u>	1.9	<u>2.0</u>	16.1	7.7	8.0
BRGScene	43.34	15.36	61.9	31.2	30.7	10.7	24.2	22.8	2.8	3.4	2.4	6.1	23.8	8.4	27.0	2.9	2.2	0.5	16.5	7.0	7.2
HTCL	44.23	17.09	64.4	34.8	33.8	12.4	25.9	27.3	5.7	1.8	2.2	5.4	25.3	10.8	31.2	1.1	3.1	0.9	21.1	9.0	8.3
CGFormer	44.41	16.63	64.3	34.2	34.1	12.1	25.8	26.1	4.3	3.7	1.3	2.7	24.5	11.2	29.3	1.7	3.6	0.4	18.7	8.7	9.3
VLScene	45.14	17.52	64.7	34.7	32.4	<u>13.1</u>	27.3	26.1	<u>6.5</u>	4.2	3.8	<u>8.3</u>	26.4	10.0	29.4	2.8	<u>5.1</u>	0.9	20.0	8.9	8.4
ScanSSC	44.54	17.40	66.2	35.9	<u>35.1</u>	12.5	25.3	<u>27.1</u>	3.5	3.5	3.2	6.1	25.2	11.0	30.6	1.8	5.3	0.7	20.5	8.4	8.9
SOAP	46.09	19.09	63.6	<u>36.2</u>	36.8	17.2	<u>28.7</u>	28.9	3.3	<u>5.0</u>	<u>5.3</u>	7.0	29.8	15.1	<u>32.3</u>	3.7	2.3	0.2	<u>22.5</u>	11.7	13.1
Ours	48.12	19.32	<u>65.7</u>	36.3	33.5	11.6	30.5	<u>28.2</u>	9.0	5.6	6.9	10.0	<u>28.6</u>	<u>13.7</u>	32.8	2.8	4.7	2.5	<u>23.1</u>	<u>9.8</u>	<u>11.8</u>

Table 1: Quantitative results on SemanticKITTI test set. Best and second-best results are in **bold** and underlined.

Method	IoU	mIoU	car	bicycle	motorcycle	truck	other-veh.	person	road	parking	sidewalk	other-grnd.	building	fence	vegetation	terrain	pole	traf.-sign	other-struct.	other-obj.
MonoScene	37.87	12.31	19.3	0.4	0.6	8.0	2.0	0.9	48.4	11.4	28.1	3.3	32.9	3.5	26.2	16.8	6.9	5.7	4.2	3.1
VoxFormer	38.76	11.91	17.8	1.2	0.9	4.6	2.1	1.6	47.0	9.7	27.2	2.9	31.2	5.0	29.0	14.7	6.5	6.9	3.8	2.4
TPVFormer	40.22	13.64	21.6	1.1	1.4	8.1	2.6	2.4	53.0	12.0	31.1	3.8	34.8	4.8	30.1	17.5	7.5	5.9	5.5	2.7
OccFormer	40.27	13.81	22.6	0.7	0.3	9.9	3.8	2.8	54.3	13.4	31.5	3.6	36.4	4.8	31.0	19.5	7.8	8.5	7.0	4.6
Symphonies	44.12	18.58	<u>30.0</u>	1.9	5.9	25.1	<u>12.1</u>	8.2	54.9	13.8	32.8	6.9	35.1	8.6	38.3	11.5	14.0	9.6	14.4	11.3
CGFormer	48.07	20.05	29.9	3.4	4.0	17.6	6.8	6.6	63.9	17.2	40.7	5.5	42.7	8.2	38.8	<u>24.9</u>	16.2	17.5	10.2	6.8
VLScene	46.08	19.10	29.0	4.7	7.7	18.3	7.6	<u>7.4</u>	60.1	17.4	39.0	<u>6.0</u>	42.1	9.6	36.5	<u>24.8</u>	17.0	18.8	10.5	6.5
ScanSSC	<u>48.29</u>	20.14	29.9	3.8	4.3	14.3	9.1	6.7	62.2	18.2	40.2	5.2	42.7	8.8	38.8	25.5	16.6	19.1	10.3	6.9
SOAP	<u>48.12</u>	<u>20.92</u>	29.9	5.6	<u>7.8</u>	14.4	7.6	6.1	60.9	17.4	<u>40.3</u>	5.4	45.3	10.6	40.5	24.8	<u>16.8</u>	21.0	<u>12.6</u>	9.9
Ours	48.61	21.78	30.5	<u>4.9</u>	8.4	<u>24.9</u>	13.0	7.2	<u>63.6</u>	<u>17.8</u>	39.8	5.7	<u>43.6</u>	<u>10.5</u>	<u>39.1</u>	24.8	16.2	<u>20.3</u>	11.5	<u>10.2</u>

Table 2: Quantitative results on SSCBench-KITTI-360 test set. Best and second-best results are in **bold** and underlined.

Method	FE	GCA	DDVM	AAF	IoU	mIoU
Baseline					45.28	16.53
(a)	✓				46.61	18.59
(b)	✓	✓			47.17	18.95
(c)	✓		✓		47.93	19.11
(d)	✓	✓	✓		47.84	19.56
(e)	✓	✓	✓	✓	47.91	20.36

Table 3: Ablation study for architectural components.

our dual-decoupling design effectively contributes to both semantic and geometric understanding.

Ablation Study for Foundation Encoder. To evaluate the impact of different VFMs, we conduct an ablation study where each model derives context features and monocular depth distributions from shared internal representations, enabling a fair comparison of semantic-geometric trade-offs. As shown in Table 4, all VFM variants outperform the EfficientNet-B7 (Tan and Le 2019) baseline, demonstrating superior representational capacity. Among

Backbone Networks	IoU	mIoU	Parameters
EfficientNetB7	44.55	15.20	66.35M
DINOv2-S	45.33	15.52	22.06M
DepthAnythingV2-S	45.31	15.59	24.79M
DINOv2-L	45.88	17.79	304.37M
DepthAnythingV2-L	46.19	17.01	335.32M

Table 4: Ablation study on the image backbone.

them, DINOv2-L (Oquab et al. 2024) achieves the highest mIoU, while the depth-specialised DepthAnythingV2-L (Yang et al. 2024b) attains higher IoU at the expense of a slightly lower mIoU. The performance contrast highlights the specialisation-generalisation trade-off and supports a decoupled design philosophy. Although pairing separate backbones (e.g. DINOv2 for semantics and another for geometry) is feasible, such configurations increase architecture complexity. Instead, we adopt the frozen DepthAnythingV2-L within the unified FoundationStereo encoder, ensuring both simplicity and efficiency.

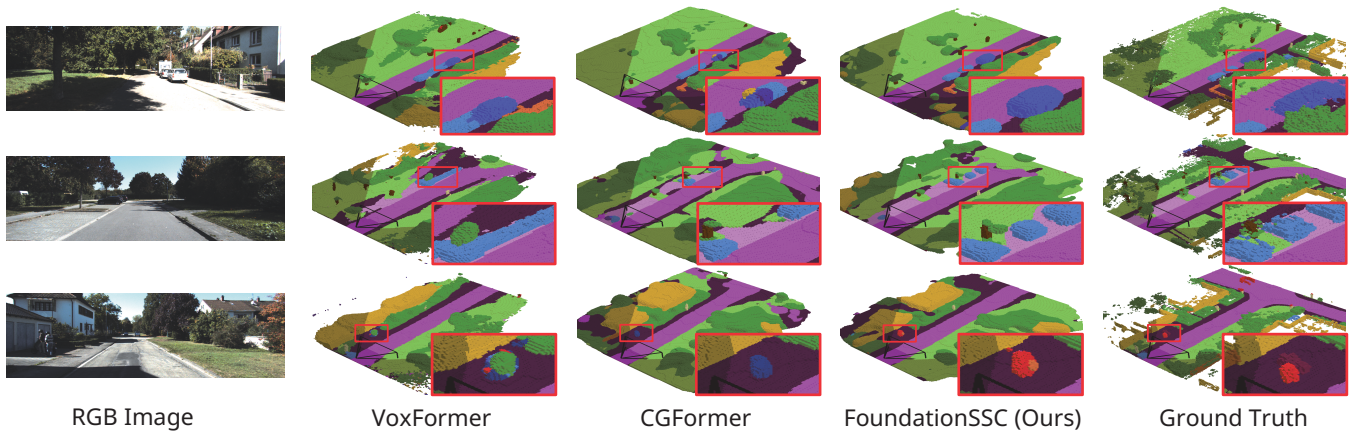


Figure 5: Qualitative visualisation results on SemanticKITTI validation set. Our FoundationSSC predicts objects with sharper geometric boundaries and more accurate semantic labels, outperforming competing works.

Setting	IoU	mIoU
Depth Refinement	47.87	19.83
Cost Volume + AR	47.76	19.59
Cost Volume + DDVM	47.91	20.36

Table 5: Ablation study on the depth generation strategies.

Setting	IoU	mIoU
w/o AAF	47.84	19.56
3D CA	48.25	20.08
AAF	47.91	20.36

Table 6: Ablation study on the 3D feature fusion strategy.

Ablation Study for DDVM. We validate our DDVM module against various depth generation strategies in Table 5. The CGFormer (Yu et al. 2024) depth refinement module, which fuses monocular distribution with a pre-processed stereo depth map using neighbour attention, achieves 19.83 mIoU and 47.87 IoU. A key limitation of such a method is its reliance on information-lossy depth maps, which discard valuable probabilistic cues from stereo matching. We then explore the direct use of the disparity cost volume. A naive approach using analytical resampling (AR) degrades performance to 19.59 mIoU, confirming that merely accessing the cost volume is insufficient. In contrast, our DDVM, by replacing rigid resampling with a learnable mapping, achieves a state-of-the-art mIoU of 20.36. The performance gain demonstrates DDVM’s superior ability to preserve and leverage fine-grained geometric cues for SSC.

Ablation Study for AAF. Table 6 compares the AAF module against alternative 3D feature fusion strategies. We establish a baseline by using only the refined volume \mathbf{F}_{vt} , mirroring prior work, which yields 19.56 mIoU. A fusion approach using a vanilla 3D channel attention (3D-CA) module improves performance to 20.08 mIoU. However, the effectiveness of 3D-CA is constrained by an isotropic global

pooling, a design that neglects directional cues critical for 3D understanding. In contrast, our AAF module is explicitly designed to be anisotropic. The module’s axis-specific fusion mechanism achieves the highest mIoU of 20.36. The substantial performance gain over the 3D-CA provides compelling evidence that modelling the anisotropic nature of 3D scenes is crucial for effective feature fusion.

Qualitative Results

Fig. 5 presents a qualitative comparison on SemanticKITTI validation set, showcasing the results of our FoundationSSC, VoxFormer (Li et al. 2023a), CGFormer (Yu et al. 2024), and the ground truth. The visual results highlight the dual advantages of our framework in both geometry and semantics. Our geometric superiority is evident with distant, partially occluded objects. While competing methods produce blurry, merged voxels, FoundationSSC renders them with crisp, well-defined boundaries. Our semantic advantage is showcased in identifying smaller or less common classes. While competing methods may misclassify objects, such as pedestrians, FoundationSSC accurately predicts their semantic labels. These visual advantages are a direct outcome of our dual-decoupling strategy, enabling the framework to leverage and process robust, disentangled priors effectively.

Conclusion

In this paper, we introduce FoundationSSC, a novel neural network designed to resolve the semantic-geometric conflict in camera-based SSC through a dual-decoupling approach. FoundationSSC features an integrated foundation encoder that provides robust and disentangled semantic and geometric priors. These priors are refined through specialised decoupled pathways and then combined using our innovative Axis-Aware Fusion module. The effectiveness of this strategy is demonstrated by achieving state-of-the-art results on SemanticKITTI and SSCBench-KITTI-360 datasets. The substantial improvements in both geometric (IoU) and semantic (mIoU) metrics confirm that our framework can properly balance the long-standing performance trade-off.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (62473276, 62573309), in part by the Natural Science Foundation of Jiangsu Province (BK20241918), and in part by the Research Fund of Horizon Robotics (H230666).

References

- Bae, J.; Ha, J.; and Kim, H. Y. 2025. Three Cars Approaching within 100m! Enhancing Distant Geometry by Tri-Axis Voxel Scanning for Camera-based Semantic Scene Completion. In *IEEE Conference on Computer Vision and Pattern Recognition*, 11939–11948.
- 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 *IEEE International Conference on Computer Vision*, 9297–9307.
- Cao, A.-Q.; and De Charette, R. 2022. Monoscene: Monocular 3d semantic scene completion. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3991–4001.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *IEEE International Conference on Computer Vision*, 9650–9660.
- Cui, Y.; Li, Z.; Wang, J.; and Fang, Z. 2025. LOMA: Language-assisted semantic occupancy network via triplane mamba. In *AAAI Conference on Artificial Intelligence*, volume 39, 2609–2617.
- Dai, Y.; Gieseke, F.; Oehmcke, S.; Wu, Y.; and Barnard, K. 2021. Attentional feature fusion. In *IEEE Winter Conference on Applications of Computer Vision*, 3560–3569.
- El Banani, M.; Raj, A.; Maninis, K.-K.; Kar, A.; Li, Y.; Rubinstein, M.; Sun, D.; Guibas, L.; Johnson, J.; and Jampani, V. 2024. Probing the 3d awareness of visual foundation models. In *IEEE Conference on Computer Vision and Pattern Recognition*, 21795–21806.
- Huang, Y.; Zheng, W.; Zhang, Y.; Zhou, J.; and Lu, J. 2023. Tri-perspective view for vision-based 3d semantic occupancy prediction. In *IEEE 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 *IEEE Conference on Computer Vision and Pattern Recognition*, 20258–20267.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *IEEE International Conference on Computer Vision*, 4015–4026.
- Lee, H.-J.; Koh, Y. J.; Kim, H.; Kim, H.; Lee, Y.; and Lee, J. 2025. SOAP: Vision-Centric 3D Semantic Scene Completion with Scene-Adaptive Decoder and Occluded Region-Aware View Projection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 17145–17154.
- Li, B.; Deng, J.; Zhang, W.; Liang, Z.; Du, D.; Jin, X.; and Zeng, W. 2024a. Hierarchical temporal context learning for camera-based semantic scene completion. In *European Conference on Computer Vision*, 131–148. Springer.
- Li, B.; Sun, Y.; Liang, Z.; Du, D.; Zhang, Z.; Wang, X.; Wang, Y.; Jin, X.; and Zeng, W. 2024b. Bridging Stereo Geometry and BEV Representation with Reliable Mutual Interaction for Semantic Scene Completion. In *International Joint Conference on Artificial Intelligence*, 965–973.
- Li, P.; Shen, W.; Huang, Q.; and Cui, D. 2024c. DualBEV: Unifying Dual View Transformation with Probabilistic Correspondences. In *European Conference on Computer Vision*, 286–302. Springer.
- Li, Y.; Li, S.; Liu, X.; Gong, M.; Li, K.; Chen, N.; Wang, Z.; Li, Z.; Jiang, T.; Yu, F.; et al. 2024d. Ssbench: A large-scale 3d semantic scene completion benchmark for autonomous driving. In *IEEE International Conference on Intelligent Robots and Systems*, 13333–13340. IEEE.
- Li, Y.; Yu, Z.; Choy, C.; Xiao, C.; Alvarez, J. M.; Fidler, S.; Feng, C.; and Anandkumar, A. 2023a. Voxformer: Sparse voxel transformer for camera-based 3d semantic scene completion. In *IEEE Conference on Computer Vision and Pattern Recognition*, 9087–9098.
- Li, Z.; Wang, W.; Li, H.; Xie, E.; Sima, C.; Lu, T.; Yu, Q.; and Dai, J. 2024e. Bevformer: learning bird’s-eye-view representation from lidar-camera via spatiotemporal transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Li, Z.; Yu, Z.; Wang, W.; Anandkumar, A.; Lu, T.; and Alvarez, J. M. 2023b. Fb-bev: Bev representation from forward-backward view transformations. In *IEEE International Conference on Computer Vision*, 6919–6928.
- Liao, Y.; Xie, J.; and Geiger, A. 2022. Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 3292–3310.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2024. DINOv2: Learning Robust Visual Features without Supervision. *Transactions on Machine Learning Research Journal*, 1–31.
- Phillion, J.; and Fidler, S. 2020. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *European Conference on Computer Vision*, 194–210. Springer.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 8748–8763. PMLR.
- Ravi, N.; Gabeur, V.; Hu, Y.-T.; Hu, R.; Ryali, C.; Ma, T.; Khedr, H.; Rädle, R.; Rolland, C.; Gustafson, L.; et al. 2025. Sam 2: Segment anything in images and videos. In *International Conference on Learning Representations*.
- Roldao, L.; De Charette, R.; and Verroust-Blondet, A. 2020. Lmscnet: Lightweight multiscale 3d semantic completion. In *International Conference on 3D Vision*, 111–119. IEEE.

- Song, S.; Yu, F.; Zeng, A.; Chang, A. X.; Savva, M.; and Funkhouser, T. 2017. Semantic scene completion from a single depth image. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1746–1754.
- Tan, M.; and Le, Q. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, 6105–6114. PMLR.
- Wang, M.; Pi, H.; Li, R.; Qin, Y.; Tang, Z.; and Li, K. 2025. VLScene: Vision-Language Guidance Distillation for Camera-Based 3D Semantic Scene Completion. In *AAAI Conference on Artificial Intelligence*, volume 39, 7808–7816.
- Wang, Y.; and Tong, C. 2024. H2gformer: Horizontal-to-global voxel transformer for 3d semantic scene completion. In *AAAI Conference on Artificial Intelligence*, volume 38, 5722–5730.
- Wen, B.; Trepte, M.; Aribido, J.; Kautz, J.; Gallo, O.; and Birchfield, S. 2025. FoundationStereo: Zero-Shot Stereo Matching. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Wu, Y.; Zheng, W.; Zuo, S.; Huang, Y.; Zhou, J.; and Lu, J. 2025. Embodiedocc: Embodied 3d occupancy prediction for vision-based online scene understanding. In *IEEE International Conference on Computer Vision*.
- 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 *IEEE Conference on Computer Vision and Pattern Recognition*, 20124–20134.
- Yang, L.; Kang, B.; Huang, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024a. Depth anything: Unleashing the power of large-scale unlabeled data. In *IEEE Conference on Computer Vision and Pattern Recognition*, 10371–10381.
- Yang, L.; Kang, B.; Huang, Z.; Zhao, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024b. Depth anything v2. *Advances in Neural Information Processing Systems*, 37: 21875–21911.
- Yin, B.-W.; Cao, J.-L.; Cheng, M.-M.; and Hou, Q. 2025. Dformerv2: Geometry self-attention for rgb-d semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 19345–19355.
- 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. In *Advances in Neural Information Processing Systems*, volume 37, 1531–1555.
- Yue, Y.; Das, A.; Engelmann, F.; Tang, S.; and Lenssen, J. E. 2024. Improving 2d feature representations by 3d-aware fine-tuning. In *European Conference on Computer Vision*, 57–74. Springer.
- Zhang, Y.; Zhu, Z.; and Du, D. 2023. Occformer: Dual-path transformer for vision-based 3d semantic occupancy prediction. In *IEEE International Conference on Computer Vision*, 9433–9443.
- Zheng, W.; Chen, W.; Huang, Y.; Zhang, B.; Duan, Y.; and Lu, J. 2024. Occworld: Learning a 3d occupancy world model for autonomous driving. In *European Conference on Computer Vision*, 55–72. Springer.