

ALOcc: Adaptive Lifting-Based 3D Semantic Occupancy and Cost Volume-Based Flow Predictions

Dubing Chen¹, Jin Fang¹, Wencheng Han¹, Xinjing Cheng,
 Junbo Yin², Chenzhong Xu¹, Fahad Shahbaz Khan^{3,4}, Jianbing Shen¹✉

¹SKL-IOTSC, CIS, University of Macau

²CEMSE Division, King Abdullah University of Science and Technology

³Mohamed bin Zayed University of Artificial Intelligence ⁴Linköping University

<https://github.com/cdb342/ALOcc>

Abstract

3D semantic occupancy and flow prediction are fundamental to spatiotemporal scene understanding. This paper proposes a vision-based framework with three targeted improvements. First, we introduce an occlusion-aware adaptive lifting mechanism incorporating depth denoising. This enhances the robustness of 2D-to-3D feature transformation while mitigating reliance on depth priors. Second, we enforce 3D-2D semantic consistency via jointly optimized prototypes, using confidence- and category-aware sampling to address the long-tail classes problem. Third, to streamline joint prediction, we devise a BEV-centric cost volume to explicitly correlate semantic and flow features, supervised by a hybrid classification-regression scheme that handles diverse motion scales. Our purely convolutional architecture establishes new SOTA performance on multiple benchmarks for both semantic occupancy and joint occupancy semantic-flow prediction. We also present a family of models offering a spectrum of efficiency-performance trade-offs. Our real-time version exceeds all existing real-time methods in speed and accuracy, ensuring its practical viability.

1. Introduction

3D occupancy prediction is a central task in modern 3D vision, focused on converting raw camera imagery into a structured, volumetric representation of the scene [15, 38, 47, 48, 50, 55, 56]. This technique produces a dense 3D grid

✉ Corresponding author: *Jianbing Shen*. This work was supported in part by the Science and Technology Development Fund of Macau SAR (FDCT) under grants 0102/2023/RIA2 and 0154/2022/A3 and 001/2024/SKL and CG2025-IOTSC, the University of Macau SRG2022-00023-IOTSC grant, and the Jiangyin Hi-tech Industrial Development Zone under the Taihu Innovation Scheme (EF2025-00003-SKL-IOTSC).

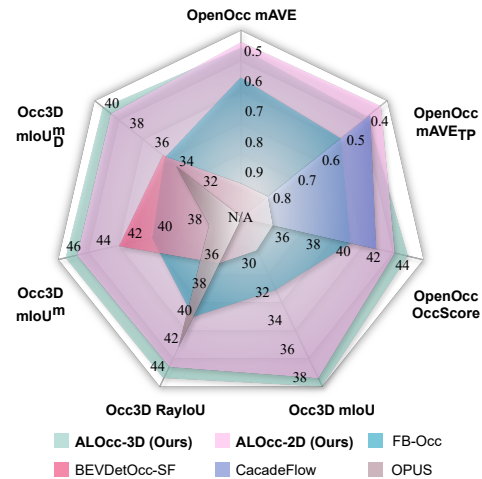


Figure 1. **3D semantic occupancy and flow prediction results of ALOcc vs. SOTAs.** ALOcc outperforms SOTAs across various benchmarks and metrics. For a fair comparison, the input image size and backbone are standardized to 256×704 and ResNet50, with per-metric comparisons conducted under identical training conditions. Since existing methods often focus on single metrics, “N/A” denotes unavailable values; we reimplemented FB-Occ and BEVDetOcc-SF to assess all metrics.

where each voxel contains a combination of attributes, like occupancy state, semantic label, and motion flow. Such a canonical data structure for scene analysis offers significant advantages in completeness and detail over conventional, object-level descriptions like bounding boxes [14, 22, 24].

The standard pipeline for this task, which typically involves 2D-to-3D view transformation and subsequent 3D feature decoding, faces significant challenges across its different stages. In view transformation, prevailing methods like depth-based Lift-Splat-Shoot (LSS) [14, 22, 44] and cross-attention [24, 26] present a critical trade-off. LSS is constrained by a strong depth inductive bias that can lead to premature convergence, poor handling of occlusions, and

the generation of sparse features for distant objects [26]. In contrast, while cross-attention circumvents these problems, it lacks an explicit geometric foundation and performs poorly. In the decoding stage, key difficulties include robustly interpreting the reconstructed 3D volume, efficiently training on dense volumetric grids, and addressing the severe class imbalance caused by both the long-tail distribution of object classes and the predominance of *empty* voxels. Furthermore, extending these frameworks to jointly predict motion flow introduces a conflicting representational demand, as features must simultaneously encode both static semantics and dynamic motion. To address these multifaceted challenges, we propose ALOcc, a unified framework that introduces targeted innovations at each stage of the occupancy prediction pipeline.

First, to overcome the limitations in view transformation, ALOcc employs an occlusion-aware adaptive lifting method. Inspired by the human cognitive ability to infer complete shapes from partial observations, this method explicitly propagates depth probabilities from visible surfaces to occluded and sparse regions, yielding a more complete and robust 3D representation. To further mitigate premature convergence from inaccurate early depth estimates, the process is stabilized by a depth denoising technique leveraging ground-truth depth data.

Next, to enhance the fidelity of the 3D decoding process, ALOcc introduces a sophisticated semantic enhancement mechanism. This mechanism correlates 3D and 2D features via a shared set of pre-defined, per-class prototypes, which transfer rich inter-class relationships from the 2D signals into the 3D representation. To address the long-tail distribution, we employ a conditional training strategy where a prototype is updated only when its class is present in the input scene. This is complemented by a hybrid sampling strategy that focuses training on hard-to-predict voxels, identified via class statistics and per-voxel prediction uncertainty. This approach significantly enhances learning efficiency.

Finally, to enable the challenging joint prediction of semantics and motion, we extend the ALOcc framework with a novel BEV cost volume-based flow head. Our core strategy here is to decouple the conflicting representational demands of this task. Specifically, we aggregate volumetric features into a compact BEV representation, from which a cost volume dedicated to motion estimation is constructed against historical features. By leveraging both this specialized cost volume and the original volumetric features for flow estimation, the primary features remain focused on occupancy semantics, thus alleviating the representational burden. A hybrid classification-regression technique is also introduced to improve adaptability to diverse motion scales.

As shown in Fig. 1, our method consistently outperforms existing approaches, establishing new SOTA performance on multiple benchmarks for both semantic occu-

pancy and joint occupancy semantic-flow prediction. Moreover, we introduce multiple variants of our ALOcc framework, employing techniques like spatial compression to achieve a superior trade-off between accuracy and computational efficiency. Our approach consistently delivers SOTA results while remaining computationally efficient, making it highly suitable for resource-constrained applications requiring real-time or on-device deployment. In summary, our main contributions are fourfold: *(i)* We introduce an occlusion-aware adaptive lifting method. This technique robustly projects 2D information into occluded and sparse 3D regions and is regularized by a depth denoising module to mitigate convergence to local optima. *(ii)* We propose shared semantic prototypes that bridge the 2D and 3D domains to enhance semantic fidelity. This approach is complemented by a conditional training and uncertainty-aware sampling scheme to address severe class imbalance. *(iii)* We introduce a novel BEV-based cost volume for occupancy flow prediction. This method alleviates the representational burden in multi-task learning and improves prediction accuracy via a hybrid classification-regression design. *(iv)* ALOcc establishes new SOTA performance on multiple benchmarks for both semantic occupancy and joint occupancy semantic-flow prediction. This superiority extends across the efficiency spectrum, with our lightweight variant outperforming all existing real-time methods, demonstrating the framework’s practical versatility.

2. Related Work

Bird’s Eye View (BEV) detection [14, 21, 24, 35] has been a cornerstone for 3D scene understanding. Foundational view transformation techniques, originally developed for BEV detection, were subsequently adapted to vision-based occupancy prediction. These notably include 2D-driven methods [14, 22, 44], which leverage depth maps [64] to project image features onto the BEV plane, and 3D-driven approaches [24, 26], which employ cross-attention mechanisms to extract spatial details from 2D images.

Semantic Scene Completion (SSC) [7, 19, 23, 31, 62] endeavors to reconstruct and semantically complete 3D scenes from given inputs. Early research focused on indoor scenes [7, 19, 31], predicting occupancy and semantic labels in limited scenarios. Subsequent studies gradually shifted attention to complex outdoor environments [1, 4]. Recent works, such as VoxFormer [23], adopt a two-stage approach that initially predicts occupancy, followed by semantic prediction of the occupied segments. OccFormer [62] introduces a dual-path Transformer submodule for 3D encoding and employs a query-based segmentation method [8].

Vision-Based 3D Occupancy Prediction [6, 15, 47, 48, 50] aligns closely with SSC but emphasizes multi-perspective image inputs. Early methods [15] pioneered this task using sparse LiDAR points for supervision. Subse-




Figure 2. **Overall Framework of ALOcc.** Image features are lifted to 3D space by adaptive lifting and encoded by a 3D encoder alongside historical features. The final representation is decoded by pluggable heads for downstream tasks (e.g., semantic occupancy, motion flow).

quently, a significant line of work [47, 48, 50, 52] advanced the field by developing denser and more accurate annotation pipelines that leverage temporal cues and instance-level labeling. Research has focused on several key technical challenges. For the critical computational efficiency demands of dense volumetric processing, methods like OctreeOcc [37], PanoOcc [51], SparseOcc [29], and OPUS [49] have introduced techniques centered on sparse 3D representations and computations. Concurrently, to mitigate the prohibitive cost of dense 3D annotation, a line of research has explored NeRF- or 3D Gaussian-based rendering to enable supervision-light learning [10, 16, 32, 42, 59]. More recently, SparseOcc [29] explored a new metric, RayIoU, to better assess prediction quality. Furthermore, the introduction of 3D occupancy flow prediction [5, 20, 28, 32, 48, 63] has added a temporal dimension by estimating per-voxel dynamics, pushing the boundaries of dynamic scene understanding. However, previous research lacked a cohesive evaluation framework, often conducting experiments on isolated benchmarks (such as Occ3D [47] or OpenOcc [48]) or comparing performance using single metrics (like mIoU or RayIoU). In this work, we present a unified framework that excels in both semantic occupancy prediction and occupancy flow prediction tasks, validated comprehensively to establish a new, robust baseline.

3. Method

Our research focuses on vision-based 3D semantic occupancy and flow prediction. Given N_I input images $\mathbf{I} \in \mathbb{R}^{N_I \times X \times Y \times 3}$, the goal is to predict the 3D semantic occupancy grid and the 3D occupancy flow grid, supervised by the corresponding ground truth $\mathbf{O}_c \in \mathbb{R}^{H \times W \times Z}$ and $\mathbf{O}_f \in \mathbb{R}^{H \times W \times Z \times 2}$. Each value in \mathbf{O}_c identifies the class of its corresponding spatial position, including *empty* (indicating unoccupied space) and semantic classes such as *pedestrian* and *car*. Each voxel in the flow grid \mathbf{O}_f contains a 2D flow vector specifying movement in the x and y directions. Fig. 2 presents a schematic view of our approach. Initially, 2D features \mathbf{f}_I are extracted from surrounding images, which are subsequently transformed into 3D space by adaptive lifting. The lifted features \mathbf{f}_{Lift} are encoded within the 3D space, together with historical frame data. The encoded volume features \mathbf{f}_v are then decoded by various task

heads to predict the 3D semantic occupancy and flow grids.

Overall, Sec. 3.2 presents an occlusion-aware adaptive lifting technique to improve 2D-to-3D view transformation. Sec. 3.3 describes a semantic prototype-based occupancy prediction head that enhances 2D-3D semantic alignment and addresses the long-tail problem. Sec. 3.4 introduces a BEV cost volume approach with a regression-classification strategy to boost occupancy flow predictions.

3.1. Revisiting Depth-Based LSS

As the core module of vision-based occupancy prediction, the 2D-to-3D view transformation process is formulated as:

$$\mathbf{f}_{Lift} = \mathbf{M}_T \cdot \mathbf{f}_I, \quad (1)$$

where \mathbf{f}_I and \mathbf{f}_{Lift} denote the image and lifted features, respectively. Flattening and reshaping operations are omitted here. Depth-based LSS initializes \mathbf{M}_T as a zero matrix, filling it with discretized depth probabilities. For each pixel feature $\mathbf{x}^{(u,v)}$, depth probability is predicted as $\{P(o_d^i|\mathbf{x})\}_{i=1}^D$ across D discretized bins along the camera sightline. Assume bin i 's true depth is d (with slight abuse of notation). The coordinate is mapped with camera parameters (provided by datasets) from the image system (ICS) to the voxel system (VCS), i.e., $(u, v, d) \rightarrow (h, w, z)$. Each depth probability is assigned to \mathbf{M}_T at the row indexed by the rounded voxel coordinate $[(h, w, z)]$ and the column indexed by (u, v) . Given \mathbf{M}_T 's sparsity, this is efficiently implemented using sparse operations, such as `sparse_coo_tensor` in `PyTorch` or `BEVPool` [14, 44].

3.2. Occlusion-Aware Adaptive Lifting

We enhance depth-based LSS by introducing probability transfer from surface to occluded areas. We first replace the hard rounding-based filling strategy in Sec. 3.1 with a probability-based soft filling approach. As detailed in Fig. 3, we use trilinear interpolation to diffuse a mapped coordinate in VCS over its eight neighboring points. Continuing the notation from the previous section, the interpolation calculates eight probability values:

$$\{p_s = (1 - |h - \hat{h}_s|) \cdot (1 - |w - \hat{w}_s|) \cdot (1 - |z - \hat{z}_s|)\}_{s=1}^8, \quad (2)$$

where $\{(\hat{h}_s, \hat{w}_s, \hat{z}_s)\}_{s=1}^8$ are neighbor coordinates. For each of them, we fill \mathbf{M}_T with $p_s \cdot P(o_d^i|\mathbf{x})$, indexing the column




Figure 3. **Comparison between (a) Depth-based LSS and (b) Our occlusion-aware adaptive lifting.** Blue and red denote distinct pixels or voxels. The dotted line represents the camera sightline, with opacity reflecting the weight of mapping 2D features to 3D voxels—higher opacity indicates greater weights.

with the ICS coordinate (u, v) and the row with VCS coordinate $(\hat{h}_s, \hat{w}_s, \hat{z}_s)$. This distributes $P(o_d^i | \mathbf{x})$ by distance, enabling differentiable 2D-to-3D lifting *w.r.t.* coordinates.

Depth-based LSS methods excel by guiding the 2D-to-3D feature transformation through explicit depth probability modeling. However, the target of depth estimation follows a δ distribution, which causes most of the weights to concentrate on surface points. As a result, occluded areas get much lower weights, creating a bias that hinders learning from these regions. We aim to resolve this informational gap by addressing two key scenarios: intra-object (within an object) and inter-object (between objects) occlusion. For both, we construct a probability transfer matrix from visible to occluded parts. For the former type of occlusion, we design conditional probabilities to transfer surface probabilities to occluded length probabilities. For notational simplicity, we redefine the discrete depth probabilities of pixel feature \mathbf{x} as $\{P(o_d^i)\}_{i=1}^D$. We employ the same binning approach to discretize occluded length predictions. To convert the discrete depth probability $P(o_d)$ into the discrete occluded length probability $P(o_{ol})$, we introduce the Bayesian conditional probability $P(o_{ol} | o_d)$:

$$P(o_{ol}^j) = \sum_{i=1}^D P(o_{ol}^j | o_d^i) \cdot P(o_d^i), \quad j = 1, \dots, D. \quad (3)$$

To determine $P(o_{ol}^j | o_d^i)$, we consider the physical meaning of depth: if the camera sightline reaches point i , it indicates that all points at shallower depths (*i.e.*, with indices smaller than i are empty. Therefore, we only need to parameterically estimate the conditional probability for bin i to positions with larger depth. $P(o_{ol}^j | o_d^i)$ is modeled as:

$$P(o_{ol}^j | o_d^i) = \begin{cases} 0, & \text{if } j < i \\ 1, & \text{if } j = i \\ f_{ol}(\mathbf{x}, j - i), & \text{if } j > i. \end{cases} \quad (4)$$

Here, $f_{ol}(\mathbf{x}, j - i)$ is the likelihood of bin j being occupied given bin i corresponds to the surface point, estimated using pixel feature \mathbf{x} . In implementation, we compute $(D - 1)$

discrete probabilities across all relative positions of i and j via convolution. The transformation from depth to occluded length probabilities is achieved by multiplying depth probabilities with the causal conditional probability matrix.

For inter-object occlusion, we design a probability transfer matrix to propagate depth probabilities to surrounding points. For each point (u, v, d) , we use MLPs to estimate offsets $(\Delta u, \Delta v)$ and weights ω from its feature \mathbf{x} . New points $(u + \Delta u, v + \Delta v, d^i; \omega \cdot P(o_{ol}))$ are then incorporated into \mathbf{M}_T via soft filling and optimized during training. For efficiency, we limit transfer to the m points with top depth probabilities. m is set to a small value of 3 in practice.

The predominant role of depth probabilities in 2D-to-3D transformation can lead to suboptimal model convergence due to initial inaccurate depth estimations. To mitigate this, we introduce a denoising operation, inspired by query denoising in object detection [12, 18, 60]. This approach utilizes ground-truth depth probabilities to guide early training. We combine ground-truth and predicted depths via a weighted average for adaptive lifting, with the ground-truth weight starting at 1 and decaying to 0 via cosine annealing:

$$P(o_d) = \frac{1}{2}[(1 + \cos(\frac{\pi e}{E})) \cdot P_{gt}(o_d) + (1 - \cos(\frac{\pi e}{E})) \cdot P_{pred}(o_d)], \quad (5)$$

where $P(o_d)$ is the final depth probability for adaptive lifting. $P_{gt}(o_d)$ and $P_{pred}(o_d)$ are ground truth and predicted depth probabilities, e is the current training step, and E is the total annealing steps. During test, $P(o_d) = P_{pred}(o_d)$.

3.3. Semantic Prototype-Based Occupancy Head

After 2D-to-3D view transformation, we align features across domains using shared semantic prototypes (Fig. 4). These prototypes are randomly initialized and function as the class weights for both 2D and 3D loss computations, thereby creating a direct semantic bridge between the 2D and 3D feature representations. Given per-class prototypes, a straightforward method for decoding semantic occupancy is to compute voxel-feature-to-prototype similarity with cross-entropy supervision. However, this is suboptimal due to the severe class imbalance in typical scenes.

To overcome this, we introduce a targeted training strategy comprising two components: conditional loss computation and uncertainty-guided voxel sampling. First, we only compute losses for classes present in each ground-truth sample, generating per-class logit maps by computing the inner product between the voxel features against the prototypes of only these classes. Second, inspired by prior works [8, 17, 62], we perform uncertainty-guided sampling to further focus the training on informative regions. We treat the computed logits as a measure of model uncertainty and combine this with class priors to form a sampling distribution. From this distribution, we sample K hard voxels, concentrating computation on low-confidence regions and under-represented classes. The final loss is then computed




Figure 4. **Illustration of the semantic prototype-based occupancy head.** Shared semantic prototypes calculate segmentation logits for both 2D and 3D features. Training is conducted only on the top K voxels (or pixels) sampled based on logit uncertainty.

exclusively on the sampled voxels and is a combination of per-class binary cross-entropy and Dice loss:

$$\mathcal{L}_{3D} = \alpha \mathcal{L}_{Dice} + \beta \mathcal{L}_{BCE}, \quad (6)$$

where α, β are balancing coefficients. During inference, the final voxel category prediction is the one with the highest response to its semantic prototype or *empty* embedding:

$$\hat{\mathbf{o}}_v = \arg \max_c (\text{MLP}(P_c) \cdot \mathbf{f}_v), \quad (7)$$

where \mathbf{f}_v denotes the volume features, P_c is the prototype of class c or *empty* embedding, and MLP further encodes P_c . To further enforce semantic consistency, we also introduce an auxiliary 2D loss, \mathcal{L}_{2D} , by supervising the 2D image features with the same shared prototypes. The 2D ground-truth masks for this loss are generated by projecting semantic LiDAR data (a component of the occupancy GT) onto the image planes. The formulation of \mathcal{L}_{2D} mirrors that of \mathcal{L}_{3D} , combining Dice and binary cross-entropy losses.

3.4. BEV Cost Volume-Based Flow Prediction

As shown in Fig. 2, occupancy flow is decoded from the same volumetric features as semantics. Conventional approaches typically predict flow using single-frame features. However, this places a significant representational bottleneck, as features must jointly encode static semantics and dynamic motion cues. Inspired by stereo matching [22], we resolve this by constructing a BEV cost volume, which provides an explicit motion prior from cross-frame correspondences to decouple the tasks. The cost volume construction (Fig. 5) begins by collapsing the volumetric features (0-4m height) into a compact, downsampled BEV map. This step is crucial as it efficiently models the dominant planar (x-y) motion while simultaneously enlarging the features' receptive fields. The prior frame's BEV map is then warped to the current coordinate system via ego-motion. Finally, the cost volume is populated by computing the cosine similarity between current features and warped historical features within a local search window [22]:

$$\text{cv}(\mathbf{f}_v^{(t)}; k) = \frac{\hat{\mathbf{f}}_v^{(t)} \cdot \text{warp}(\hat{\mathbf{f}}_v^{(t-1)}(\Delta p_k))}{\|\hat{\mathbf{f}}_v^{(t)}\|_2 \cdot \|\text{warp}(\hat{\mathbf{f}}_v^{(t-1)}(\Delta p_k))\|_2}, \quad (8)$$




Figure 5. **Illustration of the Flow Head.** The cost volume is generated from downsampled BEV features and used to predict bin probabilities along with volume features. Final flow values are computed as a weighted sum of the bin center values.

where $\hat{\mathbf{f}}_v^{(t)}$ is the current frame's BEV map, derived by compressing and downsampling the volumetric feature $\mathbf{f}_v^{(t)}$. The $\text{warp}(\cdot)$ function transforms historical features to the current coordinate system, and Δp_k denotes a predefined sampling offset in the search window. For notational simplicity, we omit post-processing steps like upsampling from this formulation; these operations are illustrated in Fig. 5. The final flow prediction head is then formulated as:

$$\hat{\mathbf{o}}_f = f_{Flow}(\text{cat}(\mathbf{f}_v^{(t)}, \text{cv}(\mathbf{f}_v^{(t)}))). \quad (9)$$

The cost volume-based flow prediction method establishes cross-frame correspondences through explicit feature matching. This approach alleviates the representational burden inherent in semantic-motion joint prediction and is highly efficient, as it reuses cached features from the prior frame, obviating the need for re-computation. Furthermore, to enhance robustness to diverse flow magnitudes, we adopt a hybrid classification-regression scheme [2, 27]. The continuous flow space is discretized into a set of bins based on the value range observed in the training data. The flow head then predicts a probability distribution over these bins, from which the final continuous flow vector is computed as the expectation over the bin centers:

$$\hat{\mathbf{o}}_f = \sum_{n=1}^{N_b} p_b^n \cdot \mathbf{b}^n, \quad (10)$$

where p_b^n is the predicted probability of the flow vector falling into the n -th bin, \mathbf{b}^n is the corresponding bin's center value over a total of N_b bins. f_{Flow} is thus redefined to predict p_b . We supervise the prediction with the ground-truth flow \mathbf{o}_f by minimizing the L2 loss and maximizing the cosine similarity:

$$\mathcal{L}_{flow}^{reg} = \|\hat{\mathbf{o}}_f - \mathbf{o}_f\|_2^2 - \frac{\hat{\mathbf{o}}_f \cdot \mathbf{o}_f}{\|\hat{\mathbf{o}}_f\|_2 \cdot \|\mathbf{o}_f\|_2}, \quad (11)$$

where the L2 loss ensures magnitude accuracy while cosine similarity guarantees directional precision. To improve flow classification across discrete bins, we additionally introduce a classification loss with the ground-truth bin index:

$$\mathcal{L}_{flow}^{cls} = - \sum_{n=1}^{N_b} \mathbb{1}(\text{Index}[\mathbf{o}_f; \mathbf{b}] = n) \log(p_b^n), \quad (12)$$

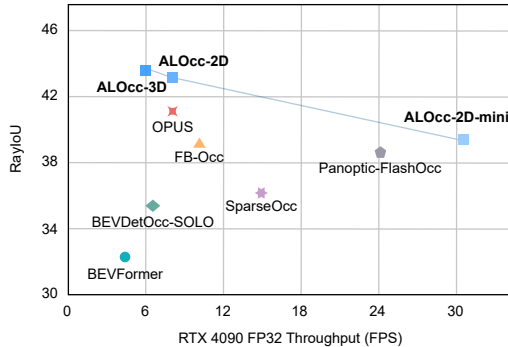


Figure 6. **Speed/accuracy trade-off of our method.** We benchmark RayIoU and speed on the Occ3D dataset using one RTX 4090 GPU with a batch size of 1. All the methods utilize a ResNet-50 backbone and 256×704 input size, except BEVFormer (900×1600). All detailed results are available in Tab. 2.

where $\mathbb{1}(\cdot)$ is the indicator function and $\text{Index}[\mathbf{o}_f; \mathbf{b}]$ represents the index of ground truth \mathbf{o}_f within the bins \mathbf{b} .

The final occupancy flow loss is:

$$\mathcal{L}_{flow} = \mathcal{L}_{flow}^{reg} + \mathcal{L}_{flow}^{cls}. \quad (13)$$

3.5. Training Objective

For only predicting 3D semantic occupancy, the overall training objective is

$$\mathcal{L}_{sem} = \mathcal{L}_{3D} + \mathcal{L}_{2D} + \mathcal{L}_{depth}, \quad (14)$$

where \mathcal{L}_{depth} is the depth loss. When extending to joint semantic occupancy and flow prediction, we incorporate flow supervision, resulting in the overall objective:

$$\mathcal{L}_{sem-flow} = \mathcal{L}_{3D} + \mathcal{L}_{2D} + \mathcal{L}_{depth} + \mathcal{L}_{flow}. \quad (15)$$

4. Experiment

4.1. Experimental Setup

Dataset. We employ the nuScenes dataset [3] for most experiments, which contains 700 training, 150 validation, and 150 test scenes. Occ3D [47] and OpenOcc [48] add voxel-wise annotations to nuScenes, spanning a range of $\pm 40m$ (X,Y) and $-1m$ to $5.4m$ (Z). The voxel resolution is $0.4m$ in all dimensions. Occ3D provides 18 semantic categories, including 17 object classes and one *empty* class that denotes unoccupied spaces. OpenOcc provides 17 categories (omitting “others”) and additionally includes per-voxel motion flow annotations for the X-Y plane.

Benchmarks. We evaluate ALOcc on three distinct benchmarks to assess its capabilities in various scenarios: *i) 3D semantic occupancy prediction with mask.* This benchmark emphasizes the prediction of observed regions. Following prior works [25, 39, 51], ALOcc is trained on Occ3D using the official camera visibility mask. Performance is measured by mIoU on 17 semantic categories. To better reflect performance on critical dynamic objects, we also report mIoU_D on a subset of 8 dynamic foreground classes.

Method	Input	Backbone	Image Size	mIoU_D^m	mIoU^m	FPS
BEVDetOcc-SF [14]	C	ResNet-50	256×704	34.4	41.9	6.5
UniOCC [41]	C	ResNet-50	256×704	-	39.7	-
FB-Occ [25]	C	ResNet-50	256×704	34.2	39.8	10.3
SurroundSDF [30]	C	ResNet-50	256×704	36.2	42.4	-
FlashOCC [57]	C	ResNet-50	256×704	24.7	32.0	29.6
COTR [39]	C	ResNet-50	256×704	<u>38.6</u>	<u>44.5</u>	0.5
ViewFormer [20]	C	ResNet-50	256×704	35.0	41.9	-
OPUS [49]	C	ResNet-50	256×704	33.3	36.2	8.2
ALOcc-2D-mini	C	ResNet-50	256×704	35.4	41.4	30.5
ALOcc-2D	C	ResNet-50	256×704	38.7 \uparrow 0.1	44.8 \uparrow 0.3	8.2
ALOcc-3D	C	ResNet-50	256×704	39.3 \uparrow 0.7	45.5 \uparrow 1.0	6.0
ALOcc-2D-mini	C	Intern-T	256×704	37.9	43.7	16.6
ALOcc-2D	C	Intern-T	256×704	40.7	46.6	8.8
ALOcc-3D	C	Intern-T	256×704	41.5	47.5	5.9
HyDRa [53]	C+R	ResNet-50	256×704	40.6	44.4	-
EFFOcc [46]	C+L	ResNet-50	256×704	<u>50.1</u>	<u>52.8</u>	-
SDGOcc [9]	C+L	ResNet-50	256×704	47.7	51.7	7.5
ALOcc-2D-mini	C+D	ResNet-50	256×704	46.2	50.0	28.1
ALOcc-2D	C+D	ResNet-50	256×704	50.3 \uparrow 0.2	53.5 \uparrow 0.7	8.1
ALOcc-3D	C+D	ResNet-50	256×704	50.6 \uparrow 0.5	54.5 \uparrow 1.7	6.0
ALOcc-2D-mini	C+D	Intern-T	256×704	48.9	52.1	16.1
ALOcc-2D	C+D	Intern-T	256×704	52.0	54.9	8.7
ALOcc-3D	C+D	Intern-T	256×704	52.4	55.6	5.8
BEVFormer [24]	C	ResNet-101	900×1600	37.2	39.2	4.4
VoxFormer [23]	C	ResNet-101	900×1600	-	40.7	-
SurroundOcc [52]	C	ResNet-50	900×1600	31.2	37.2	-
FastOcc [13]	C	ResNet-101	640×1600	34.5	39.2	-
PanoOcc [51]	C	ResNet-101	640×1600	37.3	42.1	-
OSP [45]	C	ResNet-101	900×1600	37.0	41.2	-
BEVDetOcc [14]	C	Swin-Base	512×1408	36.9	42.0	1.1
COTR [39]	C	Swin-Base	512×1408	<u>41.3</u>	<u>46.2</u>	-
ALOcc-2D	C	Swin-Base	512×1408	44.5 \uparrow 3.2	49.3 \uparrow 3.1	1.8
ALOcc-3D	C	Swin-Base	512×1408	46.1 \uparrow 4.8	50.6 \uparrow 4.4	1.5
OccFusion [40]	C+L	ResNet-101	900×1600	45.3	46.8	-
BEVFusion [35]	C+L	Swin-Base	512×1408	48.7	54.0	-
EFFOcc [46]	C+L	Swin-Base	512×1408	50.7	54.1	-
FusionOcc [61]	C+L	Swin-Base	512×1408	<u>53.1</u>	<u>56.6</u>	-
ALOcc-2D	C+D	Swin-Base	512×1408	56.8 \uparrow 3.7	58.7 \uparrow 2.1	1.8
ALOcc-3D	C+D	Swin-Base	512×1408	57.8 \uparrow 4.7	60.0 \uparrow 3.4	1.5

Table 1. **3D semantic occupancy prediction performance of mIoU_D^m and mIoU^m w.r.t. training with camera visible mask on Occ3D.** Input modalities include *Camera (C)*, *Radar (R)*, *LIDAR (L)*, and *Depth (D)*, a sparser signal mapped from LIDAR points. Best results among similar conditions are **bolded**, top results from other methods are underlined, with improvements over these marked by arrows. FPS is measured on RTX 4090.

ii) 3D semantic occupancy prediction without mask. This benchmark emphasizes the simultaneous prediction of both observed and occluded regions. We follow the setup of SparseOcc [29] and train on Occ3D without the visibility masks. We calculate mIoU and RayIoU for 17 semantic categories. *iii) Joint 3D Semantic Occupancy and Flow Prediction.* This benchmark evaluates the joint prediction of semantics and motion on OpenOcc, following the standard protocol [48, 63]. We report RayIoU for 16 semantic categories and mAVE_{TP} for flow metrics in true-positive areas. Their weighted average is denoted as Occ Score. In particular, we observed that mAVE_{TP} exhibits a strong dependence on semantic occupancy prediction, resulting in suboptimal convergence across epochs. Therefore, we additionally employ per-voxel mAVE as a supplementary metric.

Implementation Details. Our default setup uses a ResNet-50 [11] backbone with the input image size of 256×704 .

Method	Backbone	Image Size	mIoU	RayIoU	RayIoU _{1m, 2m, 4m}			FPS
OccFormer [62]	ResNet-101	640 × 960	21.9	-	-	-	-	-
TPVFormer [15]	ResNet-101	640 × 960	27.8	-	-	-	-	-
CTF-Occ [47]	ResNet-101	640 × 960	28.5	-	-	-	-	-
BEVFormer [24]	ResNet-101	900 × 1600	23.7	32.4	26.1	32.9	38.0	4.4
BEVDetOcc-SF [14]	ResNet-50	256 × 704	24.3	35.2	31.2	35.9	38.4	6.5
FB-Occ [25]	ResNet-50	256 × 704	31.1	39.0	33.0	40.0	44.0	10.3
RenderOcc [42]	Swin-Base	512 × 1408	24.4	19.5	13.4	19.6	25.5	-
SparseOcc [29]	ResNet-50	256 × 704	30.9	36.1	30.2	36.8	41.2	15.0
Panoptic-FlashOcc [58]	ResNet-50	256 × 704	<u>31.6</u>	38.5	32.8	39.3	43.4	24.2
OPUS [49]	ResNet-50	256 × 704	-	<u>41.2</u>	34.7	42.1	46.7	8.2
ALOcc-2D-mini	ResNet-50	256 × 704	33.4	39.3	32.9	40.1	44.8	30.5
ALOcc-2D	ResNet-50	256 × 704	37.4 \uparrow 5.8	43.0 \uparrow 1.8	37.1	43.8	48.2	8.2
ALOcc-3D	ResNet-50	256 × 704	38.0 \uparrow 6.4	43.7 \uparrow 2.5	37.8	44.7	48.8	6.0
ALOcc-2D-mini	Intern-T	256 × 704	35.9	42.4	35.9	43.3	47.9	16.6
ALOcc-2D	Intern-T	256 × 704	39.1	44.9	38.8	45.8	50.0	8.8
ALOcc-3D	Intern-T	256 × 704	40.0	45.9	39.8	46.9	50.9	5.9

Table 2. **3D semantic occupancy prediction performance of RayIoU and mIoU w.r.t. training without camera visible mask on Occ3D.** All methods operate on camera inputs.

Method	Sup.	Backbone	Image Size	Occ Score	mAVE	mAVE _{TP}	RayIoU	RayIoU _{1m, 2m, 4m}			FPS
OccNeRF [59]	2D	ResNet-101	256 × 704	28.5	-	1.59	31.7	16.6	29.3	49.2	-
RenderOcc [42]	2D	Swin-Base	512 × 1408	33.0	-	1.63	36.7	20.3	32.7	49.9	-
LetOccFlow [33]	2D	ConvNeXt-Base	512 × 1408	36.4	-	1.45	<u>40.5</u>	25.5	39.7	56.3	-
OccNet [48]	3D	ResNet-50	900 × 1600	35.7	-	1.61	39.66	29.3	39.7	50.0	-
BEVDetOcc-SF [14]	3D	ResNet-50	250 × 704	33.0	1.42	1.78	36.7	31.6	37.3	41.1	6.2
FB-Occ [25]	3D	ResNet-50	250 × 704	39.2	<u>0.591</u>	0.651	39.0	32.7	39.9	44.4	10.1
CascadeFlow [28]	3D	ResNet-50	250 × 704	40.9	-	<u>0.47</u>	39.6	33.5	40.6	45.3	-
F-Occ [63]	3D	Intern-T	250 × 704	<u>41.0</u>	-	0.493	39.9	33.9	40.7	45.2	-
ALOcc-Flow-2D	3D	ResNet-50	256 × 704	42.1 \uparrow 1.1	0.537 \downarrow 0.054	0.427 \downarrow 0.043	40.5	34.3	41.3	45.8	7.0
ALOcc-Flow-3D	3D	ResNet-50	256 × 704	43.0 \uparrow 2.0	0.556 \downarrow 0.035	0.481	41.9 \uparrow 1.4	35.6	42.8	47.4	5.5
ALOcc-Flow-2D	3D	Intern-T	256 × 704	44.4	0.545	0.451	43.2	36.8	44.1	48.7	6.7
ALOcc-Flow-3D	3D	Intern-T	256 × 704	45.3	0.574	0.469	44.4	38.0	45.5	49.8	5.3

Table 3. **Performance evaluation of joint 3D semantic occupancy and flow prediction on OpenOcc.** All methods listed are camera-based. The “sup.” column indicates the supervision signal, specifying that weakly supervised methods are trained using only labels.

We also test with FlashInternImage-Tiny [54] and scale up with Swin-Transformer Base [34] at an image size of 512×1408 . We follow established methods for depth estimation [14, 26], image/BEV augmentation [14], and temporal fusion (16 frames) with CBGS strategies [26, 43]. Our default 3D/2D models use voxel sizes of $200 \times 200 \times 16$ (32-dim features) and $200 \times 200 \times 1$ (80-dim features), respectively. All models are trained with the AdamW optimizer [36] with a learning rate of 2×10^{-4} and a batch size of 16. Training lasts 12 epochs for semantic occupancy and 18 for joint occupancy semantic-flow prediction, with loss weights $\alpha = 5$, $\beta = 20$ (following Mask2Former [8]). **Notably**, the flow head and its supervision are used only for *benchmark iii*). Further details are in the supplementary material.

4.2. Accuracy/Speed Trade-off Evaluation

To assess ALOcc’s deployment capabilities, we analyze the speed/accuracy trade-off of its variants on *benchmark ii*). As shown in Fig. 6, we offer multiple model versions that achieve an excellent balance between accuracy and efficiency. Compared to ALOcc-3D, ALOcc-2D (-mini) employs spatial compression techniques similar to FlashOcc [57], using height compression and 2D convolutions for vol-

ume encoding. ALOcc-2D-mini leverages monocular depth estimation [21] and reduced channel sizes, while other variants use stereo depth [22]. Results show that ALOcc-3D and ALOcc-2D surpass SOTAs with higher speeds, and ALOcc-2D-mini achieves real-time inference while maintaining a near-SOTA performance. Further details of each variant are provided in the **supplemental materials**.

4.3. Comparison with SOTAs

We compare our method with the current SOTAs on the Occ3D and OpenOcc datasets, with results summarized in Tab. 1, Tab. 2, and Tab. 3. ALOcc and ALOcc-Flow in the tables differ by the presence/absence of flow prediction. ALOcc-2D (default) uses 2D convolutions for height-compressed volume encoding with 80 channels, while ALOcc-3D (default) employs 3D convolutions for full 3D volume encoding with 32 channels. All results are reported without post-processing, such as test-time augmentation.

3D Semantic Occupancy Prediction with Mask. Tab. 1 compares ALOcc with SOTAs methods on Occ3D benchmark, focusing training on critical, camera-visible regions. Our method significantly outperforms existing approaches across various image input sizes and backbones. Notably,

Exp.	Condition	mIoU _D ^m	mIoU ^m	mIoU	RayIoU	RayIoU _{1m, 2m, 4m}}
0	ALOcc-2D-40	38.5	44.5	36.8	42.5	36.6 43.4 47.7
1	Exp. 0 w/o AL	37.5	43.5	36.1	41.3	35.2 42.1 46.6
2	Exp. 0 w/o SP	36.0	42.1	33.8	39.9	34.5 40.6 44.7
3	Exp. 2 w/o AL	34.9	41.2	33.3	39.5	34.0 40.3 44.2

Table 4. **Ablation study of semantic occupancy prediction on Occ3D.** **AL**: adaptive lifting; **SP**: semantic prototype-based occupancy head. “w/o” indicates the removal of a module.

when provided with ground-truth depth in adaptive lifting, ALOcc surpasses even multi-modal fusion methods that rely on extra LiDAR or radar backbones, demonstrating the superiority of our approach.

3D Semantic Occupancy Prediction without Mask.

Tab. 2 showcases the results of 3D semantic occupancy prediction of training without the camera visible mask, where we outperform all prior methods on both mIoU and RayIoU. This demonstrates the effectiveness of our method in capturing the complete scene context and understanding the spatial relationships between objects. Our method also exhibits superior performance in terms of RayIoU at different distances, further highlighting its robustness in capturing object details. Notably, these performance gains are achieved with minimal computational overhead. Our scaled-down variant outperforms competing real-time models in inference speed while maintaining high accuracy, which is critical for on-device applications.

Joint 3D Semantic Occupancy and Flow Prediction.

Tab. 3 presents the results of 3D semantic occupancy and flow joint prediction on OpenOcc. Due to the limited existing approaches on this benchmark, we also include weakly supervised methods for comparison. Our method excels in Occ Score, mAVE, mAVE_{TP}, and RayIoU across distances. It effectively captures semantic and motion information, demonstrating robustness in scene dynamics. ALOcc-Flow-2D outperforms ALOcc-Flow-3D in terms of mAVE and mAVE_{TP}. This can be attributed to: 1) mAVE_{TP} is computed on true positive semantic prediction areas, so higher semantic errors (ALOcc-Flow-2D) reduce the regions requiring this metric; 2) This benchmark evaluates the flow on X and Y axes, so height compression and added feature channels emphasize critical motion cues.

4.4. Ablation Study

Unless specified, we use ALOcc-2D with a BEV feature channel of 40 as our baseline for these experiments.

Components Analysis for Occ Head. Tab. 4 presents an ablation study to evaluate the effectiveness of our core components in terms of semantic occupancy results. The full model (**Exp. 0**) demonstrates superior performance across all metrics. Removing the adaptive lifting component (**Exp. 1** and **Exp. 3**) results in decreased performance, particularly in mIoU_D^m and mIoU, underscoring the efficacy of our adaptive and accurate 2D-to-3D view transfor-

Exp.	Condition	Occ Score	mAVE	mAVE _{TP}	RayIoU	RayIoU _{1m, 2m, 4m}}
0	ALOcc-2D-40	-	-	-	42.4	35.9 43.3 47.9
1	Exp. 0 + Flow	40.7	0.597	0.508	39.7	33.5 40.5 45.2
2	Exp. 1 + BC	39.9	0.565	0.464	38.3	32.7 39.0 43.3
3	Exp. 1 + CV	41.1	0.588	0.503	40.2	33.8 41.1 45.6
4	Exp. 2 + CV	41.1	0.562	0.451	39.6	33.3 40.4 45.1
5	Exp. 4 + CS	42.1	0.537	0.427	40.5	34.3 41.3 45.8

Table 5. **Ablation study of 3D semantic occupancy and flow joint prediction results on OpenOcc.** **Flow**: flow head, **BC**: bin classification, **CV**: BEV cost volume, **CS**: channel scaling from 40 to 80. “+” indicates module addition.

mation. In **Exp. 2**, we substitute the semantic prototype-based occupancy head with FBOcc’s occupancy head, leading to a more pronounced performance decline. This highlights the crucial role of the SP component in enhancing the model’s semantic occupancy prediction accuracy.

Components Analysis for Flow Head. As shown in Tab. 5, we conducted an ablation study to evaluate how different components affect flow prediction performance. In **Exp. 1**, we add a flow head implemented with a simple convolution. **Exp. 1 vs. Exp. 0** show that adding the flow head harms the prediction of semantic occupancy. This is due to the motion flow containing different information from the semantic occupancy, which increases the burden of volume features. The results of **Exp. 2 vs. Exp. 1** show that using bin classification significantly improves flow prediction performance. **Exp. 3** and **Exp. 4** demonstrate that the use of BEV cost volume alleviates pressure on scene representation, enhancing flow prediction while maintaining occupancy prediction performance. In **Exp. 5**, increasing BEV feature channels substantially improves overall performance, indicating that the information capacity of features is a bottleneck in joint learning of semantic occupancy and occupancy flow.

5. Conclusion

In this paper, we explore the challenge of vision-based 3D semantic occupancy and flow prediction. We propose an occlusion-aware adaptive lifting method, complemented by depth denoising to enhance the adaptability and robustness of the 2D-to-3D view transformation process. To further improve semantic occupancy learning, we introduce a semantic prototype-based occupancy head that aligns 2D and 3D semantics, combined with hard sample mining techniques to mitigate the long-tail problem. Additionally, we present a BEV cost volume-based approach to facilitate occupancy flow learning, reducing the burden on features to represent semantics and motion simultaneously. Experiments conducted on the Occ3D and OpenOcc datasets demonstrate that our method outperforms current SOTA solutions. Benefiting from the lightweight nature of our approach, we provide multiple model versions: our highest-performing model is faster than other methods with comparable performance, while our fastest model achieves superior performance compared to methods with similar speed.

References

- [1] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9297–9307, 2019. 2
- [2] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4009–4018, 2021. 5
- [3] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11621–11631, 2020. 6
- [4] Anh-Quan Cao and Raoul de Charette. Monoscene: Monocular 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3991–4001, 2022. 2
- [5] Dubing Chen, Wencheng Han, Jin Fang, and Jianbing Shen. Adaocc: Adaptive forward view transformation and flow modeling for 3d occupancy and flow prediction. *arXiv preprint arXiv:2407.01436*, 2024. 3
- [6] Dubing Chen, Huan Zheng, Jin Fang, Xingping Dong, Xianfei Li, Wenlong Liao, Tao He, Pai Peng, and Jianbing Shen. Rethinking temporal fusion with a unified gradient descent view for 3d semantic occupancy prediction. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1505–1515, 2025. 2
- [7] Xiaokang Chen, Kwan-Yee Lin, Chen Qian, Gang Zeng, and Hongsheng Li. 3d sketch-aware semantic scene completion via semi-supervised structure prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4193–4202, 2020. 2
- [8] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1290–1299, 2022. 2, 4, 7
- [9] ZaiPeng Duan, ChenXu Dang, Xuzhong Hu, Pei An, Junfeng Ding, Jie Zhan, YunBiao Xu, and Jie Ma. Sdgocc: Semantic and depth-guided bird’s-eye view transformation for 3d multimodal occupancy prediction. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6751–6760, 2025. 6
- [10] Wanshui Gan, Fang Liu, Hongbin Xu, Ning kai Mo, and Naoto Yokoya. Gaussianocc: Fully self-supervised and efficient 3d occupancy estimation with gaussian splatting. *arXiv preprint arXiv:2408.11447*, 2024. 3
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 6
- [12] Yulin He, Wei Chen, Tianci Xun, and Yusong Tan. Real-time 3d occupancy prediction via geometric-semantic disentanglement. *arXiv preprint arXiv:2407.13155*, 2024. 4
- [13] Jiawei Hou, Xiaoyan Li, Wenhao Guan, Gang Zhang, Di Feng, Yuheng Du, Xiangyang Xue, and Jian Pu. Fastocc: Accelerating 3d occupancy prediction by fusing the 2d bird’s-eye view and perspective view. In *IEEE International Conference on Robotics and Automation*, 2024. 6
- [14] Junjie Huang and Guan Huang. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *arXiv preprint arXiv:2203.17054*, 2022. 1, 2, 3, 6, 7
- [15] Yuanhui Huang, Wenzhao Zheng, Yunpeng Zhang, Jie Zhou, and Jiwen Lu. Tri-perspective view for vision-based 3d semantic occupancy prediction. *arXiv preprint arXiv:2302.07817*, 2023. 1, 2, 7
- [16] Yuanhui Huang, Wenzhao Zheng, Borui Zhang, Jie Zhou, and Jiwen Lu. Selfocc: Self-supervised vision-based 3d occupancy prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19946–19956, 2024. 3
- [17] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. Pointrend: Image segmentation as rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9799–9808, 2020. 4
- [18] Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M Ni, and Lei Zhang. Dn-detr: Accelerate detr training by introducing query denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13619–13627, 2022. 4
- [19] Jie Li, Yu Liu, Dong Gong, Qinfeng Shi, Xia Yuan, Chunxia Zhao, and Ian Reid. Rgbd based dimensional decomposition residual network for 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7693–7702, 2019. 2
- [20] Jinke Li, Xiao He, Chonghua Zhou, Xiaoqiang Cheng, Yang Wen, and Dan Zhang. Viewformer: Exploring spatiotemporal modeling for multi-view 3d occupancy perception via view-guided transformers. In *Computer Vision—ECCV 2024: 18th European Conference*, 2024. 3, 6
- [21] Yin hao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and Zeming Li. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. *arXiv preprint arXiv:2206.10092*, 2022. 2, 7
- [22] Yin hao Li, Han Bao, Zheng Ge, Jinrong Yang, Jianjian Sun, and Zeming Li. Bevestereo: Enhancing depth estimation in multi-view 3d object detection with temporal stereo. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1486–1494, 2023. 1, 2, 5, 7
- [23] Yiming Li, Zhiding Yu, Christopher Choy, Chaowei Xiao, Jose M Alvarez, Sanja Fidler, Chen Feng, and Anima Anandkumar. Voxformer: Sparse voxel transformer for camera-based 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9087–9098, 2023. 2, 6
- [24] Zhiqi Li, Wenhao Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Qiao Yu, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. *arXiv preprint arXiv:2203.17270*, 2022. 1, 2, 6, 7

- [25] Zhiqi Li, Zhiding Yu, David Austin, Mingsheng Fang, Shiyi Lan, Jan Kautz, and Jose M Alvarez. Fb-occ: 3d occupancy prediction based on forward-backward view transformation. *arXiv preprint arXiv:2307.01492*, 2023. 6, 7
- [26] Zhiqi Li, Zhiding Yu, Wenhai Wang, Anima Anandkumar, Tong Lu, and Jose M Alvarez. Fb-bev: Bev representation from forward-backward view transformations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6919–6928, 2023. 1, 2, 7
- [27] Zhenyu Li, Xuyang Wang, Xianming Liu, and Junjun Jiang. Binsformer: Revisiting adaptive bins for monocular depth estimation. *IEEE Transactions on Image Processing*, 2024. 5
- [28] Zhimin Liao and Ping Wei. Cascadeflow: 3d occupancy and flow prediction with cascaded sparsity sampling refinement framework. *CVPR 2024 Autonomous Grand Challenge Track On Occupancy and Flow*, 2024. 3, 7
- [29] Haisong Liu, Haiguang Wang, Yang Chen, Zetong Yang, Jia Zeng, Li Chen, and Limin Wang. Fully sparse 3d occupancy prediction. In *Computer Vision–ECCV 2024: 18th European Conference*, 2024. 3, 6, 7
- [30] Lizhe Liu, Bohua Wang, Hongwei Xie, Daqi Liu, Li Liu, Zhiqiang Tian, Kuiyuan Yang, and Bing Wang. Surroundsdf: Implicit 3d scene understanding based on signed distance field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 6
- [31] Shice Liu, Yu Hu, Yiming Zeng, Qiankun Tang, Beibei Jin, Yinhe Han, and Xiaowei Li. See and think: Disentangling semantic scene completion. *Advances in Neural Information Processing Systems*, 31, 2018. 2
- [32] Yili Liu, Linzhan Mou, Xuan Yu, Chenrui Han, Sitong Mao, Rong Xiong, and Yue Wang. Let occ flow: Self-supervised 3d occupancy flow prediction. *arXiv preprint arXiv:2407.07587*, 2024. 3
- [33] Yili Liu, Linzhan Mou, Xuan Yu, Chenrui Han, Sitong Mao, Rong Xiong, and Yue Wang. Let occ flow: Self-supervised 3d occupancy flow prediction. In *Conference on Robot Learning*, pages 2895–2912. PMLR, 2025. 7
- [34] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 7
- [35] Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. *arXiv preprint arXiv:2205.13542*, 2022. 2, 6
- [36] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. 2018. 7
- [37] Yuhang Lu, Xinge Zhu, Tai Wang, and Yuexin Ma. Ocotreeocc: Efficient and multi-granularity occupancy prediction using otree queries. *arXiv preprint arXiv:2312.03774*, 2023. 3
- [38] Junyi Ma, Xieyuanli Chen, Jiawei Huang, Jingyi Xu, Zhen Luo, Jintao Xu, Weihao Gu, Rui Ai, and Hesheng Wang. Cam4docc: Benchmark for camera-only 4d occupancy forecasting in autonomous driving applications. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21486–21495, 2024. 1
- [39] Qihang Ma, Xin Tan, Yanyun Qu, Lizhuang Ma, Zhizhong Zhang, and Yuan Xie. Cotr: Compact occupancy transformer for vision-based 3d occupancy prediction. *arXiv preprint arXiv:2312.01919*, 2023. 6
- [40] Zhenxing Ming, Julie Stephany Berrio, Mao Shan, and Stewart Worrall. Occfusion: Multi-sensor fusion framework for 3d semantic occupancy prediction. *IEEE Transactions on Intelligent Vehicles*, 2024. 6
- [41] Mingjie Pan, Li Liu, Jiaming Liu, Peixiang Huang, Longlong Wang, Shanghang Zhang, Shaoqing Xu, Zhiyi Lai, and Kuiyuan Yang. Uniocc: Unifying vision-centric 3d occupancy prediction with geometric and semantic rendering. *arXiv preprint arXiv:2306.09117*, 2023. 6
- [42] Mingjie Pan, Jiaming Liu, Renrui Zhang, Peixiang Huang, Xiaoqi Li, Hongwei Xie, Bing Wang, Li Liu, and Shanghang Zhang. Renderocc: Vision-centric 3d occupancy prediction with 2d rendering supervision. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12404–12411. IEEE, 2024. 3, 7
- [43] Jinhung Park, Chenfeng Xu, Shijia Yang, Kurt Keutzer, Kris Kitani, Masayoshi Tomizuka, and Wei Zhan. Time will tell: New outlooks and a baseline for temporal multi-view 3d object detection. *arXiv preprint arXiv:2210.02443*, 2022. 7
- [44] Jonah Philion and Sanja Fidler. 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*, pages 194–210. Springer, 2020. 1, 2, 3
- [45] Yiang Shi, Tianheng Cheng, Qian Zhang, Wenyu Liu, and Xinggang Wang. Occupancy as set of points. In *Computer Vision–ECCV 2024: 18th European Conference*, 2024. 6
- [46] Yining Shi, Kun Jiang, Jinyu Miao, Ke Wang, Kangan Qian, Yunlong Wang, Jiushi Li, Tuopu Wen, Mengmeng Yang, Yiliang Xu, et al. Effocc: Learning efficient occupancy networks from minimal labels for autonomous driving. *arXiv preprint arXiv:2406.07042*, 2024. 6
- [47] Xiaoyu Tian, Tao Jiang, Longfei Yun, Yucheng Mao, Huitong Yang, Yue Wang, Yilun Wang, and Hang Zhao. Occ3d: A large-scale 3d occupancy prediction benchmark for autonomous driving. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 2, 3, 6, 7
- [48] Wenwen Tong, Chonghao Sima, Tai Wang, Li Chen, Silei Wu, Hanming Deng, Yi Gu, Lewei Lu, Ping Luo, Dahua Lin, et al. Scene as occupancy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8406–8415, 2023. 1, 2, 3, 6, 7
- [49] Jiabao Wang, Zhaojiang Liu, Qiang Meng, Liujiang Yan, Ke Wang, Jie Yang, Wei Liu, Qibin Hou, and Mingming Cheng. Opus: occupancy prediction using a sparse set. In *Advances in Neural Information Processing Systems*, 2024. 3, 6, 7
- [50] Xiaofeng Wang, Zheng Zhu, Wenbo Xu, Yunpeng Zhang, Yi Wei, Xu Chi, Yun Ye, Dalong Du, Jiwen Lu, and Xinggang Wang. Openoccupancy: A large scale benchmark for surrounding semantic occupancy perception. *arXiv preprint arXiv:2303.03991*, 2023. 1, 2, 3

- [51] Yuqi Wang, Yuntao Chen, Xingyu Liao, Lue Fan, and Zhaoxiang Zhang. Panoocc: Unified occupancy representation for camera-based 3d panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 3, 6
- [52] Yi Wei, Linqing Zhao, Wenzhao Zheng, Zheng Zhu, Jie Zhou, and Jiwen Lu. Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving. *arXiv preprint arXiv:2303.09551*, 2023. 3, 6
- [53] Philipp Wolters, Johannes Gilg, Torben Teepe, Fabian Herzog, Anouar Laouichi, Martin Hofmann, and Gerhard Rigoll. Unleashing hydra: Hybrid fusion, depth consistency and radar for unified 3d perception. *arXiv preprint arXiv:2403.07746*, 2024. 6
- [54] Yuwen Xiong, Zhiqi Li, Yuntao Chen, Feng Wang, Xizhou Zhu, Jiapeng Luo, Wenhai Wang, Tong Lu, Hongsheng Li, Yu Qiao, et al. Efficient deformable convnets: Rethinking dynamic and sparse operator for vision applications. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5652–5661, 2024. 7
- [55] Tianyi Yan, Dongming Wu, Wencheng Han, Junpeng Jiang, Xia Zhou, Kun Zhan, Cheng-zhong Xu, and Jianbing Shen. Drivingsphere: Building a high-fidelity 4d world for closed-loop simulation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 27531–27541, 2025. 1
- [56] Shichao Yang, Yulan Huang, and Sebastian Scherer. Semantic 3d occupancy mapping through efficient high order crfs. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 590–597. IEEE, 2017. 1
- [57] Zichen Yu, Changyong Shu, Jiajun Deng, Kangjie Lu, Zongdai Liu, Jianguo Yu, Dawei Yang, Hui Li, and Yan Chen. Flashocc: Fast and memory-efficient occupancy prediction via channel-to-height plugin. *arXiv preprint arXiv:2311.12058*, 2023. 6, 7
- [58] Zichen Yu, Changyong Shu, Qianpu Sun, Junjie Linghu, Xiaobao Wei, Jianguo Yu, Zongdai Liu, Dawei Yang, Hui Li, and Yan Chen. Panoptic-flashocc: An efficient baseline to marry semantic occupancy with panoptic via instance center. *arXiv preprint arXiv:2406.10527*, 2024. 7
- [59] Chubin Zhang, Juncheng Yan, Yi Wei, Jiabin Li, Li Liu, Yansong Tang, Yueqi Duan, and Jiwen Lu. Occnerf: Self-supervised multi-camera occupancy prediction with neural radiance fields. *arXiv preprint arXiv:2312.09243*, 2023. 3, 7
- [60] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. 4
- [61] Shuo Zhang, Yupeng Zhai, Jilin Mei, and Yu Hu. Fusionocc: Multi-modal fusion for 3d occupancy prediction. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 787–796, 2024. 6
- [62] Yunpeng Zhang, Zheng Zhu, and Dalong Du. Occformer: Dual-path transformer for vision-based 3d semantic occupancy prediction. *arXiv preprint arXiv:2304.05316*, 2023. 2, 4, 7
- [63] Yun Zhao, Peiru Zheng, and Zhan Gong. 3d occupancy and flow prediction based on forward view transformatio. *CVPR* 2024 Autonomous Grand Challenge Track On Occupancy and Flow, 2024. 3, 6, 7
- [64] Huan Zheng, Wencheng Han, and Jianbing Shen. Decoupling fine detail and global geometry for compressed depth map super-resolution. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 951–960, 2025. 2