

Scene as Occupancy

Wenwen Tong^{1,2*}, Chonghao Sima^{1,3*†}, Tai Wang^{1,4}, Li Chen^{1,3}, Silei Wu²,
Hanming Deng², Yi Gu¹, Lewei Lu², Ping Luo³, Dahua Lin^{1,4}, Hongyang Li^{1†}

¹ Shanghai AI Laboratory ² SenseTime Research

³ The University of Hong Kong ⁴ The Chinese University of Hong Kong

*Equal contribution †Project lead

<https://github.com/OpenDriveLab/OccNet>

<https://github.com/OpenDriveLab/OpenScene>

Abstract

Human driver can easily describe the complex traffic scene by visual system. Such an ability of precise perception is essential for driver’s planning. To achieve this, a geometry-aware representation that quantizes the physical 3D scene into structured grid map with semantic labels per cell, termed as 3D Occupancy, would be desirable. Compared to the form of bounding box, a key insight behind occupancy is that it could capture the fine-grained details of critical obstacles in the scene, and thereby facilitate subsequent tasks. Prior or concurrent literature mainly concentrate on a single scene completion task, where we might argue that the potential of this occupancy representation might obsess broader impact. In this paper, we propose *OccNet*, a multi-view vision-centric pipeline with a cascade and temporal voxel decoder to reconstruct 3D occupancy. At the core of *OccNet* is a general occupancy embedding to represent 3D physical world. Such a descriptor could be applied towards a wide span of driving tasks, including detection, segmentation and planning. To validate the effectiveness of this new representation and our proposed algorithm, we propose *OpenOcc*, the first dense high-quality 3D occupancy benchmark built on top of *nuScenes*. Empirical experiments show that there are evident performance gain across multiple tasks, e.g., motion planning could witness a collision rate reduction by 15%-58%, demonstrating the superiority of our method.

1. Introduction

When you are driving on the road, how would you describe the scene in 3D space through your eyes? Human driver can easily describe the environment by “There is a Benz on the left side of my car in around 5 inches”, “There is a truck carrying huge protruding gas pipe on the rear,

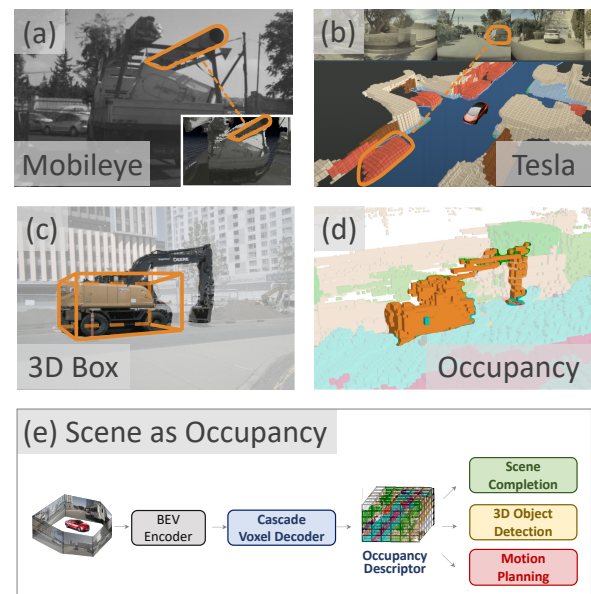


Figure 1. **Scene as Occupancy.** Representing objects as ViDAR (a) or 3D occupancy (b) has been endorsed by industry [1, 2], due to the fact that conventional 3D bounding box cannot describe in detail irregular vehicles in daily driving scenes, e.g., protruding tail in (a) or (c). Defining the 3D world as *Occupancy* in (d) serves better to represent obstacles and avoid collision. In this paper, we envision *Occupancy* as a general *Scene Descriptor* as in (e) for a wide span of driving tasks *beyond* detection, such as planning, and witness performance gain compared to previous alternatives.

in around 50 meters ahead” and so on. Having the ability to describe the real world in a “There is” form is essential for making safe autonomous driving (AD) a reality. This is non-trivial for vision-centric AD systems due to the diverse range of entities present in the *Scene*, including vehicles such as cars, SUVs, and construction trucks, as well as static barriers, pedestrians, background buildings and vegetation.

Quantizing the 3D scene into structured cells with semantic labels attached, termed as *3D Occupancy*, is an intuitive solution, and this form is also advocated in the industry communities such as Mobileye [1] and Tesla [2]. Compared to the 3D box that oversimplifies the shape of objects, 3D occupancy is geometry-aware, depicting different objects and background shapes via the 3D cube collections with different geometric structure. As illustrated in Figure 1(c-d), 3D box can only describe the main body of the construction vehicle, while 3D occupancy can preserve the detail of its crane arm. Other conventional alternatives, such as point cloud segmentation and bird’s-eye-view (BEV) segmentation, while being widely deployed in the context of AD, have their limitations in cost and granularity, respectively. A detailed comparison can be referred in Table 1. Such evident advantages of 3D occupancy encourage an investigation into its potential for augmenting conventional perception tasks and downstream planning.

Similar works have discussed 3D occupancy at an initial stage. Occupancy grid map, a similar concept in Robotics, is a typical representation in mobile navigation [29] but only serves as the search space of planning. 3D semantic scene completion (SSC) [33] can be regarded as a perception task to evaluate the idea of 3D occupancy. Exploiting temporal information as geometric prior is intuitive for the vision-centric models to reconstruct the geometry-aware 3D occupancy, yet previous attempts [19, 21, 5, 26] have failed to address this. A coarse-to-fine approach is also favorable in improving 3D geometric representation at affordable cost, while it is ignored by one-stage methods [19, 26, 5]. In addition, the community still seeks a practical approach to evaluate 3D occupancy in a full-stack autonomous driving spirit as vision-centric solutions [16, 8] prevail.

Towards these issues aforementioned, we propose *OccNet*, a multi-view vision-centric pipeline with a cascade voxel decoder to reconstruct 3D occupancy with the aid of temporal clues, and task-specific heads supporting a wide range of driving tasks. The core of *OccNet* is a compact and representative 3D occupancy embedding to describe the 3D scene. To achieve this, unlike straightforward voxel feature generation from image features or sole use of BEV feature as in previous literature [22, 7, 37], *OccNet* employs a cascade fashion to decode 3D occupancy feature from BEV feature. The decoder adopts a progressive scheme to recover the height information with voxel-based temporal self-attention and spatial cross-attention, bundled alongside a deformable 3D attention module for efficiency. Equipped with such a 3D occupancy descriptor, *OccNet* simultaneously supports general 3D perception tasks and facilitates downstream planning task, *i.e.*, 3D occupancy prediction, 3D detection, BEV segmentation, and motion planning. For fair comparison across methods, we build *OpenOcc*, a 3D occupancy benchmark with dense and high-quality annota-

Representation	Output space	Foreground object	Background & Mapping	Description Granularity	Require point cloud input
3D Box	3D	✓	-	0.4 ~ 12m	-
BEV Seg.	BEV	✓	✓	0.5m ~ 1m	-
Point cloud	3D	✓	-	~ 0.02m	✓
3D Occupancy	3D	✓	✓	0.25 ~ 0.5m	-

Table 1. **Comparison on different representations.** 3D Occupancy unifies foreground objects and background stuff into a fine-grain and dense voxel space, and is input-modality-agnostic.

tions, based on nuScenes dataset [4, 11]. It comprises 34149 annotated frames with over 1.4 billion 3D occupancy cells, each assigned to one of 16 classes to describe foreground objects and background stuff. Such dense and semantic-rich annotations leverage vision models towards superior 3D geometry learning, compared to the sparse alternative. It takes object motion into consideration with directional flow annotations as well, being extensible to the planning task.

We evaluate *OccNet* on *OpenOcc* benchmark, and empirical studies demonstrate the superiority of 3D occupancy as a scene representation over traditional alternatives from three aspects: 1) *Better perception.* 3D occupancy facilitates the acquisition of 3D geometry from vision-only models, as evidenced by the point cloud segmentation performance comparable with LiDAR-based methods and the enhanced 3D detection performance with occupancy-based pre-training or joint-training. 2) *Better Planning.* More accurate perception also translates into improved planning performance. 3) *Dense is better.* Dense 3D occupancy proves more effective than sparse form in supervising vision-only models. On the *OpenOcc* benchmark, *OccNet* outperforms state-of-the-art, *e.g.* *TPVFormer* [19], with a relative improvement of 14% in the semantic scene completion task. Compared with *FCOS3D* [38], the detection model performance pre-trained on *OccNet* increases by about 10 points when fine-tuned on small-scale data. For the motion planning task based on 3D occupancy, we can reduce the collision rate by 15%-58% compared with the planning policy based on BEV segmentation or 3D boxes.

To sum up, our contributions are two folds: (1) We propose *OccNet*, a vision-centric pipeline with a cascade voxel decoder to generate 3D occupancy using temporal clues. It better captures the fine-grained details of the physical world and supports a wide range of driving tasks. (2) Based on the proposed *OpenOcc* benchmark with dense and high-quality annotations, we demonstrate the effectiveness of *OccNet* with an evident performance gain upon perception and planning tasks. The conclusion is that 3D occupancy, as scene representation, is superior to conventional alternatives.

2. Related Work

3D object detection [32, 38, 22, 25, 40, 12] adopts 3D boxes as the objective of perception in AD since the box-

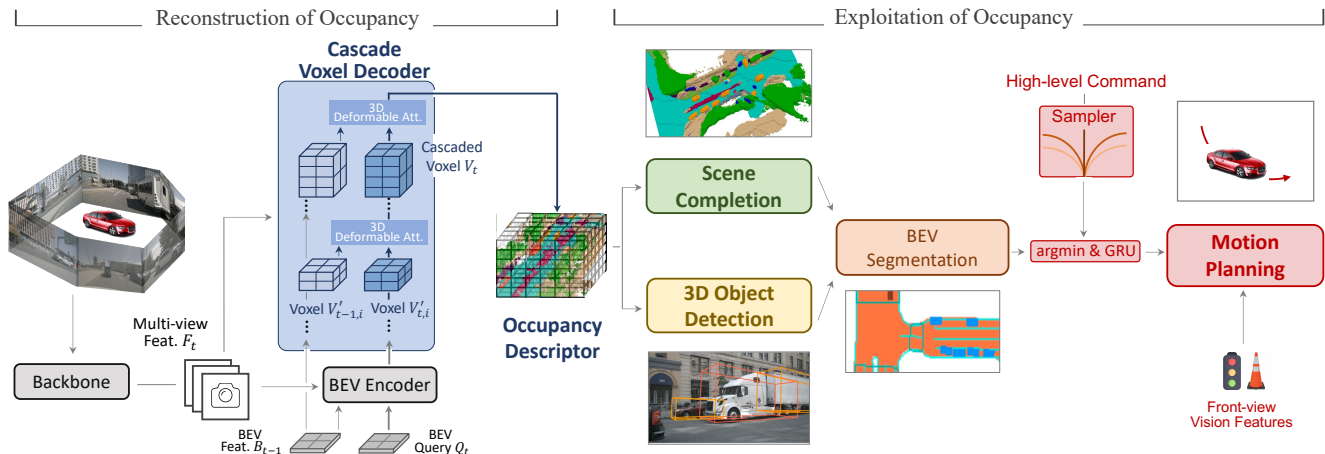


Figure 2. **OccNet pipeline.** The core of OccNet is to obtain a representative Occupancy Descriptor and apply it for various driving tasks. Our proposed algorithm consists of two stages. **I. Reconstruction of Occupancy.** Given multiple visual inputs, we first generate features from the BEV encoder. Voxel Decoder is performed in a cascade fashion where voxels are refined progressively. A 3D deformable attention (att.) unit serves similar functionality as does in 2D case. Temporal voxels V_{t-1} are also incorporated. Some connections are omitted for brevity. See context for details. **II. Exploitation of Occupancy.** Equipped with the occupancy descriptor, we can proceed tasks including semantic scene completion and 3D object detection. Compacting them in BEV space would obtain a BEV segmentation map, which can be directly fed into the planning pipeline [15]. Such a design can ensure desirable improvement in planning task.

form is well structured for downstream rule-based approaches. Such a representation abstracts 3D objects with different shapes into standardized cuboids, hence only cares about foreground objects and oversimplifies object shape. In contrast, 3D occupancy is a fine-grained description of the physical world and can differentiate objects with various shapes.

LiDAR segmentation [42, 28] is tasked as point-level 3D scene understanding. It requires point cloud as input, which is expensive and less portable. Since LiDAR inherently suffers from limited sensing range and sparsity in 3D scene description, it is not friendly to holistic 3D scene semantic understanding [33] using such a pipeline.

3D reconstruction and rendering. Inferring the 3D geometry of objects or scenes from 2D images [13, 27] is prevailing yet challenging for many years in computer vision. Most approaches in this domain [30, 6, 34, 36] cope with a *single* object or scene. For AD application, this is not feasible since it requires strong generalization ability. Note that 3D reconstruction and rendering concentrates more on the quality of the scene geometry and visual appearance. It pays less attention to model efficiency and semantic understanding.

Semantic Scene Completion. The definition of occupancy prediction discussed in this work shares the most resemblance with SSC [33]. MonoScene [5] first adopts U-Net to infer from a single monocular RGB image the dense 3D occupancy with semantic labels. There is a burst of related works released in arXiv recently. We deem them as concurrent and briefly discuss below. VoxFormer [21] utilizes the

depth estimation to set voxel queries in a two-stage framework. OccDepth [26] also adopts a depth-aware spirit in a stereo setting with distillation to predict semantic occupancy. TPVFormer [19] employs LiDAR-based sparse 3D occupancy as the supervision and proposes a tri-perspective view representation to obtain features. Wang *et al.* [39] provides a well human-crafted occupancy benchmark that could facilitate the community.

Despite different settings from ours with work conducted on Semantic-KITTI [3] and NYUv2 [31] (monocular or RGB-D), prior or concurrent literature unanimously neglect the adoption of temporal context. Utilizing history voxel features is straightforward; it is verified by Tesla [2]. Yet there is no technical details or report to the public. Moreover, we position our work to be the first to investigate occupancy as a general descriptor that could enhance multiple tasks *beyond* detection.

3. Methodology

In this paper, we propose an effective and general framework, named OccNet, which obtains robust occupancy features from images and supports multiple driving tasks, as shown in Figure 2. Our method comprises two stages, *Reconstruction of Occupancy* and *Exploitation of Occupancy*. We term the bridging part as *Occupancy Descriptor*, a unified description of the driving scene.

Reconstruction of Occupancy. The goal of this stage is to obtain a representative occupancy descriptor for supporting downstream tasks. Motivated by the fast development in BEV perception [22, 7, 23], OccNet is designed to exploit

that gain for the voxel-wise prediction task in 3D space. To achieve this, the sole usage of BEV feature in downstream tasks, as the simplest architecture, is not suitable for height-aware task in 3D space. Going from one extreme to another, directly constructing voxel feature from images has huge computational cost. We term these two extreme as *BEVNet* and *VoxelNet*, and the design of OccNet finds a balance between them, achieving the best performance with affordable cost. The reconstruction stage first extracts multi-view feature F_t from surrounding images, and feeds them into BEV encoder along with history BEV feature B_{t-1} and current BEV query Q_t to get current BEV feature. The BEV encoder follows the structure of BEVFormer [22], where history BEV feature B_{t-1} , current BEV query Q_t and image feature F_t go through a spatial-temporal-transformer block to get current BEV feature. Then, the image feature, the history and current BEV feature are together decoded into occupancy descriptor via *Cascade Voxel Decoder*. Details of the decoder is presented in Sec. 3.1.

Exploitation of Occupancy. A wide range of driving tasks can be deployed based on the reconstructed occupancy descriptor. Inspired by Uni-AD [16], an explicit design of each representation is preferred. Intuitively, 3D semantic scene completion [33] and 3D object detection are attached upon the occupancy descriptor. Squeezing 3D occupancy grid map and 3D boxes along the height generates a BEV segmentation map. Such a map can be directly fed into motion planning head, along with sampler of high-level command, resulting in the ego-vehicle trajectory via *argmin* and GRU module. Detailed illustration is provided in Sec. 3.2.

3.1. Cascade Voxel Decoder

To obtain a better voxel feature effectively and efficiently, we design a cascade structure in the decoder to progressively recover the height information in voxel feature.

From BEV to Cascaded Voxel. Based on the observation that directly using BEV feature or directly reconstructing voxel feature from perspective view suffers from performance or efficiency drop (see our ablation in Table 9), we break this reconstruction from BEV feature ($B_t \in \mathbb{R}^{H \times W \times C_{\text{BEV}}}$) to the desired voxel feature ($V_t \in \mathbb{R}^{Z \times H \times W \times C_{\text{Voxel}}}$) into N steps, named a cascade structure. Here H and W are the 2D spatial shape of BEV space, C the feature dimension and Z the desired height of voxel space. Between the input BEV feature and the desired cascaded voxel feature, we term the intermediate voxel feature with different height as $V'_{t,i} \in \mathbb{R}^{Z_i \times H \times W \times C_i}$, where Z_i and C_i are uniformly distributed between $\{1, N\}$ and $\{C_{\text{BEV}}, C_{\text{Voxel}}\}$ respectively. As shown in Figure 2, the B_{t-1} and B_t are lifted into $V'_{t-1,i}$ and $V'_{t,i}$ via feed-forward network, go through the i -th voxel decoder to obtain a refined $V'_{t,i}$, and the later steps follow the same scheme. Each voxel decoder comprises voxel-based temporal self-

attention and voxel-based spatial cross-attention modules, and refines $V'_{t,i}$ with history $V'_{t-1,i}$ and image feature F_t respectively. Step by step, the model gradually increases Z_i and decreases C_i to learn the final occupancy descriptor V_t effectively and efficiently.

Voxel based Temporal Self-Attention. The temporal information is crucial to represent the driving scene accurately [22]. Given the history voxel feature $V'_{t-1,i}$, we align it to the current occupancy features $V'_{t,i}$ via the position of ego-vehicle. For a typical self-attention, each query attends to every key and value, so the computation cost is very huge and even increases Z^2 times in 3D space compared to the 2D case. To alleviate the computation cost, we design a voxel-based efficient attention, termed as *3D Deformable Attention* (3D-DA in short), to handle the computational burden. By applying it in the voxel-based temporal self-attention, we ensure that each voxel query only needs to interact with local voxels of interest, making the computational cost affordable.

3D Deformable Attention. We extend the traditional 2D deformable attention [41] to 3D form. Given a voxel feature $V'_{t,i} \in \mathbb{R}^{Z_i \times H \times W \times C_i}$, a voxel query with feature $\mathbf{q} \in \mathbb{R}_i^C$ and 3D referent point \mathbf{p} , the 3D deformable attention is represented by:

$$\text{3D-DA}(\mathbf{q}, \mathbf{p}, V'_{t,i}) = \sum_{m=1}^M W_m \sum_{k=1}^K A_{mk} W'_k V'_{t,i}(\mathbf{p} + \Delta \mathbf{p}_{mk}), \quad (1)$$

where M is the number of attention heads, K is the sampled key number with $K \ll Z_i H W$, $W_m \in \mathbb{R}^{C_i \times (C_i/M)}$ and $W'_k \in \mathbb{R}^{(C_i/M) \times C_i}$ are the learning weights, A_{mk} is the normalized attention weight, and $\mathbf{p} + \Delta \mathbf{p}_{mk}$ is the learnable sample point position in 3D space, in which the feature is computed by trilinear interpolation from the voxel feature.

Voxel-based Spatial Cross-Attention. In the cross attention, the voxel feature $V'_{t,i}$ interacts with the multi-scale image features F_t with 2D deformable attention [41]. Each i -th decoder directly samples $N_{\text{ref},i}$ 3D points from the corresponding voxel to the image view, and interact with the sampled image feature. Such a design maintains the height information and ensures the learning of voxel-wise feature.

3.2. Exploiting Occupancy on Various Tasks

The OccNet depicts the scene in 3D space with fine-grained occupancy descriptor, which can be fed into various driving tasks without excessive computational overhead.

Semantic Scene Completion. For simplicity, we design the MLP head to predict the semantic label of each voxel, and apply the Focal loss [24] to balance the huge numerical inequality between occupied and empty voxels. In addition, the flow head with L_1 loss are attached to estimate the flow velocity per occupied voxels.

3D Object Detection. Inspired by the head design in BEVFormer [22], we compact the occupancy descriptor into BEV, then apply a query-based detection head (an invariant of Deformable DETR [41]) to predict the 3D boxes.

BEV segmentation. Following the spatial-temporal-fusion perception structure in ST-P3 [15], map representation and semantic segmentation are predicted from the BEV feature as in 3D object detection. The BEV segmentation head includes the drivable-area head and the lane head for map representation, the vehicle segmentation head and the pedestrian segmentation head for semantic segmentation.

Motion Planning. For motion planning task, either the predicted occupancy results in SSC or 3D bounding box can be transformed into the BEV segmentation, as shown in 2. The 3D occupancy results is squeezed along the height dimension and the 3D boxes as well. All the semantic labels per BEV cell from either 3D occupancy or 3D boxes are turned into a 0-1 format, where 1 indicates the cell is occupied and 0 for empty. Then, such a BEV segmentation map is applied to the safety cost function, and we compute the safety, comfort and progress cost on the sampled trajectories. Note that compared to 3D boxes, the richer background information in occupancy scene completion leads to the more comprehensive safety cost function, and thus the safety cost value is needed to be normalized between these two kinds of BEV segmentation. All candidate trajectories are sampled by random velocity, acceleration, and curvature. Under the guidance of high-level command including forward, turn left and turn right, the trajectory corresponding to the specific command with the lowest cost will be output. GRU refinement enabled with the front-view vision feature is further performed on this trajectory as ST-P3 [15] to obtain the final trajectory.

4. OpenOcc: 3D Occupancy Benchmark

To fairly evaluate the performance of occupancy across literature, we introduce the first 3D occupancy benchmark named as *OpenOcc* built on top of the prevailing nuScenes dataset [4]. Compared with existing counterparts such as SemanticKITTI [3] with only front camera, OpenOcc provides surrounding camera views with the corresponding 3D occupancy and flow annotations.

4.1. Benchmark Overview

We generate occupancy data with dense and high quality occupancy annotations utilizing the sparse LiDAR information and 3D boxes. It comprises 34149 annotated frames for all 700 training and 150 validation scenes. We annotate over 1.4 billion voxels and 16 classes in the benchmark, including 10 foreground objects and 6 background stuffs. Moreover, we take the foreground object motion into consideration with additional flow annotation of object voxels. We compare our occupancy data with other bench-

Dataset	Multi-view	Scenes	Flow	Density
SemanticKITTI [3]	-	22	-	-
SparseOcc [19]	✓	850	-	~ 0.11
OccData [10]	✓	850	-	~ 0.76
OpenOcc (Ours)	✓	850	✓	1

Table 2. **Comparison of OpenOcc with existing benchmarks.** *Multi-view* denotes the dataset that use multi-view image as input. *Flow* represent the flow annotation is given in the dataset. The density measures the voxel density in the dataset.




Figure 3. **Visual comparison on 3D occupancy annotations.** Compared to (a) sparse occupancy [19] and (b) OccData [10], we generate (c) dense and high-quality annotations with (d) the additional flow annotation of foreground objects, which can be applied for motion planning.

mark in Table 2, indicating that our benchmark can provide the most complete representation of the scene including the occupancy and flow information. As depicted in Figure 3, SparseOcc [19] only utilized the sparse key frame LiDAR data to voxelize the 3D space, which is too sparse to represent the 3D scene. In comparison, our occupancy can represent the complete scene with flow information and capture the local fine grained scene geometry with high quality.

4.2. Generating High-quality Annotation

Independent Accumulation of Background and Foreground. To generate dense representation, it is intuitive to accumulate all sparse LiDAR points from the key frame and intermediate frame to obtain the dense representation [3]. However, directly accumulating points from intermediate frame by coordinate transformation is problematic owing to the existence of moving objects. We propose to split the LiDAR point into the static background points and foreground points based on 3D box and accumulate them separately. Then we can accumulate static background points in the global world system and object points in the object coordinate system to generate dense points.

Method	Backbone	IoU _{geo}	mIoU	barrier	bicycle	bus	car	const. veh.	motorcycle	pedestrian	traffic cone	trailer	truck	driv. surf.	other flat	sidewalk	terrain	manmade	vegetation
BEVDet4D [17]	ResNet50	18.27	9.85	13.56	0.00	13.04	26.98	0.61	1.20	6.76	0.93	1.93	12.63	27.23	11.09	13.64	12.04	6.42	9.56
BEVDepth [20]	ResNet50	23.45	11.88	15.15	0.02	20.75	27.05	1.10	2.01	9.69	1.45	1.91	14.31	31.92	7.88	17.08	16.27	8.76	14.75
BEVDet [18]	ResNet50	27.46	12.49	16.06	0.11	18.27	21.09	2.62	1.42	7.78	1.08	3.4	13.76	33.89	10.84	17.55	22.03	11.72	18.15
OccNet (ours)	ResNet50	37.69	19.48	20.63	5.52	24.16	27.72	9.79	7.73	13.38	7.18	10.68	18.00	46.13	20.6	26.75	29.37	16.90	27.21
Occ3D* [35]	ResNet101	33.93	21.68	27.55	9.45	28.67	36.17	10.69	14.23	17.55	12.41	11.24	22.97	44.15	23.54	28.10	25.11	11.88	23.20
TPVFormer* [19]	ResNet101	37.47	23.67	27.95	12.75	33.24	38.70	12.41	17.84	11.65	8.49	16.42	26.47	47.88	25.43	30.62	30.18	15.51	23.12
OccNet (ours)	ResNet101	41.08	26.98	29.77	16.89	34.16	37.35	15.58	21.92	21.29	16.75	16.37	26.23	50.74	27.93	31.98	33.24	20.80	30.68

Table 3. **3D Occupancy Prediction in terms of Semantic Scene Completion.** The semantic occupancy prediction and geometric prediction metrics are compared for models with RGB input. OccNet significantly outperforms previous SOTAs in terms of mIoU and IoU_{geo}. Methods with * stands for training and evaluating on OpenOcc dataset. The Occ3D is implemented on its baseline method.

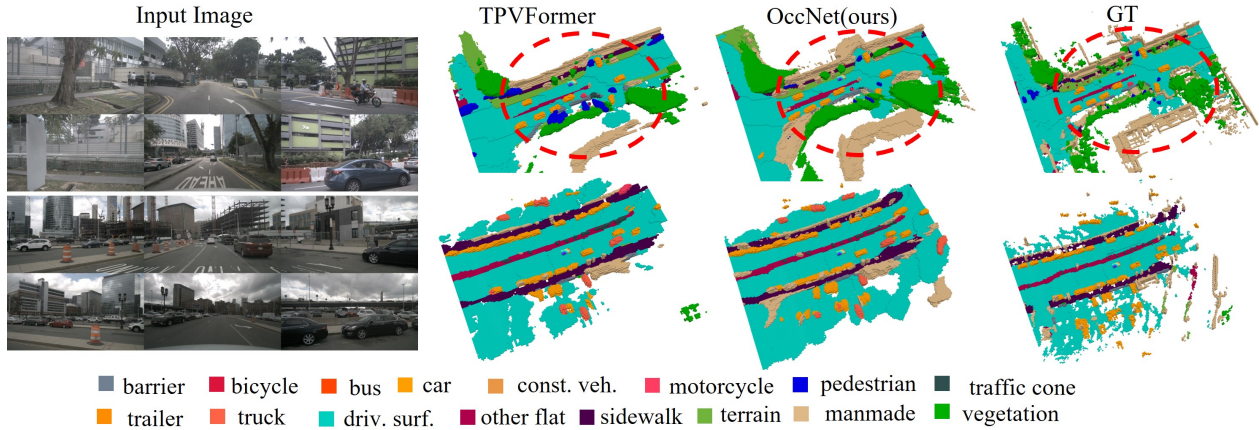


Figure 4. **Qualitative results of occupancy prediction.** Our method outperforms TPVFormer [19] in terms of scene details and the semantic classification accuracy of foreground objects, such as the pedestrian in the dashed region.

Generation of Annotation. Given dense background and object points, we first voxelize the 3D space and label the voxel based on the majority vote of labelled points in the voxel. Different with existing benchmark with only occupancy labels, we annotate the flow velocity of voxel based on the 3D box velocity to facilitate the downstream task such as motion planning. Only Using key frame will cause sparsity of generated occupancy data, thus we annotate the voxel with unlabeled LiDAR points from intermediate frame based on the surrounding labelled voxels to further improve the data density. In addition, as nuScenes has the issue of missing translation in the z-axis, we refine the occupancy data by completing the scene, such as filling the holes on the road for higher quality. Moreover, we set part of voxels as invisible from the camera view by tracing the ray, which is more applicable for the task with camera input.

5. Experiments

Benchmark Details. We select a volume of $\mathcal{V} = [-50\text{m}, 50\text{m}] \times [-50\text{m}, 50\text{m}] \times [-5\text{m}, 3\text{m}]$ in LiDAR coordinate system for occupancy data generation, and vox-

elize the 3D space by the resolution of $\Delta s = 0.5\text{m}$ into $200 \times 200 \times 16$ voxels to represent the 3D space. Evaluation metric can be referred in Supplementary.

OccNet Details. Following the experimental setting of BEVFormer [22], we use two types of backbone: ResNet50 [14] initialized from ImageNet [9], and ResNet101-DCN [14] initialized from FCOS3D [38]. We define the BEV feature as B_t with $H = 200$, $W = 200$, and $C_{\text{BEV}} = 256$. For the decoder, we design $N = 4$ occupancy feature maps $V'_{t,i} \in \mathbb{R}^{Z_i \times H \times W \times C_i}$ with $Z_i = 2^i$, $C_1 = C_2 = 128$, $C_3 = C_4 = 64$. For the voxel-spatial cross attention, we sample $N_{\text{ref},i} = 4$ points in each queried voxel. By default, we train OccNet with 24 epochs with a learning rate of 2×10^{-4} .

5.1. Main Results

Semantic Scene Completion. We compare OccNet with previous state-of-the-art methods for semantic scene completion task in Table 3 and Figure 4. We reproduce BEVDet4D [17], BEVDepth [20] and BEVDet [18] by replacing the detection head with the scene completion head built on their BEV feature maps, and OccNet outperforms

Method	Input	mIOU	barrier	bicycle	bus	car	const. veh.	motorcycle	pedestrian	traffic cone	trailer	truck	driv. surf.	other flat	sidewalk	terrain	manmade	vegetation
RangeNet++ [28]	LiDAR	65.50	66.00	21.30	77.20	80.90	30.20	66.80	69.60	52.10	54.20	72.20	94.10	66.60	63.50	70.10	83.10	79.80
Cylinder3D [42]	LiDAR	76.10	76.40	40.30	91.20	93.80	51.30	78.00	78.90	64.90	62.10	84.40	96.80	71.60	76.40	75.40	90.50	87.40
TPVFormer* [19]	Camera	58.45	65.99	24.50	80.88	74.28	47.04	47.09	33.42	14.52	53.96	70.79	88.55	61.63	59.46	63.15	75.76	74.17
OccNet (ours)	Camera	60.46	66.95	32.58	77.37	73.88	37.62	50.87	51.45	33.69	52.20	67.08	88.72	57.99	58.04	63.06	78.91	76.97

Table 4. **The performance of OccNet (ResNet101) on nuScenes validation set for LiDAR segmentation task.** OccNet with camera input is comparable with LiDAR based method. Methods with * stands for training from scratch on OpenOcc dataset.

these methods by a large margin as shown in Table 3. Compared with BEV feature map, our occupancy descriptor is better for the voxel-wise prediction task. We also compare OccNet with TPVFormer [19], which is developed for surrounding 3D semantic occupancy prediction task, and our model surpasses it over 3.31 points in terms of mIOU (26.98 vs. 23.67), indicating that occupancy descriptor is better than TPV features for scene representation. Note that TPVFormer surpasses the OccNet in car, truck and trailer, because samples of these three objects are relatively large in the benchmark and TPVFormer learns better feature on these classes from their sampling strategy. However, for the objects with small size such as pedestrian and traffic cone, our method can outperform TPVFormer [19] with a large margin of 10 points in Table 3.

Occupancy for LiDAR Segmentation. Occupancy is a voxelized representation of points in 3D space, and semantic scene completion is equivalent to semantic LiDAR prediction task when $\Delta s \rightarrow 0$. We transfer semantic occupancy prediction to LiDAR segmentation by assigning the point label based on associated voxel label, and then evaluate the model on the mIOU metric. As reported in Table 4, given camera as input without LiDAR supervision, OccNet can be comparative with the LiDAR segmentation model RangeNet++ [28] in terms of mIOU (60.46 vs. 65.50), and OccNet can even outperforms RangeNet++ in the IoU of bicycle (32.58 vs. 21.30). Compared with TPVFormer [19], OccNet also outperforms it with 2 points in mIOU.

Occupancy for 3D Detection. In the scene completion task, the location of foreground object can be coarse regressed, which can help the 3D detection task with 3D box regression. As shown in Table 5, the joint training of scene completion and 3D detection task can improve the detector performance for all our three models, including BEVNet, VoxNet and OccNet, in terms of mAP and NDS. Note that the voxelized representation of occupancy with $\Delta s = 0.5m$ is too coarse when calculating the metric dependent on the precise center distance and IoU of 3D box, and thus mATE and mASE is a little increased with joint training.

Pretrained Occupancy for 3D Detection and BEV segmentation. The OccNet trained on semantic scene completion task can obtain general representation for 3D space

Method	Joint	mAP \uparrow	NDS \uparrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow	mATE \downarrow	mASE \downarrow
BEVNet	-	0.259	0.377	0.600	0.592	0.216	0.828	0.290
	\checkmark	0.271	0.390	0.578	0.541	0.211	0.835	0.293
VoxNet	-	0.271	0.380	0.603	0.616	0.219	0.832	0.284
	\checkmark	0.277	0.387	0.586	0.614	0.203	0.828	0.285
OccNet	-	0.276	0.382	0.655	0.588	0.209	0.817	0.290
	\checkmark	0.276	0.390	0.585	0.570	0.190	0.842	0.292

Table 5. **Joint training of 3D occupancy and 3D detection.** Results reported on nuScenes validation set show that joint training of 3D occupancy and 3D detection can help the latter task.




Figure 5. **The comparison of detector performance using different pretrained models and different scale of training dataset.** OccNet (sparse) and OccNet (dense) means the OccNet trained on sparse and dense occupancy data respectively. Best view in color.

owing to the scene reconstructed in the occupancy descriptor. Thus, the learned occupancy descriptor can be directly transferred to the downstream 3D perception tasks with model fine-tuning. As described in Figure 5, the model performance on 3D detection with pretrained OccNet is superior to that pretrained on FCOS3D [38] detector in different scales of training dataset with the performance gain about 10 points for mAP and NDS. We also compare the occupancy pretraining and detection pretraining for the BEV segmentation task, indicating that the occupancy pretraining can help BEV segmentation achieve higher IoU in the fine-tuning stage on both semantic and map segmentation as shown in Table 6.

Occupancy for Planning. With the prediction results from upstream tasks, i.e., bounding box and occupancy, the final trajectory can be obtained through a cost filter and a

Task	Main value	Drivable area	Lane	Vehicle	Pedestrian
Det	18.18	44.59	13.62	12.29	2.21
Occ	19.17	47.21	13.83	12.91	2.74

Table 6. **Different pretraining tasks for BEV segmentation.** Occupancy task can help BEV segmentation task achieve higher IoU.

Input	Collision (%)↓			L2 (m)↓		
	1s	2s	3s	1s	2s	3s
Bbox GT	0.23	0.66	1.50	1.32	2.16	3.00
Occupancy GT	0.20	0.56	1.30	1.29	2.13	2.98
Segmentation pred. [15]	0.50	0.88	1.49	1.39	2.21	3.02
Bbox pred. (OccNet)	0.27	0.68	1.59	1.32	2.17	3.03
Occupancy pred. (OccNet)	0.21	0.55	1.35	1.31	2.18	3.07
Occupancy pred. (OccNet, full)	0.21	0.59	1.37	1.29	2.13	2.99

Table 7. **Planning results with different scene representations.** Occupancy representation helps the planning task achieve a lower collision rate and more accurate L2 distance in all time intervals.

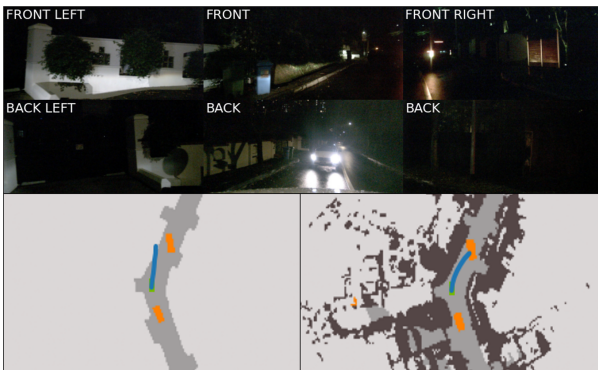


Figure 6. **Visualization of planning.** The blue line represents the planned trajectory, and the lower figures are rasterisation results of bounding box and occupancy, respectively.

GRU refinement module [15] with the BEV segmentation inputs. To obtain these segmentation results, we rasterize the outputs of our OccNet in BEV space. We compare the rasterisation results of bounding box and occupancy by using the predictions from OccNet. We also compare our results with the direct segmentation from ST-P3 [15]. For a fair comparison, we follow the same setup as ST-P3 with only vehicle and pedestrian classes kept. We also add the ground truth rasterisation inputs for better comparison. As shown in Table 7, the best performance can be obtained by using ground truth of occupancy to filter trajectories. For predicted results, the collision rate can be reduced by 15% - 58% based on the occupancy prediction from OccNet. We also conduct the experiment using all 16 classes of occupancy, which shows that full classes of occupancy can bring the performance improvement on L2 distance. As shown in the Figure 6, planning with full classes of occupancy can make decisions within the feasible areas to avoid collisions from the background objects.

Method	mIOU	IoU _{geo}	Params↓	Memory↓	FPS↑
BEVNet	17.37	36.11	39M	8G	4.5
VoxelNet	19.06	37.59	72M	23G	1.9
OccNet	19.48	37.69	40M	18G	2.6

Table 8. **Efficiency and performance analysis with model structure.** The evaluation is measured on a V100 GPU.

Task	Δs	mIoU _{ir}	truck	trailer	cons. veh.
Det	0.5m	15.92	23.19	9.57	15.00
Occ	0.5m	18.13 (+2.21)	25.59	12.29	16.51
Det	0.25m	9.90	14.84	5.10	9.75
Occ	0.25m	13.41 (+3.51)	18.85	7.14	14.25

Table 9. **The comparison of detection task and occupancy task on the recognition of irregular object.** mIoU_{ir} denotes mean IoU of truck, trailer and construction vehicle.




Figure 7. **Visualization of 3D box and occupancy prediction.** Compared with 3D box, occupancy representation can identify the irregular object well.

5.2. Discussion

Model Efficiency. In Table 8, we compare the performance of different models in the semantic scene completion task. Compared with BEVNet and VoxelNet, OccNet can obtain the best performance in terms of mIOU and IoU_{geo} with efficiency and effectiveness.

Irregular Object. Representing the irregular object such as construction vehicle with 3D box or the background stuff such as traffic sign is difficult and inaccurate as indicated in Figure 7. We transform 3D box into voxel to compare the 3D detection and occupancy task on irregular object in Table 9, verifying that occupancy can describe the irregular object better. To study the effect of voxel size, we also generate the dataset with $\Delta s = 0.25m$. With the decrease of Δs from 0.5m to 0.25m, the performance gap between 3D box and occupancy increases because the finer granularity can better depict the irregular object.

Dense v.s. Sparse Occupancy. Compared with sparse occupancy, dense occupancy can help depict the complete geometry of background and foreground object in detail as shown in Figure 3. Intuitively, dense occupancy is better for 3D perception and motion planning owing to more abundant information input. We validate that model pretrained on dense occupancy can benefit the downstream 3D detection task more as shown in Figure 5.

6. Conclusion

We dive into the potential of the 3D occupancy as scene representation and propose a general framework OccNet to evaluate the idea. The experiments on various downstream tasks validate the effectiveness of our method. The OpenOcc benchmark with dense and high-quality labels is also provided for community.

Limitations and future work. Currently, the annotation is still based on the well-established dataset. Utilizing self-supervised learning to further reduce the human-annotation cost is a valuable direction. We hope occupancy framework can be the foundation model of autonomous driving.

Acknowledgement

The paper is jointly supported by the National Key R&D Program of China No.2022ZD0161000, the General Research Fund of Hong Kong No.17200622, Shanghai Committee of Science and Technology (21DZ1100100), and NSFC (62206172). We thank reviewers for fruitful comments and the research community for participating the Occupancy Prediction Challenge 2023.

References

- [1] CES 2020 by Mobileye . <https://youtu.be/HPWGFzqd7pI>, 2020. 1, 2
- [2] Tesla AI Day. <https://www.youtube.com/watch?v=j0z4FweCy4M>, 2021. 1, 2, 3
- [3] 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 *ICCV*, pages 9297–9307, 2019. 3, 5
- [4] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nusenes: A multimodal dataset for autonomous driving. In *CVPR*, pages 11621–11631, 2020. 2, 5
- [5] Anh-Quan Cao and Raoul de Charette. Monoscene: Monocular 3d semantic scene completion. In *CVPR*, pages 3991–4001, 2022. 2, 3
- [6] Anh-Quan Cao and Raoul de Charette. Scenerf: Self-supervised monocular 3d scene reconstruction with radiance fields. *arXiv preprint arXiv:2212.02501*, 2022. 3
- [7] Li Chen, Chonghao Sima, Yang Li, Zehan Zheng, Jiajie Xu, Xiangwei Geng, Hongyang Li, Conghui He, Jianping Shi, Yu Qiao, et al. Persformer: 3d lane detection via perspective transformer and the openlane benchmark. In *ECCV*, pages 550–567. Springer, 2022. 2, 3
- [8] Li Chen, Penghao Wu, Kashyap Chitta, Bernhard Jaeger, Andreas Geiger, and Hongyang Li. End-to-end autonomous driving: challenges and frontiers. *arXiv preprint arXiv:2306.16927*, 2023. 2
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 6
- [10] Ming Fang and Zhiqi Li. occupancy-for-nuscenes. <https://github.com/FANG-MING/occupancy-for-nuscenes>, 2023. 5
- [11] Whye Kit Fong, Rohit Mohan, Juana Valeria Hurtado, Lubing Zhou, Holger Caesar, Oscar Beijbom, and Abhinav Valada. Panoptic nuscenes: A large-scale benchmark for lidar panoptic segmentation and tracking. In *ICRA*, 2022. 2
- [12] Yulu Gao, Chonghao Sima, Shaoshuai Shi, Shangzhe Di, Si Liu, and Hongyang Li. Sparse dense fusion for 3d object detection. *arXiv preprint arXiv:2304.04179*, 2023. 2
- [13] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004. 3
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [15] Shengchao Hu, Li Chen, Penghao Wu, Hongyang Li, Junchi Yan, and Dacheng Tao. St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning. In *ECCV*, pages 533–549. Springer, 2022. 3, 5, 8
- [16] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, Lewei Lu, Xiaosong Jia, Qiang Liu, Jifeng Dai, Yu Qiao, and Hongyang Li. Planning-oriented autonomous driving. In *CVPR*, 2023. 2, 4
- [17] Junjie Huang and Guan Huang. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *arXiv preprint arXiv:2203.17054*, 2022. 6
- [18] Junjie Huang, Guan Huang, Zheng Zhu, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021. 6
- [19] 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. 2, 3, 5, 6, 7
- [20] Yinhao 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. 6
- [21] 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. *arXiv preprint arXiv:2302.12251*, 2023. 2, 3
- [22] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *ECCV*, pages 1–18. Springer, 2022. 2, 3, 4, 5, 6
- [23] Tingting Liang, Hongwei Xie, Kaicheng Yu, Zhongyu Xia, Zhiwei Lin, Yongtao Wang, Tao Tang, Bing Wang, and Zhi Tang. Bevfusion: A simple and robust lidar-camera fusion framework. *arXiv preprint arXiv:2205.13790*, 2022. 3
- [24] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. 4

- [25] Zhijian Liu, Haotian Tang, Alexander Amini, Xingyu Yang, Huizi Mao, Daniela Rus, and Song Han. Bvffusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. *arXiv preprint arXiv:2205.13542*, 2022. [2](#)
- [26] Ruihang Miao, Weizhou Liu, Mingrui Chen, Zheng Gong, Weixin Xu, Chen Hu, and Shuchang Zhou. Occdepth: A depth-aware method for 3d semantic scene completion. *arXiv:2302.13540*, 2023. [2](#), [3](#)
- [27] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. [3](#)
- [28] Andres Milioto, Ignacio Vizzo, Jens Behley, and Cyrill Stachniss. Rangenet++: Fast and accurate lidar semantic segmentation. In *IROS*, pages 4213–4220. IEEE, 2019. [3](#), [7](#)
- [29] Julien Moras, J. Dezert, and Benjamin Pannetier. Grid occupancy estimation for environment perception based on belief functions and pcr6. volume 9474, 04 2015. [2](#)
- [30] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022. [3](#)
- [31] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *ECCV*, 2012. [3](#)
- [32] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaoqiang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *CVPR*, pages 10529–10538, 2020. [2](#)
- [33] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *CVPR*, pages 1746–1754, 2017. [2](#), [3](#), [4](#)
- [34] Matthew Tancik, Vincent Casser, Xinchen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P. Srinivasan, Jonathan T. Barron, and Henrik Kretschmar. Block-nerf: Scalable large scene neural view synthesis. In *CVPR*, pages 8248–8258, June 2022. [3](#)
- [35] Xiaoyu Tian, Tao Jiang, Longfei Yun, Yue Wang, Yilun Wang, and Hang Zhao. Occ3d: A large-scale 3d occupancy prediction benchmark for autonomous driving. *arXiv preprint arXiv:2304.14365*, 2023. [6](#)
- [36] Wenwen Tong, Jiangwei Xie, Tianyu Li, Hanming Deng, Xiangwei Geng, Ruoyi Zhou, Dingchen Yang, Bo Dai, Lewei Lu, and Hongyang Li. 3d data augmentation for driving scenes on camera. *arXiv preprint arXiv:2303.10340*, 2023. [3](#)
- [37] Tai Wang, Jiangmiao Pang, and Dahua Lin. Monocular 3d object detection with depth from motion. In *European Conference on Computer Vision*, pages 386–403. Springer, 2022. [2](#)
- [38] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *ICCV*, pages 913–922, 2021. [2](#), [6](#), [7](#)
- [39] Xiaofeng Wang, Zheng Zhu, Wenbo Xu, Yunpeng Zhang, Yi Wei, Xu Chi, Yun Ye, Dalong Du, Jiwen Lu, and Xingang Wang. Openoccupancy: A large scale benchmark for surrounding semantic occupancy perception. *arXiv preprint arXiv:2303.03991*, 2023. [3](#)
- [40] Jia Zeng, Li Chen, Hanming Deng, Lewei Lu, Junchi Yan, Yu Qiao, and Hongyang Li. Distilling focal knowledge from imperfect expert for 3d object detection. In *CVPR*, pages 992–1001, 2023. [2](#)
- [41] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable DETR: Deformable transformers for end-to-end object detection. In *ICLR*, 2021. [4](#), [5](#)
- [42] Xinge Zhu, Hui Zhou, Tai Wang, Fangzhou Hong, Yuexin Ma, Wei Li, Hongsheng Li, and Dahua Lin. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In *CVPR*, pages 9939–9948, 2021. [3](#), [7](#)