Fully Sparse 3D Occupancy Prediction

Haisong Liu^{1,2,*}[●], Yang Chen^{1,*}, Haiguang Wang¹, Zetong Yang², Tianyu Li², Jia Zeng², Li Chen², Hongyang Li², and Limin Wang^{1,2,⊠}[●]

 $^1\,$ State Key Laboratory for Novel Software Technology, Nanjing University $^2\,$ Shanghai AI Lab

Abstract. Occupancy prediction plays a pivotal role in autonomous driving. Previous methods typically construct dense 3D volumes, neglecting the inherent sparsity of the scene and suffering high computational costs. To bridge the gap, we introduce a novel fully sparse occupancy network, termed SparseOcc. SparseOcc initially reconstructs a sparse 3D representation from visual inputs and subsequently predicts semantic/instance occupancy from the 3D sparse representation by sparse queries. A mask-guided sparse sampling is designed to enable sparse queries to interact with 2D features in a fully sparse manner, thereby circumventing costly dense features or global attention. Additionally, we design a thoughtful ray-based evaluation metric, namely RayIoU, to solve the inconsistency penalty along depths raised in traditional voxel-level mIoU criteria. SparseOcc demonstrates its effectiveness by achieving a RayIoU of 34.0, while maintaining a real-time inference speed of 17.3 FPS, with 7 history frames inputs. By incorporating more preceding frames to 15, SparseOcc continuously improves its performance to 35.1 RayIoU without bells and whistles. Code is available at https://github.com/MCG-NJU/SparseOcc.

Keywords: 3D Occupancy Estimation · Semantic Scene Completion · 3D Reconstruction · Autonomous Driving

1 Introduction

Vision-centric 3D occupancy prediction [1] focuses on partitioning 3D scenes into structured grids from visual images. Each grid is assigned a label indicating if it is occupied or not. This task offers more geometric details than 3D detection and produces an alternative representation to LiDAR [22, 31, 32, 61–64].

Existing methods [16,27,28,45,58] typically construct dense 3D features yet suffer from computational overhead (e.g., $2 \sim 3$ FPS on Tesla A100). However, dense representations are not necessary for occupancy predictions. We statistic the geometry sparsity and find that more than 90% of the voxels are empty. This manifests a large room in occupancy prediction acceleration by exploiting the sparsity. Some works [19,26] explore the sparsity of 3D scenes, but they still

^{*:} Equal contribution.

 $[\]boxtimes$: Corresponding author.



Fig. 1: (a) SparseOcc reconstructs a sparse 3D representation from camera-only inputs by a sparse voxel decoder, and then estimates the mask and label of each segment via a set of sparse queries. (b) Performance comparison on the validation split of Occ3D-nus. FPS is measured on a Tesla A100 with the PyTorch **fp32** backend.

rely on sparse-to-dense modules for dense predictions. This inspires us to seek a pure sparse occupancy network without any dense design.

In this paper, we propose SparseOcc, the first fully sparse occupancy network. As depicted in Fig. 1 (a), SparseOcc includes two steps. First, it leverages a *sparse voxel decoder* to reconstruct the sparse geometry of a scene in a coarseto-fine manner. This only models non-free regions, saving computational costs significantly. Second, we design a *mask transformer* with sparse semantic/instance queries to predict masks and labels of segments from the sparse space. The mask transformer not only improves performance on semantic occupancy but also paves the way for panoptic occupancy. A *mask-guided sparse sampling* is designed to achieve sparse cross-attention in the mask transformer. As such, our SparseOcc fully exploits the sparse property and gets rid of any dense design like dense 3D features, sparse-to-dense modules, and global attention.

Besides, we notice flaws in popular voxel-level mean Intersection-over-Union (mIoU) metrics for occupancy evaluation and further design a ray-level evaluation, RayIoU, as the solution. The mIoU criterion is an ill-posed formulation given the ambiguous labeling of unscanned voxels. Previous methods [50] relieve this issue by only evaluating observed areas but raise extra issues in inconsistency penalty along depths. Instead, RayIoU addresses the two aforementioned issues simultaneously. It evaluates predicted 3D occupancy volume by retrieving depth and category predictions of designated rays. To be specific, RayIoU casts query rays into predicted 3D volumes and decides true positive predictions as the ray with the correct distance and class of its first touched occupied voxel grid. This formulates a more fair and reasonable criterion.

Thanks to the sparsity design, SparseOcc achieves 34.0 RayIoU on Occ3Dnus [50], while maintaining a real-time inference speed of 17.3 FPS (Tesla A100, PyTorch **fp32** backend), with 7 history frames inputs. By incorporating more preceding frames to 15, SparseOcc continuously improves its performance to 35.1 RayIoU, achieving state-of-the-art performance without bells and whistles. The comparison between SparseOcc with previous methods in terms of performance and efficiency is shown in Fig. 1 (b).

We summarize our contributions as follows:

- 1. We propose SparseOcc, the first fully sparse occupancy network without any time-consuming dense designs. It achieves 34.0 RayIoU on Occ3D-nus benchmark with an inference speed of 17.3 FPS.
- 2. We present RayIoU, a ray-wise criterion for occupancy evaluation. By querying rays to 3D volume, it solves the ambiguous penalty issue for unscanned free voxels and the inconsistent depth penalty issue in the mIoU metric.

2 Related Work

Camera-based 3D Occupancy Prediction. Occupancy Networks were originally proposed by Mescheder et al. [37, 42], focusing on continuous object representations in 3D space. Recent variations in occupancy networks [1, 4, 12, 45, 50, 55, 57, 59] mostly draw inspiration from Bird's Eve View (BEV) perception [15-18, 23, 24, 27, 33-35, 54, 56, 60] and predicts voxel-level semantic information from image inputs. For instance, MonoScene [4] completes scene occupancy through a 2D and a 3D UNet [43] connected by a sight projection module. SurroundOcc [58] proposes a coarse-to-fine architecture. However, the large number of voxel queries is computationally heavy. TPVFormer [19] proposes using triperspective view representations to supplement vertical structural information, but this inevitably leads to information loss. VoxFormer [26] initializes sparse queries based on monocular depth prediction. Nevertheless, VoxFormer is not fully sparse as it still requires a sparse-to-dense MAE [13] module to complete the scene. Some methods emerged in the CVPR 2023 occupancy challenge [10,28,40], but none of them exploits a fully sparse design. In this paper, we make the first step to explore the fully sparse architecture for 3D occupancy prediction from camera-only inputs.

Sparse Architectures for 3D Vision. Sparse architectures find widespread adoption in LiDAR-based reconstruction [48] and perception [8,61,62,64], leveraging the inherent sparsity of point clouds. However, when it comes to vision-to-3D tasks, a direct adaptation is not feasible due to the absence of point cloud inputs. A prior work, SparseBEV [33], proposes a fully sparse architecture for 3D object detection. Nevertheless, directly adapting this approach is non-trivial because 3D object detection focuses on a sparse set of objects, whereas 3D occupancy requires dense predictions for each voxel. Consequently, designing a fully sparse architecture for 3D occupancy remains a challenging task.

End-to-end 3D Reconstruction from Posed Images. As a related task to 3D occupancy prediction, 3D reconstruction recovers the 3D geometry from multiple posed images. Recent methods focus on more compact and efficient end-to-end 3D reconstruction pipelines [2, 11, 39, 46, 47]. Atlas [39] extracts features from multi-view input images and maps them to 3D space to construct the truncated



Fig. 2: SparseOcc is a fully sparse architecture since it neither relies on dense 3D feature, nor has sparse-to-dense and global attention operations. The sparse voxel decoder reconstructs the sparse geometry of the scene, consisting of K voxels ($K \ll W \times H \times D$). The mask transformer then uses N sparse queries to predict the mask and label of each segment. SparseOcc can be easily extended to panoptic occupancy by replacing the semantic queries with instance queries.

signed distance function [9]. NeuralRecon [47] directly reconstructs local surfaces as sparse TSDF volumes and uses a GRU-based TSDF fusion module to fuse features from previous fragments. VoRTX [46] utilizes transformers to address occlusion issues in multi-view images.

Mask Transformer. Recently, unified segmentation models have been widely studied to handle semantic and instance segmentation concurrently. Cheng *et al.* first propose MaskFormer [7] for unified segmentation in terms of model architecture, loss functions, and training strategies. Mask2Former [6] then introduces masked attention, with restricted receptive fields on instance masks, for better performance. Later on, Mask3D [44] successfully extends the mask transformer for point cloud segmentation with state-of-the-art performance. Open-Mask3D [49] further achieves the open-vocabulary 3D instance segmentation task and proposes a model for zero-shot 3D segmentation.

3 SparseOcc

SparseOcc is a vision-centric occupancy model that only requires camera inputs. As shown in Fig. 2, SparseOcc has three modules: an image encoder consisting of an image backbone and FPN [30] to extract 2D features from multi-view images; a sparse voxel decoder (Sec. 3.1) to predict sparse class-agnostic 3D occupancy with correlated embeddings from the image features; a mask transformer decoder (Sec 3.2) to distinguish semantics and instances in the sparse 3D space.

3.1 Sparse Voxel Decoder

Since 3D occupancy ground truth [45, 50, 55, 58] is a dense volume with dimensions $W \times H \times D$ (e.g., $200 \times 200 \times 16$), existing methods typically build a dense



Fig. 3: The sparse voxel decoder employs a coarse-to-fine pipeline with three layers. Within each layer, we utilize a transformer-like architecture for 3D-2D interaction. At the end of every layer, the voxel resolution is upsampled by a factor of $2\times$, and probabilities of voxel occupancy are estimated.

3D feature of shape $W \times H \times D \times C$, but suffer from computational overhead. In this paper, we argue that such dense representation is not necessary for occupancy prediction. As in our statistics, we find that over 90% of the voxels in the scene are free. This motivates us to explore a sparse 3D representation that only models the non-free areas of the scene, thereby saving computational resources.

Overall architecture. Our designed sparse voxel decoder is shown in Fig. 3. In general, it follows a coarse-to-fine structure but only models the non-free regions. The decoder starts from a set of coarse voxel queries equally distributed in the 3D space (e.g., 25×25). In each layer, we first upsample each voxel by $2 \times$, e.g., a voxel with size d will be upsampled into 8 voxels with size $\frac{d}{2}$. Next, we estimate an occupancy score for each voxel and conduct pruning to remove useless voxel grids. Here we have two approaches for pruning: one is based on a threshold (e.g., only keeps score > 0.5); the other is by top-k selection. In our implementation, we simply keep voxels with top-k occupancy scores for training efficiency. k is a dataset-related parameter, obtained by counting the maximum number of non-free voxels in each sample at different resolutions. The voxel tokens after pruning will serve as the input for the next layer.

Detailed design. Within each layer, we use a transformer-like [52] architecture to handle voxel queries. The concrete architecture is inspired by SparseBEV [33], a detection method using a sparse scheme. To be specific, in layer l with K_{l-1} voxel queries described by 3D locations and a C-dim content vector, we first use self-attention to aggregate local and global features for those query voxels. Then, a linear layer is used to generate 3D sampling offsets $\{(\Delta x_i, \Delta y_i, \Delta z_i)\}$ for each voxel query from the associated content vector. These sampling offsets are utilized to transform voxel queries to obtain reference points in global coordinates. We finally project those sampled reference points to multi-view image space for integrating image features by adaptive mixing [20,51]. In summary, our approach differs from SparseBEV by shifting the query formulation from pillars to 3D voxels. Other components such as self attention, adaptive sampling and mixing are directly borrowed. 6 Liu et al.

Temporal modeling. Previous dense occupancy methods [16,27] typically warp the history BEV/3D feature to the current timestamp, and use deformable attention [65] or 3D convolutions to fuse temporal information. However, this approach is not applicable in our case, given our sparse 3D features. To handle this, we leverage the flexibility of the aforementioned global sampled reference points by warping them to previous timestamps to sample history multi-view image features. The sampled multi-frame features are stacked and aggregated by adaptive mixing so as for temporal modeling.

Supervision. We compute loss for sparse voxels from each layer. We use binary cross entropy (BCE) loss as the supervision, given that we are reconstructing a class-agnostic sparse occupancy space. Only the kept sparse voxels are supervised, while the discarded regions during pruning in earlier stages are ignored.

Moreover, due to the severe class imbalance, the model can be easily dominated by categories with a large proportion, such as the ground, thereby ignoring other important elements in the scene, such as cars, people, etc. Therefore, voxels belonging to different classes are assigned with different loss weights. For example, voxels belonging to class c are assigned with a loss weight of:

$$w_c = \frac{\sum_{i=1}^C M_i}{M_c},\tag{1}$$

where M_i is the number of voxels belonging to the *i*-th class in ground truth.

3.2 Mask Transformer

Our mask transformer is inspired by Mask2Former [6], which uses N sparse semantic/instance queries decoupled by binary mask queries $\mathbf{Q}_m \in [0, 1]^{N \times K}$ and content vectors $\mathbf{Q}_c \in \mathbb{R}^{N \times C}$. The mask transformer consists of three steps: multi-head self attention (MHSA), mask-guided sparse sampling, and adaptive mixing. MHSA is used for the interaction between different queries as the common practice. Mask-guided sparse sampling and adaptive mixing are responsible for the interaction between queries and 2D image features.

Mask-guided sparse sampling. A simple baseline of mask transformer is to use the masked cross-attention module in Mask2Former. However, it attends to all positions of the key, with unbearable computations. Here, we design a simple alternative. We first randomly select a set of 3D points within the mask predicted by the previous (l-1)-th Transformer decoder layer. Then, we project those 3D points to multi-view images and extract their features by bilinear interpolation. Besides, our sparse sampling mechanism makes the temporal modeling easier by simply warping the sampling points (as done in the sparse voxel decoder).

Prediction. For class prediction, we apply a linear classifier with a sigmoid activation based on the query embeddings \mathbf{Q}_c . For mask prediction, the query embeddings are converted to mask embeddings by an MLP. The mask embeddings

 $\mathbf{M} \in \mathbb{R}^{Q \times C}$ have the same shape as query embeddings \mathbf{Q}_c and are dot-producted with the sparse voxel embeddings $\mathbf{V} \in \mathbb{R}^{K \times C}$ to produce mask predictions. Thus, the prediction space of our mask transformer is constrained to the sparsified 3D space from the sparse voxel decoder, rather than the full 3D scene. The mask predictions will serve as the mask queries \mathbf{Q}_m for the next layer.

Supervision. The reconstruction result from the sparse voxel decoder may not be reliable, as it may overlook or inaccurately detect certain elements. Thus, supervising the mask transformer presents certain challenges since its predictions are confined within this unreliable space. In cases of missed detection, where some ground truth segments are absent in the predicted sparse occupancy, we opt to discard these segments to prevent confusion. As for inaccurately detected elements, we simply categorize them as an additional "no object" category.

Loss Functions. Following MaskFormer [7], we match the ground truth with the predictions using Hungarian matching. Focal loss L_{focal} is used for classification, while a combination of DICE loss [38] L_{dice} and BCE mask loss L_{mask} is used for mask prediction. Thus, the total loss of SparseOcc is composed of:

$$L = L_{focal} + L_{mask} + L_{dice} + L_{occ}, \tag{2}$$

where L_{occ} is the loss of sparse voxel decoder.

4 Ray-level mIoU

4.1 Revisiting the Voxel-level mIoU

The Occ3D dataset [50], along with its proposed evaluation metrics, are widely recognized as benchmarks in this field. The ground truth occupancy is reconstructed from LiDAR point clouds, and the mean Intersection over Union (mIoU) at the voxel level is employed to assess performance. Due to factors such as distance and occlusion, accumulated LiDAR point clouds are not perfect. Some areas unscanned by LiDAR are marked as free, resulting in fragmented instances. This raises the problem of label inconsistency. To solve this problem, Occ3D uses a binary *visible mask* that indicates whether a voxel is observed in the current camera view. Only the observed voxels contribute to evaluation.

However, we found that solely calculating mIoU on the observed voxel positions remains vulnerable and can be hacked by *predicting a thicker surface*. Dense methods (*e.g.*, BEVFormer [27]) can easily achieve this by training with the visible mask. During training, the area behind the surface lacks supervision, causing the model to fill it with duplicated predictions, resulting in a thicker surface. As an example, consider BEVFormer, which generates a thick and noisy surface when trained with the visible mask (see Fig. 4). Despite this, its performance exhibits an unreasonably inflated improvement ($+5\sim15$ mIoU) under the current evaluation protocol.



Fig. 4: Visualization of the discrepancy between qualitative and quantitative results. We observe that training existing dense occupancy methods (*e.g.* BEVFormer) with a visible mask results in a thick surface, leading to an unreasonably inflated improvement in the current mIoU metrics. In contrast, our new RayIoU metrics provide a more accurate reflection of model performance.

The misalignment between qualitative and quantitative results is caused by the inconsistent penalty along the depth direction. A toy example in Fig. 5 reveals several issues with the current evaluation metrics:

- 1. If the model fills all areas behind the surface, it inconsistently penalizes depth predictions. The model can obtain a higher IoU by filling all areas behind the surface and predicting a closer depth. This thick surface issue is very common in models trained with visible masks or 2D supervision.
- 2. If the predicted occupancy represents a thin surface, the penalty becomes overly strict. Even a deviation of just one voxel results in an IoU of zero.
- 3. The visible mask only considers the visible area at the current moment, thereby reducing 3D occupancy to a depth estimation task with categories and overlooking the scene completion ability.

4.2 Mean IoU by Ray Casting

To address the above issues, we propose a new metric: Ray-level mIoU (RayIoU for short). In RayIoU, the set elements are query rays rather than voxels. We emulate LiDAR behavior by projecting query rays into the predicted 3D occupancy volume. For each query ray, we compute the distance it travels before intersecting any surface and retrieve the corresponding class label. We then apply the same procedure to the ground-truth occupancy to obtain the ground-truth depth and class label. In case a ray does not intersect with any voxel present in the ground truth, it will be excluded from the evaluation process.



Fig. 5: Illustration of inconsistent depth penalties caused by current metrics. Consider a scenario where we have a wall in front of us, with a ground-truth distance of d and a thickness of d_v . When the prediction has a thickness of $d_p \gg d_v$, we encounter an inconsistent penalty along depth. Specifically, if the predicted wall is d_v farther than the ground truth (total distance $d+d_v$), its IoU will be zero. Conversely, if the predicted wall is d_v closer than the ground truth (total distance $d - d_v$), the IoU remains at 0.5. This occurs because all voxels behind the surface are filled with duplicated predictions. Similarly, when the predicted depth is $d - 2d_v$, the resulting IoU is $\frac{1}{3}$, and so forth.



Fig. 6: Covered area of RayIoU. (a) The raw LiDAR ray samples are unbalanced at different distances. (b) We resample the rays to balance the weight on distance. (c) To investigate the performance of scene completion, we propose evaluating occupancy in the visible area on a wide time span, by casting rays on visited waypoints.

As shown in Fig. 6 (a), the original LiDAR rays in a real dataset tend to be unbalanced from near to far. Thus, we resample the LiDAR rays to achieve a balanced distribution across different distances (Fig. 6 (b)). In the near field, we modify the LiDAR ray channels to achieve equal-distant spacing when projected onto the ground plane. In the far field, we increase the angular resolution of the ray channels to ensure a more uniform data density across varying ranges. Moreover, our query ray can originate from the LiDAR position at the current, past, or future moments of the ego path. Temporal casting (Fig. 6 (c)) allows us to evaluate scene completion performance while maintaining a well-posed task.

A query ray is classified as a *true positive* (TP) if the class labels coincide and the L1 error between the ground-truth depth and the predicted depth is less than a certain threshold (e.g., 2m). Let C be the number of classes,

$$RayIoU = \frac{1}{C} \sum_{c=1}^{C} \frac{TP_c}{TP_c + FP_c + FN_c},$$
(3)

10 Liu et al.

Table 1: 3D occupancy prediction performance on Occ3D-nuScenes [50]. We use RayIoU to compare our SparseOcc with other methods. "8f" and "16f" mean fusing temporal information from 8 or 16 frames. SparseOcc outperforms all existing methods under a weaker setting.

Method	Backbone	Input Size	Epoch	RayIoU	RayI	oU _{1m,}	2m, 4m	mIoU	FPS
BEVFormer (4f) [27]	R101	1600×900	24	32.4	26.1	32.9	38.0	39.2	3.0
RenderOcc [40]	Swin-B	$1408{\times}512$	12	19.5	13.4	19.6	25.5	24.4	-
SimpleOcc [12]	R101	672×336	12	22.5	17.0	22.7	27.9	31.8	9.7
BEVDet-Occ $(2f)$ [15]	R50	704×256	90	29.6	23.6	30.0	35.1	36.1	2.6
BEVDet-Occ-Long (8f)	R50	704×384	90	32.6	26.6	33.1	38.2	39.3	0.8
FB-Occ (16f) [28]	R50	$704{ imes}256$	90	33.5	26.7	34.1	39.7	39.1	10.3
SparseOcc (8f)	R50	704×256	24	34.0	28.0	34.7	39.4	30.1	17.3
SparseOcc (16f)	R50	704×256	24	35.1	29.1	35.8	40.3	30.6	12.5
SparseOcc (16f)	R50	$704{ imes}256$	48	36.1	30.2	36.8	41.2	30.9	12.5

where TP_c , FP_c and FN_c correspond to the number of true positive, false positive, and false negative predictions for class c_i .

RayIoU addresses all three of the aforementioned problems:

- 1. Since the query ray calculates the distance to the first voxel it touches, the model cannot obtain a higher IoU by predicting a thicker surface.
- 2. RayIoU determines true positives based on a distance threshold, which mitigates the overly strict nature of voxel-level mIoU.
- 3. The query ray can originate from any position in the scene. This flexibility allows RayIoU to consider the model's scene completion ability, preventing the reduction of occupancy estimation to mere depth prediction.

5 Experiments

We evaluate our model on the Occ3D-nus [50] dataset. Occ3D-nus is based on the nuScenes [3] dataset, which consists of large-scale multimodal data collected from 6 surround-view cameras, 1 lidar and 5 radars. The dataset has 1000 videos and is split into 700/150/150 videos for training/validation/testing. Each video has roughly 20s duration and the key samples are annotated every 0.5s.

We use the proposed RayIoU to evaluate the semantic segmentation performance. The query rays originate from 8 LiDAR positions of the ego path. We calculate RayIoU under three distance thresholds: 1, 2 and 4 meters. The final ranking metric is averaged over these distance thresholds.

5.1 Implementation Details

We implement our model using PyTorch [41]. Following previous methods, we adopt ResNet-50 [14] as the image backbone. The mask transformer consists of 3 layers with shared weights across different layers. In our main experiments, we employ semantic queries where each query corresponds to a semantic class,

Voxel Decoder	RayIoU	${\rm RayIoU_{1m}}$	${\rm RayIoU_{2m}}$	${\rm RayIoU_{4m}}$	FPS
Dense coarse-to-fine Dense patch-based	29.9 25.8	24.0 20.4	$30.4 \\ 26.0$	35.4 30.9	6.3 7.8
Sparse coarse-to-fine	29.9	23.9	30.5	35.2	24.0

Table 2: Sparse voxel decoder vs. dense voxel decoder. Our sparse voxel decoder achieves nearly $4 \times$ faster inference speed than the dense counterparts.

rather than an instance. The ray casting module in RayIoU is implemented based on the codebase of [21].

During training, we use the AdamW [36] optimizer with a global batch size of 8. The initial learning rate is set to 2×10^{-4} and is decayed with cosine annealing policy. For all experiments, we train our models for 24 epochs. FPS is measured on a Tesla A100 GPU with the PyTorch **fp32** backend (single batch size).

5.2 Main Results

In Tab. 1 and Fig. 1 (b), we compare SparseOcc with previous state-of-the-art methods on the validation split of Occ3D-nus. Despite under a weaker setting (ResNet-50 [14], 8 history frames, and input image resolution of 704×256), SparseOcc significantly outperforms previous methods including FB-Occ, the winner of CVPR 2023 occupancy challenge, with many complicated designs including forward-backward view transformation, depth net, joint depth and semantic pre-training, and so on. SparseOcc achieves better results (+1.6 RayIoU) while being much faster and simpler than FB-Occ, which demonstrates the superiority of our solution.

We further provide qualitative results in Fig. 7. Both BEVDet-Occ and FB-Occ are dense methods and make many redundant predictions behind the surface. In contrast, SparseOcc discards over 90% of voxels while still effectively modeling the geometry of the scene and capturing fine-grained details.

5.3 Ablations

In this section, we conduct ablations on the validation split of Occ3D-nuScenes to confirm the effectiveness of each module. By default, we use the single frame version of SparseOcc as the baseline. The choice for our model is made **bold**.

Sparse voxel decoder vs. dense voxel decoder. In Tab. 2, we compare our sparse voxel decoder to the dense counterparts. Here, we implement two baselines, and both of them output a dense feature map with shape as $200 \times 200 \times 16 \times C$. The first baseline is a coarse-to-fine architecture without pruning empty voxels. In this baseline, we also replace self-attention with 3D convolution and use 3D deconvolution to upsample predictions. The other baseline is a patch-based architecture by dividing the 3D space into a small number of patches as PETRv2 [35] for BEV segmentation. We use $25 \times 25 \times 2 = 1250$ queries and each one of them



Fig. 7: Visualized comparison of semantic occupancy prediction. Despite discarding over 90% of voxels, our SparseOcc effectively models the geometry of the scene and captures fine-grained details (*e.g.*, the yellow-marked traffic cone in the bottom row).

corresponds to a specific patch of shape $8 \times 8 \times 8$. A stack of deconvolution layers are used to lift the coarse queries to a full-resolution 3D volume.

As we can see from the table, the dense coarse-to-fine baseline achieves a good performance of 29.9 RayIoU but with a slow inference speed of 6.3 FPS. The patch-based one is slightly faster with 7.8 FPS inference speed but with a severe performance drop by 4.1 RayIoU. Instead, our sparse voxel decoder produces sparse 3D features in the shape of $K \times C$ (where $K = 32000 \ll 200 \times 200 \times 16$), achieving an inference speed that is nearly $4 \times$ faster than the counterparts without compromising performance. This demonstrates the necessity and effective-ness of our sparse design.

Mask Transformer. In Tab. 3, we first ablate the effectiveness of mask transformer. The first row is a per-voxel baseline which directly predicts semantics from the sparse voxel decoder using a stack of MLPs. Introducing mask transformer with vanilla dense cross attention (as it is the common practice in Mask-Former and Mask3D) gives a performance boost of 1.7 RayIoU, but inevitably slows down the inference speed. Therefore, to speed up the dense cross-attention pipeline, we adopt a sparse sampling mechanism which brings a 50% reduction in inference time. Then, by further sampling sparse 3D points via predicted mask guidance, we finally achieve 29.2 RayIoU with 24 FPS.

MT Cross Attention	RayIoU	${\rm RayIoU_{1m}}$	${\rm RayIoU_{2m}}$	$\mathrm{RayIoU_{4m}}\big \mathrm{FPS}$
- $ -$ Dense cross attention	27.0 28.7	$20.3 \\ 22.9$	$27.5 \\ 29.3$	33.1 29.0 33.8 16.2
$\begin{array}{c c} & \text{Sparse sampling} \\ & + \text{Mask-guided} \end{array}$	25.8 29.2	20.5 23.4	26.2 29.8	30.8 24.0 34.5 24.0

Table 3: Ablation of mask transformer (MT) and the cross attention module in MT. Mask-guided sparse sampling is stronger and faster than the dense cross attention.



Fig. 8: Ablations on voxel sparsity and temporal modeling. (a) The optimal performance occurs when k is set to 32000 (5% sparsity). (b) The performance continues to increase with the number of frames, but it starts to saturate after 12 frames.

Is a limited set of voxels sufficient to cover the scene? In this study, we delve deeper into the impact of voxel sparsity on final performance. To investigate this, we systematically ablate the value of k in Fig. 8 (a). Starting from a modest value of 16000, we observe that the optimal performance occurs when k is set to 32000, which is only 5% of the total number of dense voxels $(200 \times 200 \times 16 = 640000)$. Surprisingly, further increasing k does not yield any performance improvements; instead, it introduces noise. Thus, our findings suggest that a 5% sparsity level is sufficient, and additional sparsity would be counterproductive.

Temporal modeling. In Fig. 8 (b), we validate the effectiveness of temporal fusion. We can see that the temporal modeling of SparseOcc is very effective, with performance steadily increasing as the number of frames increases. The performance peaks at 12 frames and then saturates.

5.4 More Studies

Panoptic occupancy. We then show that SparseOcc can be easily extended for panoptic occupancy predictions, a task derived from panoptic segmentation that segments images to not only semantically meaningful regions but also to detect and distinguish individual instances. Compared to panoptic segmentation, panoptic occupancy prediction requires the model to have geometric awareness in order to construct the 3D scene for segmentation. By additionally introducing instance queries to the mask transformer, we seamlessly achieve the first panoptic occupancy prediction framework using camera inputs. In Fig. 9, we visualize



Fig. 9: Panoptic occupancy prediction. Different instances are distinguished by colors. Our model can capture fine-grained objects and road structures simultaneously.

the panoptic occupancy results of SparseOcc. As illustrated, semantic regions and individual objects as well as their occupancy grids are well predicted. For more details, we refer the readers to the appendix.

Enhancing sparsity by removing the road surface. The majority of non-free occupancy data pertains to background geometry. In practice, the drivable surface occupancy can be effectively substituted with High-Definition Map (HD Map) or online mapping techniques [5, 25, 29, 53]. This replacement not only streamlines the sparsity but also enriches the semantic and structural understanding of roads. We construct experiments to investigate the effect of removing road surface in SparseOcc. The details can be found in the appendix.

6 Limitations

Accumulative errors. In order to implement a fully sparse architecture, we discard a large number of empty voxels in the early stages. However, empty voxels that are mistakenly discarded cannot be recovered in subsequent stages. Moreover, the prediction of the mask transformer is constrained within a space predicted by the sparse voxel decoder. Some ground-truth instances do not appear in this unreliable space, leading to inadequate training of the mask transformer.

7 Conclusion

In this paper, we proposed a fully sparse occupancy network, named SparseOcc, which neither relies on dense 3D feature, nor has sparse-to-dense and global attention operations. We also created RayIoU, a ray-level metric for occupancy evaluation, eliminating the inconsistency flaws of previous metric. Experiments show that SparseOcc achieves the state-of-the-art performance on the Occ3D-nuScenes dataset for both speed and accuracy. We hope this exciting result will attract more attention to the fully sparse 3D occupancy paradigm.

Acknowledgements

We thank the anonymous reviewers for their suggestions that make this work better. This work is supported by the National Key R&D Program of China (No. 2022ZD0160900), the National Natural Science Foundation of China (No. 62076119, No. 61921006), the Fundamental Research Funds for the Central Universities (No. 020214380119), and the Collaborative Innovation Center of Novel Software Technology and Industrialization.

References

- 1. Tesla AI Day. https://www.youtube.com/watch?v=j0z4FweCy4M (2021)
- 2. Bozic, A., Palafox, P., Thies, J., Dai, A., Nießner, M.: Transformerfusion: Monocular rgb scene reconstruction using transformers. In: NeurIPS (2021)
- Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: CVPR (2020)
- Cao, A.Q., de Charette, R.: Monoscene: Monocular 3d semantic scene completion. In: CVPR (2022)
- Chen, L., Sima, C., Li, Y., Zheng, Z., Xu, J., Geng, X., Li, H., He, C., Shi, J., Qiao, Y., Yan, J.: Persformer: 3d lane detection via perspective transformer and the openlane benchmark. In: ECCV (2022)
- Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. In: CVPR (2022)
- Cheng, B., Schwing, A., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. In: NeurIPS (2021)
- Choy, C., Gwak, J., Savarese, S.: 4d spatio-temporal convnets: Minkowski convolutional neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3075–3084 (2019)
- 9. Curless, B., Levoy, M.: A volumetric method for building complex models from range images. In: SIGGRAPH (1996)
- Ding, Y., Huang, L., Zhong, J.: Multi-scale occ: 4th place solution for Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition 2023 3d occupancy prediction challenge. arXiv preprint arXiv:2306.11414 (2023)
- 11. Feng, Z., Yang, L., Guo, P., Li, B.: Cvrecon: Rethinking 3d geometric feature learning for neural reconstruction. In: ICCV (2023)
- Gan, W., Mo, N., Xu, H., Yokoya, N.: A comprehensive framework for 3d occupancy estimation in autonomous driving. IEEE Transactions on Intelligent Vehicles pp. 1–19 (2024)
- 13. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: CVPR (2022)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
- Huang, J., Huang, G.: Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. arXiv preprint arXiv:2203.17054 (2022)
- Huang, J., Huang, G., Zhu, Z., Du, D.: Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. arXiv preprint arXiv:2112.11790 (2021)
- 17. Huang, L., Li, Z., Sima, C., Wang, W., Wang, J., Qiao, Y., Li, H.: Leveraging vision-centric multi-modal expertise for 3d object detection. In: NeurIPS (2024)

- 16 Liu et al.
- Huang, L., Wang, H., Zeng, J., Zhang, S., Cao, L., Ji, R., Yan, J., Li, H.: Geometric-aware pretraining for vision-centric 3d object detection. arXiv preprint arXiv:2304.03105 (2023)
- Huang, Y., Zheng, W., Zhang, Y., Zhou, J., Lu, J.: Tri-perspective view for visionbased 3d semantic occupancy prediction. In: CVPR (2023)
- Jia, X., De Brabandere, B., Tuytelaars, T., Gool, L.V.: Dynamic filter networks. In: NeurIPS (2016)
- Khurana, T., Hu, P., Held, D., Ramanan, D.: Point cloud forecasting as a proxy for 4d occupancy forecasting. In: CVPR (2023)
- 22. Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. In: CVPR (2019)
- Li, H., Li, Y., Wang, H., Zeng, J., Cai, P., Xu, H., Lin, D., Yan, J., Xu, F., Xiong, L., Wang, J., Zhu, F., Yan, K., Xu, C., Wang, T., Mu, B., Ren, S., Peng, Z., Qiao, Y.: Open-sourced data ecosystem in autonomous driving: the present and future. arXiv preprint arXiv:2312.03408 (2023)
- 24. Li, H., Sima, C., Dai, J., Wang, W., Lu, L., Wang, H., Zeng, J., Li, Z., Yang, J., Deng, H., et al.: Delving into the devils of bird's-eye-view perception: A review, evaluation and recipe. IEEE TPAMI (2023)
- 25. Li, T., Jia, P., Wang, B., Chen, L., Jiang, K., Yan, J., Li, H.: Lanesegnet: Map learning with lane segment perception for autonomous driving. In: ICLR (2024)
- Li, Y., Yu, Z., Choy, C., Xiao, C., Alvarez, J.M., Fidler, S., Feng, C., Anandkumar, A.: Voxformer: Sparse voxel transformer for camera-based 3d semantic scene completion. In: CVPR (2023)
- Li, Z., Wang, W., Li, H., Xie, E., Sima, C., Lu, T., Qiao, Y., Dai, J.: Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers. In: ECCV (2022)
- Li, Z., Yu, Z., Austin, D., Fang, M., Lan, S., Kautz, J., Alvarez, J.M.: Fb-occ: 3d occupancy prediction based on forward-backward view transformation. arXiv preprint arXiv:2307.01492 (2023)
- Liao, B., Chen, S., Wang, X., Cheng, T., Zhang, Q., Liu, W., Huang, C.: Maptr: Structured modeling and learning for online vectorized hd map construction. In: ICLR (2023)
- 30. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR (2017)
- 31. Liu, H., Lu, T., Xu, Y., Liu, J., Li, W., Chen, L.: Camliflow: Bidirectional cameralidar fusion for joint optical flow and scene flow estimation. In: CVPR (2022)
- 32. Liu, H., Lu, T., Xu, Y., Liu, J., Wang, L.: Learning optical flow and scene flow with bidirectional camera-lidar fusion. arXiv preprint arXiv:2303.12017 (2023)
- Liu, H., Teng, Y., Lu, T., Wang, H., Wang, L.: Sparsebev: High-performance sparse 3d object detection from multi-camera videos. In: ICCV (2023)
- 34. Liu, Y., Wang, T., Zhang, X., Sun, J.: PETR: position embedding transformation for multi-view 3d object detection. In: ECCV (2022)
- Liu, Y., Yan, J., Jia, F., Li, S., Gao, Q., Wang, T., Zhang, X., Sun, J.: Petrv2: A unified framework for 3d perception from multi-camera images. arXiv preprint arXiv:2206.01256 (2022)
- 36. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: ICLR (2019)
- Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: CVPR (2019)
- Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 3DV (2016)

- Murez, Z., Van As, T., Bartolozzi, J., Sinha, A., Badrinarayanan, V., Rabinovich, A.: Atlas: End-to-end 3d scene reconstruction from posed images. In: ECCV (2020)
- Pan, M., Liu, J., Zhang, R., Huang, P., Li, X., Liu, L., Zhang, S.: Renderocc: Visioncentric 3d occupancy prediction with 2d rendering supervision. arXiv preprint arXiv:2309.09502 (2023)
- 41. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, highperformance deep learning library. In: NeurIPS (2019)
- Peng, S., Niemeyer, M., Mescheder, L., Pollefeys, M., Geiger, A.: Convolutional occupancy networks. In: ECCV (2020)
- 43. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI (2015)
- 44. Schult, J., Engelmann, F., Hermans, A., Litany, O., Tang, S., Leibe, B.: Mask3d: Mask transformer for 3d semantic instance segmentation. In: ICRA (2023)
- 45. Sima, C., Tong, W., Wang, T., Chen, L., Wu, S., Deng, H., Gu, Y., Lu, L., Luo, P., Lin, D., Li, H.: Scene as occupancy. In: ICCV (2023)
- 46. Stier, N., Rich, A., Sen, P., Höllerer, T.: Vortx: Volumetric 3d reconstruction with transformers for voxelwise view selection and fusion. In: 3DV (2021)
- 47. Sun, J., Xie, Y., Chen, L., Zhou, X., Bao, H.: Neuralrecon: Real-time coherent 3d reconstruction from monocular video. In: CVPR (2021)
- Takikawa, T., Litalien, J., Yin, K., Kreis, K., Loop, C., Nowrouzezahrai, D., Jacobson, A., McGuire, M., Fidler, S.: Neural geometric level of detail: Real-time rendering with implicit 3d shapes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11358–11367 (2021)
- Takmaz, A., Fedele, E., Sumner, R.W., Pollefeys, M., Tombari, F., Engelmann, F.: Openmask3d: Open-vocabulary 3d instance segmentation. arXiv preprint arXiv:2306.13631 (2023)
- Tian, X., Jiang, T., Yun, L., Wang, Y., Wang, Y., Zhao, H.: Occ3d: A large-scale 3d occupancy prediction benchmark for autonomous driving. In: NeurIPS Datasets and Benchmarks (2023)
- Tolstikhin, I.O., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., et al.: Mlp-mixer: An all-mlp architecture for vision. In: NeurIPS (2021)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017)
- 53. Wang, H., Li, T., Li, Y., Chen, L., Sima, C., Liu, Z., Wang, B., Jia, P., Wang, Y., Jiang, S., Wen, F., Xu, H., Luo, P., Yan, J., Zhang, W., Li, H.: Openlane-v2: A topology reasoning benchmark for unified 3d hd mapping. In: NeurIPS (2023)
- Wang, S., Liu, Y., Wang, T., Li, Y., Zhang, X.: Exploring object-centric temporal modeling for efficient multi-view 3d object detection. arXiv preprint arXiv:2303.11926 (2023)
- 55. Wang, X., Zhu, Z., Xu, W., Zhang, Y., Wei, Y., Chi, X., Ye, Y., Du, D., Lu, J., Wang, X.: Openoccupancy: A large scale benchmark for surrounding semantic occupancy perception. arXiv preprint arXiv:2303.03991 (2023)
- Wang, Y., Guizilini, V.C., Zhang, T., Wang, Y., Zhao, H., Solomon, J.: Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In: CoRL (2022)
- 57. Wang, Y., Chen, Y., Liao, X., Fan, L., Zhang, Z.: Panoocc: Unified occupancy representation for camera-based 3d panoptic segmentation. arXiv preprint arXiv:2306.10013 (2023)
- Wei, Y., Zhao, L., Zheng, W., Zhu, Z., Zhou, J., Lu, J.: Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving. In: ICCV (2023)

- 18 Liu et al.
- 59. Yang, Z., Chen, L., Sun, Y., Li, H.: Visual point cloud forecasting enables scalable autonomous driving. In: CVPR (2024)
- 60. Yang, Z., Jiang, L., Sun, Y., Schiele, B., Jia, J.: A unified query-based paradigm for point cloud understanding. In: ICCV (2022)
- Yang, Z., Sun, Y., Liu, S., Jia, J.: 3dssd: Point-based 3d single stage object detector. In: CVPR (2020)
- 62. Yang, Z., Sun, Y., Liu, S., Shen, X., Jia, J.: Std: Sparse-to-dense 3d object detector for point cloud. In: ICCV (2019)
- 63. Yang, Z., Zhou, Y., Chen, Z., Ngiam, J.: 3d-man: 3d multi-frame attention network for object detection. In: ICCV (2021)
- Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3d object detection and tracking. In: CVPR (2021)
- Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)