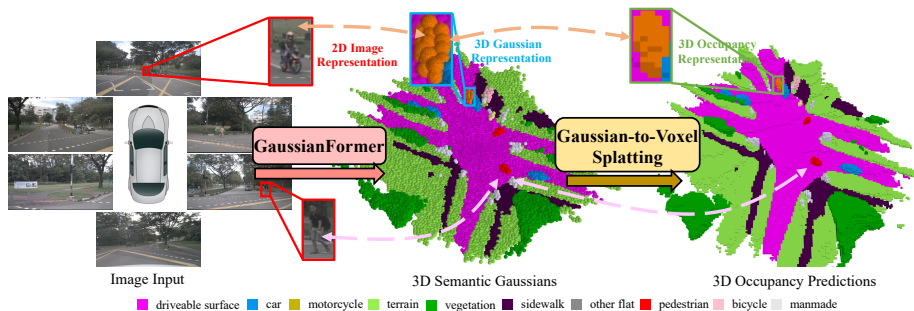


# GaussianFormer: Scene as Gaussians for Vision-Based 3D Semantic Occupancy Prediction

Yuanhui Huang<sup>1</sup>, Wenzhao Zheng<sup>1,2\*</sup>, Yunpeng Zhang<sup>3</sup>,  
Jie Zhou<sup>1</sup>, and Jiwen Lu<sup>1†</sup>

<sup>1</sup>Tsinghua University <sup>2</sup>University of California, Berkeley <sup>3</sup>PhiGent Robotics  
<https://wzzheng.net/GaussianFormer>  
huangyh22@mails.tsinghua.edu.cn; wenzhao.zheng@outlook.com;  
yunpengzhang97@gmail.com; {jzhou,lujiwen}@tsinghua.edu.cn



**Fig. 1:** Considering the universal approximating ability of Gaussian mixture [9, 12], we propose an object-centric 3D semantic Gaussian representation to describe the fine-grained structure of 3D scenes without the use of dense grids. We propose a GaussianFormer model consisting of sparse convolution and cross-attention to efficiently transform 2D images into 3D Gaussian representations. To generate dense 3D occupancy, we design a Gaussian-to-voxel splatting module that can be efficiently implemented with CUDA. With comparable performance, our GaussianFormer reduces memory consumption of existing 3D occupancy prediction methods by **75.2%** - **82.2%**.

**Abstract.** 3D semantic occupancy prediction aims to obtain 3D fine-grained geometry and semantics of the surrounding scene and is an important task for the robustness of vision-centric autonomous driving. Most existing methods employ dense grids such as voxels as scene representations, which ignore the sparsity of occupancy and the diversity of object scales and thus lead to unbalanced allocation of resources. To address this, we propose an object-centric representation to describe 3D scenes with sparse 3D semantic Gaussians where each Gaussian represents a flexible region of interest and its semantic features. We aggregate information from images through the attention mechanism and iteratively refine the properties of 3D Gaussians including position, covariance, and semantics. We then propose an efficient Gaussian-to-voxel splatting

\* Project Leader

† Corresponding Author

method to generate 3D occupancy predictions, which only aggregates the neighboring Gaussians for a certain position. We conduct extensive experiments on the widely adopted nuScenes and KITTI-360 datasets. Experimental results demonstrate that GaussianFormer achieves comparable performance with state-of-the-art methods with only 17.8% - 24.8% of their memory consumption. Code is available at: <https://github.com/huang-yh/GaussianFormer>.

**Keywords:** 3D Occupancy Prediction · 3D Gaussian splatting · Autonomous Driving

## 1 Introduction

Whether to use LiDAR for 3D perception has long been the core debate among autonomous driving companies. While vision-centric systems share an economical advantage, their inability to capture obstacles of arbitrary shapes hinders driving safety and robustness [14, 18, 26, 27]. The emergence of 3D semantic occupancy prediction methods [4, 17, 19, 36, 51, 58, 64] remedies this issue by predicting the occupancy status of each voxel in the 3D space, which facilitates a variety of newly rising tasks such as end-to-end autonomous driving [47, 60], 4D occupancy forecasting [59], and self-supervised 3D scene understanding [16].

Despite the promising applications, the dense output space of 3D occupancy prediction poses a great challenge in how to efficiently and effectively represent the 3D scene. Voxel-based methods [24, 51] assign each voxel with a feature vector to obtain dense representations to describe the fine-grained structure of a 3D scene. They employ coarse-to-fine upsampling [47, 51] or voxel filtering [24, 34] to improve efficiency considering the sparse nature of the 3D space. As most of the voxel space is unoccupied [3], BEV-based methods [28, 57] compress the height dimension and employ the bird’s eye view (BEV) as scene representations, yet they usually require post-processing such as multi-scale fusion [57] to capture finer details. TPVFormer [17] generalizes BEV with two additional planes and achieves a better performance-complexity trade-off with the tri-perspective view (TPV). However, they are all grid-based methods and inevitably suffer from the redundancy of empty grids, resulting in more complexity for downstream tasks [47]. It is also more difficult to capture scene dynamics with grid-based representations since it is objects instead of grids that move in the 3D space [59].

In this paper, we propose the first object-centric representation for 3D semantic occupancy prediction. We employ a set of 3D semantic Gaussians to sparsely describe a 3D scene. Each Gaussian represents a flexible region of interest and consists of the mean, covariance, and its semantic category. We propose a GaussianFormer model to effectively obtain 3D semantic Gaussians from image inputs. We randomly initialize a set of queries to instantiate the 3D Gaussians and adopt the cross-attention mechanism to aggregate information from multi-scale image features. We iteratively refine the properties of the 3D Gaussians for smoother optimizations. To efficiently incorporate interactions among 3D Gaussians, we treat them as point clouds located at the Gaussian means and

leverage 3D sparse convolutions to process them. We then decode the properties of 3D semantic Gaussians from the updated queries as the scene representation. Motivated by the 3D Gaussian splatting method in image rendering [21], we design an efficient Gaussian-to-voxel splatting module that aggregates neighboring Gaussians to generate the semantic occupancy for a certain 3D position. The proposed 3D Gaussian representation uses a sparse and adaptive set of features to describe a 3D scene but can still model the fine-grained structure due to the universal approximating ability of Gaussian mixtures [9, 12]. Based on the 3D Gaussian representation, GaussianFormer further employs sparse convolution and local-aggregation-based Gaussian-to-voxel splatting to achieve efficient 3D semantic occupancy prediction, as shown in Fig. 1. We conduct extensive experiments on the nuScenes and KITTI-360 datasets for 3D semantic occupancy prediction from surrounding and monocular cameras, respectively. GaussianFormer achieves comparable performance with existing state-of-the-art methods with only 17.8% - 24.8% of their memory consumption. Our qualitative visualizations show that GaussianFormer is able to generate a both holistic and realistic perception of the scene.

## 2 Related Work

### 2.1 3D Semantic Occupancy Prediction

3D semantic occupancy prediction has garnered increasing attention in recent years due to its comprehensive description of the driving scenes, which involves predicting the occupancy and semantic states of all voxels within a certain range. Learning an effective representation constitutes the fundamental step for this challenging task. A straightforward approach discretizes the 3D space into regular voxels, with each voxel being assigned a feature vector [61]. The capacity to represent intricate 3D structures renders voxel-based representations advantageous for 3D semantic occupancy prediction [4, 7, 19, 22, 36, 44, 51, 52, 58]. However, due to the sparsity of driving scenes and the high resolution of vanilla voxel representation, these approaches suffer from considerable computation and storage overhead. To improve efficiency, several methods propose to reduce the number of voxel queries by the coarse-to-fine upsampling strategy [50], or the depth-guided voxel filtering [24]. Nonetheless, the upsampling process might not adequately compensate for the information loss in the coarse stage. And the filtering strategy ignores the occluded area and depends on the quality of depth estimation. Although OctreeOcc [34] introduces voxel queries of multi-granularity to enhance the efficiency, it still conforms to a predefined regular partition pattern. Our 3D Gaussian representation resembles the voxel counterpart, but can flexibly adapt to varying object scales and region complexities in a deformable way, thus achieving better resource allocation and efficiency.

Another line of work utilizes the bird’s-eye-view (BEV) representation for 3D perception in autonomous driving, which can be categorized into two types according to the view transformation paradigm. Approaches based on lift-splat-shoot (LSS) actively project image features into 3D space with depth guid-

ance [15, 25, 26, 32, 40, 41], while query-based methods typically use BEV queries and deformable attention to aggregate information from image features [20, 27, 54]. Although BEV-based perception has achieved great success in 3D object detection, it is less employed for 3D occupancy prediction due to information loss from height compression. FB-OCC [28] uses dense BEV features from backward projection to optimize the sparse voxel features from forward projection. FlashOcc [57] applies a complex multi-scale feature fusion module on BEV features for finer details. However, existing 3D occupancy prediction methods are based on grid representations, which inevitably suffer from the computation redundancy of empty grids. Differently, our GaussianFormer is based on object-centric representation and can efficiently tend to flexible regions of interest.

## 2.2 3D Gaussian Splatting

The recent 3D Gaussian splatting (3D-GS) [21] uses multiple 3D Gaussians for radiance field rendering, demonstrating superior performance in rendering quality and speed [6]. In contrast to prior explicit scene representations, such as meshes [42, 43, 53] and voxels [11, 38], 3D-GS is capable of modeling intricate shapes with fewer parameters. Compared with implicit neural radiance field [1, 37], 3D-GS facilitates fast rendering through splat-based rasterization, which projects 3D Gaussians to the target 2D view and renders image patches with local 2D Gaussians. Recent advances in 3D-GS include adaptation to dynamic scenarios [35, 55], online 3D-GS generalizable to novel scenes [5, 63], and generative 3D-GS [8, 29, 46, 56].


Although our 3D Gaussian representation also adopts the physical form of 3D Gaussians (i.e. mean and covariance) and the multi-variant Gaussian distribution, it differs from 3D-GS in significant ways, which imposes unique challenges in 3D semantic occupancy prediction: 1) Our 3D Gaussian representation is learned in an online manner as opposed to offline optimization in 3D-GS. 2) We generate 3D semantic occupancy predictions from the 3D Gaussian representation in contrast to rendering 2D RGB images in 3D-GS.

## 3 Proposed Approach

### 3.1 Object-Centric 3D Scene Representation

Vision-based 3D semantic occupancy prediction aims to predict dense occupancy states and semantics for each voxel grid with multi-view camera images as input. Formally, given a set of multi-view images  $\mathcal{I} = \{\mathbf{I}_i \in \mathbb{R}^{3 \times H \times W} | i = 1, \dots, N\}$ , and corresponding intrinsics  $\mathcal{K} = \{\mathbf{K}_i \in \mathbb{R}^{3 \times 3} | i = 1, \dots, N\}$  and extrinsics  $\mathcal{T} = \{\mathbf{T}_i \in \mathbb{R}^{4 \times 4} | i = 1, \dots, N\}$ , the objective is to predict 3D semantic occupancy  $\mathbf{O} \in \mathcal{C}^{X \times Y \times Z}$ , where  $N$ ,  $\{H, W\}$ ,  $\mathcal{C}$  and  $\{X, Y, Z\}$  denote the number of views, the image resolution, the set of semantic classes and the target volume resolution.

The autonomous driving scenes contain foreground objects of various scales (such as buses and pedestrians), and background regions of different complexities



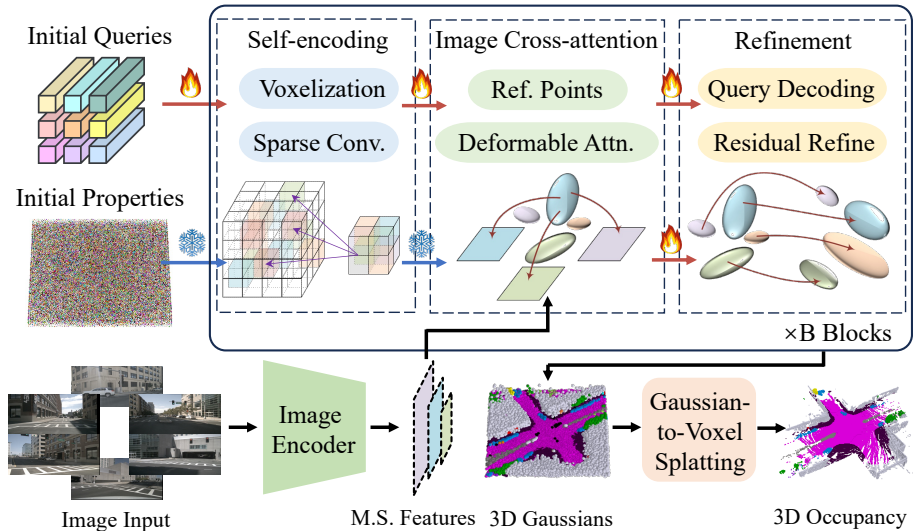
**Fig. 2: Comparisons of the proposed 3D Gaussian representation with existing grid-based scene representations (figures from TPVFormer [17]).** The voxel representation [24, 51] assigns each voxel in the 3D space with a feature and is redundant due to the sparsity nature of the 3D space. BEV [27] and TPV [17] employ 2D planes to describe 3D space but can only alleviate the redundancy issue. Differently, the proposed object-centric 3D Gaussian representation can adapt to flexible regions of interest yet can still describe the fine-grained structure of the 3D scene due to the strong approximating ability of mixing Gaussians [9, 12].

(such as road and vegetation). Dense voxel representation [4, 24, 49] neglects this diversity and processes every 3D location with equal storage and computation resources, which often leads to intractable overhead because of unreasonable resource allocation. Planar representations, such as BEV [15, 27] and TPV [17], achieve 3D perception by first encoding 3D information into 2D feature maps for efficiency and then recovering 3D structures from 2D features. Although planar representations are resource-friendly, they could cause a loss of details. The grid-based methods can hardly adapt to regions of interest for different scenes and thus lead to representation and computation redundancy.

To address this, we propose an object-centric 3D representation for 3D semantic occupancy prediction where each unit describes a region of interest instead of fixed grids, as shown in Fig. 2. We represent an autonomous driving scene with a number of 3D semantic Gaussians, and each of them instantiates a semantic Gaussian distribution characterized by mean, covariance, and semantic logits. The occupancy prediction for a 3D location can be computed by summing up the values of semantic Gaussian distributions evaluated at that location. Specifically, we use a set of  $P$  3D Gaussians  $\mathcal{G} = \{\mathbf{G}_i \in \mathbb{R}^d | i = 1, \dots, P\}$  for each scene, and each 3D Gaussian is represented by a  $d$ -dimensional vector in the form of  $(\mathbf{m} \in \mathbb{R}^3, \mathbf{s} \in \mathbb{R}^3, \mathbf{r} \in \mathbb{R}^4, \mathbf{c} \in \mathbb{R}^{|\mathcal{C}|})$ , where  $d = 10 + |\mathcal{C}|$ , and  $\mathbf{m}$ ,  $\mathbf{s}$ ,  $\mathbf{r}$ ,  $\mathbf{c}$  denote the mean, scale, rotation vectors and semantic logits, respectively. Therefore, the value of a semantic Gaussian distribution  $\mathbf{g}$  evaluated at point  $\mathbf{p} = (x, y, z)$  is

$$\mathbf{g}(\mathbf{p}; \mathbf{m}, \mathbf{s}, \mathbf{r}, \mathbf{c}) = \exp\left(-\frac{1}{2}(\mathbf{p} - \mathbf{m})^T \boldsymbol{\Sigma}^{-1}(\mathbf{p} - \mathbf{m})\right) \mathbf{c}, \quad (1)$$

$$\boldsymbol{\Sigma} = \mathbf{R} \mathbf{S} \mathbf{S}^T \mathbf{R}^T, \quad \mathbf{S} = \text{diag}(\mathbf{s}), \quad \mathbf{R} = \text{q2r}(\mathbf{r}), \quad (2)$$



**Fig. 3: Framework of our GaussianFormer for 3D semantic occupancy prediction.** We first extract multi-scale (M.S.) features from image inputs using an image backbone. We then randomly initialized a set of queries and properties (mean, covariance, and semantics) to represent 3D Gaussians and update them with interleaved self-encoding, image cross-attention, and property refinement. Having obtained the updated 3D Gaussians, we employ an efficient Gaussian-to-voxel splatting module to generate dense 3D occupancy via local aggregation of Gaussians.

where  $\Sigma$ ,  $\text{diag}(\cdot)$  and  $\text{q2r}(\cdot)$  represent the covariance matrix, the function that constructs a diagonal matrix from a vector, and the function that transforms a quaternion into a rotation matrix, respectively. Then the occupancy prediction result at point  $\mathbf{p}$  can be formulated as the summation of the contribution of individual Gaussians on the location  $\mathbf{p}$ :

$$\hat{o}(\mathbf{p}; \mathcal{G}) = \sum_{i=1}^P \mathbf{g}_i(\mathbf{p}; \mathbf{m}_i, \mathbf{s}_i, \mathbf{r}_i, \mathbf{c}_i) = \sum_{i=1}^P \exp\left(-\frac{1}{2}(\mathbf{p} - \mathbf{m}_i)^T \Sigma_i^{-1}(\mathbf{p} - \mathbf{m}_i)\right) \mathbf{c}_i. \quad (3)$$

Compared with voxel representation, the mean and covariance properties allow the 3D Gaussian representation to adaptively allocate computation and storage resources according to object scales and region complexities. Therefore, we need fewer 3D Gaussians to model a scene for better efficiency while still maintain expressiveness. Meanwhile, the 3D Gaussian representation take 3D Gaussians as its basic unit, and thus avoids potential loss of details from dimension reduction in planar representations. Moreover, every 3D Gaussian has explicit semantic meaning, making the transformation from the scene representation to occupancy predictions much easier than those in other representations which often involve decoding per-voxel semantics from high dimensional features.

### 3.2 GaussianFormer: Image to Gaussians

Based on the 3D semantic Gaussian representation of the scene, we further propose a GaussianFormer model to learn meaningful 3D Gaussians from multi-view images. The overall pipeline is shown in Fig. 3. We first initialize the properties of 3D Gaussians and their corresponding high-dimensional queries as learnable vectors. Then we iteratively refine the Gaussian properties within the  $B$  blocks of GaussianFormer. Each block consists of a self-encoding module to enable interactions among 3D Gaussians, an image cross-attention module to aggregate visual information, and a refinement module to rectify the properties of 3D Gaussians.

**Gaussian Properties and Queries.** We introduce two groups of features in GaussianFormer. The Gaussian properties  $\mathcal{G} = \{\mathbf{G}_i \in \mathbb{R}^d | i = 1, \dots, P\}$  are the physical attributes as discussed in Section 3.1, and they are de facto the learning target of the model. On the other hand, the Gaussian queries  $\mathcal{Q} = \{\mathbf{Q}_i \in \mathbb{R}^m | i = 1, \dots, P\}$  are the high-dimensional feature vectors that implicitly encode 3D information in the self-encoding and image cross-attention modules, and provide guidance for rectification in the refinement module. We initialize the Gaussian properties as learnable vectors denoted by Initial Properties in Fig. 3.

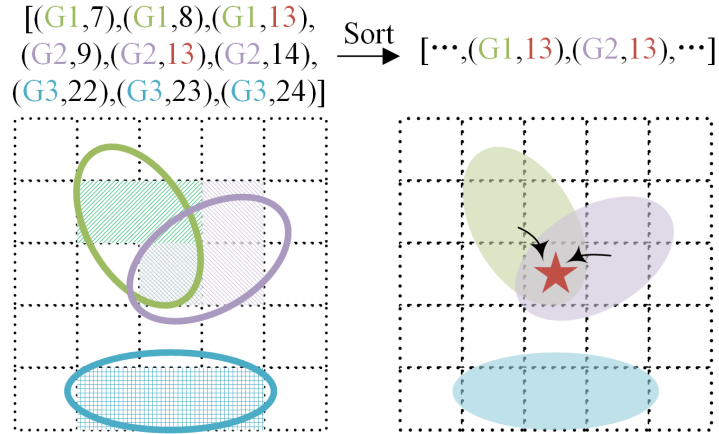
**Self-encoding Module.** Methods with voxel or planar representations usually implement self-encoding modules with deformable attention for efficiency considerations, which is not well supported for unstructured 3D Gaussian representation. Instead, we leverage 3D sparse convolution to allow interactions among 3D Gaussians, sharing the same linear computational complexity as deformable attention. Specifically, we treat each Gaussian as a point located at its mean  $\mathbf{m}$ , voxelize the generated point cloud (denoted by Voxelization in Fig. 3), and apply sparse convolution on the voxel grid. Since the number of 3D Gaussians  $P$  is much fewer than  $X \times Y \times Z$ , sparse convolution could effectively take advantage of the sparsity of Gaussians.

**Image Cross-attention Module.** The image cross-attention module (ICA) is designed to extract visual information from images for our vision-based approach. To elaborate, for a 3D Gaussian  $\mathbf{G}$ , we first generate a set of 3D reference points  $\mathcal{R} = \{\mathbf{m} + \Delta\mathbf{m}_i | i = 1, \dots, R\}$  by permuting the mean  $\mathbf{m}$  with offsets  $\Delta\mathbf{m}$ . We calculate the offsets according to the covariance of the Gaussian to reflect the shape of its distribution. Then we project the 3D reference points onto image feature maps with extrinsics  $\mathcal{T}$  and intrinsics  $\mathcal{K}$ . Finally, we update the Gaussian query  $\mathbf{Q}$  with the weighted sum of retrieved image features:

$$\text{ICA}(\mathcal{R}, \mathbf{Q}, \mathbf{F}; \mathcal{T}, \mathcal{K}) = \frac{1}{N} \sum_{n=1}^N \sum_{i=1}^R \text{DA}(\mathbf{Q}, \pi(\mathcal{R}; \mathcal{T}, \mathcal{K}), \mathbf{F}_n), \quad (4)$$

where  $\mathbf{F}$ ,  $\text{DA}(\cdot)$ ,  $\pi(\cdot)$  denote the image feature maps, the deformable attention function and the transformation from world to pixel coordinates, respectively.

**Refinement Module.** We use the refinement module to rectify the Gaussian properties with guidance from corresponding Gaussian queries which have aggregated sufficient 3D information in the prior self-encoding and image cross-attention modules. Specifically, we take inspiration from DETR [62] in object



**Fig. 4: Illustration of the Gaussian-to-voxel splatting method in 2D.** We first voxelize the 3D Gaussians and record the affected voxels of each 3D Gaussian by appending their paired indices to a list. Then we sort the list according to the voxel indices to identify the neighboring Gaussians of each voxel, followed by a local aggregation to generate the occupancy prediction.

detection. For a 3D Gaussian  $\mathbf{G} = (\mathbf{m}, \mathbf{s}, \mathbf{r}, \mathbf{c})$ , we first decode the intermediate properties  $\hat{\mathbf{G}} = (\hat{\mathbf{m}}, \hat{\mathbf{s}}, \hat{\mathbf{r}}, \hat{\mathbf{c}})$  from the Gaussian query  $\mathbf{Q}$  with a multi-layer perceptron (MLP). When refining the old properties with the intermediate ones, we treat the intermediate mean  $\hat{\mathbf{m}}$  as a residual and add it with the old mean  $\mathbf{m}$ , while we directly substitute the other intermediate properties  $(\hat{\mathbf{s}}, \hat{\mathbf{r}}, \hat{\mathbf{c}})$  for the corresponding old properties:

$$\hat{\mathbf{G}} = (\hat{\mathbf{m}}, \hat{\mathbf{s}}, \hat{\mathbf{r}}, \hat{\mathbf{c}}) = \text{MLP}(\mathbf{Q}), \quad \mathbf{G}_{new} = (\mathbf{m} + \hat{\mathbf{m}}, \hat{\mathbf{s}}, \hat{\mathbf{r}}, \hat{\mathbf{c}}). \quad (5)$$

We refine the mean of Gaussian with residual connections in order to keep their coherence throughout the  $B$  blocks of GaussianFormer. The direct replacement of the other properties is due to the concern about vanishing gradient from the sigmoid and softmax activations we apply on covariance and semantic logits.

### 3.3 Gaussian-to-Voxel Splatting

Due to the universal approximating ability of Gaussian mixtures [9, 12], the 3D semantic Gaussians can efficiently represent the 3D scene and thus can be directly processed to perform downstream tasks like motion planning and control. Specifically, to achieve 3D semantic occupancy prediction, we design an efficient Gaussian-to-voxel splatting module to efficiently transform 3D Gaussian representation to 3D semantic occupancy predictions using only local aggregation operation.

Although (3) demonstrates the main idea of the transformation as a summation over the contributions of 3D Gaussians, it is infeasible to query all Gaussians for every voxel position due to the intractable computation and storage



complexity ( $O(XYZ \times P)$ ). Since the weight  $\exp(-\frac{1}{2}(\mathbf{p} - \mathbf{m})^T \boldsymbol{\Sigma}^{-1}(\mathbf{p} - \mathbf{m}))$  in (1) decays exponentially with respect to the square of the mahalanobis distance, it should be negligible when the distance is large enough. Based on the locality of the semantic Gaussian distribution, we only consider 3D Gaussians within a neighborhood of each voxel position to improve efficiency.

As illustrated by Fig. 4, we first embed the 3D Gaussians into the target voxel grid of size  $X \times Y \times Z$  according to their means  $\mathbf{m}$ . For each 3D Gaussian, we then calculate the radius of its neighborhood according to its scale property  $\mathbf{s}$ . We append both the index of the Gaussian and the index of each voxel inside the neighborhood as a tuple  $(g, v)$  to a list. Then we sort the list according to the indices of voxels to derive the indices of 3D Gaussians that each voxel should attend to:

$$\text{sort}_{vox}([(g, v_{g_1}), \dots, (g, v_{g_k})]_{g=1}^P) = [(g_{v_1}, v), \dots, (g_{v_l}, v)]_{v=1}^{XYZ}, \quad (6)$$

where  $k, l$  denote the number of neighboring voxels of a certain Gaussian, and the number of Gaussians that contribute to a certain voxel, respectively. Finally, we can approximate (3) efficiently with only neighboring Gaussians:

$$\hat{o}(\mathbf{p}; \mathcal{G}) = \sum_{i \in \mathcal{N}(\mathbf{p})} \mathbf{g}_i(\mathbf{p}; \mathbf{m}_i, \mathbf{s}_i, \mathbf{r}_i, \mathbf{c}_i), \quad (7)$$

where  $\mathcal{N}(\mathbf{p})$  represents the set of neighboring Gaussians of the point at  $\mathbf{p}$ . Considering the dynamic neighborhood sizes of 3D Gaussians, the implementation of Gaussian-to-voxel splatting is nontrivial. To fully exploit the parallel computation ability of GPU, we realize it with the CUDA programming language to achieve better acceleration.

The overall GaussianFormer model can be trained efficiently in an end-to-end manner. For training, we use the cross entropy loss  $L_{ce}$  and the lovasz-softmax [2] loss  $L_{lov}$  following TPVFormer [17]. To refine the Gaussian properties in an iterative manner, we apply supervision on the output of each refinement module. The overall loss function is  $L = \sum_{i=1}^B L_{ce}^i + L_{lov}^i$ , where  $i$  denote the  $i$ -th block.

## 4 Experiments

In this paper, we propose a 3D semantic Gaussian representation to effectively and efficiently describe the 3D scene and devise a GaussianFormer model to perform 3D occupancy prediction. We conducted experiments on the nuScenes [3] dataset and the KITTI-360 [30] dataset for 3D semantic occupancy prediction with surrounding and monocular cameras, respectively.

### 4.1 Datasets

**NuScenes** [3] consists of 1000 sequences of various driving scenes collected in Boston and Singapore, which are officially split into 700/150/150 sequences for training, validation and testing, respectively. Each sequence lasts 20 seconds

with RGB images collected by 6 surrounding cameras, and the keyframes are annotated at 2Hz. We leverage the dense semantic occupancy annotations from SurroundOcc [51] for supervision and evaluation. The annotated voxel grid spans  $[-50\text{m}, 50\text{m}]$  along the X and Y axes, and  $[-5\text{m}, 3\text{m}]$  along the Z axis with a resolution of  $200 \times 200 \times 16$ . Each voxel is labeled with one of the 18 classes (16 semantic, 1 empty and 1 unknown classes).

**KITTI-360** [30] is a large-scale dataset covering a driving distance of 73.7km corresponding to over 320k images and 100k laser scans. We use the dense semantic occupancy annotations from SSCBench-KITTI-360 [23] for supervision and evaluation. It provides ground truth labels for 9 long sequences with a total of 12865 key frames, which are officially split into 7/1/1 sequences with 8487/1812/2566 key frames for training, validation and testing, respectively. The voxel grid spans  $51.2 \times 51.2 \times 6.4\text{m}^3$  in front of the ego car with a resolution of  $256 \times 256 \times 32$ , and each voxel is labeled with one of 19 classes (18 semantic and 1 empty). We use the RGB images from the left camera as input to our model.

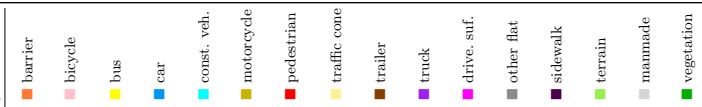
## 4.2 Implementation Details

We set the resolutions of input images as  $900 \times 1600$  for nuScenes [3] and  $376 \times 1408$  for KITTI-360 [30]. We employ ResNet101-DCN [13] initialized from FCOS3D [48] checkpoint as the image backbone for nuScenes and ResNet50 [13] pretrained with ImageNet [10] for KITTI-360. We use the feature pyramid network [31] (FPN) to generate multi-scale image features with downsample rates of 4, 8, 16 and 32. We set the number of Gaussians to 144000 and 38400 for nuScenes and KITTI-360, respectively, and use 4 transformer blocks in GaussianFormer to refine the properties of Gaussians. For optimization, we utilize the AdamW [33] optimizer with a weight decay of 0.01. The learning rate warms up in the first 500 iterations to a maximum value of  $2\text{e-}4$  and decreases according to a cosine schedule. We train our models for 20 epochs with a batch size of 8, and employ random flip and photometric distortion augmentations.


## 4.3 Results and Analysis

**Surrounding-Camera 3D semantic occupancy prediction.** In Table 1, we present a comprehensive quantitative comparison of various methods for multi-view 3D semantic occupancy prediction on nuScenes validation set, with dense annotations from SurroundOcc [51]. Our GaussianFormer achieves notable improvements over methods based on planar representations, such as BEVFormer [27] and TPVFormer [17]. Even compared with dense grid representations, GaussianFormer performs on par with OccFormer [58] and SurroundOcc [51]. These observations prove the valuable application of the 3D Gaussians for semantic occupancy prediction. This is because the 3D Gaussian representation better exploits the sparse nature of the driving scenes and the diversity of object scales with flexible properties of position and covariance.

**Table 1: 3D semantic occupancy prediction results on nuScenes validation set.** While the original TPVFormer [17] is trained with LiDAR segmentation labels, TPVFormer\* is supervised by dense occupancy annotations. Our method achieves comparable performance with state-of-the-art methods.

Method																			
	SC IoU	SSC mIoU	barrier	bicycle	bus	car	const. veh.	motorcycle	pedestrian	traffic cone	trailer	truck	drive. suf.	other flat	sidewalk	terrain	manmade	vegetation	
MonoScene [4]	23.96	7.31	4.03	0.35	8.00	8.04	2.90	0.28	1.16	0.67	4.01	4.35	27.72	5.20	15.13	11.29	9.03	14.86	
Atlas [39]	28.66	15.00	10.64	5.68	19.66	24.94	8.90	8.84	6.47	3.28	10.42	16.21	34.86	15.46	21.89	20.95	11.21	20.54	
BEVFormer [27]	30.50	16.75	14.22	6.58	23.46	28.28	8.66	10.77	6.64	4.05	11.20	17.78	37.28	18.00	22.88	22.17	<b>13.80</b>	<b>22.21</b>	
TPVFormer [17]	11.51	11.66	16.14	7.17	22.63	17.13	8.83	11.39	10.46	8.23	9.43	17.02	8.07	13.64	13.85	10.34	4.90	7.37	
TPVFormer* [17]	<b>30.86</b>	17.10	15.96	5.31	23.86	27.32	9.79	8.74	7.09	5.20	10.97	19.22	<b>38.87</b>	<b>21.25</b>	<b>24.26</b>	<b>23.15</b>	11.73	20.81	
OccFormer [58]	<b>31.39</b>	<b>19.03</b>	<b>18.65</b>	<b>10.41</b>	<b>23.92</b>	<b>30.29</b>	<b>10.31</b>	<b>14.19</b>	<b>13.59</b>	<b>10.13</b>	<b>12.49</b>	<b>20.77</b>	<b>38.78</b>	19.79	24.19	22.21	<b>13.48</b>	<b>21.35</b>	
SurroundOcc [51]	<b>31.49</b>	<b>20.30</b>	<b>20.59</b>	<b>11.68</b>	<b>28.06</b>	<b>30.86</b>	<b>10.70</b>	<b>15.14</b>	<b>14.09</b>	<b>12.06</b>	<b>14.38</b>	<b>22.26</b>	37.29	<b>23.70</b>	<b>24.49</b>	<b>22.77</b>	<b>14.89</b>	<b>21.86</b>	
<b>Ours</b>	29.83	<b>19.10</b>	<b>19.52</b>	<b>11.26</b>	<b>26.11</b>	<b>29.78</b>	<b>10.47</b>	<b>13.83</b>	<b>12.58</b>	<b>8.67</b>	<b>12.74</b>	<b>21.57</b>	<b>39.63</b>	<b>23.28</b>	<b>24.46</b>	<b>22.99</b>	9.59	19.12	

**Table 2: 3D semantic occupancy prediction results on SSCBench-KITTI-360 validation set.** Our method achieves performance on par with state-of-the-art methods, excelling at some smaller and general categories (i.e. motorcycle, other-veh.).

Method	Input																				
		SC IoU	SSC mIoU	car	bicycle	motorcycle	truck	other-veh.	person	road	parking	sidewalk	other-grnd	building	fence	vegetation	terrain	pole	traf.-sign	other-struct.	other-object
LMSNet [44]	L	47.53	13.65	20.91	0	0	0.26	0	0	62.95	13.51	33.51	0.2	43.67	0.33	40.01	26.80	0	0	3.63	0
SSCNet [45]	L	53.58	16.95	31.95	0	0.17	10.29	0.58	0.07	65.7	17.33	41.24	3.22	44.41	6.77	43.72	28.87	0.78	0.75	8.60	0.67
MonoScene [4]	C	37.87	12.31	19.34	0.43	0.58	8.02	2.03	0.86	48.35	11.38	28.13	3.22	32.89	3.53	26.15	<b>16.75</b>	6.92	<b>5.67</b>	4.20	3.09
Voxformer [24]	C	38.76	11.91	17.84	<b>1.16</b>	0.89	4.56	2.06	1.63	47.01	9.67	27.21	2.89	31.18	<b>4.97</b>	28.99	14.69	6.51	<b>6.92</b>	3.79	2.43
TPVFormer [17]	C	<b>40.22</b>	<b>13.64</b>	<b>21.56</b>	<b>1.09</b>	<b>1.37</b>	8.06	2.57	2.38	<b>52.99</b>	<b>11.99</b>	<b>31.07</b>	<b>3.78</b>	<b>34.83</b>	4.80	<b>30.08</b>	<b>17.51</b>	<b>7.46</b>	5.86	5.48	2.70
OccFormer [58]	C	<b>40.27</b>	<b>13.81</b>	<b>22.58</b>	0.66	0.26	<b>9.89</b>	<b>3.82</b>	<b>2.77</b>	<b>54.30</b>	<b>13.44</b>	<b>31.53</b>	3.55	<b>36.42</b>	4.80	<b>31.00</b>	<b>19.51</b>	<b>7.77</b>	<b>8.51</b>	<b>6.95</b>	<b>4.60</b>
Symphonies [19]	C	<b>44.12</b>	<b>18.58</b>	<b>30.02</b>	<b>1.85</b>	<b>5.90</b>	<b>25.07</b>	<b>12.06</b>	<b>8.20</b>	<b>54.94</b>	<b>13.83</b>	<b>32.76</b>	<b>6.93</b>	<b>35.11</b>	<b>8.58</b>	<b>38.33</b>	11.52	<b>14.01</b>	<b>9.57</b>	<b>14.44</b>	<b>11.28</b>
<b>Ours</b>	C	35.38	12.92	18.93	1.02	<b>4.62</b>	<b>18.07</b>	<b>7.59</b>	<b>3.35</b>	45.47	10.89	25.03	<b>5.32</b>	28.44	<b>5.68</b>	29.54	8.62	2.99	2.32	<b>9.51</b>	<b>5.14</b>

**Monocular 3D semantic occupancy prediction.** Table 2 compares the performance of GaussianFormer with other methods for monocular 3D semantic occupancy prediction on SSCBench-KITTI-360. Notably, GaussianFormer achieves comparable performance with state-of-the-art models, excelling at some smaller categories such as motorcycle and general categories such as other-vehicle. This is due to 3D Gaussians can adaptively change their positions and covariance to match the boundaries of small objects in images in contrast to rigid grid projections on images which might be misleading. Furthermore, the flexibility of 3D Gaussians also benefits the predictions for general objects (i.e. categories with other- prefix) which often have distinct shapes and appearances with normal categories.

**Efficiency comparisons with existing methods.** We provide the efficiency comparisons of different scene representations in Table 3. Notably, GaussianFormer surpasses all existing competitors with significantly reduced mem-

**Table 3: Efficiency comparison of different representations** on nuScenes. The latency and memory consumption for GaussianFormer are tested on one NVIDIA 4090 GPU with batch size one, while the results for other methods are reported in OctreeOcc [34] tested on one NVIDIA A100 GPU. Our method demonstrates significantly reduced memory usage compared to other representations.

Methods	Query Form	Query Resolution	Latency ↓	Memory ↓
BEVFormer [27]	2D BEV	200×200	<b>302</b> ms	25100 M
TPVFormer [17]	2D tri-plane	200×(200+16+16)	341 ms	29000 M
PanoOcc [50]	3D voxel	100×100 ×16	502 ms	35000 M
FBOCC [28]	3D voxel & 2D BEV	200×200×16 & 200×200	463 ms	31000 M
OctreeOcc [34]	Octree Query	91200	386 ms	26500 M
GaussianFormer	3D Gaussian	144000	372 ms	<b>6229</b> M

**Table 4: Ablation on the components of GaussianFormer.** Deep Supervision represents supervising the output of each refinement module. Residual Refine means on which properties of Gaussian to apply residual refinement as opposed to substitution.

Deep Supervision	Sparse Conv.	Residual Refine	mIoU	IoU
	✓	mean	16.36	29.32
✓		mean	15.93	28.99
✓	✓	none	-	-
✓	✓	all except semantics	16.24	29.30
✓	✓	mean	<b>16.41</b>	<b>29.37</b>

ory consumption. The memory efficiency of GaussianFormer originates from its object-centric nature which assigns explicit semantic meaning to each 3D Gaussian, and thus greatly simplifies the transformation from the scene representation to occupancy predictions, getting rid of the expensive decoding process from high dimensional features. While being slightly slower ( $\sim 70$  ms) than methods based on planar representations [17, 27], GaussianFormer achieves the lowest latency among dense grid representations. Notably, our method is faster than OctreeOcc [34] even with more queries.

**Analysis of components of GaussianFormer.** In Table 4, we provide comprehensive analysis on the components of GaussianFormer to validate their effectiveness. We conduct these experiments on nuScenes and set the number of 3D Gaussians to 51200. The strategy for refinement of Gaussian properties has a notable influence on the performance. The experiment for the all substitution strategy (denoted by none) collapses where we directly replace the old properties with new ones. This is because the positions of Gaussians are sensitive to noise which quickly converge to a trivial solution without regularization for coherence during refinement. Rectifying all properties except semantics in a residual style (denoted by all except semantics) is also harmful because the sigmoid activation used for covariance is prone to vanishing gradient. Moreover, the 3D sparse convolution in the self-encoding module is crucial for the performance because it

**Table 5: Ablation on the number of Gaussians.** The latency and memory are tested on an NVIDIA 4090 GPU with batch size one during inference. The performance improves consistently with more Gaussians while taking up more time and memory.

Number of Gaussians	Latency	Memory	mIoU	IoU
25600	<b>227</b> ms	<b>4850</b> M	16.00	28.72
38400	249 ms	4856 M	16.04	28.72
51200	259 ms	4866 M	16.41	29.37
91200	293 ms	5380 M	18.31	27.48
144000	372 ms	6229 M	<b>19.10</b>	<b>29.83</b>

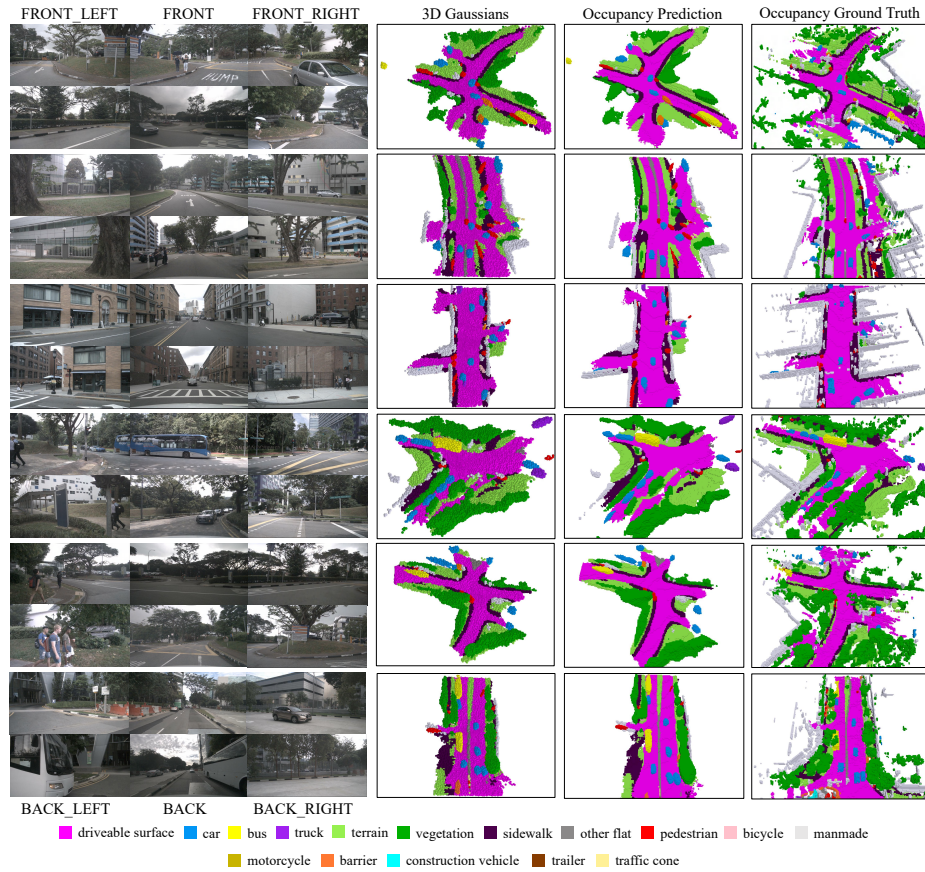
is responsible for the interactions among 3D Gaussians. On the other hand, the deep supervision strategy also contributes to the overall performance by ensuring that every intermediate refinement step benefits the 3D perception.

**Effect of the number of Gaussians.** Table 5 presents the ablation study on the number of Gaussians, where we analyze its influence on efficiency and performance. The mIoU increases linearly when the number of 3D Gaussians is greater than 38400, which is due to the enhanced ability to represent finer details with more Gaussians. The latency and memory consumption also correlate linearly with the number of Gaussians, offering flexibility for deployment.

**Visualization results.** We provide qualitative visualization results in Fig. 5. Our GaussianFormer can generate a holistic and realistic perception of the scene. Specifically, the 3D Gaussians adjust their covariance matrices to capture the fine details of object shapes such as the flat-shaped Gaussians at the surface of the road and walls (e.g. in the third row). Additionally, the density is higher in regions with vehicles and pedestrians compared with the road surface (e.g. in the first row), which proves that the 3D Gaussians cluster around foreground objects with iterative refinement for a reasonable allocation of resources. Furthermore, our GaussianFormer even successfully predicts objects that are not in the ground truth and barely visible in the images, such as the truck in the front right input image and the upper right corner of the 3D visualizations in the fourth row.

## 5 Conclusion and Discussions

In this paper, we have proposed an efficient object-centric 3D Gaussian representation for 3D semantic occupancy prediction to better exploit the sparsity of occupancy and the diversity of object scales. We describe the driving scenes with sparse 3D Gaussians each of which is characterized by its position, covariance and semantics and represents a flexible region of interest. Based on the 3D Gaussian representation, we have designed the GaussianFormer to effectively learn 3D Gaussians from input images through attention mechanism and iterative refinement. To efficiently generate voxelized occupancy predictions from 3D Gaussians, we have proposed an efficient Gaussian-to-voxel splatting method, which only aggregates the neighboring Gaussians for each voxel. GaussianFormer has achieved comparable performance with state-of-the-art methods



**Fig. 5: Visualization results for 3D semantic occupancy prediction on nuScenes.** We visualize the 3D Gaussians by treating them as ellipsoids centered at the Gaussian means with semi-axes determined by the Gaussian covariance matrices. Our GaussianFormer not only achieves reasonable allocation of resources, but also captures the fine details of object shapes.

on the nuScenes and KITTI-360 datasets, and significantly reduced the memory consumption by more than 75%. Our ablation study has shown that the performance of GaussianFormer scales well to the number of Gaussians. Additionally, the visualization has proved the abilities of 3D Gaussians to capture the details of object shapes and to reasonably allocate computation and storage resources.

**Limitations.** The performance of GaussianFormer is still inferior to state-of-the-art methods despite the much lower memory consumption. This might result from the inaccuracy of 3D semantic Gaussian representation or simply the wrong choice of hyperparameter since we did not perform much hyperparameter tuning. GaussianFormer also requires a large number of Gaussians to achieve satisfactory performance. This is perhaps because the current 3D semantic Gaussians include empty as one category and thus still can be redundant. It is interesting to only model solid objects to further improve performance and speed.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant 62125603, Grant 62321005, and Grant 62336004.

## References

1. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: ICCV. pp. 5855–5864 (2021)
2. Berman, M., Triki, A.R., Blaschko, M.B.: The lovász-softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4413–4421 (2018)
3. 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)
4. Cao, A.Q., de Charette, R.: Monoscene: Monocular 3d semantic scene completion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3991–4001 (2022)
5. Charatan, D., Li, S., Tagliasacchi, A., Sitzmann, V.: pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. arXiv preprint arXiv:2312.12337 (2023)
6. Chen, G., Wang, W.: A survey on 3d gaussian splatting. arXiv preprint arXiv:2401.03890 (2024)
7. Chen, X., Lin, K.Y., Qian, C., Zeng, G., Li, H.: 3d sketch-aware semantic scene completion via semi-supervised structure prior. In: CVPR (2020)
8. Chen, Z., Wang, F., Liu, H.: Text-to-3d using gaussian splatting. arXiv preprint arXiv:2309.16585 (2023)
9. Dalal, S., Hall, W.: Approximating priors by mixtures of natural conjugate priors. *Journal of the Royal Statistical Society* **45**(2), 278–286 (1983)
10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR. pp. 248–255. Ieee (2009)
11. Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: CVPR. pp. 5501–5510 (2022)
12. Goodfellow, I., Bengio, Y., Courville, A.: *Deep learning* (2016)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition (2016)
14. Hu, Y., Yang, J., Chen, L., Li, K., Sima, C., Zhu, X., Chai, S., Du, S., Lin, T., Wang, W., et al.: Goal-oriented autonomous driving. arXiv preprint arXiv:2212.10156 (2022)
15. 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)
16. Huang, Y., Zheng, W., Zhang, B., Zhou, J., Lu, J.: Selfocc: Self-supervised vision-based 3d occupancy prediction. In: CVPR (2024)
17. Huang, Y., Zheng, W., Zhang, Y., Zhou, J., Lu, J.: Tri-perspective view for vision-based 3d semantic occupancy prediction. In: CVPR. pp. 9223–9232 (2023)

18. Jiang, B., Chen, S., Xu, Q., Liao, B., Chen, J., Zhou, H., Zhang, Q., Liu, W., Huang, C., Wang, X.: Vad: Vectorized scene representation for efficient autonomous driving. arXiv preprint arXiv:2303.12077 (2023)
19. Jiang, H., Cheng, T., Gao, N., Zhang, H., Liu, W., Wang, X.: Symphonize 3d semantic scene completion with contextual instance queries. arXiv preprint arXiv:2306.15670 (2023)
20. Jiang, Y., Zhang, L., Miao, Z., Zhu, X., Gao, J., Hu, W., Jiang, Y.G.: Polarformer: Multi-camera 3d object detection with polar transformer. In: AAAI. vol. 37, pp. 1042–1050 (2023)
21. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics **42**(4) (2023)
22. Li, J., Han, K., Wang, P., Liu, Y., Yuan, X.: Anisotropic convolutional networks for 3d semantic scene completion. In: CVPR (2020)
23. Li, Y., Li, S., Liu, X., Gong, M., Li, K., Chen, N., Wang, Z., Li, Z., Jiang, T., Yu, F., et al.: Sscbench: A large-scale 3d semantic scene completion benchmark for autonomous driving. arXiv preprint arXiv:2306.09001 (2023)
24. 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. pp. 9087–9098 (2023)
25. Li, Y., Bao, H., Ge, Z., Yang, J., Sun, J., Li, Z.: Bevstereo: Enhancing depth estimation in multi-view 3d object detection with temporal stereo. In: AAAI. vol. 37, pp. 1486–1494 (2023)
26. Li, Y., Ge, Z., Yu, G., Yang, J., Wang, Z., Shi, Y., Sun, J., Li, Z.: Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. arXiv preprint arXiv:2206.10092 (2022)
27. Li, Z., Wang, W., Li, H., Xie, E., Sima, C., Lu, T., Yu, Q., Dai, J.: Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. arXiv preprint arXiv:2203.17270 (2022)
28. 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)
29. Liang, Y., Yang, X., Lin, J., Li, H., Xu, X., Chen, Y.: Luciddreamer: Towards high-fidelity text-to-3d generation via interval score matching. arXiv preprint arXiv:2311.11284 (2023)
30. Liao, Y., Xie, J., Geiger, A.: KITTI-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. PAMI (2022)
31. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR (2017)
32. Liu, Z., Tang, H., Amini, A., Yang, X., Mao, H., Rus, D.L., Han, S.: Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. In: ICRA. pp. 2774–2781. IEEE (2023)
33. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
34. Lu, Y., Zhu, X., Wang, T., Ma, Y.: Octreeocc: Efficient and multi-granularity occupancy prediction using octree queries. arXiv preprint arXiv:2312.03774 (2023)
35. Luiten, J., Kopanas, G., Leibe, B., Ramanan, D.: Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. arXiv preprint arXiv:2308.09713 (2023)
36. Miao, R., Liu, W., Chen, M., Gong, Z., Xu, W., Hu, C., Zhou, S.: Ocdepth: A depth-aware method for 3d semantic scene completion. arXiv preprint arXiv:2302.13540 (2023)



37. Mildenhall, B., Srinivasan, P.P., Tanck, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
38. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ToG* **41**(4), 1–15 (2022)
39. 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*. pp. 414–431. Springer (2020)
40. Park, J., Xu, C., Yang, S., Keutzer, K., Kitani, K., Tomizuka, M., Zhan, W.: Time will tell: New outlooks and a baseline for temporal multi-view 3d object detection. *arXiv preprint arXiv:2210.02443* (2022)
41. Phillion, J., Fidler, S.: Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In: *ECCV* (2020)
42. Qiao, Y.L., Gao, A., Xu, Y., Feng, Y., Huang, J.B., Lin, M.C.: Dynamic mesh-aware radiance fields. In: *ICCV*. pp. 385–396 (2023)
43. Rakotosaona, M.J., Manhardt, F., Arroyo, D.M., Niemeyer, M., Kundu, A., Tombari, F.: Nerfmeshing: Distilling neural radiance fields into geometrically-accurate 3d meshes. *arXiv preprint arXiv:2303.09431* (2023)
44. Roldão, L., de Charette, R., Verroust-Blondet, A.: Lmscnet: Lightweight multiscale 3d semantic completion. In: *ThreeDV* (2020)
45. Song, S., Yu, F., Zeng, A., Chang, A.X., Savva, M., Funkhouser, T.: Semantic scene completion from a single depth image. In: *CVPR*. pp. 1746–1754 (2017)
46. Tang, J., Ren, J., Zhou, H., Liu, Z., Zeng, G.: Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653* (2023)
47. Tong, W., Sima, C., Wang, T., Chen, L., Wu, S., Deng, H., Gu, Y., Lu, L., Luo, P., Lin, D., et al.: Scene as occupancy. In: *ICCV*. pp. 8406–8415 (2023)
48. Wang, T., Zhu, X., Pang, J., Lin, D.: Fcos3d: Fully convolutional one-stage monocular 3d object detection. In: *ICCV*. pp. 913–922 (2021)
49. 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)
50. 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)
51. Wei, Y., Zhao, L., Zheng, W., Zhu, Z., Zhou, J., Lu, J.: Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving. In: *ICCV*. pp. 21729–21740 (2023)
52. Yan, X., Gao, J., Li, J., Zhang, R., Li, Z., Huang, R., Cui, S.: Sparse single sweep lidar point cloud segmentation via learning contextual shape priors from scene completion. In: *AAAI*. vol. 35, pp. 3101–3109 (2021)
53. Yang, B., Bao, C., Zeng, J., Bao, H., Zhang, Y., Cui, Z., Zhang, G.: Neumesh: Learning disentangled neural mesh-based implicit field for geometry and texture editing. In: *ECCV*. pp. 597–614. Springer (2022)
54. Yang, C., Chen, Y., Tian, H., Tao, C., Zhu, X., Zhang, Z., Huang, G., Li, H., Qiao, Y., Lu, L., et al.: Bevformer v2: Adapting modern image backbones to bird’s-eye-view recognition via perspective supervision. *arXiv preprint arXiv:2211.10439* (2022)
55. Yang, Z., Gao, X., Zhou, W., Jiao, S., Zhang, Y., Jin, X.: Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. *arXiv preprint arXiv:2309.13101* (2023)

56. Yi, T., Fang, J., Wu, G., Xie, L., Zhang, X., Liu, W., Tian, Q., Wang, X.: Gaus-siandreamer: Fast generation from text to 3d gaussian splatting with point cloud priors. arXiv preprint arXiv:2310.08529 (2023)
57. Yu, Z., Shu, C., Deng, J., Lu, K., Liu, Z., Yu, J., Yang, D., Li, H., Chen, Y.: Flashocc: Fast and memory-efficient occupancy prediction via channel-to-height plugin. arXiv preprint arXiv:2311.12058 (2023)
58. Zhang, Y., Zhu, Z., Du, D.: Occformer: Dual-path transformer for vision-based 3d semantic occupancy prediction. arXiv preprint arXiv:2304.05316 (2023)
59. Zheng, W., Chen, W., Huang, Y., Zhang, B., Duan, Y., Lu, J.: Occworld: Learning a 3d occupancy world model for autonomous driving. arXiv preprint arXiv:2311.16038 (2023)
60. Zheng, W., Song, R., Guo, X., Chen, L.: Genad: Generative end-to-end autonomous driving. arXiv preprint arXiv:2402.11502 (2024)
61. Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. In: CVPR. pp. 4490–4499 (2018)
62. 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)
63. Zou, Z.X., Yu, Z., Guo, Y.C., Li, Y., Liang, D., Cao, Y.P., Zhang, S.H.: Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. arXiv preprint arXiv:2312.09147 (2023)
64. Zuo, S., Zheng, W., Huang, Y., Zhou, J., Lu, J.: Pointocc: Cylindrical tri-perspective view for point-based 3d semantic occupancy prediction. arXiv preprint arXiv:2308.16896 (2023)