

GaussianFlowOcc: Sparse and Weakly Supervised Occupancy Estimation using Gaussian Splatting and Temporal Flow

Simon Boeder
Robert Bosch GmbH

simon.boeder@de.bosch.com

Fabian Gigengack
Robert Bosch GmbH

fabian.gigengack@de.bosch.com

Benjamin Risse
University of Münster

b.risse@uni-muenster.de

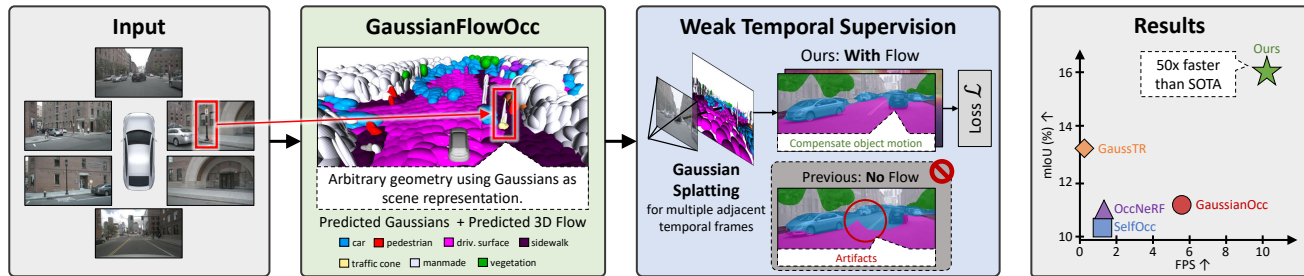


Figure 1. **Contributions of GaussianFlowOcc.** We propose an efficient and fast model for estimating dynamic scenes using 3D Gaussian distributions, trained with weak supervision. By employing a *Temporal Module*, we efficiently model scene dynamics to account for object motion during training. Our approach significantly outperforms previous methods in terms of mIoU and inference speed.

Abstract

Occupancy estimation has become a prominent task in 3D computer vision, particularly within the autonomous driving community. In this paper, we present a novel approach to occupancy estimation, termed GaussianFlowOcc, which is inspired by Gaussian Splatting and replaces traditional dense voxel grids with a sparse 3D Gaussian representation. Our efficient model architecture based on a Gaussian Transformer significantly reduces computational and memory requirements by eliminating the need for expensive 3D convolutions used with inefficient voxel-based representations that predominantly represent empty 3D spaces. GaussianFlowOcc effectively captures scene dynamics by estimating temporal flow for each Gaussian during the overall network training process, offering a straightforward solution to a complex problem that is often neglected by existing methods. Moreover, GaussianFlowOcc is designed for scalability, as it employs weak supervision and does not require costly dense 3D voxel annotations based on additional data (e.g., LiDAR). Through extensive experimentation, we demonstrate that GaussianFlowOcc significantly outperforms all previous methods for weakly supervised occupancy estimation on the nuScenes dataset while featuring an inference speed that is 50 times faster than current SOTA.

1. Introduction

Vision-based 3D object detection and occupancy estimation are fundamental challenges in the field of autonomous driving, as an accurate understanding of a vehicle’s surroundings is essential for reliable planning and navigation [46, 62, 73]. Despite recent advances, existing approaches face several limitations. First, most recent occupancy estimation models rely on fully annotated 3D ground truth data for training, which is expensive and challenging to obtain at scale. Second, these models depend on dense 3D voxel-based representations, leading to substantial computational overhead that is often unnecessary, given the natural sparsity of real-world 3D scenes. Although self-supervised methods have been developed to reduce reliance on 3D labels by leveraging temporal rendering [20, 70], these approaches fail to account for the inherent dynamics of driving scenarios. This introduces temporal inconsistencies during training, severely limiting their performance.

To address these limitations, we introduce *GaussianFlowOcc*, a novel occupancy estimation method that combines efficiency and expressiveness through several contributions. At the core of our approach is a sparse representation of the scene that utilizes 3D Gaussian distributions, replacing traditional dense voxel grids. We propose a *Gaussian Transformer* network, a novel deep learning model designed to directly transform image features into 3D Gaus-

sians. This shift not only reduces computational overhead but also provides a more flexible representation for modeling dynamic scenes. Additionally, it allows us to leverage *Gaussian Splatting* [25] for efficient training without relying on 3D voxel labels. The estimated Gaussians are rendered back into the input camera space, where losses can be computed using 2D labels such as pseudo ground truth derived from pretrained models like Grounded-SAM [45] for semantics and Metric3D [66] for depth. Another key contribution is *Temporal Gaussian Splatting*, where adjacent frames are rendered to enhance the supervisory signal, allowing the model to capture 3D geometry more effectively. To handle temporal inconsistencies caused by dynamic objects, GaussianFlowOcc extends this framework with a *Temporal Module* that estimates a 3D flow for each Gaussian to each rendered temporal frame. This additional modeling of scene dynamics mitigates errors caused by object motion during training, strengthening the supervision signal and improving the overall performance. By combining these techniques, GaussianFlowOcc achieves state-of-the-art performance while maintaining computational efficiency, rendering it highly effective for practical occupancy estimation applications. In summary, our contributions are:

- **Efficient Gaussian Transformer for sparse 3D Gaussian scene representation:** We introduce a novel *Gaussian Transformer* that leverages set-based attention mechanisms to efficiently estimate a sparse set of 3D Gaussian distributions representing the scene from multi-view images. It entirely eliminates voxelization and computationally intensive operations, leading to an inference speedup of $50\times$ compared to state-of-the-art. At the same time, by leveraging Gaussian Splatting, our method learns from 2D pseudo labels instead of relying on costly 3D annotations from other sensors like LiDAR. Our approach requires only images for training, making it more scalable than methods depending on 3D labeled datasets.
- **Modeling of scene dynamics via learned 3D flow:** GaussianFlowOcc integrates a *Temporal Module* that additionally estimates 3D flow for each Gaussian, enabling robust handling of dynamic objects during inference and during *Temporal Gaussian Splatting* at training time. GaussianFlowOcc is the first method to incorporate temporal supervision while explicitly modeling scene dynamics, without requiring ground truth flow annotations.
- **State-of-the-art performance:** On top of being highly efficient at inference, GaussianFlowOcc substantially outperforms previous weakly supervised methods.

2. Related Work

2.1. Occupancy Estimation

The 3D semantic occupancy estimation task has become an important area in the autonomous driving research com-

munity in recent years. Numerous voxel-based occupancy benchmarks have been introduced for datasets such as SemanticKITTI [3] and nuScenes [53, 58]. Early works in 3D occupancy estimation utilize established techniques from Birds-Eye-View (BEV) perception and object detection [19, 29] to lift multi-view camera images into a unified 3D voxel grid [7, 18, 21, 28, 54]. Subsequent approaches have improved model efficiency [31, 35, 47, 52, 56, 68], optimized the training procedure and label efficiency [5, 14, 16, 41, 49], and improved occupancy estimation performance through architectural innovations [23, 30, 36, 37, 50, 72, 74]. Specifically, GaussianFormer [22] and GaussianWorld [81] employ a sparse and efficient Gaussian representation similar to our work.

Despite these advancements, the majority of existing 3D occupancy estimation methods require costly voxel-based 3D ground truth labels for training. Therefore, a parallel line of research has emerged that focuses on self and weakly supervised learning for occupancy estimation models, leveraging only 2D labels. SelfOcc [20] and OccNeRF [70] employ volume rendering inspired by NeRF [38] to render estimated 3D occupancy back to the 2D image space. Photometric consistency losses can be used to train the geometry estimation, while semantic information is incorporated using pretrained foundation models like OpenSeeD [71] or GroundedSAM [45]. GaussianOcc [15] replaces the volume rendering pipeline with 3D Gaussian Splatting to accelerate training, yet it continues to model the scene with dense occupancy grids, thus failing to exploit the benefits of a fully sparse Gaussian representation. Also, while these methods eliminate the need for 3D labels, they fail to address scene dynamics, a critical aspect when relying on temporal consistency losses.

A growing body of research aims to align 3D occupancy with feature spaces of strong foundation models. OccFeat [48] distills features of CLIP [44] and DINO [8, 40] into an occupancy representation for model pretraining. POP-3D [55], LOcc [69] and OVO [51] perform open-vocabulary occupancy estimation by aligning voxel-based predictions with vision-language features extracted from pretrained encoders like MaskCLIP [78]. VEON [75] and LangOcc [4] follow up on the self-supervised methods and directly use volume rendering of vision-language features to train open-vocabulary models.

2.2. Differentiable Rendering and 3D Gaussian Splatting

Differentiable rendering has emerged as a powerful technique for learning 3D scene representations by projecting them into 2D views, followed by an optimization based on photometric or semantic consistency. Neural Radiance Fields (NeRF) [38] have been particularly influential, modeling scenes as volumetric representations that encode ra-

diance and density, enabling novel view synthesis through differentiable volume rendering. Recently, 3D Gaussian Splatting (GS) [25] has introduced a novel paradigm for 3D scene reconstruction by representing scenes as a collection of 3D Gaussians. This approach significantly reduces computational overhead while preserving expressive scene modeling, making it highly efficient for dynamic or large-scale reconstructions. NeRF and GS approaches were originally designed to reconstruct individual scenes for novel-view synthesis, with research focusing on improving efficiency [11, 39], rendering quality [1, 2] or feature enrichment [26, 43, 65, 79]. Several works have further explored the modeling of dynamic scenes for video reconstruction [13, 42, 61]. Another line of work incorporates different priors like depth, stereo-matching, or additional data like LiDAR, to train generalizable reconstruction models [9, 10, 32, 60, 63, 67, 76]. Several methods have also been developed specifically for reconstruction within autonomous driving scenarios [34, 64, 80], while GaussianAD [77] employs a 3D Gaussian representation for end-to-end driving. As mentioned above, similar ideas have been adapted to train occupancy estimation models. Methods like OccNeRF [70], SelfOcc [20] and GaussianOcc [15] employ volume rendering or GS for weakly supervised training. However, they still rely on voxel grids for scene representation, which limits efficiency and scalability. GSRender [49] additionally introduces a ray compensation method to mitigate duplicate predictions along camera rays. Lastly, the recent approach GaussTR [24] shares similarities with our method by adopting 3D Gaussians as the scene representation. However, GaussTR uses multiple pretrained feature encoders (e.g., CLIP, Metric3D, SAM) during inference, making the pipeline computationally expensive. Furthermore, GaussTR employs standard attention layers, which significantly limits the number of Gaussians it can handle.

3. Methodology

3.1. Problem Definition and Scene Representation

The objective of occupancy prediction is to accurately estimate the 3D geometry and semantics surrounding a vehicle based on a set of L multi-view images $I = \{I^1, I^2, \dots, I^L\}$ at the current time step. Previous methods have typically utilized a semantic voxel volume $V = \{c_1, c_2, \dots, c_C\}^{X \times Y \times Z}$ on a predefined grid, with C representing the number of semantic classes, as a structured representation of the scene. In contrast, our approach involves estimating the 3D scene using 3D Gaussian distributions, drawing inspiration from 3D Gaussian Splatting [25]. The scene is defined as a collection of N Gaussians $\mathcal{G} = \{G_1, G_2, \dots, G_N\}$, each characterized by a set of properties: mean $\mu \in \mathbb{R}^3$, opacity $\sigma \in [0, 1]$, scale $s \in \mathbb{R}^3$,

rotation quaternions $r \in \mathbb{R}^4$, and semantic logits $c \in \mathbb{R}^C$:

$$G_i = (\mu, \sigma, s, r, c). \quad (1)$$

We train a model $\mathbb{M}(I)$ to estimate the 3D Gaussians \mathcal{G} . If desired, the resulting scene representation can be easily converted to a voxel volume as a post-processing step to facilitate comparison with previous methods (see Sec. 3.7).

3.2. Model Architecture

The architecture of the proposed model is illustrated in Fig. 2. Initially, image features \hat{I} are extracted from the input images using an image encoder. The image features, along with the initial positions and features of the Gaussians, are then fed into the proposed *Gaussian Transformer* (Sec. 3.2.3). The *Gaussian Transformer* iteratively transforms the initial Gaussians to their final positions by employing cross-attention to the image features, self-attention between the Gaussians, and temporal attention to the Gaussians of the previous frame. Subsequently, the *Gaussian Heads* (Sec. 3.2.4) estimate the remaining Gaussian properties and their semantics. Concurrently, the *Temporal Module* (Sec. 3.6) computes a 3D flow for each Gaussian to its temporal neighbors. The final Gaussians and their temporal offsets are then utilized in the Gaussian Splatting pipeline (Sec. 3.4) to render depth and semantic maps into the input cameras and a set of temporally adjacent frames. The model is trained using losses between the rendered and ground truth depth and semantic segmentation maps.

3.2.1. Initial Gaussians

We initialize with a set of N Gaussian means $\mathcal{G}_\mu^0 \in \mathbb{R}^{N \times 3}$ and latent features $\mathcal{G}_f^0 \in \mathbb{R}^{N \times D}$, akin to the queries in previous works [29, 33, 57]. The initial positions and features of the Gaussians are learnable, enabling the model to incorporate prior knowledge of driving scenes. Notably, the other Gaussian properties o, s, r, c are left uninitialized; instead, the model learns to infer them from the Gaussian features through the *Gaussian Heads* in the output stage, a strategy that has proven to be more robust for training.

3.2.2. Image Encoder

Image features \hat{I} are extracted from the input images I using a pretrained backbone architecture like *ResNet-50* [17]. These features and the initial Gaussians are then fed into the *Gaussian Transformer*.

3.2.3. Gaussian Transformer

The *Gaussian Transformer* iteratively updates the initial Gaussian means and features over B blocks, with each block b comprising five successive modules, elaborated in the subsequent sections. An overview of a *Gaussian Transformer* block is depicted in Fig. 3.

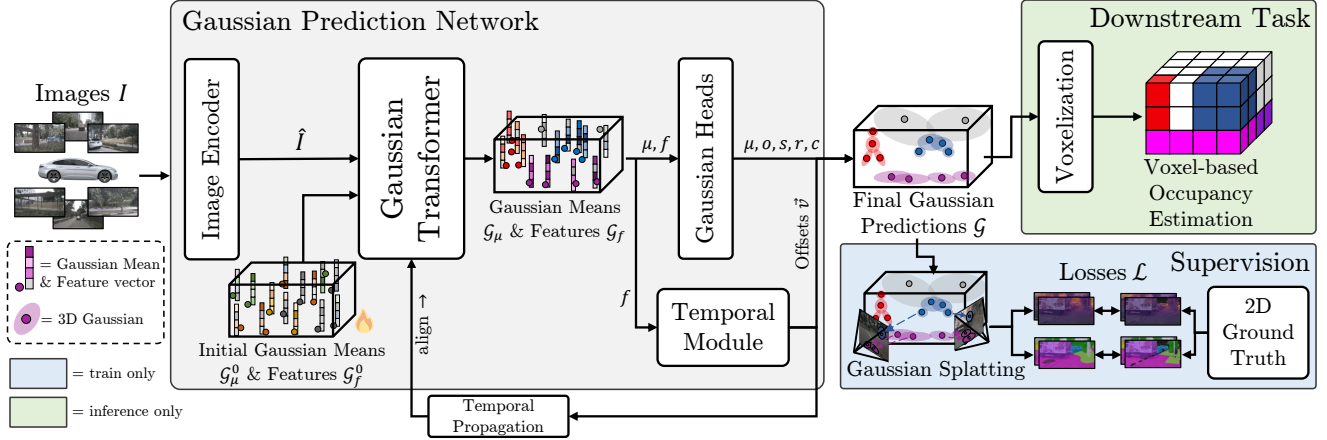


Figure 2. **Overview of GaussianFlowOcc.** After encoding a set of input images, the *Gaussian Transformer* iteratively estimates the position and features of 3D Gaussian distributions. The *Gaussian Heads* then predict opacity, scale, rotation and semantics of these Gaussians. The model is trained using Gaussian Splatting with 2D labels, generated from off-the-shelf models. The *Temporal Module* simultaneously estimates 3D temporal offsets for each Gaussian to correct temporal inconsistencies when using *Temporal Gaussian Splatting*.

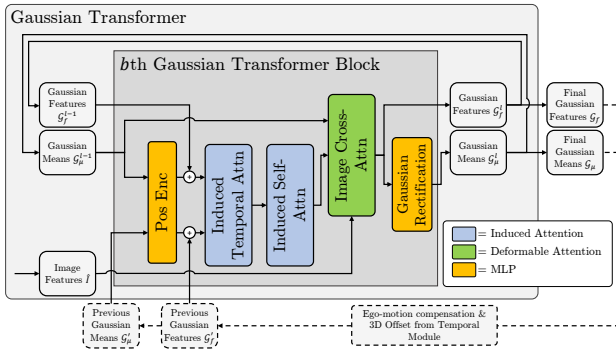


Figure 3. **Architecture of the proposed *Gaussian Transformer*.** In each block, the features and means of the Gaussians are refined using Pos Enc, ITA, ISA, GICA and Gaussian Rectification.

Positional Encoding The Gaussian means of the previous block \mathcal{G}_μ^{b-1} are encoded into the latent dimension D using a Multi-Layer Perceptron (MLP) and are added onto the previous features \mathcal{G}_f^{b-1} .

Gaussian-Image Cross-Attention To enable Gaussians to acquire meaningful scene information, we introduce the *Gaussian-Image Cross-Attention* (GICA), which facilitates interactions between Gaussians \mathcal{G} and the image features \hat{I} . For this we leverage deformable cross-attention, originally proposed by [29] and widely adopted in subsequent works. In this operation, the current positions of the Gaussians \mathcal{G}_μ are projected onto the image feature maps using the camera parameters. For each projected point, a set of image features around this point is sampled and used as keys and values for an attention operation. This allows the Gaussians to extract rich, localized information from the scene in a com-

putationally efficient manner. For a detailed description of deformable attention, we refer readers to [29].

Induced Self-Attention Computing interactions among Gaussians is a critical aspect of our approach. However, the standard attention mechanism is characterized by quadratic time and memory complexity $\mathcal{O}(N^2)$. This imposes severe constraints on the number of Gaussians that can be processed (as demonstrated in Fig. 5). To address this limitation, we draw inspiration from the Set Transformer [27] and employ an attention operation called *Induced Self-Attention* (ISA). ISA reduces the memory complexity to approximately linearly scaling with respect to the number of Gaussians, allowing to employ a significantly larger number of Gaussians. The operation is schematically illustrated in Fig. 4 (A). In essence, ISA replaces quadratic attention with two attention operations connected via a bottleneck. Instead of allowing Gaussians to directly attend to one another, we introduce a set of M trainable latent feature vectors, called inducing points $P \in \mathbb{R}^{M \times D}$ (with $M \ll N$). These inducing points aggregate information from all Gaussians into a bottleneck representation H . Subsequently, all Gaussians interact with the bottleneck features H enabling indirect yet complete Gaussian-to-Gaussian interactions. This transformation reduces the computational complexity to $\mathcal{O}(MN)$. The overall process is formally defined as:

$$\text{ISA}(\mathcal{G}_f) = \text{MHA}(\mathcal{G}_f, H, H), \quad (2)$$

$$\text{where } H = \text{MHA}(P, \mathcal{G}_f, \mathcal{G}_f), \quad (3)$$

where $\text{MHA}(Q, K, V)$ denotes the standard multi-head attention operation, which includes a skip connection and a feed-forward layer.




Figure 4. **Illustration of the Induced Attention Modules.** (A) Induced Self-Attention (ISA) (B) Induced Temporal Attention (ITA)

Induced Temporal Attention Building upon the concept of ISA, we introduce *Induced Temporal Attention* (ITA) to enable efficient temporal information propagation across frames, as depicted in Fig. 4 (B). The ITA operation leverages the structure of induced attention by letting the inducing points P first attend to the Gaussian features from the previous frame to compute the bottleneck features Z . Subsequently, Gaussian features from the current frame interact with the bottleneck features, facilitating temporal attention with linear complexity. Formally, given \mathcal{G}'_f as the final Gaussian features from the previous frame, the ITA operation is expressed as:

$$ITA(\mathcal{G}_f, \mathcal{G}'_f) = MHA(\mathcal{G}_f, Z, Z), \quad (4)$$

$$\text{where } Z = MHA(P, \mathcal{G}'_f, \mathcal{G}'_f). \quad (5)$$

Gaussian Rectification Thus far, only the Gaussian features have been updated using cross-attention, self-attention and temporal attention. In the *Gaussian Rectification* step, these features are used to estimate an update to the Gaussian means for the next block. Unlike prior approaches such as GaussianFormer [22] and GaussTR [24], our method refrains from modifying Gaussian properties other than the means within the transformer itself. Instead, we focus on gathering features that encapsulate all the necessary information. The update is performed using an MLP, denoted as $Rect(\mathcal{G}_f)$, which estimates a residual for each Gaussian and adds it to the current means to refine their position:

$$\mathcal{G}_\mu^{b+1} = \mathcal{G}_\mu^b + \Delta \mathcal{G}_\mu^b, \quad (6)$$

$$\text{where } \Delta \mathcal{G}_\mu^b = Rect(\mathcal{G}_f^b). \quad (7)$$

3.2.4. Gaussian Heads

After refining the Gaussian positions and features through B blocks of the *Gaussian Transformer*, the *Gaussian Heads* estimate the remaining properties of the Gaussians, completing the scene representation. This is accomplished using a set of individual MLP heads, each dedicated to predict a specific property from the Gaussian features. To ensure valid outputs, a sigmoid activation $\sigma(\cdot)$ function is applied to estimate the opacity o , while the rotation coefficients r are normalized to unit length. The outputs are computed as

follows:

$$o = \sigma(\text{MLP}_o(\mathcal{G}_f)), \quad s = \text{MLP}_s(\mathcal{G}_f), \quad (8)$$

$$r = \text{norm}(\text{MLP}_r(\mathcal{G}_f)), \quad c = \text{MLP}_c(\mathcal{G}_f). \quad (9)$$

The outputs generated by the *Gaussian Heads* define the final 3D scene representation, which can be used for downstream applications.

3.3. Temporal Propagation

As we show in Fig. 2 and Fig. 3, the previous Gaussian features \mathcal{G}'_f are derived by first correcting the previous final Gaussian means using ego-motion and estimated temporal flow to the current frame, computing the positional encoding and adding them to the previous final features. The ITA module then fuses the temporal information in a recurrent manner. The model always operates on inputs from a single time step while maintaining a memory of the last prediction to facilitate effective temporal fusion.

3.4. Gaussian Splatting Supervision

Representing the scene with 3D Gaussians allows us to leverage the highly efficient Gaussian Splatting (GS) [25] method during training. For a detailed explanation of GS, we refer to the original work [25]. Essentially, the estimated Gaussian properties are used to project the 3D Gaussians back onto the input image views via the corresponding camera parameters. This enables us to rasterize predicted 2D semantic logits and depth values. Notably, unlike standard GS, which rasterizes RGB values, we rasterize semantic logits instead. We generate 2D semantic and depth labels using pretrained models to serve as ground truth. The entire model is then trained using 2D rendering losses. This enables training in a weakly supervised manner directly in the 2D image space, eliminating the need for 3D labels or additional sensor data, such as LiDAR, entirely. We optimize the MSE-Loss \mathcal{L}_{depth} between rendered depth \hat{D} and pre-computed depth maps D and the binary cross-entropy loss \mathcal{L}_{seg} between rendered logits \hat{S} and pre-computed semantic segmentation maps S :

$$\mathcal{L} = \mathcal{L}_{depth}(\hat{D}, D) + \mathcal{L}_{seg}(\hat{S}, S). \quad (10)$$

3.5. Temporal Gaussian Splatting

Cameras in autonomous driving scenarios usually share a relatively small frustum overlap, which complicates learning of correct depth and scene geometry. To address this limitation and increase viewpoint overlap between training cameras, we incorporate *Temporal Gaussian Splatting*, a technique commonly employed in recent works [5, 15, 41]. During training, for a specified temporal horizon T , we load the camera parameters and 2D ground truth labels of the T previous and T subsequent frames of the sequence. We then use GS to rasterize our estimated 3D Gaussians into

predicted depth and semantic maps for all temporal cameras and compute the same 2D rendering losses as for the current frame. This strategy effectively widens the supervisory signal by incorporating data from neighboring temporal views. While naïvely applying *Temporal Gaussian Splatting* can already significantly improve prediction performance, it overlooks the fact that many objects in the scene are in motion. This motion introduces inconsistencies in supervisory signals across frames, as objects predicted in the current frame might have moved in previous/subsequent frames, so the 2D losses can induce erroneous signals.

3.6. Temporal Module and Gaussian Flow

To address temporal inconsistencies caused by the motion of dynamic objects during training, we introduce a *Temporal Module* that learns to correct object motion. This module reduces discrepancies across temporally adjacent frames by estimating 3D offsets for each Gaussian to align them with all frames within a defined temporal horizon.

The *Temporal Module* takes the Gaussian features output by the *Gaussian Transformer* as input and computes a 3D motion offset for each Gaussian relative to each temporal frame in the horizon. To achieve this, we first define a set of learnable D -dimensional *time tokens* $\Psi \in \mathbb{R}^{2T \times D}$, one token for each frame in the temporal horizon $([-T, T])$. For a given target time step, we create a copy of the Gaussians and append the corresponding time token to their features. These temporally encoded features are then passed through an MLP to estimate a 3D motion offset vector $\vec{v}(t) \in \mathbb{R}^3$:

$$\vec{v}(t) = \text{MLP}_v(\mathcal{G}_f \oplus \Psi(t)). \quad (11)$$

The estimated offsets are added to the copied Gaussian positions, effectively translating the Gaussians to their updated locations in the target time step. GS is subsequently applied in the usual manner, using the translated Gaussians for the corresponding temporal frame. The *Temporal Module* is trained jointly with the rest of the model. Importantly, it does not require additional losses or ground truth motion data. Dynamic objects are correctly rendered in temporally adjacent frames only when their motion is accurately estimated, allowing the module to learn implicitly through the existing rendering losses. The benefits of this dynamic object handling are evaluated in Sec. 4.4.3, demonstrating its impact on scene consistency and model performance.

3.7. Voxelization

The final output of our model consists solely of 3D Gaussian distributions (see Fig. 1), with no voxelization or dense operations used during the model’s architecture or training process. However, these Gaussians can be efficiently voxelized as a post-processing step to facilitate direct comparison with prior work on occupancy estimation benchmarks.

For each voxel center with position $p_i = (x, y, z)$ on a defined voxel grid, the estimated 3D Gaussian distributions are queried and their opacity and semantic logits are accumulated at each point. Each voxel is assigned a label corresponding to the highest accumulated semantic logits. Voxels with an accumulated opacity below a defined threshold are classified as free. Further details about the voxelization are provided in the supplementary material.

4. Experiments

To evaluate the performance of GaussianFlowOcc, we compare our approach with recent state-of-the-art methods on occupancy estimation. In addition, an extensive ablation study is performed to study the impact of each module.

4.1. Dataset and Metrics

We conduct experiments on the Occ3D-nuScenes dataset [53], a widely used benchmark for occupancy estimation based on the nuScenes dataset [6]. The dataset defines a voxel grid spanning $[-40m, 40m]$ along the X and Y axes and $[-1m, 5.4m]$ along the Z axis, with a voxel size of $0.4m$. Each voxel is assigned one of 17 semantic labels from nuScenes Lidarseg [12] or classified as free space. Performance on this benchmark is evaluated using the Intersection over Union (IoU) score for each of the 17 semantic categories, aggregated to the mean IoU (mIoU). Additionally, following recent trends in the literature, we also report the RayIoU metric [31], a ray-based metric focused on rewarding good 3D completion while penalizing thick surface predictions. Finally, we provide information on inference speed (FPS) on a single A100 GPU.

4.2. Implementation Details

We utilize a ResNet-50 [17] backbone for feature extraction with an image resolution of 256×704 and set the number of Gaussians to $N = 10000$. The *Gaussian Transformer* consists of three blocks, each with induced attention layers employing $M = 500$ inducing points. The *Gaussian Heads* are implemented as a single linear layer, while all layers in the model share a latent dimension of 256. The *Temporal Module* consists of three linear layers and operates with a temporal horizon $T = 6$ during training. The model is trained for 18 epochs using four A100 GPUs. The source code is available at <https://github.com/boschresearch/GaussianFlowOcc>.

4.3. 3D Occupancy Prediction Results

We evaluate our proposed method against state-of-the-art approaches for weakly supervised 3D occupancy prediction on the Occ3D-nuScenes dataset [53], with the results summarized in Tab. 1. Each of the compared methods relies on a pretrained semantic segmentation model to generate 2D semantic labels for training. Our approach signifi-

Table 1. **Occupancy estimation performance on the Occ3D-nuScenes validation dataset.** The *Training Labels* column specifies the method used for pseudo label generation. Performance is reported in terms of IoU (%). The best-performing result in each column is highlighted in **bold**, while the second-best is shown in *italic*. A dash (-) indicates that no results were provided by the original work.

Method	Backbone	Training Labels	mIoU	IoU	RayIoU	RayIoU@1	RayIoU@2	RayIoU@4	FPS
SelfOcc [20]	ResNet-50	OpenSeeD	10.54	45.01	-	-	-	-	1.15
OccNeRF [70]	ResNet-101	GroundedSAM	10.81	22.81	-	-	-	-	1.27
GaussianOcc [15]	ResNet-101	GroundedSAM	11.26	-	11.85	8.69	11.90	14.95	5.57
GaussTR [24]	2× ViT	CLIP + Metric3D + GroundedSAM	<i>13.26</i>	<i>45.19</i>	-	-	-	-	0.20
GaussianFlowOcc (Ours)	ResNet-50	GroundedSAM + Metric3D	17.08	46.91	16.47	11.81	16.58	20.98	10.2

cantly outperforms all existing models, achieving at least a 51% mIoU improvement over methods that rely on RGB reconstruction and pseudo semantic labels. Notably, we surpass the recent GaussTR, which also uses Metric3D for depth supervision, by 29% mIoU, improving from 13.26 to 17.08. Furthermore, our model achieves this superior performance with a rather lightweight ResNet-50 backbone, while GaussTR uses CLIP and Metric3D as backbones during inference, which is significantly more costly. A similar trend is observed in the RayIoU metric, where our method improves performance by approximately 39% compared to GaussianOcc. Beyond accuracy, our model also offers a substantial efficiency advantage. Due to its lightweight architecture, it achieves inference speeds up to almost 9× faster than previous voxel-based methods. As mentioned above, especially GaussTR remains computationally expensive due to its reliance on multiple ViT-L backbones; GaussianFlowOcc is 50× faster.

We attribute these performance gains primarily to two key innovations: Firstly, our *Temporal Module* significantly improves temporal consistency, yielding large performance increases as we show in Sec. 4.4.3. Secondly, the *Gaussian Transformer* allows for a much higher number of Gaussians to be processed efficiently, due to approximately linearly scaling induced attention modules. A detailed per-class performance breakdown is provided in the supplement. Additionally, qualitative results presented in Fig. 1 and in the supplement highlight our model’s strong 3D scene completion abilities. In particular, our approach excels at representing small, thin, and flat objects (e.g., traffic signs, poles, pedestrians), a major advantage of using 3D Gaussians. In contrast, voxel-based methods are inherently constrained by a grid resolution, limiting their ability to capture fine details.

4.4. Ablation Study

4.4.1. Induced Attention vs. Full Attention

To verify the effectiveness of the proposed induced attention mechanism, we compare it to a standard full attention mechanism. The results are shown in Fig. 5. The plot shows both mIoU score and memory consumption for a single sample during training for induced attention (with fixed $M = 500$) and standard attention. As can be seen, with lower num-




Figure 5. **Comparison in mIoU and GPU memory consumption between Induced and Full Attention.**

bers of Gaussians, the induced attention and full attention exhibit comparable performance and memory usage, as the quadratic complexity of the full attention is not yet significant. Increasing the number of Gaussians to $N = 5000$ already requires double the amount of memory for full attention, while not providing any significant performance increase. Using 10000 Gaussians demands 50GB of GPU memory, forcing us to reduce the batch size to 1 on our hardware. Training a model this way did not lead to convergence. In contrast, using induced attention, the memory requirement scales linearly, ultimately enabling better performance. The increase in mIoU achieved by the Induced Attention model when increasing the number of Gaussians from 5000 to 10000 indicates that a higher number of Gaussians is necessary to represent the complexity of driving scenes. We find that increasing N beyond 10000 did not improve performance.

4.4.2. Self-Attention and Temporal Attention

We also show the necessity of the self-attention and temporal attention mechanisms by training a model without each of these modules. The results are displayed in Tab. 2, indicating that both modules increase performance significantly. Note that GICA is always used, as no scene information would be used otherwise.

4.4.3. Temporal Module

We demonstrate the effectiveness of the proposed *Temporal Module* by training a model without it while keeping all other components unchanged. As visible in Tab. 3, using the *Temporal Module* leads to a relative improvement in

Table 2. **Ablation on self and temporal attention.** Performance changes when omitting self and/or temporal attention within the *Gaussian Transformer*.

Induced Self-Attention	Induced Temporal Attention	mIoU
		13.81
✓		14.60
	✓	14.47
✓	✓	17.08

Table 5. **Ablation on different horizons T .** The best performance is achieved with $T = 6$. The model diverges with $T \geq 10$.

Time Horizon	0	2	4	6	8	10
mIoU	12.67	14.17	14.71	17.08	14.82	N/A

mIoU of 20% and in RayIoU of 14%, which emphasizes the importance of object motion compensation during training. A demonstration of rasterized views of temporally adjacent frames with and without applied estimated flow is provided in the supplement and in Fig. 1.

4.4.4. Gaussian Parameters

In contrast to previous works [15, 49], we find that allowing the model to freely adjust all Gaussian parameters yields the best results. To demonstrate this, we conduct an experiment (Tab. 4) where we disable the *Gaussian Heads* for certain properties and fix their values, preventing them from being learned. We set the opacity to 1, the scale to 0.3 and the rotations to unit quaternions. When only estimating the mean and fixing all other properties, only a performance of 9.48 mIoU can be achieved, showing the importance of the other properties. When allowing the model to influence the Gaussians opacity, a large improvement to 13.12 mIoU is reached, outperforming most previous methods. Adding the other properties further increases performance, with the maximum score achieved when using all properties. As mentioned above, the positive effect of using rotation and scale is especially noticeable for thin or flat structures.

4.4.5. Time Horizon

We ablate the *Temporal Gaussian Splatting* horizon T in the range of $\{0, 2, \dots, 10\}$ to demonstrate how temporal rendering leads to better supervision, shown in Tab. 5. With horizons greater than 8, the model diverges, which is likely due to large discrepancy between the time steps, making it increasingly harder to estimate (especially non-linear) motion. Additionally, some camera poses are just too distant from the current time step.

4.4.6. Pseudo Depth Labels

To analyze the impact of pseudo depth labels generated by Metric3D, we train a model using only pseudo semantics. The results of this experiment are presented in Tab. 6. Even

Table 3. **Ablation study of the proposed *Temporal Module*.** Performance comparison when compensating for object motion during training versus ignoring temporal inconsistencies during *Temporal Gaussian Splatting*.

Temporal Module	mIoU	RayIoU
	14.18	14.46
✓	17.08	16.47

Table 4. **Ablation on the Gaussian parameters.** Performance values when only using and learning a subset of the Gaussian parameters.

Opacity	Scale	Rotation	mIoU
			9.48
✓			13.12
✓	✓		14.98
✓	✓	✓	17.08

Table 6. **Ablation on pseudo depth labels.** Depth labels substantially improve performance.

Pseudo Depth	mIoU	IoU	RayIoU
	14.07	34.12	12.75
✓	17.08	46.91	16.47

without pseudo depth supervision, our model outperforms all prior works, demonstrating its strong capability. Notably, other approaches [20, 70] incorporate additional photometric losses to enhance geometric consistency, which we did not find to be beneficial. Interestingly, GaussTR [24] fails to converge without depth labels, further highlighting the robustness of our approach. This suggests that *Temporal Gaussian Splatting* alone provides sufficient 3D cues for geometry estimation by leveraging semantic maps from multiple cameras over an extended time horizon.

5. Conclusion and Future Work

In this work, we introduced a novel, linearly scaling Transformer architecture for efficiently modeling driving scenes using 3D Gaussian distributions, trained with *Temporal Gaussian Splatting* and explicit object motion modeling. Our approach significantly outperforms previous weakly supervised methods on the Occ3D-nuScenes benchmark while maintaining substantially faster inference speeds. Despite these advances, there are still areas for improvement. We found that self-supervised learning (SSL) with RGB rendering did not provide benefits, suggesting that further investigation is needed in this direction. Moreover, a more sophisticated approach to dynamic object modeling could further improve temporal consistency. Incorporating physically informed motion priors, leveraging temporal and Gaussian dependencies as well as adapting more properties than just the mean to better represent deformable objects could enhance temporal modeling. For future work, the sparse and efficient attention mechanism proposed in this work could be applied to other downstream tasks, such as object detection, lane segmentation, and planning. This opens up exciting possibilities for using our model in autonomous driving and robotics applications.

References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 3
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. 3
- [3] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9297–9307, 2019. 2
- [4] Simon Boeder, Fabian Gigengack, and Benjamin Risse. Langocc: Self-supervised open vocabulary occupancy estimation via volume rendering. *arXiv preprint arXiv:2407.17310*, 2024. 2, 6
- [5] Simon Boeder, Fabian Gigengack, and Benjamin Risse. Occlownet: Towards self-supervised occupancy estimation via differentiable rendering and occupancy flow. *arXiv preprint arXiv:2402.12792*, 2024. 2, 5
- [6] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liang, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 6
- [7] Anh-Quan Cao and Raoul de Charette. Monoscene: Monocular 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3991–4001, 2022. 2
- [8] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9650–9660, 2021. 2
- [9] Di Chang, Aljaž Božič, Tong Zhang, Qingsong Yan, Yingcong Chen, Sabine Süsstrunk, and Matthias Nießner. Rmvsnet: unsupervised multi-view stereo with neural rendering. In *European Conference on Computer Vision*, pages 665–680. Springer, 2022. 3
- [10] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14124–14133, 2021. 3
- [11] Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. In *European Conference on Computer Vision*, pages 370–386. Springer, 2025. 3
- [12] Whye Kit Fong, Rohit Mohan, Juana Valeria Hurtado, Lubing Zhou, Holger Caesar, Oscar Beijbom, and Abhinav Valada. Panoptic nusenes: A large-scale benchmark for lidar panoptic segmentation and tracking. *arXiv preprint arXiv:2109.03805*, 2021. 6
- [13] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12479–12488, 2023. 3
- [14] Wanshui Gan, Ningkai Mo, Hongbin Xu, and Naoto Yokoya. A simple attempt for 3d occupancy estimation in autonomous driving. *arXiv preprint arXiv:2303.10076*, 2023. 2
- [15] Wanshui Gan, Fang Liu, Hongbin Xu, Ningkai Mo, and Naoto Yokoya. Gaussianocc: Fully self-supervised and efficient 3d occupancy estimation with gaussian splatting. *arXiv preprint arXiv:2408.11447*, 2024. 2, 3, 5, 7, 8, 1, 6
- [16] Adrian Hayler, Felix Wimbauer, Dominik Muhle, Christian Rupperecht, and Daniel Cremers. S4c: Self-supervised semantic scene completion with neural fields. In *2024 International Conference on 3D Vision (3DV)*, pages 409–420. IEEE, 2024. 2
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3, 6
- [18] Junjie Huang and Guan Huang. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *arXiv preprint arXiv:2203.17054*, 2022. 2
- [19] Junjie Huang, Guan Huang, Zheng Zhu, Yun Ye, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021. 2
- [20] Yuanhui Huang, Wenzhao Zheng, Borui Zhang, Jie Zhou, and Jiwen Lu. Selfocc: Self-supervised vision-based 3d occupancy prediction. *arXiv preprint arXiv:2311.12754*, 2023. 1, 2, 3, 7, 8, 6
- [21] Yuanhui Huang, Wenzhao Zheng, Yunpeng Zhang, Jie Zhou, and Jiwen Lu. Tri-perspective view for vision-based 3d semantic occupancy prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9223–9232, 2023. 2
- [22] Yuanhui Huang, Wenzhao Zheng, Yunpeng Zhang, Jie Zhou, and Jiwen Lu. Gaussianformer: Scene as gaussians for vision-based 3d semantic occupancy prediction. *arXiv preprint arXiv:2405.17429*, 2024. 2, 5
- [23] Haoyi Jiang, Tianheng Cheng, Naiyu Gao, Haoyang Zhang, Wenyu Liu, and Xinggang Wang. Symphonize 3d semantic scene completion with contextual instance queries. *arXiv preprint arXiv:2306.15670*, 2023. 2
- [24] Haoyi Jiang, Liu Liu, Tianheng Cheng, Xinjie Wang, Tianwei Lin, Zhizhong Su, Wenyu Liu, and Xinggang Wang. Gausstr: Foundation model-aligned gaussian transformer for self-supervised 3d spatial understanding. *arXiv preprint arXiv:2412.13193*, 2024. 3, 5, 7, 8, 2
- [25] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 2, 3, 5

- [26] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lrf: Language embedded radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19729–19739, 2023. 3
- [27] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosior, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International conference on machine learning*, pages 3744–3753. PMLR, 2019. 4
- [28] Yiming Li, Zhiding Yu, Christopher Choy, Chaowei Xiao, Jose M Alvarez, Sanja Fidler, Chen Feng, and Anima Anandkumar. Voxformer: Sparse voxel transformer for camera-based 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9087–9098, 2023. 2
- [29] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European conference on computer vision*, pages 1–18. Springer, 2022. 2, 3, 4, 1
- [30] Zhiqi Li, Zhiding Yu, David Austin, Mingsheng Fang, Shiyi Lan, Jan Kautz, and Jose M Alvarez. Fb-occ: 3d occupancy prediction based on forward-backward view transformation. *arXiv preprint arXiv:2307.01492*, 2023. 2, 1, 3
- [31] Haisong Liu, Yang Chen, Haiguang Wang, Zetong Yang, Tianyu Li, Jia Zeng, Li Chen, Hongyang Li, and Limin Wang. Fully sparse 3d occupancy prediction, 2024. 2, 6
- [32] Tianqi Liu, Guangcong Wang, Shoukang Hu, Liao Shen, Xinyi Ye, Yuhang Zang, Zhiguo Cao, Wei Li, and Ziwei Liu. Mvsgaussian: Fast generalizable gaussian splatting reconstruction from multi-view stereo. In *European Conference on Computer Vision*, pages 37–53. Springer, 2025. 3
- [33] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petr: Position embedding transformation for multi-view 3d object detection. In *European Conference on Computer Vision*, pages 531–548. Springer, 2022. 3
- [34] Hao Lu, Tianshuo Xu, Wenzhao Zheng, Yunpeng Zhang, Wei Zhan, Dalong Du, Masayoshi Tomizuka, Kurt Keutzer, and Yingcong Chen. Drivingrecon: Large 4d gaussian reconstruction model for autonomous driving. *arXiv preprint arXiv:2412.09043*, 2024. 3
- [35] Yuhang Lu, Xinge Zhu, Tai Wang, and Yuexin Ma. Octreeocc: Efficient and multi-granularity occupancy prediction using octree queries. *arXiv preprint arXiv:2312.03774*, 2023. 2
- [36] Junyi Ma, Xieyuanli Chen, Jiawei Huang, Jingyi Xu, Zhen Luo, Jintao Xu, Weihao Gu, Rui Ai, and Hesheng Wang. Cam4docc: Benchmark for camera-only 4d occupancy forecasting in autonomous driving applications. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21486–21495, 2024. 2
- [37] Qihang Ma, Xin Tan, Yanyun Qu, Lizhuang Ma, Zhizhong Zhang, and Yuan Xie. Cotr: Compact occupancy transformer for vision-based 3d occupancy prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19936–19945, 2024. 2
- [38] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2
- [39] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 3
- [40] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2
- [41] Mingjie Pan, Jiaming Liu, Renrui Zhang, Peixiang Huang, Xiaoqi Li, Li Liu, and Shanghang Zhang. Renderocc: Vision-centric 3d occupancy prediction with 2d rendering supervision. *arXiv preprint arXiv:2309.09502*, 2023. 2, 5, 6
- [42] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. 3
- [43] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d language gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20051–20060, 2024. 3
- [44] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [45] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, et al. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*, 2024. 2, 1
- [46] Yining Shi, Kun Jiang, Jiusi Li, Zelin Qian, Junze Wen, Mengmeng Yang, Ke Wang, and Diange Yang. Grid-centric traffic scenario perception for autonomous driving: A comprehensive review. *IEEE Transactions on Neural Networks and Learning Systems*, 2024. 1
- [47] Yiang Shi, Tianheng Cheng, Qian Zhang, Wenyu Liu, and Xinggang Wang. Occupancy as set of points. In *European Conference on Computer Vision*, pages 72–87. Springer, 2025. 2
- [48] Sophia Sirko-Galouchenko, Alexandre Boulch, Spyros Gidaris, Andrei Bursuc, Antonin Vobecky, Patrick Pérez, and Renaud Marlet. Occfeat: Self-supervised occupancy feature prediction for pretraining bev segmentation networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4493–4503, 2024. 2
- [49] Qianpu Sun, Changyong Shu, Sifan Zhou, Zichen Yu, Yan Chen, Dawei Yang, and Yuan Chun. Gsrender: Deduplicated occupancy prediction via weakly supervised 3d gaussian

- sian splatting. *arXiv preprint arXiv:2412.14579*, 2024. 2, 3, 8
- [50] Xin Tan, Wenbin Wu, Zhiwei Zhang, Chaojie Fan, Yong Peng, Zhizhong Zhang, Yuan Xie, and Lizhuang Ma. Geocc: Geometrically enhanced 3d occupancy network with implicit-explicit depth fusion and contextual self-supervision. *arXiv preprint arXiv:2405.10591*, 2024. 2
- [51] Zhiyu Tan, Zichao Dong, Cheng Zhang, Weikun Zhang, Hang Ji, and Hao Li. Ovo: Open-vocabulary occupancy. *arXiv preprint arXiv:2305.16133*, 2023. 2
- [52] Pin Tang, Zhongdao Wang, Guoqing Wang, Jilai Zheng, Xianguan Ren, Bailan Feng, and Chao Ma. Sparseocc: Rethinking sparse latent representation for vision-based semantic occupancy prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15035–15044, 2024. 2
- [53] Xiaoyu Tian, Tao Jiang, Longfei Yun, Yue Wang, Yilun Wang, and Hang Zhao. Occ3d: A large-scale 3d occupancy prediction benchmark for autonomous driving. *arXiv preprint arXiv:2304.14365*, 2023. 2, 6, 1, 3
- [54] Wenwen Tong, Chonghao Sima, Tai Wang, Li Chen, Silei Wu, Hanming Deng, Yi Gu, Lewei Lu, Ping Luo, Dahua Lin, et al. Scene as occupancy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8406–8415, 2023. 2
- [55] Antonin Vobecky, Oriane Siméoni, David Hurych, Spyridon Gidaris, Andrei Bursuc, Patrick Pérez, and Josef Sivic. Pop-3d: Open-vocabulary 3d occupancy prediction from images. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [56] Jiabao Wang, Zhaojiang Liu, Qiang Meng, Liujiang Yan, Ke Wang, Jie Yang, Wei Liu, Qibin Hou, and Mingming Cheng. Opus: occupancy prediction using a sparse set. In *Advances in Neural Information Processing Systems*, 2024. 2
- [57] Shihao Wang, Yingfei Liu, Tiancai Wang, Ying Li, and Xiangyu Zhang. Exploring object-centric temporal modeling for efficient multi-view 3d object detection. *arXiv preprint arXiv:2303.11926*, 2023. 3
- [58] Xiaofeng Wang, Zheng Zhu, Wenbo Xu, Yunpeng Zhang, Yi Wei, Xu Chi, Yun Ye, Dalong Du, Jiwen Lu, and Xingang Wang. Openoccupancy: A large scale benchmark for surrounding semantic occupancy perception. *arXiv preprint arXiv:2303.03991*, 2023. 2
- [59] Yi Wei, Linqing Zhao, Wenzhao Zheng, Zheng Zhu, Yongming Rao, Guan Huang, Jiwen Lu, and Jie Zhou. Surround-depth: Entangling surrounding views for self-supervised multi-camera depth estimation. In *Conference on robot learning*, pages 539–549. PMLR, 2023. 2
- [60] Felix Wimbauer, Nan Yang, Christian Rupprecht, and Daniel Cremers. Behind the scenes: Density fields for single view reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9076–9086, 2023. 3
- [61] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xingang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20310–20320, 2024. 3
- [62] Huaiyuan Xu, Junliang Chen, Shiyu Meng, Yi Wang, and Lap-Pui Chau. A survey on occupancy perception for autonomous driving: The information fusion perspective. *Information Fusion*, 114:102671, 2025. 1
- [63] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5438–5448, 2022. 3
- [64] Yunzhi Yan, Haotong Lin, Chenxu Zhou, Weijie Wang, Haiyang Sun, Kun Zhan, Xianpeng Lang, Xiaowei Zhou, and Sida Peng. Street gaussians: Modeling dynamic urban scenes with gaussian splatting. In *European Conference on Computer Vision*, pages 156–173. Springer, 2024. 3
- [65] Jianglong Ye, Naiyan Wang, and Xiaolong Wang. Feature-nerf: Learning generalizable nerfs by distilling foundation models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8962–8973, 2023. 3
- [66] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3d: Towards zero-shot metric 3d prediction from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9043–9053, 2023. 2, 1
- [67] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. 3
- [68] Zichen Yu, Changyong Shu, Jiajun Deng, Kangjie Lu, Zongdai Liu, Jiangyong Yu, Dawei Yang, Hui Li, and Yan Chen. Flashocc: Fast and memory-efficient occupancy prediction via channel-to-height plugin. *arXiv preprint arXiv:2311.12058*, 2023. 2
- [69] Zhu Yu, Bowen Pang, Lizhe Liu, Runmin Zhang, Qihao Peng, Maochun Luo, Sheng Yang, Mingxia Chen, Si-Yuan Cao, and Hui-Liang Shen. Language driven occupancy prediction. *arXiv preprint arXiv:2411.16072*, 2024. 2
- [70] Chubin Zhang, Juncheng Yan, Yi Wei, Jiaxin Li, Li Liu, Yansong Tang, Yueqi Duan, and Jiwen Lu. Occnerf: Self-supervised multi-camera occupancy prediction with neural radiance fields. *arXiv preprint arXiv:2312.09243*, 2023. 1, 2, 3, 7, 8, 6
- [71] Hao Zhang, Feng Li, Xueyan Zou, Shilong Liu, Chunyuan Li, Jianwei Yang, and Lei Zhang. A simple framework for open-vocabulary segmentation and detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1020–1031, 2023. 2
- [72] Yunpeng Zhang, Zheng Zhu, and Dalong Du. Occformer: Dual-path transformer for vision-based 3d semantic occupancy prediction. *arXiv preprint arXiv:2304.05316*, 2023. 2, 1, 3
- [73] Yanan Zhang, Jinqing Zhang, Zengran Wang, Junhao Xu, and Di Huang. Vision-based 3d occupancy prediction in autonomous driving: a review and outlook. *arXiv preprint arXiv:2405.02595*, 2024. 1
- [74] Linqing Zhao, Xiuwei Xu, Ziwei Wang, Yunpeng Zhang, Borui Zhang, Wenzhao Zheng, Dalong Du, Jie Zhou, and

- Jiwen Lu. Lowrankocc: Tensor decomposition and low-rank recovery for vision-based 3d semantic occupancy prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9806–9815, 2024. [2](#)
- [75] Jilai Zheng, Pin Tang, Zhongdao Wang, Guoqing Wang, Xi-angxuan Ren, Bailan Feng, and Chao Ma. Veon: Vocabulary-enhanced occupancy prediction. In *European Conference on Computer Vision*, pages 92–108. Springer, 2025. [2](#)
- [76] Shunyuan Zheng, Boyao Zhou, Ruizhi Shao, Boning Liu, Shengping Zhang, Liqiang Nie, and Yebin Liu. Gps-gaussian: Generalizable pixel-wise 3d gaussian splatting for real-time human novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19680–19690, 2024. [3](#)
- [77] Wenzhao Zheng, Junjie Wu, Yao Zheng, Sicheng Zuo, Zixun Xie, Longchao Yang, Yong Pan, Zhihui Hao, Peng Jia, Xi-angepeng Lang, et al. Gaussianad: Gaussian-centric end-to-end autonomous driving. *arXiv preprint arXiv:2412.10371*, 2024. [3](#)
- [78] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *European Conference on Computer Vision*, pages 696–712. Springer, 2022. [2](#)
- [79] Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejia Xu, Pradyumna Chari, Suyu You, Zhangyang Wang, and Achuta Kadambi. Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21676–21685, 2024. [3](#)
- [80] Xiaoyu Zhou, Zhiwei Lin, Xiaojun Shan, Yongtao Wang, Deqing Sun, and Ming-Hsuan Yang. Drivinggaussian: Composite gaussian splatting for surrounding dynamic autonomous driving scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21634–21643, 2024. [3](#)
- [81] Sicheng Zuo, Wenzhao Zheng, Yuanhui Huang, Jie Zhou, and Jiwen Lu. Gaussianworld: Gaussian world model for streaming 3d occupancy prediction. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6772–6781, 2025. [2](#)