

## Test-Time 3D Occupancy Prediction

Fengyi Zhang<sup>1</sup> Xiangyu Sun<sup>1</sup> Huitong Yang<sup>1</sup> Zheng Zhang<sup>2</sup> Zi Huang<sup>1</sup> Yadan Luo<sup>1\*</sup>  
<sup>1</sup>UQMM Lab, The University of Queensland <sup>2</sup>Harbin Institute of Technology

### Abstract

*Self-supervised 3D occupancy prediction offers a promising solution for understanding complex driving scenes without requiring costly 3D annotations. However, training dense occupancy decoders to capture fine-grained geometry and semantics can demand hundreds of GPU hours, and once trained, such models struggle to adapt to varying voxel resolutions or novel object categories without extensive retraining. To overcome these limitations, we propose a practical and flexible test-time occupancy prediction framework termed TT-Occ. Our method incrementally constructs, optimizes, and voxelizes time-aware 3D Gaussians from raw sensor streams by integrating vision foundation models (VFMs) at runtime. The flexible representation of 3D Gaussians enables voxelization at arbitrary user-specified resolutions, while the strong generalization capability of VFMs supports accurate perception and open-vocabulary recognition without requiring any network training or fine-tuning. To validate the generality and effectiveness of our framework, we present two variants: a LiDAR-based version and a vision-centric version, and conduct extensive experiments on the Occ3D-nuScenes and nuCraft benchmarks under varying voxel resolutions. Experimental results show that TT-Occ significantly outperforms existing computationally expensive pretrained self-supervised counterparts. Code is available at <https://github.com/Xian-Bei/TT-Occ>.*

### 1. Introduction

3D Occupancy prediction seeks to accurately identify regions within an environment that are occupied by objects of particular classes and those that remain free. This capability is crucial to enable collision-free trajectory planning and reliable navigation in autonomous driving systems [6, 7, 44, 49] and embodied agents [18, 31, 38, 45]. Existing solutions [4, 12, 14, 27, 35, 43, 52], however, heavily rely on dense 3D annotations obtained through labor-intensive manual labeling of dynamic driving scenes spanning up to 80 meters per frame. To mitigate this cost, recent studies have resorted to *self-supervised* alternatives [3, 10, 11, 13, 16, 39, 50, 53, 57].

\*Corresponding Author.




Figure 1. Comparison of self-supervised occupancy prediction methods in terms of pretraining time (x-axis), mIoU (y-axis), and runtime FPS (marker radius). TT-Occ achieves strong mIoU and competitive FPS without any pretraining.

While reducing labeling costs, they still require substantial computational overhead. For instance, training SelfOcc [13] on Occ3D-nuScenes [35] at a voxel resolution of 0.4m requires approximately 2 days on eight GPUs. Furthermore, once trained, adapting to finer resolution (e.g., 0.2m of nuCraft [58] dataset) or novel object classes (e.g., beyond the 17 predefined classes of Occ3D-nuScenes [35]) may necessitate extensive retraining.

Recent vision foundation models (VFMs), however, change this landscape. Models such as VGGT [37] and MapAnything [19] provide reliable multiview geometry, and REX-Omni [17] enables open-vocabulary semantic reasoning, both directly at test time. As these cues become increasingly accessible without task-specific training, it is natural to revisit a long-standing assumption of the field: *if such information no longer needs to be learned by a network, is training an occupancy model still necessary?* To answer this question, we explore a test-time occupancy estimation method termed **TT-Occ**, which progressively constructs, optimizes and voxelizes time-aware 3D Gaussians from raw sensor streams by integrating VFMs. We introduce two variants, **TT-OccCamera** and **TT-OccLiDAR**, which differ in

the sensor modality used to initialize the Gaussian primitives, respectively. Both variants eliminate costly pretraining and allow flexible adaptation to arbitrary voxel resolutions and user-specified semantic queries. Unlike previous NeRF [47, 48] and 3DGS-based [8, 56] reconstruction methods that perform offline per-scene modeling assisted by *external GT priors* (e.g., HD maps or bounding boxes), TT-Occ generates occupancy representations in an online fashion, relying solely on raw sensor streams and generally trained VFMs to instantiate Gaussians capturing object geometry and semantics in unbounded outdoor scenes.

Our framework follows a simple yet effective “*lift-track-voxelize*” symphony: (1) Lift: at each test time step, we first “*lift*” geometry and semantic information of surrounding views extracted via VFMs into time-aware 3D Gaussians on the fly. The generated Gaussians can also be splatted back onto the image plane through differentiable rasterization for parameter optimization [20]. (2) Track: next, we “*track*” dynamic Gaussians and accumulate static ones using estimated motion flow. This motion compensates for partial object visibility and prevents trailing artifacts while maintaining long-term temporal coherence. (3) Voxelize: at any given timestamp, the generated 3D Gaussians can be voxelized onto discrete occupancy grids with arbitrary user-specified resolutions. Optionally, to further mitigate the inherent noise in VFM predictions and tracking results, we introduce a Trilateral Radial Basis Function (TRBF), which jointly considers semantic, color, and spatial affinities to periodically smooth the Gaussian parameters.

Extensive experiments on Occ3D-nuScenes [35] and the recently released high-resolution nuCraft [58] demonstrate that TT-Occ achieves better performance than existing self-supervised counterparts, which typically require hundreds of GPU training hours. Qualitative analysis further highlights the superiority of TT-Occ in terms of temporal consistency and open-vocabulary generalization.

## 2. Related Work

**Occupancy Prediction.** Fully supervised occupancy methods predict voxel-level semantics using dense voxel grids [12, 21, 42], depth priors [15, 22], or sparse representations [14, 25, 33]. Despite their effectiveness, these approaches rely heavily on costly large-scale 3D annotations. To mitigate this, recent methods explore self-supervised occupancy learning. SelfOcc [13] leverages signed distance fields (SDF) and multi-view stereo embeddings to achieve temporally consistent occupancy from videos. OccNeRF [50] utilizes photometric consistency and 2D foundation model supervision for semantic occupancy estimation in unbounded scenes. In open-world scenarios, POP3D [36] jointly trains class-agnostic occupancy grids and open-vocabulary semantics using unlabeled paired LiDAR and images, but suffers from sparsity and semantic ambiguity due to low-resolution

CLIP [28] features. VEON [53] introduces a vocabulary-enhanced occupancy framework trained with LiDAR supervision, leveraging CLIP features for open-vocabulary prediction and addressing depth ambiguities via enhance depth model (MiDaS [29], ZoeDepth [2]). GaussianOcc [10] uses Gaussian Splatting [20] for cross-view optimization without pose annotations, while GaussianTR [16] aligns rendered Gaussian features with pre-trained foundation models, enabling open-vocabulary occupancy prediction without explicit annotations. Despite these advances, existing methods either rely on extensive offline training or struggle with open-vocabulary settings and fixed resolutions. In contrast, our approach overcomes these limitations by enabling occupancy prediction through temporally coherent, training-free Gaussian optimization at test time.

**3D Reconstruction of Driving Scenes.** Recent advances in dynamic scene modeling [19, 32, 37] have achieved impressive photorealism and multi-view consistency. OmniRe [8] performs real-time 3D reconstruction and simulation by building local canonical spaces for dynamic urban actors. Street Gaussians [46] separates moving vehicles from static backgrounds, enabling efficient and high-quality rendering. DrivingGaussian [56] incrementally reconstructs static scenes and dynamically integrates moving objects via Gaussian graphs for interactive editing. HUGS [55] jointly optimizes geometry, appearance, semantics, and motion to achieve real-time view synthesis and 3D semantic reconstruction without explicit bounding box annotations. Autoregressive world modeling methods [40, 54] predict future occupancy using previously estimated 3D occupancies, facilitating temporal reasoning in dynamic environments. Our test-time approach fundamentally *differs* from these methods by eliminating dependencies on external priors and annotations (e.g., HD maps and GT bounding boxes). Instead, we focus solely on raw sensor inputs, optimizing Gaussian representations independently at each frame to directly infer the accurate geometry, rather than reconstructing photorealistic scenes or predicting future occupancy.

## 3. Proposed Approach

**Task Formulation.** At each time step  $t$ , the objective of occupancy estimation is to infer the voxelized geometry and semantic labels of the current scene directly from raw sensor inputs. Formally, we define the voxel grid as  $\mathbf{O}^{(t)} \in \mathbb{C}^{\frac{X}{\delta} \times \frac{Y}{\delta} \times \frac{Z}{\delta}}$ , where  $X, Y, Z$  defines the spatial dimensions of the region of interest, and  $\delta$  is the voxel resolution (e.g., 0.2m). The input modality varies by variant. The input of the vision-centric variant is  $M$  surrounding-view camera images  $\mathcal{I}^{(t)} = \{\mathbf{I}_m^{(t)} \in \mathbb{R}^{3 \times H \times W}\}_{m=1}^M$ , while LiDAR-based variant additionally takes a LiDAR point cloud  $\mathcal{P}^{(t)} = \{\mathbf{p}_i^{(t)} \in \mathbb{R}^3\}_{i=1}^{N_t}$ . Each voxel in  $\mathbf{O}^{(t)}$  is assigned a semantic label from the set  $\mathbb{C} = \{0, 1, \dots, C\}$ , where 0 indi-

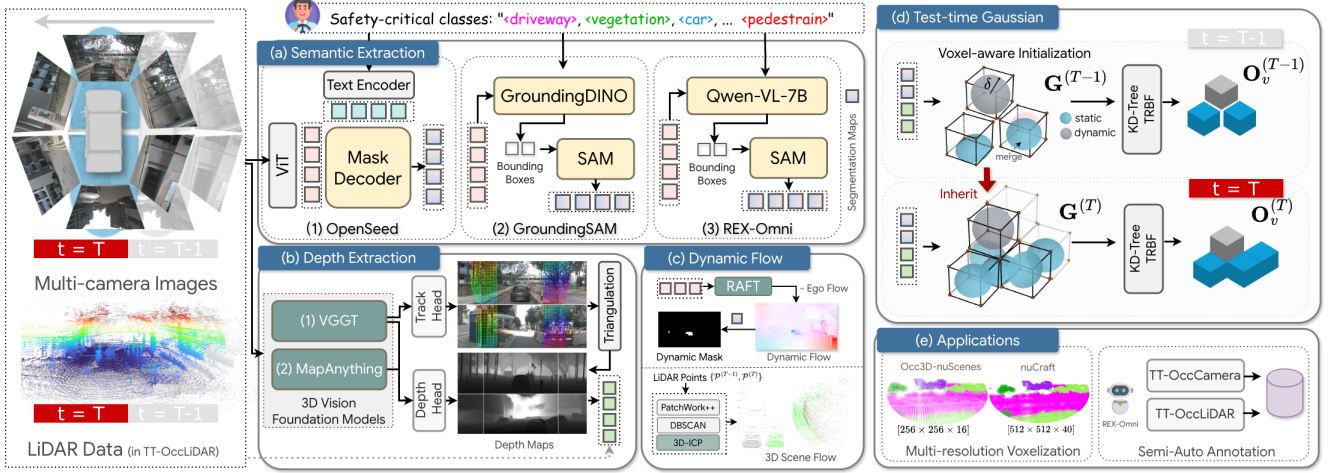


Figure 2. **Overview of the proposed TT-OccCamera and TT-OccLiDAR approaches.** TT-Occ performs test-time 3D occupancy estimation by directly integrating a suite of VFMs at runtime, avoiding any network training or fine-tuning. (a) Multi-view semantics are obtained using arbitrary open-vocabulary segmentation VFMs (e.g., OpenSeed, GroundingSAM, REX-Omni). (b) Geometry cues are extracted via depth/correspondence VFMs (e.g., VGGT, MapAnything). (c) Dynamic flow is estimated to track moving objects and prevent trailing artifacts while accumulating static structure over time. (d) The resulting features instantiate and refine time-aware 3D Gaussians, which can be voxelized at any resolutions for occupancy prediction. (e) TT-Occ supports both LiDAR-based and camera-only variants and enables multi-resolution voxelization and semi-automatic annotation. Despite leveraging multiple VFMs, TT-Occ remains highly efficient, delivering competitive occupancy performance across Occ3D-nuScenes and nuCraft.

ates an empty cell and labels 1 to  $C$  corresponds to distinct occupied categories.

### 3.1. Lift Geometry and Semantics into Gaussians

For each time step, a set of time-aware Gaussian blobs  $\mathcal{G}^{(t)} = \{\mathbf{G}_i^{(t)}\}_{i=1}^{K_t}$  are instantiated to represent the scene. Each Gaussian is parameterized by its mean position  $\boldsymbol{\mu}_i \in \mathbb{R}^3$ , opacity  $\alpha_i \in (0, 1)$ , color  $\mathbf{c}_i \in \mathbb{R}^3$ , semantic probability  $\mathbf{m}_i \in \mathbb{R}^C$ , time step  $t$ , and its spatial density is given by:

$$\mathbf{G}_i^{(t)}(\mathbf{x}) = \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_i)^\top \boldsymbol{\Sigma}_i^{-1}(\mathbf{x} - \boldsymbol{\mu}_i)\right), \quad (1)$$

where covariance matrix  $\boldsymbol{\Sigma}_i = R(\mathbf{q}_i) \text{diag}(s_i^2) R(\mathbf{q}_i)^\top$  is factorized by the orientation quaternion  $\mathbf{q}_i \in \mathbb{R}^4$  and the scale vector  $s_i \in \mathbb{R}_+^3$ . To project each Gaussian on the 2D plane, we apply perspective transformation  $\text{Proj}(\mathbf{x}; \mathbf{K}, \mathbf{E})$  with the intrinsic matrix  $\mathbf{K} \in \mathbb{R}^{3 \times 3}$  and extrinsic matrix  $\mathbf{E} \in \mathbb{R}^{3 \times 4}$ . The projected mean and covariance are:

$$\boldsymbol{\mu}_i^{2D} = \text{Proj}(\boldsymbol{\mu}_i)_{1:2}, \quad \boldsymbol{\Sigma}_i^{2D} = \mathbf{J}_{\text{Proj}}(\boldsymbol{\mu}_i) \boldsymbol{\Sigma}_i \mathbf{J}_{\text{Proj}}(\boldsymbol{\mu}_i)^\top_{1:2,1:2},$$

where  $\mathbf{J}_{\text{Proj}}$  is the Jacobian matrix. The color of the pixel  $u$  is then obtained by alpha blending.

**Modality-Specific Initialization.** In the LiDAR-based variant TT-OccLiDAR, the sparse LiDAR points are directly initialized as 3D Gaussians, inheriting the precise spatial positions from real-world measurements. In contrast, the vision-centric variant TT-OccCamera reconstructs a 3D

point cloud from 3D vision foundation models (3DVFM). Specifically, we employ a pretrained 3DVFMs such as VGGT [37] and MapAnything [19] to estimate dense depth maps from multi-view RGB inputs. When predicted depth maps suffer from inherent *scale ambiguity* due to the lack of metric supervision (e.g., in VGGT), we mitigate this issue by applying multi-view triangulation to keypoint correspondences in overlapping views. See supplementary materials for details. **VFM Semantics.** To incorporate semantic information, we extract semantic maps from  $M$  surrounding views by querying an open-vocabulary model [17, 30, 51]. Our framework is agnostic to the specific VFM used and various segmentation and grounding models were evaluated, as illustrated in Fig. 2. These semantic maps are then lifted to 3D via a visibility-weighted projection:

$$\mathbf{m}_i = \frac{1}{M} \sum_{m=1}^M \mathbb{I}_m(\boldsymbol{\mu}_i) \mathcal{M}_m(\text{Proj}(\boldsymbol{\mu}_i; \mathbf{K}_m, \mathbf{E}_m)), \quad (2)$$

where  $\mathbf{m}_i$  is a fused semantic probability vector for the 3D point. Specifically,  $\boldsymbol{\mu}_i \in \mathbb{R}^3$  denotes the 3D coordinate of the  $i$ -th point. For each camera view  $m \in \{1, \dots, M\}$ :  $\text{Proj}(\boldsymbol{\mu}_i; \mathbf{K}_m, \mathbf{E}_m)$  projects the 3D point  $\boldsymbol{\mu}_i$  onto the 2D image plane using the corresponding camera intrinsics  $\mathbf{K}_m$  and extrinsics  $\mathbf{E}_m$ ;  $\mathcal{M}_m(\cdot)$  returns the semantic probability vector from the open-vocabulary segmentation map at the projected 2D location;  $\mathbb{I}_m(\boldsymbol{\mu}_i)$  is a binary visibility indicator, equal to 1 if the point  $\boldsymbol{\mu}_i$  is visible in view  $m$ , and 0 otherwise. The use of foundation models enables compatibility




Figure 3. **Illustration of trailing artifacts caused by naïvely accumulating per-frame Gaussians.** Without handling dynamic regions, moving objects (*e.g.*, cars shown in blue voxels) leave behind smeared or duplicated structures, corrupting the occupancy field. Our tracking suppresses these artifacts by separating dynamic Gaussians from static ones.

with open-vocabulary semantic queries, allowing TT-Occ to flexibly adapt to user-specified class definitions at test time. For benchmark evaluation on nuScenes [5], we adopt the standard label space; however, our method inherently supports open-vocabulary settings without requiring retraining, in contrast to conventional self-supervised occupancy prediction approaches (*e.g.*, [13]) that depend on fixed decoder architectures and label sets.

**Voxel-aware Simplifications.** To accelerate Gaussian optimization and subsequent voxelization, we simplify the standard 3DGS [20] by initializing the scale parameters with  $\delta$  and constraining them using a sigmoid activation rather than an exponential function to prevent excessive growth. Additionally, we prune redundant Gaussians within the same voxel cell (size  $\delta$ ) while merging their semantic probabilities.

### 3.2. Track Dynamic Gaussians

Reconstructing a driving scene faithfully can be challenging due to fast-moving objects (*e.g.*, vehicles, pedestrians) that are often only partially observed. Without prior knowledge such as complete trajectories or bounding box annotations of moving instances used in [39, 56], optimizing 3D Gaussians online can often result in severe *trailing artifacts*. To address this, we propose to track dynamic Gaussians while accumulating static ones across adjacent frames. In the upper image of Fig. 3, fast-moving vehicles (blue voxels) produce noticeable trailing artifacts, while the lower image shows a clean reconstruction after the Gaussian tracking.

**Modality-Specific Tracking.** Both TT-OccCamera and TT-OccLiDAR share the same mechanism for *static* Gaussian inheritance, which enhances scene completeness by accumulating temporally consistent observations across frames. The key difference lies in how *dynamic* Gaussians are handled, particularly in their tracking strategy and purpose. For TT-OccLiDAR, we track the Gaussian motion with learning-free *Gaussian scene flow* estimation, which allows us to relocate dynamic Gaussians accordingly. Our pipeline consists of the following steps: associating instances by projecting LiDAR points onto segmentation masks; denoising with DBSCAN [9]; matching clusters across frames

based on spatial proximity and shape similarity; and finally estimating 3D flow using the iterative closest point (ICP) algorithm [1]. For TT-OccCamera, we estimate optical flow between adjacent frames of the same camera using RAFT [34], and compute ego-motion flow based on inter-frame camera poses and per-pixel depth predicted by 3DVFM (e.g., VGGT [37] or MapAnything [19]). Subtracting the ego-motion flow from the optical flow yields a residual dynamic flow, which reflects true object motion. Although this 2D dynamic flow could, in principle, guide the 3D motion of dynamic Gaussians, back-projecting it into 3D space tends to amplify noise from both RAFT and 3DVFM, resulting in unstable Gaussian motion. To mitigate this, we adopt a compromise strategy by thresholding the dynamic flow magnitude to obtain a dynamic mask that identifies likely moving regions. The corresponding 3D Gaussians projected onto these regions are treated as dynamic and excluded from static accumulation in the next frame. While this approach does not allow accumulation of dynamic objects as in the LiDAR-based variant, it effectively reduces artifacts caused by noisy motion cues and temporal inconsistencies. See supplementary materials for implementation details.

### 3.3. Gaussian Voxelization

Following 3DGS [20], our model refines Gaussian parameters at test time by minimizing a loss that enforces color consistency, with sky regions intentionally masked out as in [56]. Optionally, to further mitigate the inherent noise and errors in VFMs’ predictions and tracking results, we introduce a Trilateral Radial Basis Function (TRBF) kernel for periodic smoothing and denoising. TRBF kernel improves the spatial and temporal coherence of occupancy predictions by leveraging spatial, radiometric, and semantic affinities among Gaussians for anisotropic information propagation while preserving local structures and semantic boundaries. Formally, for each  $\mathbf{m}_i \in \mathbf{G}_i^{(t)}$ , the kernel smoothing is defined as a deformable convolution over its nearest neighbors:

$$\mathbf{m}_i \leftarrow \frac{1}{Z(i)} \sum_{j \in \text{NN}(i)} \mathbf{m}_j \cdot \mathcal{K}(i, j), \quad (3)$$

where  $\text{NN}(\cdot)$  identifies  $K$  nearest Gaussians using a KD-Tree for efficient search and  $Z(i)$  is a normalization factor  $Z(i) = \sum_{j \in \text{NN}(i)} \mathcal{K}(i, j)$  to ensure that  $\mathbf{m}_i$  sums to 1 as a valid probability. By the Schur Product Theorem, the trilateral kernel decomposes element-wise into spatial, radiometric, and semantic components:

$$\mathcal{K}(i, j) = \mathcal{K}_\mu(i, j) \cdot \mathcal{K}_c(i, j) \cdot \mathcal{K}_m(i, j), \quad (4)$$

where each term  $\text{attr} \in \{\mu, c, m\}$  is defined as the following format:

$$\mathcal{K}_{\text{attr}}(i, j) = \exp\left(-\frac{\|\text{attr}_i - \text{attr}_j\|^2}{2\sigma_{\text{attr}}^2}\right). \quad (5)$$

Table 1. Occupancy prediction performance on **Occ3D-nuScenes** [35]. Best results among self-supervised methods are highlighted in **bold**.

Method	Input	mIoU ↑	barr	bike	bus	car	c-veh	moto	ped	t-cone	trail	truck	d-surf	s-walk	terr	man	vege
BEVFormer <sub>(ECCV'22)</sub> [23]	C	26.88	37.83	17.87	40.44	42.43	7.36	23.88	21.81	20.98	22.38	30.70	55.35	36.0	28.06	20.04	17.69
CTF-Occ <sub>(NeurIPS'23)</sub> [35]	C	28.53	39.33	20.56	38.29	42.24	16.93	24.52	22.72	21.05	22.98	31.11	53.33	37.98	33.23	20.79	18.0
RenderOcc <sub>(ICRA'24)</sub> [27]	C	23.93	27.56	14.36	19.91	20.56	11.96	12.42	12.14	14.34	20.81	18.94	68.85	42.01	43.94	17.36	22.61
OccFlowNet <sub>(WACV'25)</sub> [4]	C&L	26.14	27.50	26.00	34.00	32.00	20.40	25.90	18.60	20.20	26.00	28.70	62.00	37.80	39.50	29.00	26.80
SimpleOcc <sub>(TIV'24)</sub> [11]	C	7.99	0.67	1.18	3.21	7.63	1.02	0.26	1.80	0.26	1.07	2.81	40.44	18.30	17.01	13.42	10.84
OccNeRF <sub>(Arxiv'24)</sub> [50]	C	10.81	0.83	0.82	5.13	12.49	3.50	0.23	3.10	1.84	0.52	3.90	52.62	20.81	24.75	18.45	13.19
DistillNeRF <sub>(NeurIPS'24)</sub> [39]	C	8.93	1.35	2.08	10.21	10.09	2.56	1.98	5.54	4.62	1.43	7.90	43.02	16.86	15.02	14.06	15.06
GaussianOcc <sub>(Arxiv'24)</sub> [10]	C	11.26	1.79	5.82	14.58	13.55	1.30	2.82	7.95	9.76	0.56	9.61	44.59	20.10	17.58	8.61	10.29
GaussianTR <sub>(CVPR'25)</sub> [16]	C	11.70	2.09	5.22	14.07	20.43	5.70	7.08	5.12	3.93	0.92	13.36	39.44	15.68	22.89	21.17	21.87
LangOcc <sub>(3DV'25)</sub> [3]	C	11.84	3.10	9.00	6.30	14.20	0.40	10.80	6.20	9.00	3.80	10.70	43.70	9.50	26.40	19.60	26.40
VEON-LiDAR <sub>(ECCV'24)</sub> [53]	C&L	15.14	10.40	6.20	<b>17.70</b>	12.70	8.50	7.60	6.50	5.50	8.20	11.80	54.50	25.50	30.20	25.40	25.40
SelfOcc <sub>(CVPR'24)</sub> [13]	C	10.54	0.15	0.66	5.46	12.54	0.00	0.80	2.10	0.00	0.00	8.25	<b>55.49</b>	26.30	26.54	14.22	5.60
TT-OccCamera	C	16.70	21.51	10.46	10.70	14.72	11.89	12.33	9.71	12.20	4.37	7.92	48.27	23.70	28.31	14.13	20.24
TT-OccLiDAR	C&L	<b>27.41</b>	<b>32.44</b>	<b>16.08</b>	16.88	<b>22.56</b>	<b>28.32</b>	<b>21.06</b>	<b>14.34</b>	<b>20.02</b>	<b>10.33</b>	<b>14.47</b>	52.60	<b>31.40</b>	<b>37.21</b>	<b>48.02</b>	45.37

From a signal processing perspective, the trilateral smoothing behaves as a *non-stationary low-pass* filter with locally adaptive cutoff frequencies.

For efficient occupancy estimation, we voxelize the accumulated Gaussians  $\mathcal{G}^{(t)}$  into a discrete grid  $\Omega = [\frac{X}{\delta} \times \frac{Y}{\delta} \times \frac{Z}{\delta}]$ , where each Gaussian’s contribution on a voxel is weighted based on its spatial proximity. Formally, the semantic probability of a voxel  $v \in \Omega$  is given by:

$$\mathbb{P}(\mathbf{O}_v^{(t)}) = \frac{1}{Z_v} \sum_{\mathbf{G}_i^{(t)} \in \mathcal{G}^{(t)}} (\mathbf{m}_i \cdot \mathcal{K}_\mu(i, v)), \quad (6)$$

where  $Z_v$  is the normalizing factor to ensure that  $\mathbb{P}(\mathbf{O}_v^{(t)})$  sums to 1 as a valid probability. This voxelization strategy allows flexible scaling to varying voxel resolutions during test-time, balancing efficiency and precision.

## 4. Experiments

### 4.1. Experiment Setup

Experiments were conducted on the widely used nuScenes [5] benchmark using 3D occupancy GT from Occ3D-nuScenes [35] and nuCraft [58]. The nuScenes dataset consists of 600 training scenes and 150 validation ones. Existing supervised and self-supervised methods typically require extensive offline training on the training split. In contrast, TT-Occ requires no pretraining and is directly evaluated on the validation split. In particular, Occ3D-nuScenes [35] provides voxelized occupancy annotations at  $0.4m$  resolution, covering a spatial range of  $[-40m, 40m]$  along the X and Y axes and  $[-1m, 5.4m]$  along the Z axis. nuCraft [58] offers more finer-grained annotations with a resolution of  $0.2m$ ,

covering  $[-51.2m, 51.2m]$  in the X and Y directions and  $[-5m, 3m]$  in the Z direction.

We evaluate semantic occupancy prediction using mean Intersection over Union (mIoU), computed as the average IoU across all classes. Following prior works [10, 16, 50], we exclude the “noise” and “other flat” categories, as these do not correspond to valid prompts in open-vocabulary segmentation. We primarily compare our method with self-supervised counterparts, including SimpleOcc [11], OccNeRF [50], DistillNeRF [39], GaussianOcc [10], GaussianTR [16], LangOcc [3], VEON [53], and SelfOcc [13]. These methods represent a broad range of self-supervised occupancy research and include both NeRF [26] and 3DGS [20] representation. For reference, we also include results from self-supervised methods that serve as upper bounds for performance comparison. These include BEVFormer [23] and CTF-Occ [35], which are trained with dense 3D voxel-level annotations, as well as RenderOcc [27] and OccFlowNet [4], which are trained on sparse point-level ground truth. All experiments were conducted on an NVIDIA RTX 6000 Ada (48GB) GPU. Please refer to the supplementary materials for implementation details.

### 4.2. Main Results

**Results on Occ3D-nuScenes** are shown in Table 1. It is evident that *both* variants of TT-Occ not only eliminate the need for costly offline training but also surpass the previous SOTA. Notably, TT-OccLiDAR even achieves an mIoU of 27.41, *surpassing* RenderOcc [27] (23.93), which is trained with *sparse 3D ground truth*. Moreover, our camera-only variant TT-OccCamera reaches an mIoU of 16.70, *exceeding* VEON-LiDAR [53] (15.14), which is trained using *Li-*

Table 2. 3D occupancy prediction performance on the high-resolution **nuCraft** dataset [58]. Pretraining time is reported in GPU hours.

Method	3D GT	Pretraining	mIoU $\uparrow$	barr	bike	bus	car	c-veh	moto	ped	t-cone	trail	truck	d-surf	s-walk	terr	man	vege
C-CO <sub>Net</sub> <sub>(ICCV'23)</sub> [41]	Dense	-	13.4	14.30	9.10	16.50	18.30	7.40	12.30	11.10	9.40	5.80	13.20	32.50	-	-	-	19.90
SelfOcc <sup>†</sup> <sub>(CVPR'24)</sub> [13]	$\times$	384 hrs	2.22	0.41	0.54	2.79	7.12	0.00	0.81	1.67	0.00	0.00	5.50	2.41	3.88	3.55	1.96	2.72
TT-OccCamera	$\times$	$\times$	5.95	9.21	5.58	5.35	6.19	3.93	7.28	4.86	4.73	2.89	3.44	10.15	6.17	7.18	3.92	8.45
TT-OccLiDAR	$\times$	$\times$	<b>10.92</b>	<b>14.92</b>	<b>8.15</b>	<b>12.21</b>	<b>11.93</b>	<b>10.71</b>	<b>13.82</b>	<b>6.25</b>	<b>8.39</b>	<b>6.93</b>	<b>9.05</b>	<b>12.57</b>	<b>8.41</b>	<b>10.29</b>	<b>12.89</b>	<b>17.37</b>

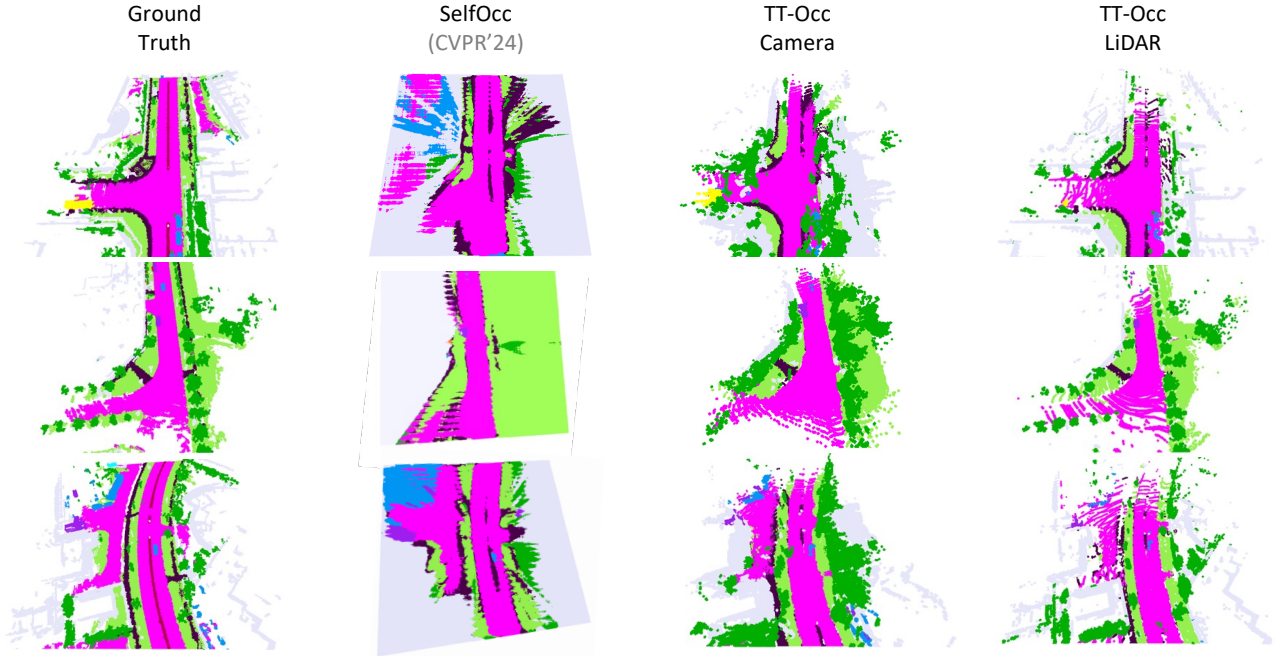


Figure 4. Qualitative comparisons on nuCraft [58] between both variants of the proposed TT-Occ and SelfOcc [13].

Table 3. Comparison of TT-Occ and SelfOcc [13] using the same OpenSeed [51] VFM under different RayIoU thresholds.

Method	mIoU	RayIoU@1	RayIoU@2	RayIoU@4
SelfOcc [13]	10.5	8.9	10.5	12.0
TT-OccCamera	13.4	10.0	12.9	15.7
TT-OccLiDAR	<b>23.6</b>	<b>20.9</b>	<b>24.1</b>	<b>25.8</b>

*DAR supervision.* In addition, compared to the SOTA self-supervised baseline SelfOcc [13], our approach achieves higher IoU not only for frequently occurring, large-area categories such as terrain and vegetation, but also shows substantial improvements on rare, dynamic, and small-area categories such as motorcycle, bus, and pedestrian.

To further evaluate the geometric quality beyond conventional mIoU, we report RayIoU under multiple thresholds, following the standard protocol in [24]. Importantly, all

methods are equipped with the same OpenSeed [51] VFM for fair comparison. As shown in Table 3, both TT-Occ variants achieve consistent gains across all RayIoU levels, indicating that our time-aware 3D Gaussians improve not only occupancy classification but also the metric accuracy of reconstructed geometry along camera rays. TT-OccCamera improves the RayIoU@4 metric by 30.8% over SelfOcc, while TT-OccLiDAR delivers the strongest performance with a substantial 115% improvement. These results highlight that TT-Occ suppresses temporal drift and produces cleaner free-space estimates, where those qualities that are not fully captured by mIoU alone. In practice, higher RayIoU translates to more reliable scene geometry for downstream planning and safety-critical perception modules, underscoring the superiority of TT-Occ given the same semantic VFM.

**Results on nuCraft** are summarized in Table 2. As no prior self-supervised methods have been trained or evaluated under this setting, we adapt SelfOcc [13] using its official

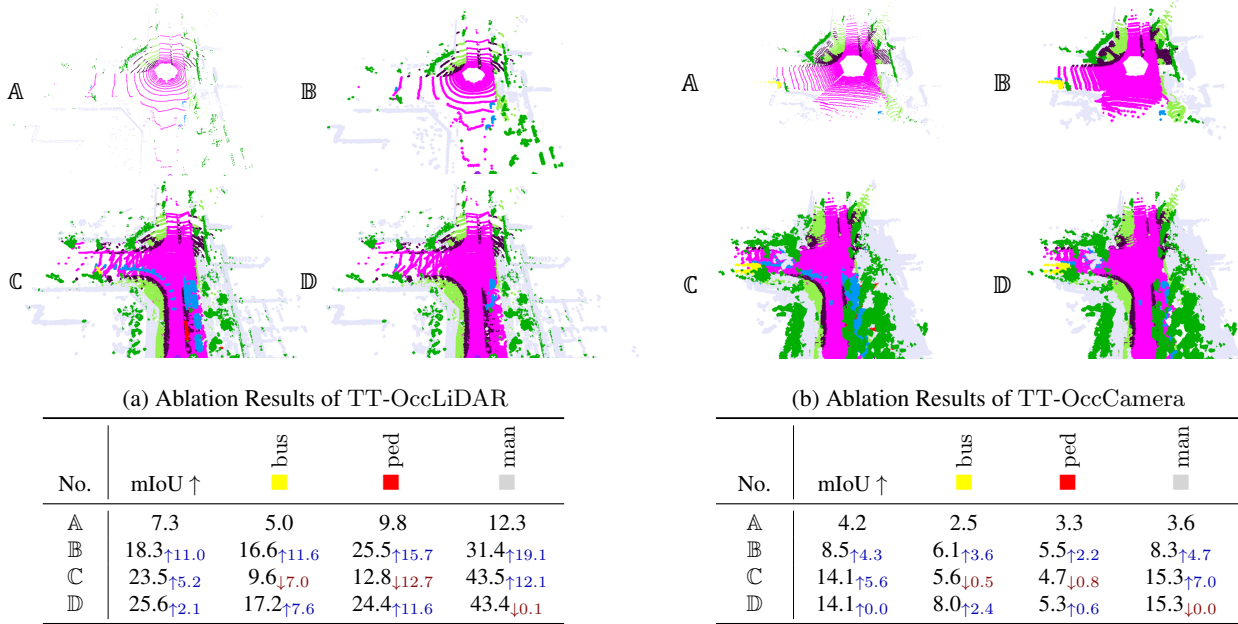


Figure 5. **Qualitative and quantitative comparisons of TT-Occ variants.** Left: TT-OccLiDAR visualization and ablation results. Right: TT-OccCamera visualization and ablation results. Each variant incrementally incorporates additional components: A: Baseline; B: Covariance-aware Voxelization; C: Inherit Previous Gaussians; D: Track Dynamic Gaussians. Please zoom in to view details.

implementation and checkpoint as a baseline for comparison. As shown in the table, TT-Occ consistently and significantly outperforms SelfOcc, demonstrating superior adaptability and robustness across *varying resolutions*.

**Qualitative comparisons on nuCraft** between both variants of TT-Occ and SelfOcc [13] are shown in Fig. 4. Several key observations emerge from these results. (1) Both our LiDAR and camera variants produce highly accurate occupancy predictions that closely align with the ground truth. In contrast, SelfOcc generates overly dense predictions, assigning occupancy to nearly all voxels, including empty regions. This not only incurs significant computational redundancy but also results in severe discrepancies with the ground truth, particularly around dynamic objects (see the radial blue regions). (2) The LiDAR-based variant produces geometrically accurate reconstructions with broad spatial coverage. However, its fidelity is inherently constrained by the sparsity of LiDAR returns, especially for small or partially scanned objects such as vehicles. (3) The camera-based variant offers denser reconstructions and better captures small objects within the field of view. Nonetheless, it may struggle with distant regions due to occlusions or limited depth resolution, and the geometry inferred from depth estimation is generally less accurate than that derived from LiDAR. Despite these challenges, TT-OccCamera still remains the state-of-the-art among vision-only occupancy methods.

### 4.3. Ablation Studies

To evaluate the effectiveness of each component in TT-Occ, we conduct ablation studies on a 10% held-out validation subset of Occ3D-nuScenes [35]. Since dynamic classes typically occupy only a small portion of the scene but play a critical role in both human perception and downstream tasks, we report not only the overall IoU and mIoU, but also the IoU of representative dynamic classes (bus, pedestrian) and a representative static class (manmade). We use 3DGS [20] as the baseline (A), where Gaussians are initialized using the “lift” strategy introduced in Section 3.1 at each time step without temporal information. Gaussians are voxelized by directly scattering their centers. As shown in the tables present in Fig. 5, this naïve approach yields poor results due to sparse observations, emphasizing the importance of using anisotropic Gaussian occupancy to better approximate scene geometry. Next, we introduce covariance-aware voxelization (Eq. (6)) and apply sigmoid-based scale regulation (B). These lead to consistent improvements across both static and dynamic classes for both LiDAR and camera inputs.

Both Variants A and B are single-frame models. Allowing Gaussians to directly accumulate across frames without tracking (C) greatly improves the overall and static class performance (*e.g.*, manmade) due to the aggregation of Gaussians for static content, which dominates the scene. However, dynamic class performance drops significantly, as untracked accumulation of moving Gaussians causes temporal inconsistency (see C in Fig. 3 for trailing artifacts). To address this,

Table 4. Impact of semantic and geometric VFMs of TT-Occ.

Method	Semantic	Geometry	mIoU $\uparrow$
SelfOcc	OpenSeeD [51]	-	11.6
TT-OccLiDAR	OpenSeeD [51]	-	23.6
	GroundedSAM2 [30]	-	21.3
	REX-Omni [17]	-	<b>27.4</b>
TT-OccCamera	OpenSeeD [51]	VGGT [37]	13.4
	OpenSeeD [51]	MapAnything [19]	14.3
	GroundedSAM2 [30]	VGGT [37]	12.2
	REX-Omni [17]	VGGT [37]	15.9
	REX-Omni [17]	MapAnything [19]	<b>16.7</b>

we incorporate tracking dynamic Gaussians as described in Section 3.2, which significantly improves the accuracy of dynamic classes while maintaining performance on static content. As shown in  $\mathbb{D}$ , this yields cleaner occupancy with trailing and ghosting artifacts largely eliminated. The ablation study on the optional TRBF fusion module is presented in supplementary materials.

**Ablation on Loosely Coupled VFM Components.** To demonstrate that our framework is loosely coupled and can support arbitrary VFMs, and to further investigate how different VFMs affect overall performance, we integrate three semantic models, namely OpenSeeD [51], GroundingSAM2 [30], and REX-Omni [17], as well as two depth-estimation backbones, VGGT [37] and MapAnything [19]. Results of all combinations on Occ3D-nuScenes [35] are summarized in Table 4. GroundingSAM2 yields the weakest performance, while REX-Omni provides the strongest semantic cues. Among depth models, MapAnything consistently outperforms VGGT, primarily because it provides metric-scale depth directly, whereas VGGT requires additional triangulation to recover metric scale, which inevitably introduces errors. Overall, these experiments confirm that our loosely coupled, modular design allows seamless integration of advanced VFMs without modification, directly benefiting from their rapid evolution and scaling trend.

**Efficiency Analysis.** We provide a detailed runtime breakdown of our pipeline for the vision-centric and LiDAR-based variants in Fig. 6 (using OpenSeeD [51]). The reported values represent the average processing time per timestep across six input images. Semantic segmentation using OpenSeeD [51] constitutes the most computationally intensive step in both pipelines, accounting for 28.5% of total runtime in the camera variant and 77.9% in the LiDAR variant. In the vision-centric scenario, the absence of LiDAR data requires additional processes such as depth estimation, triangulation-based calibration, and point cloud denoising, collectively contributing 46.5% of the overall runtime. Gaussian voxelization and the optional TRBF fusion module are relatively efficient; however, their runtime is proportional to the number of Gaussians involved. Therefore, the camera-based pipeline which has denser Gaussians experiences slightly increased computational overhead com-

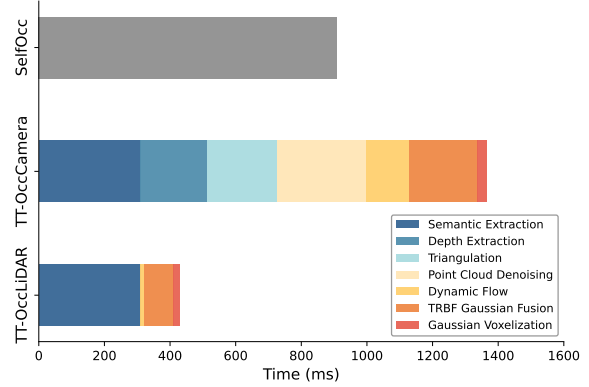


Figure 6. Inference time comparison among TT-OccCamera, TT-OccLiDAR, and SelfOcc [13]. The horizontal stacked bars illustrate the per-module runtime composition of each module.

Table 5. Memory statistics of TT-Occ.

Variant	Gaussians	Peak Mem (GB)
TT-OccLiDAR	786,874	5.6 ( <i>dom. by OpenSeeD</i> )
TT-OccCamera	2,968,647	9.9 ( <i>dom. by Gaussians</i> )

pared to its LiDAR-based counterpart. Finally, tracking dynamic Gaussians in TT-OccCamera incurs much higher computational costs compared to TT-OccLiDAR due to its reliance on dense optical flow estimation across six images using RAFT [34], whereas the LiDAR variant only applies lightweight ICP alignment for sparse foreground points.

We report the average peak GPU memory usage on high-resolution nuCraft [58] in Table 5. Because our system processes modules sequentially, the peak memory footprint is determined by the most demanding component at runtime (*e.g.*, OpenSeeD or the accumulated Gaussians). In the LiDAR variant, memory usage is primarily dominated by the semantic module, whereas in the camera variant, the growing set of Gaussians becomes the major contributor. Overall, TT-Occ maintains a peak memory consumption below 10GB for both LiDAR and camera settings.

## 5. Conclusion

In this paper, we presented TT-Occ, a practical and flexible test-time 3D occupancy framework that turns generic vision foundation models into an effective occupancy predictor through time-aware 3D Gaussians. By bypassing conventional dense occupancy decoders, TT-Occ enables arbitrary voxel resolutions, open-vocabulary object recognition, and strong temporal coherence without any additional pretraining. Extensive evaluations on Occ3D and nuCraft demonstrate that both our LiDAR-based and vision-centric variants deliver consistent accuracy gains and competitive efficiency, underscoring the promise of TT-Occ as a deployable, VFM-native solution for real-world driving environments.

## 6. Acknowledgment

This work was partially supported by ARC DE240100105, DP240101814, DP230101196, BA24006, and ARC Industrial Transformation Research Hubs IH230100013.

## References

- [1] P.J. Besl and Neil D. McKay. A method for registration of 3-d shapes. *IEEE TPAMI*, 14(2):239–256, 1992. 4
- [2] Shariq Farooq Bhat, Reiner Birkel, Diana Wofk, Peter Wonka, and Matthias Müller. Zoedepth: Zero-shot transfer by combining relative and metric depth. *CoRR*, abs/2302.12288, 2023. 2
- [3] Simon Boeder, Fabian Gigengack, and Benjamin Risse. Langocc: Self-supervised open vocabulary occupancy estimation via volume rendering. *arXiv preprint arXiv:2407.17310*, 2024. 1, 5
- [4] 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. 1, 5
- [5] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, pages 11618–11628, 2020. 4, 5
- [6] Zhuoxiao Chen, Yadan Luo, Zheng Wang, Mahsa Baktashmotlagh, and Zi Huang. Revisiting domain-adaptive 3d object detection by reliable, diverse and class-balanced pseudo-labeling. In *ICCV*, pages 3714–3726, 2023. 1
- [7] Zhuoxiao Chen, Zixin Wang, Yadan Luo, Sen Wang, and Zi Huang. Dpo: Dual-perturbation optimization for test-time adaptation in 3d object detection. In *ACM MM*, page 4138–4147, 2024. 1
- [8] Ziyu Chen, Jiawei Yang, Jiahui Huang, Riccardo de Lutio, Janick Martinez Esturo, Boris Ivanovic, Or Litany, Zan Gocic, Sanja Fidler, Marco Pavone, Li Song, and Yue Wang. Omnire: Omni urban scene reconstruction. In *ICLR*, 2025. 2
- [9] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, page 226–231, 1996. 4
- [10] 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. 1, 2, 5
- [11] Wanshui Gan, Ningkai Mo, Hongbin Xu, and Naoto Yokoya. A comprehensive framework for 3d occupancy estimation in autonomous driving. *IEEE TIV*, pages 1–19, 2024. 1, 5
- [12] Yuanhui Huang, Wenzhao Zheng, Yunpeng Zhang, Jie Zhou, and Jiwen Lu. Tri-perspective view for vision-based 3d semantic occupancy prediction. In *CVPR*, pages 9223–9232, 2023. 1, 2
- [13] Yuanhui Huang, Wenzhao Zheng, Borui Zhang, Jie Zhou, and Jiwen Lu. Selfocc: Self-supervised vision-based 3d occupancy prediction. In *CVPR*, pages 19946–19956, 2024. 1, 2, 4, 5, 6, 7, 8
- [14] Yuanhui Huang, Wenzhao Zheng, Yunpeng Zhang, Jie Zhou, and Jiwen Lu. Gaussianformer: Scene as gaussians for vision-based 3d semantic occupancy prediction. In *ECCV*, pages 376–393, 2024. 1, 2
- [15] Haoyi Jiang, Tianheng Cheng, Naiyu Gao, Haoyang Zhang, Tianwei Lin, Wenyu Liu, and Xinggang Wang. Symphonize 3d semantic scene completion with contextual instance queries. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 20258–20267. IEEE, 2024. 2
- [16] 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. 1, 2, 5
- [17] Qing Jiang, Junan Huo, Xingyu Chen, Yuda Xiong, Zhaoyang Zeng, Yihao Chen, Tianhe Ren, Junzhi Yu, and Lei Zhang. Detect anything via next point prediction. *arXiv preprint arXiv:2510.12798*, 2025. 1, 3, 8
- [18] Kapil D. Katyal, Adam Poleyov, Joseph Moore, Craig Knuth, and Katie M. Popek. High-speed robot navigation using predicted occupancy maps. In *ICRA*, pages 5476–5482, 2021. 1
- [19] Nikhil Keetha, Norman Müller, Johannes Schönberger, Lorenzo Porzi, Yuchen Zhang, Tobias Fischer, Arno Knapitsch, Duncan Zaus, Ethan Weber, Nelson Antunes, et al. Mapanything: Universal feed-forward metric 3d reconstruction. *arXiv preprint arXiv:2509.13414*, 2025. 1, 2, 3, 4, 8
- [20] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 42(4), 2023. 2, 4, 5, 7
- [21] Bohan Li, Jiajun Deng, Wenyao Zhang, Zhujin Liang, Dalong Du, Xin Jin, and Wenjun Zeng. Hierarchical temporal context learning for camera-based semantic scene completion. In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part IV*, pages 131–148. Springer, 2024. 2
- [22] Yiming Li, Zhiding Yu, Christopher B. Choy, Chaowei Xiao, José M. Álvarez, Sanja Fidler, Chen Feng, and Anima Anandkumar. Voxformer: Sparse voxel transformer for camera-based 3d semantic scene completion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 9087–9098. IEEE, 2023. 2
- [23] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *ECCV*, pages 1–18, 2022. 5
- [24] Haisong Liu, Haiguang Wang, Yang Chen, Zetong Yang, Jia Zeng, Li Chen, and Limin Wang. Fully sparse 3d panoptic occupancy prediction. *arXiv preprint arXiv:2312.17118*, 2023. 6
- [25] Yuhang Lu, Xinge Zhu, Tai Wang, and Yuexin Ma. Octreeocc:

- Efficient and multi-granularity occupancy prediction using octree queries. *CoRR*, abs/2312.03774, 2023. 2
- [26] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, pages 405–421, 2020. 5
- [27] Mingjie Pan, Jiaming Liu, Renrui Zhang, Peixiang Huang, Xiaoqi Li, Hongwei Xie, Bing Wang, Li Liu, and Shanghang Zhang. Renderocc: Vision-centric 3d occupancy prediction with 2d rendering supervision. In *ICRA*, pages 12404–12411, 2024. 1, 5
- [28] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18–24 July 2021, Virtual Event*, pages 8748–8763. PMLR, 2021. 2
- [29] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(3):1623–1637, 2022. 2
- [30] 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. 3, 8
- [31] Ziyang Song, Lin Liu, Hongyu Pan, Bencheng Liao, Mingzhe Guo, Lei Yang, Yongchang Zhang, Shaoqing Xu, Caiyan Jia, and Yadan Luo. Diver: Reinforced diffusion breaks imitation bottlenecks in end-to-end autonomous driving, 2025. 1
- [32] Stanislaw Szymanowicz, Eldar Insafutdinov, Chuanxia Zheng, Dylan Campbell, Joao F Henriques, Christian Rupprecht, and Andrea Vedaldi. Flash3d: Feed-forward generalisable 3d scene reconstruction from a single image. In *2025 International Conference on 3D Vision (3DV)*, pages 670–681. IEEE, 2025. 2
- [33] Pin Tang, Zhongdao Wang, Guoqing Wang, Jilai Zheng, Xiangxuan Ren, Bailan Feng, and Chao Ma. Sparseocc: Rethinking sparse latent representation for vision-based semantic occupancy prediction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16–22, 2024*, pages 15035–15044. IEEE, 2024. 2
- [34] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, pages 402–419, 2020. 4, 8
- [35] Xiaoyu Tian, Tao Jiang, Longfei Yun, Yucheng Mao, Huitong Yang, Yue Wang, Yilun Wang, and Hang Zhao. Occ3d: a large-scale 3d occupancy prediction benchmark for autonomous driving. In *NeurIPS*, 2023. 1, 2, 5, 7, 8
- [36] Antonín Vobecký, Oriane Siméoni, David Hurych, Spyridon Gidaris, Andrei Bursuc, Patrick Pérez, and Josef Sivic. POP-3D: open-vocabulary 3d occupancy prediction from images. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10–16, 2023*, 2023. 2
- [37] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vgg: Visual geometry grounded transformer. In *CVPR*, 2025. 1, 2, 3, 4, 8
- [38] Lizi Wang, Hongkai Ye, Qianhao Wang, Yuman Gao, Chao Xu, and Fei Gao. Learning-based 3d occupancy prediction for autonomous navigation in occluded environments. In *IROS*, pages 4509–4516, 2021. 1
- [39] Letian Wang, Seung Wook Kim, Jiawei Yang, Cunjun Yu, Boris Ivanovic, Steven Waslander, Yue Wang, Sanja Fidler, Marco Pavone, and Peter Karkus. Distillnerf: Perceiving 3d scenes from single-glance images by distilling neural fields and foundation model features. In *NeurIPS*, pages 62334–62361, 2024. 1, 4, 5
- [40] Lening Wang, Wenzhao Zheng, Yilong Ren, Han Jiang, Zhiyong Cui, Haiyang Yu, and Jiwen Lu. Occsora: 4d occupancy generation models as world simulators for autonomous driving, 2024. 2
- [41] 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. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 17804–17813, 2023. 6
- [42] Yuqi Wang, Yuntao Chen, Xingyu Liao, Lue Fan, and Zhaoxiang Zhang. Panoocc: Unified occupancy representation for camera-based 3d panoptic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16–22, 2024*, pages 17158–17168. IEEE, 2024. 2
- [43] Yi Wei, Linqing Zhao, Wenzhao Zheng, Zheng Zhu, Jie Zhou, and Jiwen Lu. Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving. In *ICCV*, pages 21672–21683, 2023. 1
- [44] Zihao Wen, Yifan Zhang, Xinhong Chen, Jianping Wang, Yung-Hui Li, and Yu-Kai Huang. Tofg: Temporal occupancy flow graph for prediction and planning in autonomous driving. *IEEE TIV*, 9(1):2850–2863, 2024. 1
- [45] Siyuan Wu, Gang Chen, Moji Shi, and Javier Alonso-Mora. Decentralized multi-agent trajectory planning in dynamic environments with spatiotemporal occupancy grid maps. In *ICRA*, pages 7208–7214, 2024. 1
- [46] Yunzhi Yan, Haotong Lin, Chenxu Zhou, Weijie Wang, Haiyang Sun, Kun Zhan, Xianpeng Lang, Xiaowei Zhou, and Sida Peng. Street gaussians for modeling dynamic urban scenes. *CoRR*, abs/2401.01339, 2024. 2
- [47] Jiawei Yang, Boris Ivanovic, Or Litany, Xinshuo Weng, Seung Wook Kim, Boyi Li, Tong Che, Danfei Xu, Sanja Fidler, Marco Pavone, et al. Emernerf: Emergent spatial-temporal scene decomposition via self-supervision. In *ICLR*, 2024. 2
- [48] Ze Yang, Yun Chen, Jingkang Wang, Sivabalan Manivasagam, Wei-Chiu Ma, Anqi Joyce Yang, and Raquel Urtasun. Unisim: A neural closed-loop sensor simulator. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1389–1399, 2023. 2

- [49] Chi Zhang, Shirui Ma, Muzhi Wang, Gereon Hinz, and Alois Knoll. Efficient pomdp behavior planning for autonomous driving in dense urban environments using multi-step occupancy grid maps. In *ITSC*, pages 2722–2729, 2022. 1
- [50] Chubin Zhang, Juncheng Yan, Yi Wei, Jiaxin Li, Li Liu, Yansong Tang, Yueqi Duan, and Jiwen Lu. Occnerf: Advancing 3d occupancy prediction in lidar-free environments. *arXiv preprint arXiv:2312.09243*, 2023. 1, 2, 5
- [51] 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 *ICCV*, pages 1020–1031, 2023. 3, 6, 8
- [52] Yunpeng Zhang, Zheng Zhu, and Dalong Du. Occformer: Dual-path transformer for vision-based 3d semantic occupancy prediction. In *ICCV*, pages 9399–9409, 2023. 1
- [53] Jilai Zheng, Pin Tang, Zhongdao Wang, Guoqing Wang, Xianguan Ren, Bailan Feng, and Chao Ma. Veon: Vocabulary-enhanced occupancy prediction. In *ECCV*, pages 92–108, 2024. 1, 2, 5
- [54] Wenzhao Zheng, Weiliang Chen, Yuanhui Huang, Borui Zhang, Yueqi Duan, and Jiwen Lu. Occworld: Learning a 3d occupancy world model for autonomous driving. In *ECCV*, pages 55–72, 2024. 2
- [55] Hongyu Zhou, Jiahao Shao, Lu Xu, Dongfeng Bai, Weichao Qiu, Bingbing Liu, Yue Wang, Andreas Geiger, and Yiyi Liao. HUGS: holistic urban 3d scene understanding via gaussian splatting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 21336–21345. IEEE, 2024. 2
- [56] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 21634–21643. IEEE, 2024. 2, 4
- [57] Xiaoyu Zhou, Jingqi Wang, Yongtao Wang, Yufei Wei, Nan Dong, and Ming-Hsuan Yang. Autoocc: Automatic open-ended semantic occupancy annotation via vision-language guided gaussian splatting. *arXiv preprint arXiv:2502.04981*, 2025. 1
- [58] Benjin Zhu, Zhe Wang, and Hongsheng Li. ncraft: Crafting high resolution 3d semantic occupancy for unified 3d scene understanding. In *ECCV*, pages 125–141, 2024. 1, 2, 5, 6, 8