AdaOcc: Adaptive 3D Occupancy Prediction for Embodied Tasks

Jinglong Wang[1,2,*,§] **Yunjie Wang**[2,3,*,§] **Zhiyang Zhang**[1,2,*,§]
Jiawei He[2,4,†] Ye Yuan[5]
Bo Qiu[6,†] Jing Zhang[1,†]

[1]School of Software, Beihang University [2]Beijing Academy of Artificial Intelligence
[3]Hebei University of Technology [4]XYZ Embodied AI
[5]ShanghaiTech University [6]University of Science and Technology Beijing

wjlzy@buaa.edu.cn 202421902001@stu.hebut.edu.cn 20377279@buaa.edu.cn
jwhe2024@gmail.com yuanye2024@shanghaitech.edu.cn
qiubo@ustb.edu.cn zhang_jing@buaa.edu.cn

*These authors contributed equally to this work.
§These authors conducted this work during an internship at XYZ Embodied AI.
†Corresponding authors: Jing Zhang, Bo Qiu, and Jiawei He.

Abstract

Embodied tasks demand accurate, flexible, and semantically rich 3D scene representations. 3D semantic occupancy is well suited to this requirement, as it can model holistic 3D spaces by encoding geometric occupancy along with semantic categories. However, existing occupancy prediction methods struggle to meet practical deployment requirements, such as adapting to varying computing budgets, sensor setups, and observation views. In this paper, we propose a point-based Adaptive 3D Occupancy Prediction method, called AdaOcc, tailored for embodied scenarios. To accommodate heterogeneous sensor inputs, AdaOcc uses an adaptive geometry-guided dual-branch encoder that can support RGB images in various numbers of views with (estimated) depth maps or LiDAR scans. AdaOcc represents occupied regions via sparse semantic points trained with a progressive query learning strategy, allowing the prediction computational budget to be flexibly adjusted through query point numbers and decoder layers. To facilitate high-fidelity geometric modeling for lightweight point-based occupancy learning, we further propose a novel containment loss that regularizes predicted points to reside within valid occupied regions. Extensive experiments show that our method achieves a new state-of-the-art on Occ-ScanNet with considerable performance improvements over previous methods. Moreover, our framework demonstrates strong practical applicability as an adaptive 3D perception module in real-world embodied systems.

1 Introduction

Accurate, flexible, and semantically rich 3D scene representations play a fundamental role in embodied AI tasks, as embodied agents typically make decisions based on their perception of the environment. 3D semantic occupancy can model holistic 3D spaces by encoding both geometric occupancy and semantic categories, providing a general-purpose representation for embodied scene

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

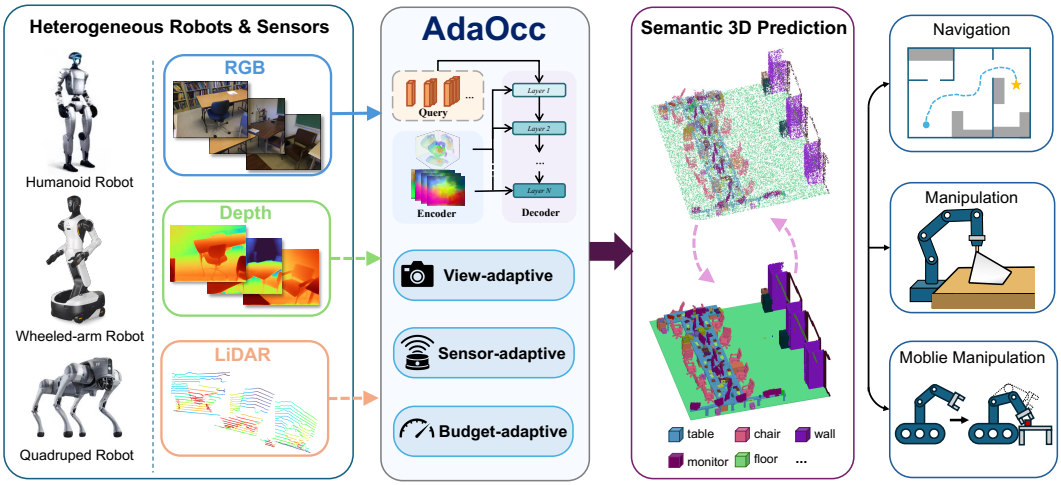

Figure 1: Overview of AdaOcc. AdaOcc targets adaptive 3D occupancy prediction for embodied scene understanding. It accepts heterogeneous visual and geometric inputs from different embodied platforms and adapts to different computational budgets through adjustable query numbers and decoder depths. AdaOcc predicts sparse semantic points, which can be converted into semantic occupancy to provide structured spatial cues for downstream embodied tasks.

understanding (Song et al., 2017; Cao and De Charette, 2022; Li et al., 2023a; Zhang et al., 2023; Wei et al., 2023; Tong et al., 2023).

Early 3D occupancy prediction methods mainly focus on outdoor driving scenes, representing 3D scenes as dense voxel grids (Zhang et al., 2023; Wei et al., 2023; Tong et al., 2023; Huang et al., 2023; Yu et al., 2023; Ma et al., 2024; Lu et al., 2024), 3D Gaussian primitives (Huang et al., 2024, 2025), or sparse points (Wang et al., 2024a). Recent research on embodied perception (Wang et al., 2024b; Wu et al., 2025; Wang et al., 2025a; Zhang et al., 2025) extends these outdoor methods by improving geometric and semantic modeling for more diverse and complex indoor scenes, with concurrent work further exploring adaptive serialization (Wang et al., 2026), geometry-prior sparse Gaussians (Zhou et al., 2026a), and open-vocabulary prediction (Zhou et al., 2026b). Despite these advances, prevailing methods still face inherent limitations in balancing accuracy and latency. Moreover, they cannot adapt to robotic platforms with heterogeneous sensor setups, and show limited generalization across diverse embodied downstream tasks.

In this work, we introduce AdaOcc, a unified point-based framework for adaptive 3D occupancy prediction in embodied tasks. As illustrated in Fig. 1, AdaOcc supports different robots and sensors, and can be applied to different downstream embodied tasks. Since robotic platforms can have different numbers of cameras and different geometric sensors, we design an adaptive geometry-guided dual-branch encoder that accepts RGB observations with varying numbers of views and incorporates geometric cues from estimated depth maps, depth-camera measurements, or LiDAR scans. We unify different geometric cues into a calibrated 3D point cloud, by calibrating the LiDAR point cloud to the RGB camera coordinate system and lifting depth maps into point clouds, enabling AdaOcc to adapt to different sensing conditions.

We observe that robotic platforms differ in computational capability, and embodied tasks demand different levels of perceptual granularity (e.g., manipulation often requires finer perception than navigation). We design AdaOcc to initialize sparse point queries from geometric inputs and reconstruct occupied semantic points via a multi-layer decoder equipped with progressive query learning. This design supports adaptive inference by adjusting the number of query points or using intermediate decoder outputs, allowing AdaOcc to balance accuracy and latency under varying computational budgets and environmental complexity.

To further enhance geometric fidelity in our point-based pipeline, we introduce a containment loss to regularize predicted points to reside within valid occupied regions, reducing surface-floating artifacts and improving spatial consistency around object boundaries and occupied structures. Experiments on Occ-ScanNet (Yu et al., 2024) show that AdaOcc achieves new state-of-the-art performance, reaching 65.29 IoU and 59.67 mIoU and outperforming the strongest prior method by 2.15 and 3.48 points,

respectively. Further efficiency, adaptability, and real-world deployment studies demonstrate that AdaOcc is not only a strong benchmark model, but also a practical spatial perception module for downstream embodied tasks. Our main contributions are summarized as follows:

- We propose AdaOcc, a point-based adaptive 3D occupancy framework for embodied scene understanding. AdaOcc enables flexible occupancy prediction under varying compute budgets, sensor setups, and inference settings.
- For robust point-based occupancy prediction, we introduce key technical designs including an adaptive geometry-guided dual-branch encoder, progressive query learning for adaptive inference, and containment-guided optimization for improved boundary fidelity.
- AdaOcc achieves state-of-the-art performance on Occ-ScanNet with large margins over previous methods. Extensive experiments and real-world embodied-system deployment demonstrate its effectiveness, efficiency, and adaptability for downstream embodied tasks.

2 Related work

Semantic occupancy prediction. Early 3D scene understanding research is largely framed as semantic scene completion, which recovers complete scene geometry and semantics from partial RGB-D observations (Song et al., 2017; Garbade et al., 2019; Li et al., 2019; Roldao et al., 2020; Chen et al., 2020; Cai et al., 2021). This formulation later moves toward vision-only scene understanding, where monocular and camera-based methods infer 3D semantic occupancy from RGB observations by learning to lift image features into 3D space (Cao and De Charette, 2022; Yao et al., 2023; Li et al., 2023a; Yu et al., 2024; Wang et al., 2026). As occupancy prediction expands to larger scenes and multi-view settings, the field further shifts from dense voxel prediction toward more efficient representations. Dense voxel and view-transformation methods provide regular 3D supervision but suffer from high memory and computation costs (Zhang et al., 2023; Wei et al., 2023; Tong et al., 2023; Wang et al., 2024c; Li et al., 2023b), motivating compact decompositions, sparse structures, Gaussian primitives, and query-based formulations for scalable 3D reasoning (Yu et al., 2023; Ma et al., 2024; Lu et al., 2024; Tang et al., 2024; Huang et al., 2025; Jia et al., 2023; Li et al., 2024a; Wang et al., 2024a; Zhou et al., 2026a,b). These works establish strong occupancy prediction backbones, but their assumptions on inputs, views, and inference budgets are usually tied to fixed evaluation protocols.

Occupancy for embodied perception. Recent embodied AI studies further extend occupancy prediction from offline scene reconstruction toward deployable spatial perception for robot navigation, interaction, and decision-making. Embodied perception benchmarks emphasize holistic 3D representations for realistic agent environments (Wang et al., 2024b; Xu et al., 2024), while robot-oriented occupancy methods begin to consider online perception, temporal scene understanding, and embodied semantic reasoning (Wu et al., 2025; Wang et al., 2025a; Zhang et al., 2025; Li et al., 2025). These works suggest that 3D occupancy is a useful spatial representation for embodied perception and robot scene understanding. However, practical robot deployment introduces additional variability that is less explored in existing studies. Most current embodied occupancy pipelines are still built around fixed sensing inputs, fixed view configurations, and fixed prediction budgets, making them difficult to adapt across platforms and tasks without redesign or retraining. Motivated by this gap, AdaOcc develops adaptive point-based occupancy prediction for embodied deployment.

3 Method

3.1 Problem Setup and Method Overview

We study semantic occupancy prediction in an agent-centric 3D space. Given one or multiple calibrated RGB images $\mathcal{I} = \{\mathbf{I}^v\}_{v=1}^{V}$ and optional geometric observations $\mathcal{G} = \{\mathbf{g}_m\}_{m=1}^{M}$, the goal is to predict the occupancy state and semantic label within a bounded local 3D region Ω. Here, $\mathcal{G}$ is a unified point set in the local agent-centric coordinate system, where M is the number of geometric points and $\mathbf{g}_m$ is the m-th point. When no independent geometric sensor is available, $\mathcal{G}$ can be estimated solely from $\mathcal{I}$, as in our Occ-ScanNet experiments. Following standard settings, Ω is discretized into a voxel grid $\mathcal{V} = \{\mathbf{p}_i\}_{i=1}^{H\times W\times D}$. Each voxel has a label $y_i \in \{0, 1, \ldots, C\}$, where 0 denotes empty space and $1, \ldots, C$ denote occupied semantic classes. Instead of directly classifying

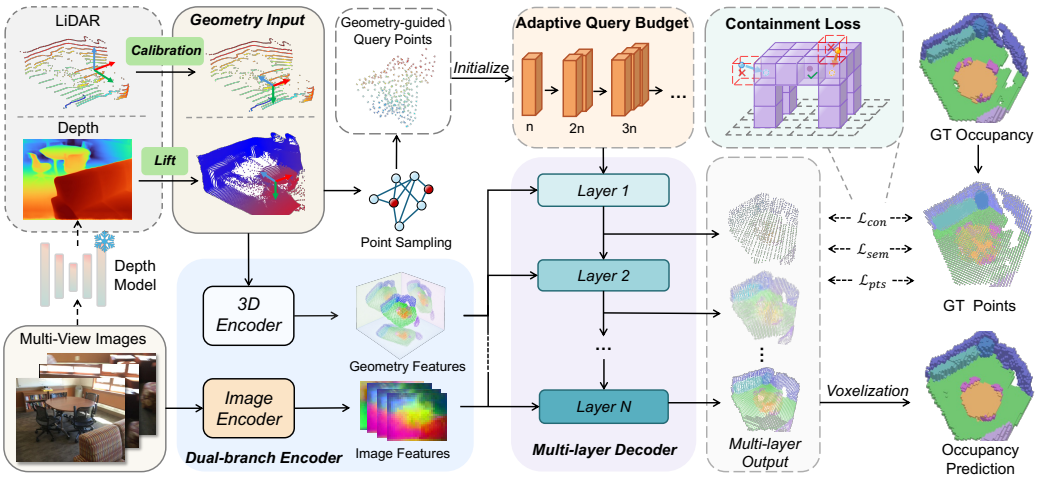

Figure 2: Overview of AdaOcc. AdaOcc is an adaptive 3D occupancy prediction framework for embodied scene understanding. It takes RGB observations as the primary input and can optionally incorporate geometric cues from estimated depth maps, depth cameras, or LiDAR scans. AdaOcc also adjusts its prediction budget under different computational constraints. The predicted semantic occupancy provides structured spatial cues for downstream embodied tasks.

all voxels, AdaOcc predicts a sparse semantic point set $\hat{\mathcal{P}} = \{(\hat{\mathbf{p}}_j, \hat{\mathbf{s}}_j)\}_{j=1}^{N}$, where $\hat{\mathbf{p}}_j$ is a predicted occupied location and $\hat{\mathbf{s}}_j$ contains its semantic logits. Finally, the predicted point set is voxelized into $\hat{\mathbf{Y}}$ for standard occupancy evaluation.

As illustrated in Fig. 2, AdaOcc is a lightweight point-based framework for adaptive 3D occupancy prediction in embodied scenes. It uses an adaptive geometry-guided dual-branch encoder to process RGB observations and optional geometric cues from estimated depth maps, depth cameras, or LiDAR scans. Geometry-guided queries are then refined by a progressive point decoder, producing sparse semantic points under adjustable query and decoder budgets. The model is trained with multi-layer point supervision and the proposed containment-guided optimization.

3.2 Adaptive Geometry-Guided Dual-Branch Encoding

Embodied occupancy prediction requires semantic cues from RGB observations and spatial grounding from geometric measurements. AdaOcc therefore adopts an adaptive geometry-guided dual-branch encoder that uses RGB observations as the primary input and optionally incorporates geometric observations when available. The two branches produce image and geometric features as:

$$\mathcal{F}^{\text{img}} = E_{\text{img}}(\mathcal{I}) = \{\mathbf{F}_r^{\text{img}}\}_{r=1}^{R_{\text{img}}}, \qquad \mathcal{F}^{\text{geo}} = E_{\text{geo}}(\mathcal{G}) = \{\mathbf{F}^{xy}, \mathbf{F}^{xz}, \mathbf{F}^{yz}\}. \tag{1}$$

Here, $\mathbf{F}_r^{\text{img}}$ denotes the image feature map at the r-th scale, and R_{img} is the number of image feature scales. The geometric observations $\mathcal{G}$ are first voxelized and encoded as follows. For depth inputs, including estimated depth maps and depth-camera measurements, AdaOcc lifts depth pixels into 3D points using camera intrinsics and transforms them with the corresponding camera extrinsics. For LiDAR scans, AdaOcc transforms points from the LiDAR coordinate frame to the RGB-camera coordinate frame using calibrated sensor extrinsics, aligning geometric measurements with visual observations. After this normalization, different geometric sources share the same point-set interface and can be processed by the same geometric branch.

To encode geometry, AdaOcc voxelizes $\mathcal{G}$ and models its local 3D structure with a sparse 3D convolutional encoder. Instead of materializing a dense 3D feature volume, the resulting sparse geometric features are compressed into three orthogonal feature planes $\{\mathbf{F}^{xy}, \mathbf{F}^{xz}, \mathbf{F}^{yz}\}$. The xy plane captures horizontal layout cues, while the xz and yz planes preserve vertical geometric profiles, providing height-aware spatial context with low memory cost.

3.3 Progressive Point Query Decoding

Inspired by OPUS (Wang et al., 2024a), AdaOcc adopts a sparse point decoder and extends it with geometry-guided query initialization and progressive query learning for budget-adaptive inference. Given an active query budget N_q, AdaOcc initializes point queries as $\mathcal{Q}_0 = \{(\mathbf{q}_i^0, \mathbf{f}_i^0)\}_{i=1}^{N_q}$ by sampling 3D locations from the geometric point set $\mathcal{G}$, where $\mathbf{q}_i^0 \in \Omega$ is the initial reference location and $\mathbf{f}_i^0$ is the query feature. This provides sensor-agnostic spatial anchors because $\mathcal{G}$ can come from estimated depth maps, depth cameras, or LiDAR scans.

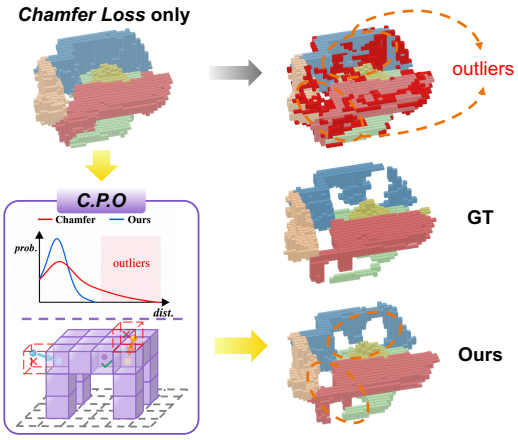

Figure 3: Illustration of Containment-Guided Point Optimization (C.P.O.), which regularizes predicted outlier points in supervised empty space to reside inside nearby occupied voxels, reducing surface-floating artifacts.

At decoder layer ℓ, each query predicts K 3D sampling coordinates around its current reference location, denoted as $\mathcal{R}_i^{(\ell)} = \{\mathbf{r}_{i,k}^{(\ell)}\}_{k=1}^{K}$. AdaOcc then samples both appearance and geometric evidence at these coordinates: $\mathbf{z}_{i,k}^{\text{img}} = \psi_{\text{img}}(\mathbf{r}_{i,k}^{(\ell)}, \mathcal{F}^{\text{img}})$ and $\mathbf{z}_{i,k}^{\text{geo}} = \psi_{\text{geo}}(\mathbf{r}_{i,k}^{(\ell)}, \mathcal{F}^{\text{geo}})$. Here, ψ_{img} projects each sampling coordinate onto calibrated image planes and samples multi-scale image features, while ψ_{geo} samples geometric features from the xy, xz, and yz planes. The sampled appearance and geometric features are aggregated by the decoder and fused with the previous query feature through a residual update:

$$\mathbf{f}_i^{(\ell)} = \mathbf{f}_i^{(\ell-1)} + \phi_\ell \left(\mathbf{f}_i^{(\ell-1)}, \{(\mathbf{z}_{i,k}^{\text{img}}, \mathbf{z}_{i,k}^{\text{geo}})\}_{k=1}^{K} \right), \tag{2}$$

where ϕ_ℓ denotes the feature aggregation and mixing operation in the ℓ-th decoder layer.

The updated query feature predicts $N_p^{(\ell)}$ semantic points at layer ℓ, denoted as $\{(\hat{\mathbf{p}}_{i,j}^{(\ell)}, \hat{\mathbf{s}}_{i,j}^{(\ell)})\}_{j=1}^{N_p^{(\ell)}}$, where $\hat{\mathbf{p}}_{i,j}^{(\ell)} = \mathbf{q}_i^{(\ell-1)} + \Delta\mathbf{p}_{i,j}^{(\ell)}$ is a predicted occupied location and $\hat{\mathbf{s}}_{i,j}^{(\ell)}$ denotes semantic logits. These predicted points serve as both the layer-wise occupancy output and the basis for spatial refinement. We recenter the next-layer query reference location by the mean predicted coordinate, $\mathbf{q}_i^{(\ell)} = \frac{1}{N_p^{(\ell)}} \sum_{j=1}^{N_p^{(\ell)}} \hat{\mathbf{p}}_{i,j}^{(\ell)}$. This iterative feature and coordinate refinement allows sparse queries to progressively align with occupied regions.

The active query number directly controls the prediction budget. To improve robustness across query budgets, AdaOcc gradually increases the active number of queries during training:

$$N_q(t) = \min\left(N_{\max}, N_{\text{init}} + \left\lfloor \frac{t - t_0}{\Delta t} \right\rfloor \Delta N \right), \tag{3}$$

where t is the training epoch, N_{init} is the initial query budget, ΔN is the query increment, Δt is the growth interval, and $N_{\max}$ is the maximum query budget. This curriculum enables the same trained decoder to adapt to different computational budgets by changing N_q or using intermediate decoder outputs at inference time.

To handle varying observation views, we apply view-level masking during training by randomly retaining a subset of RGB views and their associated geometric observations. When V_t views are retained, the total query budget becomes $N_q(t, V_t) = V_t N_q(t)$, allowing the model to process different numbers of views jointly in a single forward pass.

3.4 Containment-Guided Point Optimization

Point-based occupancy prediction is commonly supervised by Chamfer-style point-set reconstruction losses, such as Chamfer Distance (Fan et al., 2017) and Density-aware Chamfer Distance (DCD) (Wu et al., 2021). Although these losses encourage predicted points to cover the target geometry, they do not explicitly enforce whether a point lies inside an occupied region. As shown in Fig. 3, this

may produce surface-floating points that are close to occupied voxels but still located in empty space, leading to inaccurate free-space boundaries.

To address this issue, we introduce Containment-Guided Point Optimization (C.P.O.), which adds a containment loss to regularize predicted points to reside within valid occupied regions. Let $\mathcal{P}^+ = \{\mathbf{p}_i \in \mathcal{V} \mid y_i > 0\}$ denote occupied ground-truth voxel centers. Each occupied voxel center $\mathbf{p}_i$ induces an occupied box with side length $\mathbf{v}$, where $\mathbf{v}$ is the voxel size. For a predicted point $\hat{\mathbf{p}}$, its distance to the nearest occupied box is

$$d_{\text{box}}(\hat{\mathbf{p}}, \mathcal{P}^+) = \min_{\mathbf{p}_i \in \mathcal{P}^+} \left\| \text{ReLU}\left(|\hat{\mathbf{p}} - \mathbf{p}_i| - \frac{\mathbf{v}}{2} \right) \right\|_2, \tag{4}$$

where $|\cdot|$ and the subtraction by $\frac{\mathbf{v}}{2}$ are applied element-wise. The distance is zero if $\hat{\mathbf{p}}$ lies inside any occupied voxel box and increases as it drifts into empty space.

Let $\Omega_{\text{occ}} \subseteq \Omega$ denote the valid occupancy supervision region defined by the evaluation mask. For decoder layer ℓ, we collect predictions that fall into supervised empty space:

$$\mathcal{O}^{(\ell)} = \left\{ \hat{\mathbf{p}}^{(\ell)} \in \hat{\mathcal{P}}^{(\ell)} \mid \hat{\mathbf{p}}^{(\ell)} \in \Omega_{\text{occ}},\ Y(\hat{\mathbf{p}}^{(\ell)}) = 0 \right\}. \tag{5}$$

The containment loss is defined as

$$\mathcal{L}_{\text{con}} = \frac{1}{|\mathcal{S}|} \sum_{\ell \in \mathcal{S}} \frac{1}{|\mathcal{O}^{(\ell)}|} \sum_{\hat{\mathbf{p}} \in \mathcal{O}^{(\ell)}} \rho_\beta \left(d_{\text{box}}(\hat{\mathbf{p}}, \mathcal{P}^+) + m \right), \tag{6}$$

where $\mathcal{S}$ denotes the supervised decoder layers, ρ_β is the Smooth-L1 penalty, and m is a small positive offset that keeps the penalty active near occupied-box boundaries, encouraging predictions to move inside occupied voxels rather than stopping near the surface. By penalizing only predictions in supervised empty regions, the loss complements point-set reconstruction without suppressing valid occupied predictions or unsupervised space.

3.5 Training Objective

AdaOcc is trained with multi-layer supervision over semantic point predictions. At each decoder layer ℓ, the predicted point set $\hat{\mathcal{P}}^{(\ell)}$ is supervised by semantic classification, point-set reconstruction, and containment regularization. The overall objective is

$$\mathcal{L} = \sum_{\ell=1}^{L} \gamma^{L-\ell} \left(\lambda_{\text{sem}} \mathcal{L}_{\text{sem}}^{(\ell)} + \lambda_{\text{pts}} \mathcal{L}_{\text{pts}}^{(\ell)} \right) + \lambda_{\text{con}} \mathcal{L}_{\text{con}}, \tag{7}$$

where $\mathcal{L}_{\text{sem}}^{(\ell)}$ is focal loss (Lin et al., 2017b) over semantic logits, $\mathcal{L}_{\text{pts}}^{(\ell)}$ is DCD point reconstruction between $\hat{\mathcal{P}}^{(\ell)}$ and the occupied ground-truth point set $\mathcal{P}^+$, and $\mathcal{L}_{\text{con}}$ is the proposed containment loss. The containment term is applied to selected late decoder layers in C.P.O. L denotes the number of decoder layers, γ is the layer-wise decay factor, and λ_{sem}, λ_{pts}, and λ_{con} balance the semantic, reconstruction, and containment terms.

4 Experiments

4.1 Datasets and Evaluation Metrics

We evaluate AdaOcc on Occ-ScanNet, an indoor semantic occupancy prediction benchmark built upon RGB-D scans (Dai et al., 2017). The task requires predicting a semantic occupancy grid for a bounded local 3D volume from calibrated RGB observations. AdaOcc predicts over an agent-centric range that is discretized into a dense voxel grid of size $130 \times 120 \times 140$ with a voxel size of $0.08\,\text{m}$, which fully contains the official Occ-ScanNet evaluation grid of size $60 \times 60 \times 36$ at the same voxel size. Following the benchmark protocol, metrics are computed on this official grid, and predictions falling outside it are not evaluated. The semantic label space contains 11 occupied categories, including ceiling, floor, wall, window, chair, bed, sofa, table, tvs, furniture, and objects, together with an empty class; voxels with unknown ground-truth labels are ignored during evaluation.

We report results on both Occ-ScanNet and Occ-ScanNet-mini. The mini split is used for efficient comparison and ablation studies, while the full split is used to evaluate the final model against existing

Table 1: Local prediction performance on the Occ-ScanNet dataset. **Rep.** denotes the main scene representation: $\mathcal{V}$ for dense voxels, $\mathcal{T}$ for TPV features, $\mathcal{G}$ for Gaussians, and $\mathcal{P}$ for points or point queries. † indicates the AdaOcc variant with an EfficientNet image encoder. Bold and underline indicate the best and second-best results.

Dataset	Method	Rep.	mIoU	ceiling	floor	wall	window	chair	bed	sofa	table	tvs	furniture	objects	IoU
Occ-ScanNet	TPVFormer	$\mathcal{T}$	24.94	6.96	32.97	14.41	9.10	24.01	41.49	45.44	28.61	10.66	35.37	25.31	33.39
	MonoScene	$\mathcal{V}$	24.62	15.17	44.71	22.41	12.55	26.11	27.03	35.91	28.32	6.57	32.16	19.84	41.60
	ISO	$\mathcal{V}$	28.71	19.88	41.88	22.37	16.98	29.09	42.43	42.00	29.60	10.62	36.36	24.61	42.16
	SurroundOcc	$\mathcal{V}$	30.83	18.90	49.30	24.80	18.00	26.80	42.00	44.10	32.90	18.60	36.80	26.90	42.52
	GaussianFormer	$\mathcal{G}$	29.93	20.70	42.00	23.40	17.40	27.00	44.30	44.80	32.70	15.30	36.70	25.00	40.91
	OPUS	$\mathcal{P}$	38.96	39.06	45.04	34.97	28.63	35.92	49.27	54.39	37.93	23.93	45.04	34.42	45.62
	AdaSFormer	$\mathcal{V}$	45.33	37.66	57.34	40.51	29.49	43.29	60.41	63.08	47.05	29.63	54.08	36.15	54.63
	EmbodiedOcc	$\mathcal{G}$	45.48	40.90	50.80	41.90	33.00	41.20	55.20	61.90	43.80	35.40	53.50	42.90	53.95
	EmbodiedOcc++	$\mathcal{G}$	46.20	36.40	53.10	41.80	34.40	42.90	57.30	64.10	45.20	34.80	54.20	44.10	54.90
	DiScene	$\mathcal{P}$	47.17	45.21	50.63	40.38	36.73	42.28	59.68	62.04	45.60	41.17	52.42	42.72	51.99
	RoboOcc	$\mathcal{G}$	47.67	45.36	53.49	44.35	34.81	43.38	56.93	63.35	46.35	36.12	55.48	44.78	56.48
	SplatSSC	$\mathcal{G}$	51.83	49.10	59.00	48.30	38.80	47.40	62.40	67.00	49.50	42.60	60.70	45.40	62.83
	GPOcc-DPT	$\mathcal{G}$	51.88	51.42	50.35	46.97	41.84	46.98	60.39	66.16	50.51	47.97	58.88	49.23	56.96
	GPOcc-VGGT	$\mathcal{G}$	56.19	51.67	**59.93**	52.07	46.44	51.35	64.45	69.47	54.30	51.76	63.29	53.36	63.14
	AdaOcc† (Ours)	$\mathcal{P}$	59.03	53.97	57.90	53.14	47.34	55.95	70.38	74.34	59.86	52.95	66.90	56.66	64.60
	AdaOcc (Ours)	$\mathcal{P}$	**59.67**	**55.40**	58.50	**53.63**	**48.00**	**56.81**	**70.82**	**74.83**	**60.89**	**53.15**	**67.22**	**57.17**	**65.29**
Occ-ScanNet-mini	MonoScene	$\mathcal{V}$	25.90	17.00	46.20	23.90	12.70	27.00	29.10	34.80	29.10	9.70	34.50	20.40	41.90
	ISO	$\mathcal{V}$	29.40	21.10	42.70	24.60	15.10	30.80	41.00	43.30	32.20	12.10	35.90	25.10	42.90
	EmbodiedOcc	$\mathcal{G}$	45.57	29.50	49.40	41.70	36.30	41.90	60.40	59.60	46.30	34.50	58.00	43.50	55.13
	EmbodiedOcc++	$\mathcal{G}$	48.20	23.30	51.00	42.80	39.30	43.50	65.60	64.00	50.70	40.70	60.30	48.90	55.70
	SplatSSC	$\mathcal{G}$	48.87	36.60	55.70	46.50	40.10	45.60	64.50	62.40	48.60	30.60	61.20	45.39	61.47
	AdaOcc† (Ours)	$\mathcal{P}$	57.97	45.76	57.46	56.15	47.50	**59.17**	74.91	**75.16**	**57.63**	42.20	64.18	57.54	65.16
	AdaOcc (Ours)	$\mathcal{P}$	**58.74**	**48.73**	**57.61**	**56.25**	**48.18**	58.90	**75.30**	75.07	57.58	**46.52**	**64.37**	**57.61**	**65.38**

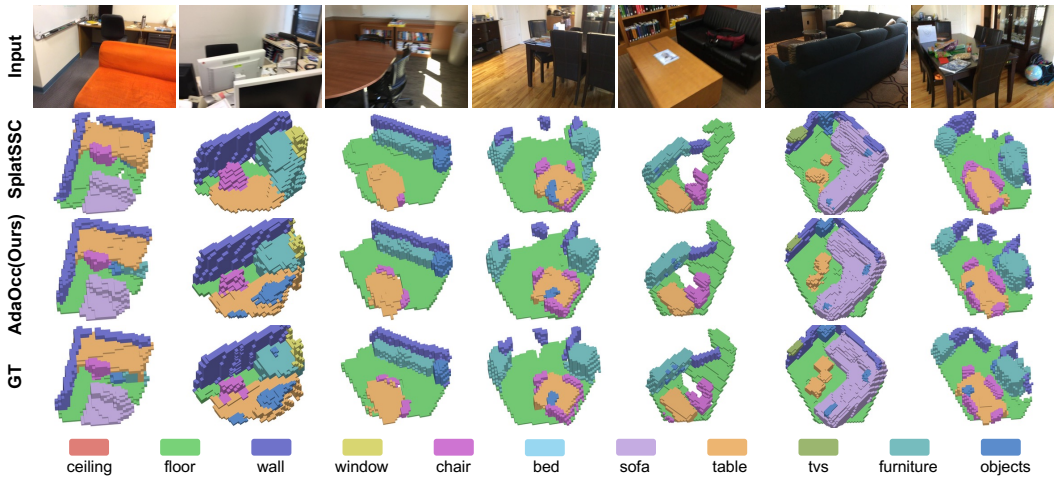

Figure 4: Qualitative comparison on Occ-ScanNet. AdaOcc produces cleaner and more complete semantic occupancy compared to the existing approaches.

methods. Following the standard protocol, we use occupied IoU to measure binary occupancy quality and mean IoU (mIoU) over the 11 occupied semantic classes to evaluate semantic occupancy prediction. We also report per-class IoU to analyze category-level performance.

4.2 Experimental Setup

By default, AdaOcc uses C-RADIOv3-B (Heinrich et al., 2025) as the image encoder and a sparse 3D convolutional encoder for geometric feature extraction. Following SplatSSC (Qian et al., 2026), our benchmark experiments use depth maps estimated from the input RGB observations by Depth Anything v2 (Yang et al., 2024); the depth maps are back-projected into 3D points and used only for geometry-guided query initialization and geometric encoding. No ground-truth depth, point cloud, or occupancy information is used as input during inference. We quantify the effect of the depth

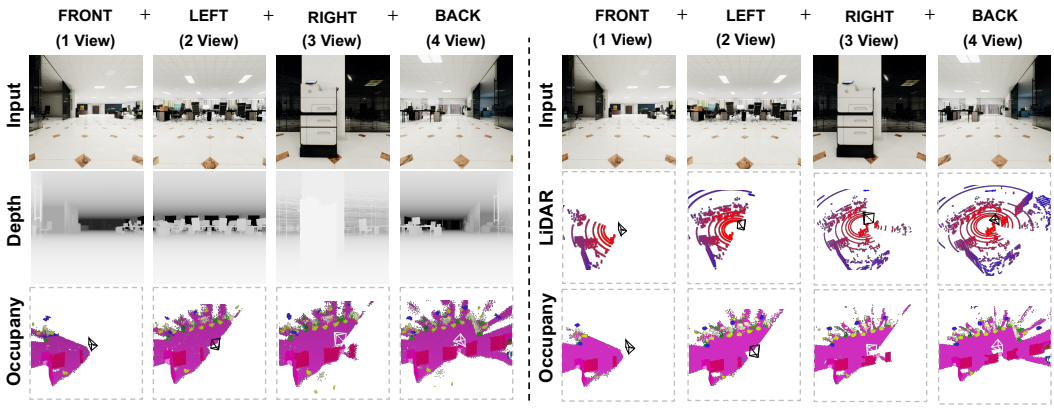

Figure 5: Qualitative adaptability analysis on TartanGround. AdaOcc accepts geometric inputs from depth maps or LiDAR scans, and processes RGB observations with varying numbers of views jointly in a single forward pass.

estimator, including a second monocular estimator and ground-truth depth, in Appendix A.5. Training follows the objective described in Sec. 3, and predicted points are voxelized for standard evaluation. AdaOcc† denotes the variant that uses the same EfficientNet image encoder as SplatSSC (Tan and Le, 2019; Qian et al., 2026), replacing the default C-RADIOv3-B encoder for a controlled comparison.

4.3 Main Results

Table 1 summarizes the main results on the Occ-ScanNet and Occ-ScanNet-mini benchmarks. On the full Occ-ScanNet split, AdaOcc achieves state-of-the-art performance with 59.67 mIoU and 65.29 IoU. The strongest prior method, GPOcc-VGGT (Zhou et al., 2026a), reaches 56.19 mIoU and 63.14 IoU, so AdaOcc improves over it by 3.48 mIoU and 2.15 IoU even though GPOcc-VGGT uses the stronger VGGT geometry (Wang et al., 2025b) prior while AdaOcc estimates depth from RGB alone. For the controlled comparison we retain SplatSSC (Qian et al., 2026) as the primary baseline because it shares our setting, estimating depth from RGB with Depth-Anything-V2, and is the strongest method under this prior: it reaches 62.83 IoU, whereas the Depth-Anything-V2 variant of GPOcc obtains 56.96, and its IoU is within 0.31 of GPOcc-VGGT despite the latter using the much stronger VGGT prior. Swapping the geometry prior inside GPOcc changes its own IoU by 6.18, indicating that the advantage of the strongest competitor is largely attributable to its prior. Relative to SplatSSC, AdaOcc improves by 7.84 mIoU and 2.46 IoU; the larger gain in mIoU than in IoU indicates stronger semantic discrimination across indoor categories. To control for the effect of the image backbone, AdaOcc† uses the same EfficientNet image encoder as SplatSSC and still reaches 59.03 mIoU and 64.60 IoU, improving over SplatSSC by 7.20 mIoU and 1.77 IoU. This suggests that the gains mainly come from the progressive point decoder, the adaptive geometry-guided dual-branch encoder, and containment-guided point optimization, rather than from the image encoder. The qualitative results further support this observation: as shown in Fig. 4, AdaOcc produces cleaner object-level structures and more complete indoor layouts.

4.4 Ablation Study

Component Ablation Table 2 ablates the key components of AdaOcc on Occ-ScanNet-mini. Starting from the RADIO image encoder baseline, adding C.P.O. and progressive training substantially improves performance from 49.65 mIoU and 56.29 IoU to 55.50 mIoU and 62.92 IoU. Adding geometric initialization further improves the result to 57.06 mIoU and 64.72 IoU, showing the benefit of sensor-guided spatial anchors. With the full 3D encoder, AdaOcc reaches 58.74 mIoU and 65.38 IoU, indicating that explicit geometric features provide

Table 2: Ablation study of key components in AdaOcc.

Image Encoder	Geometric Initialization	3D Encoder	C.P.O.	Progressive Training	mIoU	IoU
RADIO	–	–	–	–	49.65	56.29
	–	–	✓	✓	55.50	62.92
	✓	–	✓	✓	57.06	64.72
	✓	✓	–	✓	54.54	61.31
	✓	✓	✓	–	54.63	62.35
	✓	✓	✓	✓	**58.74**	**65.38**
EffNet	✓	✓	✓	✓	<u>57.97</u>	<u>65.16</u>

Table 3: Sensitivity to training hyperparameters on Occ-ScanNet-mini. All other settings are fixed within each study, and bold marks the default configuration. For the query schedule, X/Y denotes adding X queries every Y epochs, from X up to 500.

Parameter	Setting	IoU	mIoU
Margin m	0.00 / **0.04** / 0.08	64.25 / **65.38** / 65.32	57.47 / **58.74** / 58.12
C.P.O. weight	$0.5\times$ / **1**$\times$ / $2\times$	65.19 / **65.38** / 65.32	58.16 / **58.74** / 58.17
Decay γ	0.80 / **0.90** / 0.95	**65.51** / 65.38 / 65.35	58.62 / **58.74** / 58.37
Query schedule	– / 50/20 / **100/40**	62.35 / 65.17 / **65.38**	54.63 / 58.15 / **58.74**

Table 4: Efficiency and adaptability analysis on the Occ-ScanNet-mini dataset using a single NVIDIA H20 GPU. All results use the same EfficientNet image encoder as SplatSSC for a controlled comparison. (a) Query-number ablation for AdaOcc†. (b) Anchor-number ablation for SplatSSC. (c) Decoder-layer output ablation for AdaOcc†. Within each panel only the listed factor is varied while all other settings are fixed.

(a) Query (Ours)

Query	Time↓ (ms)	Mem.↓ (MiB)	mIoU	IoU
200	63.2	1710.7	55.13	62.36
500	64.4	1710.6	57.97	**65.16**
800	73.3	1798.8	**58.32**	64.58
1000	75.9	1863.7	58.25	64.29

(b) Anchor (SplatSSC)

Anchor	Time↓ (ms)	Mem.↓ (MiB)	mIoU	IoU
600	91.0	1837.2	43.30	53.75
1131	85.1	1826.1	**48.46**	**60.86**
1911	85.0	1835.1	45.26	58.85
2596	85.6	1830.3	42.95	56.33

(c) Decoder layers

Layer	Time↓ (ms)	Mem.↓ (MiB)	mIoU	IoU
3	58.2	1254.6	49.46	54.92
4	59.9	1405.9	54.51	60.96
5	62.5	1556.4	56.80	63.67
6	64.4	1710.6	**57.97**	**65.16**

complementary spatial grounding for point-based occupancy prediction. Removing either C.P.O. or progressive training causes clear performance drops, verifying the contribution of containment-guided optimization and budget-aware training. The EfficientNet variant remains competitive, confirming that the gains are not solely due to the image backbone.

Hyperparameter sensitivity. Table 3 varies the principal training hyperparameters on Occ-ScanNet-mini with all other settings fixed. The offset margin trades recall for precision: setting $m = 0$ lowers IoU and mIoU by 1.13 and 1.27 points, while increasing it to 0.08 changes them by only 0.06 and 0.62 points, so a small positive margin suffices. Halving or doubling the C.P.O. weight schedule changes IoU and mIoU by at most 0.19 and 0.58 points, and varying the decoder decay factor γ between 0.80 and 0.95 changes them by at most 0.13 and 0.37 points. By contrast, removing progressive training costs 3.03 IoU and 4.11 mIoU, while using a finer query schedule changes the result by only 0.21 IoU and 0.59 mIoU. Overall, these results show that AdaOcc is robust to moderate variations of its loss-related hyperparameters.

4.5 Efficiency and Adaptability

Table 4 evaluates the efficiency-adaptability trade-off on Occ-ScanNet-mini using a single NVIDIA H20 GPU. We use SplatSSC (Qian et al., 2026) as the main efficiency baseline, since it reports lower latency and memory usage than EmbodiedOcc (Wu et al., 2025) in its Occ-ScanNet-mini comparison. For a controlled comparison with SplatSSC, all results use the same EfficientNet image encoder, corresponding to AdaOcc†. AdaOcc† is trained with a progressive query schedule from 100 to 500 queries, while SplatSSC is trained with 1131 anchors. When varying the query budget, AdaOcc remains stable across a wide range of query numbers and provides a stronger accuracy-latency trade-off than SplatSSC. For example, AdaOcc with 500 queries achieves 57.97 mIoU and 65.16 IoU at 64.4 ms, whereas the best SplatSSC setting obtains 48.46 mIoU and 60.86 IoU at 85.1 ms. The decoder-layer ablation further shows that the same trained model can produce valid intermediate predictions without retraining, enabling flexible latency-memory-accuracy trade-offs by selecting different decoder depths.

Beyond computational adaptability, Fig. 5 shows that AdaOcc can handle heterogeneous geometric inputs and flexible multi-view observations on TartanGround (Patel et al., 2025). The same trained model produces coherent reconstructions with geometric inputs from either depth maps or LiDAR scans, and processes different numbers of RGB views jointly in a single forward pass rather than stitching per-view predictions.

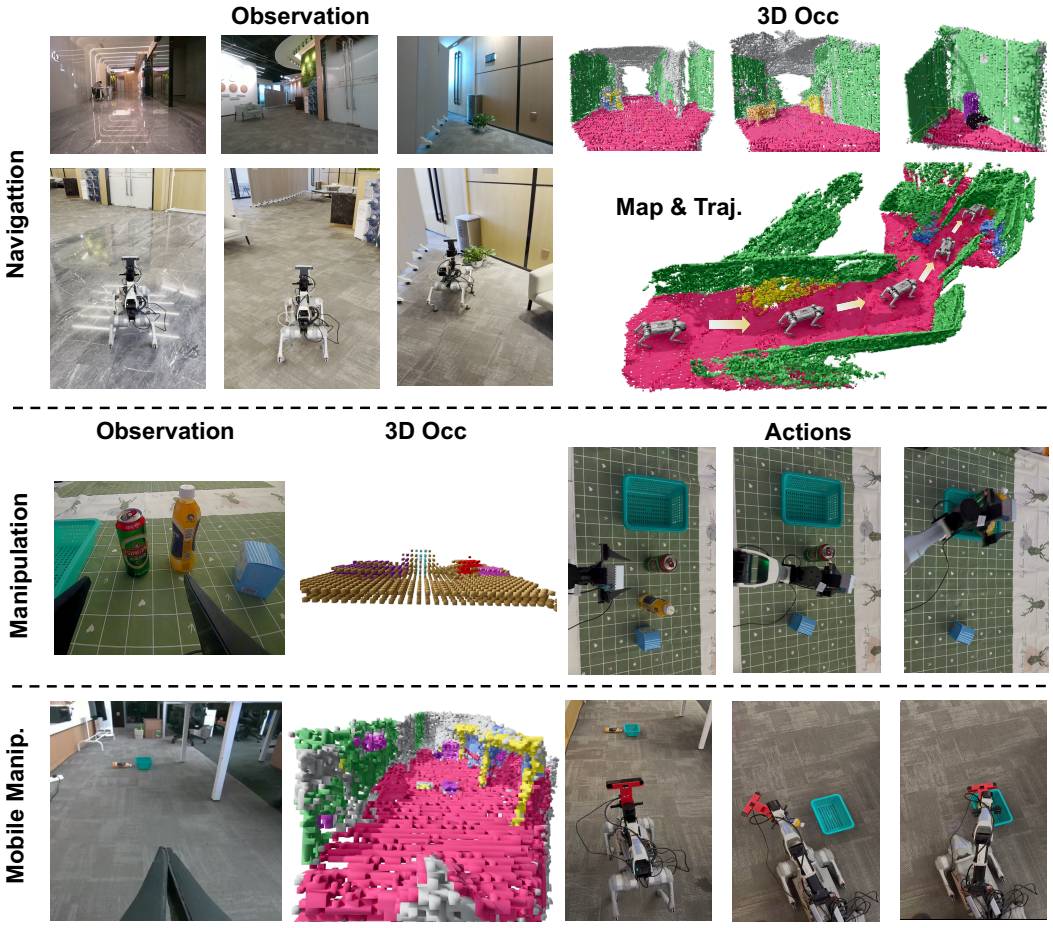

Figure 6: Real-world embodied applications of AdaOcc. AdaOcc provides semantic 3D occupancy for navigation, manipulation, and mobile manipulation. The semantic prompts are *plant* for navigation, *bottle* and *basket* for manipulation, and *trash* and *basket* for mobile manipulation.

4.6 Embodied Applications

To further examine the practical applicability of AdaOcc, we deploy it in real-world indoor embodied systems for navigation, manipulation, and mobile manipulation. As shown in Fig. 6, AdaOcc reconstructs semantic 3D occupancy from robot observations and provides structured spatial cues for downstream planning and interaction. In navigation, the reconstructed occupancy supports target-aware path planning; in manipulation, it localizes task-relevant objects and receptacles; and in mobile manipulation, it supports navigation toward objects and subsequent placement. These demonstrations suggest that AdaOcc can serve as a practical 3D perception module beyond offline benchmark evaluation. More real-robot demonstrations are available at `https://wangjl-nb.github.io/AdaOcc_web/`, and a closed-loop navigation study with quantitative results is reported in Appendix B.2.

5 Conclusions

We presented AdaOcc, a unified point-based framework for adaptive 3D occupancy prediction in embodied scenes. AdaOcc normalizes geometric cues from estimated depth maps, depth-camera measurements, or LiDAR scans into a common point-set interface, and combines them with RGB observations from varying numbers of views through an adaptive geometry-guided dual-branch encoder. Its progressive point decoder supports budget-adaptive inference by changing query numbers or using intermediate decoder outputs, while the containment loss regularizes predicted points to reside within valid occupied regions and improves spatial consistency. Experiments on Occ-ScanNet and Occ-ScanNet-mini demonstrate state-of-the-art performance, strong efficiency-adaptability trade-offs, and practical deployment in a real robot navigation system. Future work will focus on improving the reconstruction of small and thin objects in cluttered embodied scenes.

Acknowledgments and Disclosure of Funding

Funding. This work was supported by the National Natural Science Foundation of China (No. 62461160331 and No. 62132001).

Competing interests. The first three authors conducted this work during an internship at XYZ Embodied AI. The authors declare no other competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Yingjie Cai, Xuesong Chen, Chao Zhang, Kwan-Yee Lin, Xiaogang Wang, and Hongsheng Li. Semantic scene completion via integrating instances and scene in-the-loop. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 324–333, 2021.

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.

Xiaokang Chen, Kwan-Yee Lin, Chen Qian, Gang Zeng, and Hongsheng Li. 3d sketch-aware semantic scene completion via semi-supervised structure prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4193–4202, 2020.

Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017.

Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017.

Hao-Shu Fang, Chenxi Wang, Minghao Gou, and Cewu Lu. Graspnet-1billion: A large-scale benchmark for general object grasping. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11444–11453, 2020.

Martin Garbade, Yueh-Tung Chen, Johann Sawatzky, and Juergen Gall. Two stream 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.

Peter E Hart, Nils J Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.

Greg Heinrich, Mike Ranzinger, Hongxu Yin, Yao Lu, Jan Kautz, Andrew Tao, Bryan Catanzaro, and Pavlo Molchanov. Radiov2. 5: Improved baselines for agglomerative vision foundation models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 22487–22497, 2025.

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.

Yuanhui Huang, Wenzhao Zheng, Yunpeng Zhang, Jie Zhou, and Jiwen Lu. Gaussianformer: Scene as gaussians for vision-based 3d semantic occupancy prediction. In *European Conference on Computer Vision*, pages 376–393. Springer, 2024.

Yuanhui Huang, Amonnut Thammatadatrakoon, Wenzhao Zheng, Yunpeng Zhang, Dalong Du, and Jiwen Lu. Gaussianformer-2: Probabilistic gaussian superposition for efficient 3d occupancy prediction. In *Proceedings of the computer vision and pattern recognition conference*, pages 27477–27486, 2025.

Yupeng Jia, Jie He, Runze Chen, Fang Zhao, and Haiyong Luo. Occupancydetr: Using detr for mixed dense-sparse 3d occupancy prediction. *arXiv preprint arXiv:2309.08504*, 2023.

Jianing Li, Ming Lu, Juntao Liu, Hao Wang, Chenyang Gu, Wenzhao Zheng, Li Du, and Shanghang Zhang. Sliceocc: Indoor 3d semantic occupancy prediction with vertical slice representation. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 15762–15768. IEEE, 2025.

Jie Li, Yu Liu, Xia Yuan, Chunxia Zhao, Roland Siegwart, Ian Reid, and Cesar Cadena. Depth based semantic scene completion with position importance aware loss. *IEEE Robotics and Automation Letters*, 5(1):219–226, 2019.

Jinke Li, Xiao He, Chonghua Zhou, Xiaoqiang Cheng, Yang Wen, and Dan Zhang. Viewformer: Exploring spatiotemporal modeling for multi-view 3d occupancy perception via view-guided transformers. In *European Conference on Computer Vision*, pages 90–106. Springer, 2024a.

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, 2023a.

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*, 2023b.

Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Qiao Yu, and Jifeng Dai. Bevformer: learning bird's-eye-view representation from lidar-camera via spatiotemporal transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(3):2020–2036, 2024b.

Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017a.

Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017b.

Yuhang Lu, Xinge Zhu, Tai Wang, and Yuexin Ma. Octreeocc: Efficient and multi-granularity occupancy prediction using octree queries. *Advances in Neural Information Processing Systems*, 37:79618–79641, 2024.

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.

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 *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12404–12411. IEEE, 2024.

Manthan Patel, Fan Yang, Yuheng Qiu, Cesar Cadena, Sebastian Scherer, Marco Hutter, and Wenshan Wang. Tartanground: A large-scale dataset for ground robot perception and navigation. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 20524–20531. IEEE, 2025.

Rui Qian, Haozhi Cao, Tianchen Deng, Shenghai Yuan, and Lihua Xie. Splatssc: Decoupled depth-guided gaussian splatting for semantic scene completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 8520–8528, 2026.

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.

Luis Roldao, Raoul De Charette, and Anne Verroust-Blondet. Lmscnet: Lightweight multiscale 3d semantic completion. In *2020 International Conference on 3D Vision (3DV)*, pages 111–119. IEEE, 2020.

Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1746–1754, 2017.

Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.

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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15035–15044, 2024.

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.

Hao Wang, Xiaobao Wei, Xiaoan Zhang, Jianing Li, Chengyu Bai, Ying Li, Ming Lu, Wenzhao Zheng, and Shanghang Zhang. Embodiedocc++: Boosting embodied 3d occupancy prediction with plane regularization and uncertainty sampler. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 925–934, 2025a.

Jiabao Wang, Zhaojiang Liu, Qiang Meng, Liujiang Yan, Ke Wang, Jie Yang, Wei Liu, Qibin Hou, and Ming-Ming Cheng. Opus: occupancy prediction using a sparse set. *Advances in Neural Information Processing Systems*, 37:119861–119885, 2024a.

Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5294–5306. IEEE, 2025b.

Tai Wang, Xiaohan Mao, Chenming Zhu, Runsen Xu, Ruiyuan Lyu, Peisen Li, Xiao Chen, Wenwei Zhang, Kai Chen, Tianfan Xue, et al. Embodiedscan: A holistic multi-modal 3d perception suite towards embodied ai. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19757–19767, 2024b.

Xuzhi Wang, Xinran Wu, Song Wang, Lingdong Kong, and Ziping Zhao. Adasformer: Adaptive serialized transformers for monocular semantic scene completion from indoor environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 34148–34159, 2026.

Yuqi Wang, Yuntao Chen, Xingyu Liao, Lue Fan, and Zhaoxiang Zhang. Panoocc: Unified occupancy representation for camera-based 3d panoptic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17158–17168, 2024c.

Yi Wei, Linqing Zhao, Wenzhao Zheng, Zheng Zhu, Jie Zhou, and Jiwen Lu. Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21729–21740, 2023.

Tong Wu, Liang Pan, Junzhe Zhang, Tai Wang, Ziwei Liu, and Dahua Lin. Density-aware chamfer distance as a comprehensive metric for point cloud completion. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, pages 29088–29100, 2021.

Yuqi Wu, Wenzhao Zheng, Sicheng Zuo, Yuanhui Huang, Jie Zhou, and Jiwen Lu. Embodiedocc: Embodied 3d occupancy prediction for vision-based online scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 26360–26370, 2025.

Xiuwei Xu, Chong Xia, Ziwei Wang, Linqing Zhao, Yueqi Duan, Jie Zhou, and Jiwen Lu. Memory-based adapters for online 3d scene perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21604–21613, 2024.

Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *Advances in Neural Information Processing Systems*, 37:21875–21911, 2024.

Jiawei Yao, Chuming Li, Keqiang Sun, Yingjie Cai, Hao Li, Wanli Ouyang, and Hongsheng Li. Ndc-scene: Boost monocular 3d semantic scene completion in normalized device coordinates space. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9421–9431. IEEE, 2023.

Hongxiao Yu, Yuqi Wang, Yuntao Chen, and Zhaoxiang Zhang. Monocular occupancy prediction for scalable indoor scenes. In *European Conference on Computer Vision*, pages 38–54. Springer, 2024.

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.

Yunpeng Zhang, Zheng Zhu, and Dalong Du. Occformer: Dual-path transformer for vision-based 3d semantic occupancy prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9433–9443, 2023.

Zhang Zhang, Qiang Zhang, Wei Cui, Shuai Shi, Yijie Guo, Gang Han, Wen Zhao, Hengle Ren, Renjing Xu, and Jian Tang. Roboocc: Enhancing the geometric and semantic scene understanding for robots. *arXiv preprint arXiv:2504.14604*, 2025.

Changqing Zhou, Yueru Luo, and Changhao Chen. Generalizing visual geometry priors to sparse gaussian occupancy prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28578–28587, 2026a.

Changqing Zhou, Yueru Luo, Han Zhang, Zeyu Jiang, and Changhao Chen. Monocular open vocabulary occupancy prediction for indoor scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21627–21637, 2026b.

Supplemental Material

A Additional Experiments

A.1 Open-Vocabulary Semantic Occupancy Prediction

We further adapt AdaOcc to the open-vocabulary semantic occupancy setting. Different from conventional semantic occupancy prediction, which predicts a closed-set class label for each occupied voxel, our open-vocabulary occupancy prediction represents each occupied voxel with a text-aligned feature. This formulation enables arbitrary language queries over 3D occupied space, allowing semantically related concepts to activate consistent regions. This makes the occupancy representation more flexible for downstream embodied tasks, where the queried object category may be expressed with synonyms or fine-grained natural-language descriptions.

To achieve this, we modify AdaOcc with a CLIP-aligned semantic feature branch. The original image backbone is replaced by a frozen OpenAI CLIP RN101 visual encoder Radford et al. (2021). We directly connect the multi-stage ResNet features to an FPN Lin et al. (2017a). The decoder is kept query-based, but each query predicts two outputs: a binary occupancy score and a semantic feature. The binary occupancy head models whether a query corresponds to occupied space, while the feature head maps occupied queries into the CLIP text embedding space.

For semantic supervision, we pre-compute a CLIP text prototype bank for the TartanGround vocabulary. For each class name c, we encode multiple prompt templates using the CLIP RN101 text encoder and average the normalized embeddings:

$$\mathbf{t}_c = \mathrm{Norm}\left(\frac{1}{M}\sum_{m=1}^{M}\mathrm{Norm}\left(E_{\mathrm{text}}(\pi_m(c))\right)\right),$$

where $\pi_m(\cdot)$ denotes a prompt template and E_{text} is the CLIP text encoder. During training, each occupied query matched to a ground-truth semantic voxel with label y_i is supervised by the corresponding prototype $\mathbf{t}_{y_i}$.

The overall training loss is

$$\mathcal{L} = \lambda_{\mathrm{occ}}\mathcal{L}_{\mathrm{occ}} + \lambda_{\mathrm{pts}}\mathcal{L}_{\mathrm{pts}} + \lambda_{\mathrm{con}}\mathcal{L}_{\mathrm{con}} + \lambda_{\mathrm{feat}}\mathcal{L}_{\mathrm{feat}}.$$

The semantic feature loss is applied only to positive occupied queries:

$$\mathcal{L}_{\mathrm{feat}} = \lambda_{\cos}\left(1 - \cos(\mathbf{f}_i, \mathbf{t}_{y_i})\right) + \lambda_{\mathrm{ce}}\mathrm{CE}\left(\frac{\mathbf{f}_i\mathbf{T}^\top}{\tau}, y_i\right),$$

where $\mathbf{f}_i$ is the predicted voxel feature, $\mathbf{T}$ is the CLIP prototype bank, and τ is the temperature. We use cosine alignment as the main supervision and a lightweight prototype classification term to stabilize the embedding space. Class-balanced weights are used to reduce the dominance of frequent classes, while rare categories are assigned stronger weights according to voxel statistics.

We train this open-vocabulary variant on the TartanGround indoor occupancy dataset Patel et al. (2025). The voxel size is 0.08m, and the model is trained for 200 epochs with the CLIP RN101 image encoder frozen. Only the FPN, occupancy decoder, binary occupancy head, and semantic feature head are optimized. At inference time, an arbitrary text query q is encoded by the same CLIP text encoder, and each occupied voxel is scored by cosine similarity:

$$s_i(q) = \cos\left(\mathrm{Norm}(\mathbf{f}_i), \mathrm{Norm}(E_{\mathrm{text}}(q))\right).$$

Qualitatively, the trained model shows clear open-vocabulary behavior in 3D occupancy prediction. Given different but semantically related text queries, the CLIP-aligned voxel features tend to activate the same object regions in the reconstructed scene. For example, synonyms of objects such as beds and mattress produce consistent high-response areas. We visualize this by highlighting the top activated occupied points in red and rendering the remaining points in gray, as shown in Fig. 7.

A.2 Flexible Multi-View and Geometric Input

During open-vocabulary occupancy training, we further introduce a flexible input setting that supports varying numbers of camera views and geometric points from different sources, as illustrated in Fig. 5.

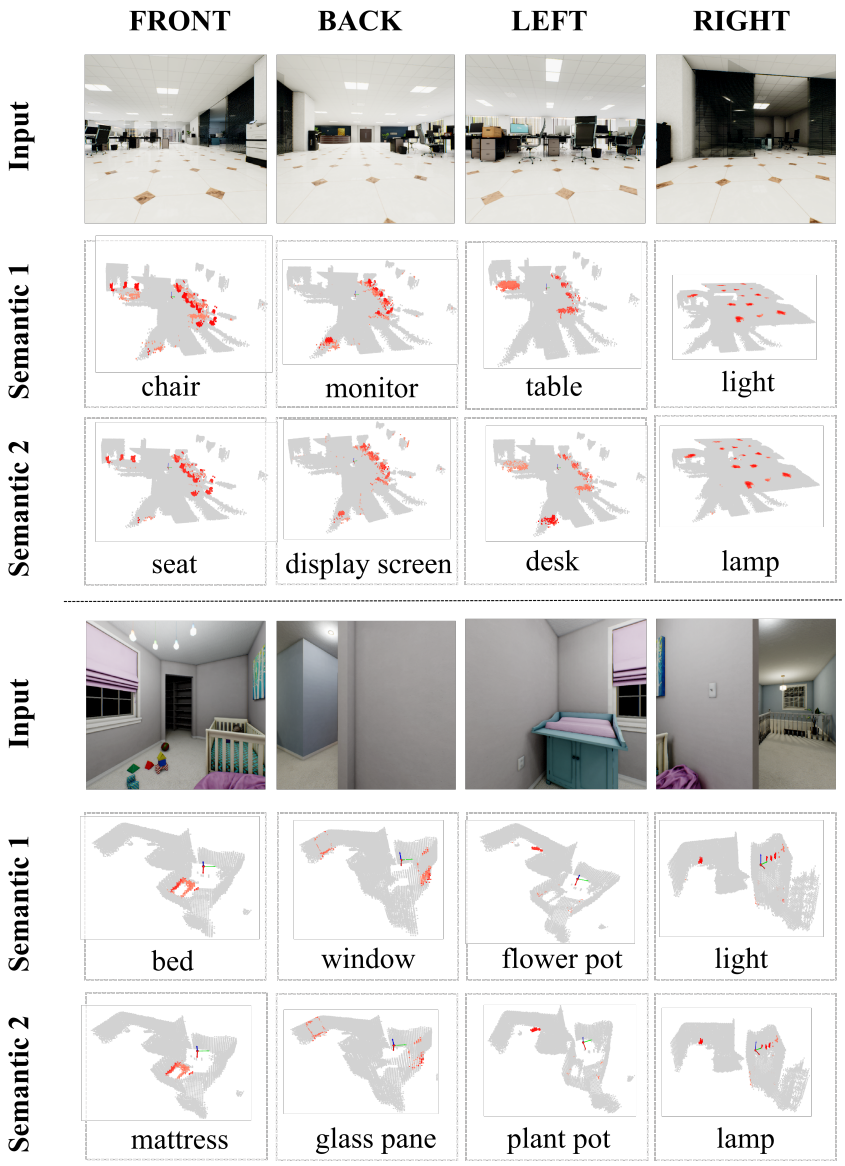

Figure 7: Open-vocabulary semantic activation in reconstructed 3D occupancy. Semantically related text queries activate consistent object regions, demonstrating that the learned CLIP-aligned voxel features can respond to free-form language.

Table 5: Comparison of the GT-depth and LiDAR input variants on the TartanGround test set.

Variant	Input Source	IoU (%)	mIoU (%)
GT-depth version	Depth-derived points	67.55	26.13
LiDAR version	LiDAR points	55.65	19.40

Although this is used in our open-vocabulary experiments, the same training strategy is general and can be applied to a broader range of experimental settings.

To enable this flexibility, we use four cameras, including the front, back, left, and right views. During training, we randomly select a subset of the available camera views as input, with the number of selected views ranging from one to four. The model therefore sees different view combinations during training and can be evaluated with one, two, three, or four input views at inference time.

The geometric input can be provided either by depth-derived points or by LiDAR points. For the RGB-D setting, points are generated from the depth maps of the selected camera views. For the RGB-LiDAR setting, the raw LiDAR point cloud is used as the geometric input instead. When only a subset of camera views is selected, the same LiDAR point cloud is paired with the selected RGB views, while the image features are computed only from those views. The point loading and deduplication pipeline is adjusted accordingly, but the occupancy decoder and prediction heads remain unchanged.

With this design, the same training framework supports randomly sampled one-to-four-view input under both RGB-D and RGB-LiDAR settings. The model can thus reconstruct occupancy from different numbers of camera views while using either depth-based points or LiDAR points as geometric support. We compare two settings on TartanGround test set. The results are shown in Table 5.

Table 6 reports the occupancy quality under one to four input views on 1,379 TartanGround test samples from 12 scenes (ground-truth depth). The model is trained once with randomly sampled view subsets and evaluated without retraining. Within the same observed regions, the reconstruction quality changes only slightly as the number of input views varies, indicating that AdaOcc is robust to different view configurations. Additional views mainly expand spatial coverage: under the fixed four-view union, the IoU increases from 17.24 with a single view to 67.55 with four views.

Table 6: Occupancy IoU / mIoU on the TartanGround test set under one to four input views, using the model trained with randomly sampled view subsets and evaluated without retraining. "Current-view union" evaluates the union of the regions covered by the views currently provided as input, whereas "Four-view union" uses the fixed union covered by all four cameras.

Input views	Front	Left	Back	Right	Current-view union	Four-view union
Front	62.87 / 23.24	–	–	–	62.87 / 23.24	17.24 / 6.93
Front + Left	65.98 / 25.56	64.43 / 23.95	–	–	65.29 / 25.06	32.34 / 13.29
Front + Left + Back	65.88 / 25.15	64.99 / 24.37	67.24 / 25.26	–	66.14 / 25.04	52.14 / 19.18
Four views	67.57 / 26.55	66.33 / 25.18	68.63 / 26.49	67.29 / 27.03	67.55 / 26.14	67.55 / 26.14

A.3 Occ3D-nuScenes Dataset Results

We conduct outdoor 3D semantic occupancy prediction experiments on Occ3D-nuScenes dataset. Compared with other SOTA methods, we achieve 3.8 RayIoU improvement. With LiDAR and surrounding RGB inputs, we have batter performance with 50.9 RayIoU, showing that we can utilize more geometry information and support different sensor and view inputs.

A.4 Hyperparameter Sensitivity

We study the sensitivity of AdaOcc to its main training hyperparameters on Occ-ScanNet-mini. All variants keep the remaining settings fixed and are trained for 200 epochs.

Margin and C.P.O. weight. Table 8 reports the effect of the C.P.O. offset margin m and of the C.P.O. loss weight schedule. A positive margin converts a moderate recall reduction ($87.88 \rightarrow 83.41$) into a clear precision gain ($70.50 \rightarrow 75.32$), and $m = 0.04$ attains the best IoU, mIoU, and F1 (79.16, against 78.24 for $m = 0$ and 79.02 for $m = 0.08$). Halving or doubling the weight schedule changes IoU and mIoU by at most 0.19 and 0.58 points, respectively.

Decay factor and query schedule. Table 9 varies the decoder loss decay factor γ and the granularity of the progressive query schedule. Both affect the final result only mildly: within $\gamma \in [0.80, 0.95]$ the variation is at most 0.13 IoU and 0.37 mIoU, and using a finer query schedule instead of the default changes the result by only 0.21 IoU and 0.59 mIoU. Removing progressive training, by contrast, costs 3.03 IoU and 4.11 mIoU, indicating that the prediction budget, rather than the loss coefficients, is the dominant factor.

Table 7: Comparison of RayIoU results on the Occ3D-nuScenes dataset.

Methods	**RayIoU**$_{1m}$	**RayIoU**$_{2m}$	**RayIoU**$_{4m}$	**RayIoU**
RenderOcc (Pan et al., 2024)	13.1	19.6	25.5	19.5
BEVFormer (Li et al., 2024b)	26.1	32.9	38.0	32.4
BEVDet-Occ	23.6	30.0	35.1	29.6
BEVDet-Occ (8f)	26.6	33.1	38.2	32.6
FB-Occ (16f) (Li et al., 2023b)	26.7	34.1	39.7	33.5
SparseOcc (8f) (Tang et al., 2024)	28.0	34.7	39.4	34.0
SparseOcc (16f) (Tang et al., 2024)	29.1	35.8	40.3	35.1
OPUS-T (8f) (Wang et al., 2024a)	31.7	39.2	44.3	38.4
OPUS-S (8f) (Wang et al., 2024a)	32.6	39.9	44.7	39.1
OPUS-M (8f) (Wang et al., 2024a)	33.7	41.1	46.0	40.3
OPUS-L (8f) (Wang et al., 2024a)	34.7	42.1	46.7	41.2
AdaOcc (Ours) (1f) (RGB+Pseudo depth map)	37.55	46.14	51.20	44.97
AdaOcc (Ours) (1f) (RGB+LiDAR)	**46.60**	**51.58**	**54.48**	**50.89**

Table 8: Effect of the C.P.O. offset margin m and of the C.P.O. loss weight schedule on Occ-ScanNet-mini. Bold marks the default configuration.

Study	Setting	IoU	mIoU	Precision	Recall	F1
Margin m	0.00	64.25	57.47	70.50	87.88	78.24
	0.04	**65.38**	**58.74**	75.32	83.41	**79.16**
	0.08	65.32	58.12	**75.36**	83.06	79.02
C.P.O. weight	$0.05 \rightarrow 0.15$	65.19	58.16	74.32	**84.14**	–
	0.10 → 0.30	**65.38**	**58.74**	**75.32**	83.41	–
	$0.20 \rightarrow 0.60$	65.32	58.17	75.17	83.29	–

A.5 Depth Estimator Sensitivity

AdaOcc estimates depth from RGB with Depth Anything v2 (DAv2) by default. To quantify the dependence on this external component, we replace DAv2 with a second monocular estimator, MoGe-2, and, separately, replace the estimated depth with ground-truth depth, keeping the architecture, the evaluation protocol, and all training settings unchanged. We first measure depth quality over all 6,646 unique Occ-ScanNet-mini frames: both estimators process every frame, and predictions are resized to the ground-truth resolution and evaluated over valid depths.

Occupancy accuracy follows depth quality consistently: replacing DAv2 with the less accurate MoGe-2 reduces IoU and mIoU by 3.06 and 2.59 points, whereas ground-truth depth improves them by a further 4.11 and 4.32 points. The remaining headroom therefore lies in the external depth estimator, whose errors concentrate around reflective or low-texture surfaces, thin structures, and occluded boundaries. The DAv2 entry of Table 10 stems from the run shared with the MoGe-2 and ground-truth variants and differs from the DAv2 results of Tables 2 and 3 (65.38 IoU / 58.74 mIoU) by training stochasticity.

B Embodied System

We use different robots and perception sensors to conduct real-world embodied tasks, as listed in Table 12. The devices used in these tasks are shown in Fig. 8. The details of each task are as following:

Table 9: Effect of the decoder loss decay factor γ and of the progressive query schedule on Occ-ScanNet-mini. The query schedule is written as queries added per epoch interval; bold marks the default configuration.

Study	Setting	IoU	mIoU
Decay γ	0.80	65.51	58.62
	0.90	**65.38**	**58.74**
	0.95	65.35	58.37
Query schedule	none	62.35	54.63
	50/20	65.17	58.15
	100/40	**65.38**	**58.74**

Table 10: Depth quality and the resulting occupancy accuracy on Occ-ScanNet-mini. All three entries are produced in a single controlled run with identical training settings.

Depth input	AbsRel ↓	RMSE (cm) ↓	IoU ↑	mIoU ↑
Ground truth	0.0000	0.00	69.61	62.67
DAv2	0.0441	13.82	65.50	58.35
MoGe-2	0.0782	23.43	62.44	55.76

B.1 Object Navigation with Go2

We implement a lightweight embodied navigation system to connect AdaOcc predictions with real-robot execution. The platform uses a Unitree Go2 quadruped robot equipped with a D455 camera. AdaOcc reconstructs the observed scene as a semantic PLY point cloud, where each point is associated with its 3D position and semantic label or text-aligned semantic response. This semantic point cloud serves as the spatial representation for downstream navigation.

Given a language-specified target, we use open-vocabulary semantic activation in the reconstructed point cloud to determine the navigation goal. The text query activates semantically relevant occupied regions, and the goal position is obtained from the corresponding activated target region. In this way, the navigation module specifies goals through semantic concepts rather than manually defined metric coordinates.

To connect 3D reconstruction with robot navigation, we convert the semantic PLY map into a 2.5D traversability costmap. Floor or otherwise traversable points are projected onto a ground-plane grid as free space, while non-floor points within the robot collision-height range are treated as obstacles. The obstacles are further inflated according to the robot radius and a safety margin, yielding a conservative costmap for quadruped navigation. Unknown cells are treated as non-traversable during planning.

On the resulting costmap, we run A* search Hart et al. (1968) from the robot start position to the semantic target. The planner snaps the start and goal positions to nearby traversable cells when necessary, uses clearance-aware costs to prefer safer regions, and resamples the planned path into robot-friendly waypoints. The final output is a waypoint sequence in the map frame, including position, heading, speed, and accumulated path distance, which can be directly used by the Go2 waypoint tracker. Fig. 9 visualizes the costmap and the planned trajectory.

B.2 Closed-Loop Navigation on FloorPlan-R2R

We evaluate AdaOcc as the occupancy module of a closed-loop navigation policy on the FloorPlan-R2R val-unseen split in Habitat. Unlike standard vision-and-language navigation, the agent receives a semantic floorplan, in which each room is represented by a 2D polygon with a semantic category and a region ID, together with a concise instruction that specifies a start region, a target region, and an object- or relation-level stopping condition (e.g., "start in office 16, go to living room 13, and stop at the couch"). At every step the agent observes egocentric RGB-D and selects among MOVE_FORWARD, TURN_LEFT, TURN_RIGHT, and STOP, so success requires both reaching the correct region and stopping where the local observation satisfies the language condition. The agent is not given its pose in floorplan coordinates or the floorplan-to-world scale.



Figure 8: The photo of devices used in real-world embodied tasks.

AdaOcc is used only as the online occupancy module: it converts accumulated egocentric RGB-D observations into a spatial representation used for floorplan grounding, target-region projection, and collision-free planning. The grounding, planning, and fixed zero-shot value-map stopping components are identical for both methods below, so the comparison isolates the utility of the occupancy representation.

Table 11: Closed-loop navigation on the FloorPlan-R2R val-unseen split. Both methods share the same grounding, planning, and stopping backend; FP-Nav additionally uses privileged floorplan alignment, whereas our pipeline estimates pose and scale online.

Method	NE ↓	OSR ↑	SR ↑	SPL ↑
FP-Nav	9.40	39.80	28.80	24.00
AdaOcc + value-map stopping	**5.38**	**63.66**	**43.53**	**25.04**

NE is the final distance to the goal, OSR records whether the agent reaches the success region at any time, SR additionally requires a correct final STOP, and SPL weighs SR by path efficiency. The AdaOcc occupancy representation reduces NE by 4.02 m and improves OSR, SR, and SPL by 23.86, 14.73, and 1.04 points. The large OSR gain indicates more reliable cross-region navigation and the SR gain shows that this advantage persists through final task completion, while the smaller SPL gain reflects that stopping accuracy, rather than path efficiency, remains the main bottleneck.

This setting differs from occupancy prediction for autonomous driving, where benchmarks typically use a fixed, calibrated surround-view rig, predict a predefined metric volume around a vehicle in structured outdoor scenes, and evaluate occupancy mainly as a perception output. Here a moving embodied agent incrementally maps an unseen indoor environment from changing egocentric views, and narrow passages, nearby clutter, dense occlusion, and fine object or room boundaries directly affect action selection and stopping. The predicted occupancy therefore acts as persistent spatial memory and traversability evidence inside a closed-loop policy rather than as an endpoint in itself.

B.3 Pick and Place with ARX X5

We further validate the practical utility of AdaOcc in a real-world manipulation setup for picking bottle and place into basket using the ARX X5 platform. Given an input scene, AdaOcc predicts three object-level point sets, including two point clouds corresponding to the two target bottles and one point cloud corresponding to the basket.

The predicted point cloud of each bottle is independently provided to GraspNet Fang et al. (2020), which outputs an executable grasp pose for the ARX X5 manipulator. This process yields the grasping poses for the two bottles shown in the figure. For placement, we compute the geometric center of the predicted basket point cloud and use it as the target placement position.

Based on the predicted bottle points, the generated grasp poses, and the basket-centered placement coordinate, the ARX X5 platform executes a complete manipulation pipeline for waste-bottle collection, including target grasping and basket placement. This experiment shows that the object-level geometric outputs of AdaOcc can be directly coupled with off-the-shelf grasp planning for real-robot manipulation.

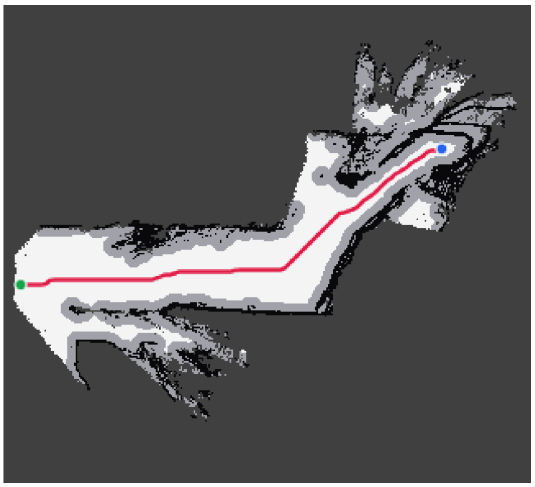

Figure 9: Costmap-based navigation from AdaOcc predictions. The semantic PLY prediction is converted into a 2.5D traversability costmap, where free space, inflated obstacles, and unknown regions are represented in different gray levels. Given an open-semantic target activation, A* plans a collision-aware path from the robot start position to the target region and exports map-frame waypoints for the Unitree Go2 robot.

Table 12: Equipment Used in Our Real-world Experiments

Device Name	Type	Embodied Tasks
Unitree Go2	Quadruped robot	Navigation/Mobile manipulation
ARX X5	Robotic arm	Maniulation/Mobile manipulation
Intel RealSense D455	Long-range stereo depth camera	Navigation
Intel RealSense D405	Short-range depth camera	Mobile manipulation
MRDVS M4	ToF RGB-D sensor	Maniulation

B.4 Mobile Manipulation with ARX X5 on Go2

Building on the navigation and manipulation pipeline above, we further demonstrate a long-range embodied mobile manipulation task using ARX X5 on Go2 platform. In this setting, the robot first observes a target bottle and a basket from a distance. The bottle location predicted by AdaOcc is used as the semantic navigation target, allowing Go2 to move from the initial observation point to the vicinity of the target object.

After reaching a suitable manipulation range, we apply the same object-centric perception and manipulation logic as above. Specifically, AdaOcc predicts the point sets of the target bottle and the basket in the local scene. The predicted bottle point set is then fed into GraspNet to generate an executable grasp pose, while the center of the predicted basket point set is used as the placement target. The robot then completes the bottle collection task by grasping the bottle and placing it into the basket.

This experiment demonstrates that AdaOcc supports a complete embodied pipeline that starts from long-range semantic target localization, continues with waypoint-based quadruped navigation, and finally enables goal-conditioned bottle grasping and basket placement in a real-world setting.

C Standard Deviation of Experiments

We repeat the best AdaOcc configuration on Occ-ScanNet-mini with five random seeds. S0–S4 correspond to seeds 0, 1785784487, 1602478221, 1383549716, and 977375669, respectively. As shown in Table 13, AdaOcc shows stable performance across runs, with a standard deviation of 0.20

Table 13: Standard deviation over five random seeds on Occ-ScanNet-mini. S0–S4 correspond to seeds 0, 1785784487, 1602478221, 1383549716, and 977375669, respectively. All results are reported at epoch 200. Bold and underline indicate the best and second-best results among the five runs.

Seed	mIoU	ceiling	floor	wall	window	chair	bed	sofa	table	tvs	furniture	objects	IoU
S0	**58.74**	**48.73**	**57.61**	56.25	**48.18**	58.90	75.30	75.07	57.58	**46.52**	64.37	57.61	65.38
S1	58.26	47.13	57.39	56.12	47.64	58.69	75.19	74.99	57.63	43.60	64.47	**58.05**	65.24
S2	<u>58.54</u>	<u>48.69</u>	57.54	<u>56.40</u>	<u>48.15</u>	<u>58.91</u>	<u>75.53</u>	<u>75.27</u>	**57.93**	42.91	64.82	57.81	**65.49**
S3	58.46	48.03	57.57	56.17	47.94	**58.94**	**75.54**	75.03	<u>57.85</u>	43.18	**65.19**	57.59	<u>65.46</u>
S4	58.29	46.71	<u>57.59</u>	**56.51**	47.52	58.88	74.71	**75.28**	57.30	<u>43.93</u>	<u>64.92</u>	<u>57.85</u>	<u>65.46</u>
Mean	58.46	47.86	57.54	56.29	47.89	58.86	75.25	75.13	57.66	44.03	64.75	57.78	65.41
Std.	0.20	0.91	0.09	0.16	0.30	0.10	0.34	0.14	0.25	1.45	0.34	0.19	0.10

mIoU and 0.10 IoU. Across the five runs, the mIoU ranges from 58.26 to 58.74 and the IoU ranges from 65.24 to 65.49.

D Implementation Details and Compute Resources

D.1 Implementation Details

We provide the main implementation details for the Occ-ScanNet-mini AdaOcc configuration in Table 14. The model is trained for 200 epochs with progressive query learning, where the active query budget increases from 100 to 500 queries in five stages. During evaluation, the full query budget is used unless otherwise specified.

Coordinate frames. Each Occ-ScanNet sample provides camera intrinsics, camera-to-world extrinsics, ego-to-world pose, and LiDAR-to-ego pose. We use the ego/occupancy frame as the canonical local 3D frame for training. In the data loader, the original metadata contains the transform from ego to LiDAR, denoted as $T_{\text{ego}\to\text{lidar}}$, and an identity transform from ego to occupancy, $T_{\text{ego}\to\text{occ}} = I$. Depth-derived points are first reconstructed in the camera frame, transformed to the LiDAR frame using the corresponding sensor-to-LiDAR extrinsic, and then transformed into the occupancy frame by

$$T_{\text{lidar}\to\text{occ}} = T_{\text{ego}\to\text{occ}} T^{-1}_{\text{ego}\to\text{lidar}}.$$

Since $T_{\text{ego}\to\text{occ}}$ is identity in this setting, the occupancy frame is aligned with the current ego frame. For image feature sampling, we construct an ego-to-image projection matrix $T_{\text{ego}\to\text{img}} = KT_{\text{ego}\to\text{cam}}$, where K is the camera intrinsic matrix. Any image-space resize/crop transform is represented by a 4×4 image data augmentation matrix A_{ida} and left-multiplied into the projection matrix, i.e.,

$$T_{\text{ego}\to\text{img}} \leftarrow A_{\text{ida}} T_{\text{ego}\to\text{img}},$$

so that the projected 3D points remain consistent with the resized input image.

Depth back-projection and pseudo points. For Occ-ScanNet dataset, we use the front camera only. The input RGB image has raw resolution 1296×968 and is resized/cropped to 960×720. The associated depth map is used to construct pseudo points. For each valid depth pixel (u, v) with depth d, where $d \in [0.1, 7.5]$ meters, we back-project using the OpenCV camera convention:

$$x = \frac{(u + 0.5 - c_x)d}{f_x}, \qquad y = \frac{(v + 0.5 - c_y)d}{f_y}, \qquad z = d.$$

The resulting camera-frame point (x, y, z) is transformed to the LiDAR frame and then to the occupancy frame as described above. Each point stores five channels:

$$(x, y, z, \text{intensity}, \Delta t),$$

where the intensity channel is set to zero and Δt is the timestamp offset. In the reported single-frame configuration, only the current frame is kept, so $\Delta t = 0$ for the selected points. Points outside the local occupancy range are removed before voxelization.

Occupancy range and voxelization. The local 3D range is

$$[-3.2, -4.8, -5.6,\ 7.2, 4.8, 5.6],$$

corresponding to $(x_{\min}, y_{\min}, z_{\min}, x_{\max}, y_{\max}, z_{\max})$ in meters. We use a voxel size of

$$\Delta = (0.08, 0.08, 0.08)\ \mathrm{m}.$$

Thus the Cartesian grid has size

$$N_x = 130, \quad N_y = 120, \quad N_z = 140.$$

This range is the agent-centric volume over which AdaOcc predicts; it fully contains the official Occ-ScanNet evaluation grid of size $60 \times 60 \times 36$ at the same voxel size.

Input points are voxelized with voxel size 0.08 m, maximum 10 points per voxel, and at most 90k/120k voxels for train/test respectively. The voxel coordinates are computed by

$$\mathbf{i} = \left\lfloor \frac{\mathbf{p}_{\mathrm{occ}} - \mathbf{p}_{\min}}{\Delta} \right\rfloor,$$

where $\mathbf{p}_{\min} = (-3.2, -4.8, -5.6)$. Only points whose voxel indices fall inside the valid grid are kept.

Ground-truth occupancy voxels. The Occ-ScanNet ground truth is loaded from compressed `labels.npz` files. Each file contains semantic occupancy labels, camera visibility masks, LiDAR visibility masks, and, for raw Occ-ScanNet supervision, the raw semantic grid with its voxel origin and voxel size. We use the raw Occ-ScanNet ground truth branch during training and evaluation. The raw semantic grid has size $60 \times 60 \times 36$ at 0.08 m together with a per-frame voxel origin and voxel size, and follows the native $[y, x, z]$ axis order. Voxels labelled 0 are empty, labels $1, \ldots, 11$ are the occupied semantic classes, and 255 denotes unknown voxels, which are excluded from both training targets and evaluation. Raw semantic labels in the range $1, \ldots, 11$ are treated as occupied semantic classes and converted to zero-based class IDs $0, \ldots, 10$; free space and unknown voxels are ignored for the point-level semantic targets.

The raw Occ-ScanNet tensor is indexed in (y, x, z) order. Therefore, when constructing 3D target points, a raw voxel coordinate (i_y, i_x, i_z) is converted to Cartesian order as (i_x, i_y, i_z). The raw voxel origin denotes the center of voxel $(0, 0, 0)$ rather than the minimum grid corner. Hence the 3D center of a raw occupied voxel is computed as

$$\mathbf{p}_{\mathrm{gt}} = \mathbf{o}_{\mathrm{raw}} + (i_x, i_y, i_z) \odot \Delta_{\mathrm{raw}},$$

where $\mathbf{o}_{\mathrm{raw}}$ is the raw voxel-origin vector and Δ_{raw} is the raw voxel size. This convention avoids a half-voxel offset when matching predicted points to raw Occ-ScanNet supervision.

Prediction decoding and voxel aggregation. The decoder predicts normalized 3D refinement points inside the configured local range. We decode them back to metric coordinates in the occupancy frame. For raw Occ-ScanNet evaluation, predicted points are further transformed from the local occupancy frame to the world frame using

$$T_{\mathrm{occ}\to\mathrm{world}} = T_{\mathrm{ego}\to\mathrm{world}} T_{\mathrm{occ}\to\mathrm{ego}}.$$

They are then mapped into the raw Occ-ScanNet grid using the raw voxel origin and voxel size. Since the raw grid origin is center-defined, the binning rule is

$$\mathbf{i}_{\mathrm{raw}} = \left\lfloor \frac{\mathbf{p}_{\mathrm{world}} - \mathbf{o}_{\mathrm{raw}}}{\Delta_{\mathrm{raw}}} + 0.5 \right\rfloor.$$

If multiple predicted points fall into the same voxel, we keep the top-k points by their maximum semantic confidence, with $k = 10$, and average their semantic scores to obtain the voxel-level prediction. During evaluation, points with semantic confidence below 0.4 are discarded, and points too far from their query center are removed using a center-distance threshold of 1.0 m. We additionally apply the configured sparse padding step to fill small holes before computing occupancy metrics.

Training target construction. For training, occupied raw voxels are converted to sparse 3D points and used as point-level targets. Predicted query points are matched to ground-truth occupied voxel centers using a bidirectional point matching objective. The point loss is a density-aware Chamfer-style loss implemented with Smooth-L1 distance, and semantic classification is supervised with focal loss. The final loss combines the last decoder layer and auxiliary decoder layers with an exponential layer decay of 0.9. A containment loss is additionally applied on decoder layers d4 and d5 to penalize predictions that fall into known empty regions and pulls them toward nearby occupied voxel boxes. Its weight is linearly increased from 0.1 to 0.3 between epochs 20 and 60.

Evaluation protocol. Evaluation uses the full 500-query budget. Predicted points are mapped from the local occupancy frame into the official Occ-ScanNet raw grid ($60 \times 60 \times 36$, voxel size 0.08 m) using the per-frame voxel origin and voxel size, and compared against its raw semantic labels. Voxels with unknown labels (255) are excluded, which corresponds to the camera-visible region of the selected front camera; predictions falling outside the official grid are ignored. We report occupied IoU and mIoU over the 11 occupied semantic classes.

Table 14: Main implementation details for AdaOcc on Occ-ScanNet-mini.

Data and Input Configuration	
Dataset	Occ-ScanNet-mini
Training / validation samples	4,639 / 2,007
RGB observations $\mathcal{I}$	Single-view RGB, current frame only
Geometric observations $\mathcal{G}$	Depth-derived 3D points for benchmark experiments
Image resolution	960×720 after resizing
Occupancy range Ω	$[-3.2, -4.8, -5.6, 7.2, 4.8, 5.6]$
Voxel size $\mathbf{v}$	$[0.08, 0.08, 0.08]$
Semantic classes	11 occupied classes + free label
Model Configuration	
Image encoder E_{img}	C-RADIOv3-B, last 4 blocks unfrozen
Image feature dimension	512
Geometric encoder E_{geo}	Sparse 3D convolutional encoder with multi-plane feature aggregation
Decoder layers L	6
Maximum query budget $N_{\max}$	500
Initial query budget N_{init}	100
Progressive query schedule $N_q(t)$	Increase by 100 queries every 40 epochs until 500
Queries at evaluation	Full query budget unless otherwise specified
Points per query refinement	$[1, 2, 4, 8, 16, 32]$ across decoder layers
Optimization	
Optimizer	AdamW
Base learning rate	2×10^{-4}
Weight decay	0.01
Batch size	8 per GPU, 64 global
Training epochs	200
Learning-rate schedule	500-iteration linear warmup + cosine decay to 2×10^{-7}
Mixed precision	Enabled with static loss scale 512
Gradient clipping	L2 norm clipping with max norm 35
Loss Configuration	
Semantic loss $\mathcal{L}_{\text{sem}}$	Focal loss with $\gamma_{\text{focal}} = 2.0$, $\alpha = 0.25$, $\lambda_{\text{sem}} = 2.0$
Point reconstruction $\mathcal{L}_{\text{pts}}$	DCD Loss with Smooth L1 regression, $\lambda_{\text{pts}} = 0.5$
Decoder loss decay	$\gamma = 0.9$
Containment loss $\mathcal{L}_{\text{con}}$	Applied to layer 5 and 6, with λ_{con} ramped from 0.10 to 0.30

Table 15: Experimental Settings on Different Datasets

Datasets	RGB Input	LiDAR Input	Open Vocabulary	Depth Map Type	#Views
Occ-ScanNet	✓			Pseudo	1
Occ3D-nuScenes (setting1)	✓	✓			6
Occ3D-nuScenes (setting2)	✓			Pseudo	6
TartanGround (setting1)	✓	✓	✓		1-4
TartanGround (setting2)	✓		✓	GT	1-4

D.2 Experimental Settings on Different Datasets

We provide a summary of the experimental settings across datasets in Table 15. The compared settings cover different combinations of RGB input, LiDAR input, open-vocabulary support, depth map sources, and the number of available views.

D.3 Compute Resources

The compute resources for the Occ-ScanNet-mini RADIO configuration are summarized in Table 16. The model was trained on one internal Linux node with 8 NVIDIA H20 GPUs using PyTorch distributed data parallelism. Each GPU has approximately 97,871 MiB of memory. The 200-epoch training process took about 8.0 hours on 8 GPUs, corresponding to approximately 64 GPU-hours for training. Including the final validation pass, the total recorded run time was about 8.2 hours. The final 500-query training stage used about 11.0 GB memory per GPU, while validation used about 4.8 GB per GPU. Validation latency measured from the training log was approximately 0.34 s per sample per GPU, including data loading and metric computation.

Table 16: Compute resources for the Occ-ScanNet-mini AdaOcc run.

Item	Value
Hardware	8 $\times$ NVIDIA H20 GPUs
GPU memory	97,871 MiB per GPU
Training framework	PyTorch DDP with NCCL backend
CUDA / cuDNN	CUDA 12.1 / cuDNN 8.9.2
PyTorch / TorchVision	2.2.0 / 0.17.0
Training time	8.0 h on 8 GPUs
Total recorded run time	8.2 h including final validation
Estimated training compute	64 GPU-hours
Steady-state training memory	11.0 GB per GPU at 500 queries
Validation memory	4.8 GB per GPU
Validation latency	0.34 s per sample per GPU
Checkpoint size	3.4 GB per checkpoint

E Limitations

AdaOcc still has limited capability in modeling very small or thin objects. Since the method represents occupied regions with a finite set of semantic points and voxelizes them into a standard occupancy grid for evaluation, fine-grained structures may be under-represented when they occupy only a few voxels or are weakly observed from the input views. This limitation can be further affected by occlusion, noisy geometric inputs, or incomplete depth/LiDAR observations in cluttered indoor scenes. Beyond fine-grained geometry, AdaOcc estimates depth from RGB with an external model (Depth Anything v2), which adds inference latency and introduces failure modes on reflective or low-texture surfaces, thin structures, and occluded boundaries; replacing the estimate with ground-truth depth quantifies

both the dependency and the remaining headroom (Appendix A.5). Our evaluation is further limited to indoor benchmarks in the single-frame or direct-aggregation setting, so broader indoor evaluation and explicit temporal memory remain open directions. Future work will mitigate these issues through finer-grained query allocation, stronger local geometric refinement, small-object-aware supervision, and a more robust treatment of the geometric input.

F Broader Impacts

AdaOcc aims to improve adaptive 3D scene perception for embodied agents. Its potential positive impacts include enabling robots and embodied systems to build more accurate spatial representations under varying sensing and computational conditions, which may benefit indoor navigation, assistive robotics, and other downstream embodied tasks. The budget-adaptive design may also make 3D perception more accessible on platforms with limited onboard computation.

At the same time, occupancy prediction systems may introduce risks when deployed in real-world environments. Incorrect occupancy or semantic predictions could affect downstream planning decisions, especially in safety-critical settings involving humans or fragile objects. In addition, 3D perception systems used in indoor environments may raise privacy concerns if raw sensor data are collected or stored. AdaOcc does not directly address these deployment-level issues. Practical use should therefore include task-specific safety checks, privacy-preserving data handling, and validation under the target sensing and operating conditions.

G Assets and Licenses

This work uses existing academic datasets, pretrained models, and baseline implementations for research evaluation. We cite the original sources for Occ-ScanNet, Occ-ScanNet-mini, TartanGround, Depth Anything v2, EfficientNet, RADIO/C-RADIO, OPUS, SplatSSC, and other compared methods in the main paper. These assets are used only for academic research and evaluation under their original licenses and terms of use. We do not redistribute the original datasets, pretrained checkpoints, or third-party code as part of this submission.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The claims in the abstract and introduction correspond to the method described in Section 3 and the results reported in Section 4. We do not make claims beyond the scope of the presented theoretical and empirical evidence.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The Supplementary Section E discusses limitations of AdaOcc in modeling very small or thin objects, and analyzes how sparse point budgets, voxelization, occlusion, and incomplete geometric observations may affect fine-grained reconstruction.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [N/A]

Justification: The paper does not present theoretical results, theorems, or formal proofs; it proposes a model architecture and validates it empirically.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Sections 3 and 4 describe the model architecture, training objective, datasets, evaluation metrics, input construction, and experimental settings used for the main results and ablations.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in

some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: Code is available at `https://github.com/wangjl-nb/AdaOcc`.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Section 3 provides the training objective, while Section 4 specifies the datasets, evaluation metrics, input construction, model variants, and ablation settings. Additional implementation details, including optimization hyperparameters and compute resources, are reported in Supplementary Section D.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: Supplementary Section C reports mean and standard deviation over five random seeds for the best AdaOcc configuration on Occ-ScanNet-mini, capturing variability due to random initialization and training stochasticity.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.

- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Supplementary Section D reports the implementation details and compute resources, including hardware configuration, software environment, batch size, optimizer, training time, approximate GPU-hours, memory usage, and validation latency.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research follows the NeurIPS Code of Ethics and uses standard academic datasets and evaluation protocols for 3D scene perception. The work does not involve human-subject experiments, private personal data, or harmful data collection.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Supplementary Section F discusses potential positive impacts for embodied perception and possible risks in safety-critical or privacy-sensitive deployment settings.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: The paper does not release high-risk generative models, scraped Internet-scale datasets, or other assets that require special controlled-release safeguards. The work focuses on 3D occupancy perception using standard academic benchmarks.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The paper cites the original works for the datasets, pretrained models, and baseline methods used in the experiments. The Supplementary Section G summarizes the existing assets used in this work and states that they are used according to their original licenses and terms, without redistribution.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not introduce or release a new dataset, benchmark, or standalone model asset as part of the current submission.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing experiments, human-subject studies, or participant data collection.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: The paper does not involve crowdsourcing experiments, human-subject studies, or participant data collection, so IRB approval or equivalent review is not applicable.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: LLMs are not used as an important, original, or non-standard component of the core method, experiments, data generation, or evaluation. Any LLM use was limited to writing, editing, or formatting assistance and does not affect the scientific contribution or results.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.