

# Multi-View Pedestrian Occupancy Prediction with a Novel Synthetic Dataset

Sithu Aung<sup>1</sup>, Min-Cheol Sagong<sup>1</sup>, Junghyun Cho<sup>1,2,3</sup>

<sup>1</sup>Korea Institute of Science and Technology

<sup>2</sup>AI-Robotics, KIST School, University of Science and Technology

<sup>3</sup>Yonsei-KIST Convergence Research Institute, Yonsei University  
{sithu, mcsagong, jhcho}@kist.re.kr

## Abstract

We address an advanced challenge of predicting pedestrian occupancy as an extension of multi-view pedestrian detection in urban traffic. To support this, we have created a new synthetic dataset called **MVP-Occ**, designed for dense pedestrian scenarios in large-scale scenes. Our dataset provides detailed representations of pedestrians using voxel structures, accompanied by rich semantic scene understanding labels, facilitating visual navigation and insights into pedestrian spatial information. Furthermore, we present a robust baseline model, termed **OmniOcc**, capable of predicting both the voxel occupancy state and panoptic labels for the entire scene from multi-view images. Through in-depth analysis, we identify and evaluate the key elements of our proposed model, highlighting their specific contributions and importance.

**Project page** — <https://sithu31296.github.io/mvpocc>

## Introduction

Recognition of crowded pedestrians in large-scale environments, particularly in surveillance systems with multiple CCTV cameras, has been an important area of research. Various multi-view tasks including multi-view crowd counting (Zhang and Chan 2019; Zhang, Lin, and Chan 2021), multi-target multi-camera tracking (Gan et al. 2021; Jeon et al. 2023; Luna et al. 2022; Nguyen et al. 2023), multi-view human pose estimation (Dong et al. 2019; Tu, Wang, and Zeng 2020; Zhang et al. 2021), and multi-view pedestrian detection (Hou, Zheng, and Gould 2020; Song et al. 2021; Hou and Zheng 2021; Aung et al. 2024) have emerged in this area. However, these tasks are still limited to the 2D ground plane when estimating crowd density and global positions of individuals. Although multi-view pose estimation methods yield 3D poses of pedestrians, they are yet constrained to small environments with a limited number of people and struggle with regressing metric-level 3D joint locations.

To address these limitations, we propose extending multi-view pedestrian detection to multi-view pedestrian occupancy prediction, as illustrated in Fig. 1, where the occupancy status has to be predicted for every voxel in the entire scene and every grid cell within the bird’s eye view plane. Predicting occupancy status within the ground plane

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

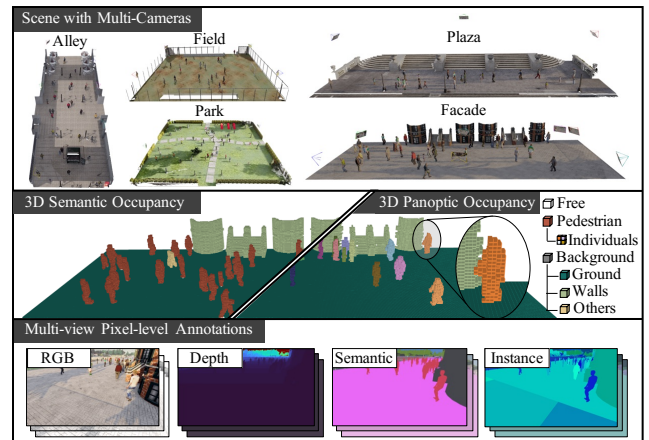


Figure 1: Visualizations of the proposed dataset. The primary objective is to predict the semantic and instance labels of the voxels and determine each pedestrian’s location within the scene. The dataset includes five expansive scenes with dense pedestrian activity. (Best viewed in color.)

allows us to identify each pedestrian instance, while voxel-level occupancy prediction provides rich geometric information about the scene and a more articulate representation of pedestrians, including their heights. Additionally, estimating the semantic label of each voxel enables us to gain insight into the surrounding contexts, which is an added advantage. However, as highlighted in Table 1, existing datasets do not provide this level of information, underscoring the need for more comprehensive datasets tailored for spacious scenes with crowded pedestrians.

Hence, we propose a novel synthetic **Multi-View Pedestrian Occupancy** dataset, **MVP-Occ**, comprising five large-scale scenes, designed to mimic real-world environments. In our dataset, the entire scene is represented by voxels, and each voxel is annotated with one of five classes, indicating whether it belongs to a pedestrian, the background environment, or is empty. Furthermore, our dataset goes beyond occupancy prediction, offering a variety of perception tasks such as segmentation, pose estimation, semantic scene completion, depth estimation, and more, spanning both 2D and 3D domains. To the best of our knowledge, our dataset

Dataset	Type	Modality	#frames	#scenes	#peds.	pose
PETS2009	Real	I	0.42k	1	40	-
CMUPanoptic	Real	I+D	297k	1	8	✓
Human-M3	Real	I+P	12.2k	3	10	✓
CityStreet	Real	I	0.5k	1	150	-
WildTrack	Real	I	0.4k	1	20	-
MultiviewX	Syn.	I	0.4k	1	40	-
GMVD	Syn.	I	6k	7	40	-
<b>MVP-Occ</b>	Syn.	I+D+P	12.5k	5	100	✓

Table 1: Comparison with available multiview pedestrian datasets. Ours is by far the most comprehensive with rich 2D/3D annotations, including *I*: image, *D*: depth, *P*: point cloud, human poses, segmentation, and occupancy labels.

is the first to provide such comprehensive 3D annotations for urban surveillance scenarios, opening avenues for further research in pedestrian recognition and scene understanding.

Alongside the newly proposed dataset, we introduce a robust baseline model, termed **OmniOcc** to tackle the challenge of multi-view pedestrian occupancy prediction. Our model is characterized by its expandability and ability to handle various combinations of camera configurations and scene dimensions during train and test time. Additionally, our model is designed to predict 2D pedestrian occupancy in the ground plane while simultaneously predicting 3D semantic occupancy for the entire scene. Using pedestrian instances as center locations, our model can further group semantic occupancy into instance and panoptic occupancies.

Furthermore, we establish a new evaluation benchmark to assess the performance of the proposed model on both 2D and 3D occupancy predictions. Our experiments cover conventional evaluations on the same scene and also address the challenging task of synthetic-to-real transfer with the WildTrack dataset (Chavdarova et al. 2018), for which we have generated the ground-truth segmentation data. Through these extensive analyses, we conduct an in-depth examination of the proposed model, dissecting the individual contributions of each component. The results underscore the superiority of our approach over previous multi-view detection methods, particularly highlighting its prowess in synthetic-to-real evaluation, where existing methods falter in transferring knowledge when confronted with disparate scenes.

## Related Work

**Multi-view human pose estimation.** Existing works (Dong et al. 2019; Tu, Wang, and Zeng 2020; Lin and Lee 2021) on multi-view 3D pose estimation methods have predominantly concentrated on constrained environments with a limited number of individuals. As summarized in Table 1, widely used datasets for this task, such as CMU-Panoptic (Joo et al. 2015) and Human-M3 (Fan et al. 2023) are not intended for expansive scenarios involving a large number of pedestrians. Consequently, there exists a demand for datasets that accurately portray complex CCTV-like situations, characterized by sparsely installed cameras that cap-

ture large scenes with dense pedestrian crowds.

**Multi-view pedestrian detection.** Recent research efforts (Song et al. 2021; Hou and Zheng 2021; Qiu et al. 2022; Aung et al. 2024) have focused on detecting dense pedestrians in large scenes using multi-view information. However, the datasets associated with these methods, such as WildTrack (Chavdarova et al. 2018), MultiviewX (Hou, Zheng, and Gould 2020), and GMVD (Vora et al. 2023), only provide pedestrian locations on the ground plane, lacking detailed information on poses and actions. This limitation may restrict their utility in certain applications. To address this gap and align with our research objectives, we aim to develop a dataset that includes detailed information about pedestrians and provides rich context at the scene level, thus facilitating a broader range of research applications.

**Semantic scene completion.** Representing scenes with volumetric occupancy and semantic labels has been actively researched in the context of semantic scene completion (SSC) tasks. SSC methods typically rely on pairs of RGB and depth images as input and require intermediate representations (Song et al. 2017; Chen et al. 2020; Wu et al. 2020; Tang et al. 2022). Moreover, they focus mainly on small indoor environments, capturing static scenes and objects, while later advancements (Roldao, de Charette, and Verroust-Blondet 2020; Cao and de Charette 2022; Yao et al. 2023) expanded the SSC to outdoor scenes by inferring dense geometry and semantics from a single monocular image. However, these methods still have limitations, including using a single RGB image as input and focusing on scene reconstruction with large objects.

**3D occupancy prediction.** There has been a surge of interest in representing the surrounding environment around autonomous vehicles through dense occupancy and semantic labels using multiple cameras, following Tesla AI Day (Tesla 2022). Subsequent research works (Huang et al. 2023; Wei et al. 2023; Li et al. 2023; Zhang, Zhu, and Du 2023; Liu et al. 2023) have aimed to address this challenge through the 3D occupancy prediction task. Concurrently, efforts (Tong et al. 2023a; Fang and Li 2023; Tong et al. 2023b; Tian et al. 2023) have been made to improve the quality of datasets by densifying the sparse occupancy labels obtained from the LiDAR sensors. Although our proposed task shares similarities with the works mentioned earlier, there are distinct differences. Our task focuses on scenarios with variable camera setups that capture static scenes with dense pedestrian crowds. In contrast, above methods address scenarios involving a fixed camera setup installed on autonomous vehicles that observe dynamic scenes and prioritize large surrounding environments rather than individual pedestrians.

## Method

First, we introduce a baseline model to tackle the multi-view pedestrian occupancy prediction task. This task entails predicting the location of each pedestrian and classifying voxels into one of the pre-defined semantic classes. The subsequent sections provide detailed information about various components of the proposed model.

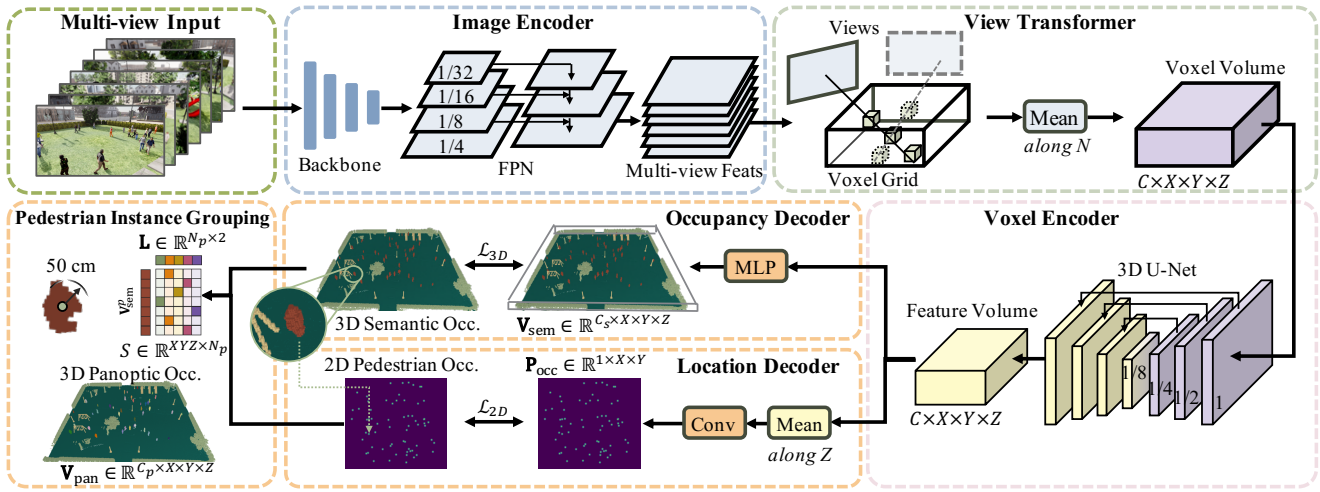


Figure 2: Overview of the proposed model. Image features are extracted using a backbone network augmented with an FPN. Next, multi-view 2D features are projected onto the voxel grid along rays and processed with a 3D U-Net to construct a feature volume. Semantic occupancy predictions are generated using a two-layer MLP network, whereas a single convolutional layer predicts the occupancy status of pedestrians. Finally, pedestrian instances are grouped on the basis of both predictions to obtain instance and panoptic occupancy labels. (Best viewed in color.)

## Model Overview

Fig. 2 shows an overview of the proposed model, which takes a set of  $N$  posed multi-view images  $\{I_n\}_{n=1}^N$ . Here,  $I_n \in \mathbb{R}^{3 \times H \times W}$  with  $H$  height and  $W$  width of the image. The goal is to estimate dense 3D semantic, instance, and panoptic occupancies, represented by  $\mathbf{V}_{\text{sem}} \in \mathbb{R}^{C_s \times X \times Y \times Z}$ ,  $\mathbf{V}_{\text{pan}} \in \mathbb{R}^{C_p \times X \times Y \times Z}$ ,  $\mathbf{V}_{\text{ins}} \in \mathbb{R}^{N_p \times X \times Y \times Z}$ , respectively, along with 2D pedestrian occupancy  $\mathbf{P}_{\text{occ}} \in \mathbb{R}^{1 \times X \times Y}$ . In these expressions,  $C_s$  and  $C_p$  indicate the number of semantic and panoptic classes, respectively,  $N_p$  denotes the total number of pedestrians within the scene, while  $X$ ,  $Y$ , and  $Z$  denote the dimensions of the voxel grid that defines the scene.

**Image encoder.** Given that the input modality is in pixel space, we employ a ResNet (He et al. 2016) model to extract multi-scale 2D features. These features are instrumental in capturing diverse semantic and contextual information across different scales, which is particularly advantageous for detecting small pedestrians within spacious scenes. The multi-scale features are effectively fused with an FPN (Lin et al. 2017) network. The spatial resolution of the fused feature map is set to one-fourth of the original image size.

**View transformer.** In estimating the occupancy status and semantic properties of the entire scene, a crucial step involves translating the multi-view features onto the 3D space. Recent query-based feature transformation methods (Huang et al. 2023; Liu et al. 2023; Wei et al. 2023) are not suitable for our context due to various camera configurations and distinct scene dimensions. As an alternative, we adopt a non-parametric feature transformation approach (Rukhovich, Vorontsova, and Konushin 2022), which has been shown to achieve comparable performance (Harley et al. 2023). This method operates under the assumption of

a uniform depth distribution along the ray, implying that all voxels along a camera ray are filled with identical features corresponding to a single pixel in a 2D space.

To implement this process, we first construct a voxel grid  $\mathbf{V}_n \in \mathbb{R}^{3 \times X \times Y \times Z}$  that contains voxel centers and corresponds to the  $n$ -th view. The value of each center of the voxel  $\mathbf{p} = (x, y, z)^T$  is interpolated bilinearly from the image feature  $F_n$  at the projected pixel coordinates  $(u_n, v_n)$  of  $\mathbf{p}$ . Afterward, these feature samples are averaged from all views, considering only the valid number of 2D projections as Eqs. (1) to (3).

$$\mathbf{V}(\mathbf{p}) = \frac{1}{N} \sum_{n=1}^N M_n(\mathbf{p}) V_n(\mathbf{p}), \quad (1)$$

$$M_n(\mathbf{p}) = \begin{cases} 1 & \text{if } 0 \leq u_n \leq \frac{W}{4} \text{ and } 0 \leq v_n \leq \frac{H}{4}, \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

$$V_n(\mathbf{p}) = \text{interp}((u_n, v_n), F_n), \quad (3)$$

where  $(u_n, v_n)^T = K'_n R_n(p, 1)^T$ ,  $K'_n$  is the scaled intrinsic matrix,  $R_n$  is the extrinsic matrix, and  $M_n$  is a binary mask indicating whether the projected  $\mathbf{p}$  is inside the camera's frustum or not.

**Voxel encoder.** A 3D U-Net (Çiçek et al. 2016) based network is applied to process the voxel volume leveraging its effectiveness in handling volumetric data across various 3D tasks (Gan et al. 2023; Miao et al. 2023). It consists of a stack of three residual blocks, each with three downsampling 3D convolutional layers followed by three transposed 3D convolutional layers with  $C = 64$  channels. This allows the network to capture and refine hierarchical features within the voxel volume, enabling precise occupancy prediction and

semantic label assignment. The channel dimension is doubled in each downsampling layer and halved in each upsampling layer, following (Murez et al. 2020).

**Occupancy decoder.** As the voxel encoder already has a sufficient learnable capacity to accurately infer geometric and semantic information, we opt for a straightforward design for the 3D occupancy decoder. We utilize a two-layer MLP network to produce semantic occupancy prediction  $\mathbf{V}_{\text{sem}}$ , encapsulating the classification of each voxel into one of five distinct semantic classes ( $C_s = 5$ ).

**Location decoder.** We compress the feature volume by averaging along the vertical dimension ( $Z$ ), resulting in a bird’s eye view (BEV) feature map. A single convolutional layer is then applied to generate a 2D occupancy map  $\mathbf{P}_{\text{occ}}$ , which indicates the presence or absence of pedestrians within each cell of the BEV grid. We find that additional processing of the feature volume does not enhance location accuracy. This suggests that the pedestrians’ locations are primarily inferred from the occupancy status of 3D voxels.

**Pedestrian instance grouping.** Taking advantage of the simultaneous prediction of  $\mathbf{V}_{\text{sem}}$  and  $\mathbf{P}_{\text{occ}}$ , we group the semantic occupancy of pedestrian class,  $\mathbf{V}_{\text{sem}}^p$  into  $\mathbf{V}_{\text{ins}}$ . Let  $\mathbf{L} \in \mathbb{R}^{N_p \times 2} = \{\mathbf{l}_j\}_{j=1}^{N_p}$  denote the set of  $N_p$  pedestrian locations, obtained by filtering  $\mathbf{P}_{\text{occ}}$  with a pre-defined confidence threshold  $\tau$  and applying non-maximum suppression. An affinity matrix  $S \in \mathbb{R}^{XYZ \times N_p}$  can be defined where each element  $S_{i,j}$  represents the euclidean distance between the  $i$ -th voxel and  $j$ -th location:

$$S_{i,j} = \|\mathbf{v}_i - \mathbf{l}_j\|, \quad (4)$$

$$W_{i,j} = \begin{cases} 1 & \text{if } S_{i,j} < r, \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

where  $\mathbf{v}_i$  is the coordinate vector with  $x$  and  $y$  coordinates of  $\mathbf{p}$ .  $W_{i,j}$  is a binary mask to check if  $S_{i,j}$  is within a threshold  $r$  of 50 cm. Then, for each voxel  $i$ , we can find the index of the closest location  $j$ :

$$D_i = \arg \min_j S_{i,:N_p}, \quad (6)$$

and if  $\exists i$  such that  $W_{i,j} = 1$ , we set the label of  $i$ -th voxel in  $\mathbf{V}_{\text{ins}}$  to  $D_i$ . We obtain panoptic occupancy prediction  $\mathbf{V}_{\text{pan}} \in \mathbb{R}^{C_p \times X \times Y \times Z}$  after combining  $\mathbf{V}_{\text{ins}} \in \mathbb{R}^{N_p \times X \times Y \times Z}$  with three background/stuff classes of  $\mathbf{V}_{\text{sem}}$ , where  $C_p = C_b + N_p$  and  $C_b = 3$ .

## Loss

**3D occupancy loss.** Discretizing the vast expanse of the scene results in a large amount of free voxels. To address the dominance of the Free class and ensure balanced optimization across all classes, we use a weighted cross-entropy loss,  $\mathcal{L}_{\text{wce}}$ . The weights are derived from the inverse of the class frequency  $f_c$ , computed as  $w = \frac{1}{\log(f_c + \epsilon)}$ , where  $\epsilon \ll 1$ , following (Roldao, de Charette, and Verroust-Blondet 2020; Milioto et al. 2019). Furthermore, we incorporate the Lovász-Softmax loss  $\mathcal{L}_{\text{lovasz}}$  (Berman, Triki, and Blaschko

Scenes	Alley	Plaza	Field	Park	Facade	WildTrack
#cameras	6	3	4	8	7	7
size(m)	18×45	15×46	29×40	48×43	36×12	12×36
#peds.	60	50	40	100	40	20

Table 2: Comparison between MVP-Occ and WildTrack.

2018) to enhance segmentation quality for pedestrians. We also apply the scene-class affinity loss  $\mathcal{L}_{\text{affinity}}$  (Cao and de Charette 2022) to further improve overall segmentation quality by optimizing geometry and semantics separately. The final 3D occupancy loss is formulated as follows:

$$\mathcal{L}_{3D} = \lambda_{\text{wce}} \cdot \mathcal{L}_{\text{wce}} + \lambda_{\text{lovasz}} \cdot \mathcal{L}_{\text{lovasz}} + \lambda_{\text{affinity}} \cdot \mathcal{L}_{\text{affinity}}, \quad (7)$$

where  $\lambda_{\text{wce}} = 0.4$ ,  $\lambda_{\text{lovasz}} = 0.3$ , and  $\lambda_{\text{affinity}} = 0.3$  are hyperparameters to balance the loss components.

**2D occupancy loss.** Occupancy prediction in the BEV plane resembles a keypoint detection problem, aiming to produce a heatmap that reflects the likelihood of occupancy at each position on the ground plane. We utilize a Gaussian kernel to splat all ground-truth pedestrian locations onto a heatmap and compute the mean squared error (MSE) between the predicted and ground-truth occupancy maps.

**Total loss.** The final loss function is a composite of both 3D and 2D occupancy losses, expressed as:

$$\mathcal{L} = (1 - \lambda) \cdot \mathcal{L}_{3D} + \lambda \cdot \mathcal{L}_{2D}, \quad (8)$$

where  $\lambda = 0.3$  is a weighting coefficient and prioritizes  $\mathcal{L}_{3D}$  due to its higher complexity and difficulty in optimization.

## Datasets

This section presents MVP-Occ, our new synthetic dataset, and describes how we added new labels to the existing real-world WildTrack dataset (Chavdarova et al. 2018) for evaluation purposes.

### MVP-Occ Dataset

Our dataset is generated using the CARLA simulator (Dosovitskiy et al. 2017), originally developed for autonomous driving research (Sun et al. 2022; Weng et al. 2021; Sima et al. 2023; Yang et al. 2024). Due to its comprehensive functionalities and rapid development capabilities, we have adapted it to suit our urban traffic scenario.

**Scene generation and simulation.** CARLA offers diverse environments, from which we select specific scenes for our study. Similarly to Human-M3 (Fan et al. 2023) and GMVD (Vora et al. 2023), we provide multiple scenes to investigate cross-scene performance, named Alley, Plaza, Field, Park and Facade, as described in Table 2. Each scene is designed with unique characteristics to challenge the model’s adaptability and robustness. Pedestrian models from the built-in assets library are utilized to simulate foot traffic, including activities such as walking and running. Pedestrians are strategically spawned within and outside the scene

Method	MVDet*			SHOT*			GMVD†			MVFP†			OmniOcc (Ours)†		
	MODA	MODP	F1	MODA	MODP	F1	MODA	MODP	F1	MODA	MODP	F1	MODA	MODP	F1
Alley	93.4	87.9	96.6	95.0	<b>91.0</b>	97.5	91.9	88.3	95.8	<u>95.8</u>	<u>90.8</u>	<u>97.8</u>	<b>96.8</b>	88.9	<b>98.0</b>
Plaza	87.5	88.5	93.3	88.9	<u>89.5</u>	94.2	85.5	87.4	92.3	<u>89.3</u>	<b>90.7</b>	<u>94.4</u>	<b>92.4</b>	86.7	<b>96.1</b>
Field	82.1	85.6	90.5	<u>85.3</u>	<b>88.9</b>	<u>92.2</u>	82.4	85.2	90.6	84.9	<u>87.1</u>	<u>92.2</u>	<b>92.6</b>	86.9	<b>96.3</b>
Park	88.9	84.9	94.2	88.5	87.3	93.9	88.4	84.5	93.9	<u>91.2</u>	<u>87.5</u>	<u>95.5</u>	<b>93.4</b>	<b>87.9</b>	<b>96.6</b>
Facade	91.4	87.4	95.5	92.0	89.5	95.8	90.1	88.1	94.8	<u>92.7</u>	<b>90.9</b>	<u>96.2</u>	<b>93.7</b>	<u>89.8</u>	<b>96.8</b>
Avg.	88.7	86.7	94.0	89.9	<u>89.2</u>	94.7	87.7	86.7	93.5	<u>90.8</u>	<b>89.4</b>	<u>95.2</u>	<b>93.8</b>	88.0	<b>96.8</b>

Table 3: 2D pedestrian occupancy prediction under same-scene evaluation on MVP-Occ. Previous methods are supervised with ground-plane locations, while our method additionally utilizes voxel-level supervision. \* methods are scene-specific, while † methods can work with variable camera setups. The best and second best results are **bolded**, and underlined, respectively.

to emulate entry and exit patterns and navigate through the scene through AI control, which enables obstacle avoidance and autonomously generates navigation paths to predefined target locations. In addition, manual intervention is applied to manipulate certain pedestrians, ensuring their presence within or departure from the scene upon reaching their designated destinations. Moreover, we control the timing and speed of certain pedestrians to simulate group motions, such as synchronized walking or following behaviors, and maintaining flock formation.

**Rendering and data generation.** We render a comprehensive suite of sensors, including RGB, depth, semantic, and instance segmentation sensors, as depicted in Fig. 1. The camera setup configuration and the degree of overlapping fields of view are meticulously adjusted for each scene. The dataset is generated at 10 FPS for 2500 frames, with an image size of  $1920 \times 1080$ . Point clouds obtained from monocular views are fused and subsequently voxelized within the area of interest to get occupancy labels. Details are provided in the supplementary material. A 10 cm voxel size is used for discretization and is categorized into one of five unified classes: Free, Pedestrian, Ground, Wall, and Others with a supercategory of Background class for all scenes. Note that in our dataset, only one instance/thing class (Pedestrian) exists with three stuff classes (Ground, Wall, and Others) for the panoptic occupancy prediction task.

### WildTrack Dataset

WildTrack (Chavdarova et al. 2018) is a real-world dataset captured using seven cameras with significant overlapped fields of view. We aim to evaluate synthetic-to-real performance on this dataset by training on the proposed synthetic scenes. However, the dataset only provides pedestrian locations on the ground plane and lacks occupancy labels, which are difficult to obtain in real-world scenes. To address this, we assess 3D occupancy prediction performance by comparing the rendered occupancy predictions with 2D segmentation masks. Due to the absence of segmentation data in WildTrack, we manually annotated each image, generating 2D semantic, instance, and panoptic segmentation labels. This approach enables us to validate our model’s performance in predicting accurate 2D segmentations, serving as a proxy for its 3D occupancy prediction capabilities.

Metric	Semantic Occ.			Inst Occ.	Panoptic Occ.		
	mIoU	IoU <sub>Ped.</sub>	IoU <sub>Bg.</sub>	AP	PQ	SQ	RQ
Alley	93.6	68.5	99.9	95.8	95.3	96.0	99.1
Plaza	93.5	67.6	99.9	92.0	94.6	96.1	98.2
Field	93.6	68.2	99.9	93.6	94.5	96.5	97.7
Park	91.8	62.9	98.9	95.5	94.5	95.9	98.3
Facade	93.7	69.5	99.8	94.8	95.9	97.0	98.8

Table 4: 3D occupancy prediction under same-scene evaluation on MVP-Occ.




Figure 3: Qualitative results of 2D and 3D occupancy predictions under same-scene evaluation on the Park scene. (Best viewed in color.)

## Experiments

In this section, we first report on the evaluation metrics and training details. The subsequent sections provide extensive experiments on both 2D and 3D occupancy prediction benchmarks and lastly, ablation studies.

**Evaluation metrics.** We use commonly adopted metrics from multi-view detection methods to gauge the performance of 2D occupancy prediction. These include MODA (Multi-Object Detection Accuracy), MODP (Multi-Object Detection Precision), and F1 score, each offering insights into the model’s proficiency in pedestrian detection within the scene. For 3D semantic occupancy, we calculate voxel-level IoU metrics to quantify the overlap between predicted and ground-truth occupancies for each semantic class and

Method	GMVD		MVFP		OmniOcc (Ours)	
	MODA	F1	MODA	F1	MODA	F1
Alley	11.8	34.1	<u>11.9</u>	<u>36.5</u>	<b>33.3</b>	<b>63.2</b>
Plaza	7.5	15.5	<u>14.3</u>	<u>35.1</u>	<b>15.0</b>	<b>54.1</b>
Field	5.2	10.0	<u>24.8</u>	<u>41.9</u>	<b>41.2</b>	<b>60.4</b>
Park	1.8	24.5	<u>18.5</u>	<u>47.6</u>	<b>48.7</b>	<b>72.1</b>
Facade	46.3	71.9	<u>52.6</u>	<u>74.9</u>	<b>75.4</b>	<b>87.5</b>

Table 5: 2D occupancy prediction under synthetic-to-real evaluation on WildTrack. Our model outperforms previous methods in all scenes.

Metric	Semantic Occ.			Insta. Occ.		Panoptic Occ.		
	mIoU	IoU <sub>P</sub>	IoU <sub>G</sub>	AP	AP <sub>25</sub>	PQ	SQ	RQ
Alley	74.5	46.1	89.5	<u>25.8</u>	<u>55.0</u>	<u>54.5</u>	<u>75.8</u>	<u>65.2</u>
Plaza	34.1	19.0	36.6	3.7	12.8	43.5	73.2	50.8
Field	62.3	42.6	78.5	20.9	41.4	43.7	69.3	57.6
Park	<u>75.2</u>	<u>47.9</u>	89.7	23.4	47.8	54.1	75.6	64.7
Facade	<b>79.8</b>	<b>58.4</b>	<b>90.7</b>	<b>40.2</b>	<b>70.0</b>	<b>58.9</b>	<b>76.2</b>	<b>72.3</b>

Table 6: 3D occupancy prediction under synthetic-to-real evaluation on WildTrack. Note that the evaluation is done with multi-view segmentation data due to the lack of occupancy ground-truths.

mean IoU across all classes. We report AP (Average Precision) as mask AP using pycocotools (Lin et al. 2014) for 3D instance occupancy and PQ (Panoptic Quality) metric (Kirillov et al. 2019), which is composed of SQ (Segmentation Quality) and RQ (Recognition Quality) for 3D panoptic occupancy. Details are provided in the supplementary material.

**Training details.** Multi-view images are resized to 1280×720, and ResNet-18 is used as the backbone network. Each scene in our dataset is divided into training and testing sets, allocating 80% of the initial frames to the training split and reserving the remaining 20% for testing purposes, while WildTrack uses a 90/10 split. A voxel size of 10 cm for dis-

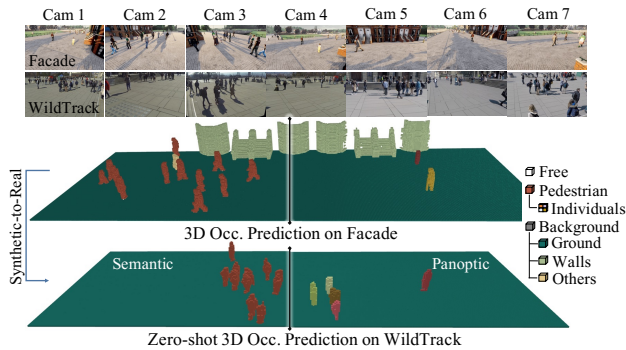


Figure 4: Qualitative results of synthetic-to-real transfer from Facade to WildTrack. (Best viewed in color.)

cretization results in variable grid dimensions, depending on the scene dimensions, as detailed in Table 2.  $\tau$  is set to 0.5 to filter out low-confidence detections in 2D occupancy prediction. The optimization process employs the AdamW optimizer with an initial learning rate of  $1 \times 10^{-3}$  and a decay rate of  $1 \times 10^{-2}$ . Training is conducted over 5 epochs, with a cosine learning rate scheduler dynamically adjusting the learning rate. The experiments were performed using four NVIDIA A100 GPUs, with a batch size of 4.

## Same-Scene Evaluation

In this evaluation scheme, all models undergo training in the train set and are evaluated on the test set of each scene in the proposed MVP-Occ dataset.

**2D occupancy prediction.** We compare our model with previous multi-view detection methods, MVDet (Hou, Zheng, and Gould 2020), SHOT (Song et al. 2021), GMVD (Vora et al. 2023), and MVFP (Aung et al. 2024), using the 2D occupancy prediction benchmark in Table 3. Our model outperforms all previous methods trained solely with 2D occupancy maps. However, despite an overall improvement in detection, our model yields a lower MODP score in most scenes, which measures the precision of true positive instances relative to the ground truth. The observed discrepancy can be attributed to the inherent misalignment between regressing each pedestrian’s location on the ground plane and the voxel-level whole-body representation.

**3D occupancy prediction.** In Table 4, we present the same-scene performance of the proposed model on the new 3D occupancy prediction benchmark, which includes semantic, instance, and panoptic occupancies. As anticipated, under the same-scene evaluation protocol, the model demonstrates exceptional performance, achieving nearly 99% IoU for the Background class and 62-68% IoU for the Pedestrian class, depending on the scene. Qualitative results for the Park scene are provided in Fig. 3, showcasing the model’s capability to effectively reconstruct the scene and accurately detect pedestrians when provided with ground-truth data for the scene of interest.

## Synthetic-to-Real Evaluation

In this evaluation scheme, all models are trained on each synthetic scene of the proposed MVP-Occ dataset and then evaluated on the real WildTrack scene.

**2D occupancy prediction.** We compare against previous cross-scene multi-view detection models, GMVD (Vora et al. 2023) and MVFP (Aung et al. 2024) in Table 5. Unlike the Facade scene, other scenes in our dataset exhibit significant differences from WildTrack, including variations in camera setup and scene characteristics. Despite these challenging scenarios, the proposed model demonstrates superior performance across all scenes compared to previous methods, which often struggle with disparate scenes. Notably, GMVD, which relies on ground plane detection, fails to transfer whereas MVFP, which utilizes 3D feature pulling, shows better generalization performance. Our OmniOcc model surpasses both by incorporating scene under-

$\mathcal{L}_{wce}$	$\mathcal{L}_{lovasz}$	$\mathcal{L}_{affinity}$	mIoU	IoU <sub>Ped.</sub>	IoU <sub>Others</sub>
-	-	-	85.3	67.1	59.9
✓	-	-	89.4	67.2	79.8
✓	✓	-	<u>93.4</u>	<u>67.9</u>	<u>99.7</u>
✓	-	✓	<u>93.4</u>	67.4	<b>99.9</b>
✓	✓	✓	<b>93.7</b>	<b>69.5</b>	<u>99.7</u>

Table 7: Ablation study on different loss functions.

standing, which helps mitigate these issues. The Park scene, with its diverse pedestrian distribution further enhances our model’s ability to focus on pedestrians, resulting in improved pedestrian detection performance. All models perform well when using the Facade scene, which was designed to closely resemble WildTrack’s conditions.

**3D occupancy prediction.** We justify the performance of the 3D occupancy prediction by evaluating the rendered 2D segmentation masks with manually labeled ground-truth 2D segmentation data. Unlike the scenes in our dataset, the WildTrack dataset has only three semantic classes: Free, Pedestrian, and Ground. Based on Table 6, the 3D semantic occupancy prediction yields conclusions similar to the 2D occupancy prediction. However, for 3D instance and panoptic occupancy predictions, we observe better segmentation quality for individual pedestrians in the Alley scene compared to the Park scene. The best performance is consistently achieved by the Facade scene, which continues to deliver outstanding results across all occupancy prediction tasks. The qualitative results in Fig. 4 further demonstrate the impressive generalization capabilities of the Facade scene in WildTrack. The segmentation quality of individual pedestrians is notably more accurate for the Facade scene, where the model was trained. Nevertheless, addressing the challenges associated with scene understanding in dissimilar scenes requires further investigation and methodological refinements.

### Ablation Study

For all experiments under the ablation study, we utilize the Facade scene unless otherwise specified.

**Loss functions.** In Table 7, we conduct an ablation study examining various combinations of loss functions using a 3D semantic occupancy prediction benchmark. The weighted cross-entropy loss  $\mathcal{L}_{wce}$  significantly enhances segmentation quality for less dominant classes, such as an advertisement board in the Others class. Direct optimization of the Jaccard index using  $\mathcal{L}_{lovasz}$  or incorporating  $\mathcal{L}_{affinity}$  proves to be effective in refining the segmentation quality for the Pedestrian and Others classes, with a trade-off between the two classes. Additionally, we find that optimizing the geometric and semantic aspects separately requires more time to converge. Ultimately, the best performance is achieved by combining all of the aforementioned loss functions.

**Semantic scene understanding.** We conduct an ablation study to assess the impact of semantic scene understanding (SSU) on 2D occupancy prediction and 3D instance occupancy prediction benchmarks. By removing the Background

SSU	2D Pedestrian Occ.			3D Instance Occ.		
	MODA	MODP	F1	AP	AP <sub>25</sub>	AP <sub>50</sub>
-	66.5	77.1	81.6	38.2	66.2	39.0
✓	<b>75.4</b>	<b>80.3</b>	<b>87.5</b>	<b>40.2</b>	<b>70.0</b>	<b>39.8</b>

Table 8: Impact of semantic scene understanding (SSU) under synthetic-to-real evaluation on WildTrack.

PIG	Testing	Park		WildTrack		
	Training	mIoU	IoU <sub>Ped.</sub>	mIoU	IoU <sub>Ped.</sub>	IoU <sub>Free</sub>
-	Park	<b>91.8</b>	<b>62.9</b>	75.1	<b>47.9</b>	87.8
✓	Park	<b>91.8</b>	<b>62.9</b>	<b>75.2</b>	<b>47.9</b>	<b>87.9</b>

Table 9: Impact of pedestrian instance grouping (PIG) on 3D semantic occupancy prediction.

class and retaining only the voxel occupancy state – indicating whether a voxel is occupied (i.e., containing a Pedestrian) or Free – the results presented in Table 8 demonstrate that integrating SSU significantly enhances performance in both tasks. This experiment highlights the critical importance of contextual scene understanding, particularly in cross-scene scenarios where conditions vary.

**Pedestrian instance grouping (PIG).** PIG is crucial for our model to predict instance and panoptic occupancy accurately. It also improves qualitative results for the prediction of semantic occupancy, as visualized in Fig. 3, where it eliminates floating voxels in the absence of pedestrian instances. However, these improvements are not reflected in the quantitative scores (Table 9) under the same-scene evaluation scheme, due to the large number of voxels compared to the number of uncertain voxels. Even so, the synthetic-to-real results indicate a slight improvement in the Free class, where noises in the free space are significantly reduced.

## Conclusions

This paper proposes a pioneering synthetic dataset specifically crafted for dense pedestrian scenarios in expansive multi-view environments. Our dataset is the first of its kind, offering extensive 2D and 3D annotations that are particularly suited for urban surveillance contexts, thereby creating new opportunities for research in this area. Through meticulous analysis, we present a robust baseline model that effectively addresses the challenges of multi-view pedestrian occupancy prediction. Our findings underscore the importance of occupancy prediction and semantic scene understanding in boosting pedestrian detection performance, enabling our model to achieve state-of-the-art accuracy. Furthermore, our model accurately reconstructs scenes with dynamic pedestrians when provided with 3D ground-truth data. Importantly, we demonstrate excellent synthetic-to-real occupancy prediction performance through scene replication, paving the way for applications in real-world scenarios.

## Acknowledgments

This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT)(RS-2023-00227592, Development of 3D Object Identification Technology Robust to Viewpoint Changes, 50%), the Korea Institute of Police Technology (KIPoT) funded by the Korean National Police Agency & Ministry of the Interior and Safety (RS-2024-00405100, 25%), and the Korea Institute of Science and Technology (KIST) Institutional Program (Project No. 2E33001, 25%).

## References

- Aung, S.; Park, H.; Jung, H.; and Cho, J. 2024. Enhancing Multi-View Pedestrian Detection Through Generalized 3D Feature Pulling. In *WACV*, 1196–1205.
- Berman, M.; Triki, A. R.; and Blaschko, M. B. 2018. The lovász-softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4413–4421.
- Cao, A.-Q.; and de Charette, R. 2022. Monoscene: Monocular 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3991–4001.
- Chavdarova, T.; Baqué, P.; Bouquet, S.; Maksai, A.; Jose, C.; Bagautdinov, T.; Lettry, L.; Fua, P.; Van Gool, L.; and Fleuret, F. 2018. Wildtrack: A multi-camera hd dataset for dense unscripted pedestrian detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5030–5039.
- Chen, X.; Lin, K.-Y.; Qian, C.; Zeng, G.; and Li, H. 2020. 3d sketch-aware semantic scene completion via semi-supervised structure prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4193–4202.
- Çiçek, Ö.; Abdulkadir, A.; Lienkamp, S. S.; Brox, T.; and Ronneberger, O. 2016. 3D U-Net: learning dense volumetric segmentation from sparse annotation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part II 19*, 424–432. Springer.
- Dong, J.; Jiang, W.; Huang, Q.; Bao, H.; and Zhou, X. 2019. Fast and robust multi-person 3d pose estimation from multiple views. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7792–7801.
- Dosovitskiy, A.; Ros, G.; Codevilla, F.; Lopez, A.; and Koltun, V. 2017. CARLA: An open urban driving simulator. In *Conference on robot learning*, 1–16. PMLR.
- Fan, B.; Wang, S.; Zheng, W.; Feng, J.; and Zhou, J. 2023. Human-M3: A Multi-view Multi-modal Dataset for 3D Human Pose Estimation in Outdoor Scenes. *arXiv preprint arXiv:2308.00628*.
- Fang, M.; and Li, Z. 2023. Occupancy Dataset for nuScenes. <https://github.com/FANG-MING/occupancy-for-nuscenes>.
- Gan, W.; Mo, N.; Xu, H.; and Yokoya, N. 2023. A Simple Attempt for 3D Occupancy Estimation in Autonomous Driving. *arXiv preprint arXiv:2303.10076*.
- Gan, Y.; Han, R.; Yin, L.; Feng, W.; and Wang, S. 2021. Self-supervised multi-view multi-human association and tracking. In *Proceedings of the 29th ACM International Conference on Multimedia*, 282–290.
- Harley, A. W.; Fang, Z.; Li, J.; Ambrus, R.; and Fragkiadaki, K. 2023. Simple-bev: What really matters for multi-sensor bev perception? In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2759–2765. IEEE.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hou, Y.; and Zheng, L. 2021. Multiview detection with shadow transformer (and view-coherent data augmentation). In *Proceedings of the 29th ACM International Conference on Multimedia*, 1673–1682.
- Hou, Y.; Zheng, L.; and Gould, S. 2020. Multiview Detection with Feature Perspective Transformation. In *ECCV*.
- Huang, Y.; Zheng, W.; Zhang, Y.; Zhou, J.; and Lu, J. 2023. Tri-perspective view for vision-based 3d semantic occupancy prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9223–9232.
- Jeon, Y.; Tran, D. Q.; Park, M.; and Park, S. 2023. Leveraging Future Trajectory Prediction for Multi-Camera People Tracking. In *CVPR*, 5398–5407.
- Joo, H.; Liu, H.; Tan, L.; Gui, L.; Nabbe, B.; Matthews, I.; Kanade, T.; Nobuhara, S.; and Sheikh, Y. 2015. Panoptic Studio: A Massively Multiview System for Social Motion Capture. In *ICCV*.
- Kirillov, A.; He, K.; Girshick, R.; Rother, C.; and Dollár, P. 2019. Panoptic segmentation. In *CVPR*, 9404–9413.
- Li, Z.; Yu, Z.; Austin, D.; Fang, M.; Lan, S.; Kautz, J.; and Alvarez, J. M. 2023. Fb-occ: 3d occupancy prediction based on forward-backward view transformation. *arXiv preprint arXiv:2307.01492*.
- Lin, J.; and Lee, G. H. 2021. Multi-view multi-person 3d pose estimation with plane sweep stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11886–11895.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*, 740–755. Springer.
- Liu, H.; Wang, H.; Chen, Y.; Yang, Z.; Zeng, J.; Chen, L.; and Wang, L. 2023. Fully Sparse 3D Panoptic Occupancy Prediction. *arXiv preprint arXiv:2312.17118*.
- Luna, E.; SanMiguel, J. C.; Martínez, J. M.; and Carballeira, P. 2022. Graph Neural Networks for Cross-Camera Data

- Association. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(2): 589–601.
- Miao, R.; Liu, W.; Chen, M.; Gong, Z.; Xu, W.; Hu, C.; and Zhou, S. 2023. Occdepth: A depth-aware method for 3d semantic scene completion. *arXiv preprint arXiv:2302.13540*.
- Milioto, A.; Vizzo, I.; Behley, J.; and Stachniss, C. 2019. Rangenet++: Fast and accurate lidar semantic segmentation. In *2019 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, 4213–4220. IEEE.
- Murez, Z.; Van As, T.; Bartolozzi, J.; Sinha, A.; Badrinarayanan, V.; and Rabinovich, A. 2020. Atlas: End-to-end 3d scene reconstruction from posed images. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*, 414–431. Springer.
- Nguyen, Q. Q.-V.; Le, H. D.-A.; Chau, T. T.-T.; Luu, D. T.; Chung, N. M.; and Ha, S. V.-U. 2023. Multi-camera people tracking with mixture of realistic and synthetic knowledge. In *CVPR*, 5495–5505.
- Qiu, R.; Xu, M.; Yan, Y.; Smith, J. S.; and Yang, X. 2022. 3d random occlusion and multi-layer projection for deep multi-camera pedestrian localization. In *ECCV*, 695–710. Springer.
- Roldao, L.; de Charette, R.; and Verroust-Blondet, A. 2020. Lmscnet: Lightweight multiscale 3d semantic completion. In *2020 International Conference on 3D Vision (3DV)*, 111–119. IEEE.
- Rukhovich, D.; Vorontsova, A.; and Konushin, A. 2022. Imvoxelnet: Image to voxels projection for monocular and multi-view general-purpose 3d object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2397–2406.
- Sima, C.; Renz, K.; Chitta, K.; Chen, L.; Zhang, H.; Xie, C.; Luo, P.; Geiger, A.; and Li, H. 2023. Drivelm: Driving with graph visual question answering. *arXiv preprint arXiv:2312.14150*.
- Song, L.; Wu, J.; Yang, M.; Zhang, Q.; Li, Y.; and Yuan, J. 2021. Stacked homography transformations for multi-view pedestrian detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6049–6057.
- Song, S.; Yu, F.; Zeng, A.; Chang, A. X.; Savva, M.; and Funkhouser, T. 2017. Semantic scene completion from a single depth image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1746–1754.
- Sun, T.; Segu, M.; Postels, J.; Wang, Y.; Van Gool, L.; Schiele, B.; Tombari, F.; and Yu, F. 2022. SHIFT: A Synthetic Driving Dataset for Continuous Multi-Task Domain Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 21371–21382.
- Tang, J.; Chen, X.; Wang, J.; and Zeng, G. 2022. Not all voxels are equal: Semantic scene completion from the point-voxel perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 2352–2360.
- Tesla. 2022. Tesla AI Day. [https://www.youtube.com/watch?v=ODSJsviD\\_SU](https://www.youtube.com/watch?v=ODSJsviD_SU).
- Tian, X.; Jiang, T.; Yun, L.; Wang, Y.; Wang, Y.; and Zhao, H. 2023. Occ3d: A large-scale 3d occupancy prediction benchmark for autonomous driving. *arXiv preprint arXiv:2304.14365*.
- Tong, W.; Sima, C.; Wang, T.; Chen, L.; Wu, S.; Deng, H.; Gu, Y.; Lu, L.; Luo, P.; Lin, D.; et al. 2023a. Scene as occupancy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8406–8415.
- Tong, W.; Sima, C.; Wang, T.; Chen, L.; Wu, S.; Deng, H.; Gu, Y.; Lu, L.; Luo, P.; Lin, D.; et al. 2023b. Scene as occupancy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8406–8415.
- Tu, H.; Wang, C.; and Zeng, W. 2020. VoxelPose: Towards Multi-Camera 3D Human Pose Estimation in Wild Environment. In *European Conference on Computer Vision (ECCV)*.
- Vora, J.; Dutta, S.; Jain, K.; Karthik, S.; and Gandhi, V. 2023. Bringing Generalization to Deep Multi-View Pedestrian Detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 110–119.
- Wei, Y.; Zhao, L.; Zheng, W.; Zhu, Z.; Zhou, J.; and Lu, J. 2023. Surroundocc: Multi-camera 3d occupancy prediction for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 21729–21740.
- Weng, X.; Man, Y.; Park, J.; Yuan, Y.; Cheng, D.; O’Toole, M.; and Kitani, K. 2021. All-In-One Drive: A Large-Scale Comprehensive Perception Dataset with High-Density Long-Range Point Clouds. *arXiv*.
- Wu, S.-C.; Tateno, K.; Navab, N.; and Tombari, F. 2020. Sc-fusion: Real-time incremental scene reconstruction with semantic completion. In *2020 International Conference on 3D Vision (3DV)*, 801–810. IEEE.
- Yang, J.; Desai, K.; Packer, C.; Bhatia, H.; Rhinehart, N.; McAllister, R.; and Gonzalez, J. 2024. CARFF: Conditional Auto-encoded Radiance Field for 3D Scene Forecasting. *arXiv preprint arXiv:2401.18075*.
- Yao, J.; Li, C.; Sun, K.; Cai, Y.; Li, H.; Ouyang, W.; and Li, H. 2023. Ndc-scene: Boost monocular 3d semantic scene completion in normalized device coordinates space. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 9421–9431. IEEE Computer Society.
- Zhang, J.; Cai, Y.; Yan, S.; Feng, J.; et al. 2021. Direct multi-view multi-person 3d pose estimation. *Advances in Neural Information Processing Systems*, 34: 13153–13164.
- Zhang, Q.; and Chan, A. B. 2019. Wide-area crowd counting via ground-plane density maps and multi-view fusion cnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8297–8306.
- Zhang, Q.; Lin, W.; and Chan, A. B. 2021. Cross-view cross-scene multi-view crowd counting. In *CVPR*, 557–567.
- Zhang, Y.; Zhu, Z.; and Du, D. 2023. OccFormer: Dual-path Transformer for Vision-based 3D Semantic Occupancy Prediction. *arXiv preprint arXiv:2304.05316*.