SSCBench: A Large-Scale 3D Semantic Scene Completion Benchmark for Autonomous Driving

Yiming Li^{1,*}, Sihang Li^{1,*}, Xinhao Liu^{1,*}, Moonjun Gong^{1,*}, Kenan Li¹, Nuo Chen¹, Zijun Wang¹, Zhiheng Li¹, Tao Jiang², Fisher Yu³, Yue Wang⁴, Hang Zhao², Zhiding Yu⁴, Chen Feng^{1,†}

Abstract—Monocular scene understanding is a foundational component of autonomous systems. Within the spectrum of monocular perception topics, one crucial and useful task for holistic 3D scene understanding is semantic scene completion (SSC), which jointly completes semantic information and geometric details from RGB input. However, progress in SSC, particularly in large-scale street views, is hindered by the scarcity of high-quality datasets. To address this issue, we introduce SSCBench, a comprehensive benchmark that integrates scenes from widely used automotive datasets (e.g., KITTI-360, nuScenes, and Waymo). SSCBench follows an established setup and format in the community, facilitating the easy exploration of SSC methods in various street views. We benchmark models using monocular, trinocular, and point cloud input to assess the performance gap resulting from sensor coverage and modality. Moreover, we have unified semantic labels across diverse datasets to simplify cross-domain generalization testing. We commit to including more datasets and SSC models to drive further advancements in this field. Our data and code are available at https://github.com/ai4ce/SSCBench.

I. INTRODUCTION

Understanding 3D scenes from a single RGB image is crucial and meaningful in vision and robotics, with monocular perception tasks like object detection [1], tracking [2], and depth estimation [3] garnering significant attention. The emerging field of 3D semantic scene completion (SSC) [4] seeks to jointly infer complete 3D semantics and geometry from a sparse and partial observation (*e.g.*, an RGB image). The resulting volumetric representation seamlessly integrates occupancy and semantic information, facilitating robotic scene understanding and planning capabilities in street views.

One critical challenge in SSC is to generate accurate ground truth labels, especially in street views. Given the current limitations of 3D sensing technology, achieving a perfectly comprehensive 3D representation is impossible. The pioneering SemanticKITTI benchmark [5] proposes to leverage the temporal information through the aggregation of different LiDAR sweeps, which can effectively reveal previously occluded 3D surfaces. Meanwhile, it excludes 3D voxels not observed from all viewpoints during driving. Consequently, SemanticKITTI provides relatively comprehensive and accurate ground truth labels for SSC tasks.

While SemanticKITTI is a valuable resource for learning sparse-to-dense mapping, its limited scale and diversity

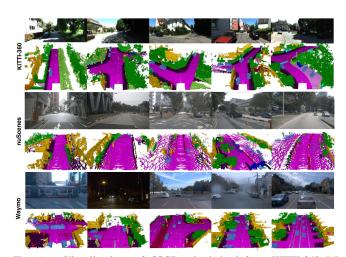


Fig. 1. **Visualizations of SSCBench** derived from KITTI-360 [6], nuScenes [7], and Waymo [8]. We showcase accurate SSC ground truth in a variety of street views.

impede the development of more powerful and generalizable SSC models. Another significant limitation of SemanticKITTI is the omission of dynamic objects during ground truth generation, resulting in inaccurate labels. Hence, there is an urgent need for a large-scale SSC dataset with reliable ground truth to advance learning-based scene understanding in street views.

To this end, we introduce SSCBench, a large-scale benchmark comprising diverse street views sourced from wellestablished automotive datasets, including KITTI-360 [6], nuScenes [7], and Waymo [8], as illustrated in Fig. 1. To enhance label accuracy, we utilize the 3D bounding box labels provided in these datasets to synchronize measurements of dynamic objects. Our key features include (a) accessibility: we provide datasets in a format compatible with SemanticKITTI, facilitating seamless usage within the community; (b) large *scale*: we offer an extensive dataset with \sim 8 times more frames than SemanticKITTI, encompassing diverse geographic locations across six cities; (c) comprehensiveness: we mainly focus on SSC methods with monocular input. Additionally, we utilize trinocular input to compare the singleview and panoramic-view methods and use point cloud input to show the gap between camera-based and LiDAR-based methods. Furthermore, we have unified semantic labels across different datasets in SSCBench, facilitating cross-domain generalization experiments. We plan to continually incorporate novel automotive datasets and SSC algorithms to drive further advancements in the field.

¹New York University, ²Tsinghua University, ³ETH Zurich, ⁴NVIDIA

^{*}Equal contribution. Yiming Li is supported by NVIDIA Fellowship.

[†]The corresponding author is Chen Feng cfeng@nyu.edu. The work was supported by NSF 2238968 and 2322242 grants; and in part through the NYU IT High Performance Computing resources, services, and staff expertise.

TABLE I

OVERVIEW OF WIDELY-USED AUTONOMOUS DRIVING DATASETS WITH MULTIMODAL SENSORS. C DENOTES CAMERA AND L DENOTES LIDAR. MOST DATASETS PROVIDE BOUNDING ANNOTATIONS FOR 3D DETECTION, YET ONLY A FEW OF THEM PROVIDE SEMANTIC LABELS FOR THE LIDAR POINT CLOUD DUE TO THE HIGH COST. NOTE THAT APOLLOSCAPE [34] ONLY PROVIDES 3D SEMANTIC LABELS FOR THE STATIC ENVIRONMENTS.

Datasets	Year	Sensors	Annotations	# Fr. with Pts Ann.	Sequential
KITTI [13] SemanticKITTI [5]	CVPR 2012 ICCV 2019	C&L C&L	3D Bbox 3D Pts.	N.A. 20K	×
nuScenes [7] Panoptic nuScenes [35]	CVPR 2019 RA-L 2022	C&L C&L	3D Bbox 3D Pts.	N.A. 40K	1
Waymo [8]	CVPR 2020	C&L	3D Bbox&Pts.	230K	/
KITTI-360 [6]	T-PAMI 2022	C&L	3D Bbox&Pts.	100K	✓
ApolloScape [34]	T-PAMI 2019	C&L	3D Bbox&Pts.	N.A.	/
Argoverse [36]	CVPR 2019	C&L	3D Bbox	N.A.	/
ONCE [37]	NeurIPS 2021	C&L	3D Bbox	N.A.	/
Lyft Level 5 [38]	2019	C&L	3D Bbox	N.A.	· /
A*3D [39]	ICRA 2020	C&L	3D Bbox	N.A.	/
A2D2 [40]	2020	C&L	3D Bbox	N.A.	X

II. RELATED WORKS

Monocular Perception and 3D Semantic Scene Completion. The simplicity, efficiency, affordability, and accessibility of monocular cameras have made monocular perception a focal point of attention in the vision and robotics community. This has resulted in extensive research into various tasks, including depth estimation [9], 3D object detection and tracking [10], as well as localization and mapping [11]. Song et al. [12] introduce the concept of monocular 3D semantic scene completion (SSC), which seeks to reconstruct and complete the semantics as well as geometry within a 3D volume from a single depth image. However, they only consider the bounded indoor scenarios due to the lack of outdoor datasets. Behley et al. [5] build the first outdoor dataset based on KITTI [13] for 3D semantic scene completion in street views. Existing approaches usually depend on 3D inputs, such as LiDAR point clouds [14, 15, 16], while recent monocular vision-based solutions also emerge [17, 18]. However, the development of outdoor SSC is hindered by the lack of datasets, with SemanticKITTI [5] being the only dataset supporting SSC in street views. Building diverse datasets is imperative to unlock the full potential of SSC for autonomous systems.

Point Cloud Segmentation in Street Views. 3D LiDAR segmentation aims to assign point-wise semantic labels for point clouds, including a range of specific tasks, like LiDAR semantic [19, 20, 21, 22], panoptic [23, 24, 25], and 4D panoptic segmentation [26, 27]. In this field, point-based methods, stemming from PointNet++ [28], perform well on small synthetic point cloud [29] rather than sparse LiDAR point cloud, with sampling and gathering disordered neighbors. Voxel-based approaches [30, 31, 32, 33] process point clouds by initially partitioning 3D space into voxels through Cartesian coordinates. Note that 3D LiDAR segmentation aims to understand the scenes based on raw LiDAR scans, while 3D semantic scene completion includes the completion of occluded areas, with the input of camera or LiDAR.

Autonomous Driving Dataset and Benchmark. Autonomous driving research thrives on high-quality datasets,

which serve as the lifeblood for training and evaluating perception [7], prediction [41], and planning algorithms [42]. In 2012, the pioneering KITTI dataset sparked a revolution in autonomous driving research, unlocking a multitude of tasks including object detection, tracking, mapping, and optical/depth estimation [13, 43, 44, 45, 46]. Since then, the research community has embraced the challenge, giving rise to a wealth of datasets. These datasets push the boundaries of autonomous driving research by addressing challenges posed by multimodal fusion [7], multi-tasking learning [47, 6], adverse weather [48], collaborative driving [49, 50, 51], repeated driving [52], and dense traffic scenarios [39, 53], etc. There are several impactful and widely-used driving datasets such as KITTI-360 [6], nuScenes [7], and Waymo [8]. They provide LiDAR and camera recordings as well as point cloud semantics and bounding annotations, as summarized in Tab. I. Therefore, we can create accurate ground truth labels for SSC by aggregating multiple semantic point clouds and leveraging the 3D boxes to handle dynamic objects.

Occ3D and OpenOccupancy. We compare SSCBench with the concurrent relevant work Occ3D [54]. The differences lie in: (a) setup: Occ3D uses surrounding-view images as input, and only considers the reconstruction of 3D voxels visible to the camera. SSCBench considers a more challenging yet meaningful setup (also a well-established one): how to reconstruct and complete 3D semantics in both visible and occluded areas only with monocular visual input. This task requires reasoning about temporal information and 3D geometric relationships to get rid of the limited field of view; (b) scale: SSCBench provides more datasets than Occ3D and plans to add more due to the abundance of monocular driving recordings; (c) accessibility: we inherit the widelyused setup from the pioneer KITTI, thus making SSCBench more accessible to the community; (d) comprehensiveness: we benchmark SSC methods with monocular, trinocular, and point cloud input and provide unified labels for crossdomain generalization tests. Another relevant benchmark, OpenOccupancy [55], exhibits similar differences, notably its exclusive use of the nuScenes dataset [7], which results in a limitation of diversity.

III. DATASET CURATION

A. Revisit of SemanticKITTI

SemanticKITTI [5] extends the odometry dataset of the KITTI vision benchmark [13] by providing point-wise semantic annotations for 22 driving sequences in Karlsruhe, Germany. SemanticKITTI not only supports 3D semantic segmentation but also serves as the first outdoor SSC benchmark. Similar to indoor SSC [12], SemanticKITTI uses voxelized 3D representation widely employed in robotics such as occupancy grid mapping [56]. SemanticKITTI generates ground truth labels via voxelization of a dense semantic point cloud given by rigid registration of multiple LiDAR scans.

Semantic KITTI has two limitations. First, rigid registration with sensor poses can only handle measurements for static environments, resulting in traces produced by dynamic objects such as moving cars as shown in Fig. 2, which can confuse the

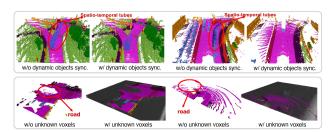


Fig. 2. **Top row: dynamic objects synchronization.** Two examples on nuScenes [7] are shown. Spatio-temporal tubes are introduced without handling dynamic objects, damaging the accuracy of labels. **Bottom row: unknown voxels exclusion.** Voxels are marked as unknown (denoted by grey color) when they are occluded or remain unprobed by the LiDAR.

3D representation learning [16]. Secondly, it is constrained by the limited scale and lack of diverse geographical coverage. The data collection is confined to a single city, resulting in training, validation, and test sets composed of 3,834, 815, and 3,992 frames, respectively, amounting to a total of 8,641 frames. However, this falls short of the large-scale benchmark necessary for comprehensive evaluation and generalization.

B. SSCBench

We aim to establish a large-scale SSC benchmark in street views that facilitates the training of robust and generalizable SSC models. To achieve this, we harness well-established and widely-used datasets and integrate them into a unified setup and format. Overall, our SSCBench, consisting of three subsets, includes 38,562 frames for training, 15,798 frames for validation, and 12,553 frames for testing respectively, amounting totally to 66,913 frames (\sim 67K), which greatly exceeds the scale of SemanticKITTI mentioned above by \sim 7.7 times. In the following, we introduce three carefully designed datasets, all based on existing data sources, that collectively contribute to our SSCBench.

SSCBench-KITTI-360. KITTI-360 [6] represents a significant advancement in autonomous driving research, building upon the renowned KITTI dataset [13]. It introduces an enriched data collection framework with diverse sensor modalities as well as panoramic viewpoints (a perspective stereo camera plus a pair of fisheye cameras) and provides comprehensive annotations including consistent 2D and 3D semantic instance labels as well as 3D bounding primitives. The dense and coherent labels not only support established tasks such as segmentation and detection but also enable novel applications like semantic SLAM [57] and novel view synthesis [58]. While KITTI-360 includes point cloud-based semantic scene completion, the prevalent methodology for SSC remains centered around voxelized representations [4], which exhibit broader applicability in robotics.

Remark. KITTI-360 covers a driving distance of 73.7km, comprising 300K images and 80K laser scans. While adhering to KITTI's forward-facing camera setup, it offers greater geographical diversity and demonstrates minimal trajectory overlap with KITTI. Leveraging the open-source training and validation set, we build SSCBench-KITTI-360 consisting of 9 long sequences. To reduce redundancy, we sample every 5 frames following the SemanticKITTI SSC benchmark. The

training set includes 8,487 frames from scenes 00, 02-05, 07, and 10, while the validation set comprises 1,812 frames from scene 06. The testing set comprises 2,566 frames from scene 09. In total, the dataset contains 12,865 (\sim 13K) frames, surpassing the scale of SemanticKITTI by \sim 1.5 times.

SSCBench-nuScenes. Unlike KITTI's forward-facing camera setup, nuScenes [7] captures a complete 360-degree view around the ego vehicle. It provides a diverse range of multimodal sensory data, including camera images, LiDAR point clouds, and radar data, gathered in Boston and Singapore. nuScenes offers meticulous annotations for complex urban driving scenarios, including diverse weather conditions, construction zones, and varying illumination. Later on, panoptic nuScenes [35] extends the original nuScenes dataset with semantic and instance labels. With comprehensive metrics and evaluation protocols, nuScenes is widely employed in autonomous driving research [59, 42, 60, 61].

Remark. The nuScenes dataset consists of 1K 20-second scenes with labels provided only for the training and validation set, totaling 850 scenes. From the available 850 scenes, we allocate 500 scenes for training, 200 for validation, and 150 for testing. This distribution results in 20,064 frames for training, 8,050 frames for validation, and 5,949 for testing, totaling 34,078 frames (~34K). This scale is approximately four times that of SemanticKITTI. As nuScenes only provides annotations for keyframes at a frequency of 2Hz, there is no downsampling in SSCBench-nuScenes.

SSCBench-Waymo. The Waymo dataset [8], collected from various locations in the US, offers a large-scale collection of multimodal sensor recordings. Waymo provides 5 cameras with a combined horizontal field of view of \sim 230 degrees, slightly smaller than nuScenes. The data is captured in diverse conditions across multiple cities, including San Francisco, Phoenix, and Mountain View, ensuring broad geographical coverage within each city. It consists of 1000 scenes for training and validation, as well as 150 scenes for testing, with each scene spanning 20 seconds.

Remark. To construct SSCBench-Waymo, we utilize the open-source training and validation scenes and redistribute them into sets of 500, 298, and 202 scenes for training, validation, and testing, respectively. We use only the annotated keyframes, which results in a training set of 14,943 frames, a validation set of 8,778 frames, and a test set of 5,946 frames, totaling 29,667 frames (\sim 30K).

C. Construction Pipeline

Prerequisites. To establish SSCBench, a driving dataset with multimodal recordings is required for LiDAR-based or camera-based SSC. The dataset should include sequentially collected 3D LiDAR point clouds with accurate sensor poses for geometry completion, per-point semantic annotations for semantic scene understanding, and 3D bounding annotations to handle dynamic instances.

Aggregation of Point Clouds. To generate a complete representation, our approach involves superimposing an extensive set of laser scans within a defined region in front of the vehicle. In short sequences like nuScenes and

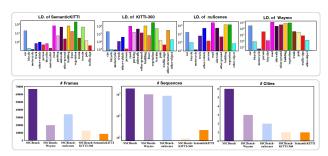


Fig. 3. **Statistical analysis.** Top row: Label Distribution (LD.) of different datasets. Bottom row: scale comparisons between SSCBench and SemanticKITTI.

Waymo, we utilize future scans with measurements from the corresponding region to create a dense semantic point cloud. In long sequences like KITTI-360, which feature multiple loop closures, we incorporate all spatial neighboring point clouds in addition to the temporal neighborhood. Accurate sensor poses, provided by advanced SLAM systems [62], greatly facilitate the aggregation of point clouds for the static environment. As for dynamic objects, we avoid the spatial-temporal tubes by synchronization. We utilize the instance label to transform dynamic objects to their spatial alignment within the current frame. As shown in Fig. 2, the spatial-temporal tubes are removed and objects have a denser shape.

Voxelization of Aggregated Point Clouds. Voxelization is to discretize a continuous 3D space into a regular grid structure composed of volumetric elements called voxels, enabling the conversion of unstructured data into a structured format that can be efficiently processed by convolutional neural networks (CNNs). Voxelization introduces a tradeoff between spatial resolution and memory consumption and offers a flexible and scalable representation for 3D perception [63, 64, 18]. For easy integration, SSCBench adheres to SemanticKITTI's setup, with a volume extending 51.2m ahead, 25.6m on each side, and 6.4m in height. Voxel resolution is 0.2m, resulting in a 256×256×32 voxel volume. Labels for each voxel are determined by majority voting among labeled points within it, while empty voxels are marked accordingly if no points are present.

Exclusion of Unknown Voxels. Capturing complete 3D outdoor scenes is nearly impossible without ubiquitous scene sensing. While it is possible to utilize spatial or prior knowledge-based inference, we intend to ensure the fidelity of ground truth by minimizing errors originating from these steps. Hence, we only consider *visible and probed* voxels from all viewpoints during training and evaluation. Specifically, we first employ ray tracing from different perspectives to identify and remove occluded voxels within objects or behind walls. Furthermore, in datasets with sparse sensing, where numerous voxels remain unprobed, we remove these unknown voxels during training and evaluation to enhance the reliability of the ground truth as shown in Fig. 2.

D. Dataset Statistics

We show dataset statistics in Fig. 3. The dataset label distribution reveals noticeable domain gaps among various cities. Specifically, KITTI-360 and SemanticKITTI exhibit similar

label distributions due to being captured within the same city in Germany. However, nuScenes and Waymo collected in the US and Singapore demonstrate distinct label distributions. Furthermore, SSCBench stands out for its larger scale in comparison to SemanticKITTI. For instance, SSCBench comprises 7.7 times more frames than SemanticKITTI, and its collection of driving sequences is also more diverse.

IV. EXPERIMENTAL SETUP

Benchmark Methods. Our benchmark employs four camera-based methods, *i.e.*, MonoScene [17], VoxFormer [18], TPVFormer [61], and OccFormer [65], as well as two LiDAR-based methods, *i.e.*, SSCNet [12] and LMSCNet [14], due to their widespread adoption and good performance. Using their public codebases, we run them under default settings but appropriately adjust data loaders and batch sizes to align with our SSCBench. We train, validate, and test these methods separately on SSCBench subsets, as reported in Sec. V. Furthermore, we provide a unified benchmark for evaluating cross-domain generalizability in Sec. VI, where models are trained on one subset and tested on others, *e.g.*, trained on SSCBench-KITTI-360 and tested on SSCBench-Waymo.

Evaluation Metrics. We adopt the intersection over union (IoU) for evaluating geometry completion and the mean IoU (mIoU) of each class for evaluating semantic segmentation, following SemanticKITTI. We also report the ratio of different classes in the dataset to better understand the relationship between IoU and mIoU. We report the performances within the volume extending 51.2m ahead, 25.6m on each side, and 6.4m in height, and the voxel resolution is 0.2m. The design of this front-view evaluation emphasizes the area directly in the vehicle's anticipated forward trajectory.

V. SEPARATE BENCHMARKING RESULTS

A. Quantitative Comparisons

Camera-based Methods. On SSCBench-KITTI-360, TPV-Former and OccFormer demonstrate superior geometry completion performance (IoU) compared to VoxFormer and MonoScene, as illustrated in Tab. II. This improved geometry completion also contributes to enhancing semantic segmentation (mIoU). However, on SSCBench-Waymo, MonoScene slightly outperforms TPVFormer in both geometry completion and semantic segmentation. On SSCBench-nuScenes, the IoU and mIoU metrics for MonoScene significantly surpass those of VoxFormer (IoU, $29.63 \rightarrow 25.16$ and mIoU, $9.60 \rightarrow 4.96$). Due to the absence of stereo data in SSCBench-nuScenes, the utilization of off-the-shelf self-supervised depth estimation modules [66, 67], primarily trained on KITTI, results in suboptimal depth knowledge, leading to less competitive performance by VoxFormer in SSC. It is evident that accurate depth estimation plays a crucial role in scene geometry estimation within camera-based methods.

LiDAR-based Methods. SSCNet consistently outperforms LMSCNet across all three subsets, mainly due to its larger number of parameters (1.03M compared to 0.35M). Additionally, it is worth noting that SSCNet exhibits superior recognition capabilities for small objects, such as motorcycles

SEPARATE BENCHMARKING RESULTS ON THREE SSCBENCH SUBSETS. WE BENCHMARK MODELS WITH monocular images and point cloud inputs. The default evaluation range is $51.2 \times 51.2 \times 6.4$ m³. Due to the label differences amongst the three subsets, missing labels are replaced with "-". The top three performances on each dataset are marked by RED, Green, and Blue respectively.

Dataset	Method	Input	loU	mIoU	= car	bicycle	motorcycle	truck	other-veh.	■ person	road	parking	■ sidewalk	other-grnd	building _	fence	vegetation	terrain errain	elod _	trafsign	other-struct.	other-object	■ bicyclist	trunk
SSCBench- KITTI-360	LMSCNet SSCNet MonoScene Voxformer TPVFormer OccFormer	L C C C	53.58 37.87 38.76 40.22	11.91	31.95 19.34 17.84 21.56	1.09	0 0.17 0.58 0.89 1.37 0.26	0.26 10.29 8.02 4.56 8.06 9.89	0 0.58 2.03 2.06 2.57 3.82	0 0.07 0.86 1.63 2.38 2.77	62.95 65.7 48.35 47.01 52.99 54.3	13.51 17.33 11.38 9.67 11.99 13.44	41.24 28.13 27.21	3.78	44.41 32.89 31.18 34.83	6.77 3.53 4.97 4.80	43.72 26.15 28.99	26.80 28.87 16.75 14.69 17.51 19.51	0 0.78 6.92 6.51 7.46 7.77	6.92	3.63 8.60 4.20 3.79 5.48 6.95	0.0003 0.67 3.09 2.43 2.70 4.60	-	-
SSCBench- nuScenes	LMSCNet SSCNet MonoScene Voxformer OccFormer	L C C	27.64 29.63 25.16	11.84		0 1.7 0.29		5.93 13.42 8.35 2.73 11.88	8.52 10.35 8.74 2.45 9.80	3.41 7.59 3.72 1.12 5.87	24.14 28.74 38.77 23.94 37.62	- - - -	10.14		18.56 21.05 7.23 3.97 9.05	-	9.02 16.33 5.50 4.58 5.92	-	- - - -	- - - -	-	0 0.01 0.03 0.06 0	- - - -	-
SSCBench- Waymo	LMSCNet SSCNet MonoScene TPVFormer	L C	70.44 38.28	39.10 12.41	71.34 17.45	1.78 3.54		- - -	- - -	67.08 64.72 7.05 4.83	71.79 72.44 49.86 50.54	- - -	24.57	46.11	51.15 14.35	-	49.40 49.47 11.29 11.32	- - -	33.82 34.72 4.19 3.43	32.64 39.18 4.34 2.71	-	2.36 19.72 6.11 5.00		25.46 30.68 2.96 2.66

(■, 2.24 ↔ 0.00 in SSCBench-nuScenes) and bicycles (■, 1.78 ↔ 0.00 in SSCBench-Waymo). This demonstrates SSCNet's advantage in handling sparse LiDAR data compared to LMSCNet. When comparing results between SSCBench-KITTI-360 and SSCBench-nuScenes to SSCBench-Waymo, SSCNet and LMSCNet consistently deliver significantly better performance on the SSCBench-Waymo dataset, with IoU values approaching 70%. This improvement can be attributed to the denser LiDAR input available in Waymo data. However, it is crucial to emphasize that while dense LiDAR input leads to satisfactory performance, implementing this 5-LiDAR setup remains costly for common autonomous driving solutions.

Camera vs. LiDAR. As demonstrated in Tab. II, on SSCBench-KITTI-360, LiDAR-based methods outperform camera-based approaches in terms of geometry metrics and most semantic metrics. This outcome is expected since camera-based methods must infer 3D scene geometry from 2D images, while LiDAR-based methods directly extract scene geometry from LiDAR input. However, the scenario changes in SSCBench-nuScenes, where camera-based methods generally surpass LiDAR-based methods in terms of IoU. This difference can be attributed to the use of a sparse LiDAR sensor (Velodyne HDL32E) in the nuScenes dataset. These results indicate that LiDAR-based methods are sensitive to the sparsity of input. Specifically, while dense input has the potential for significant performance improvement, sparse input can lead to significant degradation. This observation is further confirmed in SSCBench-Waymo, where the Waymo dataset contains point cloud data collected from one midrange and four short-range LiDARs. As seen in Tab. II, the two LiDAR-based methods outperform camera-based methods by a significant margin on all metrics in SSCBench-Waymo.

However, camera-based methods outperform LiDAR-based ones for smaller objects that comprise a minuscule fraction of samples (< 0.5%). For classes such as bicycles (\blacksquare , $0.00 \rightarrow 1.16$), persons (\blacksquare , $0.26 \rightarrow 4.54$), poles (\blacksquare , $1.09 \rightarrow 12.93$), and traffic signs (\blacksquare , $0.90 \rightarrow 14.25$) in SSCBench-KITTI-360,

TABLE III

COMPARISON BETWEEN MONOCULAR AND TRINOCULAR SETUP ON SSCBENCH-WAYMO. THE MONOCULAR SETUP UTILIZES THE FRONT CAMERA ONLY. THE TRINOCULAR SETUP UTILIZES IMAGES FROM ALL 3 FRONT-FACING CAMERAS, WHICH COLLECTIVELY PROVIDE A 180° VIEW.

Method	Camera Setting	IoU	mloU	= car	bicycle	■ motorcycle	■ person	bicyclist	road	■ sidewalk	other-grnd	building -	vegetation	■ trunk	= pole	trafsign	other-object
TPVFormer	Monocular Trinocular	37.44 39.22	11.68 13.78	16.53 20.59	1.18 1.50	0.01 0.04	4.83 9.11	2.95 2.29	50.54 53.37	25.08 28.26	22.80 25.43	14.51 16.60	11.32 16.05	2.66 5.17	3.43 5.92	2.71 4.23	5.00 4.40

as well as motorcycles (\blacksquare , $0.00 \rightarrow 3.80$) and bicycles (\blacksquare , $0.00 \rightarrow 1.70$) in SSCBench-nuScenes, camera-based methods significantly outperform LiDAR-based methods. Despite the low frequency of small objects, their identification is vitally important for collision avoidance and traffic understanding.

B. Discussions and Analyses

Impact of Point Cloud Density. Our experiments illuminate the impact of LiDAR input density on model performance. In the SSCBench-nuScenes dataset, which features relatively sparse LiDAR input (32 channels), camerabased methods outperform LiDAR-based methods on geometric metrics. However, in the SSCBench-Waymo dataset, which benefits from dense LiDAR input (64 channels, 5 LiDARs), LiDAR-based methods vastly outperform camerabased methods. The sensitivity of LiDAR-based methods to input becomes evident, with advantages observed in dense input and notable performance degradation in sparse input. This highlights the need for future research in developing robust LiDAR-based methods that can mitigate degradation while capitalizing on the benefits.

Monocular vs. Trinocular. Table III displays the performance of TPVFormer with monocular and trinocular input. While a trinocular setup offers a broader field of view that can help enhance overall performance in terms of both IoU (37.44 \rightarrow 39.22) and mIoU (11.68 \rightarrow 13.78), achieving excellent results using only a single camera remains

TABLE IV

CROSS-DOMAIN EVALUATION. WE REPORT THE EXPERIMENT RESULTS
OF TRAINING AND TESTING ON DIFFERENT DATASETS USING UNIFIED
LABELS.

Training- dataset	Testing- dataset	Method	IoU	mIoU	■ vehicle	bicycle	■ motorcycle	■ person	road	■ sidewalk	other-grnd	building	vegetation	other-object
SSCBench- KITTI-360	SSCBench- KITTI-360	LMSCNet SSCNet MonoScene	48.49 54.50 40.56	22.48 25.63 20.78	23.87 34.14 23.00	0.00 0.00 3.60	0.00 0.00 3.00	0.00 0.11 4.58	64.04 67.54 54.00	35.13 41.84 31.96	17.46 19.86 14.39	44.47 46.89 36.29	39.755 44.92 30.45	0.04 1.01 6.51
SSCE	SSCBench- Waymo	LMSCNet SSCNet MonoScene	0.11 0.13 7.16	0.02 0.04 2.14	0.38 6.35 4.89	0.00 0.00 1.05	0.00 0.00 0.00	0.00 0.00 0.67	0.31 4.02 1.15	0.07 2.31 0.81	0.06 1.49 0.58	10.69 15.56 4.85	10.27 11.44 6.84	0.00 0.62 0.55
SSCBench- Waymo	SSCBench- KITTI-360	LMSCNet SSCNet MonoScene	0.10 0.17 16.29	0.02 0.05 3.99	0.58 2.73 6.55	0.00 0.00 0.15	0.00 0.00 0.01	0.01 0.04 0.80	1.45 7.19 7.76	0.10 1.39 2.93	0.91 1.29 3.14	9.95 14.67 6.95	9.93 17.86 11.30	0.09 0.07 0.34
SSCBen	SSCBench- Waymo	LMSCNet SSCNet MonoScene	71.77 71.00 37.79	43.57 43.48 15.57	77.17 69.68 17.84	0.00 0.28 3.32	0.00 0.00 0.01	68.78 63.30 7.64	70.28 70.92 49.16	37.65 42.73 23.09	42.79 46.27 23.33	51.63 51.49 13.40	49.52 50.38 13.44	37.89 39.69 4.45

a compelling academic challenge. There is still significant research value in developing monocular methods that can match the performance of models with panoramic views, as they are memory-efficient, computationally efficient, and easy to deploy.

Comparison with SemanticKITTI. We observe significant discrepancies when comparing our experimental results on SSCBench to those from SemanticKITTI [5] (for more details, we refer readers to VoxFormer [18]). While VoxFormer performs admirably well on SemanticKITTI in metrics such as IoU and mIoU, it faces challenges with the diversity of our SSCBench dataset. This challenge primarily arises from its depth estimation module's inability to generalize beyond SemanticKITTI. Furthermore, LMSCNet, which typically exhibits superior geometric performance compared to SSCNet on SemanticKITTI, demonstrates the opposite trend on SSCBench. These discrepancies underscore two essential points. First, they highlight the significance of SSCBench, which provides diverse and demanding real-world scenarios for comprehensive evaluations. Second, they emphasize the necessity for robust methods capable of maintaining high performance across various environments.

VI. UNIFIED BENCHMARKING RESULTS

To assess the domain gap and compare the cross-domain generalizability of state-of-the-art algorithms, we established a unified benchmark for cross-validation on SSCBench. Specifically, we employed two LiDAR-based methods, LM-SCNet [14] and SSCNet [12], and one camera-based method, MonoScene [17], for experiments on SSCBench-KITTI-360 and SSCBench-Waymo. To ensure consistent evaluation metrics, we standardized the labels of SSCBench-KITTI-360 and SSCBench-Waymo to a unified set comprising 10 common objects. All other experimental settings and evaluation metrics adhere to the guidelines outlined in Sec. IV.

Overall Performance. As shown in Tab. IV, all three methods exhibit a notable decline in performance when cross-validated on another dataset across both the geometric metric (IoU) as well as the semantic metric (mIoU), regardless of the training dataset. Specifically, the model trained on SSCBench-Waymo and tested on SSCBench-KITTI-360 suffers a more severe decline for LiDAR-based methods than the other way around. This is because SSCBench-Waymo has a very dense

point cloud input from five LiDARs, which effectively reduces the performance degradation caused by domain differences. Interestingly, the deterioration trend in terms of mIoU for MonoScene is more severe when transferring from SSCBench-KITTI-360 to SSCBench-Waymo than the other way around. This can be partially explained by the higher in-domain mIoU on SSCBench-KITTI-360 than that on SSCBench-Waymo and the difference in input resolutions ($1408 \times 376 \leftrightarrow 960 \times 640$), which is magnified by the fixed model parameters, and thereby affects feature representation.

Class-Specific Performance. We observe the most significant performance drop in the "road" class () for both transfer directions in all three methods. This suggests that ground types may be represented differently across datasets, causing challenges for both camera-based and LiDAR-based methods to adapt to these variations. Cross-domain evaluations consistently show a decline in performance, highlighting the importance of our proposed SSCBench dataset. Training models on this dataset should better prepare them for the variations and complexities of cross-domain scenarios. Moreover, it also serves as motivation for the development of models that are more robust and capable of generalizing across different domains.

VII. CONCLUSIONS

Limitations and Future Work. SSCBench only encompasses 3D data following the convention of the SSC problem. This limits evaluations of 4D methods with temporal dimensions. Future work will aim to expand SSCBench to include temporal information.

Summary. In this paper, we introduce SSCBench, a large-scale benchmark composed of diverse street views, aimed at facilitating the development of robust and generalizable semantic scene completion models. Through meticulous curation and comprehensive benchmarking, we identify the bottlenecks of existing methods and offer valuable insights into future research directions. Our ambition is for SSCBench to stimulate advancements in 3D semantic scene completion, ultimately enhancing perception capabilities for the next-generation autonomous systems.

REFERENCES

- [1] T. Wang, J. Pang, and D. Lin, "Monocular 3d object detection with depth from motion," in *European Conference on Computer Vision*. Springer, 2022, pp. 386–403. 1
- [2] H.-N. Hu, Y.-H. Yang, T. Fischer, T. Darrell, F. Yu, and M. Sun, "Monocular quasi-dense 3d object tracking," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 2, pp. 1992–2008, 2022.
- [3] W. Yuan, X. Gu, Z. Dai, S. Zhu, and P. Tan, "Newcrfs: Neural window fully-connected crfs for monocular depth estimation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [4] L. Roldao, R. De Charette, and A. Verroust-Blondet, "3d semantic scene completion: a survey," *International Journal of Computer Vision*, pp. 1–28, 2022. 1, 3
- [5] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, "Semantickitti: A dataset for semantic scene understanding of lidar sequences," in *Proceedings of*

- the IEEE/CVF International Conference on Computer Vision, 2019, pp. 9297–9307. 1, 2, 6
- [6] Y. Liao, J. Xie, and A. Geiger, "Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 1, 2, 3
- [7] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11621–11631. 1, 2, 3
- [8] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine *et al.*, "Scalability in perception for autonomous driving: Waymo open dataset," in *Proceedings of the IEEE/CVF conference on computer vision* and pattern recognition, 2020, pp. 2446–2454. 1, 2, 3
- [9] X. Dong, M. A. Garratt, S. G. Anavatti, and H. A. Abbass, "Towards real-time monocular depth estimation for robotics: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 10, pp. 16940–16961, 2022.
- [10] M. Enzweiler and D. M. Gavrila, "Monocular pedestrian detection: Survey and experiments," *IEEE transactions on* pattern analysis and machine intelligence, vol. 31, no. 12, pp. 2179–2195, 2008.
- [11] G. Younes, D. Asmar, E. Shammas, and J. Zelek, "Keyframe-based monocular slam: design, survey, and future directions," *Robotics and Autonomous Systems*, vol. 98, pp. 67–88, 2017.
- [12] S. Song, F. Yu, A. Zeng, A. X. Chang, M. Savva, and T. Funkhouser, "Semantic scene completion from a single depth image," in *Proceedings of the IEEE conference on computer* vision and pattern recognition, 2017, pp. 1746–1754. 2, 4, 6
- [13] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in 2012 IEEE conference on computer vision and pattern recognition. IEEE, 2012, pp. 3354–3361. 2, 3
- [14] L. Roldao, R. de Charette, and A. Verroust-Blondet, "Lmscnet: Lightweight multiscale 3d semantic completion," in 2020 International Conference on 3D Vision (3DV). IEEE, 2020, pp. 111–119. 2, 4, 6
- [15] R. Cheng, C. Agia, Y. Ren, X. Li, and L. Bingbing, "S3cnet: A sparse semantic scene completion network for lidar point clouds," in *Conference on Robot Learning*. PMLR, 2021, pp. 2148–2161.
- [16] C. B. Rist, D. Emmerichs, M. Enzweiler, and D. M. Gavrila, "Semantic scene completion using local deep implicit functions on lidar data," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 10, pp. 7205–7218, 2021. 2,
- [17] A.-Q. Cao and R. de Charette, "Monoscene: Monocular 3d semantic scene completion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 3991–4001. 2, 4, 6
- [18] Y. Li, Z. Yu, C. Choy, C. Xiao, J. M. Alvarez, S. Fidler, C. Feng, and A. 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*, 2023, pp. 9087–9098. 2, 4, 6
- [19] Q. Hu, B. Yang, L. Xie, S. Rosa, Y. Guo, Z. Wang, N. Trigoni, and A. Markham, "Randla-net: Efficient semantic segmentation of large-scale point clouds," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11108–11117.
- [20] S. Qiu, S. Anwar, and N. Barnes, "Semantic segmentation for real point cloud scenes via bilateral augmentation and adaptive fusion," in *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, 2021, pp. 1757– 1767.

- [21] A. Milioto, I. Vizzo, J. Behley, and C. Stachniss, "Rangenet++: Fast and accurate lidar semantic segmentation," in *IEEE/RSJ* international conference on intelligent robots and systems (IROS). IEEE, 2019, pp. 4213–4220. 2
- [22] C. Xu, B. Wu, Z. Wang, W. Zhan, P. Vajda, K. Keutzer, and M. Tomizuka, "Squeezesegv3: Spatially-adaptive convolution for efficient point-cloud segmentation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020, pp. 1–19. 2
- [23] Z. Zhou, Y. Zhang, and H. Foroosh, "Panoptic-polarnet: Proposal-free lidar point cloud panoptic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 13194–13203. 2
- [24] F. Hong, H. Zhou, X. Zhu, H. Li, and Z. Liu, "Lidar-based panoptic segmentation via dynamic shifting network," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 13 090–13 099.
- [25] J. Behley, A. Milioto, and C. Stachniss, "A benchmark for lidar-based panoptic segmentation based on kitti," in 2021 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2021, pp. 13596–13603.
- [26] L. Kreuzberg, I. E. Zulfikar, S. Mahadevan, F. Engelmann, and B. Leibe, "4d-stop: Panoptic segmentation of 4d lidar using spatio-temporal object proposal generation and aggregation," in *European Conference on Computer Vision*. Springer, 2022, pp. 537–553.
- [27] M. Aygun, A. Osep, M. Weber, M. Maximov, C. Stachniss, J. Behley, and L. Leal-Taixé, "4d panoptic lidar segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 5527–5537.
- [28] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proceedings of the IEEE conference on computer vision* and pattern recognition, 2017, pp. 652–660.
- [29] L. Yi, V. G. Kim, D. Ceylan, I.-C. Shen, M. Yan, H. Su, C. Lu, Q. Huang, A. Sheffer, and L. Guibas, "A scalable active framework for region annotation in 3d shape collections," ACM Transactions on Graphics (ToG), vol. 35, no. 6, pp. 1–12, 2016.
- [30] X. Zhu, H. Zhou, T. Wang, F. Hong, Y. Ma, W. Li, H. Li, and D. Lin, "Cylindrical and asymmetrical 3d convolution networks for lidar segmentation," in *Proceedings of the IEEE/CVF* conference on computer vision and pattern recognition, 2021, pp. 9939–9948.
- [31] H. Tang, Z. Liu, S. Zhao, Y. Lin, J. Lin, H. Wang, and S. Han, "Searching efficient 3d architectures with sparse point-voxel convolution," in *European conference on computer vision*. Springer, 2020, pp. 685–702.
- [32] J. Li, H. Dai, and Y. Ding, "Self-distillation for robust lidar semantic segmentation in autonomous driving," in *European Conference on Computer Vision*. Springer, 2022, pp. 659–676.
- [33] R. Cheng, R. Razani, E. Taghavi, E. Li, and B. Liu, "2-s3net: Attentive feature fusion with adaptive feature selection for sparse semantic segmentation network," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 12547–12556.
- [34] X. Huang, P. Wang, X. Cheng, D. Zhou, Q. Geng, and R. Yang, "The apolloscape open dataset for autonomous driving and its application," *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 10, pp. 2702–2719, 2019. 2
- [35] W. K. Fong, R. Mohan, J. V. Hurtado, L. Zhou, H. Caesar, O. Beijbom, and A. Valada, "Panoptic nuscenes: A large-scale benchmark for lidar panoptic segmentation and tracking," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3795–3802, 2022. 2, 3
- [36] M.-F. Chang, J. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan et al.,

- "Argoverse: 3d tracking and forecasting with rich maps," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 8748–8757. 2
- [37] J. Mao, M. Niu, C. Jiang, X. Liang, Y. Li, C. Ye, W. Zhang, Z. Li, J. Yu, C. Xu *et al.*, "One million scenes for autonomous driving: Once dataset," 2021. 2
- [38] R. Kesten, M. Usman, J. Houston, T. Pandya, K. Nadhamuni, A. Ferreira, M. Yuan, B. Low, A. Jain, P. Ondruska, S. Omari, S. Shah, A. Kulkarni, A. Kazakova, C. Tao, L. Platinsky, W. Jiang, and V. Shet, "Woven planet perception dataset 2020," https://woven.toyota/en/perception-dataset, 2019.
- [39] Q.-H. Pham, P. Sevestre, R. S. Pahwa, H. Zhan, C. H. Pang, Y. Chen, A. Mustafa, V. Chandrasekhar, and J. Lin, "A 3d dataset: Towards autonomous driving in challenging environments," in 2020 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2020, pp. 2267–2273.
- [40] J. Geyer, Y. Kassahun, M. Mahmudi, X. Ricou, R. Durgesh, A. S. Chung, L. Hauswald, V. H. Pham, M. Mühlegg, S. Dorn et al., "A2d2: Audi autonomous driving dataset," arXiv preprint arXiv:2004.06320, 2020. 2
- [41] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. R. Qi, Y. Zhou et al., "Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9710– 9719. 2
- [42] Y. Hu, J. Yang, L. Chen, K. Li, C. Sima, X. Zhu, S. Chai, S. Du, T. Lin, W. Wang et al., "Planning-oriented autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 17853–17862. 2, 3
- [43] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *International Journal of Robotics Research (IJRR)*, 2013. 2
- [44] J. Fritsch, T. Kuehnl, and A. Geiger, "A new performance measure and evaluation benchmark for road detection algorithms," in *International Conference on Intelligent Transportation Systems (ITSC)*, 2013. 2
- [45] M. Menze and A. Geiger, "Object scene flow for autonomous vehicles," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [46] C. Chen, X. Liu, Y. Li, L. Ding, and C. Feng, "Deepmapping2: Self-supervised large-scale lidar map optimization," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 9306–9316.
- [47] X. Huang, X. Cheng, Q. Geng, B. Cao, D. Zhou, P. Wang, Y. Lin, and R. Yang, "The apolloscape dataset for autonomous driving," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2018, pp. 954–960.
- [48] M. Pitropov, D. E. Garcia, J. Rebello, M. Smart, C. Wang, K. Czarnecki, and S. Waslander, "Canadian adverse driving conditions dataset," *The International Journal of Robotics Research*, vol. 40, no. 4-5, pp. 681–690, 2021. 2
- [49] S. Agarwal, A. Vora, G. Pandey, W. Williams, H. Kourous, and J. McBride, "Ford multi-av seasonal dataset," *The International Journal of Robotics Research*, vol. 39, no. 12, pp. 1367–1376, 2020.
- [50] Y. Li, D. Ma, Z. An, Z. Wang, Y. Zhong, S. Chen, and C. Feng, "V2x-sim: Multi-agent collaborative perception dataset and benchmark for autonomous driving," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10914–10921, 2022.
- [51] R. Xu, X. Xia, J. Li, H. Li, S. Zhang, Z. Tu, Z. Meng, H. Xiang, X. Dong, R. Song et al., "V2v4real: A real-world large-scale dataset for vehicle-to-vehicle cooperative perception," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 13712–13722. 2
- [52] C. A. Diaz-Ruiz, Y. Xia, Y. You, J. Nino, J. Chen, J. Monica,

- X. Chen, K. Luo, Y. Wang, M. Emond *et al.*, "Ithaca365: Dataset and driving perception under repeated and challenging weather conditions," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 21383–21392. 2
- [53] P. Xiao, Z. Shao, S. Hao, Z. Zhang, X. Chai, J. Jiao, Z. Li, J. Wu, K. Sun, K. Jiang et al., "Pandaset: Advanced sensor suite dataset for autonomous driving," in 2021 IEEE International Intelligent Transportation Systems Conference (ITSC). IEEE, 2021, pp. 3095–3101.
- [54] X. Tian, T. Jiang, L. Yun, Y. Wang, Y. Wang, and H. Zhao, "Occ3d: A large-scale 3d occupancy prediction benchmark for autonomous driving," arXiv preprint arXiv:2304.14365, 2023.
- [55] X. Wang, Z. Zhu, W. Xu, Y. Zhang, Y. Wei, X. Chi, Y. Ye, D. Du, J. Lu, and X. Wang, "Openoccupancy: A large scale benchmark for surrounding semantic occupancy perception," arXiv preprint arXiv:2303.03991, 2023.
- [56] S. Thrun, "Probabilistic robotics," *Communications of the ACM*, vol. 45, no. 3, pp. 52–57, 2002. 2
- [57] S. L. Bowman, N. Atanasov, K. Daniilidis, and G. J. Pappas, "Probabilistic data association for semantic slam," in 2017 IEEE international conference on robotics and automation (ICRA). IEEE, 2017, pp. 1722–1729.
- [58] X. Zhang, A. Kundu, T. Funkhouser, L. Guibas, H. Su, and K. Genova, "Nerflets: Local radiance fields for efficient structure-aware 3d scene representation from 2d supervision," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 8274–8284. 3
- [59] J. Gu, C. Hu, T. Zhang, X. Chen, Y. Wang, Y. Wang, and H. Zhao, "Vip3d: End-to-end visual trajectory prediction via 3d agent queries," in *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, 2023, pp. 5496– 5506. 3
- [60] Y. Li, C. Wen, F. Juefei-Xu, and C. Feng, "Fooling lidar perception via adversarial trajectory perturbation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 7898–7907. 3
- [61] Y. Huang, W. Zheng, Y. Zhang, J. Zhou, and J. Lu, "Triperspective view for vision-based 3d semantic occupancy prediction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9223–9232. 3, 4
- [62] T. Bailey and H. Durrant-Whyte, "Simultaneous localization and mapping (slam): Part ii," *IEEE robotics & automation magazine*, vol. 13, no. 3, pp. 108–117, 2006. 4
- [63] D. Maturana and S. Scherer, "Voxnet: A 3d convolutional neural network for real-time object recognition," in 2015 IEEE/RSJ international conference on intelligent robots and systems (IROS). IEEE, 2015, pp. 922–928. 4
- [64] Y. Zhou and O. Tuzel, "Voxelnet: End-to-end learning for point cloud based 3d object detection," in *Proceedings of the IEEE* conference on computer vision and pattern recognition, 2018, pp. 4490–4499.
- [65] Y. Zhang, Z. Zhu, and D. Du, "Occformer: Dual-path transformer for vision-based 3d semantic occupancy prediction," in Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2023. 4
- [66] S. F. Bhat, I. Alhashim, and P. Wonka, "Adabins: Depth estimation using adaptive bins," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 4009–4018.
- [67] F. Shamsafar, S. Woerz, R. Rahim, and A. Zell, "Mobilestere-onet: Towards lightweight deep networks for stereo matching," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2022, pp. 2417–2426. 4