

# Anyview: General Indoor 3D Object Detection with Variable Frames

Zhenyu Wu<sup>1</sup>, Xiuwei Xu<sup>2</sup>, Ziwei Wang<sup>3</sup>, Chong Xia<sup>2</sup>, Linqing Zhao<sup>2</sup>, Jiwen Lu<sup>2</sup> and Haibin Yan<sup>\*1</sup>

**Abstract**—In this paper, we propose a novel network framework for indoor 3D object detection to handle variable input frame numbers in practical scenarios. Existing methods only consider fixed frames of input data for a single detector, such as monocular RGB-D images or point clouds reconstructed from dense multi-view RGB-D images. While in practical application scenes such as robot navigation and manipulation, the raw input to the 3D detectors is the RGB-D images with variable frame numbers instead of the reconstructed scene point cloud. However, the previous approaches can only handle fixed frame input data and have poor performance with variable frame input. In order to facilitate 3D object detection methods suitable for practical tasks, we present a novel 3D detection framework named AnyView for our practical applications, which generalizes well across different numbers of input frames with a single model. To be specific, we propose a geometric learner to mine the local geometric features of each input RGB-D image frame and implement local-global feature interaction through a designed spatial mixture module. Meanwhile, we further utilize a dynamic token strategy to adaptively adjust the number of extracted features for each frame, which ensures consistent global feature density and further enhances the generalization after fusion. Extensive experiments on the ScanNet dataset show our method achieves both great generalizability and high detection accuracy with a simple and clean architecture containing a similar amount of parameters with the baselines.

## I. INTRODUCTION

3D object detection is a fundamental scene understanding problem for robotic manipulation [15], autonomous driving [17], and AR/VR [22], which aims to detect the 3D bounding boxes and semantic labels from point clouds or images. Due to the different types of sensors used in different application scenarios, 3D object detection methods usually vary a lot for indoor [4] and outdoor [25] scenes. We focus on indoor 3D object detection, where the mainstream sensor is RGB-D cameras and the scenes are crowded with objects of multiple categories and sizes.

Although great improvement in performance has been achieved by advanced architecture design, existing methods train and evaluate 3D detectors on only fixed frames of input data, such as monocular RGB-D images (on SUN-RGBD [18] benchmark) and scene-level reconstructed point clouds (sampled from meshes) from multi-view RGB-D images (on ScanNet [2] benchmark). There is still a huge

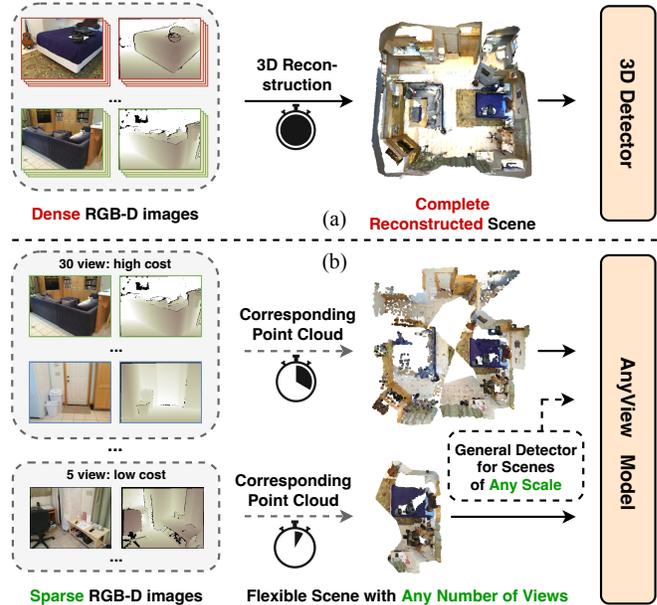


Fig. 1. (a) Previous 3D object detection methods rely on reconstructed point clouds from densely sampled RGB-D images, which requires a large amount of time to process the data and generalizes poorly to scenes in different scales. (b) Our method directly detects objects from sparse RGB-D images of any number of views, which shows good generalizability and provides more trade-off.

gap between these benchmarks and practical applications. Specifically, due to the various task budgets in practical application scenes, the number of frames input to the 3D detectors is dynamic, and previous frameworks trained on fixed frame inputs are challenging to generalize. The cost of training diverse models for tasks with various sizes of inputs is huge, severely limiting the deployment of the models on edge devices. In order to solve this problem, it is necessary to further research and develop more efficient generalized indoor 3D object detection with variable frames. For example, in online tasks such as robot navigation [6], the information captured by the agent is strictly limited by a cost budget constraint, which leads to a varying number of input frames. As we will show in Section III-A, previous models generalize poorly across different input scales, which brings a huge burden as we need to prepare a series of models trained on different scales of input data to cope with various scenarios.

In practical scenes, the raw input to the indoor 3D object detectors is the RGB-D images with variable frame numbers. The detectors require the ability to process inputs with various frame numbers at fixed parameters to complete different tasks. To simulate this, we propose a novel processing framework named AnyView for indoor 3D object detection

\*Corresponding author.

<sup>1</sup>Zhenyu Wu and Haibin Yan are with the School of Automation, Beijing University of Posts and Telecommunications. {wuzhenyu, eyanhaibin}@bupt.edu.cn.

<sup>2</sup>Xiuwei Xu, Linqing Zhao, Chong Xia, and Jiwen Lu are with the Department of Automation, Tsinghua University. {xxw21, zhaolinqing, xiac20}@mails.tsinghua.edu.cn. lujiwen@tsinghua.edu.cn.

<sup>3</sup>Ziwei Wang is with the School of Electrical and Electronic Engineering, Nanyang Technological University. {ziwei.wang}@ntu.edu.sg

as shown in Fig. 1. Previous methods require complete scene reconstruction from dense RGB-D images, limiting the model to a fixed number of inputs, which weakens generalization and wastes resources. In contrast, AnyView handles variable frame numbers, making it adaptable to different application budgets. We unify input data as sparse multi-view RGB-D images with camera parameters, compatible with datasets like ScanNet [2] and SUN-RGBD [18]. Instead of concatenating point clouds and extracting fixed global features, AnyView learns geometric clues per frame and uses a transformer-based architecture to efficiently merge and refine these features, enabling processing of variable frame inputs. We propose a geometric learner to extract local geometric features and a spatial mixtures module to combine these with global semantics, which enables AnyView to generate rich representations from scenes with any number of frames, bridging local geometric and global semantic features. We introduce randomized view and rectangular dropout strategies to help the model adapt to varying input scales and discard frames with less geometric information. Additionally, a dynamic token strategy ensures consistent feature density and parameter compatibility across point clouds. Extensive experiments on ScanNet show that our method outperforms previous methods in both accuracy and generalizability with a similar amount of parameters.

## II. RELATED WORK

**Point-based 3D object detection:** This kind of methods take in 3D scenes represented by pure point clouds, which are acquired by LIDAR sensor or 3D reconstruction from multi-view RGB-D images [8]. Early 3D object detection methods mainly include sliding-window methods [19] and template-based methods [9]. Deep learning-based end-to-end 3D object detection methods began to emerge in recent years, which are mainly based on PointNet [13] or sparse CNN [3] backbones. PointNet-based methods [23] consume point clouds directly with the set abstraction operation [14], which enables flexible receptive fields for 3D feature learning. Sparse CNN-based methods [21] project the point clouds to regular grids to be processed by advanced 2D or 3D CNN architectures. However, point-based detection models have high inference latency due to inefficient sampling algorithms.

**RGBD-based 3D object detection:** RGBD camera is the mainstream 3D sensor for indoor scene understanding tasks. Existing RGBD-based 3D object detection methods take in monocular RGB-D image [18] or point clouds reconstructed from multi-view RGB-D images (mentioned above). Here we only review the former methods which make full use of both RGB and depth information. Prior methods broadly fall into three categories: 2D-driven, 3D-driven, and modal fusion. 2D-driven methods [24] first detect object in images and then use the results to assist search in 3D space. Kundu *et al.* [7] presented a microscopic rendering and comparison loss by employing voxel patterns to represent 3D objects, which allows 3D shapes and poses to be learned via 2D supervision. However, there is still a huge performance gap between RGBD-based methods and point cloud-based methods. The

TABLE I

Generalizability of previous model. The gray cell indicates on which benchmark the model is trained.

	ScanNet-SV	ScanNet-Rec	ScanNet-MV
ImVoteNet	31.6 / 13.5	–	43.3 / 20.4
3DETR	19.8 / 6.1	64.5 / 44.0	44.7 / 25.6

most common approach to 3D-driven is to feed image information into a 3D feature extraction network as an additional channel to a point cloud or voxel. Song *et al.* [19] localized objects on a voxelized point cloud by sliding a 3D detection window, and category information was obtained from RGB pixels attached to the point cloud. However, this type of simple point cloud and image combination approach may destroy the fine-grained local geometric structural features of the point cloud, resulting in inefficient fusion between the point cloud and image modalities.

## III. ANALYSIS

### A. Problem Statement

In order to better utilize 3D detector in practical tasks where the raw input is a variable number of RGB-D images, we propose a new setting for indoor 3D object detection.

**Inference:** Unlike the input data format of the existing benchmarks, in order to simulate the inference process of the model in the practical scene, we take the sparse RGB-D images and their camera parameters as the unified input data of the 3D detector to provide the intelligent body with the scene object detection results, which can be summarized as follows:

$$\mathcal{B}_N = \mathcal{D}(\{I_k, R_k^c, T_k^c\}_{k=1}^N), R_k^c \in \mathbb{R}^{3 \times 3}, T_k^c \in \mathbb{R}^{3 \times 1} \quad (1)$$

where  $N$  represents the number of RGB-D image frames input to the 3D object detector  $\mathcal{D}$ , and  $\mathcal{B}_N$  represents the reconstructed point cloud detection results of  $N$  frames of RGB-D images.  $I_k$ ,  $R_k^c$ , and  $T_k^c$  correspond to the  $k_{th}$  frame of RGB-D image, rotation matrix, and translation matrix, respectively.  $\{I_k, R_k^c, T_k^c\}_{k=1}^N$  can represent both local region and a whole scene. To reduce the time for data collection, the maximum number of views is set to 50 ( $N \leq 50$ ), which is far less than the number required for 3D reconstruction.

**Model:** Since the number of input views will change according to the task in practical scene deployments, the 3D object detectors  $\mathcal{D}$  must be highly flexible in order to efficiently process RGB-D images with the variable number of frames input. Therefore, the training process of models  $\mathcal{D}$  needs to incorporate adaptations which enable  $\mathcal{D}$  to process point clouds with arbitrary scale inputs without additional fine-tuning. This attribute significantly improves the utility and efficiency of the 3D detection applications, enabling them to easily handle a wide range of task input scenes and provide accurate object detection results.

**Evaluation:** The performance of model should not be evaluated only on one scale of input. We prepare a variety of validation sets containing different scales of input data, such as monocular input, few-view input and scene-level input. The number of views  $N$  can range from 1 to 50. Two representative methods are chosen: ImVoteNet [11]

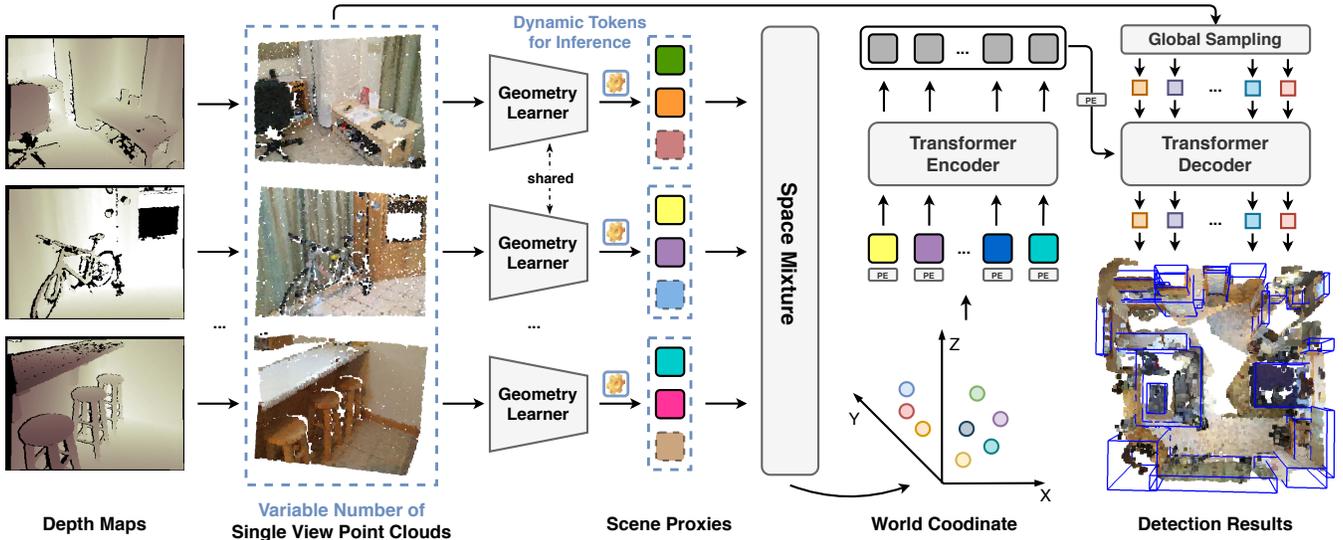


Fig. 2. The framework of *AnyView*. Given multi-view RGB-D images of variable number, we first convert the depth maps to single view point clouds according to intrinsic camera parameters. Then a shared geometry learner is applied to extract scene proxies for each view, which are local feature descriptors independent of input scale. Scene proxies from different views are mixed into world coordinates according to the extrinsic camera parameters. We adopt transformer encoder to refine the scene proxies while keeping their input scale-independent property by self-attention mechanism. A transformer decoder is used to refine the object proposals, where the initial proposals are sampled from the whole scene by FPS. The dashed blue box indicates the number of inside elements is changable during inference.

for monocular RGB-D input and 3DETR [10] for scene-level point cloud input. We prepare three benchmarks from the raw RGB-D videos provided by ScanNet: ScanNet-SV, ScanNet-Rec, and ScanNet-MV<sup>1</sup>. ScanNet-SV (single view) is a monocular RGB-D benchmark organized similar to SUN-RGBD. ScanNet-Rec (reconstruction) is the previous scene-level benchmark. ScanNet-MV (multi-view) shares the same ground truth with ScanNet-Rec but provides multi-view RGB-D images as input instead of reconstructed point clouds.

We train ImVoteNet and 3DETR on ScanNet-SV and ScanNet-Rec respectively and evaluate the models on all benchmarks. When applying 3DETR on ScanNet-MV, we fuse the multi-view depth maps into point clouds according to the camera parameters. As for applying ImVoteNet on ScanNet-MV, we predict bounding boxes for each view and fuse the results by 3D NMS. As shown in Table I, 3DETR trained on reconstructed point clouds performs poorly on the monocular RGB-D benchmark, and both ImVoteNet and 3DETR fail to achieve a satisfactory performance on the multi-view RGB-D benchmark. This experimental result indicates that the performance of the detector on previous benchmarks may not reflect its performance in practical applications. And the poor generalizability of previous models poses a huge challenge in applying the existing models in practical tasks, where time for data collection is limited and the scale of input data is changeable.

#### IV. APPROACH

##### A. Overall Framework

The overall framework of *AnyView* is illustrated in Fig. 2. Given multi-view RGB-D images as input, *AnyView* focuses

on the utilization of depth maps and camera parameters. We will introduce our method in detail as follows.

Formally, the images are represented by  $\{I_1, I_2, \dots, I_N\}$ . We sample  $K$  points for each depth map and convert them into camera coordinate by intrinsic camera calibration matrix, which are denoted as  $\{P_{c1}, P_{c2}, \dots, P_{cN}\}, P_{ci} \in \mathbb{R}^{K \times 3}$ . The rotations and translations from camera coordinate to world coordinate for each view are  $\{R_1^c, R_2^c, \dots, R_N^c\}$  and  $\{T_1^c, T_2^c, \dots, T_N^c\}$ .

**Extract scene proxies:** Previous 3D detection methods are only able to process point clouds as a whole. So a natural solution for consuming multi-view RGB-D images is to fuse the depth maps into a whole scene  $P = \mathcal{C}(R_i^c \cdot P_{ci} + T_i^c)$ , where  $\mathcal{C}$  denotes concatenation operation. However, point clouds generated from different numbers of views vary a lot in local geometry structure and global semantics, which will lead to deteriorated scene representation when numbers of view for training and inference are not the same. To decouple scene representation with the scale of input data, we propose to extract  $T$  scene proxies for each view independently with a shared geometry learner:

$$\mathcal{P}_i = G(P_{ci}), \mathcal{P}_i \in \mathbb{R}^{T \times C} \quad (2)$$

Scene proxies  $\mathcal{P}_i$  are local feature descriptors which represent the geometry structures of single view point clouds. We implement the geometry learner  $G$  as two set abstraction (SA) layers [14], the first layer with a constrained receptive field to focus on local geometric details and the second with a large receptive field to aggregate the local details into geometry structures. As  $G$  is applied on each view independently, the extracted scene proxies will ignore global semantics and be robust to the scales of input data.

**Interactions among scene proxies:** We obtain richer scene representations through interactions among scene proxies. In order to keep the input scale-independent property of

<sup>1</sup>Refer to Section V-A for more details.

scene proxies, we hope the interaction to be linear combination  $\mathcal{P}_{ij} = \sum_{i=1}^N \sum_{j=1}^T \alpha_{ij} \mathcal{P}_{ij}$ , where  $\sum_i \sum_j \alpha_{ij} = 1$ . In addition, since the number of view ( $N$ ) is variable, the number of interacting scene proxies is also variable. Benefiting from the nature of self-attention, these requirements can be elegantly achieved by transformer.

Specifically, we first mix the scene proxies in world coordinates by transforming their coordinates:

$$\{\mathbf{P}_{ij}^x, \mathbf{P}_{ij}^y, \mathbf{P}_{ij}^z\} = R_i \cdot \{\mathbf{P}_{ij}^x, \mathbf{P}_{ij}^y, \mathbf{P}_{ij}^z\} + T_i \quad (3)$$

where  $\mathbf{P}_{ij}^{x,y,z}$  represents the spatial position of  $\mathcal{P}_{ij}$  in the camera coordinate system, and we project them into the world coordinate system to obtain global information at a higher level with the camera projection matrix. Then the scene proxies are fed into the transformer encoder with their coordinates converted into Fourier positional embeddings [20]. We adopt radius mask [10] on the self-attention matrix to conduct interactions from local to global.

**Decode objects from the whole scene:** Given the  $N_T = N \times T$  features from the encoder, we adopt a transformer decoder to refine object queries layer by layer as in DETR [1]. The object queries are initial proposals, which are generated by furthest point sampling from the whole scene to ensure coverage:

$$\mathbf{Q} = \mathcal{M}(\mathcal{F}(\mathcal{C}(R_i^c \cdot P_{ci} + T_i^c))) \quad (4)$$

where  $\mathcal{M}$  represents the MLP and  $\mathcal{F}$  represents the farthest point sampling algorithm. The refined queries are converted into box parameters and supervised following the procedure in 3DETR.

### B. Scale-independent Training

**Random view dropping:** During training, the numbers of views of different inputs in a batch are kept the same in order to parallelize the computation. However, enabling the network to process variable number of views is necessary as it encourages the transformer to adapt to different scale of attention map. To this end, we randomly drop 0 to  $\frac{N}{2}$  views and pad the dropped points with 0. To cut down the interactions between the dropped scene proxies and other tokens (scene proxies for encoder, object queries for decoder), we apply a binary mask on each attention map to indicate which of scene proxies are invalid (with coordinates  $(0, 0, 0)$ ). The binary mask will set the inner products between invalid scene proxies and other tokens to  $-\infty$ , so after softmax operation these attention value will be 0.

**Global random cuboid:** The geometry structure of single view point clouds may vary from view to view. For example, when the distance between the RGB-D camera and the scene varies, the density of point clouds varies as well. When one view only contains floors and walls, the scene proxies extracted from it are uninformative. To allow the network to handle depth maps taken from various shooting situations and learn to ignore scene proxies with less information, we randomly crop a cuboid from the fused point clouds  $P = \mathcal{C}(R_i^c \cdot P_{ci} + T_i^c)$  and set other points outside the cuboid to  $(0, 0, 0)$ . After that, we transform the point clouds back

to the camera coordinates of each view and keep  $(0, 0, 0)$  unchanged. This strategy mak some views lose geometric information but still provide valid scene proxies, which effectively reduces overfitting and improves the generalizability of AnyView.

### C. Dynamic Tokens for Inference

AnyView extracts  $N_T$  scene proxies for a  $N$ -view input scene. In our setting,  $N$  may vary from 1 to 50 during inference time, which results in an order of magnitude change in the number of scene proxies. In previous work such as VoteNet [12] and 3DETR, the encoder of detector extract 1024 or 2048 seed points for one scene, no matter monocular RGB-D input or scene-level point clouds input. To reach this order of magnitude,  $T$  should be around 40 when  $N = 50$  and be around 2000 when  $N = 1$ . For the former case, AnyView only extracts 40 scene proxies for monocular input, which is not enough to fully represent the scene. While for the latter case, the number of scene proxies will be larger than 10000 when there are more than 10 views. This heavy computation cost is unaffordable in practical applications.

To our best knowledge, all previous networks which use SA layers to extract point features keep a fixed output point number during training and inference. However, as the parameters in SA layer are only relevant with the number of channels, changing the output number seems feasible. The reason why previous methods fix the output number are two-fold: 1) if the output number of a SA layer is changed, the shape of input tensor for the next layer will change and may be incompatible with the parameters; 2) even if the next layer can process input in dynamic shape, the density of point clouds in this level is different from the one during training, which makes the layer unable to extract accurate feature representation. To overcome this problem in our case, we devise a dynamic token strategy for the geometry learner by keeping the output number of the first SA layer fixed and change the output number of the second one. Thus  $T$  can be defined as:

$$T = \min\{Z/N, O_{SA_1}\} \quad (5)$$

where  $Z$  is a predefined constant and  $O_{SA_1}$  is the the output number of the first SA layer. In this way, the second SA layer receives the same number of input points as in training time and thus the output features (i.e. scene proxies) are accurate. Although the number of scene proxies varies, the transformer encoder is still able to interact and refine them due to the linear combination nature of self-attention mechanism and our scale-independent training strategy.

Meanwhile, our proposed AnyView can be well applied to the online detection task as it extracts features  $\mathcal{P}_i$  individually for each view  $V_i$  when processing RGB-D inputs and then integrates the features of different views  $\{\mathcal{P}_1, \dots, \mathcal{P}_i\}$  in the world coordinate system through an attention mechanism. The ability to detect objects online fast and accurately in dynamic 3D scenes is crucial for autonomous navigation, environment understanding, and other tasks in the field of

TABLE II

Number of parameters of different 3D detectors.

ImVoteNet	ImVoxelNet	3DETR-m	AnyView
1.8M	104.6M	7.4M	7.6M

autonomous driving and robotics. Our proposed method provides a reliable and efficient solution for online detection, which can enable robots to effectively perceive and interact with the environment.

## V. EXPERIMENT

### A. Experiments Setup

**Datasets and benchmarks:** We conduct experiments on the ScanNet [2] dataset, which is a richly annotated dataset of indoor scenes with 1201 training scenes and 312 validation scenes. For each scene, ScanNet provides RGB-D video as well as the reconstructed mesh. We uniformly sample 50 frames from the video for each scene, which serves as the maximum data available for training and evaluating in our proposed setting. Meanwhile, we further validate that the proposed approach has the ability to accurately detect objects in practical deployment scenes by verifying the performance of each model on the online detection task. We set three benchmarks as defined in Section III-A. ScanNet-Rec is the mainstream scene-level benchmark used by previous works [5], [12], [10], whose input data is reconstructed point cloud of a whole scene and ground-truth is generated by computing the axis-aligned bounding boxes [5] for objects in 18 selected categories. In addition, we designed ScanNet-Online benchmarks to explore the performance of previous 3D object detection models with AnyView in practical deployment scenes. Specifically, the input to ScanNet-Online is each frame in a sequence of RGB-D images, and the model needs to dynamically detect objects in the scene based on the new input frames. ScanNet-MV shares the same ground truth with ScanNet-Rec, while the input data are multi-view RGB-D images. If not additionally mentioned, the default number of views is 50. ScanNet-MV is also a scene-level benchmark, which means all RGB-D images from one scene are considered as a single input sample for the network. ScanNet-SV is a monocular RGB-D benchmark with a single view RGB-D image as input data. For each image, we select the bounding boxes whose center points are within the image from the corresponding scene as its ground truth.. ScanNet-Online utilizes the ScanNet-MV benchmark in the 50 views setting, and the input is each frame of sequential RGB-D images.

**Compared methods:** We train VoteNet [12] and ImVoteNet [11] on ScanNet-SV. When applying the models on ScanNet-MV, we predict bounding boxes for each view and fuse the predictions by 3D NMS. 3DETR [10] is chosen as the model trained on ScanNet-Rec, which is the mainstream transformer-based 3D detector with less inductive bias. The comparison between 3DETR and AnyView is particularly relevant since they share similar amount of parameters and both adopt DETR-like detection decoder. To apply 3DETR on ScanNet-MV, we preprocess the multi-view

TABLE III

3D object detection results (mAP@0.25 and mAP@0.5) and corresponding training/inference setting of different methods on scene-level and monocular benchmarks. Contents in bracket indicate the method modality.

	Method	Training Benchmark	Evaluation Benchmark	mAP	
				@0.25	@0.5
Scene-level	3DETR	Rec	Rec	62.7	37.5
	3DETR-m	Rec	Rec	65.0	47.0
	3DETR	MV(D)	Rec	53.6	33.0
	3DETR-m	MV(D)	Rec	56.9	37.8
	3DETR	Rec	MV(D)	37.8	22.2
	3DETR-m	Rec	MV(D)	44.7	25.6
	VoteNet	SV(D)	MV(D)	40.9	20.6
	ImVoteNet	SV(GBD)	MV(GBD)	43.3	20.4
	ImVoxelNet	MV(RGB)	MV(RGB)	46.6	25.2
	3DETR	MV(D)	MV(D)	51.7	31.0
	3DETR-m	MV(D)	MV(D)	54.7	35.3
	AnyView	MV(D)	MV(D)	60.7	35.8
Monocular	VoteNet	SV(D)	SV(D)	30.1	13.9
	ImVoteNet	SV(GBD)	SV(GBD)	31.6	13.5
	3DETR	Rec	SV(D)	14.7	3.9
	3DETR-m	Rec	SV(D)	19.8	6.1
	ImVoxelNet	MV(RGB)	SV(RGB)	21.2	8.3
	3DETR	MV(D)	SV(D)	24.7	10.9
	3DETR-m	MV(D)	SV(D)	27.7	12.8
	AnyView	MV(D)	SV(D)	32.7	15.3

RGB-D images by fusing the depth maps from each view according to the camera parameters. We train our AnyView model as well as 3DETR and ImVoxelNet [16] on ScanNet-MV, each model representing an input modality: 3DETR for scene-level point clouds, ImVoxelNet for multi-view RGB and AnyView for multi-view RGB-D. We list the number of parameters of these models in Table II.

**Implementation details:** We downsample the RGB-D videos from ScanNet to a resolution of  $320 \times 240$ . Following previous setting [12], we sample 20000 points per scene for ScanNet-SV and 40000 points per scene for ScanNet-Rec as the input point clouds. For ScanNet-MV, we sample 20000 points per view for methods that make predictions on each view independently and sample 40000 points per scene for methods that fuse multi-view point clouds as a whole. While for AnyView, we sample 5000 points per view to reduce computation cost.

The geometry learner of AnyView consists of two SA layers, the first with radius  $0.2m$ , output number of points 256 and MLP channels [3, 64, 128, 256], the second with radius  $0.8m$  and MLP channels [256, 256, 256, 256]. During training, we set  $T = 40$ . While in inference time, we set  $Z = 2000$ . Following the configurations of 3DETR-m, we adopt 3 transformer encoders and 8 decoders in AnyView. The radius mask for encoder is set to  $[0.8m, 0.8m, 1.2m]$ . In terms of data augmentation, both random view dropping and global random cuboid are applied with probability 0.75.

As for the online detection setting, in processing the online detection RGB-D input flow  $F = \{f_1, \dots, f_i\}$ , our proposed AnyView will save the extracted  $\mathcal{P}_i$  for each RGB-D image

TABLE IV

3D object detection results (mAP@0.25 and mAP@0.5) of different models on ScanNet-SV and ScanNet-MV. Here ScanNet-MV is further divided according to the number of views and how these views are combined. Uniform/Continuous means the views are sampled uniformly/adjacently from the whole 50 views. Gray cells show on which benchmark the model is trained.

	Methods	ScanNet-SV	ScanNet-MV					
			5	10	15	30	40	50
Uniform	ImVoteNet	31.6 / 13.5	30.2 / 13.7	37.3 / 17.9	40.3 / 18.5	43.5 / 20.6	42.9 / 20.4	43.3 / 20.4
	3DETR-m(10)	29.8 / 12.5	39.7 / 20.0	49.1 / 28.0	50.8 / 29.0	50.3 / 28.9	49.6 / 28.9	48.9 / 27.7
	3DETR-m(30)	29.1 / 13.1	36.8 / 19.5	50.1 / 29.6	53.1 / 31.4	55.6 / 33.9	55.7 / 34.3	55.6 / 34.7
	3DETR-m(50)	27.7 / 12.8	34.9 / 17.9	48.6 / 29.0	52.5 / 33.3	55.3 / 36.3	55.6 / 35.9	54.7 / 35.3
	AnyView(10)	31.3 / 13.7	42.2 / 20.4	51.3 / 27.2	52.8 / 29.0	53.2 / 29.9	54.9 / 30.4	54.8 / 30.1
	AnyView(30)	30.5 / 12.7	39.8 / 20.1	52.4 / 29.2	57.1 / 31.6	59.1 / 35.2	59.9 / <b>36.8</b>	60.1 / <b>37.4</b>
	AnyView(50)	<b>32.7 / 15.3</b>	<b>45.0 / 22.0</b>	<b>53.8 / 30.5</b>	<b>57.5 / 33.5</b>	<b>59.5 / 35.7</b>	<b>60.6 / 35.7</b>	<b>60.7 / 35.8</b>
Continuous	ImVoteNet	31.6 / 13.5	18.1 / 8.8	25.2 / 13.5	30.4 / 15.7	40.5 / 18.4	43.1 / 20.0	43.3 / 20.4
	3DETR-m(10)	29.8 / 12.5	19.9 / 9.9	27.7 / 15.3	34.0 / 18.1	44.2 / 24.1	48.7 / 27.8	48.9 / 27.7
	3DETR-m(30)	29.1 / 13.1	18.5 / 9.6	26.6 / 15.1	35.4 / 19.7	48.0 / 29.4	54.0 / 33.6	55.6 / 34.7
	3DETR-m(50)	27.7 / 12.8	17.9 / 9.7	26.6 / 15.8	34.9 / 19.8	46.9 / 28.8	52.7 / 33.0	54.7 / 35.3
	AnyView(10)	31.3 / 13.7	20.6 / 10.6	30.2 / 15.7	36.0 / 17.7	47.9 / 26.1	53.6 / 28.7	54.8 / 30.1
	AnyView(30)	30.5 / 12.7	19.9 / 10.1	29.9 / 16.5	36.8 / 21.4	52.3 / <b>30.0</b>	57.1 / 34.3	60.1 / <b>37.4</b>
	AnyView(50)	<b>32.7 / 15.3</b>	<b>22.1 / 11.3</b>	<b>32.9 / 18.6</b>	<b>40.6 / 21.8</b>	<b>54.4 / 29.8</b>	<b>58.7 / 34.4</b>	<b>60.7 / 35.8</b>



Fig. 3. Visual results on ScanNet. We compare the predictions of 3DETR-m and AnyView after NMS with the ground-truth bounding boxes on different scales of input. Top: the whole scene consisting of 50 views, Bottom: a part of scene consisting of 10 adjacent views. AnyView successfully detects all the doors without color information, which shows its strong ability for extracting local geometric information. The predictions of AnyView is also more consistent across different input scales.

frame  $f_i$ , at the input of the next RGB-D frame  $f_{i+1}$ , it will directly read  $\{\mathcal{P}_1, \dots, \mathcal{P}_i\}$ , which will be fed into the detector along with  $\mathcal{P}_{i+1}$  to predict the result.

### B. Results and Analysis

**On scene-level and monocular benchmarks:** We show the performance of different models on scene-level (Rec and MV) and monocular (SV) benchmarks in Table III. As discussed in Section III-A, previous models trained on ScanNet-Rec or ScanNet-SV generalizes poorly to other settings. Although 3DETR-m trained on ScanNet-Rec achieves the best performance (65.0/47.0) on scene-level benchmarks when evaluated on ScanNet-Rec, it requires reconstructed point clouds which are not available in many practical scenarios. We find 3DETR-m trained on ScanNet-MV gets lower performance (54.7/35.3) on scene-level benchmarks, which indicates the previous ScanNet-Rec benchmark is too

idealistic and the point clouds fused from multi-view depth maps are more challenging for 3D object detection. Among models trained on ScanNet-MV, ImVoxelNet gets relatively lower performance as color is less informative than depth, especially for 3D object detection. 3DETR-m generalizes better to single view inputs (54.7/35.3) when trained on ScanNet-MV, which shows ScanNet-MV is not only a more practical benchmark for evaluation, but also beneficial for training a generalizable 3D detector. Observing the rows in gray, AnyView achieves leading result (60.7/35.8) on the challenging ScanNet-MV benchmark and also generalizes well (32.7/15.3) to ScanNet-SV, which even outperforms ImVoteNet (31.6/13.5) trained on ScanNet-SV.

**On wider range of input views:** We further extend ScanNet-MV to a series of fine-grained benchmarks for more comprehensive analysis. We train models on 10/30/50 views uniformly sampled from the whole 50 views in ScanNet-

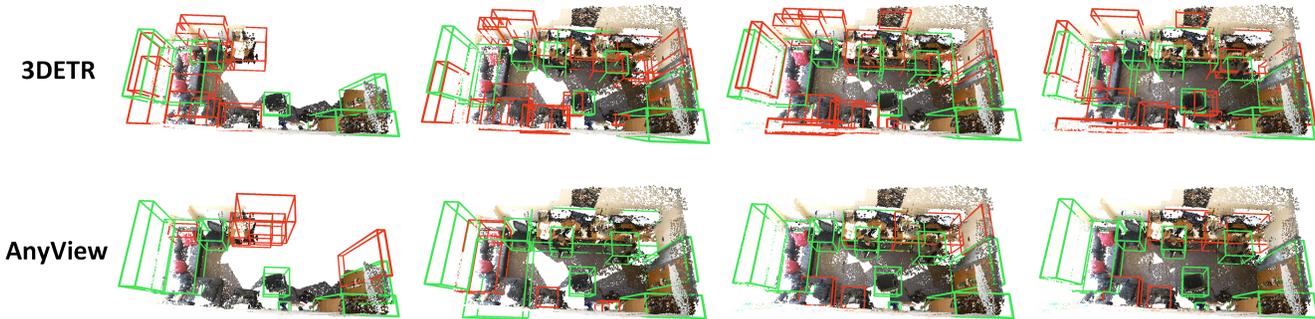


Fig. 4. Online detection visualization results. We demonstrate the detection results for 15, 30, 40, and 50 frames during the navigation(scene0474\_00 on the ScanNet dataset).Red bounding boxes and green bounding boxes represent false-positive and true-positive samples, respectively.

MV and evaluate them on different numbers of views. We devise two settings for partial scenes: Uniform and Continuous, indicating the evaluation of model is conducted on uniformly/adjacently sampled views. As shown in Table IV, 3DETR-m trained on small number of views gets low performance on large number of views, and vice versa. On the contrary, with the growing of numbers of training views, AnyView shows consistent improvement across various scales of input. Even if combining the best results of three 3DETR-m models, AnyView still surpasses them by a large margin with a single set of parameters.

The detection results of 3DETR-m and AnyView after NMS is shown in Fig. 3, where a whole scene consisting of 50 views and a part of it consisting 10 adjacent views are chosen as the input. 3DETR-m fails to detect any of the doors for the whole scene. That is because 3DETR-m fuses multi-view point clouds as a whole, which makes the detector hard to focus on thin objects like doors and windows whose geometric information is weak. It also aggregates outliers into misleading clusters, resulting in false positive. On the contrary, AnyView extracts scene proxies for each view independently, which owns better understanding of local geometry structures and successfully detects all the doors for both scenes. Benefiting from the input scale-independent property, AnyView outputs consistent predictions for different input scales.

**Online 3D object detection:** We conduct online object detection experiments to demonstrate the performance of AnyView with any-scale point cloud input as shown in Fig. 4. The views of each scene are fed into AnyView in a frame-by-frame format, and the aggregated point cloud reconstruction results are fed into the 3DETR.

### C. Ablation Study

We conduct ablation studies to show how different architecture designs and training/inference strategies influence the performance of the proposed AnyView framework. The models are trained on ScanNet-MV with 30 views. With a high degree of site cloud fragmentation (15 views), 3DETR exhibits more false positives, while AnyView maintains high detection accuracy. As the number of viewpoints increases, 3DETR’s false positives are somewhat reduced, achieving better results at 50 views.

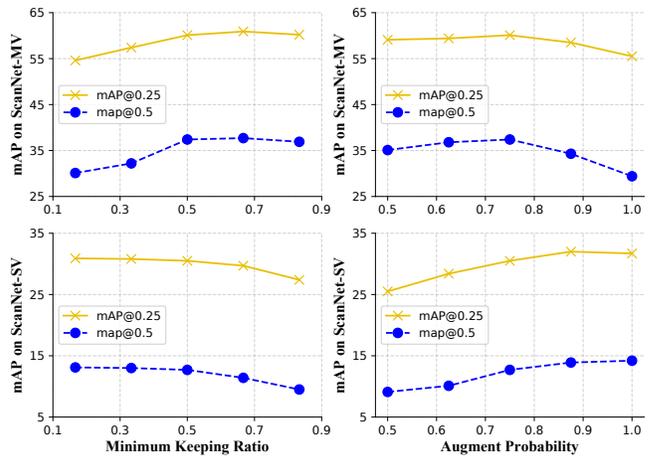


Fig. 5. The effects of minimum keeping ratio for random view dropping and probability for data augmentation on the final performance of AnyView on ScanNet-MV and ScanNet-SV.

**Performance w.r.t. architecture design:** We ablate three parts of architecture designs of AnyView, as shown in Table V. GlobalQuery indicates the object queries are sampled from the concatenated scene ( $\checkmark$ ) instead of the coordinates of scene proxies. CoordsTrans means whether to apply geometry learner in the camera coordinate ( $\checkmark$ ) of each view or in the world coordinate. PE<sub>enc</sub> indicates whether to use positional embeddings for the transformer encoder. The first two rows show that high object query coverage of the scene is beneficial for the transformer decoder. The third row shows unifying the coordinate makes the geometry learner extract more robust scene proxies. As the features of point clouds already contain positional information, 3DETR finds positional embeddings are not necessary. However, in AnyView the scene proxies are independently extracted, so the spatial relationship between scene proxies from different views is weakened. We find Fourier positional embeddings with a MLP performs best.

**Performance w.r.t. data augmentations:** We further investigate the effects of two proposed data augment strategies. ScanNet-MV measures the detection accuracy and ScanNet-SV measures the generalizability of the detector. As shown in Fig. 5, with the growing of minimum keeping ratio for random view dropping, the performance of AnyView on ScanNet-MV grows but its generalizability to single view input has deteriorated. In terms of the probability of augmentation, it is shown that low probability leads to overfitting

TABLE V

The effects of global query sampling, coordinate transformation and encoder positional embeddings on the AnyView with ScanNet-MV.

GlobalQuery	CoordsTrans	PE <sub>enc</sub>	mAP	
			@0.25	@0.5
			58.5	33.1
✓			59.2	34.4
✓	✓		59.6	34.9
✓	✓	Fourier	60.6	35.8
✓	✓	MLP◦Fourier	60.1	37.4

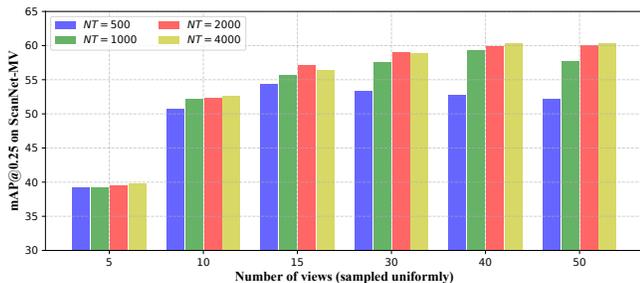


Fig. 6. The effects of token number during inference on the final performance of AnyView on ScanNet-MV.

and poor generalizability, while very high probability hurts the performance of AnyView. Therefore we choose dropping  $[0, \frac{N}{2}]$  views and augmenting with probability 0.75.

**Performance w.r.t. number of tokens:** Fig. 6 illustrates the effects of our dynamic token strategy, where the number of scene proxies for each view is  $T = \min\{X/N, 256\}$ ,  $X \in \{500, 1000, 2000, 4000\}$ . We find when the number of scene proxies is too small ( $N_T \leq 1000$ ), the performance of AnyView will even drop with the growth of view numbers. With the increasing number of  $N_T$ , the performance of AnyView is significantly improved for any view number. However, when  $N_T$  is greater than 2000, the rate of performance increase decreases significantly, introducing additional computational cost without sufficient performance improvement, and  $N_T = 4000$  even leads to a decrease in accuracy. So finally we set  $N_T = 2000$  for a better trade-off between performance and cost.

## VI. CONCLUSION

In this paper, we challenge the existing benchmarks for indoor 3D object detection in a perspective of input data. We propose a new practical setting for this task, which unifies the input modality as multi-view RGB-D images with variable input frame numbers and evaluates 3D detectors on various scales of input data. We design a new transformer-based framework named AnyView for practical applications, which is able to process scenes consisting of any number of frames and extract input scale-independent scene representations. Benefiting from the nature of the self-attention mechanism and our scale-independent training strategy, AnyView is able to change the number of scene representations extracted for each input frame during inference and flexibly handle various numbers of input frames. Extensive experiments on the ScanNet dataset show that AnyView achieves both great generalizability and high detection accuracy while containing a similar amount of parameters with the baselines.

## REFERENCES

- [1] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229. Springer, 2020.
- [2] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, pages 2432–2443, 2017.
- [3] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *CVPR*, pages 9224–9232, 2018.
- [4] Dinh-Cuong Hoang, Johannes A Stork, and Todor Stoyanov. Voting and attention-based pose relation learning for object pose estimation from 3d point clouds. *RAL*, 7(4):8980–8987, 2022.
- [5] Ji Hou, Angela Dai, and Matthias Nießner. 3d-sis: 3d semantic instance segmentation of rgb-d scans. In *CVPR*, pages 4421–4430, 2019.
- [6] Ammar Karkour, Khaled A Harras, and Eduardo Feo-Flushing. Text2map: From navigational instructions to graph-based indoor map representations using llms. In *IROS*, pages 1153–1160. IEEE, 2024.
- [7] Abhijit Kundu, Yin Li, and James M Rehg. 3d-rcnn: Instance-level 3d object reconstruction via render-and-compare. In *CVPR*, pages 3559–3568, 2018.
- [8] Chengyang Li, Yulai Zhang, Zhiqiang Yu, Xinming Liu, and Qing Shi. A robust visual slam system for small-scale quadruped robots in dynamic environments. In *IROS*, pages 321–326. IEEE, 2024.
- [9] Or Litany, Tal Remez, Daniel Freedman, Lior Shapira, Alex Bronstein, and Ran Gal. Asist: automatic semantically invariant scene transformation. *CVIU*, 157:284–299, 2017.
- [10] Ishan Misra, Rohit Girdhar, and Armand Joulin. An end-to-end transformer model for 3d object detection. In *ICCV*, pages 2906–2917, 2021.
- [11] Charles R Qi, Xinlei Chen, Or Litany, and Leonidas J Guibas. Invotenet: Boosting 3d object detection in point clouds with image votes. In *CVPR*, pages 4404–4413, 2020.
- [12] Charles R. Qi, Or Litany, Kaiming He, and Leonidas J. Guibas. Deep hough voting for 3d object detection in point clouds. In *ICCV*, pages 9277–9286, 2019.
- [13] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, pages 652–660, 2017.
- [14] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, pages 5099–5108, 2017.
- [15] Junjing Qiu, Fei Wang, and Zheng Dang. Multi-source fusion for voxel-based 7-dof grasping pose estimation. In *IROS*, pages 968–975. IEEE, 2023.
- [16] Danila Rukhovich, Anna Vorontsova, and Anton Konushin. Imvoxelnet: Image to voxels projection for monocular and multi-view general-purpose 3d object detection. In *WACV*, pages 2397–2406, 2022.
- [17] Jonas Schramm, Niclas Vödisch, Kürsat Petek, B Ravi Kiran, Senthil Yogamani, Wolfram Burgard, and Abhinav Valada. Bevcars: Camera-radar fusion for bev map and object segmentation. In *IROS*, pages 1435–1442. IEEE, 2024.
- [18] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *CVPR*, pages 567–576, 2015.
- [19] Shuran Song and Jianxiong Xiao. Sliding shapes for 3d object detection in depth images. In *CVPR*, pages 634–651, 2014.
- [20] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 33:7537–7547, 2020.
- [21] Guofeng Tong, Zheng Li, Hao Peng, and Yaqi Wang. Multi-source features fusion single stage 3d object detection with transformer. *RAL*, 8(4):2062–2069, 2023.
- [22] Vincent van der Brugge, Marc Pollefeys, Joshua B Tenenbaum, Krishna Murthy Jatavallabhula, and Ayush Tewari. Pickscan: Object discovery and reconstruction from handheld interactions. In *IROS*, pages 10167–10173. IEEE, 2024.
- [23] Shuhua Wang, Ke Lu, Jian Xue, and Yang Zhao. Da-net: Density-aware 3d object detection network for point clouds. *TMM*, 2023. doi:10.1109/TMM.2023.3245359.
- [24] Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *CVPR*, pages 244–253, 2018.
- [25] Zheyuan Zhou, Liang Du, Xiaoqing Ye, Zhikang Zou, Xiao Tan, Li Zhang, Xiangyang Xue, and Jianfeng Feng. Sgm3d: stereo guided monocular 3d object detection. *RAL*, 7(4):10478–10485, 2022.