

Open Panoramic Segmentation

Junwei Zheng¹, Ruiping Liu¹, Yufan Chen¹,
Kunyu Peng¹, Chengzhi Wu¹, Kailun Yang²,
Jiaming Zhang^{1,†}, and Rainer Stiefelhagen¹

¹ Karlsruhe Institute of Technology

² Hunan University

Abstract. Panoramic images, capturing a 360° field of view (FoV), encompass omnidirectional spatial information crucial for scene understanding. However, it is not only costly to obtain training-sufficient dense-annotated panoramas but also application-restricted when training models in a close-vocabulary setting. To tackle this problem, in this work, we define a new task termed **Open Panoramic Segmentation (OPS)**, where models are trained with FoV-restricted pinhole images in the source domain in an open-vocabulary setting while evaluated with FoV-open panoramic images in the target domain, enabling the zero-shot open panoramic semantic segmentation ability of models. Moreover, we propose a model named OOOPS with a Deformable Adapter Network (DAN), which significantly improves zero-shot panoramic semantic segmentation performance. To further enhance the distortion-aware modeling ability from the pinhole source domain, we propose a novel data augmentation method called Random Equirectangular Projection (RERP) which is specifically designed to address object deformations in advance. Surpassing other state-of-the-art open-vocabulary semantic segmentation approaches, a remarkable performance boost on three panoramic datasets, WildPASS, Stanford2D3D, and Matterport3D, proves the effectiveness of our proposed OOOPS model with RERP on the OPS task, especially **+2.2%** on outdoor WildPASS and **+2.4%** mIoU on indoor Stanford2D3D. The source code is publicly available at OPS.

1 Introduction

Panoramic imaging systems [26, 27] have advanced significantly in recent years, which has fostered a wide variety of panoramic vision applications [1, 17]. Due to the comprehensive 360° Field of View (FoV), omnidirectional panoramas provide more informative visual cues when perceiving surroundings in a broad spectrum of scene understanding tasks [3, 36, 51, 52, 54, 79], enabling a more complete and immersive capture of environmental data, which is crucial for in-depth scene understanding. This wide perspective surpasses the limited scope of pinhole images, significantly enhancing the capability of computer vision systems to perceive and interpret surroundings in a variety of applications. While the benefits of utilizing panoramic images in computer vision applications are apparent compared

[†] Correspondence: jiaming.zhang@kit.edu

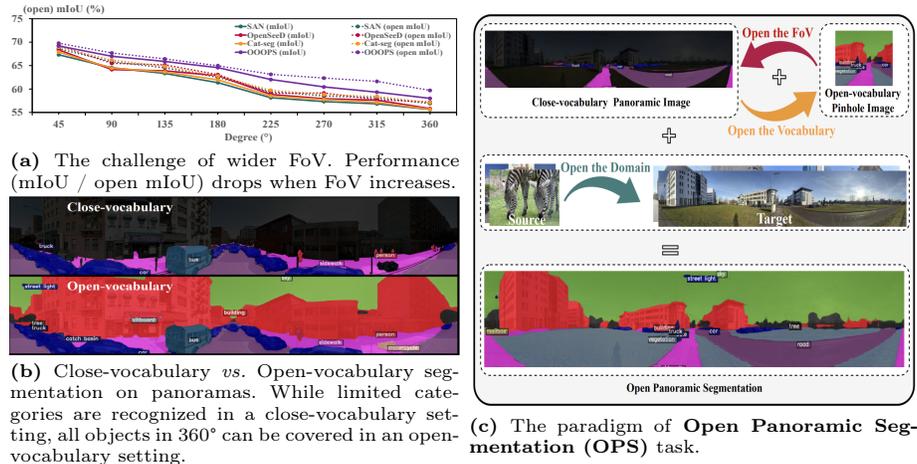


Fig. 1: (a) The challenge of existing state-of-the-art segmentation models. (b) The limitation of categories in traditional close-vocabulary panoramic segmentation tasks. (c) Our newly defined **Open Panoramic Segmentation (OPS)** task aims at tackling the above challenges. OPS consists of three important elements: **Open the FoV** targeted at the challenge of 360° FoV, **Open the Vocabulary** targeted at the drawback of close-vocabulary panoramic segmentation and **Open the Domain** targeted at the challenge of scarcity of panoramic labels.

with pinhole images, it is imperative to give increasing consideration to some noteworthy challenges as follows: (1) *The challenge of wider FoV*. For example, Fig. 1a illustrates the degraded performance of state-of-the-art open-vocabulary semantic segmentation approaches [9, 68, 80] with an increasing FoV, *i.e.*, from pinhole images to 360° panoramic images. Over 12% mIoU performance degradation can be observed, indicating the challenges introduced by the large divergence of semantic and structural information across narrow-wide imagery. (2) *The limitation of categories*. The traditional close-vocabulary segmentation task paradigm [74, 82] only provides a limited amount of labeled categories, which cannot handle the inestimable amount of categories in real-world applications. Fig. 1b illustrates the difference between close- and open-vocabulary segmentation. Compared to the close-vocabulary setting (first row) the open-vocabulary setting (second row) is not limited by the number of categories of the datasets. Only four predefined categories (highlighted in different colors) are recognized in the close-vocabulary setting in contrast to the open one, where every pixel of the panorama has its own semantic meaning even though the categories are not annotated in the dataset.

To further release the great potential of panoramic images three critical problems are eager to be solved: **How** to obtain a holistic perception via a single image? **How** to break the barrier of limited recognizable categories in existing panoramic datasets so that downstream vision applications can benefit from unrestricted informative visual hints? **How** to deal with the scarcity of

panoramic labels? Based on the aforementioned three questions, we propose a new task named **Open Panoramic Segmentation (OPS)** that comprehensively addresses these challenges, aiming to better leverage the advantages brought by panoramic images. The new task paradigm is shown in Fig. 1c. The open panoramic segmentation task consists of three important elements considering the three problems. **The concept of “Open” is three-fold**: omnidirectional panoramic images (**Open the FoV**), an unrestricted range of recognizable classes (**Open the Vocabulary**), and additionally, models are trained using pinhole images in the source domain while evaluated with panoramic images in the target domain (**Open the Domain**). Given that densely annotated pinhole segmentation labels incur lower costs compared to panoramic ones [74], it is cost-efficient to open different domains. Note that OPS is different from Domain Adaptation (DA). In the context of training in DA, the utilization of data encompasses both the source and target domains while in OPS, exclusive reliance is placed on data originating solely from the source domain during the whole training process.

Apart from the new task, we propose a new model named OOOPS to address the aforementioned three openness-related challenges in the OPS task. It consists of a frozen CLIP model [49] and a key component named *Deformable Adapter Network (DAN)*, which serves two critical purposes: (1) efficiently adapting the frozen CLIP model to the panoramic segmentation task and (2) addressing object deformation and image distortion inherent in panoramas. More specifically, a novel Deformable Adapter Operator (DAO), the key component of DAN, is designed to tackle panoramic distortion. To advance the modeling capabilities for distortion awareness in the pinhole source domain, we further introduce *Random Equirectangular Projection (RERP)* explicitly crafted to tackle object deformations and image distortions. A pinhole image is divided into four image patches that are randomly shuffled. Equirectangular Projection, one of the most common methods for mapping a globe into a panoramic plane, is then introduced in the shuffled image. Our OOOPS model with RERP outperforms other state-of-the-art open-vocabulary segmentation methods with **+2.2%**, **+2.4%**, and **+0.6%** mIoU on WildPASS [76], Stanford2D3D [2], and Matterport3D [6], respectively.

To summarize, we present the following contributions:

- We introduce a new task termed open panoramic segmentation, **OPS** for short, including **Open FoV**, **Open Vocabulary** and **Open Domain**. Models are trained with FoV-restricted pinhole images in the source domain in an open-vocabulary setting while evaluated with FoV-open panoramic images in the target domain.
- We put forward a model called OOOPS with the aim to address three openness-related challenges at once. A **Deformable Adapter Network (DAN)** is proposed to transfer the zero-shot learning ability of a frozen CLIP model from the pinhole domain to different panoramic domains.
- A novel data augmentation strategy named **Random Equirectangular Projection (RERP)**, which is specifically designed for the proposed OPS

task, further boosts the accuracy of the OOOPS model, achieving state-of-the-art performance in the open panoramic segmentation task.

- For benchmarking OPS, we conduct comprehensive evaluations on both indoor and outdoor datasets (WildPASS, Stanford2D3D and Matterport3D), involving more than 10 close- and open-vocabulary segmentation models.

2 Related Work

2.1 Open-vocabulary Semantic Segmentation

Motivated by the success of vision-language pre-training models, Open-Vocabulary Semantic Segmentation (OVSS) has been recently explored to achieve open-set pixel-wise understanding [7, 11, 12, 18, 19, 29, 31, 32, 34, 34, 37, 39, 40, 44, 53, 56, 58, 59, 61, 64, 70, 71]. Most works leverage the open-vocabulary capabilities of CLIP [49] to reason about relations between dense pixel semantics and arbitrary class labels [9, 12, 32, 38, 61, 71]. In [78, 89], multiple datasets are unified by transferring heterogeneous labels into text embeddings for training to achieve OVSS. In [35, 67], image-caption datasets are employed for learning alignment between visual regions and text entities. To reduce the computational costs of using a frozen pre-trained CLIP model for identifying novel classes, latest studies [21, 69] have also proposed efficient architectures for fast OVSS. SAN [68, 69] has designed an asymmetric input resolution strategy to decouple mask proposal prediction and attention bias prediction applied to CLIP for mask class recognition. Furthermore, open-vocabulary panoptic segmentation has been tackled in [8, 12, 13, 34, 45, 66, 66, 71, 80] to render universal image understanding. In this work, we focus on the open-vocabulary zero-shot learning on **panoramas** instead of pinhole images discussed in the aforementioned works.

2.2 Panoramic Semantic Segmentation

Panoramic semantic segmentation enables more informative decision-making and interaction within complex and dynamic environments, advancing the capabilities of intelligent systems for navigation and perception [4, 16, 24, 25, 28, 33, 42, 46, 75]. PASS [72, 73] achieves panoramic segmentation on annular images, whereas 360BEV [54] tackles panoramic semantic mapping. ECANets [76] address omnidirectional semantic segmentation with multi-source omni-supervised learning to enhance generalization. Zhang *et al.* [81] propose to use attention-augmented designs for panoramic representation learning. Guttikonda *et al.* [20] propose a single-frame panoramic semantic segmentation approach using multi-modal spherical images. Hu *et al.* [22] propose a distortion convolution module for panoramic semantic segmentation. An important cluster of works revisits panoramic semantic segmentation from an unsupervised domain adaptation perspective [23, 30, 41, 81, 82, 86, 87], by adapting from label-rich pinhole images to the panoramic images. Unlike these works, we propose a method to tackle **open-vocabulary** panoramic semantic segmentation for unconstrained surrounding understanding.

2.3 Deformable Neural Networks

In the realm of deformable neural networks, notable advancements have been made through various iterations. The pioneering work [10] introduces the concept of deformable convolution, laying the foundation for subsequent developments. Building upon this foundation, DCNv2 [90] further refines deformable convolution by introducing a modulation mechanism that expands the scope of deformation modeling. The evolution of deformable neural networks continues with the introduction of DCNv3 [57], which combines the group convolution with DCNv2 for the sake of imitating multi-head mechanism of Transformer [55]. DCNv4 [65] aims to speedup DCNv3, achieving a more efficient model. Additionally, Trans4PASS [82,83] proposes Deformable Patch Embedding (DPE) and Deformable MLP (DMLP) to handle the distortion that occurs in panoramas. Incorporating the principles of deformable convolution into vision transformers [14], Xia *et al.* propose DeformableViT [62,63], extending the applicability of deformable neural networks to the domain of transformer architectures. After rethinking a series of DCN designs, we propose **DAO** to improve the model’s ability to handle image distortion in the OPS task.

3 Methodology

3.1 Open Panoramic Segmentation

Open Panoramic Segmentation (OPS) task is proposed to deal with three challenging problems: ❶ the narrow FoV, ❷ the restricted range of classes, and ❸ the scarcity of panoramic labels. OPS gives three answers to the aforementioned issues: ❶ Open the FoV, ❷ Open the Vocabulary, and ❸ Open the Domain. The OPS task paradigm is shown in Fig. 1c. Models are trained in the narrow-FoV pinhole source domain in an open-vocabulary setting while evaluated in the wide-FoV panoramic target domain.

3.2 Model Architecture

A foundation model [49] can be efficiently transferred to downstream tasks by using an adapter [68]. To improve the modeling ability for panoramas, we design the OOOPS model. As shown in Fig. 2, it comprises a frozen CLIP model and a proposed Deformable Adapter Network (DAN) that incorporates multiple Transformer Layers and the novel DAO, which will be detailed in Sec. 3.3. Feature fusion takes place between the intermediate layers of CLIP and DAN. One of the two outputs from DAN consists of mask proposals, while the other serves as deep supervised guidance for CLIP to generate proposal logits. In the training phase, pinhole images are forwarded into OOOPS, generating mask proposals and proposal logits for the loss calculation. In the inference phase, panoramas are forwarded into OOOPS to generate segmentation predictions via the multiplication of mask proposals and the corresponding proposal logits. The frozen CLIP is necessary for the zero-shot learning ability of OOOPS.

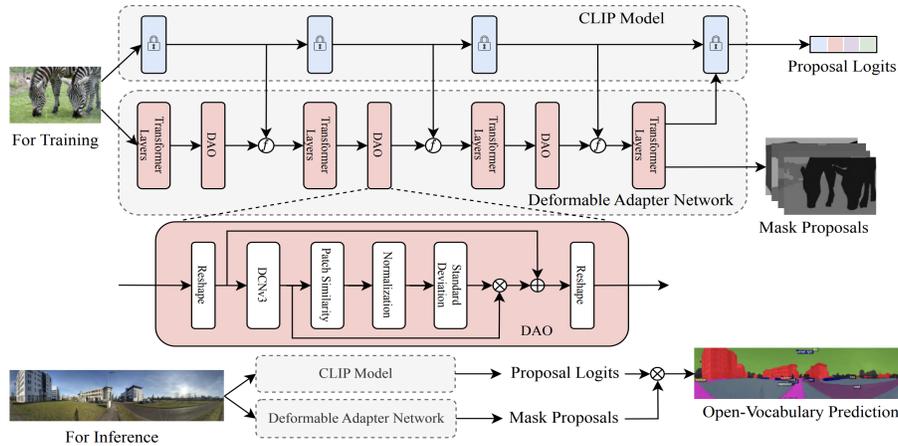


Fig. 2: Overview of the OOOPS model architecture. It consists of a frozen CLIP model and a Deformable Adapter Network (DAN) which includes Transformer Layers and the proposed DAO.

3.3 Deformable Adapter Network

The Deformable Adapter Network is a combination of multiple Transformer [55] Layers and the proposed DAO. Since distortion exists in panoramic images, which is a big challenge when utilizing informative panoramas, we delve deeper into the deformable design [57] and sampling methods like APES [60], MateRobot [85], proposing DAO to tackle the image distortions and object deformations of panoramas. To provide a detailed rethinking process of the deformable convolution design, we introduce the preliminary of the DCN series as follows.

Revisiting DCN Series. The seminal work DCN [10] can enable the traditional CNN with spatial deformation-aware ability. Given a convolutional kernel of K sampling locations, let \mathbf{w}_k and \mathbf{p}_k denote the weight and pre-specified offset for the k -th location, respectively. For example, $K=9$ and $\mathbf{p}_k \in \{(1, 1), \dots, (-1, -1)\}$ defines a 3×3 convolutional kernel of dilation 1. Let $\mathbf{x}(\mathbf{p})$ and $\mathbf{y}(\mathbf{p})$ denote the features at location \mathbf{p} from the input feature maps \mathbf{x} and output feature maps \mathbf{y} , respectively. DCN is formulated as:

$$\mathbf{y}(\mathbf{p}) = \sum_{k=1}^K \mathbf{w}_k \mathbf{x}(\mathbf{p} + \mathbf{p}_k + \Delta \mathbf{p}_k), \quad (1)$$

where $\Delta \mathbf{p}_k$ is the learnable offset for the k -th location. Although DCN is capable of capturing spatial deformation, every sampling location is treated equally when calculating the local features. DCNv2 [90] is proposed with an additional term called modulation scalar. Specifically, DCNv2 can be formulated as:

$$\mathbf{y}(\mathbf{p}) = \sum_{k=1}^K \mathbf{w}_k \mathbf{m}_k \mathbf{x}(\mathbf{p} + \mathbf{p}_k + \Delta \mathbf{p}_k), \quad (2)$$

where \mathbf{m}_k is the learnable modulation scalar for the k -th location. Inspired by Transformer, DCNv3 [57] is proposed with a grouping operation, further boosting the deformation-aware ability of DCNv2. DCNv3 can be addressed with the following formula:

$$\mathbf{y}(\mathbf{p}) = \sum_{g=1}^G \sum_{k=1}^K \mathbf{w}_g \mathbf{m}_{gk} \mathbf{x}_g(\mathbf{p} + \mathbf{p}_k + \Delta \mathbf{p}_{gk}), \quad (3)$$

where G denotes the total number of aggregation groups. DCNv4 [65] is similar to DCNv3, achieving similar performance while significantly reducing the runtime.

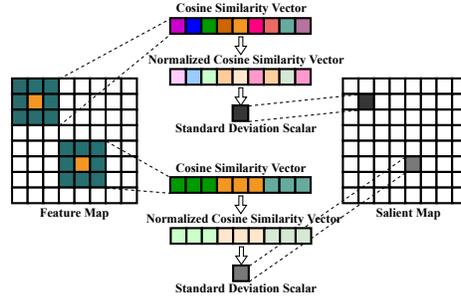


Fig. 3: Salient map generation in DAO.

Deformable Adapter Operator (DAO). When dealing with distortions in panoramas, DCNv3 and DCNv4 fall short of meeting the distortion-aware requirements. This is illustrated by the visualization of Fig. 6 in Sec. 4.5. Therefore, DAO is proposed to tackle the distortion problem in panoramic images using the following expression:

$$\mathbf{y}(\mathbf{p}) = \mathbf{s}(\mathbf{p}) \sum_{g=1}^G \sum_{k=1}^K \mathbf{w}_g \mathbf{m}_{gk} \mathbf{x}_g(\mathbf{p} + \mathbf{p}_k + \Delta \mathbf{p}_{gk}), \quad (4)$$

where $\mathbf{s}(\mathbf{p})$ is the learnable salient scalar at location \mathbf{p} . Inheriting from DCNv3, DAO comes up with an additional term called salient scalar indicating the importance of each pixel within the whole panorama. It's worth noting that DCNv3 and DCNv4 share the same mathematical expression, but DCNv3 is **more robust** according to our experiments before designing DAO. Therefore, we adopt DCNv3 as a part of DAO rather than DCNv4. Illustrated in Fig. 2, the feature map outputted by DCNv3 passes through the Patch Similarity Layer, Normalization Layer, and Standard Deviation Layer sequentially to form a salient map. The intuition for such a design is straightforward. The salient pixels in an image are those that differ significantly from their neighboring pixels, *e.g.*, edge pixels. If all pixels within an image patch are different, the standard deviation of the

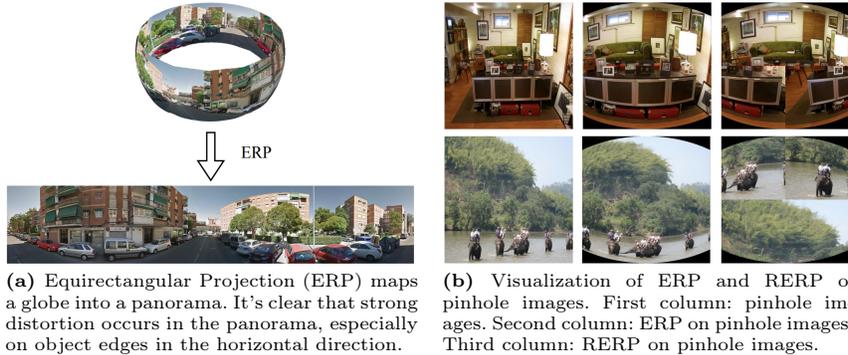


Fig. 4: (a) Visualization of ERP on a panoramic image and (b) Our proposed RERP on pinhole images.

pixel similarity for this patch is higher than the one containing similar pixels, resulting in a higher salient scalar. Fig. 3 gives a more detailed explanation of the salient map generation. Given a feature map, DAO first calculates the cosine similarity between the center pixel and all pixels within a kernel, *e.g.*, 9 pixels within a 3×3 kernel in Fig. 3, resulting in a 9-dimensional cosine similarity vector. Softmax normalization is then applied to the vector. Afterward, DAO calculates the standard deviation of this normalized cosine similarity vector, indicating the importance of the center pixel. By traversing every pixel across the entire feature map, a salient map is generated for the sake of enhancing the salient pixels, which are usually the edge pixels of an image where strong panoramic distortion occurs.

3.4 Random Equirectangular Projection

The Equirectangular Projection [50] (ERP) is one of the most common methods for mapping a globe into a panoramic plane, which transforms spherical coordinates into planar coordinates expressed as follows:

$$x = R(\lambda - \lambda_0)\cos(\varphi_1), \quad (5)$$

$$y = R(\varphi - \varphi_0), \quad (6)$$

where λ and φ are the longitude and latitude of the location to project, respectively. φ_1 are the standard parallels. λ_0 and φ_0 are the central meridian and central parallel of the map, respectively. R is the radius of the globe. x denotes the horizontal coordinate, while y represents the vertical coordinate of the projected location on the map. Fig. 4a visualizes the ERP on a panoramic plane. It can be observed that after equirectangular projection strong distortion occurs in the panorama, *e.g.*, a straight line is transformed into a curve. To further boost the performance, we propose Random Equirectangular Projection (RERP) on pinhole images since the OPS task requires the model to train with pinhole images instead of panoramas.

We divide the pinhole image into four parts and randomly shuffle the image patches. Afterward, we apply equirectangular projection to the distortion-free pinhole image. Fig. 4b visualizes pinhole images after RERP. The first column is the pinhole images without any data augmentation. The second column is ERP on pinhole images without random shuffling. The last column is our proposed RERP. It can be observed that the panorama-like distortion also occurs in the pinhole images after RERP. Random shuffling is employed to enhance robustness and promote generalization.

4 Experiments

4.1 Datasets

We train our model using the pinhole general-object COCOStuff-164k [5] while evaluating the model performance on three different panoramic datasets: panoramic street-view WildPASS [76], panoramic indoor-scene Stanford2D3D [2] and panoramic indoor-scene Matterport3D [6].

COCOStuff-164k. The pinhole dataset comprises 164k images across 171 annotated classes, partitioned into training, validation, and test sets, with 118k, 5k, and 41k images, respectively. In our experiments, we utilize the entire 118k-image training set as the training data by default.

WildPASS. The panoramic street-view dataset contains 2,500 panoramas in 65 cities from all continents Asia, Europe, Africa, Oceania, and North and South America (excluding only Antarctica) with 8 dense-annotated classes. All images have a $70^\circ \times 360^\circ$ FoV in the size of 400×2048 .

Stanford2D3D. The panoramic indoor dataset consists of 1,413 images in the size of 512×1024 and 13 object classes. All results are averaged over 3 cross-validation folds.

Matterport3D. The panoramic indoor dataset includes 10,800 panoramic views from 194,400 RGB-D images. According to the data preprocessing pipeline of 360BEV [54], a subset of 772 images in the size of 512×1024 with 20 classes is used to evaluate the segmentation performance.

4.2 Metrics

Close-vocabulary Metric. Mean Intersection over Union (mIoU) is used as the evaluation metric in close-vocabulary semantic segmentation tasks. Specifically, IoU is calculated as the area of overlap between the predicted segmentation and the ground truth, divided by the area of union between them. mIoU is a hard metric designed for close-vocabulary segmentation ignoring the semantic meanings of similar classes referring to the open-vocabulary setting.

Open-vocabulary Metric. Open mean Intersection over Union [88], open mIoU for short, is a soft metric compared to mIoU since the similarity scores between different classes are also taken into account, which are calculated by the WordNet approach [43]. In this work, we will report both close and open metrics for a comprehensive understanding of all models.

Table 1: Results on WildPASS. The mIoU, open mIoU, and per-class IoU are in percentage (%). #PARAMs is the number of learnable parameters in millions.

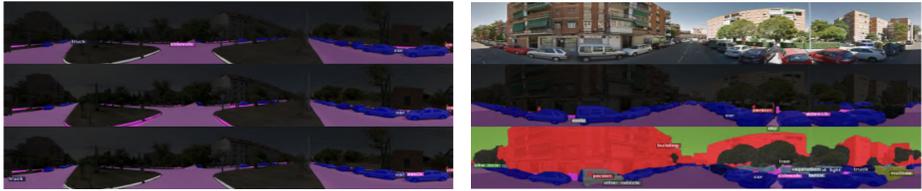
Network	#PARAMs	mIoU	Open mIoU	Car	Road	Sidewalk	Crosswalk	Curb	Person	Truck	Bus
Close-Vocabulary											
Fast-SCNN [48]	1.1	24.8	-	45.9	60.0	31.7	9.7	17.1	6.0	14.2	13.8
PSPNet18 [84]	17.5	28.2	-	58.2	66.8	28.4	13.0	19.2	6.2	18.2	15.6
SwiftNet [47]	11.8	30.0	-	55.7	64.1	29.2	16.2	22.9	8.5	21.1	22.2
DenseASPP [77]	8.3	33.2	-	51.1	69.1	38.1	16.4	26.3	8.7	27.4	28.4
ERF-PSPNet [73]	2.5	34.0	-	66.3	70.5	36.5	6.4	24.1	9.4	26.5	32.0
PSPNet50 [84]	53.3	46.1	-	80.0	74.9	51.7	23.9	31.4	19.8	38.9	48.1
DANet [15]	47.4	47.2	-	74.8	72.2	49.9	28.9	23.8	25.4	51.9	50.6
OOSS (ERF-PSPNet) [74]	2.5	56.1	-	87.2	79.3	60.8	28.0	38.1	54.5	48.8	52.2
PASS (ERF-PSPNet) [73]	2.5	64.7	-	87.3	80.0	61.4	71.1	49.9	72.2	37.5	57.9
ERF-PSPNet (omni-sup) [76]	2.5	66.8	-	90.5	82.7	65.6	70.5	51.5	58.2	62.0	53.1
ECANet (with attention) [76]	2.6	67.7	-	90.4	83.7	68.4	67.8	52.1	61.4	54.5	63.3
ECANet (with fusion) [76]	2.6	69.0	-	90.6	85.7	68.0	67.9	52.1	66.0	59.3	62.3
Open-Vocabulary											
SAN [68]	8.4	55.6	56.9	84.3	81.9	69.6	10.1	2.8	72.2	54.9	69.1
CAT-seg [9]	59.5	55.7	57.1	84.1	81.5	70.0	10.2	3.0	72.5	54.8	69.3
OpenSeeD [80]	65.4	55.8	57.2	84.3	81.6	69.7	10.3	3.0	73.6	54.9	69.1
OOOPS (ours, w/o RERP)	8.7	57.0	58.5	84.9	81.6	68.7	15.1	8.4	71.3	55.3	70.4
OOOPS (ours, w/ RERP)	8.7	58.0	59.7	84.9	81.6	68.7	19.4	8.3	73.8	54.9	72.2

4.3 Implementation Details

We utilize CLIP ViT-B/16 as the frozen CLIP model. All models are trained with the pinhole general-object COCOStuff-164k training set with $4 \times A40$ GPUs while evaluated on three panoramic datasets which are in totally different domains compared with the pinhole dataset. The AdamW optimizer is applied with an initial learning rate of 0.0001, weight decay of 0.0001, batch size of 32, and a total of 60,000 training iterations. Throughout the training, the learning rate follows a polynomial schedule with a power of 0.9. Data augmentations, including random image resizing within the short-side range of $[320, 1024]$ and a crop size of 640×640 are applied. Additionally, when verifying the effectiveness of the proposed RERP, it is applied as an extra data augmentation. Following the practice of SAN [68], apart from the supervision of mask recognition with cross-entropy loss, we utilize the dice loss and binary cross-entropy loss for mask proposal generation.

4.4 Comparison with the State of the Art

Results on WildPASS. Comparisons with both close- and open-vocabulary state-of-the-art methods on WildPASS are listed in Table 1. It is not surprising that ECANet (with fusion) [76] achieves the best performance among all methods since it is a close-vocabulary method, meaning that ECANet is only able to predict the classes predefined by the dataset. However, our OOOPS model is qualified for predicting classes that are not annotated in the dataset. The OOOPS method (with RERP) surpasses other state-of-the-art open-vocabulary panoramic segmentation methods with over **+2.2%** mIoU while introducing



(a) Comparison with baseline method SAN [68] on the WildPASS dataset with predefined 8 classes. First row: ground truth. Second row: prediction from SAN. Third row: prediction from OOOPS with RERP. **(b)** Visualization of the prediction from OOOPS in close- and open-vocabulary settings. First row: RGB. Second row: prediction with close vocabularies. Third row: prediction with open vocabularies. Zoom in for a better view.

Fig. 5: (a) Comparison on the WildPASS dataset and (b) Visualization of the prediction from OOOPS in close- and open-vocabulary settings.

only a minor increment of 0.3 million learnable parameters compared to the lightweight baseline SAN [68], offering a promising step towards bridging the gap between close- and open-vocabulary settings. Moreover, our OOOPS method enhanced by RERP demonstrates significant advancements in challenging categories such as *crosswalk* and *curb*, achieving **+9.3%** and **+5.5%** mIoU compared with the baseline SAN, respectively. It is worth noting that in some categories, *e.g.*, *sidewalk*, *person*, and *bus*, our proposed OOOPS (both with and without RERP) even outperforms the models trained in a close-vocabulary setting. All open-vocabulary methods achieve better performance with open mIoU metric compared to the close mIoU since similar semantic meanings are also considered in the calculation. Fig. 1a showcases the performance degradation on both metrics. Even though all methods have degraded performance when FoV increases, our OOOPS achieves a smaller drop (**-10.0%**) while others have a larger performance decrease (**-12.0%**), indicating the effectiveness of the OOOPS model.

Fig. 5a visualizes the prediction of the challenging *curb* category from baseline SAN and our proposed OOOPS with RERP. From top to bottom are the ground truth, prediction from SAN, and prediction from OOOPS with RERP, respectively. The baseline SAN cannot recognize the challenging *curb* category while our proposed OOOPS with RERP is capable of detecting the *curb* in the middle of the panorama highlighted in gray. Fig. 5b showcases the visualization of the prediction from OOOPS on WildPASS in close- and open-vocabulary settings. From top to bottom are input RGB, prediction from OOOPS with 8 predefined close vocabularies from WildPASS, and prediction from OOOPS with unlimited open vocabularies. It is evident that our OOOPS with RERP is capable of tackling unseen categories in the open-vocabulary setting. Every pixel on the panorama is assigned a meaningful semantic label. For example, “mailbox” is not a predefined category from WildPASS. While existing models can not recognize the *mailbox* under a close-vocabulary setting, the rare classes can be easily detected by our model in an open-vocabulary one, indicating the great potential of open-vocabulary panoramic segmentation for parsing unconstrained traffic surroundings.

Table 2: Results on Stanford2D3D. The mIoU, open mIoU, and per-class IoU are in percentage (%). #PARAMs is the number of learnable parameters in millions.

Network	#PARAMs	mIoU	Open mIoU	Beam	Board	Bookcase	Ceiling	Chair	Clutter	Column	Door	Floor	Sofa	Table	Wall	Window
Close-Vocabulary																
Trans4PASS [82]	28.2	52.1	-	11.4	60.6	51.1	81.8	59.7	37.9	23.5	54.6	88.8	34.6	53.9	68.1	51.0
360BEV [54]	28.4	54.3	-	06.0	63.6	55.5	82.7	62.7	39.9	16.4	58.8	90.8	45.2	57.1	71.1	57.6
Open-Vocabulary																
SAN [68]	8.4	38.5	39.5	0.2	51.5	39.6	66.5	52.1	10.7	0.0	45.0	49.3	31.8	49.2	50.0	53.4
CAT-seg [9]	59.5	38.6	39.6	0.0	51.6	42.3	66.8	51.4	8.0	0.0	47.4	50.7	31.2	47.5	51.4	53.5
OpenSeeD [80]	65.4	38.7	40.0	0.0	51.8	42.0	67.0	51.6	8.4	0.2	47.5	50.0	31.4	47.1	51.7	53.6
OOOPS (w/o RERP)	8.7	39.5	41.0	0.0	52.3	44.1	67.1	53.7	9.7	0.0	50.1	52.8	30.5	47.2	52.2	53.8
OOOPS (w/ RERP)	8.7	41.1	42.6	0.0	50.7	44.8	68.8	51.8	8.1	0.0	48.2	71.1	33.6	47.2	55.4	54.5

Table 3: Results on Matterport3D. The mIoU, open mIoU, and per-class IoU are in percentage (%). #PARAMs is the number of learnable parameters in millions.

Network	#PARAMs	mIoU	Open mIoU	Wall	Floor	Chair	Door	Table	Picture	Furniture	Objects	Window	Sofa	Bed	Sink	Stairs	Ceiling	Toilet	Mirror	Shower	Bathroom	Counter	Shelving
Close-Vocabulary																							
Trans4PASS+ [83]	28.5	42.6	-	63.4	79.1	39.1	40.3	32.8	36.0	31.0	31.5	37.5	44.0	63.2	20.6	41.8	77.6	40.7	24.3	23.7	58.3	34.3	32.9
360BEV [54]	28.4	46.4	-	64.1	83.1	45.8	45.0	38.0	41.1	32.3	35.1	40.6	48.7	69.8	25.1	47.8	80.2	46.0	28.7	22.3	60.1	38.6	34.8
Open-Vocabulary																							
SAN [68]	8.4	30.5	31.4	54.4	76.1	18.0	34.1	22.3	21.1	13.0	1.7	32.7	34.2	51.0	11.0	19.9	74.2	23.4	12.0	31.6	31.9	19.0	28.4
CAT-seg [9]	59.5	30.2	31.1	54.3	74.8	17.0	34.1	21.9	20.8	12.6	1.9	32.9	34.0	50.1	10.7	18.8	73.5	22.4	11.6	31.9	32.9	19.2	28.5
OpenSeeD [80]	65.4	30.6	31.6	54.2	76.3	18.1	34.3	22.5	21.0	13.0	2.1	32.8	34.1	50.9	11.2	20.2	74.1	23.5	12.2	31.1	31.8	19.1	28.6
OOOPS (w/o RERP)	8.7	31.1	32.4	55.2	77.0	17.1	35.7	25.0	17.5	13.5	2.3	33.0	32.4	51.8	12.6	20.4	74.5	22.6	10.8	34.1	36.8	19.4	29.7
OOOPS (w/ RERP)	8.7	31.2	32.5	55.4	76.4	18.4	34.8	23.9	21.2	13.2	2.6	33.3	34.5	52.9	12.1	20.3	74.6	28.2	12.5	31.9	29.3	19.4	28.6

Results on Stanford2D3D. In Table 2, the comparison is conducted between state-of-the-art methods on the Stanford2D3D dataset. All reported results are averaged over 3 cross-validation folds. Without RERP, OOOPS achieves **39.5%** mIoU, outperforming other open-vocabulary methods like SAN [68], CAT-seg [9], and OpenSeeD [80], which score between 38.5% and 38.7%. With RERP, OOOPS’s mIoU increases to **41.1%** with a **+2.4%** gain, narrowing the gap with advanced close-vocabulary methods.

Results on Matterport3D. In Table 3, the comparison is conducted between state-of-the-art methods on the Matterport3D dataset. Our OOOPS method showcases consistent performance. Without augmentation, OOOPS achieves mIoU of **31.1%**, slightly outperforming other open-vocabulary methods like SAN [68] (30.5%), CAT-seg [9] (30.2%), and OpenSeeD [80] (30.6%). With augmentation, OOOPS’s mIoU marginally increases to **31.2%**, further reinforcing its lead over these open-vocabulary methods. Both variants of OOOPS still lag behind close-vocabulary methods like Trans4PASS+ [83] and 360BEV [54], which however, are trained in a supervised fashion in a close-vocabulary setting on panoramas and have mIoUs ranging from 42.6% to 46.4%.

Table 4: Task difference between Domain Adaptation (DA) and the proposed OPS.

Task	Domain		#Classes	mIoU (%)		
	Source	Target		WildPASS	Matterport3D	Stanford2D3D
DA	✓	✓	limited	72.2	47.2	55.4
OPS	✓	✗	unlimited	58.0	31.2	41.1

Table 5: Ablation study of OOOPS components and the number of shuffling patches.

(a) Ablation study of OOOPS components on the WildPASS Dataset. The mIoU is reported in percentage (%). #PARAMs stands for the number of learnable parameters in millions.

Method	ERP	RERP	#PARAMs	mIoU
SAN [68]	✗	✗	8.4	55.6
SAN [68]	✓	✗	8.4	56.0
SAN [68]	✗	✓	8.4	56.5 (+0.9)
+ DCN	✗	✗	8.6	55.6
+ DCN	✓	✗	8.6	56.0
+ DCN	✗	✓	8.6	56.6 (+1.0)
+ DCNv2	✗	✗	8.7	55.7
+ DCNv2	✓	✗	8.7	56.4
+ DCNv2	✗	✓	8.7	56.6 (+0.9)
+ DCNv3	✗	✗	8.7	55.8
+ DCNv3	✓	✗	8.7	56.4
+ DCNv3	✗	✓	8.7	56.8 (+1.0)
+ DCNv4	✗	✗	8.7	55.8
+ DCNv4	✓	✗	8.7	56.4
+ DCNv4	✗	✓	8.7	56.8 (+1.0)
+ DAO (ours)	✗	✗	8.7	57.0
+ DAO (ours)	✓	✗	8.7	57.6
+ DAO (ours)	✗	✓	8.7	58.0 (+1.0)

(b) Ablation study of different numbers of shuffling patches on WildPASS, Matterport3D and Stanford2D3D datasets. The mIoU is reported in percentage (%).

Dataset	#Patches			
	4	9	16	25
WildPASS	58.0	58.5	58.2	58.0
Matterport3D	31.2	31.6	31.4	31.0
Stanford2D3D	41.1	41.9	41.5	41.1

Results in DA Subtask. The proposed OPS task differs from Domain Adaptation since no data from the target domain is available to models in OPS while models have access to the data in DA during training. Moreover, models are required to recognize an unlimited amount of classes in OPS while only a limited number of classes can be recognized in DA. DA is a special case of OPS, meaning that models in OPS can be easily adapted to DA and achieve better performance. The converse doesn't hold. Table 4 showcases the task difference between OPS and DA. It can be observed that over +14% gain in mIoU is achieved when adapting our OOOPS model in the DA subtask.

4.5 Ablation Study

We replace DAO in OOOPS, as illustrated in Fig. 2, with DCN [10], DCNv2 [90], DCNv3 [57] and DCNv4 [65]. When referring to the component substitution without data augmentation in Table 5a, it is clear that our DAO surpasses other counterparts with a maximum performance gain of +1.4% mIoU. As for the component substitution with ERP, all models achieve a boost in the range of +0.3% to +0.7% compared with the one without data augmentation. The gain is even bigger considering the component substitution with RERP, achieving +1.0% mIoU. It is worth noting that RERP is able to boost the model performance stably among all models. The best model performance is achieved by the combination of DAO and RERP, yielding **58.0%** mIoU with only a slight increment in parameters. Table 5b showcases the influence of different numbers of shuffling patches. The performance is not proportional to the number of patches, indicating the importance of moderate content continuity during training.



(a) Deformable offsets **without** RERP. From top to bottom are deformable offsets of DCN, DCNv2, DCNv3, DCNv4, and proposed DAO. (b) Deformable offsets **with** RERP. From top to bottom are deformable offsets of DCN, DCNv2, DCNv3, DCNv4, and proposed DAO.

Fig. 6: Visualization of deformable offsets. The green points \bullet are sampling locations. The red points \bullet are deformable offsets in 2 levels, indicating a deformable receptive field (*e.g.*, each level has a 3×3 kernel size, resulting in $(3 \times 3)^2 = 81$ red points).

To further illustrate the deformable ability of DAO and the superiority of RERP when tackling the distortion of panoramas, we visualize the deformable offset in Fig. 6. Fig. 6a presents the deformable offsets without RERP. It can be observed that DAO is better aware of the panoramic distortion compared with the rest in the DCN series. Since the strongest distortion occurs on the edge of objects, DAO focuses more on edge pixels for better distortion-aware feature extraction. The standard deviation of patch similarity in DAO implies the salience of the image patch center. By multiplicatively combining salient scalars and features, significant pixels, such as those representing object edges, are enhanced in their representation. In Fig. 6b, it is evident that the inclusion of RERP improves the distortion information capturing ability of all models in the DCN series, as opposed to their counterparts without RERP. This observation is consistent with the results in Table 5a. The incorporation of both DAO and RERP in OOOPS showcases a remarkable distortion-aware modeling capability.

5 Conclusion

In this work, we formulate the novel task of Open Panoramic Segmentation (OPS) for unfolding holistic scene understanding, which reaches Open FoV, Open Vocabulary, and Open Domain in surrounding scene segmentation. We have designed the Deformable Adapter Network (DAN) and Random Equirectangular Projection (RERP) for enhancing distortion-aware segmentation capacities of the novel proposed OOOPS model. Extensive experiments on WildPASS, Stanford2D3D, and Matterport3D benchmarks show that the proposed method significantly elevates the performance of state-of-the-art panoramic segmenters.

Acknowledgements

This work was supported in part by the Ministry of Science, Research and the Arts of Baden-Württemberg (MWK) through the Cooperative Graduate School Accessibility through AI-based Assistive Technology (KATE) under Grant BW6-03, in part by BMBF through a fellowship within the IFI programme of DAAD, in part by the Helmholtz Association Initiative and Networking Fund on the HAICORE@KIT and HOREKA@KIT partition, in part by the National Key RD Program under Grant 2022YFB4701400, and in part by Hangzhou SurImage Technology Company Ltd.

References

1. Ai, H., Cao, Z., Zhu, J., Bai, H., Chen, Y., Wang, L.: Deep learning for omnidirectional vision: A survey and new perspectives. arXiv preprint arXiv:2205.10468 (2022)
2. Armeni, I., Sax, S., Zamir, A.R., Savarese, S.: Joint 2D-3D-semantic data for indoor scene understanding. arXiv preprint arXiv:1702.01105 (2017)
3. Athwale, A., Afrasiyabi, A., Lagüe, J., Shili, I., Ahmad, O., Lalonde, J.F.: DarSwin: Distortion aware radial swin transformer. In: ICCV (2023)
4. Berenguel-Baeta, B., Bermudez-Cameo, J., Guerrero, J.J.: FreDSNet: Joint monocular depth and semantic segmentation with fast fourier convolutions from single panoramas. In: ICRA (2023)
5. Caesar, H., Uijlings, J.R.R., Ferrari, V.: COCO-Stuff: Thing and stuff classes in context. CVPR (2018)
6. Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., Zhang, Y.: Matterport3D: Learning from RGB-D data in indoor environments. 3DV (2017)
7. Chen, J., Zhu, D., Qian, G., Ghanem, B., Yan, Z., Zhu, C., Xiao, F., Elhoseiny, M., Culatana, S.C.: Exploring open-vocabulary semantic segmentation without human labels. arXiv preprint arXiv:2306.00450 (2023)
8. Chen, X., Li, S., Lim, S.N., Torralba, A., Zhao, H.: Open-vocabulary panoptic segmentation with embedding modulation. arXiv preprint arXiv:2303.11324 (2023)
9. Cho, S., Shin, H., Hong, S., An, S., Lee, S., Arnab, A., Seo, P.H., Kim, S.: CAT-Seg: Cost aggregation for open-vocabulary semantic segmentation. In: CVPR (2023)
10. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: ICCV (2017)
11. Dao, S.D., Shi, H., Phung, D., Cai, J.: Class enhancement losses with pseudo labels for open-vocabulary semantic segmentation. TMM (2023)
12. Ding, Z., Wang, J., Tu, Z.: Open-vocabulary universal image segmentation with MaskCLIP. In: ICML (2023)
13. Dong, B., Gu, J., Han, J., Xu, H., Zuo, W.: Towards universal vision-language omni-supervised segmentation. arXiv preprint arXiv:2303.06547 (2023)
14. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houshy, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021)
15. Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., Lu, H.: Dual attention network for scene segmentation. In: CVPR (2019)

16. Fu, X., Zhang, S., Chen, T., Lu, Y., Zhou, X., Geiger, A., Liao, Y.: PanopticNeRF-360: Panoramic 3D-to-2D label transfer in urban scenes. arXiv preprint arXiv:2309.10815 (2023)
17. Gao, S., Yang, K., Shi, H., Wang, K., Bai, J.: Review on panoramic imaging and its applications in scene understanding. TIM (2022)
18. Ghiasi, G., Gu, X., Cui, Y., Lin, T.Y.: Scaling open-vocabulary image segmentation with image-level labels. In: ECCV (2022)
19. Guo, J., Wang, Q., Gao, Y., Jiang, X., Tang, X., Hu, Y., Zhang, B.: MVP-SEG: Multi-view prompt learning for open-vocabulary semantic segmentation. arXiv preprint arXiv:2304.06957 (2023)
20. Guttikonda, S., Rambach, J.: Single frame semantic segmentation using multi-modal spherical images. In: WACV (2024)
21. Han, K., Liu, Y., Liew, J.H., Ding, H., Wei, Y., Liu, J., Wang, Y., Tang, Y., Yang, Y., Feng, J., Zhao, Y.: Global knowledge calibration for fast open-vocabulary segmentation. arXiv preprint arXiv:2303.09181 (2023)
22. Hu, X., An, Y., Shao, C., Hu, H.: Distortion convolution module for semantic segmentation of panoramic images based on the image-forming principle. TIM (2022)
23. Jang, S., Na, J., Oh, D.: DaDA: Distortion-aware domain adaptation for unsupervised semantic segmentation. In: NeurIPS (2022)
24. Jaus, A., Yang, K., Stiefelhagen, R.: Panoramic panoptic segmentation: Towards complete surrounding understanding via unsupervised contrastive learning. In: IV (2021)
25. Jaus, A., Yang, K., Stiefelhagen, R.: Panoramic panoptic segmentation: Insights into surrounding parsing for mobile agents via unsupervised contrastive learning. T-ITS (2023)
26. Jiang, Q., Gao, S., Gao, Y., Yang, K., Yi, Z., Shi, H., Sun, L., Wang, K.: Minimalist and high-quality panoramic imaging with PSF-aware transformers. TIP (2024)
27. Jiang, Q., Shi, H., Sun, L., Gao, S., Yang, K., Wang, K.: Annular computational imaging: Capture clear panoramic images through simple lens. TCI (2022)
28. Jiang, W., Wu, Y.: DFNet: Semantic segmentation on panoramic images with dynamic loss weights and residual fusion block. In: ICRA (2018)
29. Jiayun, L., Khandelwal, S., Sigal, L., Li, B.: Plug-and-play, dense-label-free extraction of open-vocabulary semantic segmentation from vision-language models. arXiv preprint arXiv:2311.17095 (2023)
30. Kim, J., Jeong, S., Sohn, K.: PASTS: Toward effective distilling transformer for panoramic semantic segmentation. In: ICIP (2022)
31. Li, B., Weinberger, K.Q., Belongie, S.J., Koltun, V., Ranftl, R.: Language-driven semantic segmentation. In: ICLR (2022)
32. Li, J., Chen, P., Qian, S., Jia, J.: TagCLIP: Improving discrimination ability of open-vocabulary semantic segmentation. arXiv preprint arXiv:2304.07547 (2023)
33. Li, X., Wu, T., Qi, Z., Wang, G., Shan, Y., Li, X.: SGAT4PASS: Spherical geometry-aware transformer for panoramic semantic segmentation. In: IJCAI (2023)
34. Li, Z., Zhou, Q., Zhang, X., Zhang, Y., Wang, Y., Xie, W.: Open-vocabulary object segmentation with diffusion models. In: ICCV (2023)
35. Liang, F., Wu, B., Dai, X., Li, K., Zhao, Y., Zhang, H., Zhang, P., Vajda, P., Marculescu, D.: Open-vocabulary semantic segmentation with mask-adapted CLIP. In: CVPR (2023)
36. Ling, Z., Xing, Z., Zhou, X., Cao, M., Zhou, G.: PanoSwin: A pano-style swin transformer for panorama understanding. In: CVPR (2023)

37. Liu, Y., Ge, P., Liu, Q., Huang, D.: Multi-grained cross-modal alignment for learning open-vocabulary semantic segmentation from text supervision. arXiv preprint arXiv:2403.03707 (2024)
38. Luo, H., Bao, J., Wu, Y., He, X., Li, T.: SegCLIP: Patch aggregation with learnable centers for open-vocabulary semantic segmentation. In: ICML (2023)
39. Ma, C., Yang, Y.H., Wang, Y., Zhang, Y., Xie, W.: Open-vocabulary semantic segmentation with frozen vision-language models. In: BMVC (2022)
40. Ma, C., Yang, Y., Ju, C., Zhang, F., Zhang, Y., Wang, Y.: Open-vocabulary semantic segmentation via attribute decomposition-aggregation. arXiv preprint arXiv:2309.00096 (2023)
41. Ma, C., Zhang, J., Yang, K., Roitberg, A., Stiefelwagen, R.: DensePASS: Dense panoramic semantic segmentation via unsupervised domain adaptation with attention-augmented context exchange. In: ITSC (2021)
42. Mei, J., Zhu, A.Z., Yan, X., Yan, H., Qiao, S., Chen, L., Kretzschmar, H.: Waymo open dataset: Panoramic video panoptic segmentation. In: ECCV (2022)
43. Miller, G.A.: Wordnet: a lexical database for english. *Communications of the ACM* **38**(11), 39–41 (1995)
44. Mukhoti, J., Lin, T.Y., Poursaeed, O., Wang, R., Shah, A., Torr, P.H.S., Lim, S.N.: Open vocabulary semantic segmentation with patch aligned contrastive learning. In: CVPR (2023)
45. Oin, J., Wu, J., Yan, P., Li, M., Ren, Y., Xiao, X., Wang, Y., Wang, R., Wen, S., Pan, X., Wang, X.: FreeSeg: Unified, universal and open-vocabulary image segmentation. In: CVPR (2023)
46. Orhan, S., Bastanlar, Y.: Semantic segmentation of outdoor panoramic images. SIVP (2021)
47. Orsic, M., Kreso, I., Bevandic, P., Segvic, S.: In defense of pre-trained ImageNet architectures for real-time semantic segmentation of road-driving images. In: CVPR (2019)
48. Poudel, R.P.K., Liwicki, S., Cipolla, R.: Fast-SCNN: Fast semantic segmentation network. In: BMVC (2019)
49. Radford, A., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021)
50. Ray, B., Jung, J., Larabi, M.: A low-complexity video encoder for equirectangular projected 360 video content. In: ICASSP (2018)
51. Shen, Z., Lin, C., Liao, K., Nie, L., Zheng, Z., Zhao, Y.: PanoFormer: Panorama transformer for indoor 360° depth estimation. In: ECCV (2022)
52. Shi, H., Zhou, Y., Yang, K., Yin, X., Wang, Z., Ye, Y., Yin, Z., Meng, S., Li, P., Wang, K.: PanoFlow: learning 360° optical flow for surrounding temporal understanding. T-ITS (2023)
53. Takmaz, A., Fedele, E., Sumner, R.W., Pollefeys, M., Tombari, F., Engelmann, F.: OpenMask3D: Open-vocabulary 3D instance segmentation. arXiv preprint arXiv:2306.13631 (2023)
54. Teng, Z., Zhang, J., Yang, K., Peng, K., Shi, H., Reiß, S., Cao, K., Stiefelwagen, R.: 360BEV: Panoramic semantic mapping for indoor bird’s-eye view. In: WACV (2024)
55. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017)
56. Wang, J., Li, X., Zhang, J., Xu, Q., Zhou, Q., Yu, Q., Sheng, L., Xu, D.: Diffusion model is secretly a training-free open vocabulary semantic segmenter. arXiv preprint arXiv:2309.02773 (2023)

57. Wang, W., Dai, J., Chen, Z., Huang, Z., Li, Z., Zhu, X., Hu, X., Lu, T., Lu, L., Li, H., Wang, X., Qiao, Y.: InternImage: Exploring large-scale vision foundation models with deformable convolutions. In: CVPR (2023)
58. Wang, X., Li, S., Kallidromitis, K., Kato, Y., Kozuka, K., Darrell, T.: Hierarchical open-vocabulary universal image segmentation. arXiv preprint arXiv:2307.00764 (2023)
59. Wei, M., Yue, X., Zhang, W., Kong, S., Liu, X., Pang, J.: OV-PARTS: Towards open-vocabulary part segmentation. arXiv preprint arXiv:2310.05107 (2023)
60. Wu, C., Zheng, J., Pfommer, J., Beyerer, J.: Attention-based point cloud edge sampling. In: CVPR (2023)
61. Wysoczańska, M., Ramamonjisoa, M., Trzeciński, T., Siméoni, O.: CLIP-DIY: CLIP dense inference yields open-vocabulary semantic segmentation for-free. arXiv preprint arXiv:2309.14289 (2023)
62. Xia, Z., Pan, X., Song, S., Li, L.E., Huang, G.: Vision transformer with deformable attention. In: CVPR (2022)
63. Xia, Z., Pan, X., Song, S., Li, L.E., Huang, G.: DAT++: Spatially dynamic vision transformer with deformable attention. arXiv preprint arXiv:2309.01430 (2023)
64. Xie, B., Cao, J., Xie, J., Khan, F.S., Pang, Y.: SED: A simple encoder-decoder for open-vocabulary semantic segmentation. arXiv preprint arXiv:2311.15537 (2023)
65. Xiong, Y., Li, Z., Chen, Y., Wang, F., Zhu, X., Luo, J., Wang, W., Lu, T., Li, H., Qiao, Y., Lu, L., Zhou, J., Dai, J.: Efficient deformable convnets: Rethinking dynamic and sparse operator for vision applications. arXiv preprint arXiv:2401.06197 (2024)
66. Xu, J., Liu, S., Vahdat, A., Byeon, W., Wang, X., De Mello, S.: Open-vocabulary panoptic segmentation with text-to-image diffusion models. In: CVPR (2023)
67. Xu, J., Hou, J., Zhang, Y., Feng, R., Wang, Y., Qiao, Y., Xie, W.: Learning open-vocabulary semantic segmentation models from natural language supervision. In: CVPR (2023)
68. Xu, M., Zhang, Z., Wei, F., Hu, H., Bai, X.: SAN: Side adapter network for open-vocabulary semantic segmentation. TPAMI (2023)
69. Xu, M., Zhang, Z., Wei, F., Hu, H., Bai, X.: Side adapter network for open-vocabulary semantic segmentation. In: CVPR (2023)
70. Xu, M., Zhang, Z., Wei, F., Lin, Y., Cao, Y., Hu, H., Bai, X.: A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In: ECCV (2022)
71. Xu, X., Xiong, T., Ding, Z., Tu, Z.: MasQCLIP for open-vocabulary universal image segmentation. In: ICCV (2023)
72. Yang, K., Hu, X., Bergasa, L.M., Romera, E., Huang, X., Sun, D., Wang, K.: Can we PASS beyond the field of view? Panoramic annular semantic segmentation for real-world surrounding perception. In: IV (2019)
73. Yang, K., Hu, X., Bergasa, L.M., Romera, E., Wang, K.: PASS: Panoramic annular semantic segmentation. T-ITS (2020)
74. Yang, K., Hu, X., Fang, Y., Wang, K., Stiefelhagen, R.: Omnisupervised omnidirectional semantic segmentation. T-ITS (2022)
75. Yang, K., Hu, X., Stiefelhagen, R.: Is context-aware CNN ready for the surroundings? Panoramic semantic segmentation in the wild. TIP (2021)
76. Yang, K., Zhang, J., Reiß, S., Hu, X., Stiefelhagen, R.: Capturing omni-range context for omnidirectional segmentation. In: CVPR (2021)
77. Yang, M., Yu, K., Zhang, C., Li, Z., Yang, K.: DenseASPP for semantic segmentation in street scenes. In: CVPR (2018)

78. Yin, W., Liu, Y., Shen, C., Hengel, A.v.d., Sun, B.: The devil is in the labels: Semantic segmentation from sentences. arXiv preprint arXiv:2202.02002 (2022)
79. Yu, H., He, L., Jian, B., Feng, W., Liu, S.: PanelNet: Understanding 360 indoor environment via panel representation. In: CVPR (2023)
80. Zhang, H., Li, F., Zou, X., Liu, S., Li, C., Gao, J., Yang, J., Zhang, L.: A simple framework for open-vocabulary segmentation and detection. In: ICCV (2023)
81. Zhang, J., Ma, C., Yang, K., Roitberg, A., Peng, K., Stiefelwagen, R.: Transfer beyond the field of view: Dense panoramic semantic segmentation via unsupervised domain adaptation. T-ITS (2022)
82. Zhang, J., Yang, K., Ma, C., Reiß, S., Peng, K., Stiefelwagen, R.: Bending reality: Distortion-aware transformers for adapting to panoramic semantic segmentation. In: CVPR (2022)
83. Zhang, J., Yang, K., Shi, H., Reiß, S., Peng, K., Ma, C., Fu, H., Torr, P.H.S., Wang, K., Stiefelwagen, R.: Behind every domain there is a shift: Adapting distortion-aware vision transformers for panoramic semantic segmentation. TPAMI (2024)
84. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: CVPR (2017)
85. Zheng, J., Zhang, J., Yang, K., Peng, K., Stiefelwagen, R.: Materobot: Material recognition in wearable robotics for people with visual impairments. In: 2024 IEEE International Conference on Robotics and Automation (ICRA) (2024)
86. Zheng, X., Pan, T., Luo, Y., Wang, L.: Look at the neighbor: Distortion-aware unsupervised domain adaptation for panoramic semantic segmentation. In: ICCV (2023)
87. Zheng, X., Zhu, J., Liu, Y., Cao, Z., Fu, C., Wang, L.: Both style and distortion matter: Dual-path unsupervised domain adaptation for panoramic semantic segmentation. In: CVPR (2023)
88. Zhou, H., Shen, T., Yang, X., Huang, H., Li, X., Qi, L., Yang, M.H.: Rethinking evaluation metrics of open-vocabulary segmentation. arXiv preprint arXiv:2311.03352 (2023)
89. Zhou, Q., Liu, Y., Yu, C., Li, J., Wang, Z., Wang, F.: LMSeg: Language-guided multi-dataset segmentation. In: ICLR (2023)
90. Zhu, X., Hu, H., Lin, S., Dai, J.: Deformable ConvNets V2: More deformable, better results. In: CVPR (2019)