

360 Layout Estimation via Orthogonal Planes Disentanglement and Multi-view Geometric Consistency Perception

Zhijie Shen, Chunyu Lin, Junsong Zhang, Lang Nie, Kang Liao, Yao Zhao, *Fellow, IEEE*

Abstract—Existing panoramic layout estimation solutions tend to recover room boundaries from a vertically compressed sequence, yielding imprecise results as the compression process often muddles the semantics between various planes. Besides, these data-driven approaches impose an urgent demand for massive data annotations, which are laborious and time-consuming. For the first problem, we propose an orthogonal plane disentanglement network (termed DOPNet) to distinguish ambiguous semantics. DOPNet consists of three modules that are integrated to deliver distortion-free, semantics-clean, and detail-sharp disentangled representations, which benefit the subsequent layout recovery. For the second problem, we present an unsupervised adaptation technique tailored for horizon-depth and ratio representations. Concretely, we introduce an optimization strategy for decision-level layout analysis and a 1D cost volume construction method for feature-level multi-view aggregation, both of which are designed to fully exploit the geometric consistency across multiple perspectives. The optimizer provides a reliable set of pseudo-labels for network training, while the 1D cost volume enriches each view with comprehensive scene information derived from other perspectives. Extensive experiments demonstrate that our solution outperforms other SoTA models on both monocular layout estimation and multi-view layout estimation tasks. Code can be available at <https://github.com/zhijieshen-bjtu/MV-DOPNet>.

Index Terms—360 vision, Manhattan world assumption, horizontal depth, layout estimation.

1 INTRODUCTION

INDOOR panoramic layout estimation refers to reconstructing 3D room layouts from omnidirectional images. Since the panoramic vision system captures the whole-room contextual information, we can estimate the complete room layout with a single panoramic image. However, the 360 field-of-view (FoV) inherent to panoramas introduces significant distortions that intensify with latitude, producing adverse effects on the perception of layout structures. Furthermore, current room-layout models are predicated on the Manhattan world assumption and trained in a supervised manner. For complex situations like non-cuboid rooms, these models might encounter undesired challenges, especially in the absence of annotations. In view of this, we revisit this field from two perspectives: developing an explainable fundamental model and constructing an unsupervised adaption strategy with multi-view geometric consistency.

Previous approaches [1], [2], [3], [4], [5] prefer to estimate the layouts from a 1D sequence. They advocate compressing the extracted 2D feature maps in the height dimension to obtain a 1D sequence, of which every element shares the same distortion magnitude (Fig. 1a). However, such a compressed

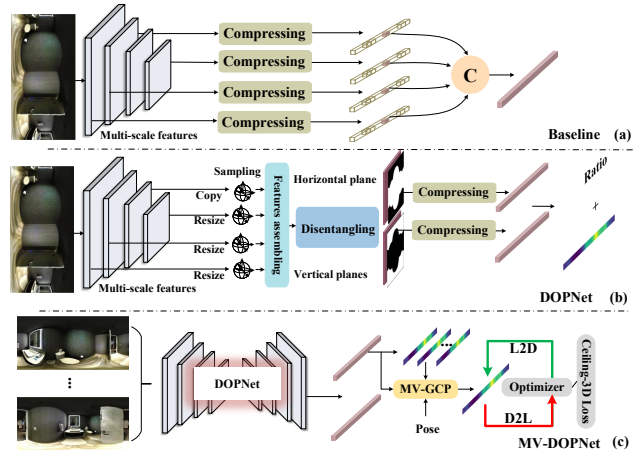


Fig. 1. (a) The commonly used architecture. (b) DOPNet. (c) MV-DOPNet. Compared with the traditional pipeline, ours has the following advantages: (1) Disentangling the orthogonal planes to get distortion-free, and semantic-clean, detail-sharp representations. (2) Leveraging multi-view geometric consistency at the feature level (MV-GCP module). (3) Optimization strategy suitable for horizon-depth/ratio layout representations.

representation does not eliminate panoramic distortions; instead, it merely distributes the distortion uniformly. Moreover, it roughly mixes the vertical and horizontal planes together, confusing the semantics of different planes that are crucial for layout estimation.

To address the above problems, we propose a novel architecture (Fig.1b) that disentangles the orthogonal planes to obtain distortion-free, semantics-clean, and detail-sharp horizon-depth/ratio representations. We termed the pro-

- Corresponding author: Chunyu Lin
- Zhijie Shen, Chunyu Lin, Junsong Zhang, Lang Nie, and Yao Zhao are with the Institute of Information Science, Beijing Jiaotong University, Beijing 100044, China, and also with the Beijing Key Laboratory of Advanced Information Science and Network Technology, Beijing 100044, China (e-mail: zhjshen@bjtu.edu.cn, cylin@bjtu.edu.cn, jszhang@bjtu.edu.cn, nielang@bjtu.edu.cn, yzhao@bjtu.edu.cn.)
- Kang Liao is with the School of Computer Science and Engineering, Nanyang Technological University, Singapore (e-mail: kang.liao@ntu.edu.sg).

posed architecture DOPNet, which comprises three modules to disentangle the inherent 1D representation step-by-step. Firstly, we propose a distortion-aware feature aggregation module to assemble cross-scale features without distortion. Secondly, we propose to pre-segment the vertical plane (*i.e.*, walls) and horizontal planes (*i.e.*, floors and ceilings) from the whole-room contextual information to incorporate 3D geometric cue to the layout model. Given the weak symmetry properties in general layout rooms, a soft-flipping fusion strategy with learnable offsets is designed to facilitate this process. The segmented orthogonal planes are then compressed into two 1D sequences. Considering the potential inaccuracy of pre-segmentation, in the third step, we present a sequence reconstruction module to carry out intra- and inter-sequence interaction. The horizon-depth and ratio are then separately regressed from the vertical and horizontal sequences to avoid mutual interference.

On the other hand, we devise an unsupervised adaptation strategy for multi-view room layout without annotations, introducing both a horizon-depth optimizer and a multi-view perception module (Fig.1c, termed MV-DOPNet). The proposed optimizer, to our knowledge, stands as a pioneering step for horizon-depth representation in the field. It addresses the direct multi-view horizon-depth alignment challenge by shifting the optimization objective from depth to 2D layout boundaries. However, optimization horizon-depth alone brings a mismatch with the unoptimized ratio. To this end, we introduce a ceiling-3D loss function, which implicitly constrains the ratio via ceiling-depth (converted from horizon-depth), ensuring that the optimization process also benefits the ratio representation. Furthermore, we develop a multi-view geometric consistency perceptron module. It employs a coarse horizon-depth sequence combined with differentiable polar coordinate warping to map different features from other views into a unified reference view to construct a 1D cost volume. It explicitly integrates multi-view layout geometry into the model, as well as ensures a more comprehensive and accurate representation.

In this enhanced version of our previous work [6], we explore unsupervised multi-view layout estimation. We introduce a novel optimization approach and a 1D cost volume construction method for horizon-depth/ratio layout representation. Our focus is on leveraging multi-view geometric consistency at both decision and feature levels.

Finally, we conduct extensive experiments on four popular datasets — MatterportLayout [7], Zind [8], Stanford 2D-3D [9], and PanoContext [10] to demonstrate the effectiveness of DOPNet. Both qualitative and quantitative results show that the proposed solution outperforms other SoTA methods. Furthermore, experimental results on the MVL challenge dataset¹ demonstrate our solution significantly outperforms the SoTA methods. Our contributions can be summarized as follows:

- We propose to disentangle ambiguous semantics to obtain spotless (distortion-free, semantics-clean, and detail-sharp) horizon-depth/ratio representations, which benefit the subsequent layout recovery.

- We incorporate two mutual transformative functions to a horizon-depth optimization strategy, complemented by a ceiling-3D loss to share the optimization gains and enhance network fine-tuning.
- We develop a multi-view perception to enhance latent features of the reference view with feature-level geometric consistency, digging into multi-view layout geometry represented by a 1D cost volume.
- On popular benchmarks, our solution outperforms other SoTA schemes on both monocular layout estimation and multi-view layout estimation tasks.

2 RELATED WORK

2.1 Supervised 360 Room Layout Estimation

Based on Manhattan World assumption [11] that all walls and floors are aligned with global coordinate system axes and are perpendicular to each other, many approaches utilize convolutional neural networks (CNNs) to extract useful features to improve layout estimation accuracy. Specifically, Zou *et al.* [12] propose LayoutNet to predict the corner/boundary probability maps directly from the panoramas and then optimize the layout parameter to generate the final predicted results. Furthermore, they annotate the Stanford 2D-3D dataset [9] with awesome layouts for training and evaluation. Yang *et al.* [13] propose DulaNet to predict a 2D floor plane semantic mask from both the equirectangular view and the perspective view of the ceilings. The modified version, LayoutNet v2, and DulaNet v2, which have better performance than the original version, are presented by Zou *et al.* [7]. Fernandez *et al.* [14] present to utilize equirectangular convolutions (EquiConvs) to generate corner/edge probability maps. Sun *et al.* propose HorizonNet [1] and HoHoNet [2] to simplify the layout estimation processes by representing the room layout with a 1D representation. Besides, they use Bi-LSTM and multi-head self-attention to build long-range dependencies and refine the 1D sequences. Recently, several approaches have adopted this 1D representation and achieved impressive performance. For example, Rao *et al.* [15] establish their network based on HorizonNet [1]. They replace standard convolution operation with spherical convolution to reduce distortions and take Bi-GRU to reduce computational complexity. Wang *et al.* [5] combine the geometric cues across the layout and propose LED2-Net that reformulates the room layout estimation as predicting the depth of walls in the horizontal direction. Not constrained to Manhattan scenes, Pintore *et al.* [4] introduce AtlantaNet to predict the room layout by combining two projections of the floor and ceiling planes. Driven by the self-attention mechanism, many transformer-based methods [16], [17], [18] proposed to build long-range dependencies. For example, Jiang *et al.* [3] represent the room layout by horizon depth and room height and introduce a Transformer to enforce the network to learn the geometric relationships. However, the popular 1D representations could confuse the semantics of different planes, thus causing the layout estimation challenging. We disentangle this 1D representation by pre-segmenting orthogonal planes.

1. https://github.com/mvlchallenge/mvl_toolkit

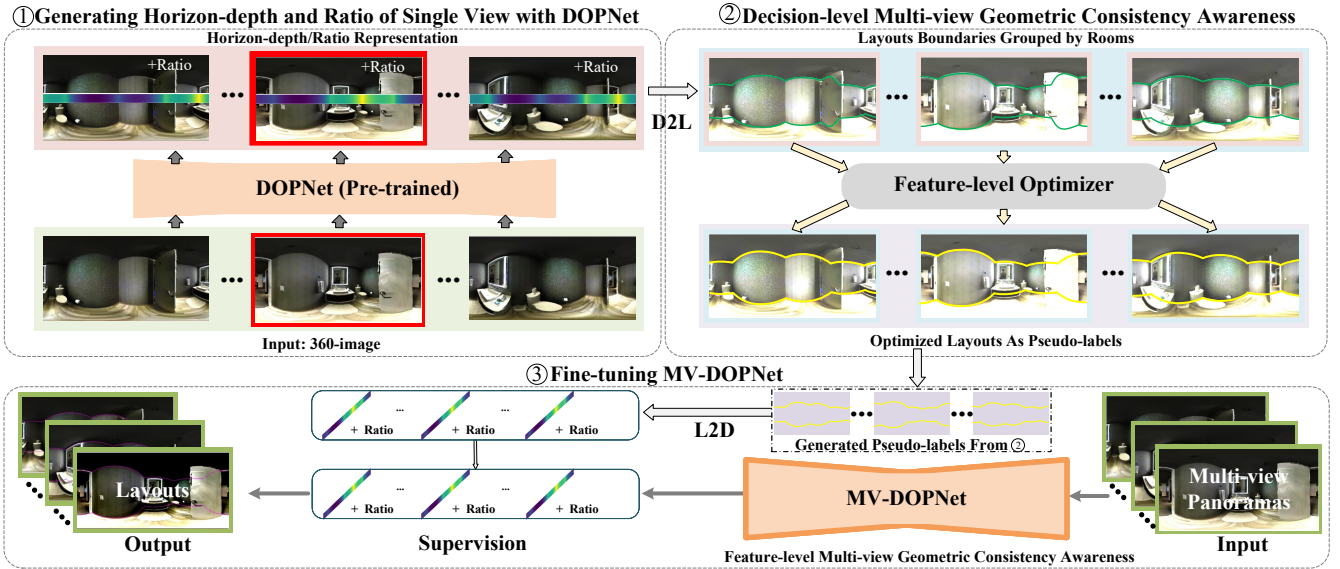


Fig. 2. Workflow of the proposed unsupervised multi-view layout framework. Our approach adopts panoramic images as input and predicts the corresponding horizon-depth map and the height (ratio). Then they are employed to recover the layout boundaries via D2L transformation. We leverage the multi-view layout consistency to optimize the boundaries and convert them back to the original format via L2D transformation. Finally, the optimized pseudo-labels are employed to fine-tune the MV-DOPNet (extending DOPNet by designing a multi-view geometric consistency perceptron module).

2.2 Unsupervised 360 Room Layout Estimation

In new target domains, such as more general rooms, existing models' performance declines due to different viewpoints, lighting conditions, or occlusions. Addressing these issues typically requires labeled data for fine-tuning. Additionally, the diversity of indoor rooms needs even more labeled data for room layout models to adapt to new domains. Recent studies devote themselves to developing unsupervised adaptation approaches for the new target domain. However, room layout representation is inherently sparse and topological, making it challenging to apply unsupervised constraints used in dense prediction tasks directly. This sparse representation can easily lead to local optima when applying image reprojection losses. To address these challenges, existing methods rely on pre-training techniques to obtain initial layout estimates for the subsequent self-training operation. The SSLayout360 [19] innovatively applies the traditional self-training technique, typically utilized for classification or object detection, to geometric prediction tasks like layout estimation. Specifically, this approach employs a pre-trained Mean Teacher [20] model to generate pseudo-labels for unlabeled data. These pseudo-labels are subsequently incorporated into the training of existing layout estimation models. Since this method emphasizes monocular unsupervised layout estimation, it neglects the geometric constraints that exist between different views within a single room. In contrast, 360-MLC [21] advocates for leveraging geometric consistency among multi-views within the same room to generate more accurate pseudo-labels. Recent studies [5] have highlighted the potential advantages of employing differentiable depth rendering techniques in layout estimation. Through the integration of depth prediction loss, it becomes feasible to incorporate geometrical insights into the model. However, the 360-MLC method designed for layout boundaries is unsuitable for these advanced approaches. Besides, the 360-MLC method still depends on a monocular

layout model without any multi-view perception designs during the fine-tuning process. To address the challenges, we investigate a horizon-depth customized optimization methodology for these advanced models and explore an effective strategy for fine-tuning these advanced schemes on unlabeled data.

2.3 360 Image Distortion Mitigation

When converting spherical data to the equirectangular projection format, severe distortions are introduced. Some recent studies [22], [23], [24] focus on designing spherical customized convolutions to make the network adapt to panoramic distortions and offer superior results compared to standard convolutional networks. Su *et al.* [25] introduce SphConv to remove distortions by changing the regular planar convolutions kernel size towards each row of the 360 images. Cohen *et al.* [26] propose a spherical CNN that transforms the domain space from Euclidean S^2 space to a $SO(3)$ group to reduce the distortion and encodes full rotation invariance in their network. Coors *et al.* [27] addresses it by defining the convolution kernels on the tangent plane and keeping the convolution domain consistent in each kernel. Therefore, the non-distorted features can be extracted directly on 360 images. Rao *et al.* [15] first applied spherical convolution operation to room layout task and showed significant improvements. Different from the above approaches, Shen *et al.* [28] propose the first panorama Transformer (PanoFormer) and divide the patches on the tangent plane to remove the negative effect of distortions. However, many of the latest studies in layout estimation tend to overlook distortions, operating under the assumption that 1D representations effectively address these issues. On the contrary, we propose a feature assembling mechanism with cross-scale distortion awareness to deal with distortions. Extensive experimental results demonstrate the

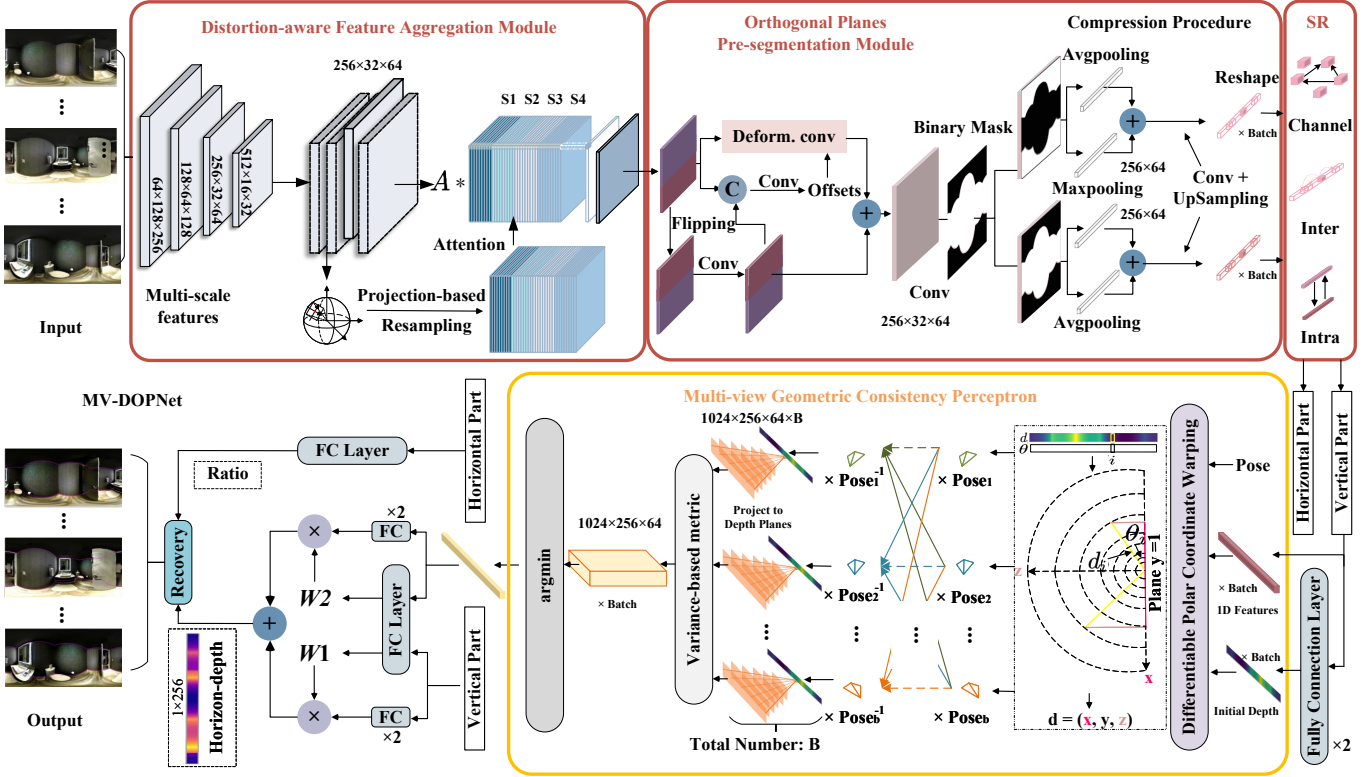


Fig. 3. MV-DOPNet. The feature assembling mechanism is designed to deal with distortions as well as integrate shallow and deep features. Then we pre-segment orthogonal planes to produce two 1D representations with distinguished plane semantics. A sequence reconstruction module is deployed to reconstruct the 1D representations. We design multi-view geometric consistency perceptron to incorporate multi-view insights. Our scheme leverages a sequence of panoramas with a resolution of 512×1024 as input and predicts the horizon-depth and the associated ratio value.

necessity of handling distortions before generating 1D representations.

3 APPROACH

3.1 Overview

As illustrated in Fig. 2, our proposed framework is tailored for unlabeled data from general room environments. This framework includes pseudo-label generation, multi-view perception, and model fine-tuning. We outline the sequential steps to the framework’s design as follows.

Generating pseudo-labels. For the unlabeled MVL data, we first follow the previous method [19] [21] to generate pseudo-labels. Specifically, we employ our DOPNet to generate initial horizon-depth maps and ratios for the multiple views.

Multi-view Geometric Consistency Awareness. These horizon-depth sequences are grouped by rooms and undergo refinement using Algorithm 1 to create pseudo-labels for fine-tuning. Additionally, during the fine-tuning period, we present MV-DOPNet, enabling feature-level multi-view insights of layout models.

Fine-tuning with Pseudo-labels. The optimized pseudo-labels are employed to fine-tune the proposed multi-view layout model (MV-DOPNet). Our model, upon receiving multi-view panoramas as input, generates the horizon-depth and ratio for each view, which are then utilized to recover the corresponding room layouts. With our customized constraints, this fine-tuning process taps into a

broader data spectrum, enhancing MV-DOPNet’s inference and generalization capabilities.

3.2 Disentangle Orthogonal Planes

We first briefly review the details of DOPNet [6]. As illustrated in Fig. 3 (circled in red), the proposed framework consists of three stages. Specifically, we adopt a backbone (basically ResNet) to extract a series of features at 4 different scales from a single panorama. Then a distortion-aware feature aggregation module is designed to fuse the multi-scale features and free them from panoramic distortions. In the next stage, we introduce a soft-flipping fusion strategy to explore the weak symmetry properties between the floor boundary and the ceiling boundary. Subsequently, we segment the vertical and horizontal planes from the fused features followed by vertical compression, yielding two 1D sequences. In the last stage, the segmented features are reconstructed to be more discriminative and informative. Finally, the horizon-depth map and ratio value are regressed from the precise 1D representations.

Distortion-aware Feature Aggregation Module. In previous works [1], [2], [3], [5], researchers prefer to follow the architecture of HorizonNet [1] to extract multi-scale feature maps of an equirectangular image. These feature maps are merged together via vertical compression and concatenation [29]. However, the popular architecture just concatenates multi-scale features and neglects the inherent distortion distribution, resulting in an inferior panoramic feature extraction capability. To address the above issues, we propose

a novel distortion-aware feature aggregation module to deal with distortions, as well as ingeniously integrate shallow and deep features.

Motivated by [14], [28], we first polymerize the most relevant features via a virtual tangent plane on the panoramic sphere projection. Based on the projection formula [27], [28], [30], we can get the distortion-free sampling positions in the 2D feature maps. Then we employ the learnable offsets to adjust the sampling locations adaptively and obtain gathered features from the distorted features. Then we select the 3rd scale as the reference scale and resize other feature maps to this scale. Specifically, we utilize the learnable sampling positions to gather the distortion-free feature maps for every feature with resized scales.




Fig. 4. Illustration of the misaligned case when leveraging the symmetry property.

Orthogonal Planes Pre-segmentation Module. To explicitly capture the geometric cues, we propose to segment the vertical and horizontal planes from the whole scenario before compressing the 2D feature maps. Specifically, we pre-segment the horizon/vertical planes in advance to capture an explicit geometric cue in 3D space and encourage the network to learn a binary mask via a simple segmentation head. Following previous works [1], [2], we further vertically compress 2D feature maps to two 1D sequences with different plane semantics.

Besides, we observe that it is challenging to learn accurate layout boundaries in a complex scene due to the disturbance of furniture and illumination, *e.g.*, occlusion. In Dula-Net [13], Yang *et al.* propose a dual-branch network, leveraging a perspective ceiling view branch to assist in accurate layout prediction. However, the ceiling-view solves wall occlusion issues, but it compounds the occlusion effect from the indoor decorations (*e.g.*, furniture and lighting). Besides, Dula-Net aims to predict a plan probability map, while our goal is to estimate a horizon-depth map and an associated ratio value. The view conversion operation can result in inconsistent depth, hindering accurate horizon-depth estimation. To address these challenges, we propose to leverage the symmetry property of indoor rooms (*e.g.*, the floor boundary and ceiling boundary are mostly symmetric in 3D space) to provide complementary information. We introduce a soft-flipping fusion strategy to exploit this symmetry property. We first vertically flip the feature map to get their symmetrical version. In fact, the floor boundary and ceiling boundary are not strictly symmetric in a general room scene (illustrated in Fig. 4). To this end, we adopt a deformable convolution with 3×3 kernel size to adaptively adjust the symmetrical version. Next, we fuse the original feature and its “soft” symmetrical version to provide more informative boundary cues.

Sequences Reconstruction Module. Previous segmentation and compression operations disrupt inherent dependencies in 2D features. To this end, we advocate generating discrimi-

native channels, establishing solid long-range dependencies, and enhancing helpful residuals. Li *et al.* [31] find that there are generally unconfined dependencies among channels, resulting in confused semantic cues. To this end, we employ graph convolution to enforce each channel to concentrate on distinguishing features. Different from normal pixel-wise [32] or object-wise [33] graphs, the introduced channel-wise one tends to guide the node to subtract the information from the neighbor nodes rather than aggregation.

Moreover, we employ the standard self-attention to reconstruct the intra-sequence long-range dependencies. It captures global interactions to adapt to the large FoV of panoramas, contributing to better performance. To compensate for inevitable errors in pre-segmentation, we introduce cross-attention to provide the missing useful residuals. Then we add the learned residuals with the two 1D sequences to implement the reconstruction. After that, two separate fully connected layers are employed to output the predicted horizon depth $d \in \mathbb{R}^{B \times W}$ from the vertical plane sequence and the ratio value from the horizontal plane one, respectively.

3.3 Multi-view Geometric Consistency Awareness

The strategies proposed in this section are specifically tailored for multi-view scenarios. We align the multi-view horizon-depths to fully harness the potential of geometric consistency across different views, thereby enhancing the precision of layout estimation tasks. Specifically, the alignment is performed in two steps: decision-level optimization and feature-level perception. At the decision level, estimated horizon depths are projected onto a 2D canvas for optimization, which avoids the issues of scale and shifting caused by changes in poses, lighting conditions, and occlusions across multiple views. In the feature-level perception, multi-view horizon-depths are initially aligned with the camera pose information. Subsequently, through fine-tuning, the model can leverage the complementary features from other views to deal with the occlusions and produce more accurate horizon-depths for each view.

3.3.1 Decision-level Optimization

Horizon-depth Optimization. Our approach utilizes the advanced horizon-depth and ratio representation. However, under the assumption in [5], horizon-depth sequences generated from the single-view layout model (DOPNet) are not robust due to the varying lighting conditions and occlusions across different views. Directly optimizing depth using camera pose information is challenging. To this end, we propose to project these depth values onto a canvas, thereby shifting the optimization object from 3D space into 2D space. Specifically, we introduce the Depth-to-Layout (D2L) conversion, the reverse of the Layout-to-Depth (L2D) conversion [5], ensuring the method’s compatibility within the advanced layout estimation framework. Following refinement, we employ the L2D transformation to maintain the consistency of the model’s original representation schema. This entire cycle is described in Algorithm 1. Notably, we incorporate the calculation of standard deviations among multi-view horizon-depth maps to formulate the loss function.

Ceiling-3D Constraint for Ratio. The layout model’s output implies a correlation wherein the horizon-depth d is

Algorithm 1 Pseudo-labels Generation

Input: Registered horizon-depth sequences D within a room and the corresponding camera pose P relative to the world coordinates

Output: Optimized pseudo-labels \hat{D} , standard deviation σ

```

1: for  $d$  in  $D$  do
2:   horizon-depth to layout:  $l = \text{D2L}(d)$ 
3:   layout in unified world coordinates:  $l^{wc} = l \times p_d$ 
4: end for
5: All the unified layouts:  $L^{wc}$ 
6: for  $p$  in  $P$  do
7:   for  $l^{wc}$  in  $L^{wc}$  do
8:     layout at the reference view (associated with  $p$ ):
        $l_p = l^{wc} \times (p)^{-1}$ 
9:   end for
10: All the layouts at the reference view:  $L_p$ .
11:  $\hat{l}_p \leftarrow$  select high-confidence layouts point by point
    within  $L_p$  (majority principle)
12:  $\sigma_p \leftarrow$  standard deviation of  $L_p$ 
13: layout to horizon-depth:  $\hat{d}_p = \text{L2D}(\hat{l}_p)$ 
14: end for
15: return  $\hat{D}, \sigma$ 
  
```

related to the ceiling-depth d^c . Upon transforming the depth readings into 3D cartesian coordinates, we ascertain the spatial 3D points designated as $\mathbf{d} = (\mathbf{x}, \mathbf{y}, \mathbf{z})$ and $\mathbf{d}^c = (\mathbf{x}^c, \mathbf{y}^c, \mathbf{z}^c)$. These coordinates sustain a definable relationship, expressed by the following expression:

$$\mathbf{y}^c = -R\mathbf{y}, \quad (1)$$

where $\mathbf{d}, \mathbf{d}^c \in \mathbb{R}^{B \times 3 \times W}$; R indicates the predicted ratio value. Consequently, we formulate constraints on the ceiling depth rather than imposing direct constraints on the ratio. This improvement facilitates a more direct control over the model’s geometric fidelity, thereby refining the overall accuracy of the layout estimation.

3.3.2 Feature-level Perception

We introduce a multi-view geometric consistency percepton, depicted in Fig. 3 (highlighted in Orange), designed to integrate multi-view insights at the feature level by constructing 1D cost volumes.

Differentiable Polar Coordinate Warping. Given the width of the sequence W , we first calculate the longitude values θ :

$$\forall i \in \{0, \dots, W-1\} : \theta = \left(\frac{i+0.5}{W} - 0.5\right) 2\pi. \quad (2)$$

With the coarse horizon-depth d generated, we apply a differentiable polar coordinate warping method to obtain its 3D coordinates $\mathbf{d} = (\mathbf{x}, \mathbf{y}, \mathbf{z})$ in the cartesian coordinate system:

$$\begin{cases} \mathbf{x} &= d \sin \theta \\ \mathbf{y} &= 1 \\ \mathbf{z} &= d \cos \theta \end{cases}, \quad (3)$$

where all the operations are element-wise. Afterward, we align the 3D points \mathbf{d} of other views to the reference view by employing the relative camera poses:

$$\mathbf{d}^{align} = p^{ref} \times \mathbf{d}, \quad (4)$$

where p^{ref} indicates the matrix of the relative camera pose between the current view and the target view.

Depth Plane Sampling. Given 64 depth planes which are defined linearly across the normalized depth range $[0, 1]$, yielding intervals $1/64$, masks \mathbf{m} are created for each depth plane by checking if the z-coordinate of the aligned points \mathbf{d}^{align} lies within the depth plane. For each view, we apply the mask to the expanded features $\mathbf{f} \in \mathbb{R}^{B \times C \times W \times 64}$ to get the volumes:

$$\mathbf{v} = \mathbf{f} \odot \mathbf{m}, \quad (5)$$

where \odot represents the element-wise product. Eventually, for each reference view, we can get all the volumes $\mathbf{V} \in \mathbb{R}^{B \times C \times W \times 64 \times B}$.

1D Cost Volumes. Similar to [34], we employ the variance-based metric to aggregate \mathbf{V} to one 1D cost volume \mathbf{c} . The process can be written as follows:

$$\mathbf{c} = \left(\sum_{i=1}^B \mathbf{V}_i - \bar{\mathbf{V}}_i^2\right) / B, \quad (6)$$

where $\mathbf{c} \in \mathbb{R}^{B \times C \times W \times 64}$, and $\bar{\mathbf{V}}_i$ is the average volume among all features volumes. We finally apply the *argmin* operation among the depth planes to get a precise 1D feature:

$$\mathbf{f}^c = \text{argmin}_{\mathbf{c}} \mathbf{c}, \quad (7)$$

where $\mathbf{f}^c \in \mathbb{R}^{B \times C \times W}$. Then we employ a fully connected layer to get a depth sequence d^{cost} from \mathbf{f}^c . Finally, we fuse the two depths to ensure an optimized depth representation. The final depth can be expressed as:

$$d^{final} = W_1 d^{cost} + W_2 d, \quad (8)$$

where d^{cost} represents the depth obtained from the cost volume; W_1 and W_2 are the weights determined by the fully connected layers.

3.4 Loss Function

3.4.1 Objective Function for Pre-training

Our objective function consists of two parts: one for room layout estimation and the other for plane segmentation. For the first part, we strictly follow [3] and denote it as \mathcal{L}_{lay} . For the segmentation part, we apply binary cross-entropy loss \mathcal{L}_{seg} as follows:

$$\mathcal{L}_{seg} = -\frac{1}{N} \sum_{i=1}^N (d_i \log \hat{d}_i + (1 - d_i) \log (1 - \hat{d}_i)), \quad (9)$$

where d_i is the ground truth of the generated segmentation label, and \hat{d}_i is the predicted value. Ultimately, we formulate the objective function as follows:

$$\mathcal{L} = \mu \mathcal{L}_{seg} + \mathcal{L}_{lay}, \quad (10)$$

where μ is set to 0.75 to balance different constraints.

3.4.2 Objective Function for Fine-tuning

The fine-tuning phase of our model employs a loss function comprised of four distinct elements: normals, gradients of normals, horizon-depth, and ratio. Firstly, we employ σ

which is calculated in Algorithm 1 as weights to design the losses:

$$d = \frac{d}{\sigma^2}, \quad (11)$$

$$\hat{d} = \frac{\hat{d}}{\sigma^2}, \quad (12)$$

where the operations are element-wise. Then we carefully follow the methodology outlined in [3] for the first three components, symbolized by \mathcal{L}_n^w , \mathcal{L}_g^w , and \mathcal{L}_d^w . The subsequent ratio loss (ceiling-3D loss) is as follows,

$$\mathcal{L}_r = \frac{1}{3N} \sum_{i=1}^3 \sum_{j=1}^N \left| \frac{\mathbf{d}_{i,j}^c - \hat{\mathbf{d}}_{i,j}^c}{\sigma_{i,j}^2} \right|, \quad (13)$$

where $\mathbf{d}_{i,j}^c$ indicates the ceiling-depth converted from the pseudo-label, and $\hat{\mathbf{d}}_{i,j}^c$ is the one converted from the predicted horizon-depth. Combining all these components, the cumulative objective function is given by:

$$\mathcal{L} = \lambda_1(\mathcal{L}_n^w + \mathcal{L}_g^w) + \lambda_2\mathcal{L}_d^w + \lambda_3\mathcal{L}_r. \quad (14)$$

In our implementation, we empirically set λ_1 to 0.1, λ_2 to 0.9, and λ_3 to 0.08 to ensure a harmonious balance between various constraints.

4 EXPERIMENTS

In this section, we validate the effectiveness of our solution and compare it with existing SoTA approaches on five popular datasets. Concretely, we conduct experiments on a single GTX 3090 GPU, and the batch size is set to 16 for training. The proposed approach is implemented with PyTorch. We choose Adam [35] as the optimizer and keep the default settings. The initialized learning rate is 1×10^{-4} . As in previous works [1], [3], we adopt standard left-right flipping, panoramic horizontal rotation, luminance change, and pano stretch for data augmentation during training. Furthermore, we compare our framework against the SoTA baseline method, 360-MLC. To ensure an equitable assessment, we retain the same parameters and settings.

4.1 Datasets

Monocular Layout Estimation Datasets. Four datasets are used for our experimental validation: Stanford 2D-3D [9], PanoContext [10], MatterportLayout [7] and ZInd [8].

PanoContext [10] and Stanford 2D-3D [9] are two commonly used datasets for indoor panoramic room layout estimation that contain 514 and 552 cuboid room layouts, respectively. Especially, Stanford 2D-3D [9] is labeled by Zou *et al.* [12] and has a smaller vertical FoV than other datasets. Besides, MatterportLayout [7] is also annotated by Zou *et al.* [7], which contains 2,295 general room layouts. The final ZInd [8] dataset includes cuboid, general Manhattan, non-Manhattan, and non-flat ceilings layouts, which better mimic the real-world data distribution. The split of ZInd [8] consists of 24,882 (for training), 3,080 (for validation), and 3,170 (for testing) panoramas. We use the official splits of the four datasets for the monocular layout estimation task and follow the same training/validation protocols as previous works [3] for a fair comparison.

TABLE 1
Quantitative comparison results with the current SoTA solutions evaluated on Stanford 2D-3D and PanoContext [8] dataset.

Method	3DIoU(%)	CE(%)	PE(%)
Train on PanoContext + Whole Stnfd.2D3D datasets			
LayoutNetv2 [7]	85.02	0.63	1.79
Dula-Netv2 [7]	83.77	0.81	2.43
HorizonNet [1]	82.63	0.74	2.17
LGT-Net [3]	85.16	-	-
LGT-Net [w/ Post-proc] [3]	84.94	0.69	2.07
Ours	85.46	-	-
Ours [w/ Post-proc]	85.00	0.69	2.13
Train on Stnfd.2D3D + Whole PanoContext datasets			
LayoutNetv2 [7]	82.66	0.83	2.59
Dula-Netv2 [7]	86.60	0.67	2.48
HorizonNet [1]	82.72	0.69	2.27
AtlantaNet [4]	83.94	0.71	2.18
LGT-Net [3]	85.76	-	-
LGT-Net [w/ Post-proc] [3]	86.03	0.63	2.11
Ours	85.47	-	-
Ours [w/ Post-proc]	85.58	0.66	2.10

Multi-view Layout Estimation Datasets. To validate the efficiency of our proposed framework, MV-DOPNet, we introduce a dataset specifically tailored for the Multi-View Layout Challenge. This challenge dataset is divided into two subsets: MP3D-FPE and HM3D-MVL. MP3D-FPE originates from the Matterport3D [36] dataset and is amassed using the MINOS [37] simulator. It contains 50 diverse scenes, aggregating 687 rooms in all. Conversely, HM3D-MVL is a recently curated dataset derived from HM3D [38], utilizing the Habitat [39] simulator. Echoing the methodology in MP3D-FPE, for HM3D-MVL, the organizer replicated user scanning patterns to encapsulate more genuine camera movement dynamics. For a fair comparison, we employ the training splits from the challenge phase for model training and leverage the pilot splits to evaluate the performance. Notably, since the test splits of the MVL dataset from the challenge phase are not publicly released, we do not use them.

4.2 Evaluation Metrics

We employ the commonly used standard evaluation metrics for a fair comparison, including corner error (CE), pixel error (PE), intersection over the union of floor shapes (2DIoU), and 3D room layouts (3DIoU). Among them, 3DIoU yields a better reflection of the accuracy of the layout estimation in 3D space. RMSE and $\delta_{1.25}$ indicate the performance of depth estimation, *e.g.* the horizon-depth map.

4.3 Cuboid Room Results

We follow LGT-Net [3] to use the combined dataset scheme mentioned in Zou *et al.* [7] to evaluate our network on cuboid datasets. We denote the combined dataset that contains training splits of PanoContext [8] and whole Stanford 2D-3D datasets as ‘‘Train on PanoContext +

TABLE 2

Quantitative comparison results with the current SoTA solutions evaluated on MatterportLayout [7] dataset and ZInD [8] dataset.

Method	2DIOU(%)	3DIOU(%)	RMSE	δ_1
MatterportLayout datasets				
LayoutNetv2 [7]	78.73	75.82	0.258	0.871
Dula-Netv2 [7]	78.82	75.05	0.291	0.818
HorizonNet [1]	81.71	79.11	0.197	0.929
AtlantaNet [4]	82.09	80.02	-	-
HoHoNet [2]	82.32	79.88	-	-
LED ² -Net [5]	82.61	80.14	0.207	0.947
DMH-Net [40]	81.25	78.97	-	0.925
LGT-Net [3]	83.52	81.11	0.204	0.951
Ours	84.11	81.70	0.197	0.950
ZInD datasets				
HorizonNet [1]	90.44	88.59	0.123	0.957
LED ² -Net [5]	90.36	88.49	0.124	0.955
LGT-Net [3]	91.77	89.95	0.111	0.960
Ours	91.94	90.13	0.109	0.960

Whole Stnfd.2D3D datasets” in Tab. 1. Similarly, “Train on Stnfd.2D3D + Whole PanoContext” in Tab. 1 represents that we train our network on the combined dataset that contains training splits of Stanford 2D-3D and the whole PanoContext dataset. The scheme is commonly used in previous works [3], [12]. We also report the results with a post-processing strategy in DuLa-Net [13] that is denoted as “Ours [w/ Post-proc]”.

Comparison Results. We exhibit the quantitative comparison results on cuboid room layouts in Tab. 1. From the first group in Tab. 1, we can observe that our approach outperforms all the other SoTA schemes with respect to 3DIOU. But from the second group, Dula-Net v2 [7] offers better 3DIOU than ours. Dula-Net v2 [7] employs a perspective view (*i.e.*, cubemap) that is more effective for panoramas with small vertical FoV. However, once applied to more general room layout datasets (MatterportLayout [7] dataset), its performance degrades significantly. On the contrary, the proposed approach shows more general performance.

4.4 General Room Results

MatterportLayout [7] dataset and ZInD [8] dataset provide more general indoor room layouts, which is much more challenging than the cuboid room layout datasets. Tab. 2 exhibits the evaluation on MatterportLayout [7]/ZInD [8] datasets. The results of LED²-Net [5] and HorizonNet [1] are from [3] that we strictly followed. Particularly, Jiang *et al.* utilize their official code to re-train and re-evaluate with the standard evaluation metrics.

Comparison Results. We observe that Dula-Net [13] demonstrates much worse performance on the MatterportLayout [7] dataset, indicating that these perspective view-based methods are hard to adapt to the general indoor scenarios. Moreover, these 1D sequence-based methods are better than those 2D convolution-based schemes. However, they cannot produce accurate 3D layouts from a 1D sequence due to confused plane semantics, even if they introduce more

TABLE 3

Comparison results with the baseline evaluated on MVL challenge dataset.

Method	2DIOU(%)	3DIOU(%)	RMSE	δ_1
Group 1: pre-trained on MatterportLayout				
HorizonNet	78.53	73.93	0.342	0.861
DOPNet	80.71	75.13	0.301	0.839
360-MLC	79.95	76.47	0.350	0.883
360-MLC-Iter	80.87	77.39	0.334	0.889
Ours	83.59	79.42	0.268	0.877
Group 2: pre-trained on ZInD				
HorizonNet	63.57	58.37	0.578	0.701
DOPNet	65.35	57.51	0.574	0.740
360-MLC	80.90	76.05	0.329	0.878
Ours	83.41	78.82	0.284	0.890

powerful relationship builders (*e.g.*, Bi-LSTM, Transformer). Hence, these methods give worse 3DIOU. Compared with them, our approach offers better performance than all other approaches with respect to 3DIOU because ours captures an explicit 3D geometric cue by disentangling orthogonal planes, which benefits recovering layouts from clear semantics (evidenced by Fig. 8).

We show the qualitative comparisons in Fig. 5/6. From the figures, we can observe that our method offers better 3D results (floor plan). In Fig. 7, our results are very similar to the post-processing version. Both indicate our proposal that disentangling orthogonal planes to capture an explicit 3D geometric cue is an effective strategy.

4.5 MVL Results

Diverging from most conventional layout datasets like MatterportLayout [7] and Zillow [8], the dataset employed to evaluate the proposed framework stands out with its expansive rooms, augmented count of rooms within each scene, and a pronounced inclusion of non-Manhattan spaces. These intricacies emphasize a novel spectrum of challenges for room layout estimation, particularly in the absence of labeled data. In this section, we compare the current SoTA methods’ performance in facing these dilemmas. All the models are evaluated on the warm-up phase pilot splits. The fine-tuning process is based on the unlabeled challenge phase training splits. Since there is an overlap between the used MVL dataset and the MatterportLayout dataset, we also conduct an additional experiment by applying the parameters pre-trained on the ZInD dataset to demonstrate the multi-view performance more convincingly.

The comparative results are presented in Table 3. In the table, both “HorizonNet” and “DOPNet” indicate that we just employ the pre-trained parameters without the fine-tuning process. Additionally, the 360-MLC method is enhanced with an iterative optimization strategy (iteratively generating refined pseudo-labels and fine-tuning², denoted as “360-MLC-Iter” in Group 1, Table 3). In the first group, we can observe that our DOPNet performs better than the HorizonNet. This exceptional performance is attributable to

2. Our solution that won the 1st Place in the MVL Challenge at the CVPR 2023 Workshop.



(a) Our qualitative results on Stanford 2D-3D [9] dataset.

(b) Our qualitative comparison on PanoContext [10] dataset.

Fig. 5. Qualitative comparison results evaluated on cuboid layout datasets, Stanford 2D-3D [9] and PanoContext [10]. The boundaries of the room layout on a panorama are shown on the left and the floor plan is on the right. Ground truth is best viewed in blue lines and the prediction in green. The predicted horizon depth, normal, and gradient are visualized below each panorama, and the ground truth is in the first row.

two primary factors: first, the incorporation of an advanced horizon-depth representation in our model. Second, the integration of meticulously designed modules within our network. Besides, we observe that our iterative strategy was found to narrow the performance disparities among base layout models. Particularly, models that initially displayed suboptimal performance witnessed notable enhancement. Yet, after the 2nd iteration, a discernible influx of noise disrupts the estimated boundary, preventing it from forming seamless and complete layouts. However, it is worth noting that even with iterative optimization, the baseline method lagged behind our approach. Our method achieves the top performance, demonstrating the superior effectiveness of the proposed approach.

In Tab.3, the depth metric (*e.g.*, RMSE, and δ_1) are also reported to demonstrate our multi-view performance. To calculate the depth accuracy metrics, we follow the previous works [3] [1] to convert the ground truth into horizon-depth representation by setting the camera height as 1.6 meters. Since the horizon-depth representation of layouts is sparse and topological, it is sensitive to the inferior values. Hence, from Group 1 in Table 3, we can observe that while our approach ranks second on δ_1 , it achieves a lower RMSE, resulting in more accurate layouts.

Occlusion analysis. From Fig. 9, we can observe that due to the occlusion (*e.g.*, the stairs, and the walls), the DOPNet fails to produce the correct layouts in these challenging regions. The current state-of-the-art approach that only optimizes the pseudo-labels is also disturbed by the occlusions. In contrast, our approach can benefit from the pro-

TABLE 4
Ablation study on MatterportLayout [7] dataset and MVL dataset.

Method	2DIoU(%)	3DIoU(%)
Group 1: ablation on MatterportLayout (Monocular)		
w/o Cross-scale interaction	83.24	80.71
w/o Feature assembling	82.36	80.10
w/o Disentangling planes	83.23	80.60
w/o Flipping fusion	83.01	80.35
w/o Discriminative channels	82.74	80.23
w/o Long-range dependencies	83.04	80.97
w/o Residuals	82.89	80.45
Ours [Full]	83.46	81.34
Group 2: ablation on MVL (Multi-view)		
DOPNet[mp3d]	80.71	75.13
DOPNet[fine-tuning]	81.64	51.91
MV-DOPNet[fine-tuning]	82.15	53.99
DOPNet[w/ \mathcal{L}_r]	81.70	76.02
MV-DOPNet[w/ \mathcal{L}_r]	82.44	77.72
MV-DOPNet[w/ $\mathcal{L}_r, \mathcal{L}_n^w, \mathcal{L}_g^w$]	83.59	79.42


posed decision-level optimization and feature-level multi-view perception approach to produce more accurate layouts even in occlusion cases.

4.6 Ablation study

We exhibit ablation studies in Tab. 4 group 1, where each component of our model is evaluated on MatterportLay-



(a) Qualitative comparison on MatterportLayout [7] dataset.



(b) Qualitative comparison on ZInD [8] dataset.

Fig. 6. Qualitative comparison results evaluated on general layout datasets, MatterportLayout [7] and ZInD [8]. We compare our method with LED²-Net [5] and LGT-Net [3]. All the compared methods did not employ the post-processing strategy. The boundaries of the room layout on a panorama are shown on the left and the floor plan is on the right. Ground truth is best viewed in blue lines and the prediction in green. The predicted horizon depth, normal, and gradient are visualized below each panorama, and the ground truth is in the first row. We labeled the significant differences with dashed lines.

out [7]. To demonstrate the effectiveness of the proposed feature assembling mechanism, we first ablate cross-scale interaction in the distortion-aware feature aggregation module (denoted as “w/o Cross-scale interaction”). Then, we further remove the distortion elimination (denoted as “w/o Feature assembling”). To exhibit the benefit of orthogonal plane disentanglement, we remove the proposed disentangling orthogonal planes procedure (denoted as “w/o Disentangling planes”) and soft-flipping fusion, respectively. Finally, we show the effect of each attention mechanism (denoted as “w/o Discriminative channels”, “w/o Long-range dependencies”, and “w/o residuals”, respectively). For a fair comparison, we train each model with the same epochs (500).

Cross-scale Distortion Awareness. From Tab. 4, the pipeline without cross-scale interaction shows inferior per-

formance to the complete model. Further removing the distortion elimination part, the pipeline yields worse results. It is proved that both dealing with distortions and integrating cross-scale features are essential for layout estimation.

Pre-segmenting Orthogonal Planes. Since the geometric cues are essential for inferring 3D information from 2D images, we propose to disentangle orthogonal planes. To prove its effectiveness, we remove that stage but preserve the flipping fusion strategy (the effectiveness of this strategy is validated separately). The results in Tab. 4 show that this stage can make the performance more competitive. Besides, the embedding of symmetry property also contributes to the promotion of the predicted results.

Reconstructing 1D Representations. The dependencies of the 1D sequences have changed when disentangling the orthogonal planes. Hence, we propose to reconstruct the

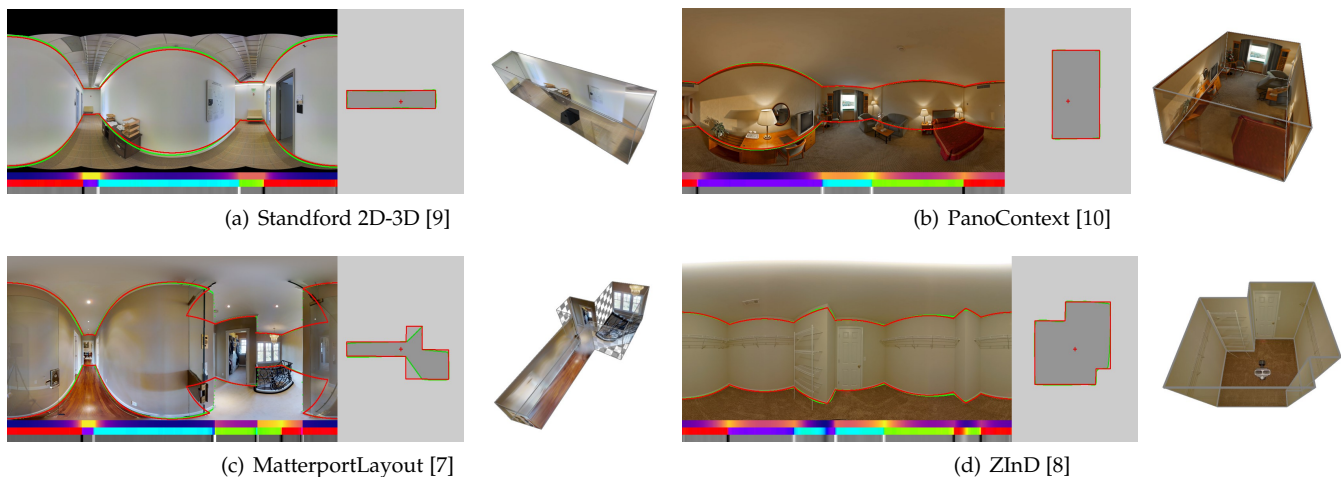


Fig. 7. The 3D visualization results. We exhibit the predicted boundaries (best viewed in green lines and the ones with post-processing of the prediction (best viewed in red lines) in the panoramas.




Fig. 8. We exhibit the heat maps of the features of the 1D representation. The one without orthogonal plane disentanglement is shown on the top, and ours is on the bottom. Without disentangling, the top contains redundant and confusing features that are not good for layout estimation. In contrast, our disentangled vertical plane features are more discriminative, showing more attention to the layout corners.




Fig. 9. Qualitative comparison on MVL dataset. We label the significant differences with solid red lines.

1D representations. We verify the effectiveness of each attention mechanism in turn. From Tab. 4, we can observe that all three reconstruction operations (generating discriminative channels, rebuilding long-range dependencies, and providing the missing residuals) can benefit the overall performance. Intuitively, when removing the channel-wise graph, the pipeline’s performance decreases significantly, demonstrating that enforcing the network concentration on discriminative channel information can capture effective information by avoiding redundancy.

As shown in Fig. 8, we can observe that the 1D representation of the layout produced by our DOPNet is much more precise than that produced by the baseline method.

MVL Specific Designs. From Tab. 4 group 2, we can observe

that multi-view insights elevate our model’s 2DIoU score to the forefront, while our formulated ratio constraints simultaneously enhance both 3DIoU and 2DIoU scores. Additionally, the improved $\mathcal{L}_n^w, \mathcal{L}_g^W$ losses further boost overall model performance, underscoring our approach’s effectiveness in advancing state-of-the-art technologies.

5 LIMITATION

Although the proposed MV-DOPNet can achieve effective multi-view perception, our approach still has some limitations. One limitation is that the advanced horizon-depth/ratio representation adopted in our methodology is based on an important assumption (*i.e.*, a normalized and

fixed camera height, 1.6 meters) which we must follow to recover the layouts. Besides, our approach follows the 360-MLC [21] method, which still relies on a pre-trained model. Since the room layout boundary representation is inherently sparse and topological, when we attempted to leverage multi-view projection consistency constraints for completely unsupervised learning, the model failed to converge and could not recover the room layout from the generated layout representation. In the future, we will try to solve this problem by proposing more suitable layout representations (e.g., an equivalent but dense representation).

6 CONCLUSION

In this paper, we propose a framework, for multi-view layout estimation. Current approaches generate a 1D representation by a vertical compression operation. We argue that this strategy confuses the semantics of different planes. We propose to disentangle orthogonal planes to get spotless 1D representations, as well as capture geometric cues in 3D space. Specifically, we obtain disentangled semantics by eliminating distortions, pre-segmenting orthogonal plans, and reconstructing dependencies of sequences. Besides, we propose both a horizon-depth optimizer and a multi-view specific module to fully exploit the multi-view layout consistency both at the decision level and the feature level. Our experimental results demonstrate that dealing with distortion, as well as integrating shallow and deep features, can enhance the performance. Experiments demonstrate that our algorithm significantly outperforms current SoTA methods on both monocular layout estimation and multi-view layout estimation tasks.

ACKNOWLEDGMENTS

We would like to thank Wei Wang, Shujun Huang, Yue Zhan, and Shuai Zheng for valuable discussions. Thanks to all authors for their insightful research and open source projects. This work was supported by the National Natural Science Foundation of China (Nos. 62172032, 62120106009).

REFERENCES

- [1] C. Sun, C.-W. Hsiao, M. Sun, and H.-T. Chen, "Horizonnet: Learning room layout with 1d representation and pano stretch data augmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1047–1056.
- [2] C. Sun, M. Sun, and H.-T. Chen, "Hohonet: 360 indoor holistic understanding with latent horizontal features," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 2573–2582.
- [3] Z. Jiang, Z. Xiang, J. Xu, and M. Zhao, "Lgt-net: Indoor panoramic room layout estimation with geometry-aware transformer network," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 1654–1663.
- [4] G. Pintore, M. Agus, and E. Gobbetti, "Atlantnet: inferring the 3d indoor layout from a single 360 image beyond the manhattan world assumption," in *European Conference on Computer Vision*. Springer, 2020, pp. 432–448.
- [5] F.-E. Wang, Y.-H. Yeh, M. Sun, W.-C. Chiu, and Y.-H. Tsai, "Led2-net: Monocular 360deg layout estimation via differentiable depth rendering," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 12956–12965.
- [6] Z. Shen, Z. Zheng, C. Lin, L. Nie, K. Liao, S. Zheng, and Y. Zhao, "Disentangling orthogonal planes for indoor panoramic room layout estimation with cross-scale distortion awareness," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 17 337–17 345.
- [7] C. Zou, J.-W. Su, C.-H. Peng, A. Colburn, Q. Shan, P. Wonka, H.-K. Chu, and D. Hoiem, "Manhattan room layout reconstruction from a single 360 image: A comparative study of state-of-the-art methods," *International Journal of Computer Vision*, vol. 129, pp. 1410–1431, 2021.
- [8] S. Cruz, W. Hutchcroft, Y. Li, N. Khosravan, I. Boyadzhiev, and S. B. Kang, "Zillow indoor dataset: Annotated floor plans with 360deg panoramas and 3d room layouts," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 2133–2143.
- [9] I. Armeni, S. Sax, A. R. Zamir, and S. Savarese, "Joint 2d-3d-semantic data for indoor scene understanding," *arXiv preprint arXiv:1702.01105*, 2017.
- [10] Y. Zhang, S. Song, P. Tan, and J. Xiao, "Panocontext: A whole-room 3d context model for panoramic scene understanding," in *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part VI 13*. Springer, 2014, pp. 668–686.
- [11] J. M. Coughlan and A. L. Yuille, "Manhattan world: Compass direction from a single image by bayesian inference," in *Proceedings of the seventh IEEE international conference on computer vision*, vol. 2. IEEE, 1999, pp. 941–947.
- [12] C. Zou, A. Colburn, Q. Shan, and D. Hoiem, "Layoutnet: Reconstructing the 3d room layout from a single rgb image," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2051–2059.
- [13] S.-T. Yang, F.-E. Wang, C.-H. Peng, P. Wonka, M. Sun, and H.-K. Chu, "Dula-net: A dual-projection network for estimating room layouts from a single rgb panorama," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 3363–3372.
- [14] C. Fernandez-Labrador, J. M. Facil, A. Perez-Yus, C. Demonceaux, J. Civera, and J. J. Guerrero, "Corners for layout: End-to-end layout recovery from 360 images," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1255–1262, 2020.
- [15] S. Rao, V. Kumar, D. Kifer, C. L. Giles, and A. Mali, "Omnilayout: Room layout reconstruction from indoor spherical panoramas," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 3706–3715.
- [16] L. Yuan, Y. Chen, T. Wang, W. Yu, Y. Shi, Z.-H. Jiang, F. E. Tay, J. Feng, and S. Yan, "Tokens-to-token vit: Training vision transformers from scratch on imagenet," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 558–567.
- [17] W. Wang, E. Xie, X. Li, D.-P. Fan, K. Song, D. Liang, T. Lu, P. Luo, and L. Shao, "Pyramid vision transformer: A versatile backbone for dense prediction without convolutions," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 568–578.
- [18] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10 012–10 022.
- [19] P. V. Tran, "Sslayout360: Semi-supervised indoor layout estimation from 360 panorama," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, 2021, pp. 15 348–15 357.
- [20] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," *Advances in neural information processing systems*, vol. 30, 2017.
- [21] B. Solarte, C.-H. Wu, Y.-C. Liu, Y.-H. Tsai, and M. Sun, "360-mlc: Multi-view layout consistency for self-training and hyperparameter tuning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 6133–6146, 2022.
- [22] Y. Lee, J. Jeong, J. Yun, W. Cho, and K.-J. Yoon, "Spherephd: Applying cnns on 360 images with non-euclidean spherical polyhedron representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 2, pp. 834–847, 2020.
- [23] Y. Xu, Z. Zhang, and S. Gao, "Spherical dnns and their applications in 360 images and videos," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 10, pp. 7235–7252, 2021.
- [24] C. Zhuang, Z. Lu, Y. Wang, J. Xiao, and Y. Wang, "Spdet: Edge-aware self-supervised panoramic depth estimation transformer with spherical geometry," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 10, pp. 12 474–12 489, 2023.
- [25] Y.-C. Su and K. Grauman, "Learning spherical convolution for

- fast features from 360 imagery," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [26] T. S. Cohen, M. Geiger, J. Köhler, and M. Welling, "Spherical cnns," *arXiv preprint arXiv:1801.10130*, 2018.
- [27] B. Coors, A. P. Condurache, and A. Geiger, "Spherenet: Learning spherical representations for detection and classification in omnidirectional images," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 518–533.
- [28] Z. Shen, C. Lin, K. Liao, L. Nie, Z. Zheng, and Y. Zhao, "Panoforner: Panorama transformer for indoor 360 depth estimation," in *European Conference on Computer Vision*. Springer, 2022, pp. 195–211.
- [29] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.
- [30] Z. Shen, C. Lin, L. Nie, K. Liao, and Y. Zhao, "Distortion-tolerant monocular depth estimation on omnidirectional images using dual-cubemap," in *2021 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2021, pp. 1–6.
- [31] Y. Li and A. Gupta, "Beyond grids: Learning graph representations for visual recognition," *Advances in neural information processing systems*, vol. 31, 2018.
- [32] X. Wang, S. Liu, H. Ma, and M.-H. Yang, "Weakly-supervised semantic segmentation by iterative affinity learning," *IJCV*, vol. 128, no. 6, pp. 1736–1749, 2020.
- [33] C. Lin, S. Zheng, Z. Liu, Y. Li, Z. Zhu, and Y. Zhao, "Sgt: Scene graph-guided transformer for surgical report generation," in *International conference on medical image computing and computer-assisted intervention*. Springer, 2022, pp. 507–518.
- [34] Y. Yao, Z. Luo, S. Li, T. Fang, and L. Quan, "Mvsnet: Depth inference for unstructured multi-view stereo," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 767–783.
- [35] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014.
- [36] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niebner, M. Savva, S. Song, A. Zeng, and Y. Zhang, "Matterport3d: Learning from rgb-d data in indoor environments," in *2017 International Conference on 3D Vision (3DV)*. IEEE, 2017, pp. 667–676.
- [37] M. Savva, A. X. Chang, A. Dosovitskiy, T. Funkhouser, and V. Koltun, "Minos: Multimodal indoor simulator for navigation in complex environments," *arXiv preprint arXiv:1712.03931*, 2017.
- [38] S. K. Ramakrishnan, A. Gokaslan, E. Wijmans, O. Maksymets, A. Clegg, J. Turner, E. Undersander, W. Galuba, A. Westbury, A. X. Chang *et al.*, "Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai," *arXiv preprint arXiv:2109.08238*, 2021.
- [39] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik *et al.*, "Habitat: A platform for embodied ai research," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9339–9347.
- [40] Y. Zhao, C. Wen, Z. Xue, and Y. Gao, "3d room layout estimation from a cubemap of panorama image via deep manhattan hough transform," in *European Conference on Computer Vision*. Springer, 2022, pp. 637–654.